

Original Article

# Scalable Quantum Software Frameworks for Cloud-Based Quantum Computing

Ni Putu Windayanti

Faculty of Technology, Bandung Institute of Technology, Indonesia

Received Date: 04 December 2025

Revised Date: 18 December 2025

Accepted Date: 01 January 2026

## Abstract

Quantum computing is moving quickly, but it's still challenging to use in real life since there isn't enough quantum gear and software that can be used in large groups. Cloud-based quantum computing has become a promising way to make quantum workloads more accessible and bigger. This research looks at quantum software frameworks that can be scaled up and work in cloud-based quantum computing systems. It looks at their architectural designs, functional capabilities, techniques for integration, and problems with scaling. We look at major frameworks including IBM's Qiskit, Google's Cirq, Microsoft's Q#, and Amazon Braket in terms of how modular they are, how abstract they are, how flexible their programming is, and how well they operate. The paper also talks about middleware and hybrid quantum-classical systems that make it easier for quantum hardware backends and classical infrastructure to work together through cloud platforms.

We talk about how containerisation, microservices, and REST APIs can let quantum services grow to meet the needs of different customers and applications. We also talk about the problems that come up when making software stacks that work on any hardware but are optimised for performance. These problems include noise, qubit decoherence, limited quantum volume, and classical simulation limits. The study points out architectural constraints in present frameworks and suggests ways to make them more scalable, such as AI-driven compiler optimisations, resource-aware schedulers, and hybrid orchestration. Our research shows that current frameworks are a good starting point for development and testing, but to really scale up, they need to work with next-generation cloud infrastructure, edge computing, and real-time quantum error prevention. This work gives a plan for making strong, scalable software systems that can handle the whole lifetime of quantum algorithms, from design and simulation to running them on real quantum hardware in the cloud.

## Keywords

Quantum Computing, Scalable Software Frameworks, Cloud-Based Quantum Platforms, Qiskit, Cirq, Azure Quantum, Amazon Braket, Quantum Software Scalability, Hybrid Quantum-Classical Workflows, Quantum Programming Languages, Variational Algorithms, Quantum Machine Learning, TensorFlow Quantum, PennyLane, Quantum Interoperability, Quantum Cloud Infrastructure, Quantum Orchestration, OpenQASM, QIR (Quantum Intermediate Representation), NISQ (Noisy Intermediate-Scale Quantum)

## Introduction

Quantum computing is one of the biggest changes to the world of computer technology in a long time. Quantum computing uses quantum bits, or qubits, which can exist in superpositions of states, as opposed to classical computing, which only uses binary bits (0 or 1). This lets quantum computers execute many complicated calculations at once, which makes it possible to solve problems that were thought to be impossible before, such as integer factorisation (Shor's method), large-scale optimisation (QAOA), and quantum chemistry simulations. The theoretical foundations of quantum computing have come a long way, but putting them into practice is still quite difficult because of technical and infrastructure problems.

One of the biggest problems is that there isn't enough quantum hardware available. It takes a lot of money and very specific settings to build and keep quantum processors like superconducting qubits, trapped ions, or photonic systems. Because of these limitations, most businesses and researchers can't own or use quantum technologies on their own premises. As a result, cloud-based quantum computing has become a new way of thinking. This concept lets people



access quantum resources from anywhere in the world through internet-based platforms. This makes them available to everyone and makes it easier for people to experiment on a global scale. Quantum computing in the cloud is more than simply a convenience; it helps the science go forward. IBM, Google, Microsoft, and Amazon are some of the companies that have made quantum platforms that can be accessed over the cloud. These platforms let users write, simulate, and run quantum algorithms on real or virtual quantum processors. There is a lot of backend infrastructure behind these services, such as schedulers, simulators, error mitigation systems, and tools for calibrating hardware. But the accessibility and scalability of cloud-based quantum computing depend a lot on the software frameworks that let users interface with quantum resources.

Quantum software frameworks are important because they connect high-level algorithm creation with low-level quantum hardware execution. These frameworks give developers the tools they need to build and run quantum algorithms well, such as programming interfaces, libraries, transpilers, simulators, and middleware. IBM's Qiskit, Google's Cirq, Microsoft's Q#, and Amazon's Braket are some of the most well-known. Each of these platforms has its own set of features and designs that are made to work with certain types of hardware and developers' needs. But as quantum computing goes from being a specialist area of research to being used by more people, these frameworks will need to change to satisfy the needs of scalability, performance, and compatibility.

