

TROY TECHNICAL

QUANTUM COMPUTING WEEKLY REPORT

October 17 - October 23, 2026 | Issue 2026-W42

**Verified computational advantage, error control
and usable access**


17

ANALYZED ARTICLES

10

PRIORITY THEMES

Issued 2026-10-23 | English edition | fixed approved cover artwork



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

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

SIGNAL 01 | QUBIT DEVICE

Google Quantum AI: Magic State Cultivation and Willow Processor

Google Quantum AI has demonstrated "magic state cultivation" on its Willow processor, reducing errors by approximately 40x compared to un-cultivated states.

SIGNAL 02 | CONTROL STACK

IBM and Rutgers: Real-Time Control Across 100 Qubits Explained

Researchers from Rutgers, IBM, and collaborating institutions demonstrated real-time quantum feedback control on up to 100 qubits using an IBM Heron processor.

SIGNAL 03 | ERROR CORRECTION

TRIO Pulses: Robust Single-Qubit Gates for Superconducting Processors

Researchers from Chalmers University and partners have validated Time-Rescaling Robust Inverse-Optimal (TRIO) pulses to mitigate control-amplitude and frequency detuning errors in superconducting qubits.

SIGNAL 04 | SYSTEM ACCESS

NVIDIA Quantum Simulation: 104-Qubit Fault-Tolerant Testing on JUPITER

NVIDIA simulated a 104-qubit fault-tolerant quantum processor on the JUPITER supercomputer to validate a magic-state cultivation approach that reduces required qubits by 90%.

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.

Five numbers or evidence anchors that require conditions


40x A02	SOURCE-BOUND CONDITION Google Quantum AI has demonstrated "magic state cultivation" on its Willow processor, reducing errors by approximately 40x compared to un-cultivated states.
100 qubit s A04	SOURCE-BOUND CONDITION Researchers from Rutgers, IBM, and collaborating institutions demonstrated real-time quantum feedback control on up to 100 qubits using an IBM Heron processor.
99.90% A09	SOURCE-BOUND CONDITION The technique achieved a 99.90% IRB fidelity for a 90-nanosecond X-gate, expanding the high-fidelity parameter range by an order of magnitude compared to standard DRAG Gaussian pulses.
90% A03	SOURCE-BOUND CONDITION NVIDIA simulated a 104-qubit fault-tolerant quantum processor on the JUPITER supercomputer to validate a magic-state cultivation approach that reduces required qubits by 90%.
30x A13	SOURCE-BOUND CONDITION Researchers from Google Quantum AI, Google DeepMind, and MIT have developed a modular quantum error correction (QEC) architecture that reduces physical qubit overhead by over 30x compared to standard 2D surface codes.

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.



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	Google Quantum AI: Magic State Cultivation and Willow Processor Business / rules cue: fid
P02	IBM and Rutgers: Real-Time Control Across 100 Qubits Explained Tech / process cue: no system/business cue
P03	TRIO Pulses: Robust Single-Qubit Gates for Superconducting Processors Business / rules cue: fid
P04	NVIDIA Quantum Simulation: 104-Qubit Fault-Tolerant Testing on JUPITER Tech / process cue: no system/business cue
P05	Correlated Decoding: Leiden and Riverlane Fix Distance-Losing Flaw Business / rules cue: fid
P06	Kipu Quantum: Neutral-Atom Feature Extraction for Toxicity Tech / process cue: no system/business cue
P07	Modular Hyperbolic Codes: Google and MIT's 30x QEC Overhead Reduction Business / rules cue: standard
P08	NVIDIA and Riverlane: GPU-Accelerated QEC Decoder Calibration Business / rules cue: acquired
P09	University of Oxford: Multi-Species Atom Network and 10s Memory Tech / process cue: no system/business cue
P10	Quantinuum Helios: Progress Toward Megaquop Error Correction Tech / process cue: no system/business cue

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

Google Quantum AI: Magic State Cultivation and Willow Processor

A02 | Qubit device | Research

Date not stated

FACT

Google Quantum AI has demonstrated "magic state cultivation" on its Willow processor, reducing errors by approximately 40x compared to un-cultivated states. This breakthrough addresses a critical resource bottleneck in universal quantum computing by enabling the efficient creation of high-fidelity magic states.

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 20/25 | MATURITY 1/5 | DECISION LOCUS: BUSINESS / RULES

P02

IBM and Rutgers: Real-Time Control Across 100 Qubits Explained

A04 | Control stack | Pilot

Date not stated

FACT

Researchers from Rutgers, IBM, and collaborating institutions demonstrated real-time quantum feedback control on up to 100 qubits using an IBM Heron processor. By performing nearly 5,000 measurements and resets to counter chaotic behavior, the team validated a critical capability required for future fault-tolerant quantum computing.

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

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

P03

TRIO Pulses: Robust Single-Qubit Gates for Superconducting Processors

A09 | Error correction | Product adoption

Date not stated

FACT

Researchers from Chalmers University and partners have validated Time-Rescaling Robust Inverse-Optimal (TRIO) pulses to mitigate control-amplitude and frequency detuning errors in superconducting qubits. The technique achieved a 99.90% IRB fidelity for a 90-nanosecond X-gate, expanding the high-fidelity parameter range by an order of magnitude compared to standard DRAG Gaussian pulses.

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 20/25 | MATURITY 4/5 | DECISION LOCUS: BUSINESS / RULES

P04

NVIDIA Quantum Simulation: 104-Qubit Fault-Tolerant Testing on JUPITER

A03 | System access | Research

Date not stated

FACT

NVIDIA simulated a 104-qubit fault-tolerant quantum processor on the JUPITER supercomputer to validate a magic-state cultivation approach that reduces required qubits by 90%. Utilizing 8,192 GPUs, the simulation models realistic hardware noise to accelerate the design of useful quantum applications in drug discovery and materials science.

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 20/25 | MATURITY 1/5 | DECISION LOCUS: TECH / PROCESS

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

P05

Correlated Decoding: Leiden and Riverlane Fix Distance-Losing Flaw

A07 | Applications | Pilot

Date not stated

FACT

Leiden University and Riverlane have identified a structural flaw in standard re-weighting protocols for two-pass correlated decoders that causes a unit loss of effective code distance below threshold ($p < 10^{-3}$). The researchers introduced an error-rate-dependent regularization parameter ($\alpha^*(p)$) that restores nominal distance scaling without adding classical runtime overhead. This ensures that error suppression in rotated surface codes scales predictably as qubit gate fidelities improve.

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: BUSINESS / RULES

P06

Kipu Quantum: Neutral-Atom Feature Extraction for Toxicity

A17 | Qubit device | Research

Date not stated

FACT

Kipu Quantum has demonstrated analog quantum feature extraction using QuEra's 256-qubit Aquila processor to improve molecular toxicity classification. The approach increased Balanced Accuracy from 0.55 to 0.70 and doubled the F1-score from ~0.30 to ~0.60 on imbalanced datasets. This result indicates a performance gain over traditional classical molecular descriptors for drug discovery and chemical safety.

