



Enabling quantum-HPC Convergence for Future Scientific Workflows

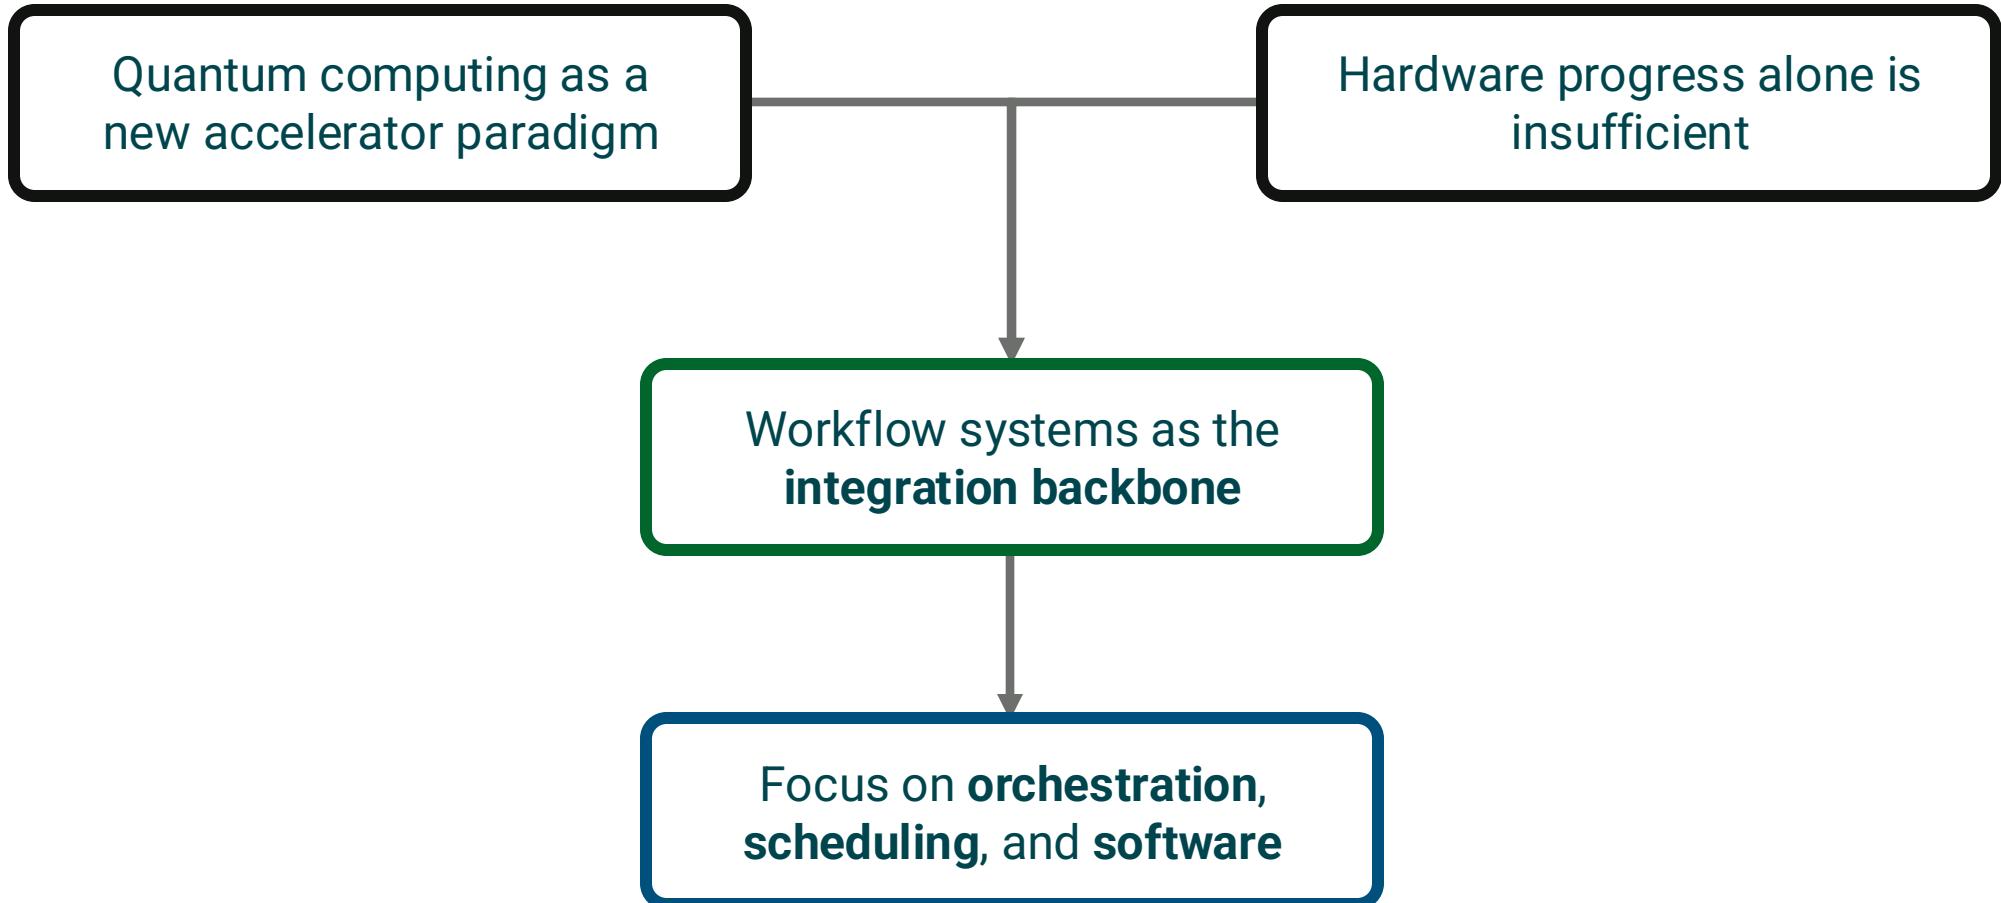
RAFAEL FERREIRA DA SILVA
SECTION HEAD (INTERIM)
COMPUTER SCIENCE AND MATHEMATICS DIVISION



ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY



Synopsis



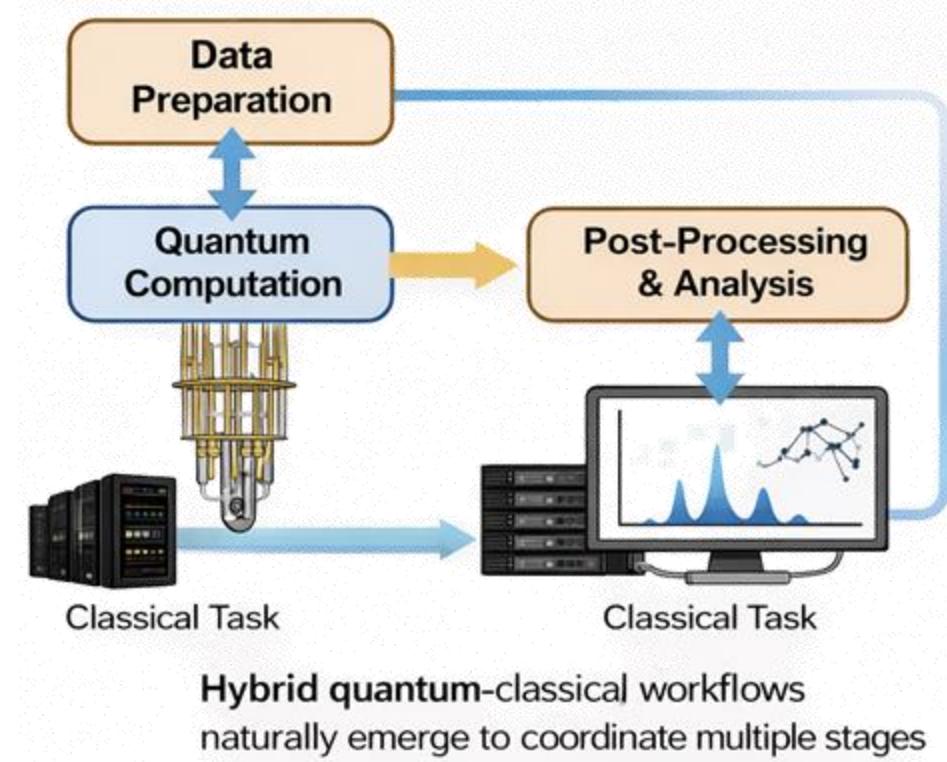
Why This Is Fundamentally a Workflow Problem