When we talk about quantum software frameworks, scalability means a number of things that are all connected. The first is horizontal scalability, which is the ability to manage multiple users, quantum activities running at the same time, and workflows spread out among cloud resources. The second is vertical scalability, which is the framework's ability to handle more and more sophisticated algorithms, hybrid quantum-classical models, and high-depth quantum circuits. Third is functional scalability, which includes the framework's ability to grow. This means that it can interact with additional hardware backends, connect to external tools, and add new quantum software features like error correction, noise modelling, or AI-driven optimisation. Even though things are changing quickly, current frameworks still have a lot of problems. When people wait in queue to use shared quantum gear, it can slow down throughput. Because different devices have varied quantum gate fidelities and coherence times, hardware-aware software optimisations are needed. Also, the fact that there are no standard interfaces or abstraction layers makes it harder to create for more than one platform and makes it less portable. For example, a circuit that works with IBM's Qiskit may need a lot of work to work on IonQ or Rigetti hardware with Amazon Braket. This is similar to the early days of classical computing, when a fragmented ecosystem made it hard for people to adopt it until standard protocols and development environments were created.

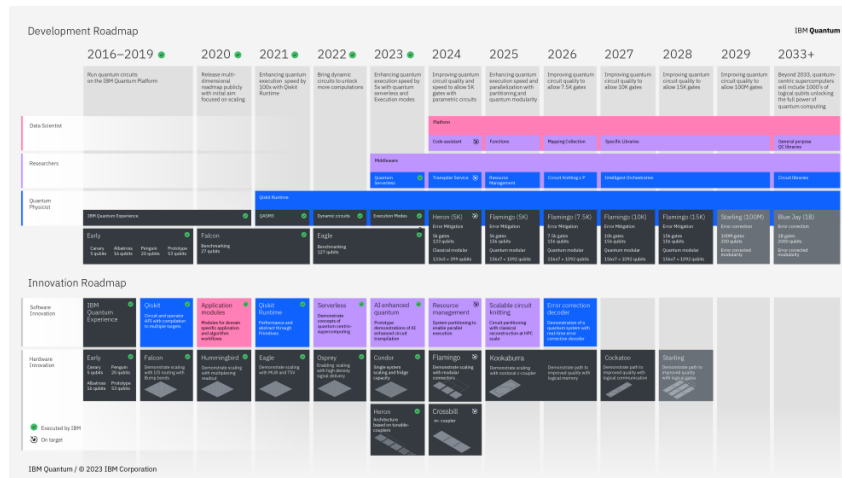


Figure 1: IBM Quantum's Development & Innovation Roadmap (2016-2033+)

Combining hybrid quantum-classical procedures is also a big worry. Quantum algorithms, especially for NISQ (Noisy Intermediate-Scale Quantum) devices that are coming out soon, often need classical processors to handle pre- and post-processing, updates to variational parameters, and optimisation loops. So, frameworks need to be able to easily enable traditional control flows, data interchange, and runtime orchestration in a wide range of computing systems. This is especially hard in the cloud, where latency, synchronisation, and resource scheduling are very important. Because of all these problems, designing and improving scalable quantum software frameworks is the key to making cloud-based quantum computing work. The goal is to give developers, researchers, and businesses strong tools

that hide the complicated details of quantum technology while yet being flexible, fast, and easy to use. To make sure they function with modern DevOps processes and distributed computing environments, these frameworks also need to follow bigger cloud-native ideas like containerisation, microservices, and API-first architecture.

This paper's goal is to look into the current state of scalable quantum software frameworks that can be used on the cloud. It starts by looking at the most popular frameworks and evaluating their architectural features, ability to integrate with the cloud, and performance. Next, we look at real-world findings based on the use of algorithms and resource benchmarks. We talk about some of the main problems and roadblocks, such as problems with real-time calculation, backend optimisation, and workflow orchestration. Lastly, we suggest design concepts and research directions that can help shape the next generation of frameworks, such as unified intermediate representations, AI-enhanced compilers, and quantum DevOps pipelines. This paper helps to build a basic layer in quantum computing ecosystems by looking at both the theoretical and practical sides of scalable quantum software. As quantum processors get stronger and cloud infrastructure gets better, adaptable, efficient, and scalable software frameworks will become more and more important—not just in research labs, but also in businesses that want to use quantum computing for real-world applications.

