Q2 2026

IBM **Quantum**

The Coherence Times

A Quarterly Readout

10 years
on the cloud

About IBM Quantum

IBM Quantum is the world's leading full-stack provider in quantum computing, dedicated to bringing useful quantum computing to the world. IBM Quantum stands out through its robust and scalable quantum hardware, a comprehensive suite of tools and services, and one of the largest fleets of quantum systems available via the cloud. With more than 75 quantum computers deployed globally, users can access these systems to explore, experiment, and build quantum solutions. The platform has supported over 4 trillion quantum circuit executions and fostered a vibrant global community of more than 240,000 users and 300+ ecosystem partners. This community has contributed to over 5,900 research publications, highlighting IBM Quantum's impact across academia, industry, and research.

Today, IBM offers a variety of access plans for quantum-interested organizations of all sizes. This includes the Open Plan offering ten free minutes per month, the Pay-As-You-Go Plan, our Flex Plan for project-based work, and our Premium Plan or On-Premises systems for clients requiring the highest capacity. Through continuous innovation and collaboration, IBM Quantum isn't just advancing computing—it's redefining what's possible.

Our quarterly leadership update

A note from Jay

Hi all, and welcome to the Q2 2026 edition of the Coherence Times: 10 years on the cloud.

On May 4, 2016, our team made history by launching the IBM Quantum Experience and putting the first quantum processor on IBM Cloud. The intervening years have brought an unthinkable amount of progress. We introduced Qiskit, the open source software development kit for programming quantum computers. We built the first integrated quantum computing system with IBM Quantum System One and the platform for scaling to fault tolerance with IBM Quantum System Two. We broke the 100-qubit barrier, and we demonstrated quantum utility—the ability to run quantum circuits beyond the ability of any brute-force classical method.

By "we," I don't just mean IBM. We accomplished incredible things as a quantum computing community, as evidenced by another incredible milestone hit this past May: users have run four trillion circuits on IBM quantum computers. These circuits represent boundary-pushing experiments and proofs-of-concept bringing quantum to real-world use cases. They represent our partners at Cleveland Clinic and RIKEN modeling proteins at the level of 12,000 atoms and at Oak Ridge National Laboratory simulating molten salts inside of fusion reactors. They represent our clients across industries prototyping scalable use cases, our startups delivering products that significantly improve outputs and handle errors, and our researchers uncovering algorithms with the potential to turbocharge some of the world's most important workflows. Things have only accelerated as we continue driving along our roadmap. Higher-performing chips have helped scale our partners' workflows, while the community has continued growing alongside ever-increasing interest. And new generative AI tools including the IBM Bob development partner have lowered the barrier-to-entry for developers hoping to get started.

Then there's quantum advantage. Some of our users are already seeing runtime and accuracy improvements by adding quantum to their workflows, while the computational science community is working to find verifiable examples of quantum separations for certain problems. Back in 2024 I predicted that we'd see quantum advantage in 2026 if the quantum and classical communities worked together—and based on what I'm seeing, I'm feeling very confident with my prediction.

It's incredible looking back over the past ten years—but it's even more exciting to see our vision for the next ten years come into focus. By creating a quantum foundry with Anderon, we're laying the foundation for an entire quantum computing industry in the United States. And we're starting to blueprint concepts for a quantum datacenter with networked quantum computers, complete with sensing and a quantum computing internet.

I'm thankful for everything this community has done to get us here. Now let's set our sights on quantum advantage!

Quantum for Fusion

The Department of Energy launched the Genesis Mission in late 2025 to accelerate scientific discovery by connecting supercomputers, AI systems, and quantum computers across the DOE's 17 national laboratories with the country's major scientific instruments. The mission is led by Darío Gil, the DOE Under Secretary for Science and former director of IBM Research.

IBM is one of the mission's industry partners, with a focus on quantum-centric supercomputing: CPUs, GPUs, and QPUs working together, each running the part of a calculation it handles best. As part of the ongoing Genesis Mission project, IBM and Oak Ridge National Laboratory (ORNL) are applying quantum computing and AI to questions involved in fuel for fusion reactors.

Today, all useful nuclear power comes from fission reactors, which harness the energy of splitting large radioactive atoms like uranium. Fusion reactors get their power from the fusion of much lighter atoms: deuterium and tritium, both isotopes of hydrogen. Deuterium is abundant in seawater but tritium is not—across planet Earth humans synthesize just a few pounds each year, mainly using fission reactors. There are no practical natural sources because tritium's half life is just a bit more than 12 years, so any stockpiles rapidly decay into helium. Practical fusion power requires new methods. But there is a plan: have fusion reactors make their own tritium.

One could start by lining a donut-shaped fusion reactor called a tokamak with a blanket of molten salt called FLiBe (rhymes with "tribe"), a mixture of lithium fluoride (LiF) and beryllium fluoride (BeF_2). As atoms fuse together in the tokamak's plasma, they release excess neutrons, which strike the lithium-6 in the salt and shatter it into a tritium nucleus and a helium nucleus. That tritium, in theory, can then be extracted and fed back in to the reactor. The molten salt holds other roles, too—as a radiation shield and to aid in various thermal tasks.

But there's a key question: How hard will it be to actually extract the tritium from the salt blanket once produced? That depends on how it behaves in the salt—behaviors which are challenging for classical methods to predict accurately. Predicting the cost of extracting the tritium from the salt blanket means computing the amount of free energy that binds tritium to the molten salt. Right now, IBM and ORNL are exploring a quantum-centric supercomputing (QCSC) approach to the problem.

In their workflow, AI agents explore, synthesize, clean, and structure 70 years' worth of experimental data on molten salts. Then, another set of AI agents uses that data to propose the best candidate salts for deeper exploration. That's where quantum enters the picture: Quantum and classical computers work together to calculate the electronic structure of the molten salts. They are aided by AI agents to reveal properties, like thermal conductivity, that would be invisible without AI in the workflow. In FLiBe, atomic nuclei do not occupy static lattice positions. Instead, under thermal fluctuations at typical operating temperatures (500-700 celsius) nuclei continuously rearrange into random, dynamic, temperature-dependent configurations. Sampling these random configurations through AI is key for accurate binding energy computations. Those results feed back into the nation's molten salt data sets to refine and improve them, closing the loop. An orchestration agent oversees the entire process.

In the work shown so far, the team computed the ground-state energies of nine FLiBe configurations using up to 64 qubits and 2,654 two-qubit gates—and the quantum-centric calculation matched leading classical methods without the approximations such calculations usually require.

This study reveals where we are in terms of computational characterization of FLiBe, where we need to go to improve, and opportunities for the inclusion of AI to boost accuracy and efficiency. There's still latency in the back-and-forth between classical and quantum, so developing tighter, more integrated workflows will be key to practical utility.

The collaboration is ongoing. The team is building toward a demonstration of the full loop, with the eventual goal of creating a tool the fusion industry can use to design and verify its own blanket salts.

