Optimizing Energy Efficiency in Heterogeneous Computing via Multi-Objective Scheduling with Reinforcement Learning

28th Workshop on Job Scheduling Strategies for Parallel Processing (JSSPP 2025)

Ezgi Nur ALISAN
PhD Student
at Özyeğin University

Asst. Prof. Ismail AKTURK
at Özyeğin University

Agenda

- Introduction
- Contributions
- Literary Review
- Methodology
- Experimental Setup
- Results
- Conclusion & Future Work

Back to Agenda Page

Introduction

- Energy constraints under growing workload demand
- Limitations with single-metric optimization
- Failure of static scheduling in dynamic, heterogeneous environments
- Need for adaptive, multi-objective, learning based scheduling

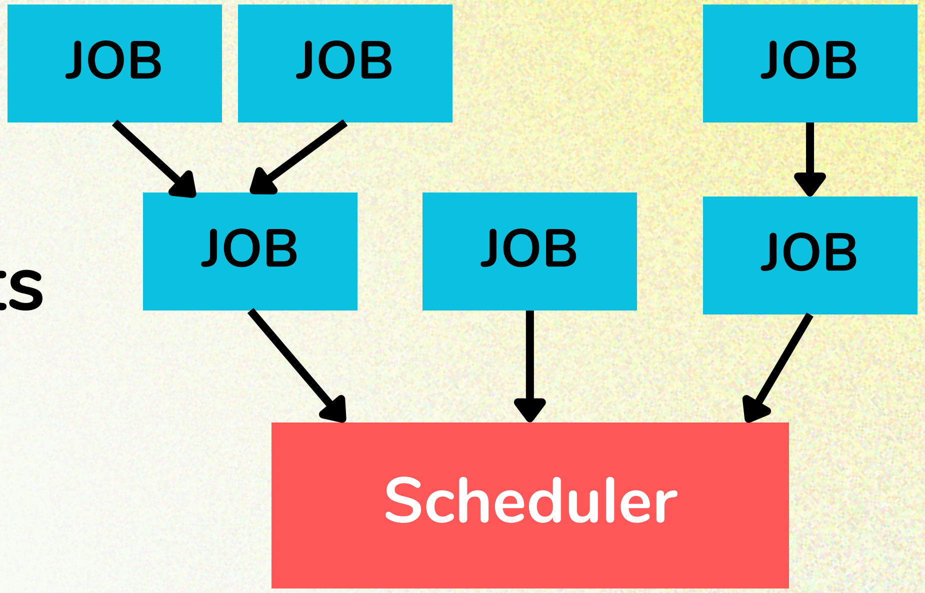

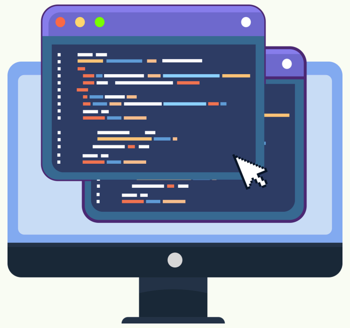

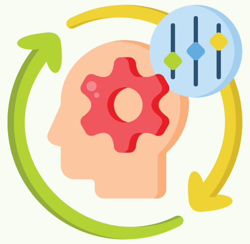

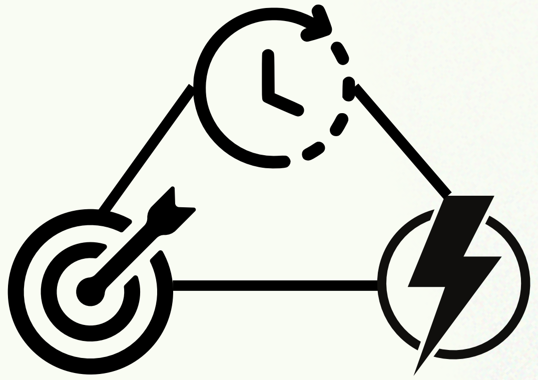

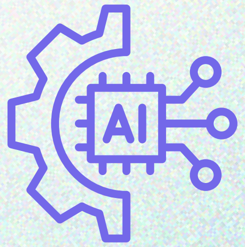

Back to Agenda Page

Contributions

Multi-objective scheduling

- Jointly Optimization of performance-energy-accuracy

Reinforcement learning-based scheduling

Evaluation of proposed method

- Across diverse workloads
- Comparison with state-of-art heuristics

Back to Agenda Page

Literary Review

Scheduling Algortihms

Heuristic-based

ETF, SJF, FIFO etc.

- Priorization of metric
- Difficulty on balancing multiple competing objectives
- Lack of adaptability

Machine Learning-based

Decision Tree, SVM, etc.

- Extensive offline training
- Limitations of real-time decision-making in dynamic systems

Evaluationary Algorithms

Particle Swarm Opt., etc.

- High computational overhead, less practical

Back to Agenda Page

Literary Review

Scheduling Algortihms

↓

Reinforcement Learning-based

Q-learning, DQN, SARSA, PPO,

- Single-objective, especially performance
- Energy-Aware RL
- Multi-objective scheduling in heterogeneous computing

Methodology

[Back to Agenda Page](#)

Scheduling on Heterogeneous Systems

Characteristics

- Workloads with Complex Task Structures
- System-Level Constraints
- Dynamic Resource Allocation
- Online Decision-Making