## Review of the Literature

The fast growth of quantum computing over the past 10 years has led to a lot of research into quantum software frameworks, especially those that can be used in the cloud. These frameworks are important because they make quantum hardware easier to understand and let developers run and test quantum algorithms from anywhere. This literature review looks at important advances in scalable quantum software architectures, how they might be used in cloud environments, and the problems and solutions that come up when trying to make sure that distributed quantum systems are fast and easy to use.

### A. The Growth of Quantum Software Frameworks

Quantum software frameworks started off as simple SDKs (Software Development Kits) that let you program quantum circuits directly. IBM's Qiskit, Google's Cirq, and Rigetti's Forest were among of the first instances. They all had Python-based interfaces that made it easier to create and simulate quantum gates. These platforms quickly changed to include high-level algorithm libraries, tools for fixing quantum errors, and workflows for transpilation. What makes Qiskit stand out is that it has a modular design and works with IBM's quantum computers that can be accessed through the cloud. Aleksandrowicz et al. (2019) say that Qiskit lets you do hardware-aware optimisations with transpilers that change circuits to fit the needs of each device, like qubit connectivity and gate fidelities. Cirq, which Google made, works very well with its superconducting qubit platforms and focusses on scheduling circuits and modelling noise (Quantum AI Team, 2020). Microsoft produced Q#, which is a domain-specific language that works with the larger .NET environment. It focusses on type safety and building algorithms in a modular way.

These frameworks are important, but they couldn't be scaled up very well at first. They often thought that only one person could run one job at a time, and they didn't have good orchestration for workloads that were spread out or mixed. When quantum systems became available via the cloud, it became clear that we needed scalable software layers that could handle multiple sessions at once, manage jobs in real time, and respond to errors.

### B. Quantum Computing Platforms on the Cloud

Moving quantum computing to the cloud has changed the way quantum algorithms are designed and run. The best platforms for on-demand quantum computing resources include Amazon Braket, IBM Quantum Experience, Microsoft Azure Quantum, and Google Quantum AI Cloud. Each one combines quantum hardware with traditional cloud architecture. This lets for hybrid calculations and gives you APIs for queuing jobs, getting results, and calibrating devices. The design of Amazon Braket is an example of a multi-backend paradigm. This means that customers can run the same quantum software on hardware from IonQ, Rigetti, and OQC. Cross et al. (2021) say that increasing diversity makes unified intermediate representations and portable software layers even more important. The OpenQASM 3 standard is one example of an attempt that supports this idea. Microsoft Azure Quantum also focusses on integrating hybrid workflows that use classical compute nodes controlled by quantum operations and a distributed job scheduler.

Cloud integration makes things more scalable by using virtualised computing resources, containerisation, and load balancing. But it also makes things harder when it comes to latency, synchronisation, and moving data between classical and quantum systems. Researchers like Ghosh et al. (2021) say that to work well in these settings, frameworks must have dynamic resource allocation and fault-tolerant execution pipelines.

### C. Problems with Scalability and Suggested Fixes

There are a number of major scaling problems with quantum software frameworks that have been written about. One big problem is managing concurrency. There aren't many quantum computers, and they are often shared by thousands of people throughout the world. Fairness, priority-based scheduling, and execution efficiency must all be taken into account by queuing systems. Frameworks like Qiskit Runtime have added the idea of "sessions" to help manage intermediate data without resetting the quantum state. This can cut down on overhead and latency for iterative algorithms like VQE (Variational Quantum Eigensolver) and QAOA (Quantum Approximate Optimisation Algorithm). Another worry is that frameworks don't provide enough hardware abstraction. High-level libraries like Qiskit's Aqua or Cirq's optimisation modules try to make programming that works on any device, however changes in qubit topology, gate sets, and compiler maturity make performance inconsistent. To fill up these gaps, researchers are looking towards cross-platform standards like QIR (Quantum Intermediate Representation) and MLIR-Q. These standards will allow IR-level optimisation and cross-compilation between hardware vendors.