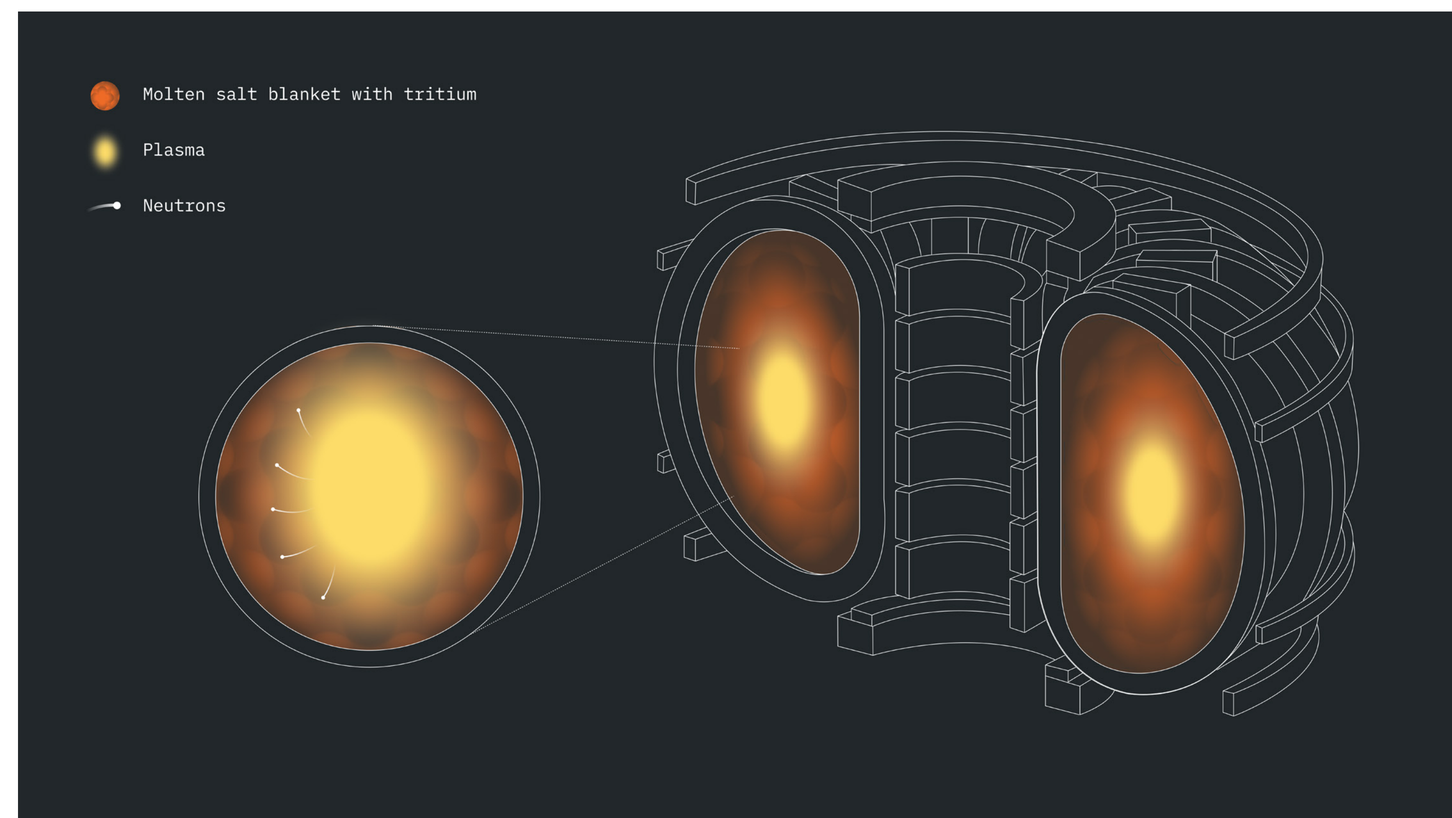

Diagram showing neutrons from the plasma interacting with the molten salt blanket in a tokamak.

Use AI for Quantum!

The emergence of high-performing AI code generators has revolutionized programming workflows. Rather than programming from scratch, developers can build applications simply by having a conversation with an AI agent—even Qiskit applications. That means programming a quantum computer has never been easier.

IBM Quantum researchers are increasingly relying on AI code generation tools like Bob, Claude, and Codex throughout their workflows. At present, they're typically using these tools to aid in the tedious parts—generating plots, performing tests, and integrating tools into other experiments. However, they're also beginning to generate circuits, edit circuits, and troubleshoot the outputs of quantum computers.

As these tools mature, we expect researchers to rely more and more on AI agents to assist in the creation of their quantum circuits and running quantum workflows. As a leading AI company, we're proud to show off how these tools are accelerating our workflows, and we'll increasingly be spotlighting interesting LLM-accelerated use cases. But most importantly, with AI tools making quantum development easier than ever, there's no excuse not to get started running your first quantum circuit. But first, here are a few guidelines:

- **Apply real development discipline.** Think about your program before you create it—documentation, design, and security review still matter. Include as much information as you can in your prompt about your program and your desired outputs.

- **Ask for context-aware explanations alongside code.** If something is unclear, ask the model to explain why it made decisions in context. For example, why did it structure its code the way that it did, why did it construct the circuits in a certain way, or why did it apply specific transformations? This reinforces your understanding and improves downstream code quality.

- **Build test harnesses early and aggressively.** Because quantum results rely on statistics, correct-looking outputs might not be correct. Build in statistical tests that compare against expected distributions.

- **Use AI as an accelerator, not an authority.** Expect hallucinations and subtle errors; never assume AI-generated Qiskit code is correct without validation.

- **Use the best model you can afford for non-trivial work.** Cheaper models may save money short-term but often cost time, correctness, and confidence in quantum applications.

So, how are our IBM developers actually using AI agents in their workflows?

Developer **Samantha Barron** on the large-scale demonstrations team uses Claude Developer to expand the scale that she operates. For example, she frequently runs layer fidelity experiments based on the well-trusted implementation in Qiskit. She was able to spin this section of Qiskit off to fold it into other experiments and make it more interoperable so she could more efficiently do her research.

Quantum Algorithms and Demonstrations researcher **Bibek Pokharel** uses IBM Bob as if it were an assistant developer—for example, if he's already crafted a circuit but wants to perform repeatable modifications to it, he'll ask Bob to create a function with documentation. He also uses Claude to troubleshoot Qiskit code that's produced incorrect circuits to find the inaccuracy—Claude was able to scan a git history, look at the code, and see exactly when the error was introduced to a circuit during the development process. Finally, he uses it to aid in creating, comparing, and analyzing plots, drastically reducing the amount of time spent with image formatting for presentations.

Quantum Theory and Capabilities researcher **Andrea D'Urbano** uses IBM Bob to develop experiments for quantum simulations using tensor networks in Julia. He typically creates the structure of his project itself, writing snippets of code and function signatures. Then in a long markdown file he specifies how to complete the sketch he did, the architecture to use, the specifics of the simulation to run, data to save, and the plots to output. He makes tweaks and updates to this file with Bob in Orchestrator mode to produce a more structured code for his purpose, which he then runs.