Reinforcement Learning Approach

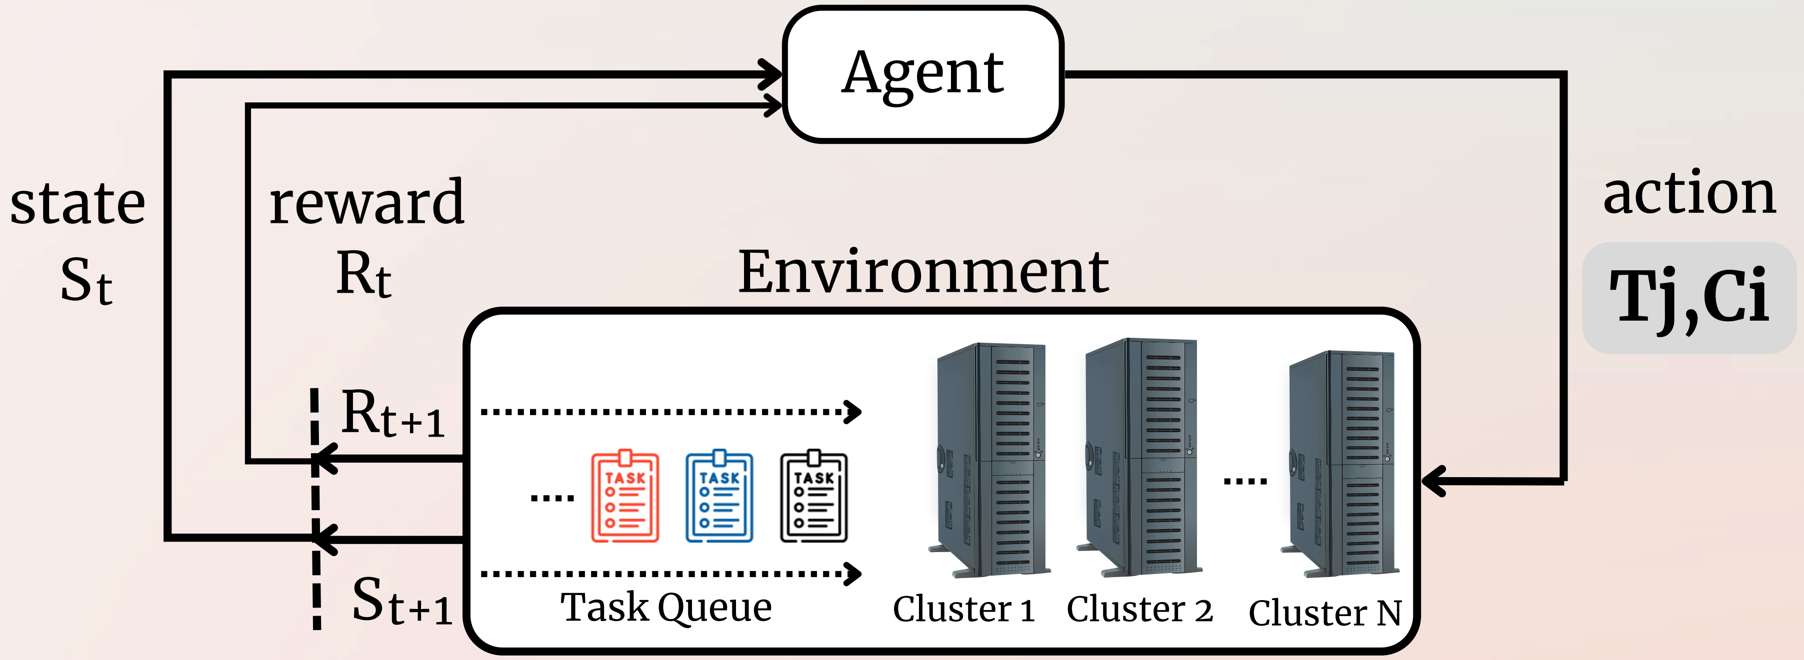

Problem Formulation

Back to Agenda Page

State Space

$$\mathbf{S} = \left[\mathbf{ID}^{\mathbf{T}}, \mathbf{T}_{\mathbf{c}1}, \mathbf{T}_{\mathbf{c}2}, \ldots, \mathbf{T}_{\mathbf{c}C}\right]$$

Action Space

$$A = \{a_1, a_2, \ldots, a_C\}$$

Proximal Policy Optimization

$$A(s_t, a_t) = r_t + \gamma V_\phi(s_{t+1}) - V_\phi(s_t)$$

$$r_t(\theta) = \frac{\pi_\theta(a_t \mid s_t)}{\pi_{\theta_{\text{old}}}(a_t \mid s_t)}$$

$$L^{CLIP}(\theta) = \mathbb{E}_t\left[\min\left(r_t(\theta)\hat{A}_t, \text{clip}\left(r_t(\theta), 1-\epsilon, 1+\epsilon\right)\hat{A}_t\right)\right]$$

Multi Objective Proximal Policy Optimization

$$\mathbf{R}_{multi-obj} = [R_1, R_2, \ldots, R_n]$$

Problem Formulation

Back to Agenda Page

$$r(s,a) = \begin{cases} \text{reward}, & \text{if termination condition is met} \\ 0, & \text{otherwise} \end{cases}$$

Reward Function

$$R_{Multi-Obj} = [R_{perf}, R_{energy}, R_{accuracy}]$$

$$R_{perf} = \frac{P_{max} - T_{total}}{P_{max}}$$

$$R_{energy} = \frac{E_{max} - E_{total}}{E_{max}}$$

$$R_{accuracy} = \frac{A - A_{min}}{1 - A_{min}}$$

Back to Agenda Page

Experimental Setup

Domain-Specific System-on-Chip Simulation Framework (DS3) [1]

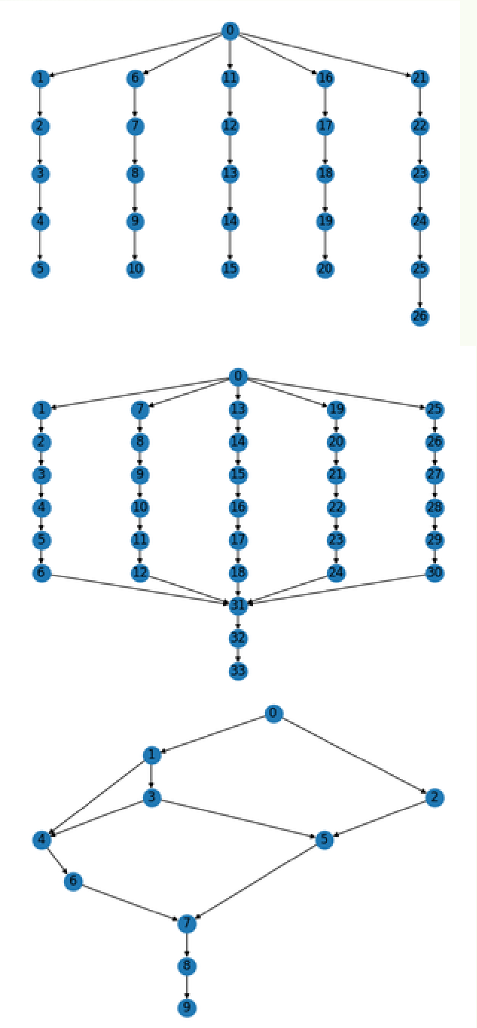

6 benchmark applications

- Wifi Transmitter
- Wifi Receiver
- Single Carrier-Transmitter
- Single Carrier-Receiver
- Range Detection
- Temporal Mitigation