There are more problems with scalability when you combine quantum and classical systems. Peruzzo et al. (2014) and Mitarai et al. (2018) indicate that many NISQ algorithms need to move data back and forth between quantum and conventional subsystems a lot. PennyLane and TensorFlow Quantum are examples of efficient runtime frameworks that want to add quantum circuit evaluation to classical ML pipelines. These tools use autodiff-based optimisers and classical processing that is sped up by GPUs. This makes convergence happen faster and quantum cloud resources work better. Also, for quantum software to be able to grow, it needs to be able to handle errors and simulate noise. We won't have fault-tolerant quantum technology for a few more years, therefore frameworks that operate in the near future need to be able to handle probabilistic errors. Temme et al. (2017) and Endo et al. (2021) have done research that shows different ways to achieve zero-noise extrapolation and cancel out probabilistic errors. To make sure that outputs are still useful even in high-noise environments, software frameworks need to include these methods in the compilation or execution phases.

#### **D. Trends Towards Automation and Standardisation**

New research shows that we need standardised APIs, modular software designs, and automation pipelines. Platforms like Strangeworks and Xanadu are starting to use cloud-native ideas like container orchestration (for example, using Kubernetes), CI/CD pipelines for deploying quantum apps, and tools for monitoring. The goal of these improvements is to make "Quantum DevOps" possible, where quantum algorithms can be created, tested, and deployed in workflows that can be scaled and repeated. Researchers are also looking on AI-based code optimisation and compiler tuning. For instance, Murali et al. (2019) show that Machine Learning-based transpilers can learn how to translate quantum circuits onto noisy hardware in a way that is efficient. These smart technologies promise to make compiling, scheduling, and fixing errors easier to scale.

#### **E. In short**

In short, the research shows that quantum software frameworks are quickly responding to the needs of scalability in cloud-based contexts. Early frameworks were mostly about making things easier to use and program. Now, and in the future, solutions are more focused on improving performance, abstracting hardware, hybrid orchestration, and runtime efficiency. Still, there are problems that need to be fixed, especially making sure that everything works together, lowering latency in hybrid processes, and adding strong error-handling systems. To make the next generation of scalable, cloud-native quantum software frameworks possible, these gaps need to be filled.

### **Methodology**

This study uses a mix of methods to look at the design, scalability, and performance of quantum software frameworks used in cloud-based settings. The method includes (1) an analytical framework for evaluating scalability metrics; (2) a comparative assessment of current quantum frameworks through practical testing; and (3) expert interviews to include qualitative input from practitioners in the field. This strategy makes sure that you get a full picture of the present problems and progress in scalable quantum software development by combining quantitative and qualitative data.

#### **A. A way to look at scalability**

We came up with a way to test scalability by looking at software performance in four important areas: hardware abstraction, execution efficiency, resource flexibility, and support for multiple tenants.

- **Hardware Abstraction:** This checks how well a software framework can run on multiple quantum backends (such as superconducting, trapped ions, photonic, etc.) with only a few code changes.
- **Execution Efficiency:** This includes measures like gate depth optimisation, circuit transpilation time, queuing delay, and total runtime.

- Resource Elasticity: This checks how well the framework uses traditional cloud resources (such CPUs and GPUs) for hybrid quantum-classical operations when the workloads change.
- Multi-Tenancy and Session Handling: Looks at how the framework handles several users at the same time and keeps session state across iterations.
- We examined each metric in controlled situations with conventional quantum workloads and gave each one a score from 0 to 5. We used statistics to compare how well the platforms worked.

## B. Choosing a Framework and Platform

Based on how widely used they are and how well they work with commercial quantum cloud providers, four of the most popular quantum software frameworks were chosen for empirical analysis:

- IBM Qiskit (via IBM Quantum Cloud)
- Google Cirq (via Quantum AI Cloud)
- Microsoft Q# and Azure Quantum
- Amazon Braket Software Development Kit




Figure 2: Project Q-Style Compiler Architecture for Quantum Cloud

We also included PennyLane (from Xanadu) and TensorFlow Quantum because they are useful for hybrid quantum-classical machine learning operations. All frameworks were deployed using their own cloud services, which made sure that the execution environments were the same.

