National Federated Compute Services NetworkPlus

Spring Conference
26 - 27 February 2026

FAIR-Compute

A roadmap for Fair and Efficient Allocation of Federated Digital Research Infrastructure

Konstantinos (Kostas) E. Zachariadis

Professor of Financial Economics, Queen Mary University of London

Queen Mary University of London

ΕΠΙΣΕΥ ICCS

UKRI

Project Team

Project Leads

Kostas Zachariadis
Professor, School of Economics and Finance, QMUL

Co-leads

Ahmed Sayed
Senior Lecturer, Electronic Engineering and Computer Science, QMUL

Project Team Members

Angeliki Mathioudaki
PhD Candidate, Electrical Engineering and Computer Science, NTUA and ICCS, Greece

Wan Shuen Siaw
MSc Student, School of Mathematical Sciences, QMUL

Dimitrios Fotakis
Professor, Electrical Engineering and Computer Science, NTUA and ICCS, Greece

How do we solve the scarce resource allocation problem under competing objectives, especially when users act strategically or face genuine uncertainty?

UK Research Computing Ecosystem

A Bird's Eye View

The Infrastructure

National HPC: Flagship systems like ARCHER2, alongside specialized clusters like DiRAC and Cirrus.

University HPC: Local institutional clusters, such as QMUL's Apocrita, supporting faculty and teaching.

Cloud HPC: Commercial, scalable pay-per-use workloads via AWS, Azure, and GCP.

The Resource Bottleneck

Researchers submit jobs requesting specific CPU, GPU, memory, and runtime allocations.

Schedulers (e.g., Slurm, Grid Engine) act as the gatekeepers, attempting to manage queues and dispatch resources fairly.

The system's goal is to **maximize throughput**, **prevent monopolisation**, and **ensure fair access** to shared computational resources.

Private HPC vs Research Computing

Same Scheduling Problem, Different Objectives

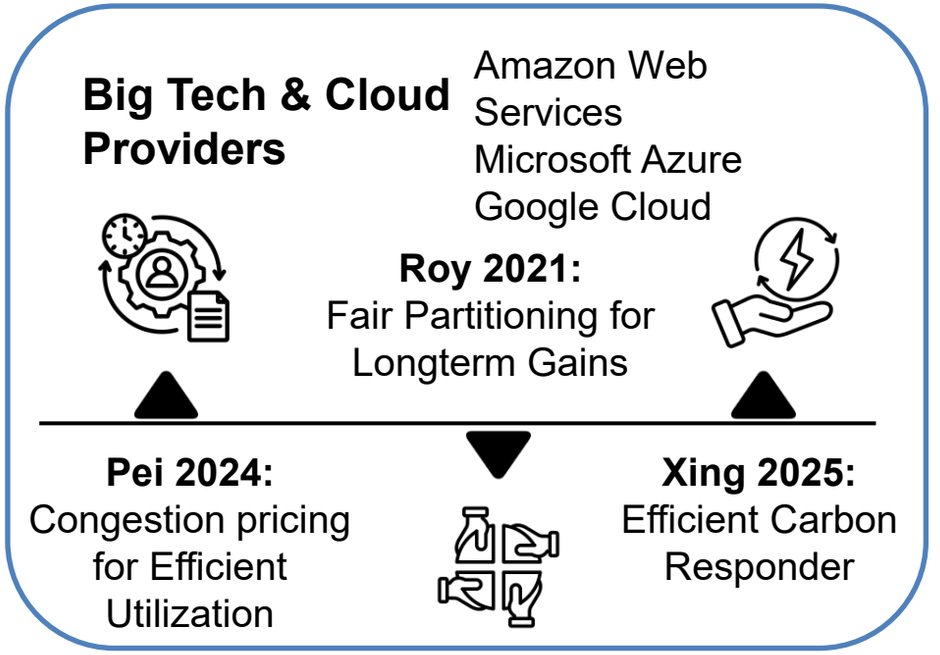

Companies offer HPC as a **paid service**.