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 1/5 | DECISION LOCUS: TECH / PROCESS

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

P07

Modular Hyperbolic Codes: Google and MIT's 30x QEC Overhead Reduction

A13 | Control stack | Pilot

Date not stated

FACT

Researchers from Google Quantum AI, Google DeepMind, and MIT have developed a modular quantum error correction (QEC) architecture that reduces physical qubit overhead by over 30x compared to standard 2D surface codes. By networking small Euclidean planar chips in a hyperbolic tiling geometry, the system achieves high encoding rates without requiring complex on-chip long-range wire crossings.

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 2/5 | DECISION LOCUS: BUSINESS / RULES

P08

NVIDIA and Riverlane: GPU-Accelerated QEC Decoder Calibration

A11 | Error correction | Research

Date not stated

FACT

NVIDIA has acquired a strategic minority stake in Riverlane to accelerate real-time quantum error correction (QEC) and decoder calibration. By utilizing H100 GPUs, the partners reduced surface code training cycles from 3 minutes on a single CPU core to 2.8 seconds, enabling the mitigation of dynamic qubit drift.

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 1/5 | DECISION LOCUS: BUSINESS / RULES

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

P09

University of Oxford: Multi-Species Atom Network and 10s Memory

A05 | System access | Research

Date not stated

FACT

Physicists at the University of Oxford have demonstrated a heterogeneous quantum network using strontium and calcium ions to generate multipartite entangled states. The team extended entanglement storage in calcium memory qubits to over 10 seconds, exceeding the creation time by more than 100x.

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 19/25 | MATURITY 1/5 | DECISION LOCUS: TECH / PROCESS

P10

Quantinuum Helios: Progress Toward Megaquop Error Correction

A16 | Applications | Research

Date not stated

FACT

Quantinuum has demonstrated a quantum computer capable of performing one million operations with only approximately five errors. This achievement utilizes a high lambda value in logical qubit scaling, bringing the industry closer to the "megaquop" regime of reliable quantum calculation.

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 1/5 | DECISION LOCUS: TECH / PROCESS

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 01-09

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	Quantinuum Helios: 99.9995% Logical Fidelity and 2029 Roadmap	https://quantumzeitgeist.com/quantinuum-cuts-quantum-errors-10x/
A02	Google Quantum AI: Magic State Cultivation and Willow Processor	https://blog.google/innovation-and-ai/models-and-research/quantum-computing/magic-state-cultivation-quantum-computing/
A03	NVIDIA Quantum Simulation: 104-Qubit Fault-Tolerant Testing on JUPITER	https://thequantuminsider.com/2026/10/19/nvidia-runs-simulation-of-fault-tolerant-quantum-computing-at-julich-supercomputing-centre/
A04	IBM and Rutgers: Real-Time Control Across 100 Qubits Explained	https://thequantuminsider.com/2026/10/21/rutgers-and-ibm-demonstrate-real-time-control-across-100-qubits/
A05	University of Oxford: Multi-Species Atom Network and 10s Memory	https://quantumcomputingreport.com/university-of-oxford-entangles-multi-species-atom-network-and-extends-memory-storage-beyond-10-seconds/
A06	D-Wave Advantage2: 3D Topological Insulator Simulation via 728 Qubits	https://quantumcomputingreport.com/d-wave-demonstrates-3d-analog-digital-quantum-simulation-of-higher-order-topological-insulators-using-728-qubits/
A07	Correlated Decoding: Leiden and Riverlane Fix Distance-Losing Flaw	https://quantumcomputingreport.com/leiden-university-and-riverlane-research-reveals-distance-losing-flaw-in-correlated-decoding-below-threshold/
A08	Iceberg Quantum Pinnacle 1.2: RSA-2048 Factoring Resource Specs	https://quantumcomputingreport.com/iceberg-quantum-validates-pinnacle-1-2-architecture-with-crevasse-qldpc-decoder-to-reduce-rsa-2048-factoring-resources/
A09	TRIO Pulses: Robust Single-Qubit Gates for Superconducting Processors	https://quantumcomputingreport.com/chalmers-university-demonstrates-trio-pulses-for-control-error-robust-single-qubit-gates-in-superconducting-processors/

ANALYZED ARTICLES | ITEMS 10-17

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
A10	IQM and NVIDIA: Automating 20-Qubit Chip Calibration via VLMS	https://quantumcomputingreport.com/iqm-and-nvidia-automate-20-qubit-chip-calibration-via-nvidia-ising-vision-language-models-and-qcaleval-benchmark/
A11	NVIDIA and Riverlane: GPU-Accelerated QEC Decoder Calibration	https://quantumcomputingreport.com/nvidia-takes-strategic-stake-in-riverlane-to-accelerate-real-time-quantum-error-correction-and-gpu-decoder-calibration/
A12	Quantinuum Helios: 99.9995% Logical Memory Fidelity and QEC Specs	https://quantumcomputingreport.com/quantinuum-demonstrates-99-9995-logical-memory-fidelity-on-helios-trapped-ion-processor/
A13	Modular Hyperbolic Codes: Google and MIT's 30x QEC Overhead Reduction	https://quantumcomputingreport.com/google-quantum-ai-deepmind-demonstrate-modular-hyperbolic-codes-for-over-30x-qec-overhead-reduction/
A14	IBM Research: AI Agent Factories for QEC Decoder Chip Design	https://quantumcomputingreport.com/ibm-research-automates-qec-decoder-chip-design-using-open-ai-agent-factories-at-11-5x-lower-inference-cost/
A15	Silicon Spin Qubits: RIKEN achieves >99.9% parallel gate fidelity	https://quantumcomputingreport.com/riken-demonstrates-high-fidelity-simultaneous-control-of-silicon-spin-qubits-above-99-9-threshold/
A16	Quantinuum Helios: Progress Toward Megaquop Error Correction	https://sciencenews.org/article/quantum-computer-megaquops-error-correction
A17	Kipu Quantum: Neutral-Atom Feature Extraction for Toxicity	https://quantumcomputingreport.com/kipu-quantum-demonstrates-neutral-atom-quantum-feature-extraction-for-molecular-toxicity-classification-on-quera-hardware/

QuantumComputing — Selected Articles

Date: 2026-10-24


Articles: 17

Table of Contents

- #01 Quantinuum Helios: 99.9995% Logical Fidelity and 2029 Roadmap
- #02 Google Quantum AI: Magic State Cultivation and Willow Processor
- #03 NVIDIA Quantum Simulation: 104-Qubit Fault-Tolerant Testing on JUPITER
- #04 IBM and Rutgers: Real-Time Control Across 100 Qubits Explained
- #05 University of Oxford: Multi-Species Atom Network and 10s Memory
- #06 D-Wave Advantage2: 3D Topological Insulator Simulation via 728 Qubits
- #07 Correlated Decoding: Leiden and Riverlane Fix Distance-Losing Flaw
- #08 Iceberg Quantum Pinnacle 1.2: RSA-2048 Factoring Resource Specs
- #09 TRIO Pulses: Robust Single-Qubit Gates for Superconducting Processors
- #10 IQM and NVIDIA: Automating 20-Qubit Chip Calibration via VLMs
- #11 NVIDIA and Riverlane: GPU-Accelerated QEC Decoder Calibration
- #12 Quantinuum Helios: 99.9995% Logical Memory Fidelity and QEC Specs
- #13 Modular Hyperbolic Codes: Google and MIT's 30x QEC Overhead Reduction
- #14 IBM Research: AI Agent Factories for QEC Decoder Chip Design
- #15 Silicon Spin Qubits: RIKEN achieves >99.9% parallel gate fidelity
- #16 Quantinuum Helios: Progress Toward Megaquop Error Correction
- #17 Kipu Quantum: Neutral-Atom Feature Extraction for Toxicity

#01 Quantinuum Helios: 99.9995% Logical Fidelity and 2029 Roadmap

Published October 19, 2026 Quantum Zeitgeist USA



OVERVIEW

Quantinuum has achieved a world-record 99.9995% logical fidelity on its Helios computer, reducing errors to five per million operations. This result demonstrates the scalability of its trapped-ion error-correction architecture, supporting the company's goal of a universal fault-tolerant quantum computer by 2029.