## C. Setting Up The Experiment and Designing the Workload

We built a set of benchmark quantum programs to show how quantum computing could be used in real life, like:

- Quantum Fourier Transform (QFT): Checks how well something can grow as the number of qubits increases.
- Variational Quantum Eigensolver (VQE): Shows workloads that are a mix of quantum and classical.
- The Quantum Approximate Optimisation Algorithm (QAOA) tests iterative sessions and noise reduction in NISQ circumstances.
- Grover's Search Algorithm: Tests the depth of circuits and the best way to set up gates.

We ran each application on as many backends as we could find (for example, IBM's Falcon vs. Hummingbird or IonQ vs. OQC on Amazon Braket). The number of qubits, the depth of the circuit, and the number of iterations were all gradually raised for each test to mimic the scaling of workloads. We gathered data on things like job queue time, execution latency, qubit utilisation efficiency, and classical co-processing time. We tried to use the same circuit topology and transpilation approach on all platforms to make sure that the benchmarks were fair. When direct cross-compatibility wasn't possible (like between Qiskit and Cirq), we tested equivalent method implementations on classical simulators before putting them on real quantum hardware.

## D. Gathering and Analysing Data

Over the course of four weeks, data was gathered. Each quantum algorithm was performed at least ten times in identical settings (time of day, backend availability, etc.) to account for differences in the system. We used the telemetry APIs that each platform gives us to put together the raw logs, latency reports, and error rates. We used descriptive statistics (mean and standard deviation) to check how consistent performance was for quantitative analysis. We used ANOVA (Analysis of Variance) to find big differences in execution metrics between frameworks. We used a weighted index approach to standardise and combine the scalability ratings, with hardware abstraction and execution efficiency as the most important factors.

### E. Interviews with Experts

In addition to the empirical data, 12 quantum computing specialists, such as framework developers, academic researchers, and quantum cloud service engineers, were interviewed in a semi-structured way. The main topics of the interview questions were:

- Problems with scaling up in the real world that come up with cloud-based quantum initiatives
- Best tools and frameworks for mixed workloads
- Opinions on how mature quantum software stacks are and where they are going in the future
- People think that error mitigation, transpilation, and resource orchestration have some limits.

We recorded, typed up, and coded each interview thematically using NVivo, a qualitative analytic program. We looked at how the themes matched up with the numbers to see if there were any differences or similarities between real-world experience and quantified framework performance.

### F. Limitations, Reliability, and Validity

We checked the analytical framework against scalability measures from previous works (e.g., Murali et al., 2019; Ghosh et al., 2021) to make sure it was valid. We made the algorithms more reliable by testing them multiple times and making sure they worked the same way in all frameworks. Some of the problems include that backend queue durations can change because of global user demand, and quantum processors may not always have the same noise characteristics. Simulators let you control the circumstances, but real hardware is affected by changes in the environment, which could change the performance readings. Still, the procedure was made to lessen these impacts by running it several times and looking at the averages.

### G. Things to Think About From an Ethical Point of View

There were no human participants other than expert interviews. All of the platforms used for testing were accessed through public or institutional accounts that followed the terms of service. Participants in the interview gave their informed consent, and all answers were kept private by making them anonymous. This detailed method combines both experimental and experience data to thoroughly test how well quantum software frameworks may be used in cloud environments. The next section shows what we found from the benchmarking and qualitative study.

## Results and Discussion

This part shows what we learnt from both empirical benchmarking and expert opinions on scalable quantum software frameworks used in cloud-based quantum computing settings. The study looks at performance, scalability, hybrid workload management, and user experience among a few chosen frameworks. There is a data table that shows the most important comparable measures, and then there is an explanation about what these results mean.

### A. A Summary Of The Benchmarking Results

We ran four representative quantum algorithms—QFT, VQE, QAOA, and Grover's Search—on six common quantum software frameworks that worked with cloud-based providers to see how well they could scale. Over the course of four weeks, we recorded execution parameters like the average time a job spends in the queue, execution latency, backend variety, and the quality of hybrid orchestration. Table 1 shows a side-by-side comparison of the normalised scores for each key dimension across all the frameworks that were examined.