Hybrid quantum-classical execution is unavoidable

Applications span multiple stages and resources

Coordination dominates time-to-solution

Workflows already address these challenges



Most realistic quantum application today is **hybrid**

The challenge is managing this at **scale**

Lessons from GPU-HPC Integration

GPUs were once **disruptive accelerators**

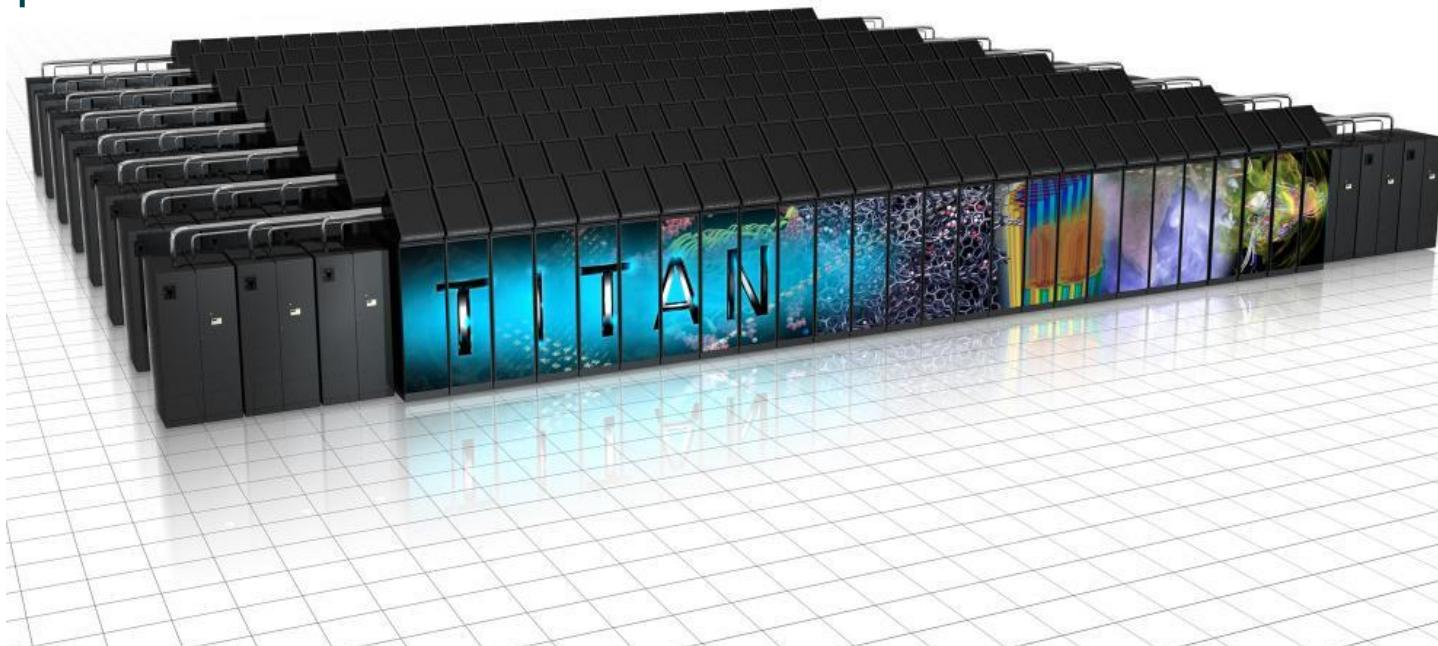
Abstractions enabled widespread adoption

Runtimes, APIs, and schedulers mattered

Quantum-HPC integration can follow this path

"The Cray XK7 system contains 18,688 nodes, with each holding a 16-core AMD Opteron 6274 processor and an **NVIDIA Tesla K20 graphics processing unit (GPU) accelerator**"

<https://www.ornl.gov/news/ornl-debuts-titan-supercomputer>



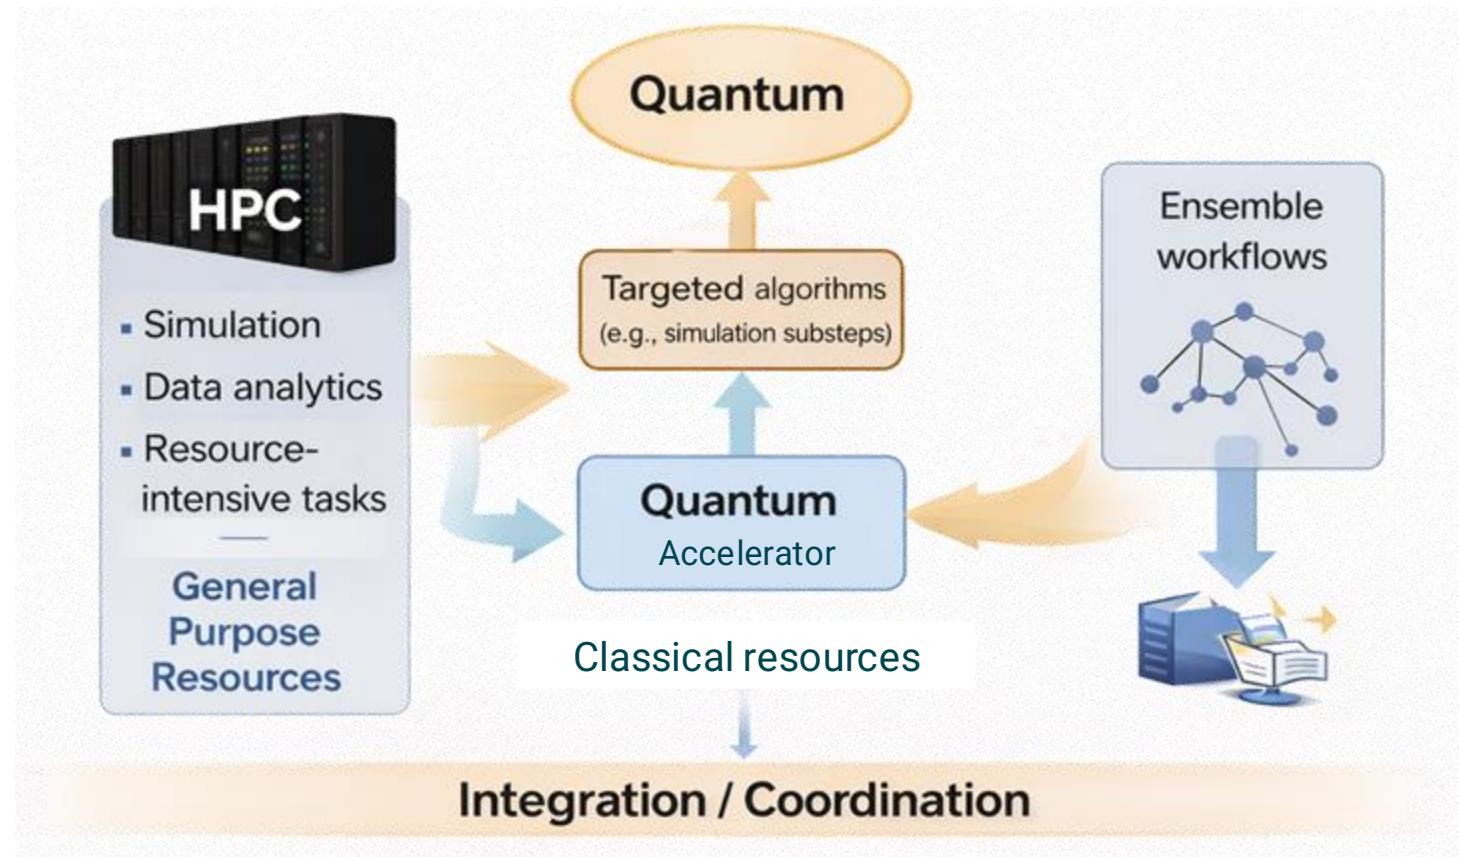
Quantum as an Accelerator, Not a Replacement

