University of Glasgow | School of Computing Science

BARCLAYS

Predicting Lakehouse Performance in Clouds: An Empirical Exploration of Query Runtime Variance

James Nurdin j.nurdin.1@research.gla.ac.uk
Wei Liu wei.liu@barclays.com
Richard Mccreadie richard.mccreadie@glasgow.ac.uk
Lauritz Thamsen lauritz.thamsen@glasgow.ac.uk

Read the full paper here!

1. Problem Setting

- Runtime performance of Lakehouse queries can vary substantially across repeated executions, even under fixed conditions.
- QPP models rely on observed runtime labels, therefore run-to-run variance introduces uncertainty into model training.
- This uncertainty can increase prediction error and weaken prediction-based orchestration, including low-carbon scheduling.

Objective

Characterise runtime variance in distributed lakehouses, evaluate how variance affects QPP accuracy, and the implications this has on low-carbon scheduling.

2. Experimental Design

Lakehouse architecture:

- Open-source lakehouse stack using Trino, Apache Iceberg, and Hive Metastore.

Deployment platforms:

- Experiments across AWS, Azure, GCP, and private cloud resources.
- Services are containerised and deployed through K8s.
- Resource allocations are kept consistent across deployments.

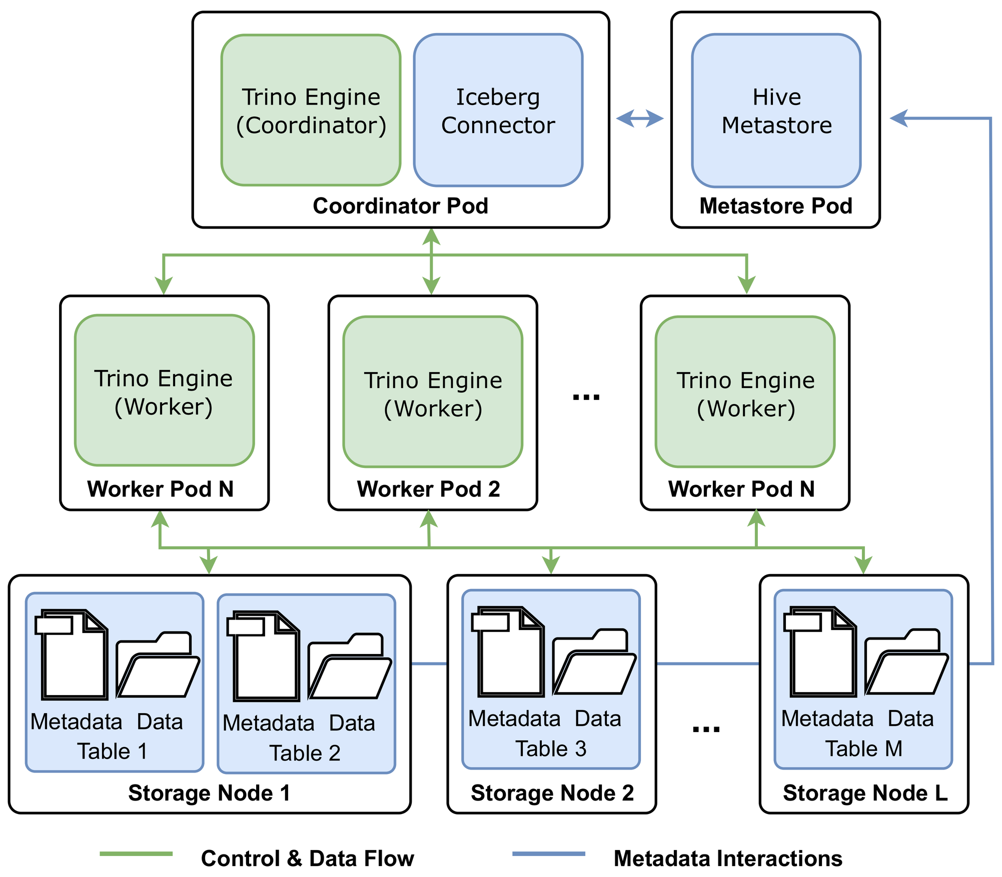

Workloads and repetitions:

- We execute analytical SQL workloads using TPC-DS at SF10, SF100, and SF1000.
- SSB and JOB used for later QPP evaluation.
- Each workload is executed repeatedly 5 times.