Research scientist **Brad Mitchell** uses IBM Bob to prototype new analyses and spin up tweaks to existing experiments. For example, he wanted to compare the results between several quantum experiments across multiple days with minor differences between them. Bob was able to handle the file parsing work, combine the data into a sensible format, and produce plots comparing the datasets. In short, he found it to be a useful tool for handling a lot of the data wrangling and plotting work that usually takes up a fair bit of time for researchers.

Introducing Anderon

On May 21, 2026, IBM, with support from the United States Department of Commerce, announced the launch of Anderon, the United States' first pure-play quantum foundry. Anderon will be a standalone IBM company dedicated to manufacturing quantum wafers to catalyze quantum innovation for the United States and strengthen the domestic quantum supply chain.

Headquartered in Albany, New York, Anderon will operate as a state-of-the-art 300-millimeter quantum wafer foundry, representing one of the most significant commitments by the U.S. government to date in quantum research and development. This initiative aims to position the United States to lead quantum wafer production globally and helps secure America's place at the center of a thriving quantum industry. Boston Consulting Group estimates that quantum computing could create between $450–850 billion in economic value by 2040, underscoring the importance of scaling manufacturing capabilities now.

Anderon will offer quantum wafer foundry services to multiple quantum hardware vendors. Initially, the foundry will focus on wafers built with superconducting qubits, with plans to expand to other quantum modalities as industry needs evolve. Quantum chip fabriaction requires bespoke equipment and processing—by leveraging IBM's proven 300mm wafer technology and quantum manufacturing expertise, Anderon will be able to serve this growing ecosystem from day one.

The Department of Commerce selected IBM for this initiative following a rigorous evaluation, recognizing IBM as the only organization with a manufacturable 300mm superconducting qubit wafer process, state-of-the-art fabrication tools, a credible quantum roadmap, and industry-leading experience supporting more than 90 deployed quantum systems. IBM's demonstrated leadership spans advanced architectures such as IBM Quantum Heron and Nighthawk, decades of semiconductor R&D, and major federal collaborations, including NIST's post-quantum cryptography standards and DARPA's quantum benchmarking efforts.

As a standalone entity, Anderon will maintain a separate workforce and isolated IT infrastructure, ensuring customer intellectual property is secure and inaccessible to IBM. IBM will transfer baseline quantum wafer processing IP to Anderon to support all customers, while next-generation IP will be licensed for exclusive IBM use and co-developed under a joint development agreement.

From day one, Anderon will engage commercial companies, laboratories, and universities, with the first process design kits expected within the first year of operation. Planning is ongoing, with transitions anticipated to begin between Q3 and Q4 of 2026, reinforcing a long-term commitment to national -scale quantum manufacturing and future growth.