Key Findings

Quantinuum has reported a new world record in quantum error correction, achieving 99.9995% logical fidelity on its Helios system. This performance represents a tenfold reduction in error rates compared to the company's previous record from September 2026. According to Quantinuum, this result is 100 times lower than the best published result from Google Quantum AI's superconducting system.

Technical Details

The achievement utilizes a "color code" implemented within an end-to-end framework that integrates hardware-level error correction, all-to-all qubit connectivity, and tailored AI decoders. The company observed an exponential decrease in logical errors as the code distance increased, with each increase reducing remaining logical errors by a factor of 12.

The Helios hardware has demonstrated flexibility by supporting different error-correction codes in rapid succession. While the recent color code used up to 37 physical qubits per logical qubit, a previous code called Helix used 10 physical qubits per logical qubit. To optimize performance, the architecture employs:

- Adaptive syndrome extraction to dynamically adjust error-checking measurements.
- Chain-mapping to connect codes suited for different computational tasks.
- Real-time decoding and efficient code structures.

Background and Context

The ability to deploy different codes on the same commercial hardware allows Quantinuum to prototype rapidly and tailor its approach to specific applications. This flexibility is a core component of the company's trapped-ion technology and error-correction architecture.

Significance and Outlook


Quantinuum asserts that the observed scaling behavior suggests fault-tolerance can be achieved with fewer qubits than currently projected for other quantum computing modalities. The company is now focused on integrating its hardware, error-correction architecture, and full-stack software platform into a cohesive system. This trajectory is intended to resolve central technical risks in the field and reach the goal of a universal fault-tolerant quantum computer by 2029.

Source: <https://quantumzeitgeist.com/quantinuum-cuts-quantum-errors-10x/>

Collected: October 24, 2026 | Automated Research System (AI)

#02 Google Quantum AI: Magic State Cultivation and Willow Processor

Published October 15, 2026 Google (The Keyword / blog.google) USA



OVERVIEW

Google Quantum AI has demonstrated "magic state cultivation" on its Willow processor, reducing errors by approximately 40x compared to un-cultivated states. This breakthrough addresses a critical resource bottleneck in universal quantum computing by enabling the efficient creation of high-fidelity magic states.

Key Findings

Google Quantum AI has successfully implemented magic state cultivation on the Willow processor, achieving a 40-fold reduction in errors relative to un-cultivated states. The team reported a record logical fidelity of 99.99% for T-states. These results, published in Science, provide the first practical demonstration of a solution to a severe resource bottleneck in quantum computing.

Technical Details

The cultivation process acts as a filter to produce high-purity quantum states. The protocol involves injecting a noisy state into a small region of the quantum chip and executing a series of coordinated measurements to identify and remove flawed attempts. This allows the system to "cultivate" a low-error state that can be adapted into a larger encoding.

During development, researchers identified that electrical equipment drift and temperature variation caused noise that interfered with benchmarking. To resolve this, the team implemented additional system controls and developed a new benchmarking technique that uses cultivation to measure itself. Additionally, the team demonstrated the ability to switch between different error correction methods, grafting states from a distance-3 color code to a distance-5 surface code.

Background & Context

Universal quantum computing requires specific quantum operations that rely on "magic states." Historically, creating high-quality, error-corrected magic states has been resource-intensive; some estimates suggested up to 95% of a quantum computer's capacity could be consumed by their production. Furthermore, magic states are a scarce resource because they cannot be reused once consumed.

Significance & Outlook

By reducing the overhead required for magic state production, this approach could substantially decrease the physical size of the quantum computer needed for practical tasks. While the team notes that cultivation must still be demonstrated at a larger scale and further technical challenges remain, the experiment turns a theoretical concept from 2024 into a demonstrated reality. This moves Google Quantum AI closer to building a useful, error-corrected quantum computer capable of solving otherwise unsolvable problems.

Source: <https://blog.google/innovation-and-ai/models-and-research/quantum-computing/magic-state-cultivation-quantum-computing/>

Collected: October 24, 2026 | Automated Research System (AI)

#03 NVIDIA Quantum Simulation: 104-Qubit Fault-Tolerant Testing on JUPITER

Published October 19, 2026 The Quantum Insider USA



OVERVIEW

NVIDIA simulated a 104-qubit fault-tolerant quantum processor on the JUPITER supercomputer to validate a magic-state cultivation approach that reduces required qubits by 90%. Utilizing 8,192 GPUs, the simulation models realistic hardware noise to accelerate the design of useful quantum applications in drug discovery and materials science.

Key Findings

NVIDIA has completed a record-breaking simulation of a 104-qubit fault-tolerant quantum processor using the JUPITER supercomputer at the Jülich Supercomputing Centre. The primary achievement is the validation of a magic-state cultivation approach that requires 90% fewer qubits than previous methods to produce the magic states necessary for fault-tolerant computation.

Technical Details

The simulation was executed on a system of NVIDIA GH200 superchips, utilizing a total of 8,192 GPUs. To achieve optimized performance and scale, the team employed NVIDIA cuQuantum software and the NVIDIA CUDA-Q platform's fault-tolerant logical layer to express the magic-state circuit. Unlike previous research from Jülich that simulated 50 qubits under ideal, noiseless conditions, this project modeled the process under realistic hardware noise to assess the impact of errors.

The development workflow involved a tiered approach to scaling:

- Initial development and testing were conducted at the NVIDIA Accelerated Quantum Research Center (NVAQC) using an NVIDIA GB200 NVL72 system.
- Tests were scaled from four GPUs to hundreds of GPUs at NVAQC before moving to the full 8,192 GPU deployment on JUPITER.
- This pipeline reduced simulation timelines from weeks to days.

Background & Context

Fault-tolerant quantum computing is required because qubit interactions introduce unavoidable, error-inducing noise. This process relies on "magic states"—specific qubit arrangements consumed during computation. Because the memory required for these simulations grows exponentially with the number of qubits, large-scale validation typically exceeds the capacity of conventional computing systems.

Significance & Outlook

By providing high-fidelity data on how fault-tolerant systems behave at scale, this simulation reduces the reliance on hardware trial and error. This capability allows researchers to predict the performance of future processor designs before they are physically built, accelerating the path toward practical applications in financial modeling, materials science, and drug discovery.

