TROY TECHNICAL

QUANTUM COMPUTING
WEEKLY REPORT

November 8 - November 14, 2026 | Issue 2026-W45

Verified computational advantage, error control and usable access

12	10
ANALYZED ARTICLES	PRIORITY THEMES

Issued 2026-11-14 | English edition | fixed approved cover artwork



Quantum Computing: the boundary moved from isolated claims to qualification evidence

12 English articles were reviewed independently from the Japanese edition. Ten themes are retained for management decisions.

SIGNAL 01 | QUBIT DEVICE

IBM Quantum Nighthawk: Specs and 1,000-Qubit Roadmap for 2027

IBM announced on November 11 that Nighthawk r3 is designed with 520 qubits that can be programmed and 10,000 operations, with a plan to link two processors into a 1,040-qubit system in 2027.

SIGNAL 02 | CONTROL STACK

IBM Starling: Fault-Tolerant Quantum Computing Roadmap to 2029

IBM reports reducing the Josephson-junction count required in a Starling processor design from roughly 5,000 to 2,000. At an assumed 99.95% junction yield, its model raises fully functional-chip yield from about one in twelve to one in three.

SIGNAL 03 | ERROR CORRECTION

KDDI: Quantum-AI Optimization for Uplink Speed Improvements

KDDI says it deployed quantum-annealing-based optimization across 683 Tokyo base stations from October 2, 2026. It reports about 80% fewer congested periods with uplink speeds below 1.5 Mbps in its commercial network.

SIGNAL 04 | SYSTEM ACCESS

QuEra, Harvard, and Anthropic: AI-Driven Quantum Circuit Optimization

QuEra, Harvard, and Anthropic paired Claude-generated circuits with an independent evaluator and reported up to a 39-fold reduction in a combined gate-and-qubit resource metric for specific building blocks.

DECISION

Focus this week's management attention on verified computational advantage, error control and usable access. Move only when the linked evidence can be reproduced under your own operating conditions.

THREE MANAGEMENT MOVES

1. Buyers: benchmark valuable workloads against classical alternatives, including total execution and validation cost.
2. Suppliers: tie components and control systems to fidelity, uptime, calibration and scaling evidence.
3. Executives: track error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EVIDENCE DISCIPLINE

FACT states what the collected English source reports. BUSINESS READ and ACTION are editorial interpretations. UNKNOWN lists information that the source file does not establish. A linked article is not treated as independent confirmation of every claim.

Numbers and evidence anchors that require conditions

520 A02	SOURCE-BOUND CONDITION Programmable qubits in IBM's planned Nighthawk r3; this is a design target, not delivered system capacity.
99.95% A07	SOURCE-BOUND CONDITION IBM's assumed Josephson-junction yield in its Starling chip model, not an observed production yield.
80% A09	SOURCE-BOUND CONDITION KDDI-reported reduction in the frequency of uplink speeds below 1.5 Mbps at 683 deployed Tokyo base stations.
39× A11	SOURCE-BOUND CONDITION Upper bound reported by QuEra and collaborators for a combined gate-and-qubit metric on specific circuit building blocks.
80 A10	SOURCE-BOUND CONDITION Qubits in Qunova's company-reported iron-sulfur-cluster chemistry benchmark; end-to-end advantage is unverified.

HOW TO USE THESE NUMBERS

Each value is shown with the condition stated in the collected English article. Do not compare values across different test methods, device sizes, operating temperatures or commercial stages without normalizing those conditions.

Where the value chain moved

Each stage combines one leading theme and one supporting theme. The report treats handoffs as evidence transfers, not decorative arrows.

1

Qubit device

P01/P06

IBM announced on November 11 that Nighthawk r3 is designed with 520 qubits that can be programmed and 10,000 operations, with.

2

Control stack

P02/P07

IBM reports reducing the Josephson-junction count required in a Starling processor design from roughly 5,000 to 2,000.

3

Error correction

P03/P08

KDDI says it deployed quantum-annealing-based optimization across 683 Tokyo base stations from October 2, 2026.

4

System access

P04/P09

QuEra, Harvard, and Anthropic paired Claude-generated circuits with an independent evaluator and reported up to a 39-fold.

5

Applications

P05/P10

Google researchers used AI agents and computer-checked mathematical proofs to significantly reduce the computing resources.

EVIDENCE REQUIRED AT EACH HANDOFF

Qubit device to Control stack	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.
Control stack to Error correction	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.
Error correction to System access	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.
System access to Applications	Transfer test conditions, acceptance criteria, failure data, owner and decision date. Recheck error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

COMMON CONTROL POINT

No theme moves to a commercial commitment until a named owner has reproduced the relevant evidence and documented error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

Priority map: maturity and next decision locus



P01	IBM Quantum Nighthawk: Specs and 1,000-Qubit Roadmap for 2027 System / adoption | cue: system
P02	IBM Starling: Fault-Tolerant Quantum Computing Roadmap to 2029 System / adoption | cue: system
P03	KDDI: Quantum-AI Optimization for Uplink Speed Improvements System / adoption | cue: deployed
P04	QuEra, Harvard, and Anthropic: AI-Driven Quantum Circuit Optimization Tech / process | cue: no system/business cue
P05	Google Quantum AI: AI-Driven Resource Reduction for Programs Tech / process | cue: no system/business cue
P06	Qunova HI-VQE: 80-Qubit Simulation Matches Classical Benchmark Tech / process | cue: no system/business cue
P07	SEAFOOD Platform: Osaka University's Quantum-HPC Integration Timeline System / adoption | cue: Platform
P08	QodeX Quantum and Partners: Forecasting Unstructured PDEs on 52-Qubit IBM Hardware System / adoption | cue: Partners
P09	Molecular Qubits: NVision and Ulm University's PIQC Roadmap Tech / process | cue: no system/business cue
P10	PSI's proof of the Aaronson-Ambainis conjecture System / adoption | cue: system

REPRODUCIBLE POSITIONING RULE

Horizontal position uses five discrete stages: Research, Prototype, Joint evaluation, Product adoption and Scale. The vertical bands classify the next decision locus: Tech / process, System / adoption, or Business / rules (including supply systems). Article-title cues take precedence; the overview is used only when the title is silent. These are categories, not impact ranks. Bubble diameter reflects the editorial score.

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P01 IBM Quantum Nighthawk: Specs and 1,000-Qubit Roadmap for 2027

A02 | Qubit device | Pilot

2026-11-11

FACT

IBM announced on November 11 that Nighthawk r3 is designed with 520 qubits that can be programmed and 10,000 operations, with a plan to link two processors into a 1,040-qubit system in 2027. Both r3 and the linked system are roadmap targets, not delivered capabilities.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P01.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

P02 IBM Starling: Fault-Tolerant Quantum Computing Roadmap to 2029

A07 | Control stack | Pilot

2026-11-12

FACT

IBM reports reducing the Josephson-junction count required in a Starling processor design from roughly 5,000 to 2,000. At an assumed 99.95% junction yield, its model raises fully functional-chip yield from about one in twelve to one in three. These are modeled yields, not a completed fault-tolerant system.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P02.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 18/25 | MATURITY 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P03 KDDI: Quantum-AI Optimization for Uplink Speed Improvements

A09 | Error correction | Commercial deployment

Prior deployment 2026-10-02

FACT

KDDI says it deployed quantum-annealing-based optimization across 683 Tokyo base stations from October 2, 2026. It reports about 80% fewer congested periods with uplink speeds below 1.5 Mbps in its commercial network. This is not evidence of gate-model quantum-computing advantage.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P03.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 5/5 | DECISION LOCUS: SYSTEM / ADOPTION

P04 QuEra, Harvard, and Anthropic: AI-Driven Quantum Circuit Optimization

A11 | System access | Pilot

2026-11-10

FACT

QuEra, Harvard, and Anthropic paired Claude-generated circuits with an independent evaluator and reported up to a 39-fold reduction in a combined gate-and-qubit resource metric for specific building blocks. The factor does not generalize to all algorithms or measured hardware runtime.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P04.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: TECH / PROCESS

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P05 Google Quantum AI: AI-Driven Resource Reduction for Programs

A04 | Applications | Pilot

2026-11-10

FACT

Google researchers used AI agents and computer-checked mathematical proofs to significantly reduce the computing resources required for quantum programs. The team achieved resource reductions of up to elevenfold for fusion simulations, potentially accelerating the timeline for practical scientific quantum applications.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P05.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: TECH / PROCESS

P06 Qunova HI-VQE: 80-Qubit Simulation Matches Classical Benchmark

A10 | Qubit device | Joint evaluation

2026-11-12

FACT

Qunova said on November 12 that its HI-VQE workflow combining IBM quantum hardware and Fugaku reached chemical accuracy against a classical reference for one 80-qubit iron-sulfur-cluster calculation. This is a company-reported single-problem benchmark; end-to-end speed and cost advantage are unverified.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P06.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 3/5 | DECISION LOCUS: TECH / PROCESS

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P07 SEAFOOD Platform: Osaka University's Quantum-HPC Integration Timeline

A06 | Control stack | Joint evaluation

2026-11-13

FACT

The University of Osaka has developed SEAFOOD, a hybrid computing infrastructure linking the OCTOPUS supercomputer with domestic superconducting quantum processors. Scheduled for formal operation in December 2026, the platform enables partitioned workloads to address the computational costs of simulating strongly correlated electronic structures.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P07.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 3/5 | DECISION LOCUS: SYSTEM / ADOPTION

P08 QodeX Quantum and Partners: Forecasting Unstructured PDEs on 52-Qubit IBM Hardware

A05 | Error correction | Pilot

2026-11-13

FACT

A research collaboration has demonstrated gradient-free forecasting of nonlinear partial differential equations (PDEs) on unstructured engineering meshes using a 52-qubit IBM Quantum Heron r3 processor. The hybrid quantum reservoir network (QRN) reduced forecast error by 21% to 41% compared to classical latent linear baselines. This approach decouples complex boundary geometries from quantum circuit topology, enabling the simulation of fluid dynamics on non-uniform grids.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P08.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

FACT is source-bound. BUSINESS READ, ACTION and UNKNOWN are explicit editorial layers.

