SC26 DOCTORAL SHOWCASE · DISSERTATION RESEARCH SUMMARY

Machine Learning-Based Multi-Objective Optimization for HPC Workload Scheduling

A Graph Neural Network & Reinforcement Learning Approach

Kyrian C. Adimora Advisor: Prof. Hongyang Sun Dept. of EECS, University of Kansas

UNIVERSITY OF KANSAS
Department of Electrical Engineering & Computer Science

SC26
DOCTORAL SHOWCASE
Nov 17–19, 2026 · Chicago, IL

377,234
Production HPC Jobs
Polaris, Mira, Cooley (GNN-RL Study)

56.6%
Peak Energy Reduction
vs. Classical Schedulers

50.0%
Peak Throughput Gain
vs. EASY-Backfilling

32%
Performance Variability
Reduction (HARMONIC)

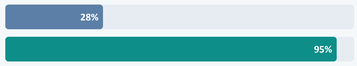
95.8%
Resource Utilization
Achieved (EA-GATSched)

01 The HPC Resource Management Crisis

Exascale systems draw up to 40 MW per facility, yet schedulers still optimize a single objective, ignoring **heterogeneity**, **uncertainty**, and **energy** until it's too late to act, closing these gaps is what lets gains reach production systems like SLURM, PBS, and LSF.

40 MW	Tens of $M	30–50%	15–40%
Peak Facility Power	Annual Energy Cost	Resource Estimation Error	Runtime Variability

STATE OF THE ART FALLS SHORT



Capability integration across 6 dimensions: multi-objective optimization, uncertainty, energy-awareness, graph modeling, adaptive learning, scalability, surveyed across 22 SOTA approaches.

TEN RESEARCH GAPS (G1–G10)

G1 Single-objective optimization dominance	**G2** Inaccurate estimation (30–50% error)
G3 Deterministic workload assumptions	**G4** Heterogeneous resource handling
G5 Limited job-dependency modeling	**G6** Static, non-adaptive production rules
G7 Theory–practice disconnect	**G8** Single-metric evaluation only
G9 Limited exascale-scale analysis	**G10** Resilience excluded from energy decisions

CENTRAL HYPOTHESIS

Graph-structured representations of workloads and system state, combined with adaptive learning policies and explicit uncertainty quantification, yield significantly better scheduling decisions than traditional methods. Tested via H1 (multi-objective superiority), H2 (uncertainty-robust allocation), and H3 (scalability & generalization).

A UNIFIED THREE-COMPONENT FRAMEWORK

GNN-RL — closes G1 G4 G5 G6 G9
EA-GATSched — closes G1 G4 G8 G10
HARMONIC — closes G2 G3 G7

▼ ▼ ▼

Unified Multi-Objective Scheduling Policy

Performance · Energy Efficiency · Resilience — detailed in Columns 2–4 →

PRIMARY RESEARCH QUESTIONS

- RQ1 — How can graph learning + RL jointly optimize performance, energy, and resilience?
- RQ2 — How can workload uncertainty be quantified and embedded into scheduling decisions?
- RQ3 — How should competing objectives be balanced under dynamic priorities and constraints?

DISSERTATION TIMELINE

Phase I–II — **Months 1–36**
GNN-RL foundations, problem formulation, preliminary multi-system validation

Phase III — **Months 37–48**
EA-GATSched implementation · GEOPM power-model calibration

Phase IV — **Months 49–54**
HARMONIC integration · comprehensive cross-framework evaluation

✓ PhD comprehensive exam successfully defended December 4, 2025 — dissertation defense forthcoming

EVALUATION FACILITIES

Polaris	ALCF · NVIDIA A100 GPU · Dragonfly+	229,404–241,772 jobs
Mira	ALCF · IBM Blue Gene/Q · 5D Torus	52,154 jobs
Cooley	ALCF · Intel Haswell · Fat-Tree	95,676–95,678 jobs
Theta / GEOPM	ALCF · Intel KNL · Power calibration	125–215W sweep

KEY TERMS FOR A MIXED AUDIENCE

GAT / GATv2	Graph Attention Network — learns which neighboring jobs/resources matter most
PPO / DDPG	Reinforcement learning algorithms that turn system state into scheduling actions
Aleatoric vs. Epistemic	Irreducible system noise vs. uncertainty from limited training data
Pareto-Optimal	No objective (energy, speed, reliability) can improve without worsening another
MDP	Markov Decision Process — the formal state/action/reward model behind every scheduler here
Cohen's d	Standardized effect size; d > 0.8 is conventionally considered a "large" effect

02 GNN-RL: Graph-Attention Scheduling Engine

FOUNDATIONAL COMPONENT · SC'24 + JOURNAL

A 3-layer Graph Attention Network encodes four job-relation types: dependency, resource-competition, temporal, and user-affinity, into per-node embeddings; a PPO policy then learns scheduling actions with meta-learned, adaptive multi-objective reward weights.

