AQADOC - Journée Scientifique

October 9, 2025

Summary of the afternoon discussion panel

Disclaimer: *This document is based on a transcription of recorded discussions held during the event. While care has been taken to ensure the accuracy and integrity of the content, the transcription is not a verbatim record. Certain passages have been edited, condensed, or omitted to improve clarity and readability, without altering the meaning or intent of the participants' remarks.*

Question 1: Do you think some of the results achieved so far in the AQADOC project could be useful to some use-cases? If so, in which domain of application?

a. Solving partial differential equations in free-flow simulations

Solving PDEs using distributed quantum computing could be interesting for use-cases in the oil and gas industry, in particular in simulating free-flow situations. Although the hardware (i.e. modular quantum computing) is not ready today, it is interesting to learn more about the methods developed in this project, and work at implementing them in the context of flow-simulations.

b. From Parallelization Challenges to Quantum Opportunities

Mohammed Hibti from EDF shared that in his work on reliability assessments, classical parallelization approaches had reached their limits as they produce complex, inefficient structures when scaling up computations. He found that some of the distributed computing concepts presented in AQADOC could offer more elegant and scalable solutions to these challenges and that discussions with the AQADOC team are already ongoing to explore these possibilities.

While distributed quantum algorithms are still in their early days, identifying relevant industrial use cases now will be crucial to ensure that, once quantum systems can be networked, the community is ready to move straight into impactful applications.

c. Translating Complex Problems into Quantum Algorithms

Kilian Müller, from Welinq's hardware team, asked how AQADOC researchers actually map complex, domain-specific problems onto quantum algorithms. Do they design entirely new algorithms, or adapt existing ones that already show a quantum advantage?

Riccardo Mengoni from Welinq's algorithm team explained that their initial motivation came from the probabilistic nature of the systems they study (for example, large process trees and Markov graphs). They first experimented with established quantum algorithms, such as Grover's search algorithm, but ran into practical limits: it required prior knowledge of the number of solutions and demanded too many qubits for large problems. This experience encouraged them to look for alternative approaches and refine how quantum routines could be applied to real-world problems.

d. Networking Quantum Computers for Matrix Inversion

Finally, a PhD student proposed a use case idea involving quantum matrix inversion (using the HHL algorithm). The concept would involve connecting two quantum computers so that one prepares a complex quantum state and the other performs the computation. This would likely find applications such as error detection or real-time data processing in high-throughput environments.

The moderator welcomed this idea, highlighting that all such suggestions will be collected and explored further as AQADOC continues to run, and expand on its projects.

Question 2: What are the advantages of distributed vs local quantum operations (i.e. working with networked devices or within a single machine)? Same question for heterogeneous vs homogeneous architecture.

a. Balancing Fidelity and Scalability

Riccardo Mengoni (Welinq, Algorithm team) explained that, at present, non-local two-qubit gates performed between separate quantum processors would inevitably suffer from lower fidelity compared to local gates executed within a single quantum processing unit (QPU). This is due to imperfections in the entangled pairs used to connect the processors and the error-prone protocols involved in maintaining those links.

The goal is therefore to minimize the use of non-local operations, since each additional step in the protocol increases the probability of error. However, scaling up quantum computers will inevitably require the use of distributed architectures, which is why projects such as AQADOC are already running today.

In the short term, researchers can rely on error mitigation techniques to compensate for lower fidelity, while in the long term, fault-tolerant architectures and error-correcting codes could help make distributed gates more reliable.

b. Exploring Heterogeneous Interconnections

The discussion then shifted toward the hardware challenges of connecting heterogeneous quantum systems, for instance, linking superconducting and neutral atom processors. One participant pointed out that efficient transduction, which converts information between different physical carriers (e.g. microwave to optical photons), remains a major bottleneck for such heterogeneous interconnection.

Another participant suggested that such hybrid connections might need to rely on offline data preparation or localized computation at each node, rather than relying on continuous real-time communication, since each transmission currently carries a probabilistic overhead. The trade-off, as noted, is between centralizing all computation in a single, error-prone quantum processor versus distributing the workload and managing interconnection losses.

The most optimal method to interconnect such devices is still being investigated, but it is important for such discussion to take place today.

c. Time Scales and Task Partitioning Across Platforms

Building on this, a new point was raised: different quantum technologies operate on vastly different time scales. Superconducting qubits can perform operations at nanosecond speeds, while neutral atom systems typically operate much more slowly but with higher coherence and fidelity. This mismatch in timing could be either a challenge or an opportunity, depending on how computations are distributed.

One proposed hybrid architecture aims to exploit this difference: Fast operations could be executed on the superconducting side, while high-fidelity or memory-intensive tasks could be performed on the neutral atom side.

However, designing algorithms or workflows that can intelligently partition problems between such systems remains non-trivial.

d. Hybrid Error Correction Across Architectures

The conversation then turned toward error correction. A participant suggested that the motivation behind hybrid architectures might not only be scalability but also the possibility of distributing error-correcting codes across different platforms with each type of qubit technology optimized for a specific kind of error.

This raised an open question: could a distributed error correction scheme be designed to leverage the strengths of multiple architectures, rather than enforcing one-size-fits-all codes? While no definitive answer emerged, participants agreed that this is an exciting direction for future exploration.

e. Beyond Quantum–Quantum Links: Quantum–Sensor Integration

The discussion concluded with an analogy from the oil and gas industry, illustrating how computation and sensing could interact dynamically. In that example, a central computer models underground structures, while sensors deployed in the field collect live data that continuously update the model, guiding subsequent operations.

