Enhancing HPC Job Scheduling with Synthetic Data Generation for Reinforcement Learning - Based Schedulers

Monish Soundar Raj, UNC Charlotte
Dr. Dong Dai, CCI

UNIVERSITY OF NORTH CAROLINA CHARLOTTE

Abstract

Importance:

- Essential for optimizing job scheduling in HPC cluster management.
- High-quality job traces are crucial for HPC cluster management, especially for RL-based schedulers.

Problem:

- Data is scarce and obtaining it is challenging due to privacy concerns. RL schedulers require vast amounts of high-quality data, which is not readily available.

Addressing the Problem:

- Using ML methods to generate synthetic job traces.
- We use ML methods like GANs, VAEs, and Transformers to generate synthetic job traces, enhancing HPC scheduling and resource management with RL-based schedulers.

Method

Collected and Restructured Data:

- Efficiently gathered and reorganized data using pandas to lay the foundation for training advanced machine learning models.Some collected datasets include Theta, Theta GPU, and SC.

Implementation:

- Used the SDV library to train ML models like CTGAN, TVAE, and CopulaGAN for generating synthetic tabular job traces.
- Additionally, employed RealTabFormer and Findiff models to further enhance our synthetic data generation.

Evaluation:

- Used trained models to generate synthetic data and visualized it to evaluate effectiveness.
- This involved assessing both marginal and joint distributions to ensure accuracy in producing synthetic job trace data.

Background

Cluster - A group of interconnected computers that work together to perform tasks.

Fig 1 - Theta Cluster (Argonne Leadership Computing Facility) [2]

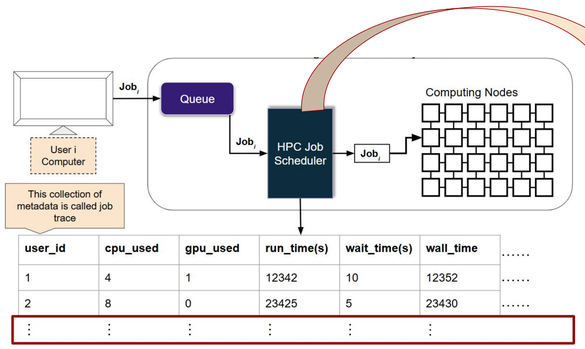

Fig 2 - Job Submission and Logging Workflow

Job Traces - Record of executed Jobs in a system.

Job	User	CPU-Num	Node-Num	Submit-Time	Wait-Time	Run-Time	Wall-Time	Job-Status
639487	0	512	8	1672529301	71	3632.92	3600	Killed
639448	0	16384	256	1672507846	7772	21654.66	21600	Pass
639135	0	262144	4096	1672174846	340972	21655.97	21600	Killed
638720	7.02963E+13	262144	4096	1671905072	632522	3648.7	3600	Killed
639463	0	8192	128	1672513950	23424	5995.18	10800	Pass
638722	7.02963E+13	262144	4096	1671905120	636224	3634.13	3600	Killed
638723	7.02963E+13	262144	4096	1671905152	639947	3669.52	3600	Killed
639488	0	8192	128	1672543325	45	5880.94	10800	Pass
639489	0	8192	128	1672549223	66	5829.52	10800	Pass
639490	0	8192	128	1672555053	46	5235.68	10800	Pass
639492	0	8192	128	1672556039	4395	57.07	10800	Pass
639497	0	8192	128	1672560303	490	5247.99	10800	Pass
639491	0	12288	192	1672555205	33	10862.96	10800	Killed
639498	0	8192	128	1672566007	175	5288.25	10800	Pass
639499	0	8192	128	1672571423	47	5236.58	10800	Pass
639500	0	8192	128	1672576677	70	5264.65	10800	Pass
639501	0	8192	128	1672581984	69	5203.48	10800	Pass
637860	0	163840	2560	1671479504	1069365	42388.64	46800	Pass
638453	0	40960	640	1671747802	801087	43245.01	43200	Killed
639502	0	8192	128	1672587200	37	5313.58	10800	Pass
638659	0	51328	802	1671868236	680633	43721.95	86400	Pass
639505	0	8192	128	1672592513	37	5293.23	10800	Pass
639508	0	8192	128	1672600471	39	38.53	10800	Pass
639509	0	512	8	1672600992	41	797.94	3600	Killed

Fig 3 - Theta Job Trace [1]

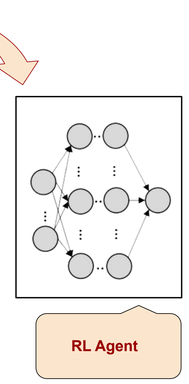

Classic Scheduler:

- Operates on fixed policies.
- May not adapt well to varying workloads or system changes.