P09 Molecular Qubits: NVision and Ulm University's PIQC Roadmap

A03 | System access | Pilot

2026-11-09

FACT

Researchers from NVision Quantum Technologies and Ulm University have developed a single-molecule qubit that can be individually addressed and connected via light. This breakthrough enables a spin-photon interface at the single-molecule level, providing a scalable path toward quantum computers with hundreds of thousands of qubits.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P09.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 20/25 | MATURITY 2/5 | DECISION LOCUS: TECH / PROCESS

P10 PSI's proof of the Aaronson-Ambainis conjecture

A08 | Applications | Pilot

Prior release 2026-11-06

FACT

Physical Superintelligence said on November 6 that its research system, using a non-public Anthropic model, found a proof of the Aaronson-Ambainis conjecture. The paper and Lean formalization are mathematical evidence; they do not demonstrate a speedup on a commercial quantum computer.

WHY IT MATTERS

The decision relevance is verified computational advantage, error control and usable access. The signal should be tested at the application and operating-system level, not accepted from a headline alone.

BUSINESS READ

For operating companies, compare the reported change with current qualification, sourcing and product-roadmap assumptions. For suppliers, tie components and control systems to fidelity, uptime, calibration and scaling evidence.

ACTION

Within 30 days, benchmark valuable workloads against classical alternatives, including total execution and validation cost; record the evidence and owner against P10.

UNKNOWN

Still unproven or incompletely stated: error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.

EDITORIAL SCORE 19/25 | MATURITY 2/5 | DECISION LOCUS: SYSTEM / ADOPTION

Playbook and 0-5 year roadmap

STAKEHOLDER	NOW | 0-30 DAYS	NEXT | 1-12 MONTHS	LATER | 1-5 YEARS
Operating companies	benchmark valuable workloads against classical alternatives, including total execution and validation cost	Run production-like qualification with explicit acceptance and stop criteria.	Build a portfolio and architecture that can absorb technology and supplier changes.
Materials, equipment and technology suppliers	tie components and control systems to fidelity, uptime, calibration and scaling evidence	Create shared reliability, process-window and failure-analysis data with customers.	Standardize interfaces, traceability, lifecycle support and recovery services.
Trading, investment and legal teams	Map error rates, logical qubits, workload advantage, uptime, cost and roadmap timing.	Contract for capacity, quality, change notice, liability, alternatives and exit.	Diversify regional, technical and customer concentration risk.

ROADMAP

NOW | 0-30 DAYS
Evidence ledger
Owner: Strategy + quality
Deliverable: claim, source, condition, owner

NOW | 0-30 DAYS
Risk screen
Owner: Procurement + legal
Deliverable: unknowns, alternatives, stop

NEXT | 1-12 MONTHS
Joint qualification
Owner: R&D + supplier
Deliverable: process window and failure data

NEXT | 1-12 MONTHS
Commercial gate
Owner: Business + finance
Deliverable: unit economics and capacity

LATER | 1-5 YEARS
Platform strategy
Owner: Executive team
Deliverable: portfolio and interface roadmap

LATER | 1-5 YEARS
Service layer
Owner: Operations + channel
Deliverable: monitoring, maintenance and recovery

GO / NO-GO GATES

GATE	OWNER	REQUIRED EVIDENCE	NO-GO CONDITION
G1 Evidence	**Quality**	Source, test conditions and reproducibility	No traceable evidence
G2 Integration	**R&D**	Interface, process window and failure analysis	Cannot meet system conditions
G3 Commercial	**Business**	Capacity, total cost, contract and alternatives	Economics or supply cannot be supported
G4 Lifecycle	**Operations**	Monitoring, maintenance, change notice and exit	No accountable operating plan

ANALYZED ARTICLES | ITEMS 01-12

Every title and URL is displayed in full. URLs wrap without shortening and remain clickable. A verified-reference label means the translated HTML did not retain its original source URL.

ID	FULL ARTICLE TITLE	FULL SOURCE / VERIFIED REFERENCE URL - CLICKABLE
A01	Quantum Process Tomography: Optimal Query Complexity for 2026	https://doi.org/10.1038/s41567-026-03499-7
A02	IBM Quantum Nighthawk: Specs and 1,000-Qubit Roadmap for 2027	https://jp.newsroom.ibm.com/2026-11-11-ibm-quantum-nighthawk-demonstrates-new-levels-of-speed-and-quality,-with-plans-to-scale-beyond-1,000-qubits-in-2027
A03	Molecular Qubits: NVision and Ulm University's PIQC Roadmap	https://www.chemistryworld.com/news/light-connected-single-molecule-qubit-hints-at-disruption-to-quantum-computing-tech/4024467.article
A04	Google Quantum AI: AI-Driven Resource Reduction for Programs	https://thequantuminsider.com/2026/11/10/google-researchers-use-ai-and-mathematical-proofs-to-cut-quantum-computing-requirements/
A05	QodeX Quantum and Partners: Forecasting Unstructured PDEs on 52-Qubit IBM Hardware	https://quantumcomputingreport.com/qodex-quantum-nvidia-rtx-bbn-and-u-s-navy-researchers-forecast-unstructured-pdes-on-52-qubit-ibm-hardware/
A06	SEAFOOD Platform: Osaka University's Quantum-HPC Integration Timeline	https://quantumcomputingreport.com/university-of-osaka-integrates-domestic-superconducting-quantum-processors-with-octopus-supercomputer-via-seafood-platform/
A07	IBM Starling: Fault-Tolerant Quantum Computing Roadmap to 2029	https://www.solnews.co.kr/news/articleView.html?idxno=85172
A08	Anthropic AI Model: Proof of Aaronson-Ambainis Conjecture	https://www.businessinsider.jp/article/2611-anthropic-quantum-computing-breakthrough-aaronson-ambainis-math-physics-psi/
A09	KDDI: Quantum-AI Optimization for Uplink Speed Improvements	https://ai.watch.impress.co.jp/docs/news/2155817.html
A10	Qunova HI-VQE: 80-Qubit Simulation Matches Classical Benchmark	https://thequantuminsider.com/2026/11/12/qunovas-hybrid-computing-method-matches-classical-benchmark-for-iron-sulfur-cluster/
A11	QuEra, Harvard, and Anthropic: AI-Driven Quantum Circuit Optimization	https://thequantuminsider.com/2026/11/10/quera-harvard-anthropic-ai-quantum-circuit-resource-requirements/
A12	IBM, RIKEN, and Cleveland Clinic: 12,000-Atom HQC Simulation	https://quantumcomputingreport.com/cleveland-clinic-riken-and-ibm-cross-12000-atom-barrier-in-heterogeneous-quantum-classical-drug-binding-simulations/

QuantumComputing — Selected Articles

Date: 2026-11-15

Articles: 12

Table of Contents

#01 Quantum Process Tomography: Optimal Query Complexity for 2026

#02 IBM Quantum Nighthawk: Specs and 1,000-Qubit Roadmap for 2027

#03 Molecular Qubits: NVision and Ulm University's PIQC Roadmap

#04 Google Quantum AI: AI-Driven Resource Reduction for Programs

#05 QodeX Quantum and Partners: Forecasting Unstructured PDEs on 52-Qubit IBM Hardware