3. Results

(Study 1) Characterising Runtime Variance:

- Across all platforms and scales, repeated executions show substantial variance, with public-cloud median CVs of **4–9%** and P99 CVs of **20–52%.**
- Private-cloud lakehouses show the strongest variance, with average CVs above **20%** and P99 CV reaching **98.65%** at SF1000.

Variance Characterisation Across Lakehouse Deployments

Lakehouse Context		Mean Query Runtime Across Workloads (s)		Variance Across 5 Repeated Runs					
				Std. (s)			CV (%)		
Cluster Platform	TPC-DS SF	Avg.	Std.	Avg.	Med.	P99	Avg.	Med.	P99
AWS	10	5.589	0.060	0.304	0.215	1.429	7.944	6.633	31.356
Azure		3.739	0.034	0.176	0.134	0.668	8.430	7.312	24.642
GCP		6.893	0.111	0.415	0.272	2.814	7.012	5.929	22.378
Private-Cloud		3.314	0.345	0.538	0.233	4.267	20.029	17.686	76.848
AWS	100	13.067	0.086	0.538	0.257	4.065	6.212	4.260	23.967
Azure		10.579	0.045	0.407	0.231	2.887	9.069	4.707	52.269
GCP		17.612	0.355	1.123	0.507	5.690	8.319	5.458	40.658
Private-Cloud		17.596	4.295	5.104	1.587	72.227	**29.560**	**27.677**	54.990
AWS	1000	43.964	6.325	6.418	0.496	73.611	6.349	3.217	32.594
Azure		39.755	3.466	4.860	0.855	39.002	10.520	4.819	51.134
GCP		27.999	1.696	3.087	1.103	27.958	10.835	8.545	41.058
Private-Cloud		**109.432**	**12.352**	**19.033**	**11.995**	**142.283**	21.634	17.895	**98.652**

(Study 2) Impact on QPP accuracy:

- We train two QPP models, a zero-shot cost GNN and a Random Forest, using runtime labels collected from higher- and lower-variance lakehouse regimes.
- Across GNN and RF models, the largest gains appeared for analytical workloads, with MAE reduced by up to **79.51%** and P99 QError reduced by up to **88.24%**.

(Study 3) Low-Carbon Scheduling

- Higher-variance labels skew runtime predictions, producing schedules **3.1–15.0%** above Oracle emissions.
- With lower-variance predictions, overhead falls to **0.1–1.8%**, bringing schedules closer to Oracle performance.