QPUs excel at narrow tasks

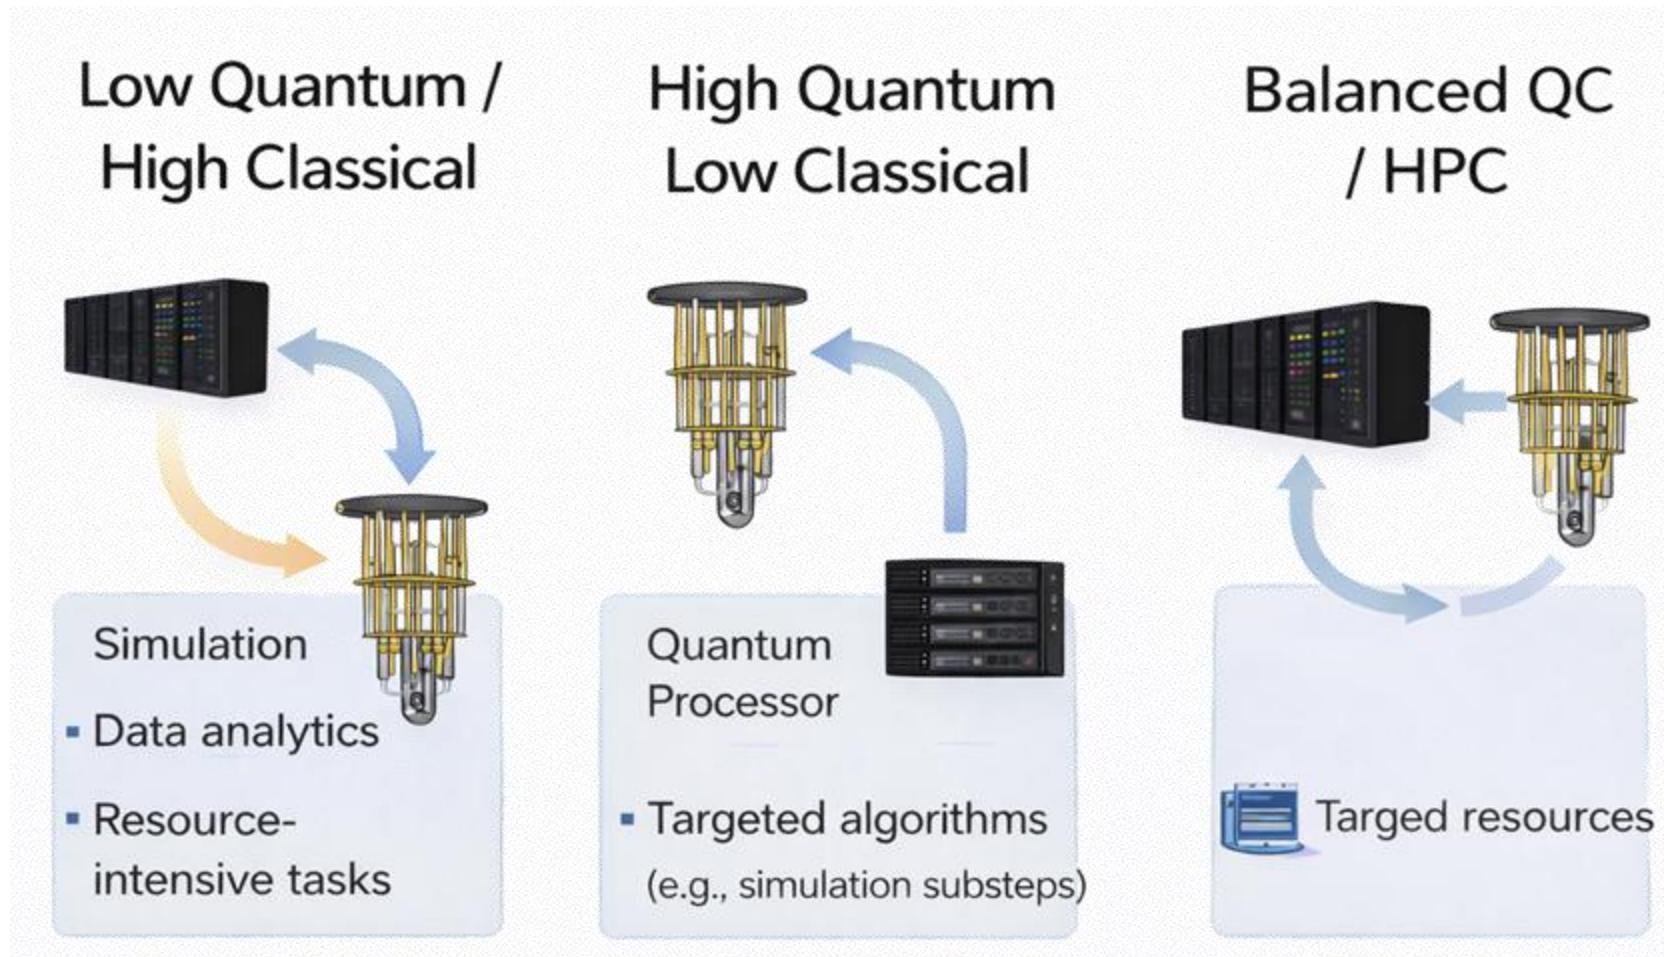
Classical HPC remains dominant

Hybrid execution will be the norm

Integration matters more than disruption



Hybrid Application Patterns



There is a need for transparent support

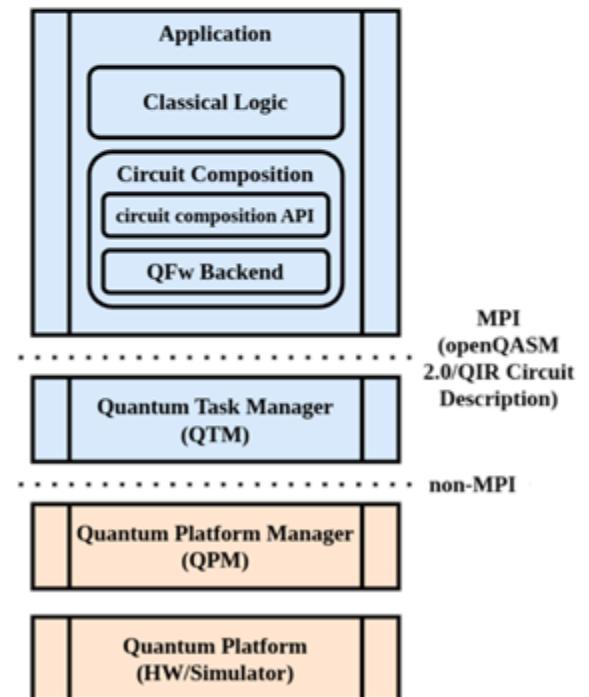
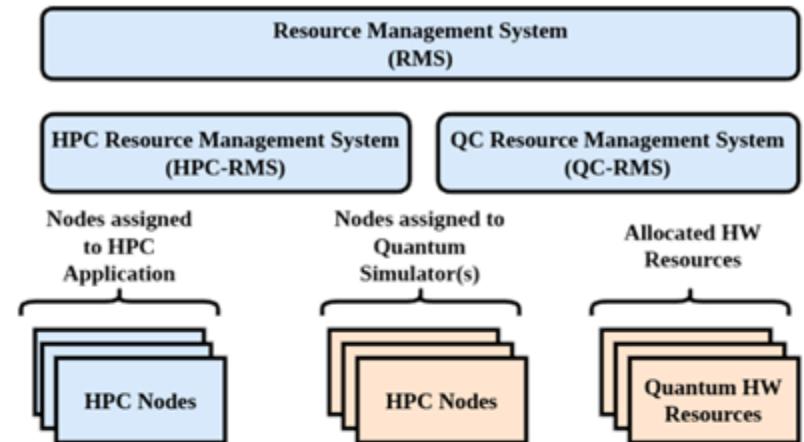
Execution and Allocation Models in Hybrid Workflows

Simultaneous allocation of HPC and QPU resources

Interleaved allocation for improved utilization

Workflow engines encode dependencies

Schedulers alone are insufficient



Beck, T., Baroni, A., Bennink, R., Buchs, G., Coello Perez, E. A., Eisenbach, M., Ferreira da Silva, R., Gopalakrishnan Meena, M., Gottiparthi, K., Groszkowski, P., Humble, T. S., Landfield, R., Maheshwari, K., Oral, S., Sandoval, M. A., Shehata, A., Suh, I.-S., & Zimmer, C. (2024). **Integrating Quantum Computing Resources into Scientific HPC Ecosystems**. Future Generation Computer Systems, 161, 11–25.