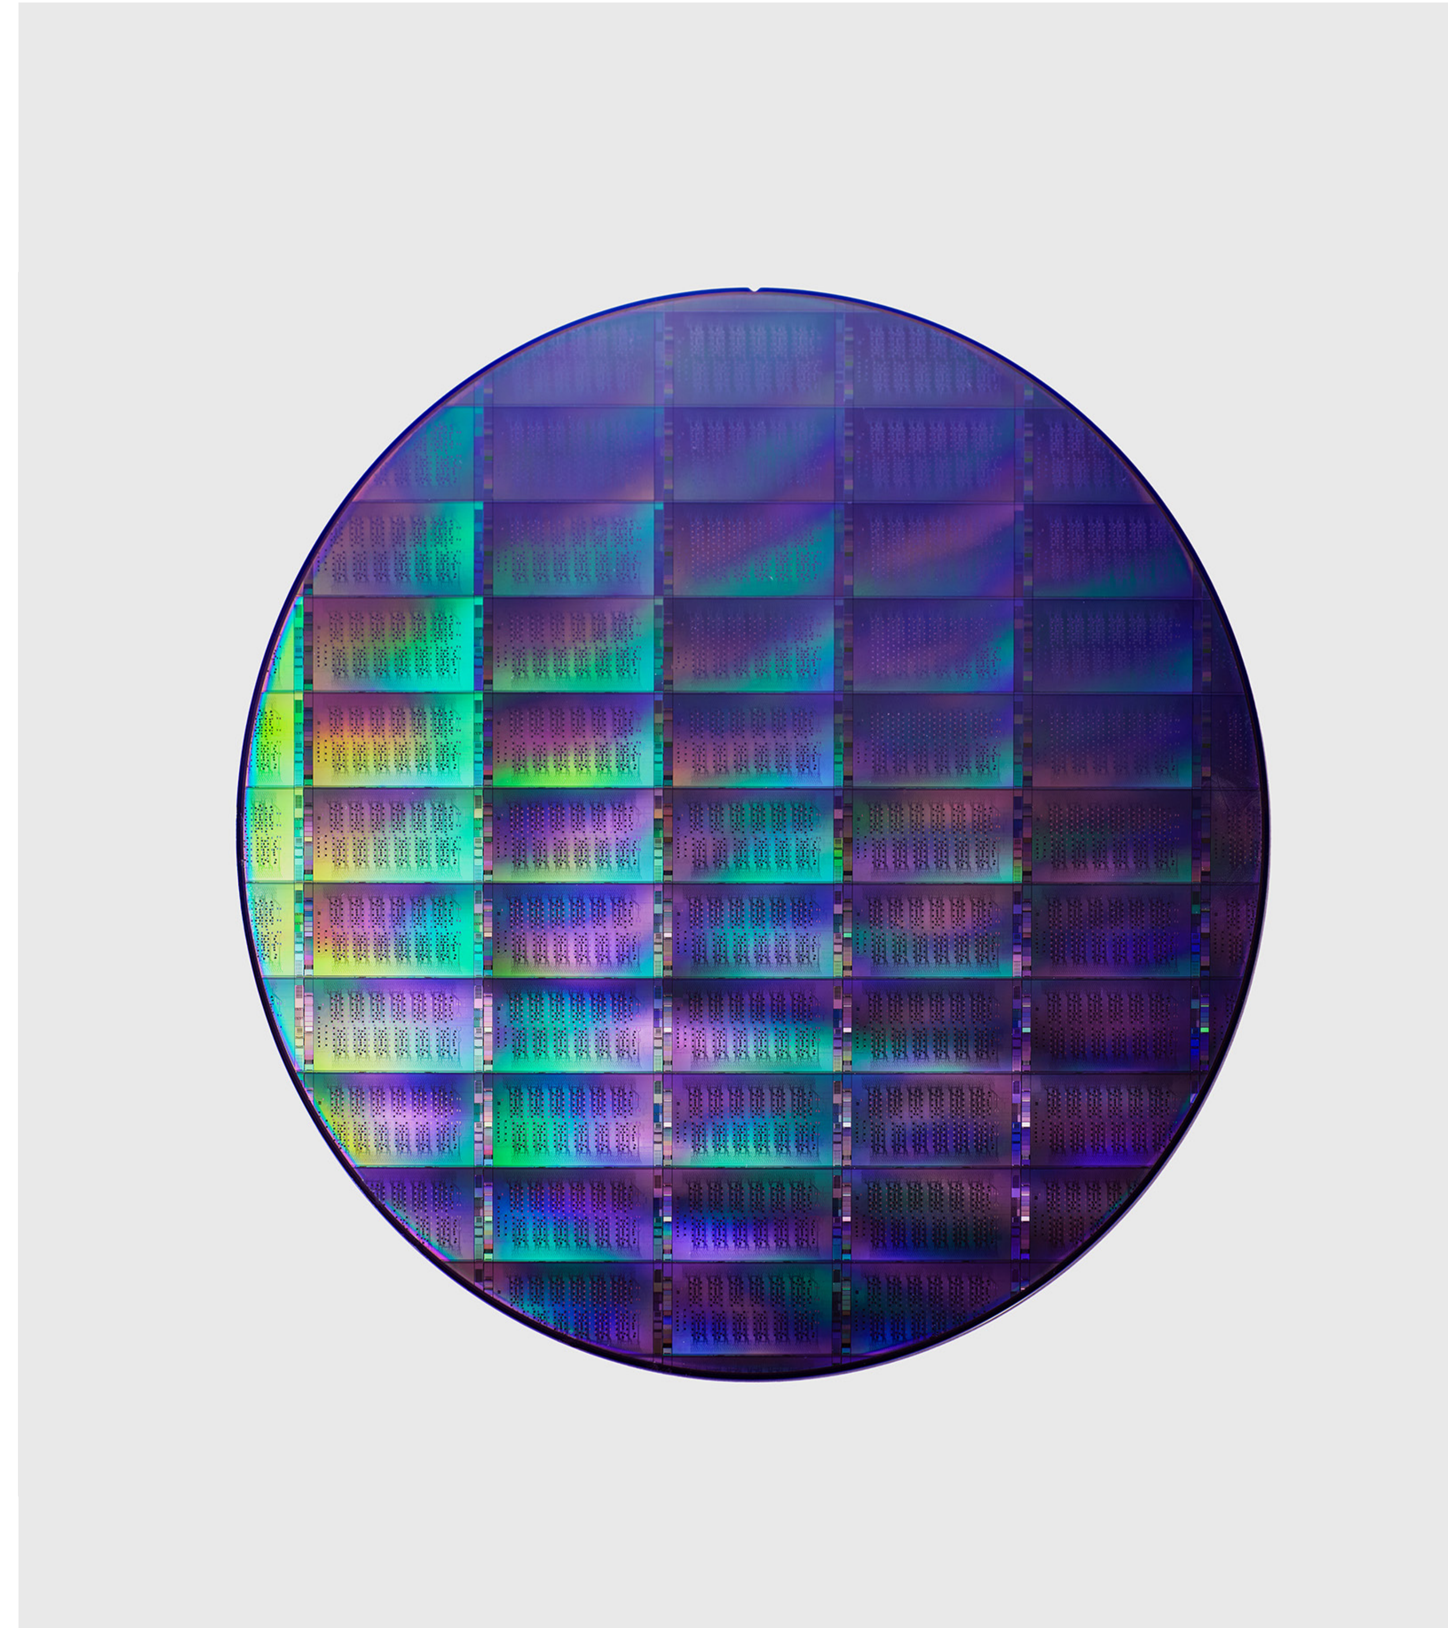

A Quantum Tool for Electrical Circuits

Differential equations study the mathematics of change—everything from waves crashing against a shoreline to the roiling plasma on the surface of the sun. They are immensely valuable with applications in drug discovery, financial modeling, structural engineering, and beyond—but to date, much work studying quantum differential equations solvers has remained in the realm of mathematical abstraction. Now, recent work from IBM researchers Arkopal Dutt, Anirban Chowdhury, Kristan Temme, and Hari Krovi has uncovered a provable speedup using a quantum differential algebraic equation solver to tackle a concrete, real-world problem: simulating the dynamics of electrical circuits consisting of resistors, inductors, and capacitors (aka RLC circuits).

Differential algebraic equations (DAEs) are essentially systems of differential equations incorporating the operations of linear algebra. The mathematical structure of RLC circuits makes them a prime example of DAEs. Current or voltage changes with time, while representations of the circuit still must satisfy specific constraints through these changes. However, the computational cost of turning an RLC circuit into a crunchable equation scales quickly as the size of the circuit increases, making it inefficient to simulate classically. At the same time, existing quantum algorithms designed for differential equations couldn't crack DAEs. So the team asked whether they could create such a quantum algorithm themselves.

In their recent paper, the team explored mathematical tools called projector-based methods that let them rewrite DAEs in a way that first runs the differential part, which extracts a simpler equation based on the initial conditions and the information about how the system changes over time, and the algebraic part, which uses this equation plus the available information about the system to calculate the value of unknown variables. They run each part as a quantum circuit, then combine the outputs, altogether allowing them to measure properties of interest.

The result was a technique that succeeds at solving DAE better than any classical method and demonstrates that under certain assumptions, solving DAEs is a BQP-hard problem, the hardest kind of classical problem that a quantum computer can tackle efficiently.

This work remains in the realm of near fault-tolerant quantum computers, but still has immense potential. Most notably, it's a step toward a quantum algorithm to improve SPICE, a widely used tool for simulating electrical circuits. Given an electrical circuit with values for resistance (R), inductance (L), and capacitance (C) and sources of current and voltage, this algorithm allows us to compute quantities like currents, voltages, energy, and power over time much faster than SPICE can today. Further, the classical algorithms in SPICE can simulate many other kinds of outputs, including nonlinear elements, which the team plans to investigate next.

Algorithm discovery is among our most critical efforts in the emerging quantum era as we establish real-world use cases that can expect quantum speedups. This new work is among the most significant algorithms yet uncovered by our team, and one which offers a whole lot to be excited about as quantum processors improve.

What's new at IBM Quantum?

From growing software capabilities to new open-access community initiatives, IBM Quantum has achieved a lot during the second quarter. Here's a quick snapshot of some key developments.

Scale molecular simulations with SQD
Discover the workflow behind a 12,635-atom result—the largest of its kind—built with our partners at Cleveland Clinic and RIKEN. Dive into the story behind the work and learn how to build fragment-based macromolecular simulations yourself. Read the blog and try out the new tutorial.

Bring real quantum computing into your classroom
With new Classroom Accounts, educators can incorporate real quantum hardware access into their coursework for up to 100 students. Request a free account for your course or share with an educator in your network. Learn more on the IBM Quantum blog.

Error detection with Qiskit Paulice
Improve the accuracy of your quantum computations with Qiskit Paulice, a new toolkit for boosting the fidelity of sampled distributions through hardware-efficient Pauli checks based on spacetime codes. Explore the Github repo to learn more.

New tools for error mitigation workflows
Build flexible error mitigation workflows with the new directed execution model, which lets you efficiently combine techniques like SLC, PEC, TREX, and post-selection in a single flow. Upgrade to qiskit-ibm-runtime 0.47.0 and read the tutorial to get started.

Qiskit SDK v2.4 adds speed and extensibility
Qiskit v2.4 introduces compiled Python extensions, expands the C API, and marks continued progress toward scalable fault-tolerant workflows. Learn more in the v2.4 release notes.