Source: <https://thequantuminsider.com/2026/10/19/nvidia-runs-simulation-of-fault-tolerant-quantum-computing-at-julich-supercomputing-centre/>

Collected: October 24, 2026 | Automated Research System (AI)

#04 IBM and Rutgers: Real-Time Control Across 100 Qubits Explained

Published October 21, 2026 The Quantum Insider USA



OVERVIEW

Researchers from Rutgers, IBM, and collaborating institutions demonstrated real-time quantum feedback control on up to 100 qubits using an IBM Heron processor. By performing nearly 5,000 measurements and resets to counter chaotic behavior, the team validated a critical capability required for future fault-tolerant quantum computing.

Key Findings

Researchers have successfully demonstrated that a quantum computer can monitor its internal state and respond in real-time while running. Using a 156-qubit IBM Quantum Heron processor, the team implemented a connected chain of up to 100 qubits to show that repeated resets can bring chaotic behavior under control. The study, published in *Nature Physics*, observed a sharp phase transition from chaos to order when the frequency of resets outpaced the process of information scrambling.

Technical Details

The experiment involved a competition between two opposing tasks performed nearly 5,000 times each:

- **Scrambling:** A process that mixed information among neighboring qubits, spreading chaos through the system similarly to shuffling a deck of cards.
- **Control:** A process where the computer checked individual qubits and reset them to a specific orderly state.

The team found that when scrambling occurred more than half the time, the system remained chaotic. However, when checking and resetting occurred more frequently, the system shifted abruptly into a controllable state. This demonstration involved nearly 5,000 two-qubit operations and nearly 5,000 measurements and resets while the processor continued to operate.

Background & Context

Quantum information is extremely sensitive to heat, electromagnetic signals, and hardware flaws, which cause errors to accumulate over time. Unlike classical data, quantum information cannot be copied, and measurement can destroy the state. To achieve fault tolerance—the ability to compute correctly despite individual qubit errors—a system must identify and correct errors repeatedly without losing the underlying information.

Significance & Outlook

This experiment represents one of the largest demonstrations of real-time quantum feedback control to date. According to Jedediah Pixley of Rutgers and Maika Takita of IBM, the ability to coordinate thousands of operations, measurements, and resets across a large system is a landmark step toward building larger, more reliable, and fault-tolerant quantum computers capable of simulating complex molecules and running advanced AI programs.

Source: <https://thequantuminsider.com/2026/10/21/rutgers-and-ibm-demonstrate-real-time-control-across-100-qubits/>

Collected: October 24, 2026 | Automated Research System (AI)

#05 University of Oxford: Multi-Species Atom Network and 10s Memory

Published October 16, 2026 Quantum Computing Report UK



OVERVIEW

Physicists at the University of Oxford have demonstrated a heterogeneous quantum network using strontium and calcium ions to generate multipartite entangled states. The team extended entanglement storage in calcium memory qubits to over 10 seconds, exceeding the creation time by more than 100x.

Key Findings

Researchers at the University of Oxford have successfully generated multipartite entangled states across separate trapped-ion modules using a heterogeneous network of two distinct atomic elements: strontium (88 Sr^+) and calcium (43 Ca^+). The team achieved 4-qubit Greenberger-Horne-Zeilinger (GHZ) states across macroscopic distances and stored this distributed entanglement in calcium memory qubits for more than 10 seconds ($\tau = 14 \pm 4\text{ s}$).

Technical Details

The network architecture utilizes microfabricated surface-electrode Paul traps at each node, co-trapping one 88 Sr^+ ion as an optical network interface and one 43 Ca^+ ion as a hyperfine ground memory qubit. The process operates as follows:

- **Entanglement Generation:** Strontium ions spontaneously emit photons at 422 nm, which interfere at a central heralding station to create remote matter-matter entanglement.
- **State Transfer:** Local 402-nm Raman laser pulses drive geometric phase gates and error-detected iSWAP operations to transfer the entangled state from the strontium interface to the calcium memory qubits.
- **Storage Performance:** The 43 Ca^+ ground hyperfine manifold exhibits first-order magnetic field insensitivity at 0.5 mT, which is approximately two orders of magnitude lower than that of strontium.

Background & Context

In standard trapped-ion quantum networks, using a single atomic species for both photonic coupling and quantum storage presents a technical trade-off. The laser pulses and photon emissions necessary for networking typically induce resonant light scattering and decoherence in nearby memory qubits. By employing a mixed-species approach, the Oxford team isolated the networking functions (strontium) from the storage functions (calcium).

Significance & Outlook


The ability to preserve remote entanglement for over 10 seconds represents a significant improvement over the average entanglement generation time of approximately 100 ms. By extending entanglement lifetimes by more than 100x relative to creation time, this research clears a major hurdle for the development of distributed quantum computing and quantum sensing networks.

Source: <https://quantumcomputingreport.com/university-of-oxford-entangles-multi-species-atom-network-and-extends-memory-storage-beyond-10-seconds/>

Collected: October 24, 2026 | Automated Research System (AI)

#06 D-Wave Advantage2: 3D Topological Insulator Simulation via 728 Qubits

Published October 17, 2026 Quantum Computing Report USA



OVERVIEW

D-Wave Quantum has demonstrated the realization of 1D, 2D, and 3D higher-order topological insulators (HOTIs) using an Advantage2 superconducting quantum annealing processor. By employing 728 qubits and new analog-digital control protocols, the team successfully measured 3D bulk topological invariants and fractionalized charge deficits. This confirms the hardware's capability as a programmable simulator for multi-dimensional condensed matter physics.

Key Findings

Scientists at D-Wave Quantum, including Pratik Sathe and Andrew D. King, have successfully simulated higher-order topological insulators (HOTIs) across one, two, and three dimensions. The team achieved the first experimental extraction of 3D bulk topology by measuring the quantized 3D electric octupole moment ($1/2$ invariant), thereby verifying the bulk-boundary correspondence within a 3D quantum simulator.

Technical Details

The research utilized the Advantage2 processor with up to 728 qubits, implementing a hybrid analog-digital architecture that departs from traditional static ground-state sampling. Key technical implementations include:

- **Qubit Dimerization:** Physical qubits were paired into strongly coupled dimers ($J_d = 2.5$ GHz, $\Delta = 3.0$ GHz) to act as effective two-level systems, enabling the realization of Benalcazar-Bernevig-Hughes (BBH) tight-binding models.
- **Wavefunction Tomography:** The system reconstructed the complex-valued propagator $\Psi_{\mathbf{r}}(t) = \sigma_{\mathbf{r}x}(t) + i\sigma_{\mathbf{r}y}(t)$ by time-shifting $\sigma_{\mathbf{r}x}(t)$ measurements by $1/(4\Delta) = 125$ ps, achieving a temporal resolution of $\delta t = 62.5$ ps.
- **Topological Measurement:** The team calculated three successive nested Wilson loops from experimentally reconstructed occupied-band projectors $\{P_v(\mathbf{k})\}$ to extract the 3D bulk topology.
- **Charge Deficits:** Local density of states (LDOS) integrals revealed fractionalized charge deficits ($\approx 1/2$) localized at zero-dimensional corner modes.

Background and Context

The study, published on arXiv under the title “Measuring higher-order bulk topology in a programmable quantum simulator,” focuses on the use of band projector tomography to measure localized boundary modes. Unlike standard quantum annealing, this approach combines pulse-level state initialization with coherent static Hamiltonian evolution and high-frequency time-series measurements.