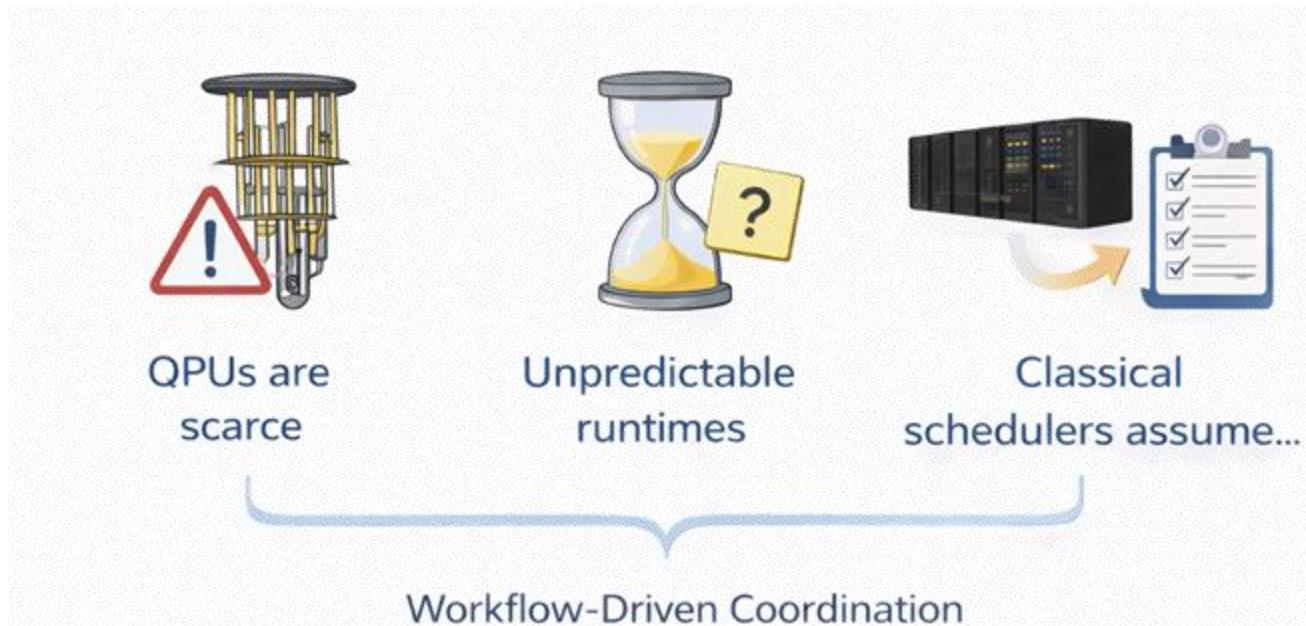
Why Scheduling Gets Hard

Quantum devices are scarce and exclusive

Execution times are unpredictable

Traditional HPC assumptions break down

Idle resources are inefficient



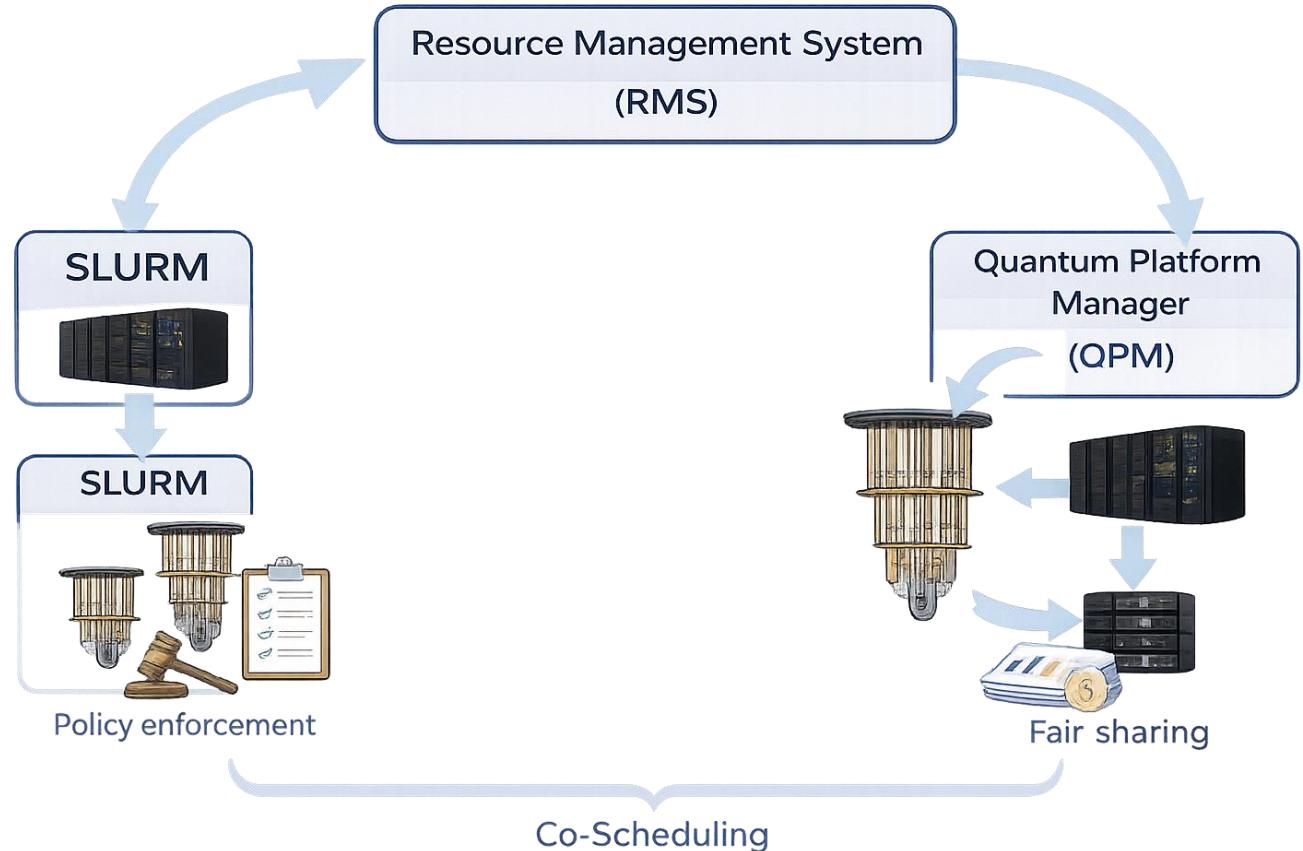
Toward Unified Resource Management

Treat QPUs as managed resources

Integrate with existing schedulers

Enable coordinated reservations

Avoid siloed quantum queues



Shehata, A., Groszkowski, P., Naughton, T., Gopalakrishnan Meena, M., Wong, E., Claudio, D., Ferreira da