Prototype fermionic circuits with ffsim
Build and test fermionic quantum circuits with ffsim, a Python library designed for fast, symmetry-aware classical simulations of fermionic systems. Run larger problems and validate results before executing on quantum hardware. Explore the docs and try the tutorial.

Meet our Qiskit advocates
Our most engaged Qiskit advocates are now featured on the IBM Quantum website, highlighting their contributions and leadership in the global quantum community. Get to know them and share with others who may want to get involved.

Whether you want to experiment with the latest tools or share resources with your network, there are more ways than ever to engage with IBM Quantum. Explore these updates and stay connected as we continue building what's next in computing.

Latest updates to Qiskit

Over the past year, Qiskit development has focused on strengthening the SDK's core foundations—making it faster, more extensible, and better-equipped to handle the complexity of advantage-scale and, eventually, fault-tolerant workloads. For this past quarter's v2.4 release, updates centered on delivering "progress without disruption." With v2.5, we're continuing that same trajectory while pushing further to deliver more customizable, high-performance workflows. Let's take a look at some of the highlights:

Qiskit v2.4: progress without disruption

Compiled Python extensions.
Build and distribute high-performance Python extensions using Qiskit's C API. That means faster tooling with no custom setup required.

Improved fault-tolerant compilation.
A faster, more scalable pipeline delivers lower T gate counts and better performance for discrete-basis workflows, supporting ongoing research into fault-tolerant architectures.

Expanded C API capabilities.
Continued investment into Qiskit's C API broadens access across programming environments while enabling high-performance extensions and deeper integrations with external tools.

Faster QPY serialization.
Improved circuit serialization by rewriting it in Rust. This speeds up saving and loading for large workloads.

Qiskit v2.5: programmable performance

Dynamic circuit support in the C API.
New capabilities allow inspection of classical control-flow instructions, enabling hybrid Python–C workflows for dynamic circuits.

Pauli-based compilation pipeline.
A new preset pass manager to target Pauli-based computation representations, supporting emerging approaches to fault-tolerant compilation workflows.

Custom compiler pipelines.
New abstractions make it possible to build multi-stage compilation flows using different intermediate representations, giving users greater control over how circuits are optimized and transformed while supporting multiple levels of abstraction in quantum compilation.

Transpiler performance gains.
Updates across the transpiler—including algorithmic improvements to LightSabre, increased usage of multithreading, and new optimization passes—deliver faster compilation and higher-fidelity results, especially for larger and more complex circuits.

Qiskit developers at IBM and across the open-source community are continually evolving the stack, from delivering foundational improvements that "just work" to building a more flexible, high-performance platform that users can configure, extend, and optimize to suit their needs. As quantum workflows progress towards advantage—growing in scale and sophistication—these capabilities will play an essential role in bridging the gap between algorithm design and hardware execution.

The Coherence Times podcast is back!

What it is: IBM has launched season two of its quantum podcast, hosted by IBM Quantum editor in chief Ryan Mandelbaum. In each episode, Ryan takes listeners on a deep dive into the world of quantum computing through interviews with some of the industry's leading experts.

Why it matters: Whether you're an experienced quantum computational scientist or just a curious observer, the Coherence Times podcast is a great way to get an inside view of the latest developments in quantum hardware, software, algorithms research, and more. Subscribe wherever you listen to podcasts or check it out on the IBM Research YouTube channel.

Quantum computing in healthcare prize announced

What it is: The first Wellcome Leap Quantum for Bio (Q4Bio) challenge announced its winners with quantum startup Algorithmiq receiving the top $2 million prize. The Q4Bio initiative, launched in 2023, provided $40 million in research funding to teams developing quantum algorithms that could provide a near-term quantum advantage for healthcare applications. Five of the six finalists choose to use IBM quantum computers to support their research.

Why it matters: The work of the Q4Bio cohort reflects quantum computing's maturing ability to be applied to real-world, impactful problems and provides a glimpse at some of the healthcare advances we could see in the next few years. The approaches developed by the finalists were diverse and included work related to cancer treatments, genome encoding, and identifying bio-markers, among others.

Santander and IBM launch quantum challenge

What it is: Santander, in collaboration with IBM, Bluzec, and Oxentia Foundation, launched the Santander X Global Challenge | The Quantum AI Leap. The project will provide up to six startups with grants alongside access to AI and quantum computing expertise from IBM. Applicants are encouraged to build technologies at the intersection of recent advances in both quantum and AI.

Why it matters: The challenge reflects the potential for quantum and AI technologies to reinforce each others' impact across major industries by speeding up development and opening up new approaches to algorithms. The organizers call out use cases including applying AI to help develop quantum systems, using both technologies together to build novel applications, and AI-based approaches to mitigating the risk posed from quantum computing to common forms of encryption like RSA.

NYU partners with IBM for quantum postdoc program

What it is: New York University and IBM launched a postdoctoral program focused on quantum computing research in chemistry, computer science, engineering, materials science, physics, and optimization. The program will center on developing quantum algorithms and applications for quantum-centric supercomputing architectures. Participating students will be able to conduct research using IBM's quantum hardware at the Thomas J. Watson Research Center in Yorktown Heights, New York.

Why it matters: The New York quantum ecosystem is growing and this new program will support that growth while providing an expanded local quantum talent pipeline. The postdoctoral program builds on the work of the NYU Quantum Institute, which the university launched last year to research quantum information sciences subjects.

Startups make quantum advantage claims

What it is: Quantum computing startups Q-CTRL, Qedma, Kipu, and Algorithmiq made claims of quantum advantage using today's IBM Quantum hardware:

- Q-CTRL investigated Fermionic simulation on 120 qubits with over 9000 2-qubit gates and benchmarked them against classical tools, a significant quantum speedup over classical (2 minutes versus 100 hours)
- QEDMA studied Ising dynamcs with 51 qubits on ibm_boston with their QESEM software, finding a solution using 7 QPU + 250 CPU hours versus 720,896 CPU hours on the supercomputer Fugaku.
- Kipu demonstrated experimentally that their bias-field digitized counterdiabatic quantum optimization (BF-DCQO) algorithm on IBM's 156-qubit quantum computers can outperform the simulated annealing (SA) classical method. This advantage scales as the problem size grows.
- Algorithmiq is measuring a significant improvement in observable estimation for perturbed mirror circuits using QPUs versus CPU-only methods.

Why it matters: Achieving quantum advantage is essential to deliver broad real-world value from quantum computing. IBM expects quantum advantage to be achieved at some point in 2026 and it's possible that some of the claims already announced will withstand scrutiny. Expect the number of quantum advantage claims and counterclaims — a sign of healthy dialogue in the quantum community — to continue increasing throughout the year.

Cleveland Clinic 12,000+ Atom Experiment

What it is: Cleveland Clinic, RIKEN, and IBM debuted the largest heterogeneous quantum-classical (HQC) electronic-structure calculation to date, using up to 94 qubits. Researchers used a quantum-centric supercomputing workflow to simulate protein–ligand chemistry, with the largest simulation reaching over 12,000 atoms (30,000 orbitals).