Significance and Outlook


The results demonstrate that the system is insensitive to local disorder; bulk topological invariants remained quantized near $1/2$ even without ideal spatial symmetries, proving topological protection against mild hardware disorder. This demonstration establishes that analog-digital quantum annealing hardware can function as a programmable simulator for non-interacting topological materials and multi-dimensional condensed matter physics.

Source: <https://quantumcomputingreport.com/d-wave-demonstrates-3d-analog-digital-quantum-simulation-of-higher-order-topological-insulators-using-728-qubits/>

Collected: October 24, 2026 | Automated Research System (AI)

#07 Correlated Decoding: Leiden and Riverlane Fix Distance-Losing Flaw

Published October 22, 2026 Quantum Computing Report USA



OVERVIEW

Leiden University and Riverlane have identified a structural flaw in standard re-weighting protocols for two-pass correlated decoders that causes a unit loss of effective code distance below threshold ($p < 10^{-3}$). The researchers introduced an error-rate-dependent regularization parameter ($\alpha^*(p)$) that restores nominal distance scaling without adding classical runtime overhead. This ensures that error suppression in rotated surface codes scales predictably as qubit gate fidelities improve.

Key Findings

A joint research team from Leiden University (Instituut-Lorentz and Applied Quantum Algorithms) and Riverlane has discovered that standard correlated decoding protocols lead to higher-than-expected logical failure rates on rotated surface codes. Specifically, under realistic circuit-level noise, these protocols cause a unit loss of effective code distance when operating deep below threshold ($p < 10^{-3}$). To counteract this, the team developed a regularization parameter, $\alpha^*(p)$, which restores nominal distance scaling and maintains the accuracy advantages of correlated decoding in the sub-threshold regime.

Technical Details

Correlated decoding is utilized in Union-Find (UF / Local Clustering Decoder) and minimum-weight perfect matching (MWPM / Sparse Blossom) implementations. The process involves two passes: an initial pass identifies errors in one subgraph (e.g., X-type), and the results are used to discount edge weights in the opposite subgraph (e.g., Z-type) for the second pass.

The flaw arises because the standard re-weighting rule discounts multiple partner edges simultaneously without penalizing the joint occurrence of independent secondary faults. This makes long, logically incorrect chains appear artificially cheap, causing size- $(d-1)/2$ fault configurations to trigger logical failures in a distance- d surface code. The proposed fix applies a smooth interpolation parameter $\alpha^*(p) = \alpha_0 / [1 + (p_0 / p)^k]$ to conditional edge probabilities ($p'v = \alpha P(ev | e\mu)$).

- **Near threshold ($p \sim 10^{-2}$):** The optimal re-weighting strength remains near unity ($\alpha \approx 1$).
- **Deep below threshold ($p \leq 10^{-4}$):** $\alpha^*(p)$ shifts toward weaker cross-subgraph discounting to prevent over-discounted corrections.

Significance and Outlook

Because the regularization update can be pre-computed during decoder initialization, the fix introduces zero latency or classical runtime overhead during real-time decoding.

These findings establish a critical design principle for quantum low-density parity-check (qLDPC) codes, belief-matching, and real-time decoding architectures.

Source: <https://quantumcomputingreport.com/leiden-university-and-riverlane-research-reveals-distance-losing-flaw-in-correlated-decoding-below-threshold/>

Collected: October 24, 2026 | Automated Research System (AI)

#08 Iceberg Quantum Pinnacle 1.2: RSA-2048 Factoring Resource Specs

Published October 22, 2026 Quantum Computing Report USA



OVERVIEW

Iceberg Quantum has validated its Pinnacle 1.2 architecture, demonstrating logical error rates below 10^{-12} per logical qubit-cycle. By utilizing the new Crevasse QLDPC decoder, the company has reduced the physical qubit requirements for RSA-2048 factorization to as few as 48,000 qubits.

Key Findings

Iceberg Quantum has released Pinnacle 1.2, an architecture that reduces the spacetime overhead required to break RSA-2048 compared to the previous Pinnacle 1.0. The company reports that RSA-2048 factorization can be achieved with 72,000 physical qubits in 3.7 days at a physical error rate of $p = 10^{-3}$. At a physical error rate of $p = 10^{-4}$ and a 1-microsecond code cycle time, these requirements drop to 48,000 physical qubits and a runtime of 0.98 days.

Technical Details

The architecture's performance is driven by Crevasse, a new quantum low-density parity-check (QLDPC) decoder. Crevasse achieves an average decoding latency of 4 microseconds per error-correction round, allowing for direct circuit-level simulations of trillions of cycles rather than statistical extrapolations. Simulation results showed:

- For $d = 16$ ($[[254, 14, 16]]$) code at $p = 10^{-3}$: three logical errors observed, resulting in a logical error rate of 2×10^{-13} .
- For $d = 10$ ($[[126, 12, 10]]$) code at $p = 10^{-4}$: zero logical errors observed ($\leq 3 \times 10^{-13}$).

These smaller code distances reduce spatial overhead by approximately 40% and temporal overhead by approximately 30% relative to Pinnacle 1.0 benchmarks.

Background and Context

To verify the architecture, Iceberg Quantum used forced-gap post-selection for state injection into magic state distillation blocks. This process achieved injection error rates of $(1.3 \pm 0.3) \times 10^{-6}$ for $d = 10$, with output magic state failure rates near $p_{\text{out}} \approx 3.5 \times 10^{-11}$. The team also tested "Frozen Pinnacle," a variant designed for superconducting systems with fixed 8-degree connectivity. In this configuration, resource requirements are 99,000 physical qubits with a 5.0-day execution runtime.

Significance and Outlook

The validation of logical error rates below 10^{-12} per logical qubit-cycle meets the threshold necessary to execute 2048-bit RSA integer factorization. By combining efficient QLDPC decoding with compiled modular exponentiation circuits, Iceberg Quantum has significantly lowered the physical hardware barrier for fault-tolerant quantum cryptanalysis.

Source: <https://quantumcomputingreport.com/iceberg-quantum-validates-pinnacle-1-2-architecture-with-crevasse-qldpc-decoder-to-reduce-rsa-2048-factoring-resources/>

Collected: October 24, 2026 | Automated Research System (AI)

#09 TRIO Pulses: Robust Single-Qubit Gates for Superconducting Processors

Published October 22, 2026 Quantum Computing Report USA



OVERVIEW

Researchers from Chalmers University and partners have validated Time-Rescaling Robust Inverse-Optimal (TRIO) pulses to mitigate control-amplitude and frequency detuning errors in superconducting qubits. The technique achieved a 99.90% IRB fidelity for a 90-nanosecond X-gate, expanding the high-fidelity parameter range by an order of magnitude compared to standard DRAG Gaussian pulses.

Key Findings

A research team led by the Wallenberg Centre for Quantum Technology (WACQT) at Chalmers University of Technology, in collaboration with Université Bourgogne Europe and Aalto University, has demonstrated a quantum control framework called Time-Rescaling Robust Inverse-Optimal (TRIO) pulses. In tests on a superconducting transmon processor, a 90-nanosecond TRIO X-gate reached an Interleaved Randomized Benchmarking (IRB) fidelity of 99.90%. The system maintained fidelities above 99.5% even when subjected to simultaneous qubit detuning shifts of up to 4 MHz and drive-amplitude variations of up to 20%.