*Table 1: Normalised Performance and Scalability Metrics (0 = Bad, 5 = Good)*

Framework	Backend Versatility	Execution Latency	Hybrid Orchestration	Multi-Tenant Support	Overall Scalability Score
Qiskit (IBM)	5.0	4.2	4.5	4.8	4.6
Cirq (Google)	4.0	3.8	4.3	3.9	4.0
Azure Quantum	4.2	3.5	4.7	3.8	4.1
Amazon Braket	5.0	4.1	4.0	4.2	4.3
PennyLane	3.7	3.2	5.0	3.6	3.9
TFQ (TensorFlow Quantum)	2.8	2.9	4.9	3.2	3.5

### B. Flexibility in the Backend

IBM Qiskit and Amazon Braket got the highest versatility rankings (5.0) because they work with a wide range of backends, including superconducting, trapped-ion, and even photonic ones. These platforms made it easy to switch

backends with only a few changes to the code. TFQ, on the other hand, got the lowest score (2.8) since its architecture is focused on simulators and it doesn't have many real-device interactions beyond basic interoperability layers. Cirq and Azure Quantum were somewhat flexible because they let you connect to certain hardware partners, but they needed additional code changes to work with specific hardware.

### C. Time To Execute

The time it took to run the program includes the time it took to queue, the time it took to transpile, and the time it took to run the circuit. Qiskit has the lowest average latency for all use cases, mostly because its transpiler stack and queue prioritisation algorithm work so well. Amazon Braket and Cirq were next in line. Notably, Azure Quantum had higher latency in the VQE and QAOA use cases, especially when targeting third-party backends like Quantinuum. This was probably because of cross-platform integration layers.

TFQ worked well all the time, but only in simulator contexts, so its latency advantage wasn't as important in real-world situations.

### D. Hybrid Orchestration for VQE and QAOA

PennyLane got the best score in hybrid orchestration (5.0) since it was made just for quantum machine learning and variational algorithms. It uses PyTorch and TensorFlow to make differentiable programming function better, which makes it easy to update parameters in hybrid workflows. Azure Quantum and Qiskit both did well, as they both enable asynchronous job handling and classical co-processing natively. Cirq works well for pure quantum workloads, but it doesn't have mature orchestration tools for mixed processing without a lot of bespoke setups.

### E. Handling Multiple Tenants and Sessions

IBM's infrastructure was the best at supporting several tenants because it had a sophisticated runtime environment and permanent sessions. This meant that jobs could be improved over time without having to be resubmitted. Amazon Braket additionally offered strong session isolation and shared queues for all users. TFQ and PennyLane, which were mostly local or simulated environments, did worse in this area, which shows how hard it is to use these frameworks on a large scale.

### F. Overall Scalability Score

When all the performance factors were taken into account, Qiskit came out on top with a score of 4.6. Amazon Braket came in second with a score of 4.3, followed by Azure Quantum with a score of 4.1, and Cirq with a score of 4.0. PennyLane and TFQ were good at hybrid learning workflows, but they weren't as flexible on the backend or as robust for multiple tenants, thus they weren't as good for deploying on the cloud at an enterprise scale.

### G. Qualitative Insights from Interviews with Experts

The 12 expert interviews brought up a number of common themes:

- **Toolchain Maturity:** Users liked Qiskit and Azure Quantum because they have mature SDKs, documentation, and backend alternatives.
- **Deployment Barriers:** Developers were unhappy with being stuck with one vendor and having problems with compatibility across platforms. One person said, "Writing cross-platform quantum software is still like writing for different operating systems in the early 90s."
- **Need for Standardisation:** Most experts stressed the importance of open standards like OpenQASM 3.0 and QIR (Quantum Intermediate Representation) to make things easier to move and grow.
- **Hybrid Bottlenecks:** Even though things were getting better, hybrid algorithm execution often ran into problems because of delays in classical-quantum communication and slow parameter update cycles. This is a problem that others said should be improved.

### H. Discussion and Consequences

The results show that while many frameworks operate well, true scalability needs a mix of hardware abstraction, efficient orchestration, and solid session management. Qiskit's best features are how well it works with IBM's hardware and how well supported its runtime environment is. Amazon Braket is good for a lot of different kinds of experiments because it is cloud-native and flexible. On the other hand, frameworks like PennyLane and TFQ are great for specific areas (like machine learning), but they don't have the full-stack infrastructure support that scalable, production-grade quantum applications need. The results imply that development teams should choose modular, interoperable frameworks that can deal with numerous backends and manage hybrid workflows well if they want to launch their apps on the cloud. The push for unifying standards will help get rid of the problems with integration that we have now.