6 different compute clusters

- LITTLE (4)
- BIG (4)
- FF ACCELERATOR-1(4)
- FF ACCELERATOR-2(4)
- FFT(4)
- VITERBI DECODER(1)

FF* : Fixed Function

[1] Arda, S.E., Krishnakumar, A., Goksoy, A.A., Kumbhare, N., Mack, J., Sartor, A.L., Akoglu, A., Marculescu, R., Ogras, U.Y.: DS3: A system-level domain-specific system-on-chip simulation framework. IEEE Transactions on Computers 69(8), 1248–1262 (2020)

Back to Agenda Page

Experimental Setup

Integration of RL-Based Scheduler to Simulation Framework

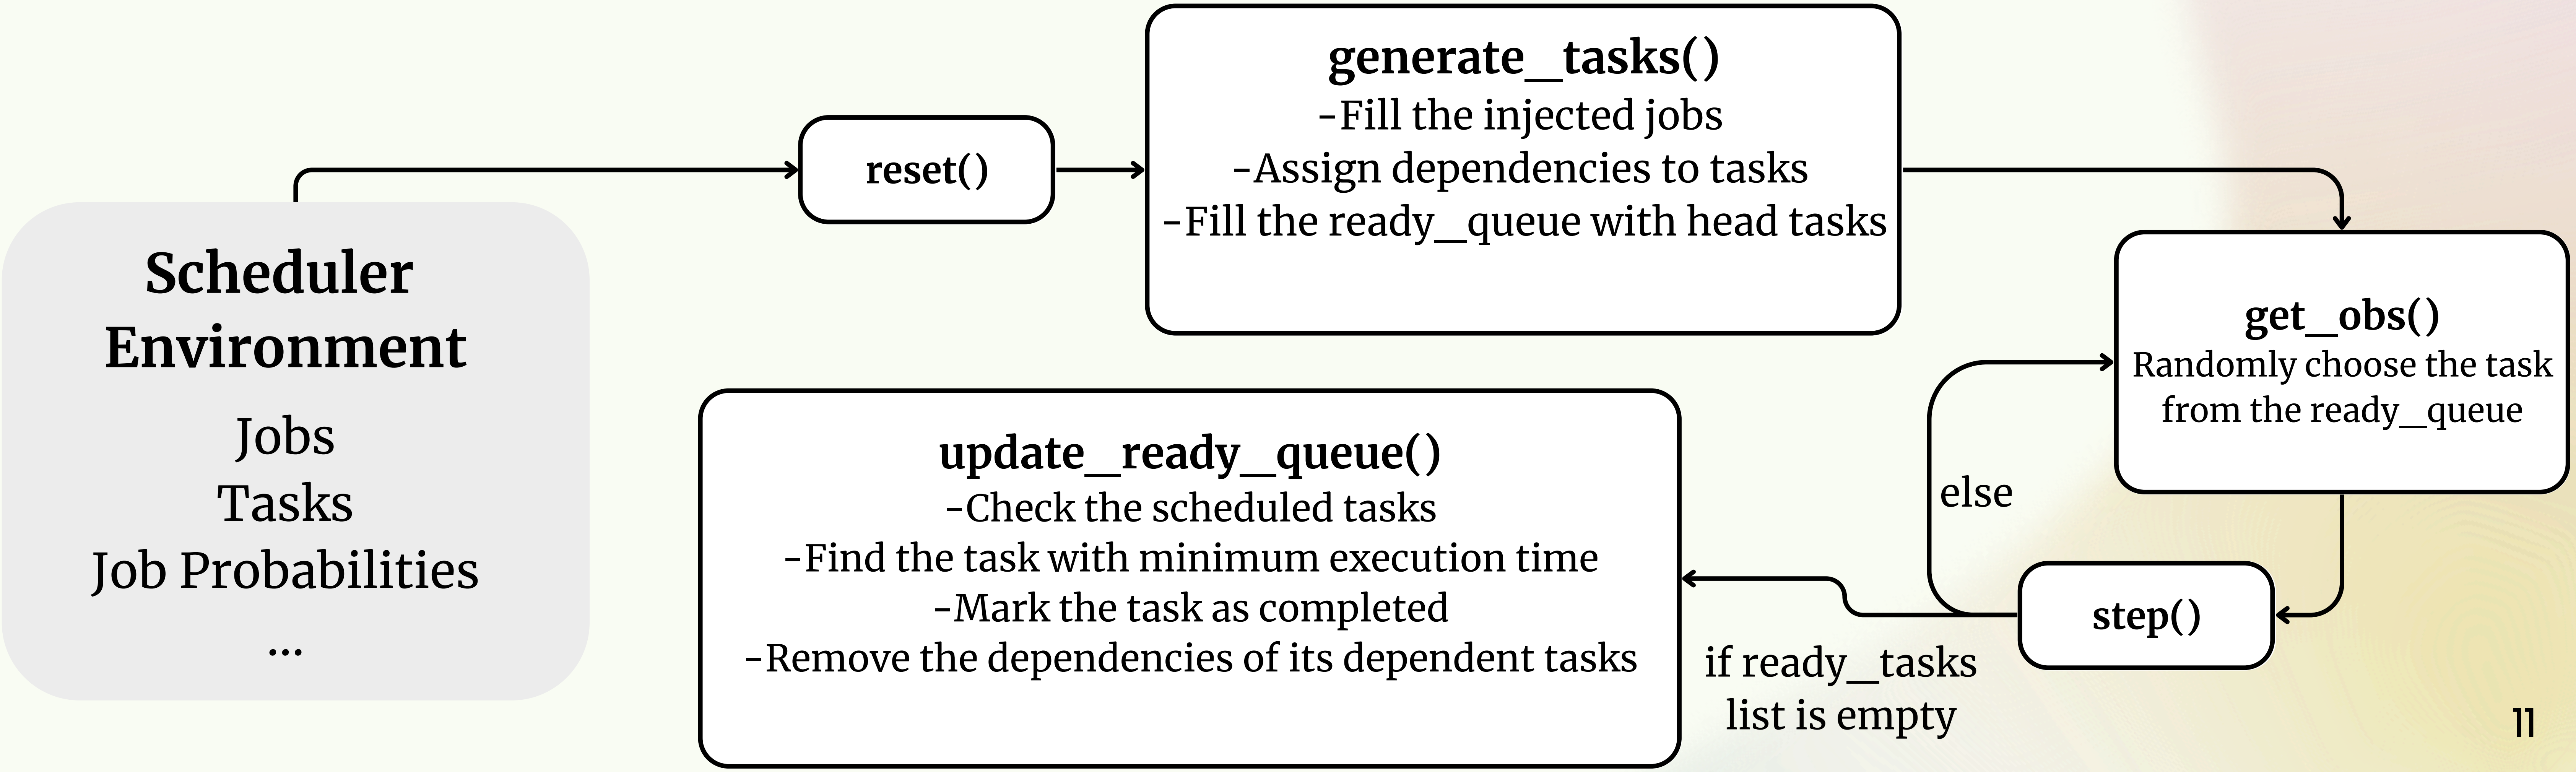

Back to Agenda Page

Experimental Setup

Integration of RL-Based Scheduler to Simulation Framework

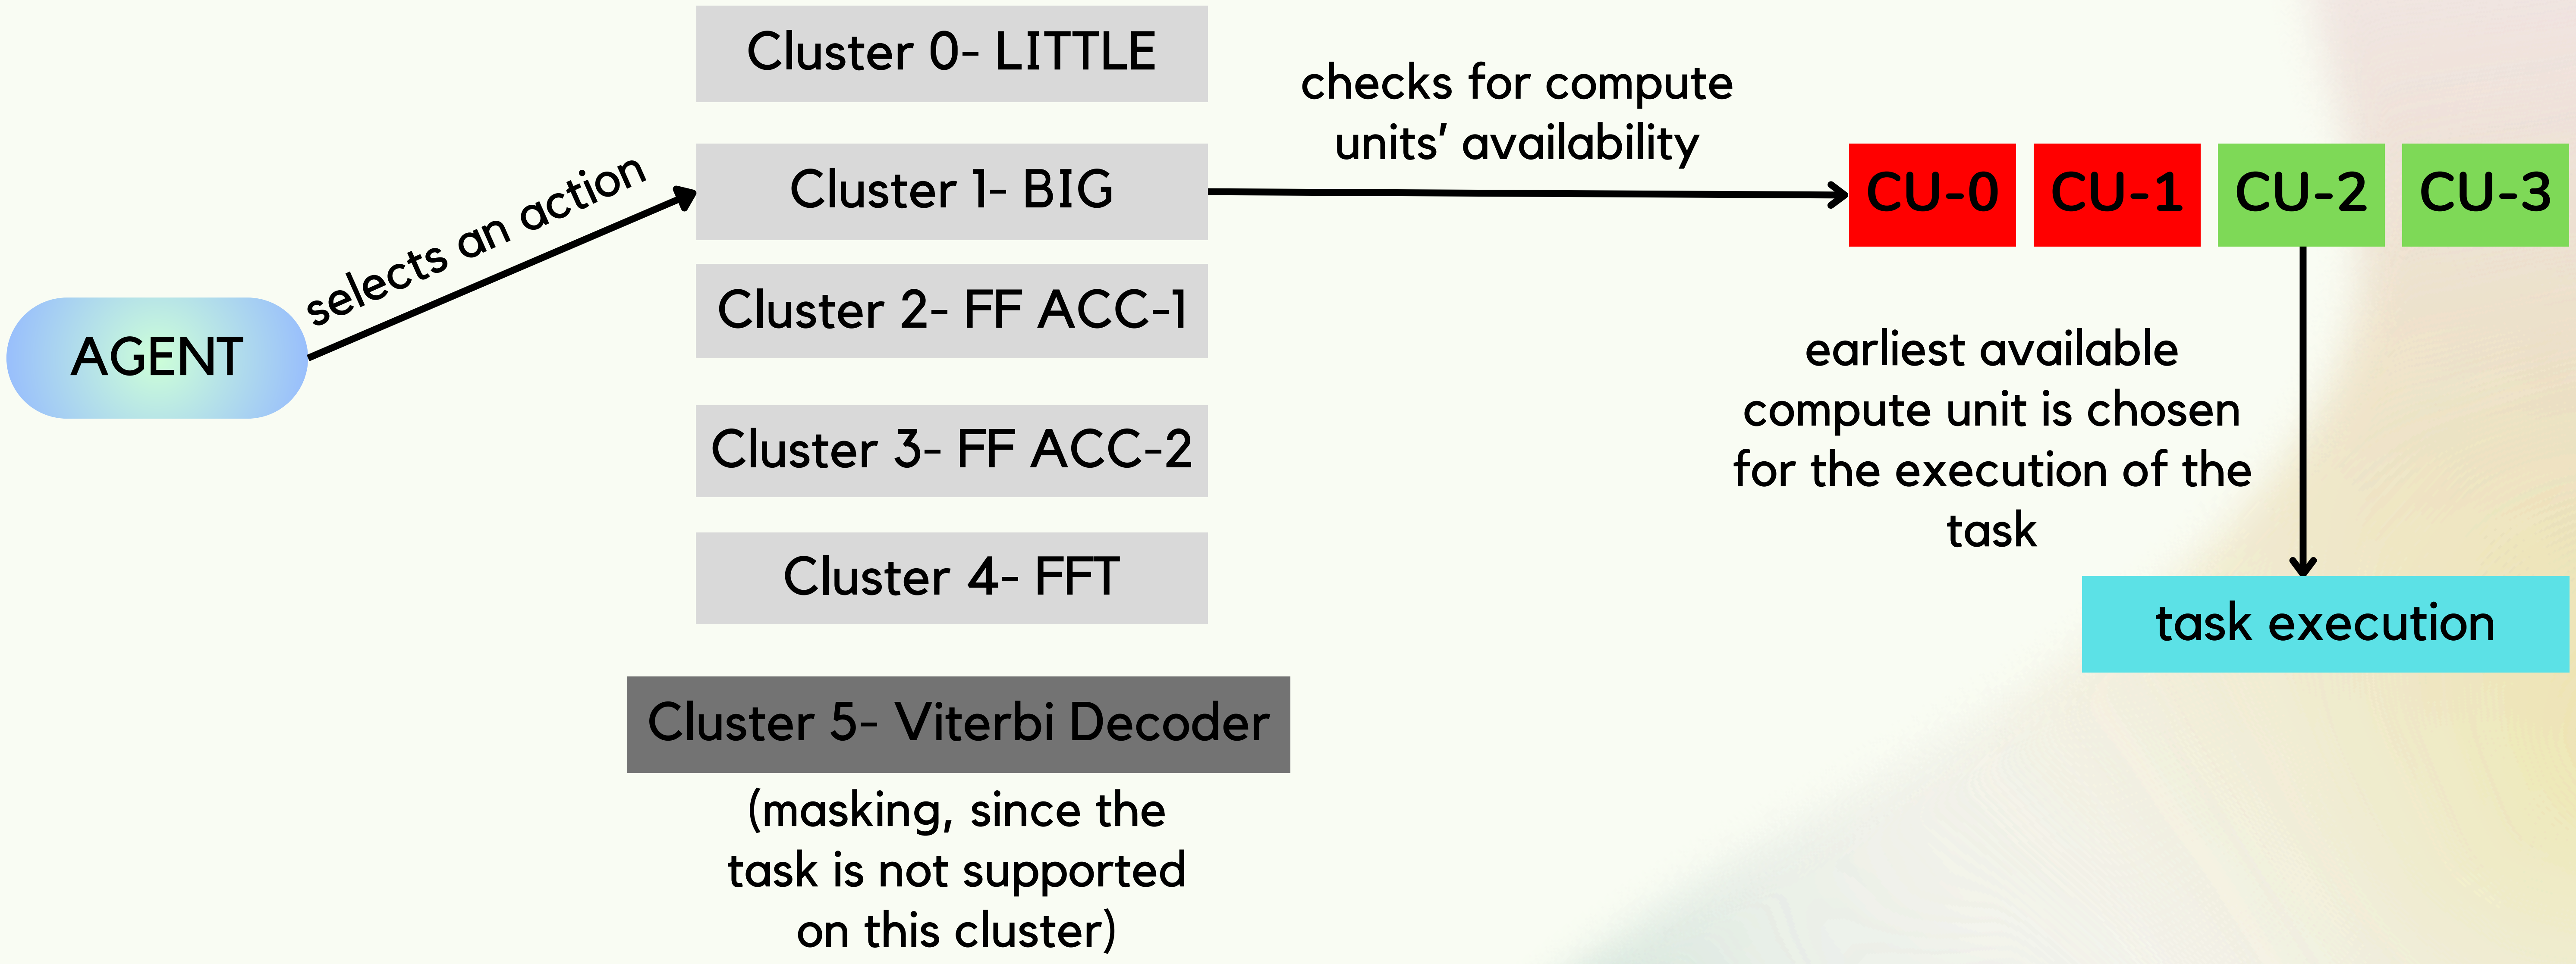

Back to Agenda Page

Experimental Setup

Performance Optimization

- Task execution times
- Communication overheads

Energy Optimization

- Energy model from DS3

Accuracy Optimization

- Predefined accuracy levels
- Ranging from 0.8 to 1.0

Back to Agenda Page

Experimental Setup

Time Prioritized	Energy/Power Prioritized	Accuracy-Prioritized
• Shortest Device Queue • Earliest Available Resource • Minimum Execution Time • Earliest Finish Time (EFT) • Shortest Task First (STF) • Earliest Task First (ETF) • ETF - Load Balancing • Longest Task First	• Earliest Task First - Energy • Earliest Task First - Power • Earliest Task First - EDP	• Maximum Accuracy • Minimum Accuracy

Results

[Back to Agenda Page](#)