ARCHITECTURE

Job Queue → Graph Construction
4 typed edges: dependency, resource-competition, temporal, user-affinity

▼

Multi-Head Graph Attention (GAT)
3 layers · hidden [64,32,16] · heads [4,2,1] · typed attention α(i,j)

▼

PPO Policy + Adaptive Reward
R = w1·U + w2·T + w3·F + w4·M, weights meta-learned online

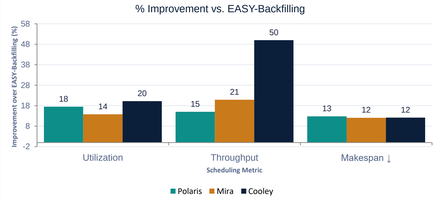
▼

Scheduling Decision + Experience Replay

Why GAT? (1) learnable attention highlights the resource conflicts that matter, ignoring noisy signals; (2) inductive design generalizes to a queue that changes size every step; (3) multi-head attention reasons about dependency, conflict, and timing in parallel, much like a human scheduler.

EVALUATION: 377,234 JOBS · 3 ALCF SYSTEMS

Polaris · Mira · Cooley (GPU, Blue Gene/Q, Visualization)



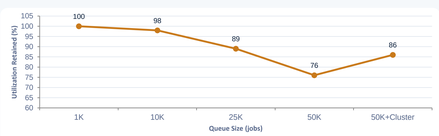
All gains significant at p<0.001 (Wilcoxon, Bonferroni-corrected); Cohen's d > 1.2 (large effect) on every metric and system, up to d = 16.98 for Cooley throughput.

- Medium jobs (16–128 nodes) gain most: +15.3 pp utilization, +28.4% throughput
- Maintains 76% effectiveness at 50K concurrent jobs; 86% with hierarchical (Louvain) clustering
- Reservation-aware: only −2.0% utilization loss under 12% reserved capacity, zero violations

JOB-LEVEL ANALYSIS (POLARIS)

Job Size	Δ Util. (pp)	Δ Throughput
Small (≤16 nodes)	+8.2	+18.4%
Medium (17–128)	+15.3	+28.4%
Large (>128 nodes)	+5.1	+9.3%

SCALABILITY — QUEUE SIZE



Decision latency: 0.05s at 1K jobs → 4.1s at 50K jobs; hierarchical Louvain clustering recovers utilization beyond 50K-job queues.

ABLATION — COMPONENT CONTRIBUTION

Configuration	Util.	Fairness	Makespan
Full GNN-RL	78.4%	0.92	10.9h
No Attention	71.2%	0.85	12.4h
No Graph Structure	69.8%	0.81	13.2h
PPO Only (MLP)	65.4%	0.76	14.7h

GNN-RL bottom line: validated on 377,234 production jobs — up to +20.2 pp utilization, +50% throughput

−1.6 pp
Utilization loss on Polaris under 12% reserved capacity

Zero
Reservation violations on Polaris (RV=0) under 12% reserved capacity

Demonstrated on Polaris under a 12% reservation load with no manual tuning, the policy learns to route around existing reservations on its own.

03 HARMONIC: Uncertainty-Aware Optimization

PUBLISHED · IEEE TPDS 2025

Bayesian neural networks decompose total predictive uncertainty into aleatoric (σ^2_a, inherent variability) and epistemic (σ^2_e, model limitations) components, feeding confidence-aware allocation and a multi-objective DDPG agent with formally ε-bounded Pareto-front guarantees.

ARCHITECTURE

Bayesian Workload Model
$\sigma^2_total = \sigma^2_aleatoric + \sigma^2_epistemic$ · variational inference

▼

Heterogeneous System Graph
Jobs + resources · GAT message passing · uncertainty propagation O(diam(G))

▼

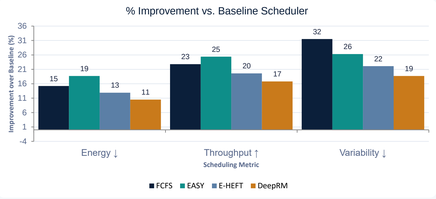
Multi-Objective DDPG
R = w_p·R_perf + w_e·R_energy + w_r·R_resil · domain-adaptive weights

▼

Confidence-Aware Resource Allocation

EVALUATION: SIMULATION + 64-NODE TESTBED

4 Workloads Compute · Memory · Comm · AI-HPC hybrid



Averaged across baselines: 14.3% energy reduction, 21.2% throughput gain, 24.7% variability reduction — benefits largest for communication- and memory-intensive workloads (WL-3, WL-2).