#06 SEAFOOD Platform: Osaka University's Quantum-HPC Integration Timeline

#07 IBM Starling: Fault-Tolerant Quantum Computing Roadmap to 2029

#08 Anthropic AI Model: Proof of Aaronson-Ambainis Conjecture

#09 KDDI: Quantum-AI Optimization for Uplink Speed Improvements

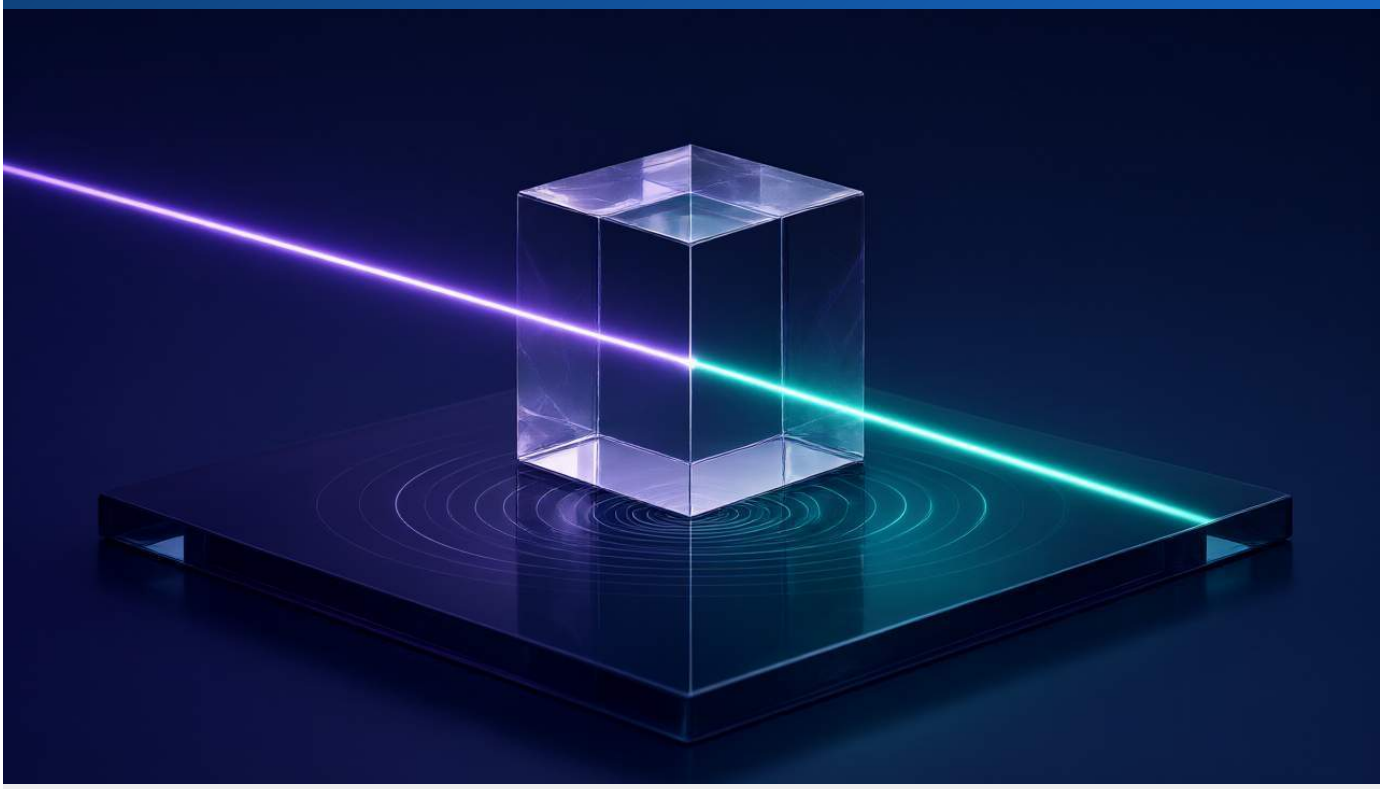
#10 Qunova HI-VQE: 80-Qubit Simulation Matches Classical Benchmark

#11 QuEra, Harvard, and Anthropic: AI-Driven Quantum Circuit Optimization

#12 IBM, RIKEN, and Cleveland Clinic: 12,000-Atom HQC Simulation

#01 Quantum Process Tomography: Optimal Query Complexity for 2026

Published November 12, 2026 Nature Physics



OVERVIEW

Researchers have established the optimal query complexity for learning unknown quantum channels using the diamond distance metric. The study proves that the number of channel uses required scales with the degrees of freedom of the channel, effectively reducing process tomography to pure-state tomography. This provides a unified framework for calibrating and validating noisy quantum devices across various hardware platforms.

IN DEPTH

Key Findings

The research identifies the precise number of invocations required to learn an unknown quantum channel to a fixed constant accuracy in diamond distance. The study demonstrates that the query complexity scales with the degrees of freedom of the unknown channel, proving that this parameter-count dependence is both sufficient and necessary. Specifically, for channels with input dimension d_{in}, output dimension d_{out}, and a Kraus rank of at most k, the algorithm achieves a complexity scaling of $\Theta(d_{\text{in}} d_{\text{out}} k)$.

Technical Details

The proposed algorithm utilizes a strategy of performing optimal tomography on the Choi-state representation of the quantum process. By doing so, the researchers show that quantum channel tomography can be reduced to the problem of pure-state tomography. This approach resolves a significant gap in previous knowledge; while the number of independent parameters needed to describe a general channel is $\Theta(d^4)$, previous tomography schemes for arbitrary d-dimensional channels required $\widetilde{O}(d^6)$ uses. The new findings align the upper bound with the information-theoretic lower bounds of $\tilde{\Omega}(d^4)$ for constant accuracy.

Background and Context

Quantum process tomography is a critical primitive for the calibration, validation, and debugging of noisy quantum devices, including those based on superconducting qubits, trapped ions, and photonic architectures. While quantum state tomography—learning a state in trace distance—has been well understood for roughly a decade, the corresponding question for general quantum channels remained open except for special cases like unitary channels. The diamond distance is used as the standard metric because it captures the worst-case distinguishability between two quantum processes across arbitrary input states and measurements.

Significance and Outlook

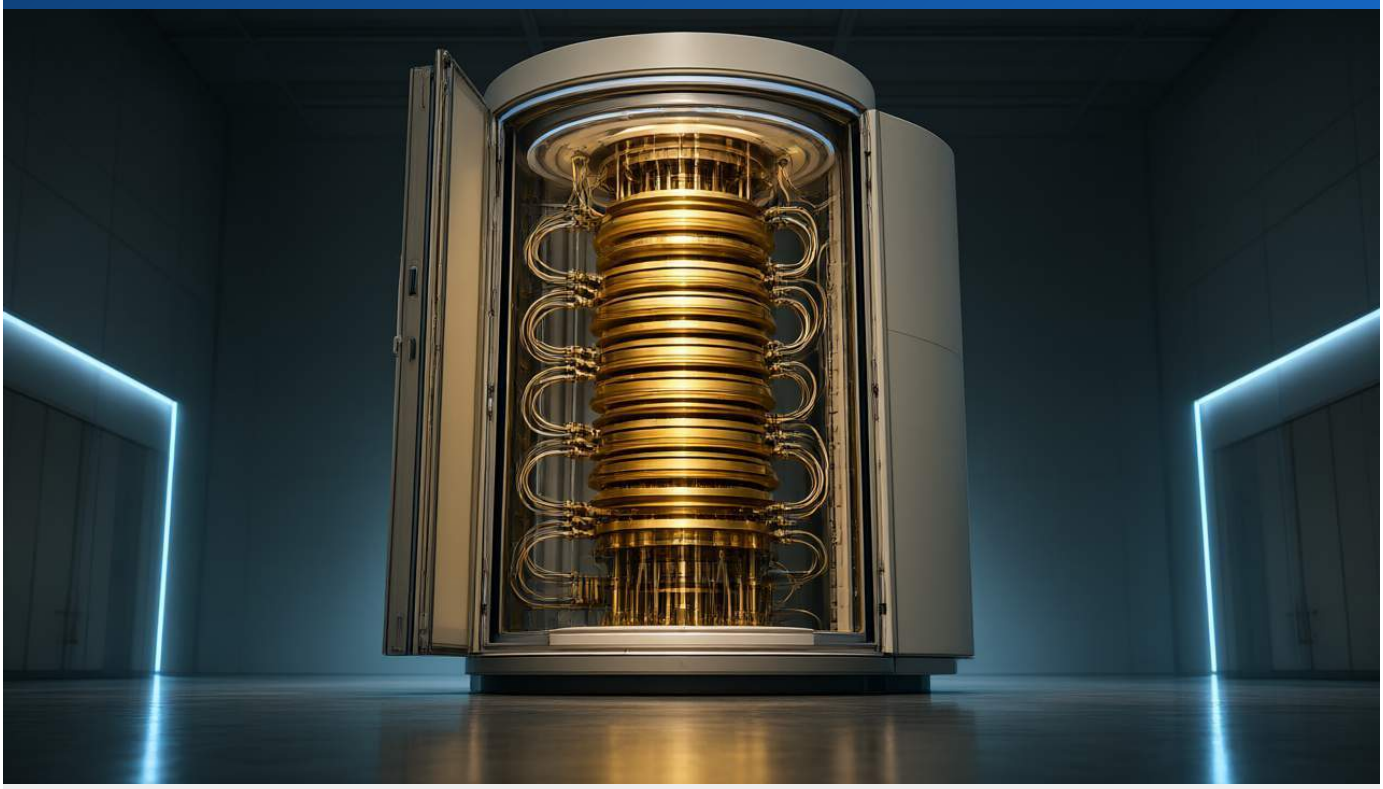
Because quantum channels subsume measurements, unitaries, and states as special cases, this protocol provides a unified framework for the tomography of all these tasks. By establishing a classical surrogate of unknown open-system dynamics that can replace the true process in future experiments with a prescribed worst-case error, this work optimizes the resources required for quantum device characterization.