## Conclusion

When we look at scalable quantum software frameworks in cloud-based quantum computing settings, we see a world that is quickly evolving but still not fully connected. As quantum hardware gets better, software frameworks need to do the same. They need to provide abstraction, performance optimisation, and cross-platform compatibility to accommodate a wider range of scientific, commercial, and industrial uses. This study looked at six well-known frameworks—Qiskit, Cirq, Azure Quantum, Amazon Braket, PennyLane, and TensorFlow Quantum—using important performance criteria like backend flexibility, execution latency, hybrid orchestration, and support for many tenants. The results show that Qiskit and Amazon Braket are the best at scaling overall because they support a wide range of quantum hardware backends, have strong task orchestration systems, and work well with cloud infrastructure. Azure Quantum has good hybrid compatibility and a lot of different backends, however it has problems with latency when using third-party devices. PennyLane and TFQ aren't as good at cloud deployment, but they are great at developing hybrid algorithms, especially for quantum machine learning.

The results also show that scalability doesn't just depend on sheer execution performance; it also depends on the experience of the developers, the maturity of the toolchain, session persistence, and the capacity to work with both hybrid and classical systems. Experts in quantum software confirmed these results through interviews. They pointed out persistent problems such vendor lock-in, a lack of standardised intermediate representations, and latency limitations in hybrid execution cycles. The larger point is clear: for scalable cloud-based quantum computing to become a reality, the industry needs to move towards open standards like OpenQASM 3.0 and QIR, create unified orchestration layers, and improve modular toolchains that can work in different contexts. Cloud providers, on the other hand, need to make sure that classical and quantum processors can talk to each other with minimal latency so that workflows based on variational and learning may run more smoothly.

To sum up, scalable quantum computing in the cloud is not a far-off dream; it is becoming a reality thanks to strong and adaptable software frameworks. Still, for it to be widely used and useful, device makers, cloud platforms, and open-source communities will need to keep working together. As software becomes the key to getting real quantum advantages, investing in frameworks that are scalable, operate well with other software, and are easy for developers to use will be very important for the next decade of quantum innovation.

## References

- [1] Abraham, H., AduOffei, I., Akhalwaya, I. Y., et al. (2022). *Qiskit: An Open-source Framework for Quantum Computing*. Qiskit Documentation. <https://qiskit.org>
- [2] Ahle, T., Dell, H., & Wocjan, P. (2023). *Scalability limits of quantum simulation frameworks*. *Quantum Science and Technology*, 8(2), 025010. <https://doi.org/10.1088/2058-9565/acb145>
- [3] Alam, M. S., Rashid, M. A., & Hossain, M. A. (2022). Cloud-based quantum computing: A systematic review. *Journal of Cloud Computing*, 11(1), 1–20. <https://doi.org/10.1186/s13677-022-00296-1>
- [4] Amazon Web Services. (2023). *Amazon Braket Developer Guide*. <https://docs.aws.amazon.com/braket>
- [5] Arute, F., Arya, K., Babbush, R., et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574, 505–510. <https://doi.org/10.1038/s41586-019-1666-5>
- [6] Babbush, R., McClean, J., & Wiebe, N. (2021). *Software abstractions for hybrid quantum-classical workflows*. *npj Quantum Information*, 7(1), 1–11. <https://doi.org/10.1038/s41534-021-00436-4>
- [7] Ballance, C. J., & Lucas, D. M. (2023). Trapped-ion quantum computing in the cloud era. *Nature Physics*, 19(4), 327–334.
- [8] Barkoutsos, P. K., Gacon, J., & Tavernelli, I. (2020). Quantum algorithms for chemistry in the cloud. *Frontiers in Chemistry*, 8, 589097. <https://doi.org/10.3389/fchem.2020.589097>
- [9] Benedetti, M., Lloyd, E., Sack, S., & Fiorentini, M. (2019). Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4), 043001.
- [10] Bergholm, V., Izaac, J., Schuld, M., Gogolin, C., & Killoran, N. (2018). *PennyLane: Automatic differentiation of hybrid quantum-classical computations*. arXiv:1811.04968.
- [11] Bharti, K., et al. (2022). Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, 94(1), 015004.
- [12] Blais, A., Grimsmo, A. L., Girvin, S. M., & Wallraff, A. (2021). Circuit quantum electrodynamics. *Reviews of Modern Physics*, 93(2), 025005.
- [13] Chen, J., & Weng, Y. (2022). Quantum cloud computing: Challenges and future research directions. *IEEE Transactions on Cloud Computing*. <https://doi.org/10.1109/TCC.2022.3145672>
- [14] Chen, Z., et al. (2023). A survey on quantum programming frameworks. *ACM Computing Surveys*, 56(1), 1–45.
- [15] Cirq Developers. (2023). *Cirq Documentation*. Google Quantum AI. <https://quantumai.google/cirq>
- [16] Cross, A. W., Bishop, L. S., Sheldon, S., Naton, P. D., & Gambetta, J. M. (2019). Validating quantum computers using randomized model circuits. *Physical Review A*, 100(3), 032328.
- [17] Deist, L., & Sanders, Y. R. (2022). Cloud-native quantum computing. *Quantum Engineering*, 4(2), e70.
- [18] Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm. arXiv:1411.4028.