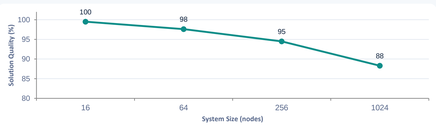
- Energy–performance balance point: ~85% of peak performance at only 60% of baseline energy (R^2=0.94)
- 23% downtime reduction via predictive failure modeling + energy-aware job migration

Three Fault-Tolerance Mechanisms: predictive failure modeling, intelligent checkpointing, and energy-aware job migration — jointly informed by the same uncertainty estimates used for scheduling.

UNCERTAINTY-QUANTIFICATION BENEFIT BY WORKLOAD

Workload	Energy↓	Variability↓
WL-1 Compute	8.3%	15.2%
WL-2 Memory	17.5%	28.4%
WL-3 Comm.	21.3%	32.1%
WL-4 AI-HPC	14.8%	23.7%

SCALABILITY — NODE COUNT



Decision time scales linearly (0.02s at 16 nodes → 3.21s at 1024); theoretical analysis projects >73% quality at 50,000 nodes.

HYBRID ALGORITHMIC-LEARNING GUARANTEE

$$C(\text{hybrid}) \le \min\{\, C(ALG),\ (1+\eta)\cdot C(ML) \,\}$$

Formal worst-case bound combining algorithmic guarantees with learned performance — defaults safely to the algorithmic solution when learned predictions are unreliable.

DOMAIN-ADAPTIVE REWARD WEIGHTS

Domain	Perf.	Energy	Resil.
AI / ML workloads	0.4	0.4	0.2
CFD (energy-critical)	0.3	0.5	0.2
Genomics (deadline)	0.5	0.2	0.3

HARMONIC bottom line: ε-bounded Pareto guarantees — 24.7% avg. variability cut, $multi-M/yr potential savings

23%
Downtime reduction via predictive migration

85% @ 60%
Peak performance retained at reduced energy (R^2=0.94)

Both gains flow from the same uncertainty estimates that drive HARMONIC's scheduling decisions — no separate subsystem required.

04 EA-GATSched: Physics-Informed Energy Scheduling

SC'25 POSTER · FULL PAPER SUBMITTED

GATv2 learns dynamic (not static) attention over conflict + dependency job graphs via three specialized heads (energy, deadline, load-balance); a CMOS power model calibrated on ALCF's public GEOPM telemetry feeds a state-conditioned Pareto-weight network and PPO policy.

ARCHITECTURE

Dynamic Job Graph G(t)
Conflict edges + dependency edges · rebuilt every 60s epoch

▼

GATv2 Encoder (3 Heads)
Energy-compatibility · deadline-urgency · load-balance attention

▼

GEOPM-Calibrated Power Model
CMOS physics model · MAE ≤ 3.1% · R^2 ≥ 0.95 (125–215W sweep)

▼

Pareto-Weight Network + PPO
State-conditioned (α_t, β_t, δ_t) · minimizes Pareto regret

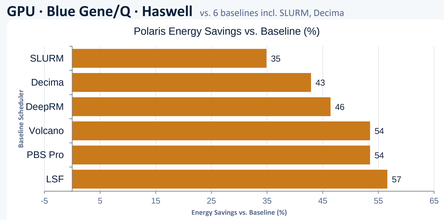
▼

Power-Budgeted Scheduling Action

Why GATv2, not GAT? Standard GAT assigns the same attention to jobs with identical features even when their conflict neighborhoods differ. GATv2's dynamic attention tells them apart.

EVALUATION: 389,604 JOBS · 3 ARCHITECTURES

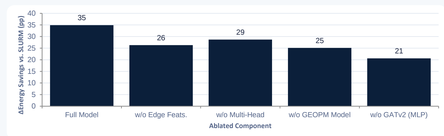
GPU · Blue Gene/Q · Haswell vs. 6 baselines incl. SLURM, Decima



27.3–56.6% energy savings over classical schedulers; 14.3–30.2% over ML baselines (DeepRM, Decima) — all p<0.01, Bonferroni-corrected, paired t-tests.

- Up to 95.8% utilization, 15.5–17.7 GFLOPS/W; sub-200ms inference latency at 5,000 nodes
- Cohen's d = 8.1 (GATv2) and d = 5.9 (power model), effects robust across 10 random seeds

ABLATION — Δ ENERGY SAVINGS (PP)



GATv2 encoder is the largest single contributor (−14.3pp if removed); the GEOPM-calibrated power model ranks 2nd (−9.8pp vs. a core-utilization proxy).

PER-ARCHITECTURE RESULTS

System	Energy↓	Util.	GFLOPS/W
Polaris (GPU)	34.9%	95.8%	17.72
Mira (BG/Q)	27.5%	87.6%	15.54
Cooley (Haswell)	27.3%	66.8%	15.88

GEOPM POWER-MODEL CALIBRATION

System	MAE	R^2
Theta / KNL (GEOPM)	2.1%	0.98
Polaris (A100 GPU)	2.3%	0.97
Mira (Blue Gene/Q)	3.1%	0.95
Cooley (Haswell)	1.8%	0.98

Calibrated on ALCF's public GEOPM traces (power governor + balancer agents, 125–215W), then architecture-rescaled via vendor peak-power specs.