Source: https://doi.org/10.1038/s41567-026-03499-7

Collected: November 15, 2026 | Automated Research System (AI)

#02 IBM Quantum Nighthawk: Specs and 1,000-Qubit Roadmap for 2027

Published November 11, 2026　IBM 日本 ニュースルーム　USA



OVERVIEW

IBM has introduced the Quantum Nighthawk r3 processor, designed to execute 10,000 quantum operations across 520 programmable qubits. The company plans to scale to a single system exceeding 1,000 qubits by 2027 by connecting two Nighthawk r3 processors. This progression focuses on combining high-fidelity operations with the execution speed necessary for industrial-scale workloads.

IN DEPTH

Key Findings

The IBM Quantum Nighthawk r3 processor is engineered to support 10,000 quantum operations on 520 programmable qubits, utilizing a total of 2,034 physical qubits. IBM intends to connect two r3 processors to create a single system with 1,040 qubits by 2027. This follows the deployment of Nighthawk r2, which supports over 7,500 operations on 120 qubits and generates up to 100,000 samples per second.

Technical Details

To enable scaling, IBM is implementing "L-couplers," which provide 10 chip-to-chip connections while maintaining the square-lattice connectivity and high-speed operations of the r2. Previous tests with Nighthawk r1 using four L-couplers achieved 99.3% fidelity in quantum state transfer. Additionally, IBM has demonstrated the ability to link two cryogenic modules into a single system, reaching temperatures below 15 millikelvin in under five days.

Background and Context

Industrial applications require high sample volumes—often millions to billions—to be practical. Nighthawk r2 has already demonstrated significant performance gains over the previous Heron processor:

- E.ON reported a 15x speedup in quantum machine learning for power grid optimization.
- Experiments with Oak Ridge National Laboratory showed a 12x speedup in simulating magnetic materials.
- Simulations of non-homogeneous quantum materials with Algorithmiq were 8x faster than on Heron, requiring an estimated 240x more memory for classical tensor network methods to achieve equivalent accuracy.

Furthermore, the IBM Quantum Platform now integrates High-Performance Computing (HPC) with quantum workflows. This integration allowed researchers from IBM and RIKEN to execute sample-based quantum diagonalization (SQD) 5x faster and perform 100x more calculations than in previous studies.

Significance and Outlook

The shift toward modularity and HPC integration is designed to facilitate fault-tolerance and error correction. IBM has established a co-design platform with AMD and Oak Ridge National Laboratory, physically co-locating Nighthawk processors with the Umbra HPC system, which features AMD Instinct MI355X GPUs and 5th Gen AMD EPYC CPUs.

Source: https://jp.newsroom.ibm.com/2026-11-11-ibm-quantum-nighthawk-demonstrates-new-levels-of-speed-and-quality,-with-plans-to-scale-beyond-1,000-qubits-in-2027

Collected: November 15, 2026 | Automated Research System (AI)

#03 Molecular Qubits: NVision and Ulm University's PIQC Roadmap

Published November 09, 2026 Chemistry World Germany



OVERVIEW

Researchers from NVision Quantum Technologies and Ulm University have developed a single-molecule qubit that can be individually addressed and connected via light. This breakthrough enables a spin-photon interface at the single-molecule level, providing a scalable path toward quantum computers with hundreds of thousands of qubits.

IN DEPTH

Key Findings

Researchers in Germany have designed a molecule, di([1,1′-biphenyl]-4-yl)carbene, that serves as a quantum bit (qubit) capable of being individually addressed and connected by light. This achievement overcomes a previous limitation where scientists could only read groups of molecules together, averaging their values. The development provides a manufacturable route toward large-scale quantum computing by utilizing photonic integrated circuits (PIQC) and existing foundry processes.

Technical Details

The molecular qubit exhibits four critical properties required for individual addressability:

- Fluorescence occurs only when electrons are in one of three specific spin states.
- These spin states remain stable long enough for quantum computation.
- The fluorescence is bright enough for reliable readout.
- The color is stable enough to enable quantum entanglement.

The molecules were synthesized from a precursor and embedded in a di([1,1′-biphenyl]-4-yl)methanone crystal to ensure precise positioning on chips. NVision intends to use electron spins for light-based connectivity and nuclear spins within the molecules as long-lived memories, exploring the inclusion of nitrogen atoms since the most common carbon isotope lacks nuclear spin.

Background & Context

While conventional computers use binary bits (0 or 1), qubits utilize quantum superposition to handle information more efficiently, which is particularly valuable for simulating molecules without compromising accuracy for computational cost. Current qubit technologies—including trapped ions, superconducting circuits, and neutral atoms—can connect up to 6100 qubits but have only detected and corrected errors on a maximum of eight. To be truly useful, systems must scale to hundreds of thousands of qubits.

Significance & Outlook

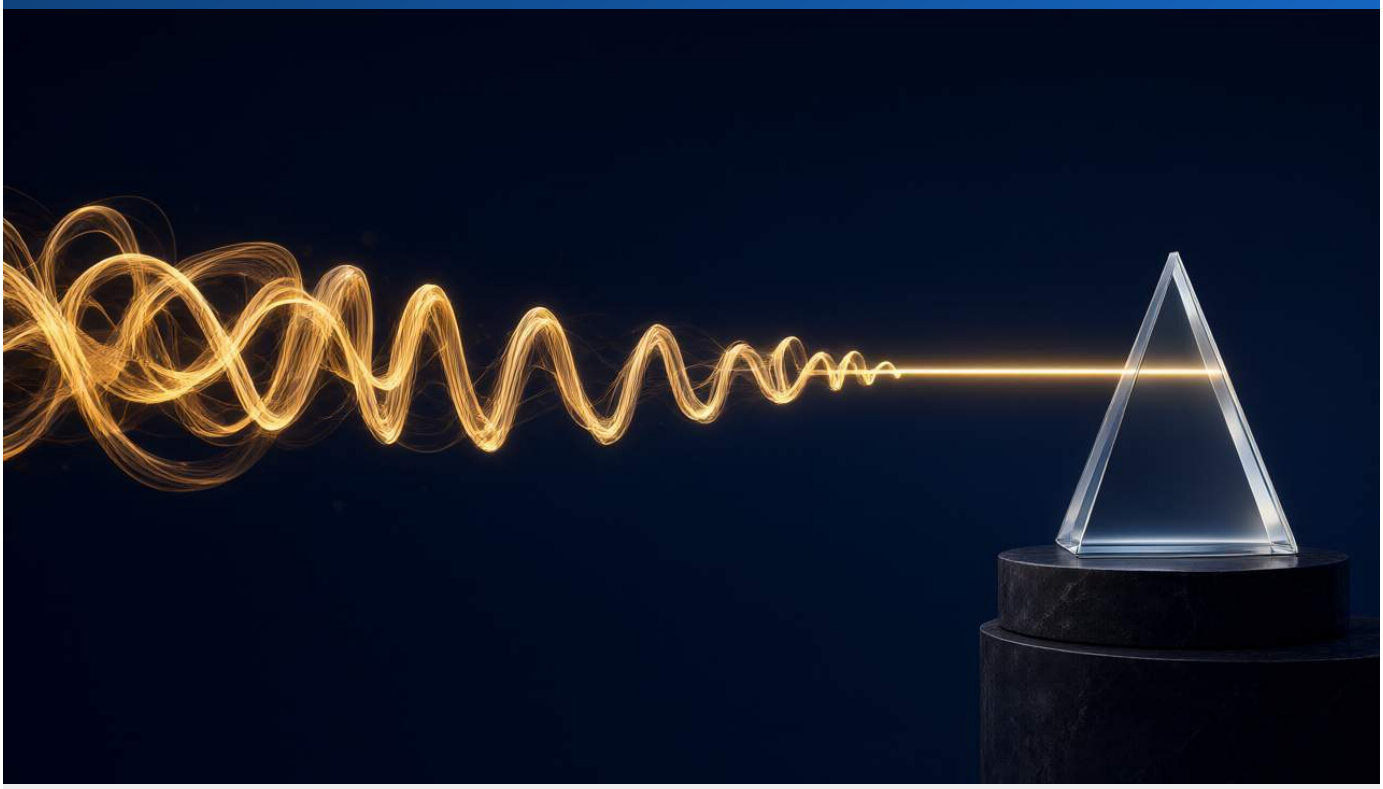
The ability to deposit these molecules as thin films on photonic integrated circuits allows for a scalable architecture. NVision is currently working with partner HyperLight to integrate these molecules onto photonic chips. This approach aims to solve the scaling challenge facing other quantum modalities by leveraging existing foundry processes rather than requiring new factory construction.

Source: https://www.chemistryworld.com/news/light-connected-single-molecule-qubit-hints-at-disruption-to-quantum-computing-tech/4024467.article

Collected: November 15, 2026 | Automated Research System (AI)

#04 Google Quantum AI: AI-Driven Resource Reduction for Programs

Published November 10, 2026 The Quantum Insider USA



OVERVIEW

Google researchers used AI agents and computer-checked mathematical proofs to significantly reduce the computing resources required for quantum programs. The team achieved resource reductions of up to elevenfold for fusion simulations, potentially accelerating the timeline for practical scientific quantum applications.