Why it matters: This workflow performs at a level comparable with established classical methods and offers enormous potential for scaling. A better method for computational chemistry could provide enormous benefits to society; if researchers had a reliable method for predicting the behavior of new molecules before synthesizing and testing them in laboratories, the pace of pharmaceutical development, new materials science, and general chemistry research could all significantly accelerate.

IBM announces a new innovation center in Chicago

What it is: IBM and Governor JB Pritzker of Illinois announced the FutureNow Chicago delivery center at the Illinois Quantum and Microelectronics Park (IQMP). The new innovation and delivery hub will host an IBM Quantum System Two.

Why it matters: The center is expected to create around 750 jobs in the Chicago area related to AI, quantum, and other emerging tech fields. Chicago is also an important hub for quantum computing and the installation of an IBM Quantum System Two within the new center will support quantum research in the area.

IBM named to TIME's Most Influential Companies of 2026

What it is: IBM was named to TIME's Most Influential Companies of 2026 list. TIME singled out IBM quantum computing research and the significant real-world traction that the company has enabled. The publication also recognized IBM's future quantum plans and its steady progress toward delivering a large-scale fault-tolerant quantum computer, IBM Quantum Starling, by 2029.

Why it matters: Being named to the list demonstrates a growing recognition of quantum computing's importance to mainstream audiences. The publication highlighted IBM's quantum achievements across numerous vectors, including the large number of quantum systems deployed, Qiskit's expansive ecosystem, the development of post-quantum cryptography algorithms, and recent scientific advances IBM's quantum computers have supported.

MIT-IBM Computing Research Lab launches in Cambridge

What it is: IBM and MIT announced the launch of the MIT-IBM Computing Research Lab. The new lab evolves from the MIT-IBM Watson AI Lab, launched in 2017, expanding its scope to include quantum computing alongside artificial intelligence research. The goal of the lab is to develop new mathematical and computational approaches that go beyond purely classical methods.

Why it matters: The MIT-IBM Computing Research Lab will expand academic research in quantum computing and its intersection with AI. The MIT-IBM Watson AI Lab, the precursor to the new lab, helped fund over 210 research projects that resulted in more than 1,500 peer-reviewed articles.

Quantum computers used to simulate spin in materials

What it is: Researchers working with the Quantum Science Center (QSC) — based at Oak Ridge National Laboratory (ORNL) — released a paper on how spin currents behave in a one-dimensional quantum spin material. The team, led by IBM and Purdue University, used quantum computers to simulate spin transport behavior across a system, as opposed to relying on information gleaned from indirect measurements.

Why it matters: The work is another example of quantum computing contributing directly to useful scientific research. Understanding how spins behave in a material provides valuable information about its properties. Researchers are planning to use a similar approach to study more quantum effects in materials and expand their work to two-dimensional spin systems.

IBM develops LLM-guided framework to discover QEC codes

What it is: Researchers from IBM demonstrated a new LLM-guided evolutionary framework for discovering quantum error correction codes. The method is intended to help solve the challenge of finding the relatively few useful codes among an enormous number of potential algebraic formulations. The new workflow, which includes several layers of verification to filter for the best codes, has already been used to discover 465 novel candidates.

Why it matters: Quantum error correction codes are essential tools to support the development of fault-tolerant quantum computers — the better our understanding of the codes available, the more potential pathways we'll have for delivering useful quantum computing. The new workflow allows researchers to more quickly assess the potential of different families of codes and helps inform them where to look next.

Quantum Networking

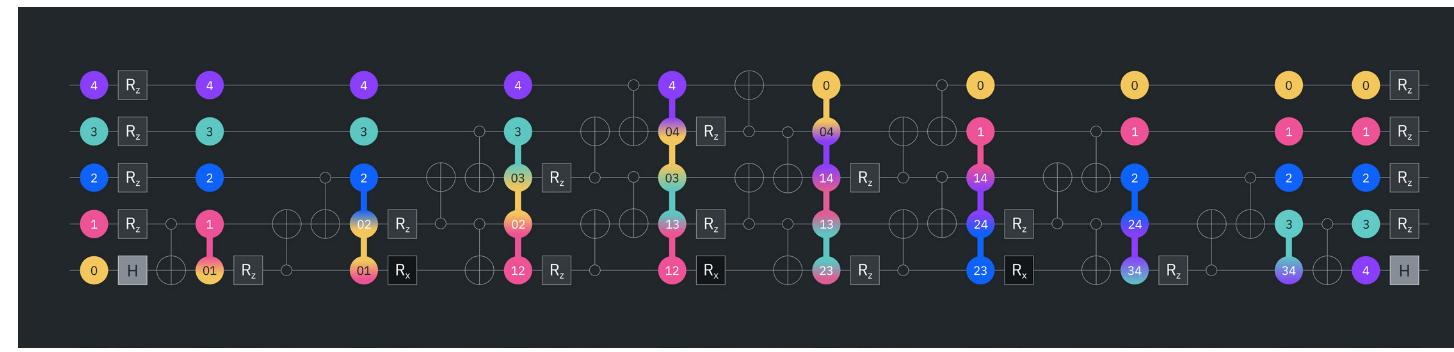

This image shows how ParityQC's Parity Twine method transfers information across a chip using efficient CNOT gates.

The quantum architecture company ParityQC has demonstrated a 52-qubit quantum Fourier transform (QFT) on an IBM Quantum Heron r3 processor, the largest reported to date. The work, described in a recent preprint, nearly doubles a 2024 QFT benchmark set on a trapped-ion quantum computer.

QFT is a fundamental building block for many quantum algorithms, including Shor's algorithm and quantum phase estimation. It is also hard to scale because QFT requires high connectivity between qubits, including distant ones. Superconducting processors mostly support interactions between neighboring qubits, so they rely on SWAP operations to move quantum states across the chip. Each SWAP adds gates, increasing circuit depth and accumulated error. At larger sizes, this routing overhead can dominate the computation.

ParityQC's method, called Parity Twine, changes what the circuit tracks. Rather than following the state of individual qubits, it tracks parity information—the relationships between qubits. That information moves along physical qubits through sequences of CNOT gates instead of SWAPs, which lowers gate count and circuit depth. At the level of how the computation is represented, all multi-qubit rotations reduce to single-qubit operations.

The team measured the results using a metric called process fidelity, which captures how accurately the full QFT was implemented across circuit depth, noise, and compilation strategy. They compared Parity Twine against the best previous implementation strategy and found higher process fidelity, with the gap widening super-exponentially as the number of qubits grew.

The demonstration drew on both algorithmic and hardware advances, and ParityQC credited the Heron r3 processor for enabling the results. This reflects how platform tools like Qiskit and the Qiskit transpiler provide a foundation that researchers across academia, startups, and industry build on.

These results do not yet translate into an application-level advantage. They do indicate what scalable performance on current hardware could look like, particularly for problems involving all-to-all interactions, such as optimization and simulation workloads. ParityQC will release an application function built upon the Parity Twine method later this year. It will feature the most efficient algorithm implementation yet for sparse and highly connected optimization problems and allow users to efficiently solve the problems end to end.

Q&A:

Paco Martin

Fran Cabrera

Jorge Carballo

A proof of concept of the IBM Quantum Experience from January, 2016.