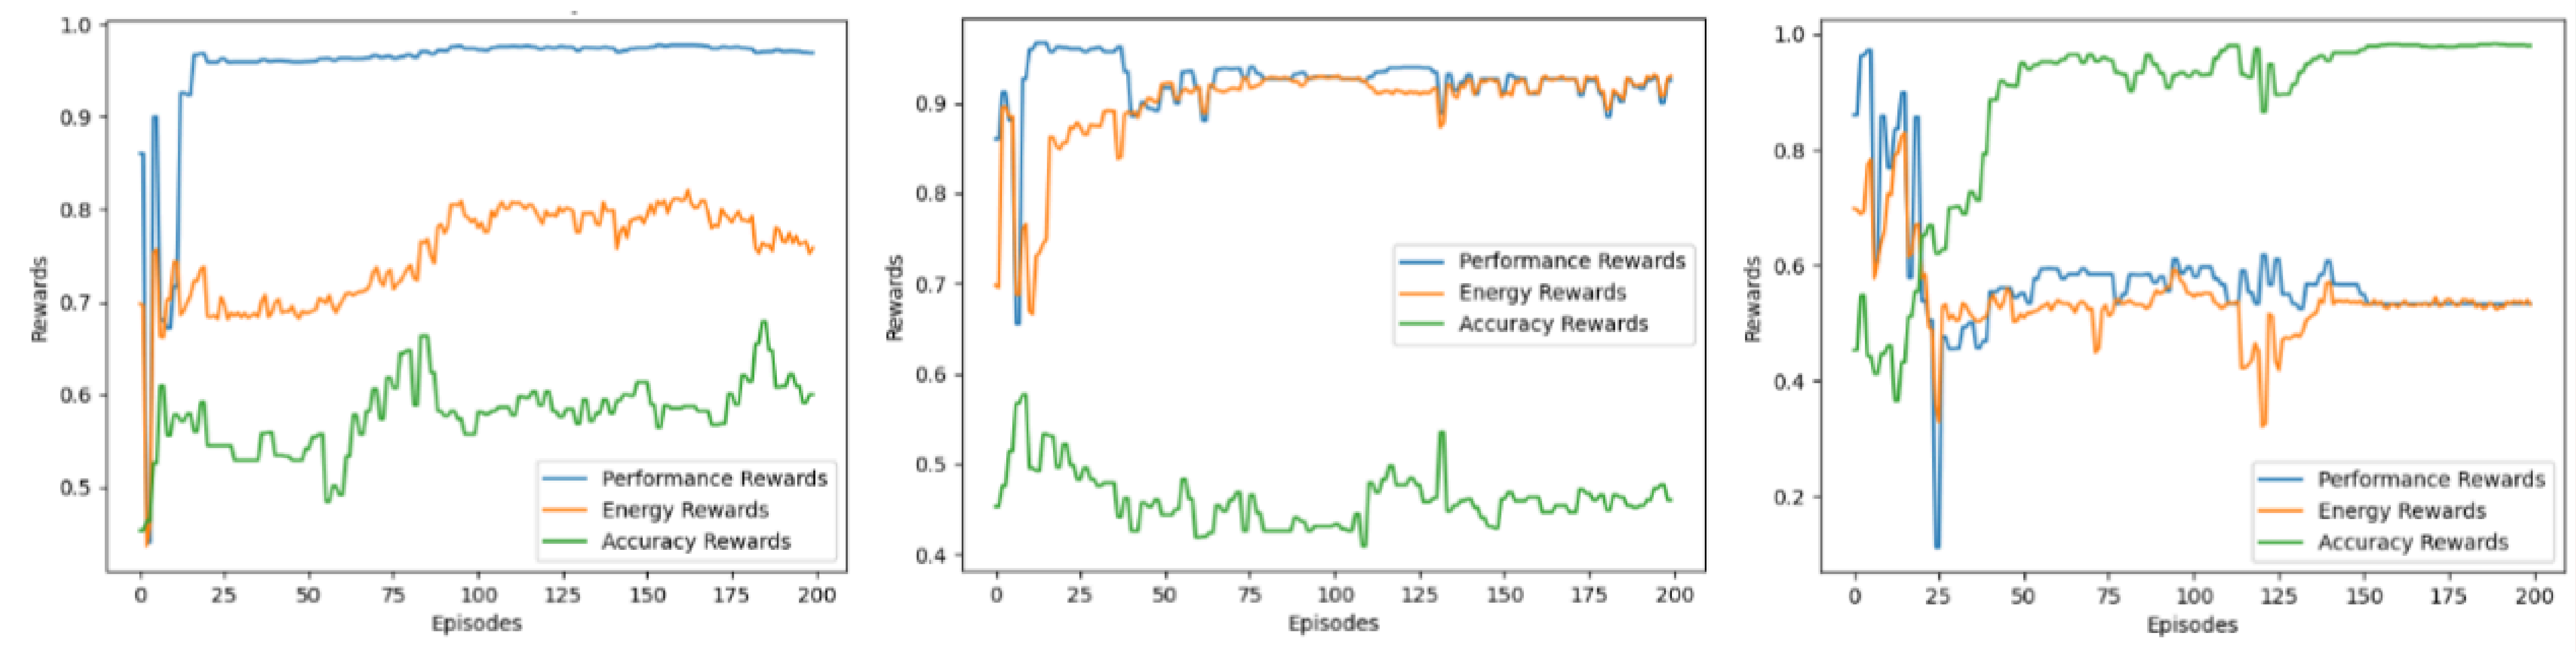

(a) Performance-only (b) Energy-only (c) Accuracy-only

Results

Back to Agenda Page

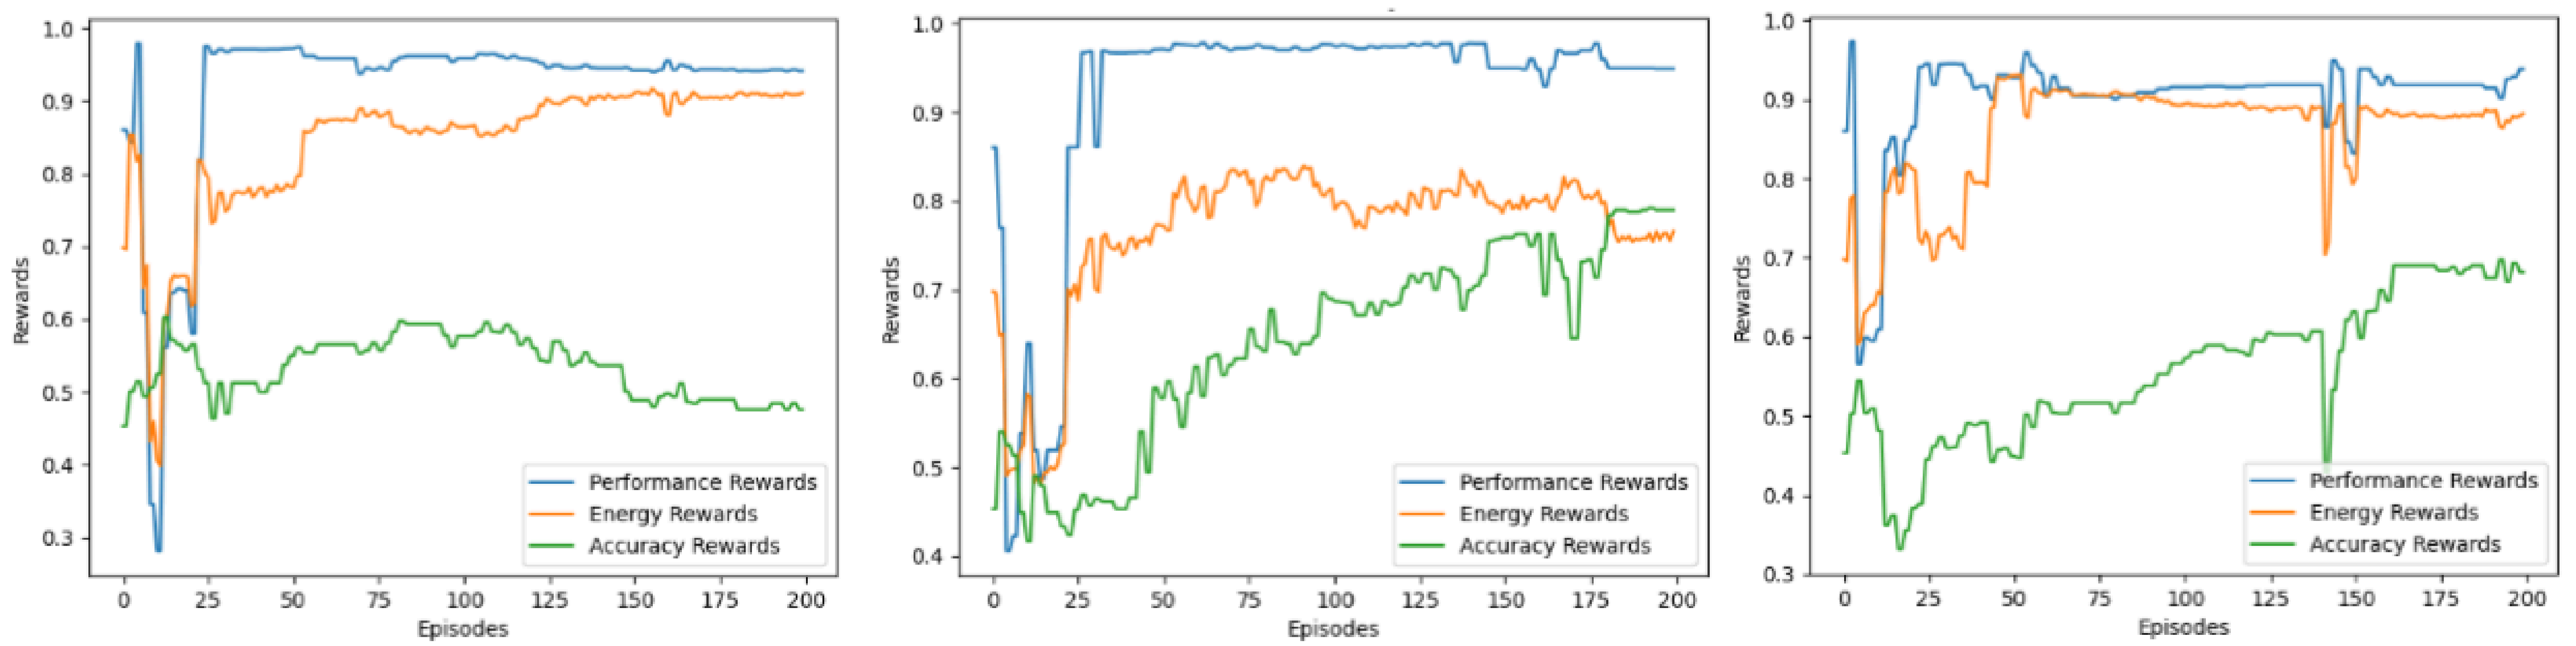

(d) Performance-Energy (e) Performance-Accuracy (f) Energy-Accuracy

Results

Back to Agenda Page

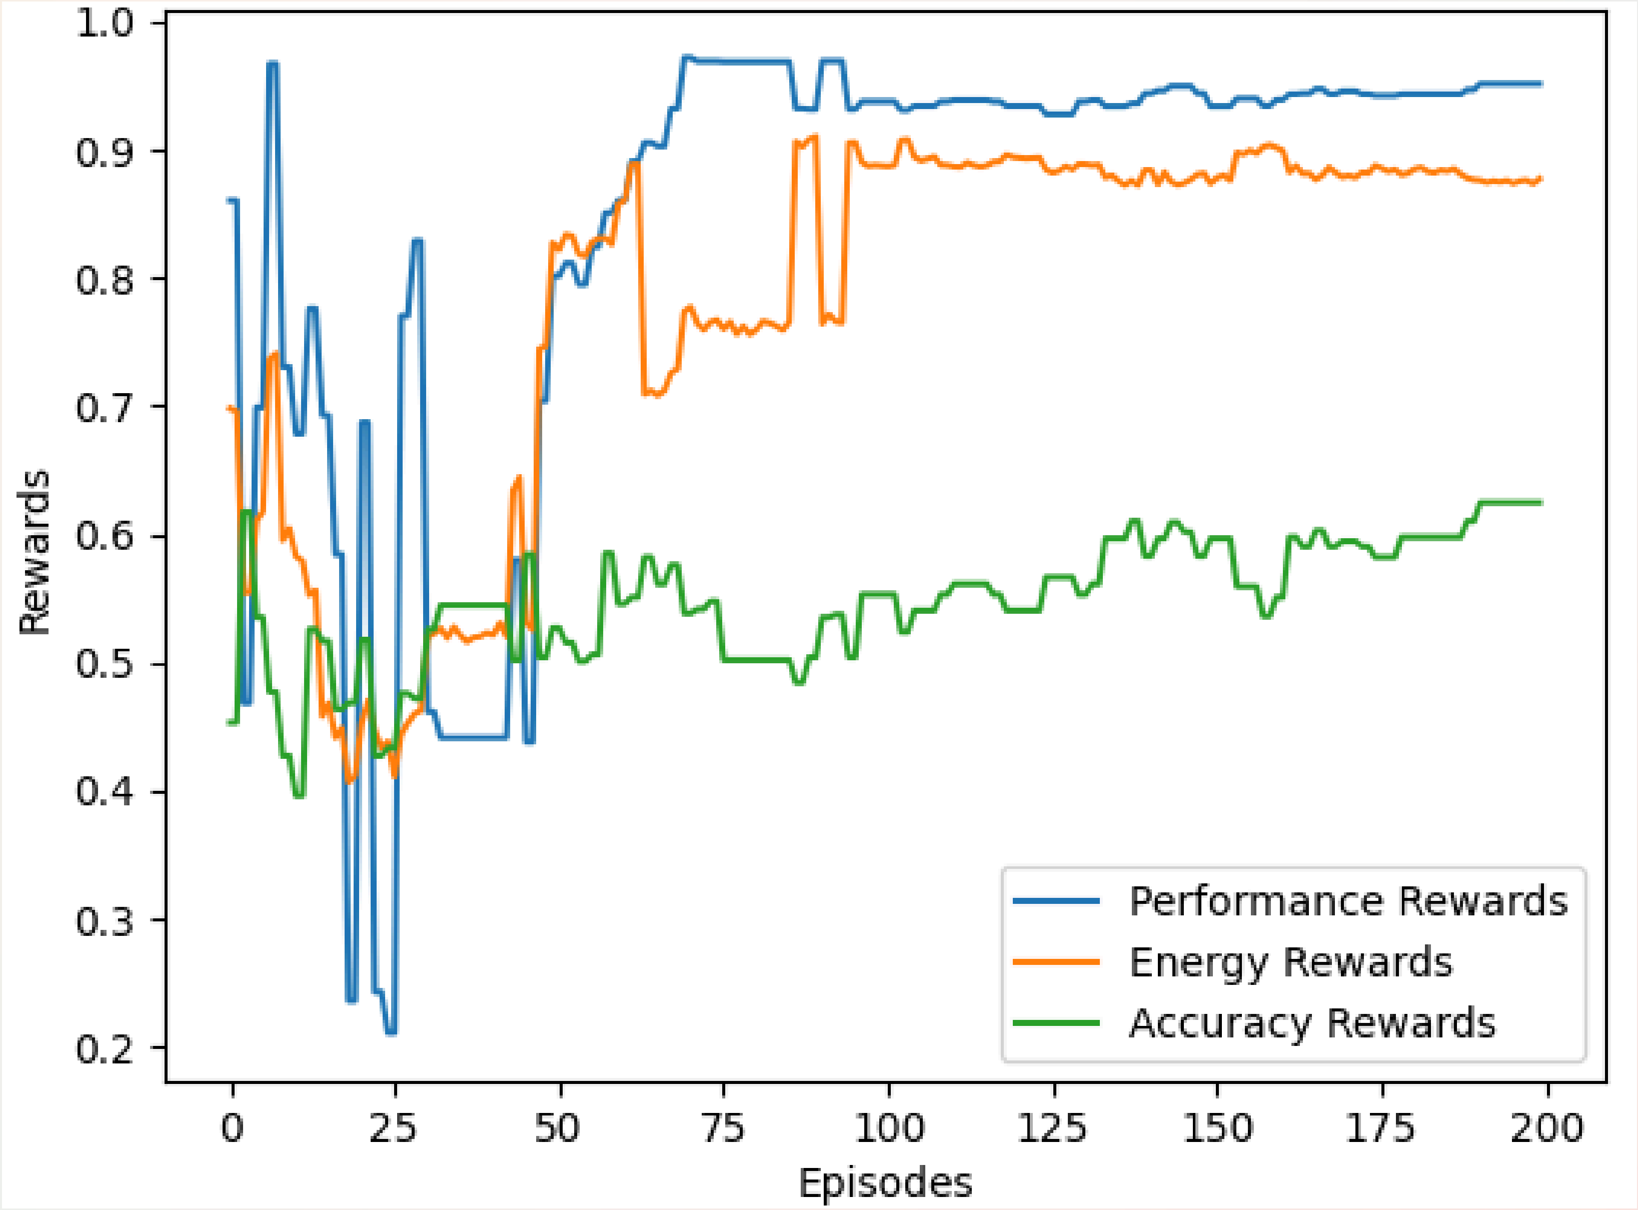

Performance-Energy-Accuracy
each weighted with 33%

Results

Back to Agenda Page

Scheduler	Execution Time(ns)	Energy (J)	EDP (J × ns)	Avg. (%) Accuracy
Random	2181748	0.00201	4450.0	0.94
Random-Available	1721085	0.00154	2817.9	0.94
Round Robin	1325748	0.00187	2677.1	0.94
Round Robin-Available	801197	0.00141	1392.2	0.94
Shortest Device Queue	211199	0.00090	191.1	0.94
Earliest Available Resource First	246722	0.00114	281.4	0.93
Minimum Execution Time	178434	0.00084	149.3	0.94
Earliest Finish Time	147299	0.00076	112.0	0.94
Shortest Task First	290282	0.00102	299.6	0.94
Longest Task First	147299	0.00076	111.9	0.94
ETF	**146699**	0.00077	112.2	0.94
ETF-Load Balancing	147001	0.00076	111.9	0.94
ETF-Power	8168870	0.00469	39176.6	0.95
ETF-Energy	537193	0.00036	198.4	0.94
ETF-EDP	332731	0.00033	108.1	0.94
Maximum Accuracy	6494036	0.00419	28131.8	**1.00**
Minimum Accuracy	1154061	0.00178	2057.8	0.89
RL Performance	158343	0.00077	121.4	0.95