IN DEPTH

Key Findings

Google researchers from Google DeepMind and Google Quantum AI have demonstrated a method to reduce the combined qubit-and-gate requirements for complex scientific calculations. The most significant improvement was seen in a fusion-related simulation, where resource costs were reduced roughly elevenfold. Additionally, the team achieved an eightfold reduction for photochemistry calculations and a 2.5-fold reduction for a drug-metabolism model.

Technical Details

The researchers utilized a system where AI agents proposed quantum programs and provided mathematical proofs of correctness and cost. To ensure accuracy, the team built a quantum programming library in Lean, a system that checks mathematical proofs. This approach allows for the verification of programs that are too large for conventional computers to simulate.

- **Cost Metric:** Resource cost was measured by multiplying the number of logical qubits by the number of T gates.
- **Fusion Simulation Specs:** In a simulation involving one alpha particle, 218 electrons, and 216 fixed protons, the improved program required 4,374 logical qubits, down from a baseline of 5,650.
- **Verification:** A separate checking system was implemented to prevent AI agents from achieving apparent improvements by weakening theorems or changing the original task.

Background & Context

The study, posted on arXiv, addresses the gap between current hardware and the requirements of complex algorithms. While early fault-tolerant processors are projected to handle roughly a million operations, many scientific applications require billions to quadrillions of logical operations. The development process involved teams of AI agents sharing successful programs and partial proofs via a common wiki and virtual conferences.

Significance & Outlook

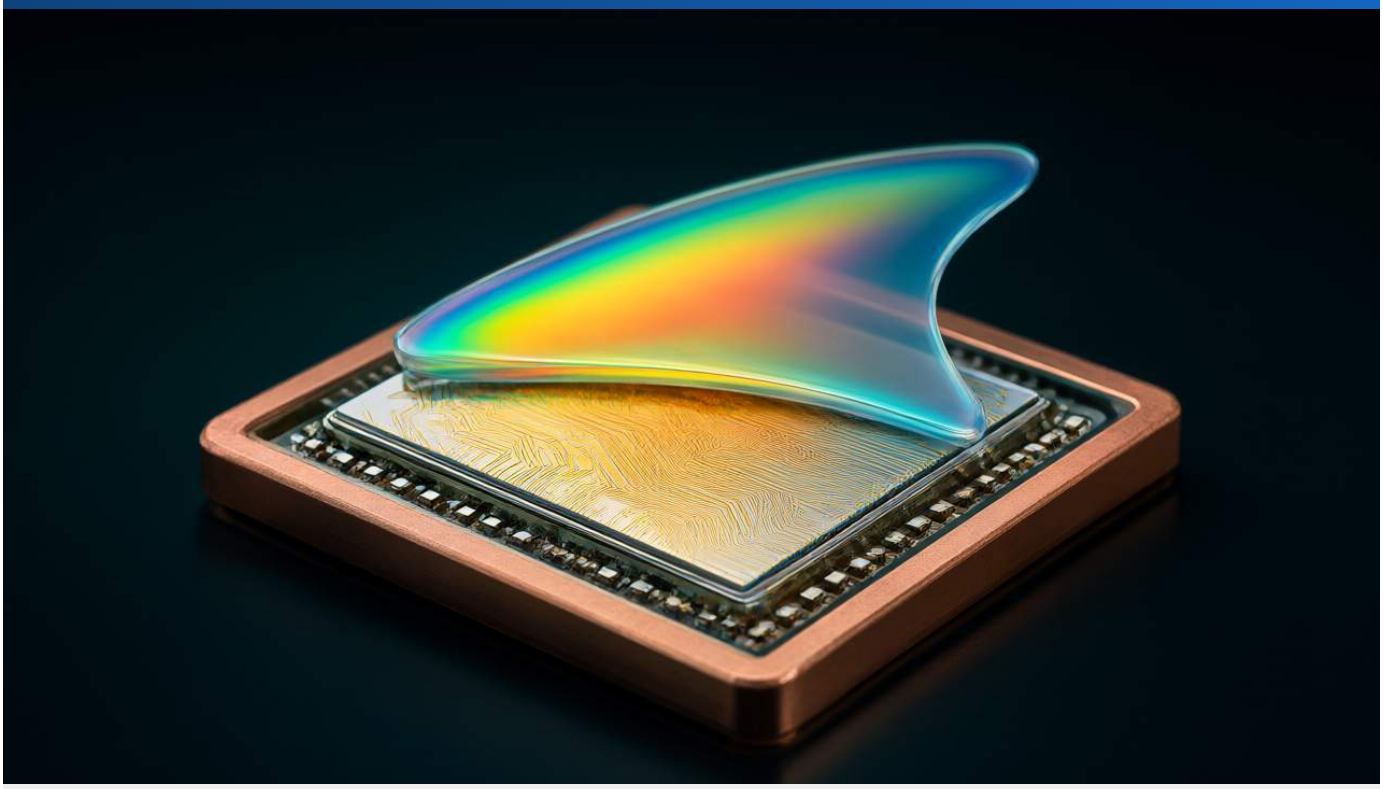
By using AI to optimize quantum programs, Google has shown that software improvements can complement hardware advances. While the current scientific programs still exceed the capabilities of projected early fault-tolerant quantum computers, these reductions bring demanding chemistry and fusion calculations closer to practical implementation.

Source: https://thequantuminsider.com/2026/11/10/google-researchers-use-ai-and-mathematical-proofs-to-cut-quantum-computing-requirements/

Collected: November 15, 2026 | Automated Research System (AI)

#05 QodeX Quantum and Partners: Forecasting Unstructured PDEs on 52-Qubit IBM Hardware

Published November 13, 2026 Quantum Computing Report USA



OVERVIEW

A research collaboration has demonstrated gradient-free forecasting of nonlinear partial differential equations (PDEs) on unstructured engineering meshes using a 52-qubit IBM Quantum Heron r3 processor. The hybrid quantum reservoir network (QRN) reduced forecast error by 21% to 41% compared to classical latent linear baselines. This approach decouples complex boundary geometries from quantum circuit topology, enabling the simulation of fluid dynamics on non-uniform grids.

IN DEPTH

Key Findings

The research team, comprising QodeX Quantum, NVIDIA, RTX BBN, the University of South Carolina, and the Naval Surface Warfare Center Panama City Division, achieved a test NRMSE of 0.257 in 10-step closed-loop autoregressive rollouts on the cylinder flow dataset. This performance significantly outperformed a classical persistence baseline (1.126) and a noiseless CUDA-Q simulation (0.707), suggesting that physical hardware noise acted as an effective regularizer.

Technical Details

The team utilized a hybrid quantum reservoir network (QRN) on the ibm_boston processor. To handle unstructured meshes, the architecture pairs a pretrained graph neural network autoencoder with a fixed, untrained quantum reservoir. The reservoir is divided into dedicated memory and readout registers. The process involves:

- Encoding latent flow variables into the memory register via data-reuploading rotations.
- Coupling memory to readout qubits using tunable partial-SWAP gates.
- Implementing mid-circuit measurement and reset of the readout register to create a tunable fading memory.
- Applying a post-hoc Pauli-weight truncation method to manage readout shot requirements.

The system executed circuits up to 24 times longer than median qubit coherence times and sustained continuous evolution for up to 28 times median T1 and 24 times median T2 without collapsing the memory register.

Background & Context

Traditional gate-based quantum solvers for computational fluid dynamics have been limited by requirements for uniform grids, low Reynolds numbers, or shallow variational circuits. These limitations often lead to heavy measurement overheads or "barren plateaus." The QRN approach bypasses these issues by using a fixed reservoir rather than a variational circuit.

Significance & Outlook

This study validates the ability to forecast nonlinear PDEs on complex engineering meshes using physical quantum hardware. Additionally, QodeX Quantum used this research to develop and validate modules for AlphaPsi, its agentic quantum workflow platform. The implementation code has been hosted on GitHub and the technical manuscript is available on arXiv.

Source: https://quantumcomputingreport.com/qodex-quantum-nvidia-rtx-bbn-and-u-s-navy-researchers-forecast-unstructured-pdes-on-52-qubit-ibm-hardware/

Collected: November 15, 2026 | Automated Research System (AI)

#06 SEAFOOD Platform: Osaka University's Quantum-HPC Integration Timeline

Published November 13, 2026 Quantum Computing Report USA



OVERVIEW

The University of Osaka has developed SEAFOOD, a hybrid computing infrastructure linking the OCTOPUS supercomputer with domestic superconducting quantum processors. Scheduled for formal operation in December 2026, the platform enables partitioned workloads to address the computational costs of simulating strongly correlated electronic structures.

IN DEPTH

Key Findings

The University of Osaka's D3 Center and Center for Quantum Information and Quantum Biology (QIQB) have integrated domestic superconducting quantum processors with the OCTOPUS supercomputer via the SEAFOOD (Supercomputing Entangled by Ambitious Fusion of OQTOPUS and OCTOPUS at D3C) platform. This marks the first time a Japanese-developed quantum computer and an on-campus supercomputer have been integrated at a single academic institution in Japan. The system is scheduled to enter formal operation in December 2026.