LIMITATIONS

- Simulator-only validation to date, no live production deployment yet
- ~20-epoch cold start before outperforming SLURM; GEOPM re-calibration (~2h) needed per new architecture

EA-GATSched bottom line: hardware-grounded power model — up to 56.6% energy saved, 95.8% utilization

05 Contributions, Impact & What Comes Next

✓ Comprehensive exam successfully defended — Dec. 4, 2025, University of Kansas

KEY CONTRIBUTIONS

- First unified GNN + RL + uncertainty-quantification framework jointly optimizing performance, energy, and resilience for HPC scheduling
- Multi-relational graph encoding (4 edge types) with GAT / GATv2 attention over heterogeneous job–resource dependencies
- Bayesian aleatoric / epistemic uncertainty decomposition with formal ε-bounded Pareto-optimality guarantees
- Hardware-calibrated (GEOPM) physics power model closing the proxy-metric gap in energy-aware RL rewards
- Validation on 377K–390K production jobs per study across 4 real ALCF systems (Polaris, Mira, Cooley, Theta), with rigorous statistical testing

Statistical Rigor: n ≥ 5–10 independent runs · Wilcoxon signed-rank (Bonferroni α=0.001–0.01) · 95% confidence intervals · Cohen's d > 1.2 (large effects) throughout

BROADER IMPACT

- Sustainability: up to 18 MW potential savings/exascale facility; ~$214K+/yr cost avoidance per system modeled
- Scientific productivity: higher utilization & throughput translate to more completed simulations per facility-year
- Open, reproducible artifacts released for community adoption and SLURM/PBS integration

"Closing the gap between theoretical scheduling models and the realities of exascale computing: one graph, one policy, one watt at a time."

ASK ME ABOUT…

- Live walkthrough of the GAT attention weights on a real ALCF job queue
- How the uncertainty math changes your resource-allocation decision
- What it would take to pilot this on your own SLURM/PBS cluster

PROPOSED / FUTURE WORK

Joint 3-component integration & validation	Production SLURM plugin, shadow-mode at ALCF
Exascale validation (>10^5 nodes, Aurora/Frontier)	Federated multi-site scheduling
DVFS-integrated action space	Quantum–classical hybrid scheduling

RESEARCH ARTIFACTS, SCAN TO ACCESS

Source Code — GitHub
Publications — Google Scholar
Contact — adimora.kyrian@ku.edu

SELECTED REFERENCES

- [1] Adimora & Sun. HARMONIC: Uncertainty-Aware Multi-Objective Optimization for Energy-Efficient HPC Resource Management. IEEE TPDS, 2025. doi:10.1109/TPDS.2025.3610354
- [2] Adimora, Gundla & Sun. ML Approaches for Optimizing HPC Scheduling: A Comprehensive Survey. Cluster Computing 28:831, 2025. doi:10.1007/s10586-025-05521-8
- [3] Adimora & Sun. GNN-RL: An Intelligent HPC Resource Scheduler. ACM SRC Poster, SC'24.
- [4] Adimora & Sun. GNN-RL: Multi-objective HPC scheduling with graph neural networks and reinforcement learning (journal extension of [3]). *Future Generation Computer Systems*, 186:108763, 2027. doi:10.1016/j.future.2026.108763
- [5] Adimora & Sun. GATSched: Multi-Objective GATs for Energy-Efficient HPC Job Scheduling. Poster, SC'25.
- [6] Adimora & Sun. AIMS: Adaptive Intelligent Multi-Objective Scheduler Powered by Digital Twins. IEEE HPEC, 2025.

About the Researcher

Kyrian C. Adimora is a Ph.D. candidate in Computer Science at the University of Kansas, working at the intersection of graph machine learning, reinforcement learning, and sustainable systems for exascale HPC.

Acknowledgments

NSF Award #2441633 · Argonne Leadership Computing Facility (ALCF), a U.S. DOE Office of Science User Facility, for production workload traces and GEOPM power telemetry · advisor Prof. Hongyang Sun and the KU EECS dissertation committee · and the University of Kansas for supporting my doctoral studies.

Thank you for visiting, questions welcome!

This poster summarizes Ph.D. dissertation research conducted in the Department of Electrical Engineering & Computer Science, University of Kansas, advised by Prof. Hongyang Sun. Supported in part by NSF Award #2441633. Workload traces and GEOPM power telemetry provided by the Argonne Leadership Computing Facility (ALCF), a U.S. DOE Office of Science User Facility. © 2026 Kyrian C. Adimora.