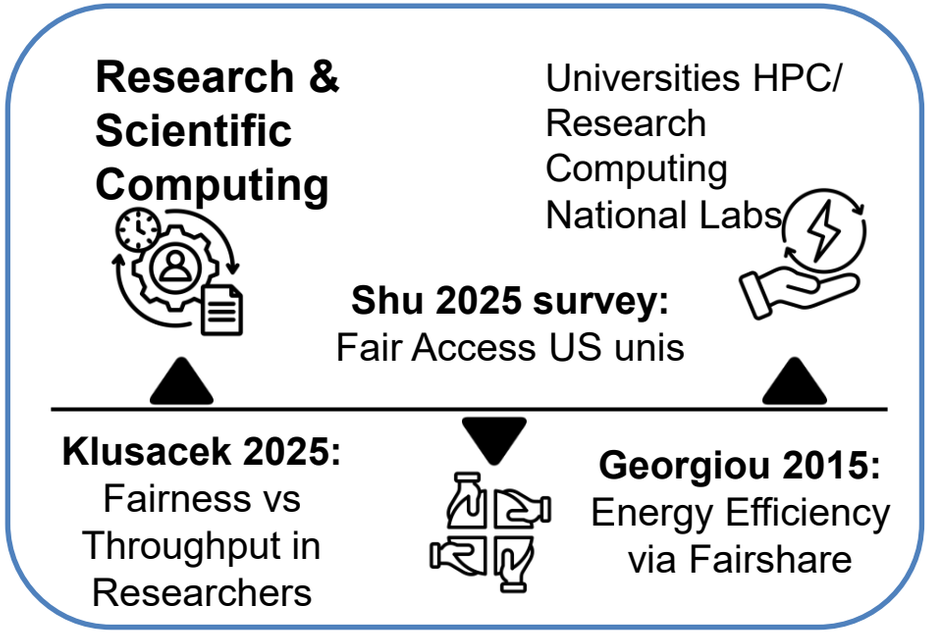

Universities offer researchers *mostly* **free access**.

The University Cell

Slurm Priority Calculation

slurm workload manager Version 25.11

w1(Fairshare)
+ w2(Age) +
w3(Job Size) +
w4(QoS) +
w5(Partition) +
…

Fairshare
Age
Job Size
QoS
Partition

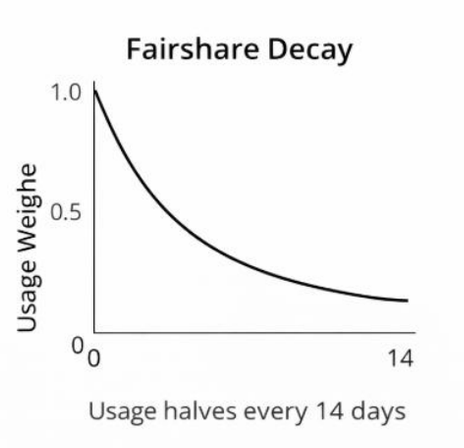

Priority Queue

Job A
Job B
Job C
Job D
…

Highest Priority

Lowest Priority

Fairshare

User A (Low use)	User B (Heavy recently)
Usage > Share	Usage > Share
Priority ↑	Priority ↓

Backfill

Lowest Priority Job may run if it fits and not delay higher priority job

Examples From a Few Universities

Fragmented Approaches

Proportional Share
e.g., Harvard

Relies on self-reported sizes.

Leads to **strategic** banking where heavy users consume account shares.

Property Rights
e.g., MIT

Uses hardware ownership shares and preemption.

Results in groups with small shares being **preempted** more often.

Funding Tiers
e.g., Cambridge

Allocates based on credits and paid tiers.

The issue of "**wealth**" influencing priority.

Is any of these suited for a national, federated system?

The Limits of Current HPC Scheduling

Scarce Resource Allocation Under Conflicting Objectives

The Behavioral Limit
Gaming

The Reality: Schedulers rely on heuristics and users to self-report their needs.

The Impact: Users act "strategically" (e.g., padding walltime), leading to artificial scarcity and idle nodes.

Our Goal: More Incentive-compatible mechanisms that reward (near) truthful reporting.

The Structural Limit
The Federation Wall

The Reality: Fairshare heuristics are "tailored" for single-institution budgets.

The Impact: They may be insufficient when cross-charging across different institutions.

Our Goal: Equitable, scalable allocation rules across the National Federated Compute Service.

The Paradigm Shift
Engineering2Economics

The Reality: Treating scheduling purely as an engineering problem.

The Impact: Human strategic behavior treated as a "bug" and not a feature of the system.

Our Goal: Shift from reactive hardware optimization to proactive (socio)economic mechanism design.

Bridging Theory & Practice

Answering the NFCS Roadmap Challenge

Stakeholders

Process
We are engaging Research Computing Admins/RSEs, PIs/Researchers, and funders, via surveys/interviews to capture real-world constraints.

Beneficiaries
Universities and researchers seeking equitable access through scalable allocation models (e.g., compute credits and grant accounting) rather than strict commercial pricing.

Epistemic Approach

The Economics
We apply mechanism design modelling to capture key desiderata of a federated research computing socioeconomic system.

The Engineering
We simulate different models under realistic workloads and compare performance against used heuristics.

Outcome

The Present
Deliver critical description of current approaches to resource allocation at the university and national level.

The Future
A roadmap of theoretically robust and simulation supported suite of policy recommendations for fair allocation in a federated, multi-institutional level.

A Positive and Normative Abstraction

Framing HPC as a Mechanism Design Problem

Resource allocation
With strategic users

Set Partition Problem (SPP)
Allocating bundles of CPUs, GPUs, and storage.

Users may misreport their true resource needs either strategically (to gain priority) or due to *genuine uncertainty*.

Mechanisms
Map reported preferences to allocations and "transfers"

Without money
Priority queues, Delay Externalities.

Virtual currency
Credits, tokens, **ColdFront.**

With money
Auctions, pricing mechanisms.

Metrics
Assess the behavior each mechanism induces

Welfare
Total value of completed jobs.

Fairness
Balanced access across users.

Shapley Value
Allocation based on marginal contribution.

Each mechanism may face an efficiency-fairness trade-off.

Thank you

Stakeholder Survey

https://forms.office.com/e/3Qb0WuNT6E?origin=lprLink