Silva, R., & Beck, T. (2026). **Bridging Paradigms: Designing for HPC-Quantum Convergence**. Future Generation Computer Systems, 174. <https://doi.org/10.1016/j.future.2025.107980>

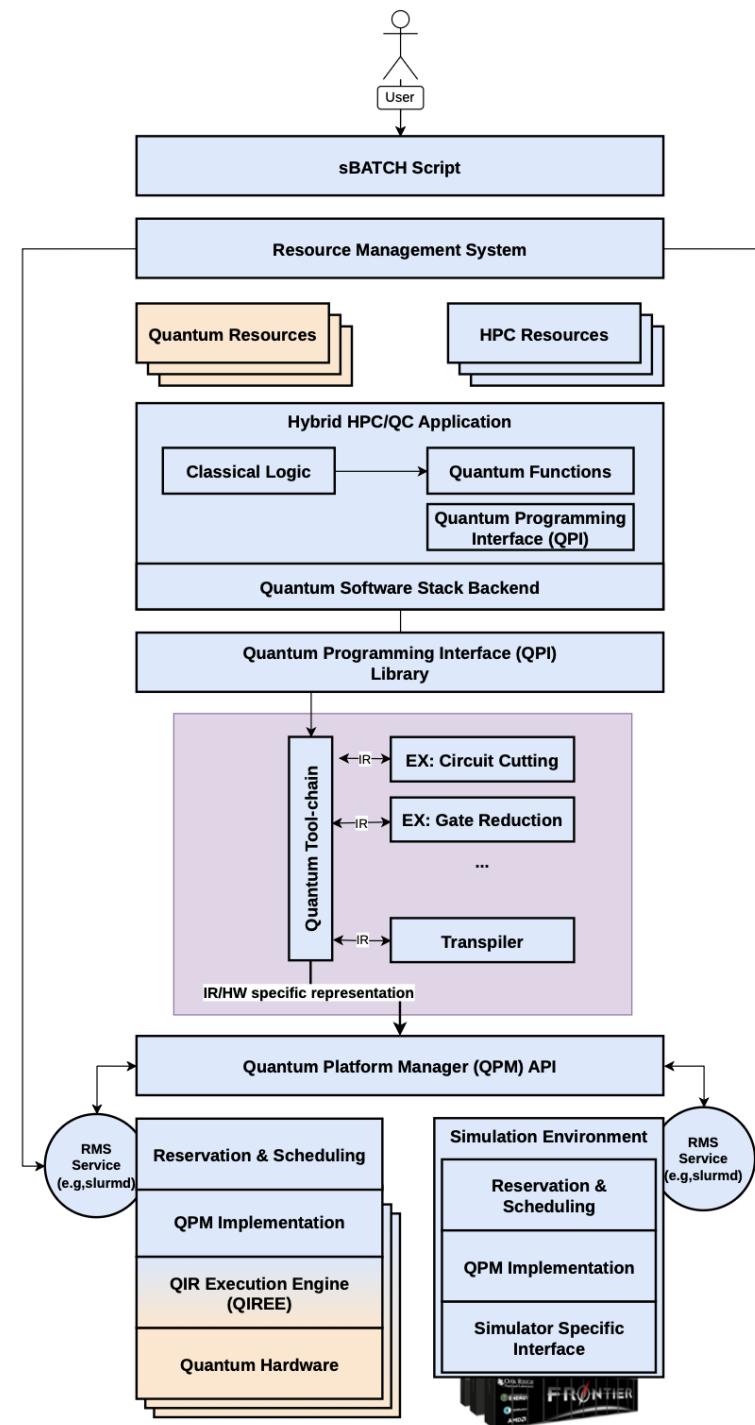


QC/HPC Software Architecture

Shehata, A., Groszkowski, P., Naughton, T., Gopalakrishnan, Meena, M., Wong, E., Claudino, D., Ferreira da Silva, R., & Beck, T. (2026). Bridging Paradigms: Designing for HPC-Quantum Convergence. *Future Generation Computer Systems*, 174. <https://doi.org/10.1016/j.future.2025.107980>