Ten years ago, we put a quantum computer on the cloud for the first time. What's happened since is incredible—developers around the world are now using cloud-based quantum processing to explore new algorithms and applications as they work to bring useful quantum computing to the world. So rather than spotlighting just one person, we decided to gather three of the developers tasked with hooking these early processors up and building the first iteration of the IBM Quantum Experience back in 2016.

Coherence Times Observations | Q+A with IBM Quantum leadership

How would you describe your current role?

Paco: I'm the current global lead for Independent Software Vendors (ISVs) and Technical Integrations. I work between the business and engineering teams, and I'm responsible for making sure our partners can integrate with our services. I manage some accounts like SAP, ServiceNow, and other kinds of software-embedded vendors at IBM Quantum. Jorge, Fran, and I were the first software developers that IBM Quantum hired in 2016.
Fran: I'm a product manager, so I manage the priorities and all the features that you see on IBM Quantum Platform. I get feedback from users and design solutions for them.
Jorge: I'm an engineering manager for quantum services, and right now I'm located in New York. I'm responsible for the API that connects Qiskit and other libraries with IBM Quantum Platform. My team enables users to run quantum circuits and workflows at scale.

What was your experience like developing IBM Quantum Experience?

Fran: It started with very crazy conversations with Ismael Faro. Ismael was part of the research team at the Zurich lab at the time, and he was trying to form a team to make quantum accessible through the web. He contacted us, and I remember that the first conversations with him, Paco, and Jorge sounded crazy to me. Initially, I couldn't believe that a quantum computer existed and could be put online. All these ideas about programming a quantum computer through a visual interface — which now is IBM Quantum Composer — sounded like fantasy. It took me one or two weeks of reading papers to really understand what we were trying to build. We released the first version of the platform after developing for two and a half months.
Paco: In the beginning, we didn't know that Ismael's project was for quantum. It was like a secret; we needed to sign an NDA to see what it was. But we committed to the project and then Ismael explained to us that IBM wanted to put the five-qubit quantum system on the cloud. Step by step, we built IBM Quantum Experience and a community website similar to Stack Overflow, but for quantum. Then we started to build the first version of Qiskit. Year by year, we were releasing new software for the quantum systems. It was a very beautiful journey.
Jorge: It was a lot of work. We were starting at 7 a.m. and we were finishing at 8 p.m., doing crazy stuff every single day. But we enjoyed it because we were working in the same office.
Fran: The three of us rented an office in Corcoles, Spain, so we were all together. There was a lot of work, but also fun parts. This was like a secret project; you felt like you were participating in something that no one else was doing.

Looking back, what are you most proud of?

Paco: We built the front end, APIs, and software to connect with the IBM Quantum system. I'm most proud of [developing] this new, visual way to interact with a quantum system. For the first time, everyone could run something in a on a real quantum computer. I think that was a challenge for IBM in general: they needed to stabilize the hardware, they needed maintenance, and they needed it to be accessible for the user. The interface was my proudest accomplishment because it was something new that we created from scratch.
Fran: I would highlight three things: first, we put out the first platform that made accessible quantum computers; second, we built the first visual way to create and run a circuit on a quantum computer; and third, we participated in giving shape to the initial version of the Qiskit SDK.
Jorge: I would say the same. That first version — building this system to run something on a quantum computer — was crazy for us. The first three months, and probably the first year, we were building a lot of things with a velocity that wouldn't seem possible today.
Paco: We were like a small startup inside of IBM Research. We were releasing new features every week.

What do you do for fun outside of work?

Jorge: Right now, I have a five-year-old girl who probably takes up all my time outside of work; I spend most of my time with her.
Paco: I enjoy time with my family. I have two kids — one in university at UConn, another one in high school — so family time is the most exciting thing to me. I also love retro gaming. I have a collection with very old and very rare games; my favorite is Metal Slug.
Fran: I play video games, I read, and I have a dog. I love walking my dog outside of work.

Phase Kickback

Differential equations are the language we use to describe change. They govern the evolution of quantum systems, the flow of heat, the motion of fluids, and the dynamics of engineered systems ranging from control circuits to chemical reactions.

For decades, solving these equations efficiently has been a central challenge in scientific computing, driving advances in numerical methods and computer architectures alike. Quantum computing introduces a fundamentally new approach: rather than approximating dynamics step-by-step on a classical machine, a quantum system can, in principle, directly encode and evolve according to the same rules. But how far can we extend this approach? The research community has spent 30 years answering that very question.

1996

Seth Lloyd showed that quantum systems can efficiently simulate other quantum systems governed by the Schrödinger equation—establishing the concept of a universal quantum simulator. **More than a decade later, the HHL algorithm (2009) extended this idea**, showing that quantum computers could solve certain linear equations exponentially faster than classical methods. This raised a natural question: could quantum systems directly simulate the differential equations that appear in engineering and physics?

2014

Dominic Berry presents one of the first quantum algorithms tailored to solving linear inhomogeneous differential equations, using high-order numerical methods. This work demonstrates that quantum algorithms can go beyond static linear algebra and directly tackle time-dependent problems, while also introducing techniques to improve accuracy through higher-order approximations.

2017

Berry and collaborators introduce a significant advance to this approach with an algorithm based on truncated Taylor series expansions. This work achieves an exponential improvement in the dependence of the solution error on precision, marking an important step toward practical quantum differential equation solvers by reducing the cost of achieving high-accuracy solutions.

2020

While these early efforts focus primarily on linear systems, attention soon turns to the more challenging problem of nonlinear dynamics. **In 2020, Ilon Joseph proposes using the Koopman–von Neumann formulation to recast nonlinear classical phase-space dynamics as linear Schrödinger-type evolution on a Hilbert space.** This approach provides a pathway to simulate nonlinear and even non-Hamiltonian classical systems using quantum computers by embedding them into a linear quantum framework, effectively extending the reach of quantum differential equation solvers beyond traditional linear regimes.

Phase Kickback

2021

Jin-Peng Liu and collaborators introduce the first quantum algorithm capable of handling nonlinear differential equations more directly by utilizing Carleman linearization. This technique transforms nonlinear systems into higher-dimensional linear systems, enabling quantum algorithms originally designed for linear systems to be applied to nonlinear dynamics, including dissipative processes.

2022

At the same time, researchers begin to more carefully assess the practical limits and advantages of quantum approaches. **In 2022, Noah Linden, Ashley Montanaro, and Changpeng Shao analyze the state-of-the-art in quantum algorithms for solving the heat equation, a canonical partial differential equation.** Their results show that while quantum algorithms can offer advantages in certain settings, at most polynomial speedups can be achieved when extracting a complete classical description of the solution, highlighting an important distinction between quantum and classical outputs.

2023

The following year brings a series of refinements and new frameworks that expand both the scope and efficiency of these methods. **In February 2023, future IBM researcher Hari Krovi improves earlier analyses of quantum algorithms based on truncated Taylor series and extends their applicability to Lyapunov-stable systems.** This work broadens the class of differential equations that can be efficiently addressed while strengthening guarantees on stability and error.