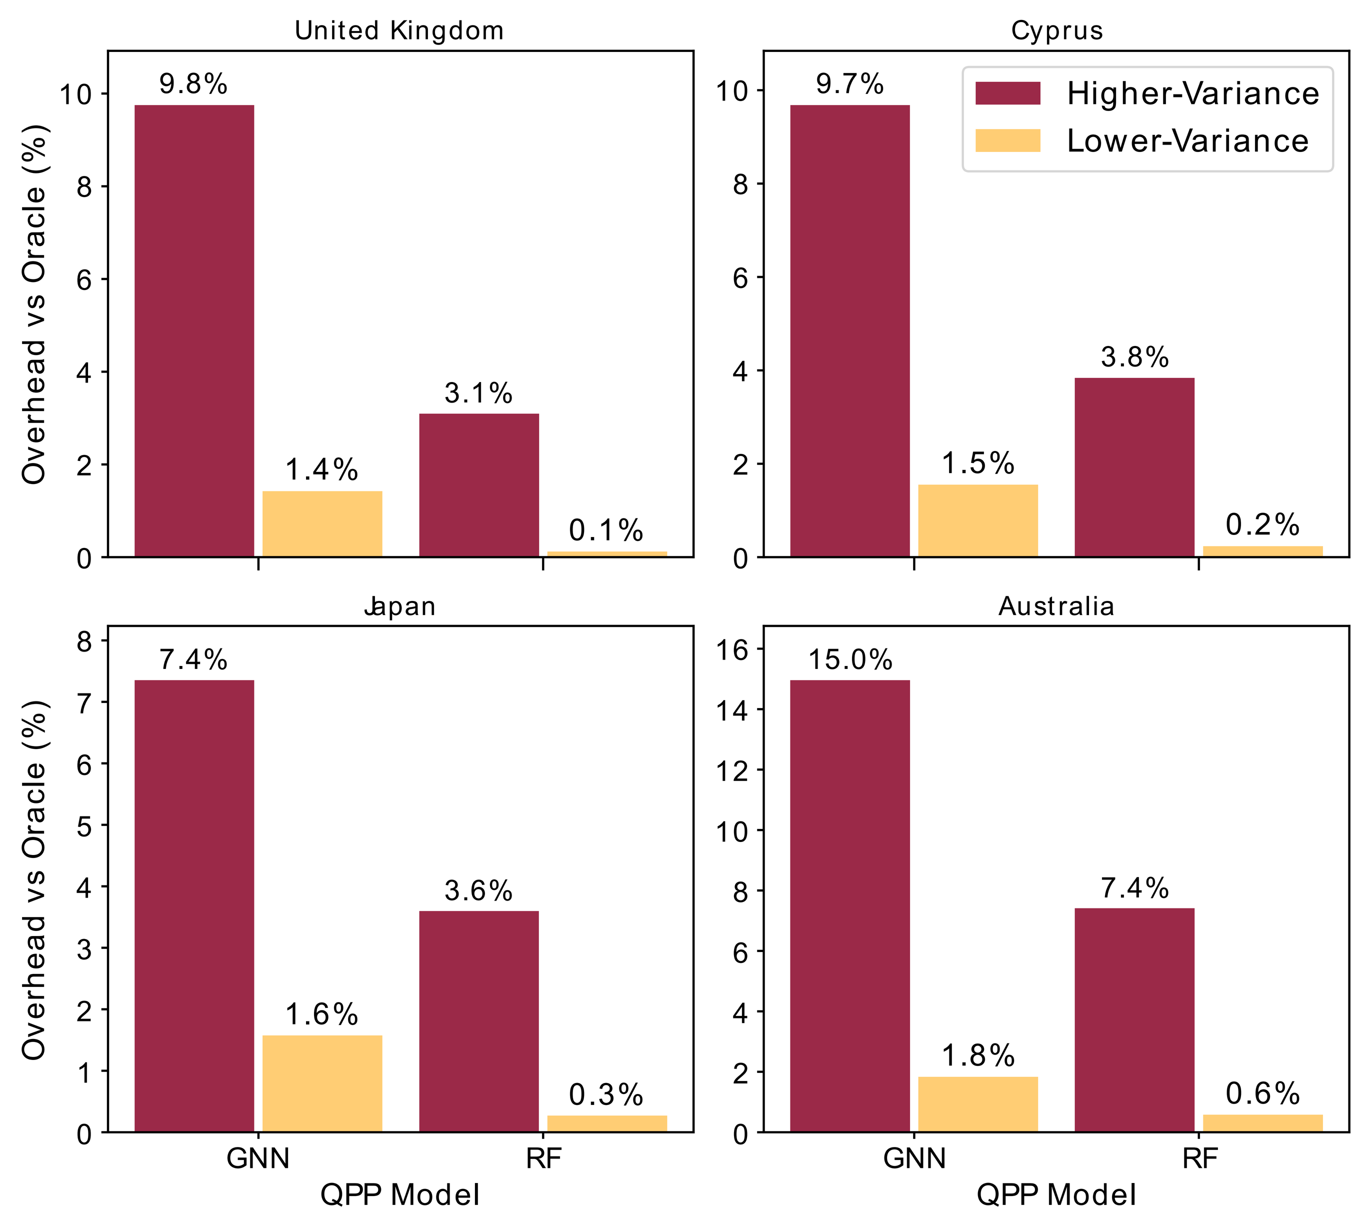

Towards uncertainty-aware scheduling

- Safety margins reduce the risk of queries overrunning into higher carbon-intensity periods, improving carbon savings across both QPP models.
- However, larger margins can overprovision execution windows and reduce these benefits.