Technical Details

Led by Professor Susumu Date and Associate Professor Hiroshi Ueda, the SEAFOOD architecture utilizes a three-layer stack to partition workloads, ensuring each processor executes routines suited to its specific hardware design. This approach addresses the noise and qubit count constraints of near-term quantum processors and the high computational costs of classical high-performance clusters.

In a Quantum-Selected Configuration Interaction (QSCI) demonstration, the workload was split as follows:

- The quantum processor sampled dominant electron configurations.
- The high-memory compute nodes of OCTOPUS constructed and diagonalized the resulting subspace Hamiltonian.

Background & Context

The project focuses on simulating strongly correlated electronic structures for materials engineering and drug discovery. Development was supported by several Japanese initiatives, including JST Moonshot Goal 6, MEXT Q-LEAP, JST COI-NEXT, CSTI SIP, and NEDO's Post-5G quantum-supercomputer hybrid platform program. The software stack, OQTOPUS, is open-source, and the QSCI methodology has been published in Physical Review Research.

Significance & Outlook

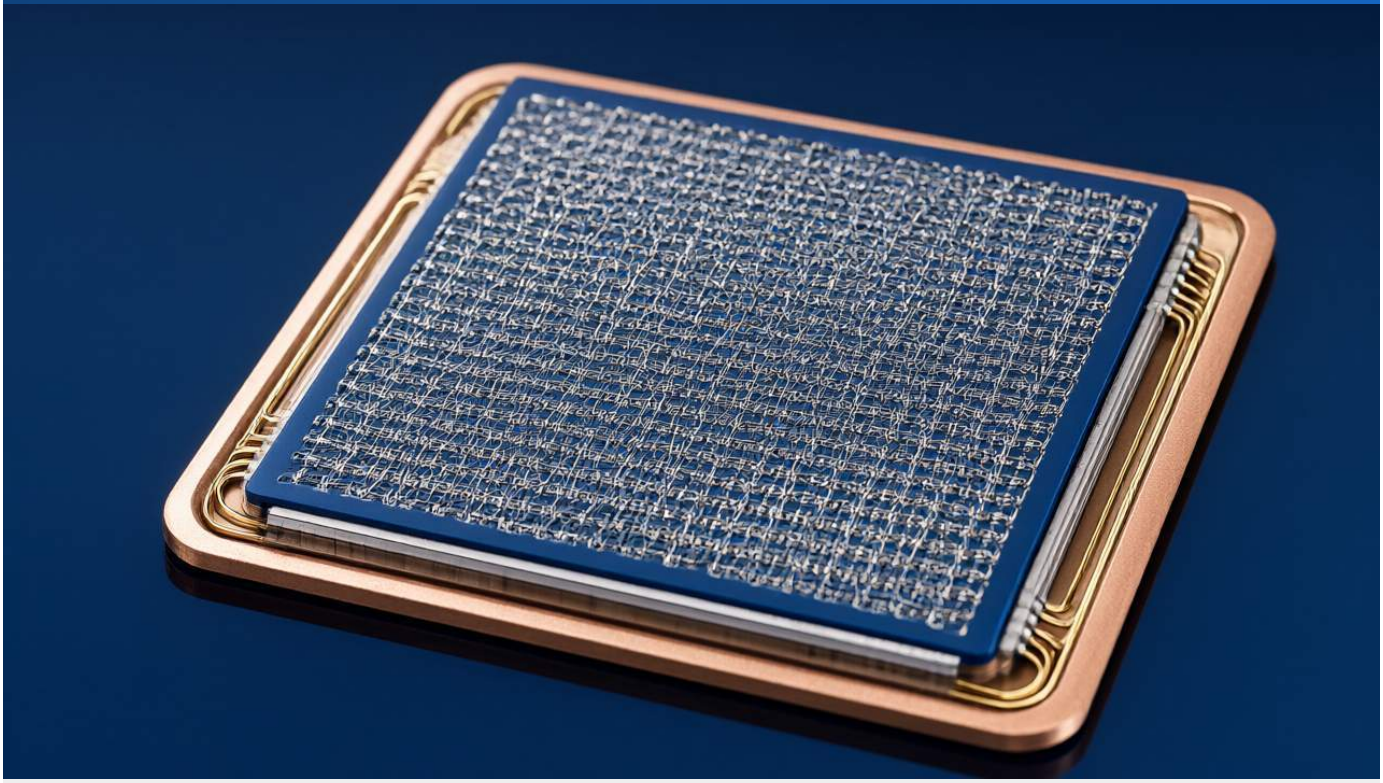
In fiscal year 2027, the D3 Center will grant trial access to up to 10 research teams from The University of Osaka through a proposal-based program. This phase is intended to refine hybrid job scheduling before the platform opens for broader external access via the nationwide JHPC-quantum initiative.

Source: https://quantumcomputingreport.com/university-of-osaka-integrates-domestic-superconducting-quantum-processors-with-octopus-supercomputer-via-seafood-platform/

Collected: November 15, 2026 | Automated Research System (AI)

#07 IBM Starling: Fault-Tolerant Quantum Computing Roadmap to 2029

Published November 12, 2026 솔루션뉴스 South Korea



OVERVIEW

IBM has reduced the number of Josephson junctions in its quantum processor design by 60%, potentially increasing the ratio of fully functional chips from 8.3% to 33.3%. These advancements support the development of "Starling," a system targeting 200 logical qubits and 100 million quantum operations by 2029.

IN DEPTH

Key Findings

IBM has reported significant technical progress toward its 2029 goal of deploying "Starling," a fault-tolerant quantum computer. The most critical achievement is a redesigned processor that reduces the number of Josephson junctions from approximately 5,000 to 2,000. Based on a 99.95% individual junction success rate, IBM projects that the proportion of fully functional chips will improve fourfold, rising from roughly 1 in 12 to 1 in 3.

Technical Details

- **Processor Architecture:** By reducing qubit connectivity from 6 to 5, IBM improved logical error rates by 20 to 100 times. The "c-coupler" length was extended to 20mm, and wiring layers were increased from 2 to 3.
- **Error Correction:** Utilizing quantum Low-Density Parity-Check (qLDPC) codes, IBM aims to reduce physical qubit requirements by a factor of 10 compared to previous methods. The "Relay-BP" decoder has reached the performance levels required for Starling.
- **Hardware Integration:** IBM is developing a dedicated ASIC using a 2nm nanosheet process to execute the decoder algorithm in real-time. For inter-processor connectivity, "l-couplers" have demonstrated an average error rate of 1% (best result 0.7%), though the target for Starling is 0.1%.
- **Cryogenic Infrastructure:** A new rectangular modular cryostat has been introduced to allow side-by-side expansion. This system supports flexible wiring, with one section accommodating up to 1.54km of cabling.

Background & Context

Quantum computers utilize qubits to process information via superposition, but they are highly sensitive to environmental noise. To combat this, IBM uses "logical qubits," which distribute information across multiple physical qubits to detect and correct errors. The Starling project follows a phased roadmap: the "Kookaburra" modular processor in 2026 and the "Cockatoo" multi-module system in 2027.

Significance & Outlook

The transition to a modular architecture—combining quantum memory and Logical Processing Units (LPU)—is designed to lower manufacturing risks and increase yield. The ultimate success of Starling depends on integrating these individual breakthroughs in cryogenic control, real-time error correction, and large-scale cooling into a single stable system by 2029.

Source: https://www.solnews.co.kr/news/articleView.html?idxno=85172

Collected: November 15, 2026 | Automated Research System (AI)

#08 Anthropic AI Model: Proof of Aaronson-Ambainis Conjecture

Published November 09, 2026 Business Insider Japan Japan



OVERVIEW

Physical Superintelligence Inc. (PSI) has used an unreleased Anthropic AI model to prove the Aaronson-Ambainis conjecture, a long-standing problem in quantum computing. The proof defines the boundary conditions for quantum supremacy, establishing that exponential speedups require specific regular patterns in input data.

IN DEPTH

Key Findings

Physical Superintelligence Inc. (PSI) announced on November 6 that it has completely proven the Aaronson-Ambainis conjecture. The proof was achieved using a proprietary research system integrated with an unreleased base model from Anthropic. The system completed the proof within three days, resolving a fundamental question regarding the limits of quantum computing that had remained unsolved since it was first proposed in 2008.

Technical Details

The proof clarifies the conditions under which quantum computers exhibit "quantum supremacy"—the ability to perform calculations significantly faster than classical computers. According to PSI, the proof demonstrates that the "exponential computational advantage" enjoyed by quantum computing is contingent upon the existence of specific regular patterns or structures within the input data. This confirms that quantum computers are not universal in their superiority and possess clear computational limits, although they maintain dramatic processing power for specific tasks.

Background and Context

The development was led by David Turturean, a technical staff member at PSI. While the specific name and specifications of the Anthropic model remain undisclosed, the result has been praised by researchers from Stanford University, including Guy Blanc, Jordan Docter, Carmen Strassle, and Li-Yang Tan, who described the proof's mathematical elegance. This follows a trend of AI-driven mathematical breakthroughs, such as OpenAI's September announcement regarding the smoothness of the Navier-Stokes equations.