Technical Details

The TRIO protocol addresses limitations found in classical Robust Inverse-Optimal (RIO) pulse designs, which often produce square-like waveforms that cause distortion in finite-bandwidth hardware and induce leakage into upper transmon states. TRIO utilizes a time-rescaling function $g(\tau)$ to map energy-optimal Bloch-sphere trajectories onto smooth, hyper-Gaussian envelopes ($\Omega \text{ TR}(\tau)$), removing high-frequency spectral components while preserving trajectory geometry.

To further optimize performance, the framework integrates first-order Derivative Removal by Adiabatic Gate (DRAG) corrections to mitigate AC Stark shifts and higher-level transmon leakage. The resulting complex envelope is synthesized via in-phase and quadrature (IQ) microwave phase modulation.

Background and Context

The team expanded the framework to a universal single-qubit gate set, including $X \pi/2$ and $Y -\pi/2$ Hadamard-type operations. These operations demonstrated average fidelities of 99.62% and 99.69%, respectively. Additionally, numerical simulations suggest that applying TRIO control to systems with higher anharmonicity, such as 1 GHz fluxonium qubits, could result in gate errors below 10^{-6} .

Significance and Outlook

By expanding the parameter range for high-fidelity execution by an order of magnitude over standard DRAG Gaussian pulses, TRIO pulses provide a method to reduce the overhead associated with periodic re-calibration. This robustness against simultaneous control errors offers a pathway toward more stable operation in large-scale quantum processors.

Source: <https://quantumcomputingreport.com/chalmers-university-demonstrates-trio-pulses-for-control-error-robust-single-qubit-gates-in-superconducting-processors/>

Collected: October 24, 2026 | Automated Research System (AI)

#10 IQM and NVIDIA: Automating 20-Qubit Chip Calibration via VLMs

Published October 22, 2026 Quantum Computing Report USA



OVERVIEW

IQM and NVIDIA have demonstrated agentic AI calibration of a 20-qubit superconducting quantum processor using specialized NVIDIA Ising Calibration vision-language models (VLMs). This integration automates parameter adjustments and diagnostic logging, reducing manual engineering bottlenecks for scaling quantum hardware.

Key Findings

IQM Quantum Computers and NVIDIA have successfully implemented agentic AI to automate the calibration of a 20-qubit superconducting quantum processor. By utilizing specialized NVIDIA Ising Calibration vision-language models (VLMs), the system can autonomously analyze experimental tuning charts and execute dynamic bring-up routines across all 20 physical qubits.

Technical Details

The VLM stack is integrated directly into IQM's proprietary hardware control and calibration software. The workflow performs the following functions:

- Analyzes raw experimental tuning charts to determine optimal parameter adjustments.
- Identifies out-of-bounds qubit parameter settings.
- Dynamically adjusts measurement parameter ranges.
- Generates plain-English diagnostic logs to ensure human auditability of calibration decisions.

The QCalEval Benchmark

To standardize AI-driven quantum processor bring-up, NVIDIA and IQM released QCalEval, an open benchmark dataset. This benchmark consists of 243 calibration charts across 22 experiment types, derived from simulated neutral-atom systems and real superconducting hardware. Each chart includes six diagnostic questions designed to test a model's ability to select corrective actions and identify anomalies.

According to the public QCalEval leaderboard, specialized NVIDIA Ising Calibration models outperformed general-purpose vision-language models by up to 19 points.

Significance and Outlook


The open-weights architecture of the Ising models allows High-Performance Computing (HPC) centers and quantum laboratories to deploy compressed versions locally on edge devices, such as the NVIDIA DGX Station. This framework is intended to eliminate manual engineering bottlenecks as superconducting quantum processing units (QPUs) scale beyond hundreds of qubits, providing a scalable and auditable method for hardware bring-up.

Source: <https://quantumcomputingreport.com/iqm-and-nvidia-automate-20-qubit-chip-calibration-via-nvidia-ising-vision-language-models-and-qcaleval-benchmark/>

Collected: October 24, 2026 | Automated Research System (AI)

#11 NVIDIA and Riverlane: GPU-Accelerated QEC Decoder Calibration

Published October 22, 2026 Quantum Computing Report USA



OVERVIEW

NVIDIA has acquired a strategic minority stake in Riverlane to accelerate real-time quantum error correction (QEC) and decoder calibration. By utilizing H100 GPUs, the partners reduced surface code training cycles from 3 minutes on a single CPU core to 2.8 seconds, enabling the mitigation of dynamic qubit drift.

Key Findings

NVIDIA has invested in Cambridge-based startup Riverlane as part of a broader \$250 million UK computing commitment. This strategic minority equity stake aligns NVIDIA with Riverlane's Deltaflow QEC physical decoder stack. The partnership has already demonstrated a significant performance leap in QEC decoder recalibration workflows by integrating NVIDIA H100 GPUs.

Technical Details

The collaboration focused on adapting Riverlane's Local Clustering Decoder (LCD) for high-throughput GPU parallel processing. The results show a drastic increase in processing speed and throughput:

- **Processing Speed:** A surface code training cycle was reduced from 3 minutes (single CPU core) to 2.8 seconds.
- **Throughput:** The H100 GPU processed over 4 million surface code error syndromes per second, compared to approximately 59,000 on a CPU.

This acceleration allows for routine, pre-experiment decoder recalibration to address hardware defects and dynamic qubit drift.

Background & Context

Scaling fault-tolerant quantum computing (FTQC) is currently hindered by hardware instability, including shifting qubit frequencies and fluctuating two-qubit gate fidelities. Standard QEC decoders rely on pre-computed error likelihood weights; however, these weights become inaccurate as physical qubit behavior drifts, which subsequently degrades logical qubit fidelities.

Significance & Outlook


By using NVIDIA GPUs to process large batches of experimental syndrome data in parallel, Riverlane's AI-driven calibration loop updates the decoder's physical noise model in real time without requiring manual tuning. This creates a tightly coupled QPU-GPU control loop where classical GPU acceleration continuously recalibrates low-latency hardware decoders, such as Riverlane's ASIC/FPGA matrix chips. This integration establishes a technical foundation for the development of reliable, multi-modal utility-scale quantum systems.

Source: <https://quantumcomputingreport.com/nvidia-takes-strategic-stake-in-riverlane-to-accelerate-real-time-quantum-error-correction-and-gpu-decoder-calibration/>

Collected: October 24, 2026 | Automated Research System (AI)

#12 Quantinuum Helios: 99.9995% Logical Memory Fidelity and QEC Specs

Published October 19, 2026 Quantum Computing Report USA



OVERVIEW

Quantinuum has achieved a logical memory fidelity of 99.9995% on its Helios trapped-ion processor using a distance-7 color code. This result demonstrates exponential error suppression, providing a technical pathway toward a universal fault-tolerant system by 2029.