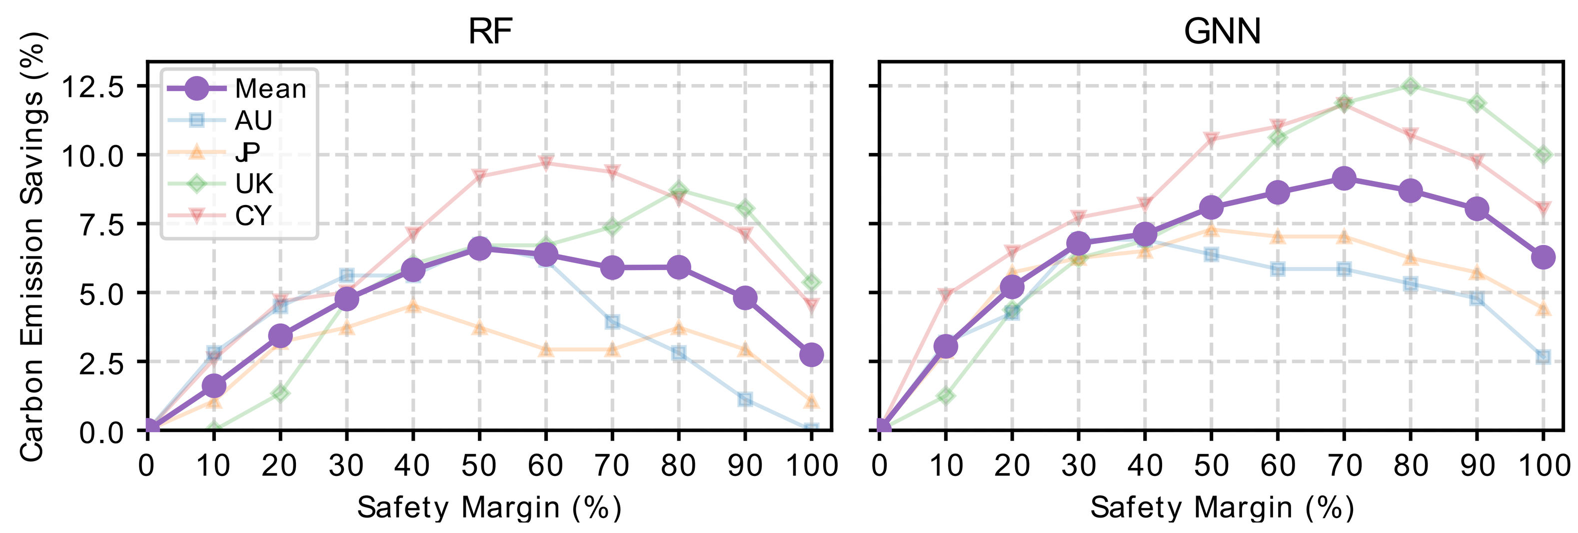

Findings

Runtime variance is not just experimental noise: it introduces label uncertainty for QPP models, increasing prediction error and weakening prediction-based orchestration decisions such as low-carbon scheduling.