RL Energy	333009	0.00028	92.6	0.94
RL Accuracy	1275632	0.00172	2224.2	**1.00**
RL Performance-Energy	326387	0.00039	126.3	0.95
RL Performance-Accuracy	182668	0.00085	155.3	0.97
RL Energy-Accuracy	387224	0.00057	224.0	0.97
RL Performance-Energy-Accuracy	314426	0.00048	150.1	0.96
RL Performance-biased	166837	0.00073	121.7	0.94
RL Energy-biased	333595	**0.00027**	**90.6**	0.94
RL Accuracy-biased	1001327	0.00170	1713.4	0.99

Our Approach

Results

Back to Agenda Page

- Up to **25%** energy reduction
- Maintains execution time within **9%** of the best heuristic
- **Max accuracy** in accuracy-critical tasks
- **0.3 ms** average decision-making time

Our Approach

Scheduler	Execution Time(ns)	Energy (J)	EDP (J × ns)	Avg. (%) Accuracy
Random	2181748	0.00201	4450.0	0.94
Random-Available	1721085	0.00154	2817.9	0.94
Round Robin	1325748	0.00187	2677.1	0.94
Round Robin-Available	801197	0.00141	1392.2	0.94
Shortest Device Queue	211199	0.00090	191.1	0.94
Earliest Available Resource First	246722	0.00114	281.4	0.93
Minimum Execution Time	178434	0.00084	149.3	0.94
Earliest Finish Time	147299	0.00076	112.0	0.94
Shortest Task First	290282	0.00102	299.6	0.94
Longest Task First	147299	0.00076	111.9	0.94
ETF	**146699**	0.00077	112.2	0.94