RL Scheduler:

- Dynamically adjusts decisions through continuous learning.
- Optimizes resource utilization and minimizes wait times in real-time.

Results

CDF: Marginal Distribution

Goal: To Minimize the Discrepancy between Actual and Synthetic Data.

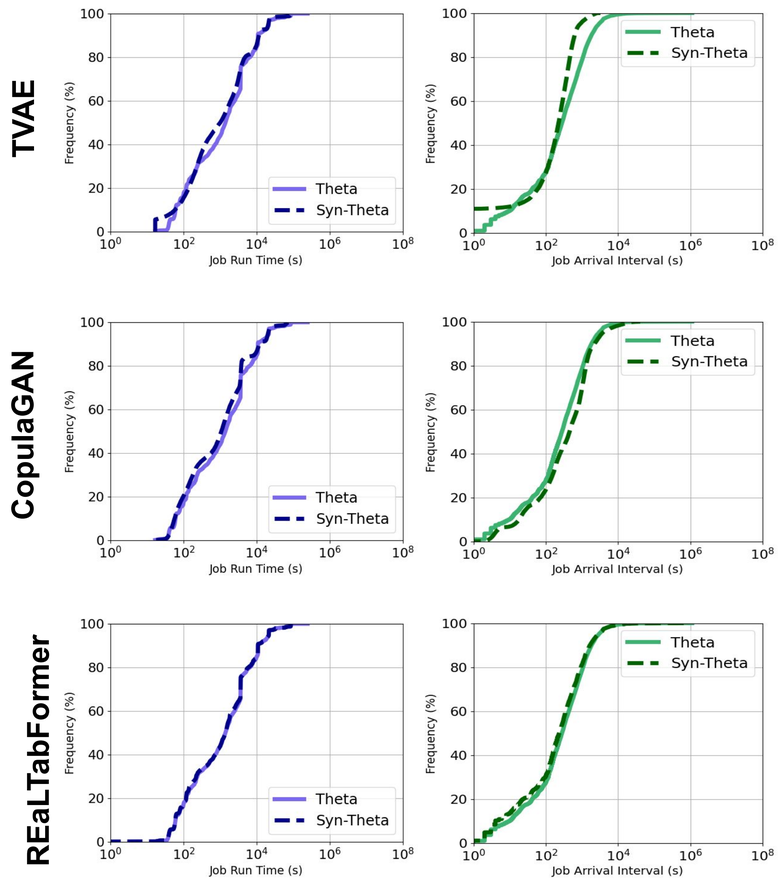

Job Run Time: CDFs comparing actual and synthetic job durations on a log scale.

Job Arrival Interval: CDFs of actual versus synthetic job arrival times.

Corr H-Maps: Joint Distribution

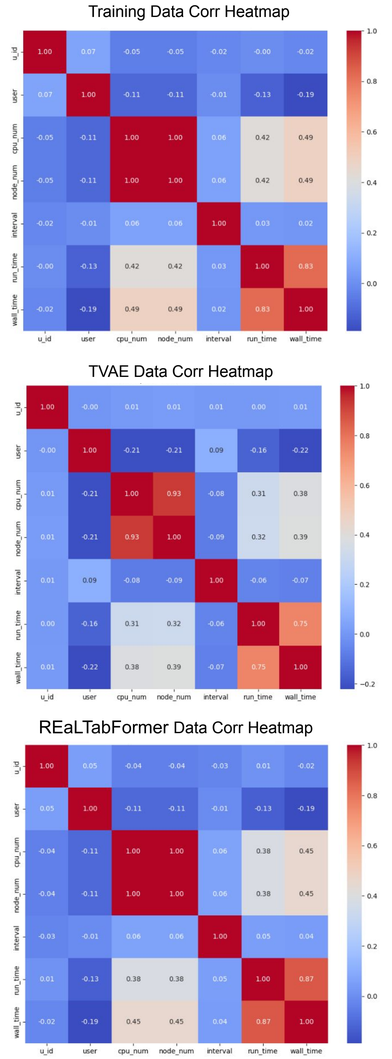

Future Work

Preliminary findings: RealTabFormer performs best in marginal CDFs and correction heatmaps for joint distribution. Our next steps involve further testing the synthetic data by training RL schedulers to evaluate performance. We aim to compare RL schedulers trained with synthetic data to those trained with original data to see if the synthetic data can achieve the same or better performance.

[1] Argonne Leadership Computing Facility. (2023). Theta Cluster Job Trace Data. Retrieved from https://reports.alcf.anl.gov/data/theta.html#DIM_JOB_COMPOSITE (Accessed Jun 2024)
[2] Argonne Leadership Computing Facility. (2023). Image of Theta Supercomputer. Retrieved from https://www.alcf.anl.gov/alcf-resources/theta (Accessed Jul 2024)