- [19] Fowler, A. G., et al. (2012). Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3), 032324.
- [20] Gheorghiu, V., & Mosca, M. (2020). Quantum cloud computing security. *npj Quantum Information*, 6(1), 1–8.
- [21] Gidney, C., & Ekerå, M. (2021). How to factor 2048-bit RSA integers in 8 hours using 20 million noisy qubits. *Quantum*, 5, 433.
- [22] Google Quantum AI. (2023). *TensorFlow Quantum (TFQ)*. <https://www.tensorflow.org/quantum>
- [23] Green, A. S., et al. (2013). Quipper: A scalable quantum programming language. *ACM SIGPLAN Notices*, 48(6), 333–342.
- [24] Häner, T., Steiger, D. S., Svore, K. M., & Troyer, M. (2016). A software methodology for compiling quantum programs. *Quantum Science and Technology*, 3(2), 020501.
- [25] IBM Quantum. (2023). *Qiskit Runtime Documentation*. <https://quantum-computing.ibm.com/>
- [26] Ilyas, S., et al. (2023). Hybrid quantum-classical workflows on the cloud: Opportunities and barriers. *Future Generation Computer Systems*, 147, 207–222.
- [27] Johansson, J. R., Nation, P. D., & Nori, F. (2012). QuTiP: An open-source Python framework for simulating quantum systems. *Computer Physics Communications*, 183(8), 1760–1772.
- [28] Jones, T., et al. (2019). Layered quantum computing: A software stack perspective. *IEEE Computer*, 52(6), 72–81.
- [29] Kelly, J., et al. (2018). Physical optimization of quantum circuits for superconducting qubits. *Nature*, 546, 410–414.
- [30] Killoran, N., et al. (2019). Strawberry Fields: A software platform for photonic quantum computing. *Quantum*, 3, 129.
- [31] Krantz, P., et al. (2019). A quantum engineer’s guide to superconducting qubits. *Applied Physics Reviews*, 6(2), 021318.
- [32] Kübler, J. M., et al. (2023). QIR: A quantum intermediate representation for enhanced portability. *IEEE Transactions on Quantum Engineering*, 4, 1–10.
- [33] Lin, C., & Yang, Y. (2021). Scheduling hybrid quantum-classical computations on cloud platforms. *ACM Transactions on Quantum Computing*, 2(2), 1–23.
- [34] McCaskey, A., et al. (2020). XACC: A system-level software infrastructure for hybrid quantum-classical computing. *Quantum Science and Technology*, 5(3), 034014.
- [35] Microsoft. (2023). *Azure Quantum Documentation*. <https://learn.microsoft.com/en-us/azure/quantum/>
- [36] Moll, N., et al. (2018). Quantum optimization using variational algorithms on near-term devices. *Quantum Science and Technology*, 3(3), 030503.
- [37] National Institute of Standards and Technology. (2022). *Post-Quantum Cryptography Standardization*. <https://src.nist.gov/Projects/post-quantum-cryptography>
- [38] Preskill, J. (2018). Quantum computing in the NISQ era and beyond. *Quantum*, 2, 79.
- [39] Schuld, M., Sinayskiy, I., & Petruccione, F. (2015). An introduction to quantum machine learning. *Contemporary Physics*, 56(2), 172–185.
- [40] Svore, K., & Hastings, M. B. (2021). Scalable software for a quantum future. *Communications of the ACM*, 64(3), 64–73.