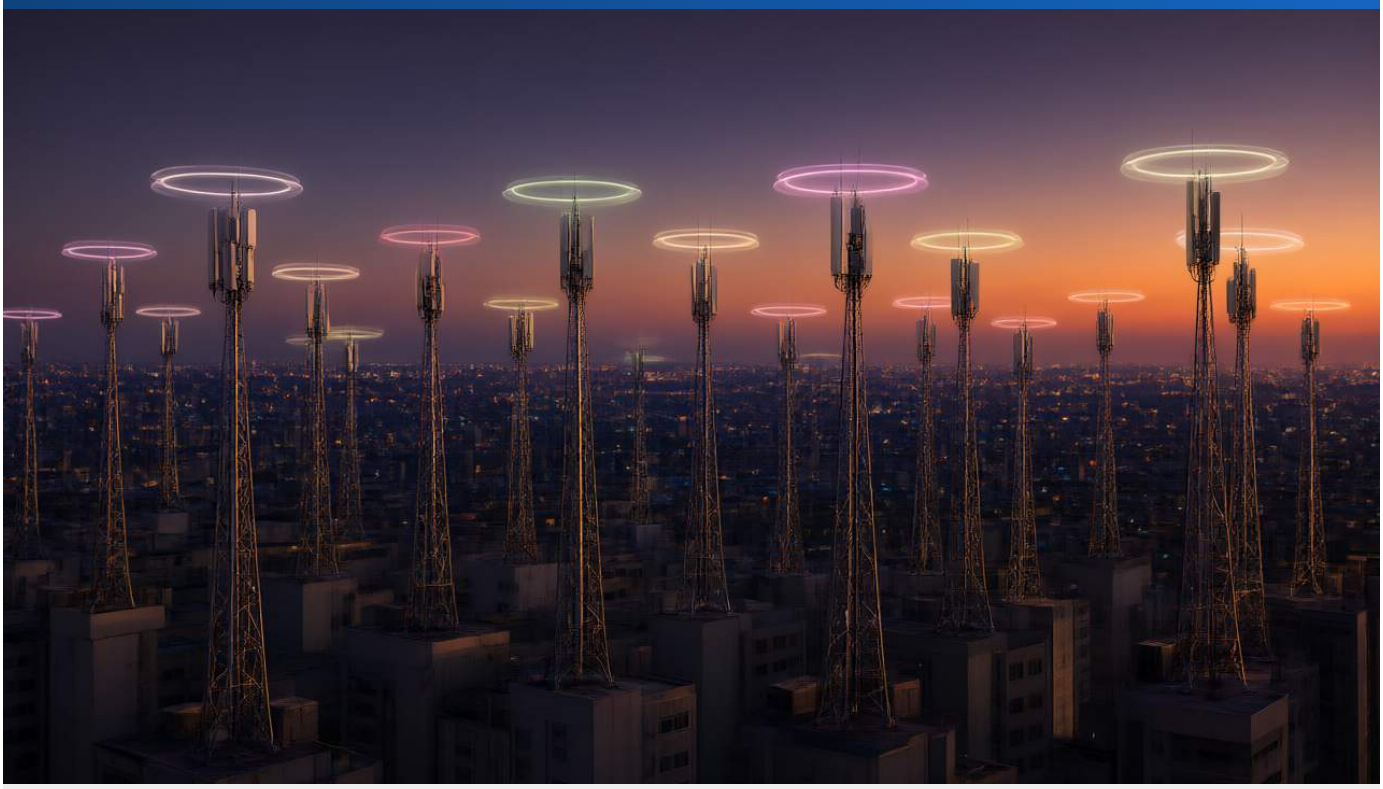
Significance and Outlook

The rapid advancement of AI in mathematics is prompting a shift in academic perspectives. Terence Tao, a mathematics professor at UCLA, noted that the "era of romanticism" in mathematics has ended. This shift is reflected in a MIT survey of graduate students and postdoctoral researchers, where 76% stated they are unlikely to remain in academia and 76% reported a decrease in motivation to conduct research due to AI's progress. Despite these academic shifts, the proof is regarded as a major milestone in the early stages of quantum computing practical application.

Source: https://www.businessinsider.jp/article/2611-anthropic-quantum-computing-breakthrough-aaronson-ambainis-math-physics-psi/

#09 KDDI: Quantum-AI Optimization for Uplink Speed Improvements

Published November 12, 2026 AI Watch Japan



OVERVIEW

KDDI and KDDI Research Institute have reduced the frequency of uplink speeds falling below 1.5Mbps by approximately 80% during congestion. By applying quantum technology to optimize radio interference across 683 base stations, the company has successfully managed a combinatorial space of approximately 10^1200 possibilities.

IN DEPTH

Key Findings

KDDI has successfully implemented a quantum-based optimization system to mitigate radio interference between base stations in select areas of Tokyo. Starting October 2, this deployment targeted 683 base stations, resulting in an approximately 80% reduction in the occurrence of uplink speeds dropping below 1.5Mbps—the benchmark for comfortable image uploads—during peak congestion periods.

Technical Details

The optimization process involved calculating the ideal configuration from a massive set of approximately 10^1200 possible combinations. While KDDI previously utilized decentralized AI for communication area optimization, the company integrated quantum technology to handle large-scale combinatorial optimization, which is a specific strength of quantum computing compared to decentralized AI.

To enable commercial application, KDDI developed a proprietary method to model radio interference between base stations into a format that can be efficiently processed by a quantum platform.

Background and Context

This initiative is supported by the "AI and Quantum Common Platform," a system developed by KDDI that allows users to combine AI and quantum computing resources without requiring specialized quantum knowledge. Furthermore, the project is part of the "Post-5G Information and Communication System Infrastructure Strengthening R&D Project (g5-3) Middleware Development for the Industrialization of Quantum Computers," managed by the New Energy and Industrial Technology Development Organization (NEDO) under the Ministry of Economy, Trade and Industry.

Significance and Outlook

Following the initial deployment in Tokyo, KDDI intends to expand this optimization technology to urban areas in Tokyo, Nagoya, and Osaka to further enhance communication quality. Beyond telecommunications networks, KDDI envisions expanding the use of the AI and Quantum Common Platform into other industrial sectors, including power grids, logistics, and materials development.

Source: https://ai.watch.impress.co.jp/docs/news/2155817.html

Collected: November 15, 2026 | Automated Research System (AI)

#10 Qunova HI-VQE: 80-Qubit Simulation Matches Classical Benchmark

Published November 12, 2026 The Quantum Insider USA



OVERVIEW

Qunova Computing's HI-VQE algorithm achieved chemical accuracy in an 80-qubit simulation of a 4Fe-4S iron-sulfur cluster using IBM quantum hardware and the Fugaku supercomputer. The method matched the accuracy of the Density Matrix Renormalization Group (DMRG) classical benchmark while reducing computational costs by approximately 300 times compared to previous quantum diagonalization approaches. This result demonstrates that quantum-based workflows can now compete with the strongest classical methods for highly demanding industrial chemistry problems.

IN DEPTH

Key Findings

Qunova Computing has reported that its HI-VQE algorithm is the first quantum algorithm to reach chemical accuracy at a scale of 80 qubits. In a simulation of the 4Fe-4S iron-sulfur cubane cluster, the algorithm's extrapolated ground-state energy estimate of −327.2482 ± 0.0007 Hartree was consistent with the leading classical reference calculated via the Density Matrix Renormalization Group (DMRG) method (−327.2471 Hartree). This result surpasses the latest extrapolated Trimmed Configuration Interaction (TrimCI) results.

Technical Details

The computation was executed through a hybrid quantum/classical supercompute workflow utilizing an IBM Quantum System Two quantum computer and RIKEN’s Fugaku supercomputer. A key performance metric highlighted is the reduction in computational overhead: HI-VQE reached a ground-state energy of −326.9821 Ha by the second iteration using 2,518 Fugaku node-hours. This represents a ~300x reduction in cost compared to previously published sample-based quantum diagonalization (SQD) results for the same molecule.

Background & Context

The 4Fe-4S iron-sulfur cluster is a computationally demanding system that serves as a proving ground for chemistry algorithms. These clusters act as catalytic and electron-transfer cores for metalloproteins, including nitrogenase, hydrogenases, and ferredoxins. Accurate electronic structure computation of these molecules is essential for understanding biological nitrogen fixation, electron transport, and hydrogen production.

