

Advanced High-Performance Computing Infrastructure Systems Research Division


Performance Evaluation of the Campus-level GPUaaS with ExpEther

Background

- Our supercomputer offers high-performance computing resources such as high-end GPUs.
- However, some researchers are unable to use our supercomputer because they handle sensitive data prohibited to be stored on a shared storage.

Proposal


- We propose a campus-level GPU-as-a-Service (GPUaaS), which connects GPUs in a GPU pool to user-managed servers on a campus.
- GPUaaS dynamically allocates and deallocates GPUs to user servers.
 - Servers can interact with GPUs through a dedicated on-campus high-bandwidth network.
 - We use ExpEther to connect CPUs and GPUs over the network.



Overview of the campus-level GPUaaS.

ExpEther


- ExpEther is a virtualization technology that tunnels PCIe packets (PCIe TLP) over an Ethernet network.
- An ExpEther card encapsulates PCIe TPLs into Ethernet frames, and decapsulates them at the other end.



An example of an ExpEther connection configuration.

Evaluation results

- We evaluate the runtime overhead incurred by GPUaaS. The cable length between the server and the GPU box is approximately 200 meters.
- Round Trip Time (RTT) between two ExpEther cards is 3.64μs.
 - The runtime overhead is smaller than 6.2% except for SSD (inference).




Self-Learning Job Scheduler for Cloud Bursting

Overview

This study focuses on optimizing job scheduling for cloud bursting, aiming to minimize both cloud cost and job waiting time.

Cloud bursting shortens job waiting time by offloading jobs to external cloud resources, with operators paying for the additional usage fees. Therefore, careful selection of which jobs to offload is essential.

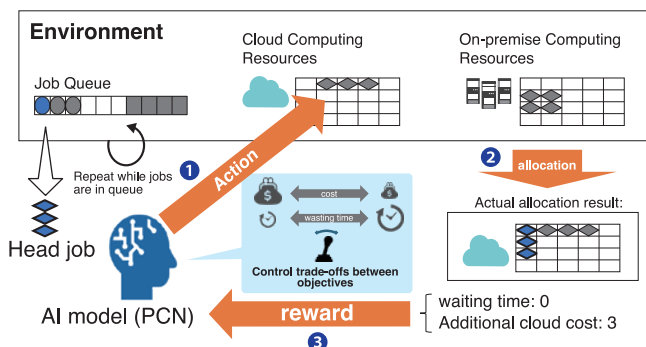
To tackle this problem, we apply **Pareto Conditioned Networks (PCN)** to job scheduling, enabling flexible control of trade-off according to operator policies.



Algorithm > Training Flow


- 1 PCN selects which job to offload considering resource availability.
- 2 Rewards are calculated based on the allocation.
- 3 Neural networks are updated using multi-objective rewards. Rewards measure decision quality.

Through ①-③, PCN learns to produce diverse and highly optimized schedules.



Result > Exploring policies

- We evaluated four scenarios defined by different numbers of jobs and nodes.
- Compared with NSGA-II and single-objective RL (DQN), the proposed method (PCN-based) achieved more diverse, higher-quality schedules.



Across all scenarios, the proposed method achieves the highest HV. This is because it yields balanced schedules between cost and waiting time, not just single-objective extremes.

PF (Pareto front): The set of non-dominated solutions across multiple objectives.
HV (Hypervolume): The area covered by the PF, measured with respect to a reference point. A larger HV indicates more diverse, higher-quality solution sets.

Future Work

- Evaluate scalability on larger clusters and longer real traces.
- Analyze and evaluate how input conditioning on features allows the neural network to control the cost-waiting-time trade-off of the resulting schedules.