Applied to quantum technologies, this could inspire hybrid quantum systems that integrate computers and sensors, where quantum processors, sensors, and classical control systems form real-time feedback loops. Such systems could play a transformative role not only in computation but also in quantum-enhanced sensing, metrology, and control.

Question 3: As a follow up to the previous question, what are the challenges and limits associated with connecting multiple QPU together?

Participants highlighted that, in theory, adding more QPUs should increase computational power, provided the quantum links between them are perfect. In practice, however, link imperfections and connectivity complexity quickly become limiting factors. As the network grows, maintaining efficient communication between every processor becomes increasingly difficult, potentially offsetting the gains from adding more QPUs.

From a hardware perspective, Welinq explained that their efforts aim to minimize the performance cost of interconnection, while improvements in software are also expected to help manage errors and optimize performance. Others in the discussion noted that the network architecture (e.g., linear, circular, or hybrid topologies) and entanglement-swapping techniques could help reduce the number of direct links required between processors.

The exchange concluded on a pragmatic note: while large-scale interconnection is a long-term goal, current efforts should focus on smaller, meaningful demonstrations where connecting a few QPUs already provides added computational value.

Question 4: What would be a realistic quantum advantage timeline? Why do you think so?

For context, quantum advantage here is defined as when a quantum computer will be able to solve a real-world problem that no classical computer can handle.

Participants generally agreed that the timeline remains highly uncertain. Some highlighted that progress in classical computing continues in parallel, constantly pushing back the threshold for what counts as a genuine advantage. Others emphasized that the definition itself matters: solving a problem faster than a classical computer is one form of advantage, but solving something classically impossible is another, much higher bar.

There was also debate about which architecture might first reach quantum advantage. While some expect progress on monolithic systems (single large processors), others suggested that distributed or hybrid approaches, combining classical and quantum systems, may reach practical milestones sooner.

Several participants expressed cautious optimism, noting how perceptions have evolved with some sceptical figures in the field having become more open to the idea that meaningful quantum advantage could occur within our lifetime. Still, many acknowledged that even when the technology matures, deploying it for real-world problems will require significant engineering, integration, and operational advances, meaning the path to widespread advantage is likely to be gradual.

Question 5: How does intra-QPU connectivity influence the need for entanglement distribution in distributed quantum computing?

a. Trade-offs in keeping everything within a single device

The discussion began with a clear recognition that intra-QPU connectivity directly affects whether qubits need to be entangled across multiple QPUs. Jesus from Le Lab Quantique noted, "If my neutral atom QPU can move qubits almost anywhere inside the module, I might not even need to go distributed for certain algorithms." Yet, other participants were quick to point out that physical qubit movement is not free of errors. One panelist remarked, "Moving qubits introduces errors that can be just as damaging as having lower connectivity. You might gain flexibility but lose fidelity."

b. The specific case of superconducting qubit platforms

Superconducting QPUs, with their inherently lower connectivity, sparked debate. Panelists observed that their low noise levels can sometimes make up for limited intra-QPU connections, especially for tasks where entanglement distribution can be managed with high fidelity. The group collectively noted that in a modular architecture, even a highly connected QPU can benefit from inter-QPU links when scaling up beyond a single module. There was consensus that connectivity reduces but does not eliminate the need for inter-QPU entanglement, particularly when considering error propagation in large-scale fault-tolerant operations. One example mentioned was probabilistic inter-module entanglement in superconducting systems, where even near-perfect intra-QPU gates cannot fully compensate for the uncertainty of inter-QPU links.

Question 6: How prepared is the current workforce to develop and deploy practical quantum computing applications? How do France, and the EU in general, compare to the rest of the world on that front?

a. Theory vs practical experience

The discussion highlighted that even as hardware capabilities advance, workforce readiness remains a bottleneck. Panelists observed that European and French academic programs provide strong theoretical foundations in quantum mechanics and mathematics but often lack practical exposure to real-world quantum computing systems. Short courses, industry internships, and hands-on workshops were suggested as essential to bridge this gap.

b. Retaining talents

The panel also discussed talent retention challenges. Students with strong quantitative backgrounds may exit quantum tracks early due to limited career visibility or insufficient exposure to applied projects. Panelists noted that skills in fast matrix computation, error correction algorithms, and hybrid classical-quantum workflows are highly transferable (for example in AI, or computational science in general), which could help motivate more students to join such programs designed with quantum technologies in mind and hopefully retain them.

A recurring recommendation was to integrate problem-focused projects and access to experimental platforms within educational programs, ensuring that graduates are ready to contribute immediately to industrial-scale quantum projects.

Question 7: Beyond raw computational speed, what factors would motivate the adoption of quantum computing technologies by HPC or data centers?

Deployment decisions hinge not only on qubit count or speed but on a balanced evaluation of performance, cost, and workforce capability. Panelists emphasized that solution quality, operational reliability, and talent availability are critical for achieving meaningful impact. For instance, France, while strong in hardware, faces gaps in software, middleware, and practical training, highlighting the need for integrated ecosystem development. This is one of the missions of the AQADOC consortium.

Participants agreed that strategic deployment requires weighing multiple factors: the capabilities of each QPU, the efficiency of distributed architectures, economic feasibility, and readiness of personnel. Realistic adoption scenarios were considered, such as hybrid QPU clusters for specific simulation tasks or interconnecting heterogeneous QPUs to exploit complementary strengths. The consensus was that long-term competitiveness in quantum computing depends on harmonizing technical performance with human and economic resources.