Significance & Outlook

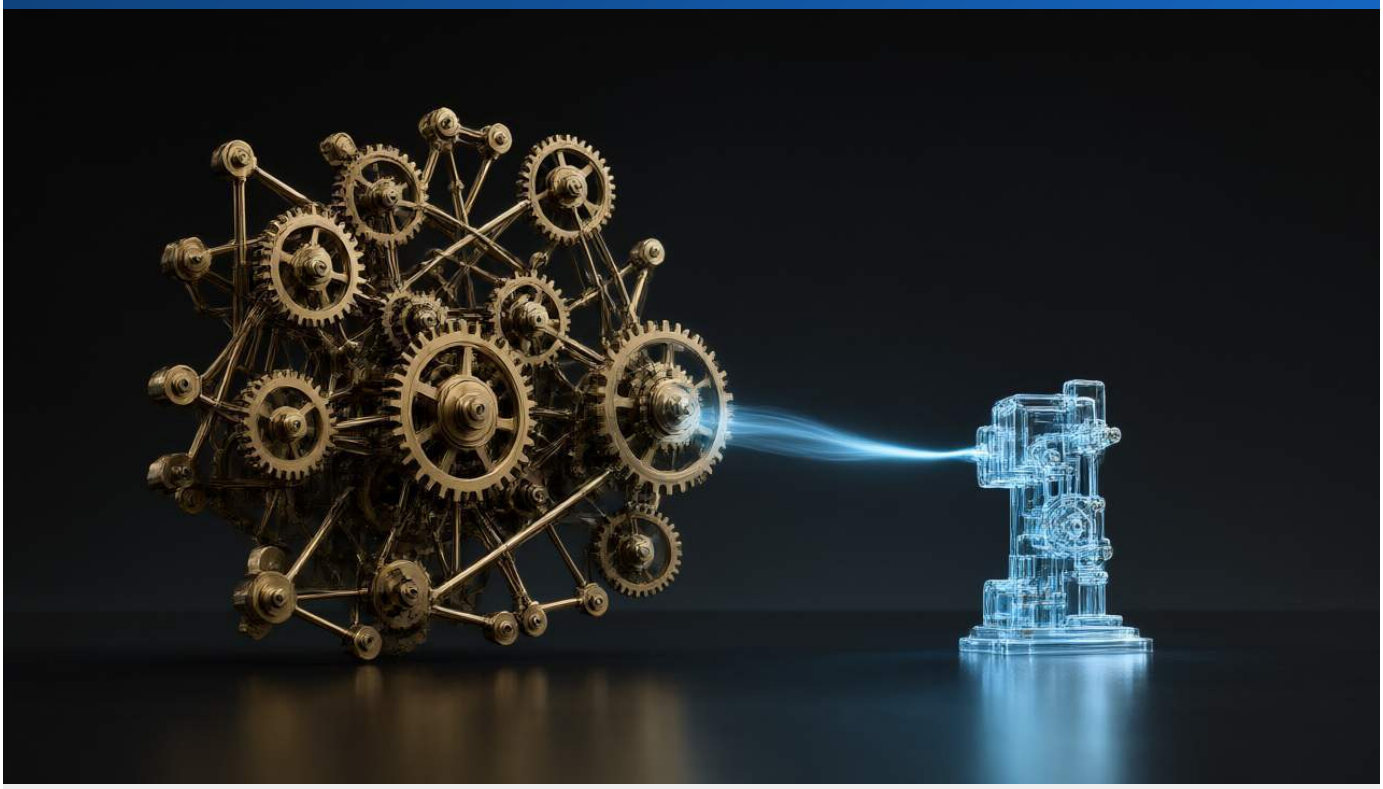
By achieving reproducible chemical accuracy on an industrially relevant molecule, Qunova aims to move quantum-enhanced chemistry from institutional demonstrations to practical industrial application. The HI-VQE solver does not assume linearity or locality, making it applicable to drug targets, electrode materials, and catalysts. Qunova plans to add orbital optimization and orbital rotation functionalities to the algorithm to reach the next accuracy milestone, which involves processing between 500 million and 3 billion unique states. This research was supported by the New Energy and Industrial Technology Development Organization (NEDO) under Japan's Ministry of Economy, Trade and Industry (METI).

Source: https://thequantuminsider.com/2026/11/12/qunovas-hybrid-computing-method-matches-classical-benchmark-for-iron-sulfur-cluster/

Collected: November 15, 2026 | Automated Research System (AI)

#11 QuEra, Harvard, and Anthropic: AI-Driven Quantum Circuit Optimization

Published November 10, 2026 The Quantum Insider USA



OVERVIEW

Researchers from QuEra, Harvard, and Anthropic used the Claude AI model to optimize quantum circuits, achieving up to a 39-fold reduction in resource requirements for specific computational building blocks. By reducing gate counts and qubit requirements, the study demonstrates a path toward executing complex algorithms on hardware with limited capacity.

IN DEPTH

Key Findings

The collaboration between QuEra Computing, Harvard University, and Anthropic has demonstrated that AI can significantly reduce the computational resources required for fault-tolerant quantum circuits. In a study of eight computational building blocks, AI-designed circuits reduced both gate counts and qubit requirements in seven cases. The most significant improvement was seen in a fixed-point Gaussian function, which achieved a 39-fold reduction in a combined resource metric (Toffoli-gate count multiplied by peak logical qubits).

Technical Details

The team paired Anthropic's Claude with an independent automated evaluator that scored proposed designs based on correctness, Toffoli-gate count, and qubit count. Claude used this feedback to iteratively improve the designs. Specific results include:

- **Optimization:** The Reed-Solomon decoder saw a 12-fold reduction in Toffoli gates, dropping from 5.1 million to approximately 440,000.
- **Chemistry:** The Coulomb potential calculation achieved a 9.9-fold reduction in combined resources, with Toffoli gates decreasing from 6,355 to 1,140 and qubits from 252 to 142.
- **Numerical Computation:** The fixed-point Gaussian function saw Toffoli gates drop from about 24,000 to 1,583 and qubits from 362 to 139.

Additionally, the AI discovered a new reversible approach to modular inversion. In a point-addition benchmark, one design required 17x fewer gates than the previous minimum-qubit design, while using 12% more qubits.

Background and Significance

The research, detailed in the preprint "Discovering state-of-the-art quantum circuits with large language models" (arXiv:2610.22367), was led by Andi Gu as part of Anthropic's STEM Fellows Program. This work focuses on the "resource curve"—reducing the efficiency requirements of algorithms to meet the rising capabilities of quantum hardware.

This project follows a previous pilot where Claude developed control software for laser stability in QuEra's systems, reducing the time to recover laser lock from 5–10 minutes to seconds. Together, these efforts highlight AI's role in improving algorithms and supporting reliable hardware operation.

Source: https://thequantuminsider.com/2026/11/10/quera-harvard-anthropic-ai-quantum-circuit-resource-requirements/

Collected: November 15, 2026 | Automated Research System (AI)

#12 IBM, RIKEN, and Cleveland Clinic: 12,000-Atom HQC Simulation

Published November 12, 2026 Quantum Computing Report USA



OVERVIEW

A collaboration between Cleveland Clinic, RIKEN, and IBM has simulated full solvated protein–ligand complexes exceeding 12,000 atoms using heterogeneous quantum-classical (HQC) electronic-structure calculations. The team achieved a 40× increase in molecular size and up to a 210× improvement in fragment energy accuracy compared to previous benchmarks. This advancement addresses the memory and accuracy limitations of classical methods in predicting drug-binding affinities.

IN DEPTH

Key Findings

Researchers from the Cleveland Clinic Foundation, the RIKEN Center for Computational Science, and IBM Quantum have executed the largest heterogeneous quantum-classical (HQC) electronic-structure calculation to date. The simulation involved full solvated protein–ligand complexes exceeding 12,000 atoms. Compared to previous quantum-embedding benchmarks, this effort resulted in a greater than 40× increase in molecular size and an improvement in fragment energy accuracy of up to 210×.

Technical Details

The hardware configuration paired two 156-qubit IBM Heron r2 processors with three leadership supercomputers. The team collected 3.0 billion measurement shots across 21,006 circuits. To manage the computational load, the researchers implemented the following:

- **Embedded Wavefunction (EWF) Formalism:** A localized approach that partitions macromolecules into spatially bounded atomic neighborhoods, reducing fragment construction costs from cubic $O(M^3)$ scaling to constant $O(1)$ scaling per fragment.
- **TrimSQD:** An extension of sample-based quantum diagonalization used for heavy-atom fragments spanning 13 to 45 molecular orbitals. It iteratively filters and merges electronic configurations across distributed GPU subgroups before final subspace diagonalization.
- **Workflow Automation:** A Prefect-based controller orchestrated asynchronous job queues between the ROQUO supercomputer and IBM Quantum Runtime, processing 608 million measurement samples in 62.4 hours of wall-clock time.

Background & Context

Predicting drug binding requires solving the many-electron Schrödinger equation across thousands of atoms. Standard density functional theory approximations can misrank binding affinities, while full-system wavefunction methods often exceed classical memory limits. This project was supported by the U.S. National Science Foundation, the National Institutes of Health, Japan’s New Energy and Industrial Technology Development Organization (NEDO), and the MEXT Supercomputer Fugaku program.

Significance & Outlook

By bridging the gap between classical memory limits and quantum accuracy, this HQC approach allows for the simulation of significantly larger molecular systems. The integration of distributed GPU subgroups and quantum processors demonstrates a scalable pipeline for complex drug-binding simulations.

Source: https://quantumcomputingreport.com/cleveland-clinic-riken-and-ibm-cross-12000-atom-barrier-in-heterogeneous-quantum-classical-drug-binding-simulations/

Collected: November 15, 2026 | Automated Research System (AI)