Key Findings

Quantinuum reported a logical memory fidelity of 99.9995% on its commercial Helios trapped-ion quantum computer. This performance corresponds to a logical error rate of 5×10^{-6} per syndrome-extraction round when utilizing a distance-7 code. The results, based on a 137-barium ($^{137}\text{Ba}^+$) QCCD processor, show that logical error rates are suppressed exponentially as the code distance increases from $d = 3$ to $d = 7$.

Technical Implementation

The achievement relies on three primary architectural pillars:

- **Exponential Distance Scaling:** With an experimental suppression factor of $\Lambda = 12$, Quantinuum projects that target logical error rates of 10^{-12} across 100 logical qubits can be reached using approximately 30,000 physical qubits.
- **Leakage Reduction:** The system employs a dual-layer approach to prevent state leakage. This includes hardware-level 3-cycle mid-circuit 1762 nm optical repumping, which reduces leakage by 6.5x, and software-driven SWAP-superdense leakage reduction circuits.
- **Neural Network Decoders:** To enable real-time streaming decoding, Quantinuum implemented 3D Convolutional Neural Networks (3D CNN) and Round-Token Transformers (RTT) on NVIDIA H100 hardware. These models achieved per-round latencies between 0.2 ms and 0.5 ms.

Comparative Analysis

Quantinuum distinguished its logical memory fidelity from Google Quantum AI's magic-state cultivation results on the Willow processor. While Google achieved 99.99% magic-state fidelity via post-selection (retaining ~8% of attempts), Quantinuum's Helios benchmark measures end-to-end logical memory accuracy. Quantinuum noted that in 2025, it previously demonstrated 99.993% magic-state fidelity with an 85% retention rate using its QCCD architecture, requiring smaller space-time overheads.

Significance and Outlook


These empirical results align with the performance targets for Apollo, Quantinuum's planned universal fault-tolerant quantum system scheduled for 2029.

Source: <https://quantumcomputingreport.com/quantinuum-demonstrates-99-9995-logical-memory-fidelity-on-helios-trapped-ion-processor/>

Collected: October 24, 2026 | Automated Research System (AI)

#13 Modular Hyperbolic Codes: Google and MIT's 30x QEC Overhead Reduction

Published October 15, 2026 Quantum Computing Report USA



OVERVIEW

Researchers from Google Quantum AI, Google DeepMind, and MIT have developed a modular quantum error correction (QEC) architecture that reduces physical qubit overhead by over 30x compared to standard 2D surface codes. By networking small Euclidean planar chips in a hyperbolic tiling geometry, the system achieves high encoding rates without requiring complex on-chip long-range wire crossings.

Key Findings

The joint research team has demonstrated a modular QEC architecture that achieves a reduction in physical qubit overhead of more than 30x relative to standard 2D surface codes. This is accomplished by connecting small, flat 2D Euclidean planar chips—containing approximately 100 qubits—using sparse, static inter-module boundary couplers arranged in a hyperbolic tiling geometry with a genus $g \geq 2$.

Technical Details

The architecture partitions the global QEC code into flat modules of either 72 qubits or 116 qubits by fine-graining closed hyperbolic surfaces with Euclidean lattice patches of side-length L . To manage errors, the system utilizes cat-flagged syndrome extraction circuits to measure stabilizers spanning adjacent modules across boundary seams.

Performance was validated through circuit-level simulations using the SI1000 noise model ($p = 0.1\%$), decoded via the AlphaQubit neural network and correlated minimum-weight perfect matching (MWPM). The results showed high tolerance to elevated inter-module seam noise ($p_{\text{seam}} = 1\%$), maintaining low logical error rates even when inter-chip cable error rates were 10x higher than on-chip gate errors.

To enable universal fault-tolerant logic without additional space overhead, the team implemented:

- Fault-tolerant walking circuits that execute logical gates via code automorphisms (global tiling permutations).
- Modular ribbon extractors—closed ancilla loops attached along module boundaries that are approximately 4.5x smaller than the memory itself.

Background & Context

The study, titled “Low-Overhead Quantum Error Correction with Boundary-Connected Planar Modules,” addresses the challenge of implementing Quantum Low-Density Parity-Check (qLDPC) codes. While qLDPC codes offer high encoding rates, they typically require complex long-range wire crossings; this modular hyperbolic approach bypasses that requirement by using a networked chip geometry.

Significance & Outlook


By reducing the qubit overhead by over 30x and demonstrating resilience against high inter-chip noise, this architecture provides a scalable path toward fault-tolerant quantum computing using smaller, manageable hardware modules rather than monolithic, complex chips.

Source: <https://quantumcomputingreport.com/google-quantum-ai-deepmind-demonstrate-modular-hyperbolic-codes-for-over-30x-qec-overhead-reduction/>

Collected: October 24, 2026 | Automated Research System (AI)

#14 IBM Research: AI Agent Factories for QEC Decoder Chip Design

Published October 15, 2026 Quantum Computing Report USA



OVERVIEW

IBM Research has automated the architectural design of classical hardware for Quantum Error Correction (QEC) decoders using AI "agent factories." This approach achieved up to a 204x hardware speedup while reducing AI inference costs by 11.5x through the use of open-weight models.

Key Findings

IBM Research has demonstrated an autonomous multi-agent search framework that automates the design of specialized classical hardware for QEC decoders. By utilizing open-weight models and a distributed inference stack, the team achieved a 52x hardware speedup for an inference cost of \$26. The system reduced blended serving costs to approximately \$0.11 per million tokens, representing an 11.5x reduction in inference costs.

Technical Details

The study focused on High-Level Synthesis (HLS) C++ implementations of the Relay-BP decoding algorithm, specifically for the $[[144, 12, 12]]$ bivariate bicycle (Gross) qLDPC code. The "Agent Factory" framework employs a coordinator model paired with four parallel worker agents to propose architectural transformations, including:

- Restructuring memory into check-major layouts.
- Unrolling loop reductions into compile-time recursive trees.
- Creating dual-layout message stores to achieve Initiation Interval $II = 1$ execution.

To optimize costs, IBM served open-weight models on-premises using the open-source llm-d distributed inference stack, which resulted in a 98% prefix-cache hit rate.

Background & Context

High-rate quantum low-density parity-check (qLDPC) codes are used to significantly reduce the physical-to-logical qubit overhead compared to standard 2D surface codes. However, the dense and non-local check matrices of qLDPC codes require complex, parallel classical decoding hardware to function effectively.

Significance & Outlook


The research, authored by Abhishek Bhandwalder, Mihir Choudhury, Ruchir Puri, and collaborators, establishes that combining open frontier models with efficient inference infrastructure makes autonomous QEC chip design economically practical. This automation is critical for the development of scalable fault-tolerant quantum systems, as it allows for rapid hardware optimization without prohibitive AI operational costs.

Source: <https://quantumcomputingreport.com/ibm-research-automates-qec-decoder-chip-design-using-open-ai-agent-factories-at-11-5x-lower-inference-cost/>

Collected: October 24, 2026 | Automated Research System (AI)

#15 Silicon Spin Qubits: RIKEN achieves >99.9% parallel gate fidelity

Published October 17, 2026 Quantum Computing Report USA



OVERVIEW

Researchers from RIKEN and international partners have demonstrated simultaneous high-fidelity single-qubit operations across a five-qubit silicon quantum dot array. By mitigating AC-Stark shift crosstalk, the team maintained gate fidelities above the 99.9% fault-tolerant threshold using a single shared microwave control line. This advancement reduces wiring density and thermal loads, supporting the scalability of quantum error correction (QEC).