Application-facing programming interfaces
Hardware-facing platform managers
Clear separation of concerns
Hardware-agnostic workflows



From Scientific Campaigns to User Adoption

Scientific workloads are campaigns, not single runs

Ensembles, sweeps, calibration, and analysis loops

Reproducibility and provenance are essential

Most users are **not quantum experts**

Real scientific use of quantum computing is not about running one circuit once

It's about campaigns: ensembles of runs, parameter sweeps, iterative refinement, and post-analysis

Open Research Challenges

Co-scheduling across heterogeneous resources

Coordinating classical HPC resources and scarce quantum devices requires new scheduling models that go beyond traditional queue-based assumptions

Quality-of-service guarantees for hybrid execution

Hybrid workflows need predictable access to quantum resources to avoid idle classical allocations and excessive end-to-end latency

Fault tolerance and resilience across hybrid stages

Failures can occur at many layers (e.g., classical tasks, quantum execution, networks, or middleware) requiring workflow-level recovery strategies

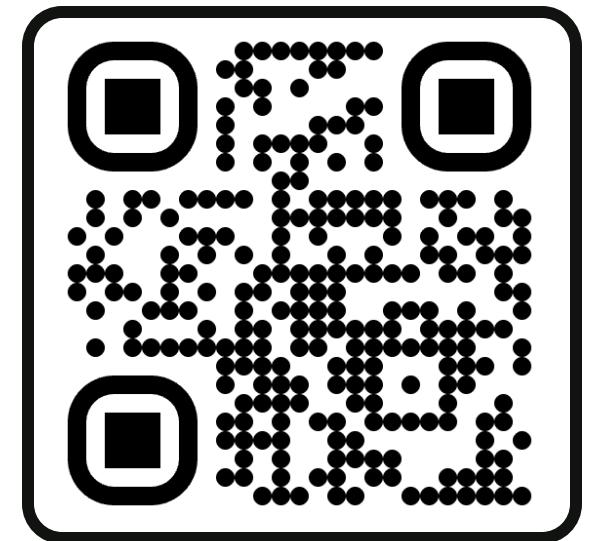
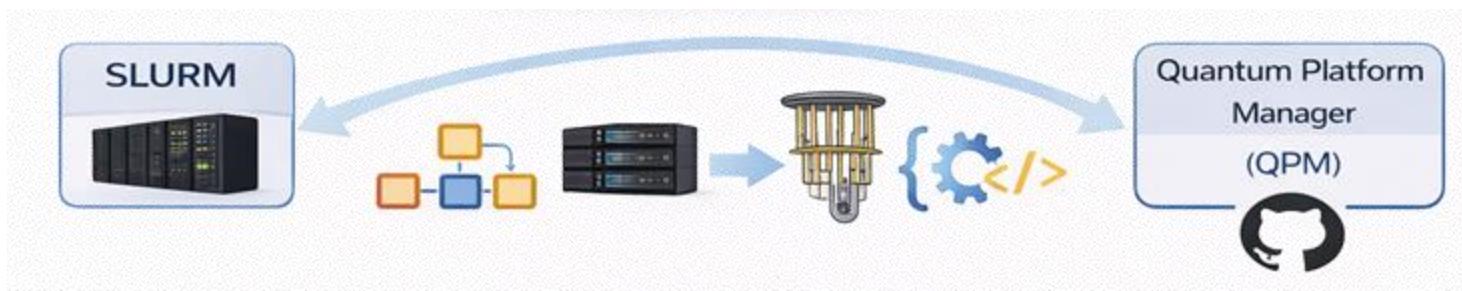
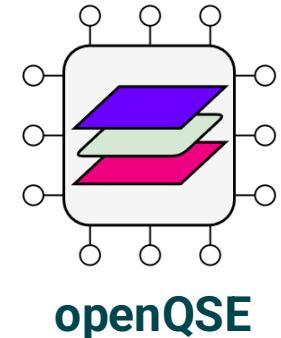
Scalable workflow execution for quantum-enhanced campaigns

Supporting ensembles, parameter sweeps, and iterative algorithms at scale remains an open challenge as quantum resources are introduced

OpenQSE

Join the Open Quantum Software Ecosystem

- Workflow and HPC-focused
- Middleware stack, APIs, demos
- **Contributions welcome**



Thank You!

Questions?



This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

Rafael Ferreira da Silva

Section Head (Interim) / Senior Research Scientist
Data and AI Systems
Computer Science and Mathematics Division
<https://rafaelsilva.com>