In October 2023, researchers Dong An, Jin-Peng Liu, and Lin Lin introduce the Linear Combination of Hamiltonian Simulation (LCHS) framework. This method simulates non-unitary dynamics by expressing them as linear combinations of Hamiltonian simulation problems, avoiding dilated linear-system formulations and achieving optimal dependence on state-preparation cost. The approach provides a more direct and resource-efficient route to simulating dissipative and driven systems.

In December 2023, Ryan Babbush and collaborators demonstrate that Hamiltonian simulation techniques can be used to model networks of coupled classical oscillators. They further show that the corresponding energy estimation problem is BQP-complete, establishing both the power and the computational complexity of quantum approaches to certain classes of differential equations.

Phase Kickback

2024

These ideas are being extended to increasingly realistic physical systems and control problems. **In April 2024, Krovi develops quantum algorithms for damped and forced quadratic classical Hamiltonian systems and extends the framework to optimal-control problems. This includes applications to Riccati equations and linear-quadratic regulator problems**, linking quantum differential equation solvers to control theory and engineering.

2025

Recent work continues to refine the treatment of nonunitary and driven systems, particularly those with external forcing or stochasticity. **In August 2025, Shi Jin, Nana Liu, and Chuwen Ma extend and analyze Schrödingerization techniques for linear dynamical systems with inhomogeneous source terms.** Their work addresses key challenges in recovering original variables and maintaining stability in the presence of forcing or boundary terms.

In October 2025, IBM Fellow Sergey Bravyi and collaborators introduce a quantum algorithm for stochastic dissipative nonlinear systems with quadratic nonlinearities. They prove exponential speedups over classical algorithms under suitable assumptions and establish BQP-completeness for the simulation task, extending quantum advantage claims into noisy and nonlinear regimes.

In December 2025, Dong An, Andrew Childs, and Lin Lin refine the LCHS framework to obtain quantum algorithms for general linear nonunitary dynamics with optimal state-preparation cost and near-optimal dependence on matrix-query complexity and other key parameters, bringing these methods closer to theoretical optimality.

2026

IBMers Arkopal Dutt, Anirban Chowdhury, Kristan Temme, and Hari Krovi introduce a quantum algorithm for linear differential-algebraic equations and apply it to the simulation of RLC circuits. Their results demonstrate exponential speedups for well-structured circuits and establish a corresponding BQP-complete problem, underscoring the growing relevance of quantum differential equation solvers to real-world engineering systems.

Phase Kickback

References

1996 Universal Quantum Simulators, Lloyd, S.

2009 Quantum Algorithm for Linear Systems of Equations, Harrow, A. et al.

2014 High-order quantum algorithm for solving linear differential equations, Berry, D. W.

2017 Quantum Algorithm for Linear Differential Equations with Exponentially Improved Dependence on Precision, Dominic W. Berry et al.

2020 Koopman–von Neumann approach to quantum simulation of nonlinear classical dynamics, Joseph, I.

2021 Efficient quantum algorithm for dissipative nonlinear differential equations, Liu, J. et al.

2022 Quantum vs. Classical Algorithms for Solving the Heat Equation, Linden, N. et al.

Phase Kickback

References

2023

Improved quantum algorithms for linear and nonlinear differential equations, Krovi, H.

Linear combination of Hamiltonian simulation for nonunitary dynamics with optimal state preparation cost, An, D. et al.

Exponential Quantum Speedup in Simulating Coupled Classical Oscillators, Babbush, R. et al.

2024

Quantum algorithms to simulate quadratic classical Hamiltonians and optimal control, Krovi, H

2025

On Schrödingerization-Based Quantum Algorithms for Linear Dynamical Systems with Inhomogeneous Terms, Jin, S. et al.

Quantum simulation of a noisy classical nonlinear dynamics, Bravyi, S. et al.

Quantum Algorithm for Linear Non-unitary Dynamics with Near-Optimal Dependence on All Parameters, An, D. et al.

2026

Simulating dynamics of RLC circuits with a quantum differential-algebraic equations solver, Dutt, A. et al.

Q2 Events

July 13–24
Qiskit Global Summer School (Virtual)

July 14
ISMB - Q4Life Sciences Reception (Washington D.C. US)

QED-C Quantum Technology Showcase on the Hill (Washington D.C., US)

July 22-23
Global Quantum Forum (Chicago, IL, US)

July 22
Genesis Consortium Kickoff (Washington D.C., US)

August 8-14
Fault-Tolerant Quantum Technologies at Benasque Science Center (Benasque, Spain)

August 6-7
2026 California Technology & Policy Summit (Napa Valley, CA, US)

August 10-14
QSIM (Amsterdam, NL)

August 23-27
ACS Fall (Chicago, IL, US)

August 24-27
Superconducting Qubits and Algorithms (SQA) Conference (Gothenberg, Sweden)

September 9-10
Q2B Copenhagen (Copenhagen, Denmark)

September 13-18
IEEE Quantum Week (Toronto, Canada)

September 24-27
AFA Air Space & Cyber Conference (Washington D.C., US)

September 16
IBM Quantum Connect: Prague (Prague, Czech Republic)

September 21
IBM Quantum Connect: Brazil (Niterói, Brazil)

September 21-25
CBCTQ - Congresso Brasileiro de Ciências e Tecnologias Quânticas (Niterói, Brazil)

September 23-25
Quantum World Congress 2026 (College Park, Maryland, US)

September 28-30
Quantum Innovation Summit, Dubai (Dubai, United Arab Emirates)

September 28- October 1
Sibos 2026 (Miami, FL, US)

Error Mitigation

This summer marks the 2026 FIFA World Cup. Since our readers have heard a lot about IBM Quantum researchers—and researchers in the United States—we wanted to show our support for teams around the world by spotlighting a selection of physicists from some other representative nations.

01
(Canada)
Donna Strickland
Nobel laureate in Physics for developing chirped pulse amplification, enabling ultra-intense laser pulses.

02.
(France)
Alain Aspect
Nobel laureate in Physics for his foundational studies of entanglement, like establishing the violation of Bell inequalities.

03
(Ghana)
Francis Allotey
Made contributions to x-ray spectroscopy such as the “Allotey formalism” and was a founding fellow of the African Academy of Sciences.

04
(Brazil)
Cesar Lattes
Contributed to the discovery and characterization of the pion particle.

05
(Argentina)
Juan Martín Maldacena
Key figure in quantum gravity who first proposed AdS/CFT correspondence, linking quantum gravity with quantum field theories.

06
(Germany)
Max Planck
Nobel Laureate in Physics and a father of quantum theory via energy quantization.

07
(Egypt)
Ahmed Hassan Zewail
Nobel laureate in Chemistry for studying chemical reactions at the femtosecond timescale using ultra-fast laser flashes.

08
(South Korea)
Philip Kim
Performed groundbreaking research on the physics of carbon nanotubes and graphene.

09
(Japan)
Shinichiro Tomonaga
Nobel laureate in Physics for co-discovery of the renormalization method used in quantum electrodynamics.

10
(Australia)
Michelle Simmons
Pioneering physicist in the field of atomic electronics and CEO/ Founder of Silicon Quantum Computing.

ibm.com/quantum