ETF-Load Balancing	147001	0.00076	111.9	0.94
ETF-Power	8168870	0.00469	39176.6	0.95
ETF-Energy	537193	0.00036	198.4	0.94
ETF-EDP	332731	0.00033	108.1	0.94
Maximum Accuracy	6494036	0.00419	28131.8	**1.00**
Minimum Accuracy	1154061	0.00178	2057.8	0.89
RL Performance	158343	0.00077	121.4	0.95
RL Energy	333009	0.00028	92.6	0.94
RL Accuracy	1275632	0.00172	2224.2	**1.00**
RL Performance-Energy	326387	0.00039	126.3	0.95
RL Performance-Accuracy	182668	0.00085	155.3	0.97
RL Energy-Accuracy	387224	0.00057	224.0	0.97
RL Performance-Energy-Accuracy	314426	0.00048	150.1	0.96
RL Performance-biased	166837	0.00073	121.7	0.94
RL Energy-biased	333595	**0.00027**	**90.6**	0.94
RL Accuracy-biased	1001327	0.00170	1713.4	0.99

Conclusion & Future Work

[Back to Agenda Page]

Reinforcement learning-based, multi-objective scheduler

Effectively balances trade-offs

Outperforms heuristic-based schedulers

- Up to **25% energy reduction**
- Maintains **execution time within 9%** of the best heuristic
- **Max accuracy** in accuracy-critical tasks
- 0.3 ms average decision-making time

Enhancing scalability and robustness

Efficiency improvement

Dynamic workload adaptation

Runtime scheduling decision overhead reduction

Thank you!