Key Findings

A collaborative team from RIKEN, National Taiwan University, ETH Zurich, and TU Delft has achieved simultaneous single-qubit gate fidelities of $\geq 99.9\%$ across a five-qubit silicon ($^{28}\text{Si}/\text{SiGe}$) quantum dot array. The researchers successfully utilized a single, shared impedance-matched microwave transmission line to control the qubits, crossing the minimum threshold required for 2D surface code quantum error correction (QEC).

Technical Details

The team identified off-resonant microwave-drive-induced AC-Stark shifts as the primary cause of parallel drive crosstalk, which previously degraded gate fidelities from $>99.9\%$ down to 99.5% . To resolve this, they implemented an $N(N-1)$ pairwise calibration protocol. Because the resulting phase shifts ($\Delta\phi_{ij}$) sum linearly, the crosstalk can be compensated by measuring pairwise phase accumulation between individual qubit pairs.

- **Performance:** Up to three qubits operated simultaneously at 99.99% accuracy; all five qubits maintained $\geq 99.9\%$ fidelity.
- **Residual Exchange:** Decoupled CZ sequence measurements confirmed that residual exchange interactions (J_{off}) between neighboring quantum dots remained below 16.6 kHz during single-qubit driving.

Background and Context

Simultaneous multi-qubit operations are essential for cyclic QEC protocols because idle silicon spin qubits dephase significantly faster ($T_2^* \approx 10\text{ }\mu\text{s}$) than driven qubits ($T_2, R \approx 100\text{ }\mu\text{s}$). Previous attempts to drive multiple qubits in parallel suffered from interference that compromised accuracy.

Significance and Outlook

The use of a single shared control line reduces physical wiring density and thermal loads, enabling modular silicon spin-qubit unit cells of approximately 10 qubits per line. The researchers indicate that replacing surface micromagnets with integrated back-end-of-line (BEOL) nanomagnets could double the local field gradient. This modification would allow for faster gate operations (2–3 MHz) at lower microwave power, providing a pathway toward 10-qubit parallel control modules in commercial silicon manufacturing.

Source: <https://quantumcomputingreport.com/riken-demonstrates-high-fidelity-simultaneous-control-of-silicon-spin-qubits-above-99-9-threshold/>

Collected: October 24, 2026 | Automated Research System (AI)

#16 Quantinuum Helios: Progress Toward Megaquop Error Correction

Published October 20, 2026 Science News USA



OVERVIEW

Quantinuum has demonstrated a quantum computer capable of performing one million operations with only approximately five errors. This achievement utilizes a high lambda value in logical qubit scaling, bringing the industry closer to the "megaquop" regime of reliable quantum calculation.

Key Findings

Quantinuum has reported a significant advancement in quantum error correction using its Helios computer. The system achieved an error rate low enough to perform a million operations with only about five errors. This performance is a critical building block for reaching the "megaquop" regime, defined as the ability to execute a million reliable operations.

Technical Details

The Helios system utilizes 98 qubits consisting of electrically charged atoms, known as ions, hovering over a chip. To mitigate the inherent instability of physical qubits, Quantinuum employs quantum error correction to group multiple physical qubits into a single "logical qubit."

A key metric in this process is λ , which represents the factor by which the error rate drops as the logical qubit is scaled up. Quantinuum's system demonstrated a λ of about 10. In comparison, a 2024 demonstration by Google reported a λ of about 2. By utilizing a logical qubit composed of 37 physical qubits, Quantinuum achieved the reported five errors per million operations.

Background & Context

Physical qubits are highly delicate and prone to unintentional state flips (such as a 0 randomly switching to a 1), which can compromise calculations. Logical qubits solve this by integrating multiple physical qubits into a more stable unit. While the current experiment focused on a single logical qubit, the megaquop regime may require the integration of 50 to 100 logical qubits.

Significance & Outlook


The high lambda value of 10 is significant because it reduces the total number of physical qubits required to reach a target error rate. According to Ken Brown, a quantum engineer at Duke University and scientific adviser for Logiqal, this represents a "huge savings" in the overall qubit count. To fully realize the megaquop regime, the current architecture will need to be scaled up from the single logical qubit demonstrated in the research.

Source: <https://sciencenews.org/article/quantum-computer-megaquops-error-correction>

Collected: October 24, 2026 | Automated Research System (AI)

#17 Kipu Quantum: Neutral-Atom Feature Extraction for Toxicity

Published October 17, 2026 Quantum Computing Report USA



OVERVIEW

Kipu Quantum has demonstrated analog quantum feature extraction using QuEra's 256-qubit Aquila processor to improve molecular toxicity classification. The approach increased Balanced Accuracy from 0.55 to 0.70 and doubled the F1-score from ~0.30 to ~0.60 on imbalanced datasets. This result indicates a performance gain over traditional classical molecular descriptors for drug discovery and chemical safety.

Key Findings

Kipu Quantum's experimental results show that analog quantum feature extraction significantly outperforms traditional classical molecular descriptors in classifying molecular toxicity. When tested on imbalanced dataset benchmarks, the quantum features raised Balanced Accuracy from 0.55 to 0.70 and increased the F1-score from approximately 0.30 to 0.60. Among the classical machine learning models used for evaluation, the Support Vector Machine (SVM) classifier achieved the highest overall AUC score of approximately 0.81.

Technical Details

The study, authored by Antonio Ferrer Sánchez, Carlos Flores, and Alejandro Gómez Cadavid, utilized QuEra Computing's 256-qubit Aquila neutral-atom quantum processor via the Kipu Quantum Hub platform. The technical implementation involved the following specifications:

- **Hardware Configuration:** 200-qubit neutral-atom arrays arranged on a 14×15 grid with atoms spaced a few micrometers apart.
- **Process:** The system applied smooth driving pulses for under 4 microseconds to leverage analog Rydberg-state continuous dynamics, avoiding the noise associated with deep, gate-based digital quantum machine learning (QML) circuits.
- **Data Generation:** For each molecule in the UCI Molecular Toxicity dataset, the system generated approximately 1,000 quantum features across 500 independent readout shots per sample.
- **Evaluation:** Features were tested across 100 random dataset splits using XGBoost, CatBoost, Random Forest, SVM, and Google Research's TabFM zero-shot tabular foundation model.

Background & Context

Traditional digital QML algorithms often struggle on current hardware due to the noise inherent in deep, gate-based circuits. By utilizing analog dynamics, Kipu Quantum bypassed discrete gate overhead. This is particularly relevant for imbalanced datasets, where classical models often produce misleading raw accuracy by simply predicting the majority non-toxic class.

Significance & Outlook

The ability to generate high-dimensional quantum features that outperform classical descriptors provides a potential advantage for chemical safety and drug discovery applications. The results demonstrate the utility of neutral-atom quantum processors in enhancing the predictive power of classical machine learning classifiers.

Source: <https://quantumcomputingreport.com/kipu-quantum-demonstrates-neutral-atom-quantum-feature-extraction-for-molecular-toxicity-classification-on-quera-hardware/>