Tech & AI Systems | Nov 20, 2026 | 12 min read | By Bryan White

Harvard and Cornell Join DOE Genesis Mission Quantum AI Consortium to Codesign Error-Corrected Scientific Computing

Tech & AI Systems * Natural Science * Research Dispatch

In November 2026, researchers at Harvard University and Cornell University announced a major multi-institution collaboration funded by a 12 million dollar U.S. Department of Energy (DOE) award under the DOE Office of Science's Phase II Genesis Mission umbrella -- reported alongside the broader 75 million dollar federal Quantum AI initiative. Titled Application-Aware Error Correction Codesign for Scientific Quantum Computing (ASQC) and administered through the DOE's Advanced Scientific Computing Research (ASCR) division, the consortium unites Harvard Principal Investigator Mikhail Lukin and Cornell co-investigator Eun-Ah Kim with partners at MIT, Los Alamos National Laboratory, Lawrence Berkeley National Laboratory, Sandia National Laboratories, QuEra Computing, NVIDIA, and Anthropic.

Rather than treating quantum hardware, error-correcting codes, classical machine-learning decoders, and scientific simulation algorithms as isolated engineering layers, ASQC attacks the central bottleneck of fault-tolerant quantum computing through vertical codesign. Physical qubits remain fragile against environmental decoherence, laser phase noise, and control-pulse drift; bridging the gap from noisy intermediate-scale devices to scientifically useful quantum computers requires shrinking the physical-to-logical qubit overhead of quantum error correction (QEC) by orders of magnitude.

The DOE Genesis Mission and the ASQC Consortium Architecture

The DOE's Genesis Mission is designed to fuse national-laboratory supercomputers, experimental scientific facilities, and frontier artificial intelligence systems into unified discovery engines. Within that framework, ASQC targets a foundational systems challenge: how to run high-value scientific workloads -- such as strongly correlated quantum materials and complex many-body physics -- on early fault-tolerant quantum processors without waiting decades for millions of uniform physical qubits.

Each partner brings a distinct layer of the quantum-AI stack to the consortium. Harvard's Lukin group -- building on its reconfigurable neutral-atom logical processor architecture -- and neutral-atom hardware developer QuEra Computing anchor physical and logical qubit operations. Cornell's Eun-Ah Kim leads machine-learning research for quantum compilers and hybrid quantum-classical materials simulation, while MIT and the three DOE national laboratories (Los Alamos, Lawrence Berkeley, and Sandia) contribute algorithmic benchmarks, verification theory, and supercomputing integration. Industry partners NVIDIA and Anthropic supply GPU-accelerated high-throughput simulation infrastructure and frontier AI models for automated code discovery, compiler optimization, and real-time syndrome decoding.

| Consortium Layer | | Lead Institutions / Partners | | Core Technical Responsibility |

| Program & Award | | DOE ASCR (Genesis Mission Phase II, $12M ASQC award) | | Integrates AI, supercomputers, & quantum hardware for discovery |

| Neutral-Atom Qubits & QEC | | Harvard (M. Lukin, PI), QuEra Computing, MIT | | Reconfigurable atom arrays & logical-qubit error correction |

| AI Compilers & Materials | | Cornell (E.-A. Kim, co-PI) | | Machine-learning quantum compilers & hybrid materials simulation |

| National Lab Benchmarks | | Los Alamos, Lawrence Berkeley, & Sandia National Labs | | Scientific workloads, verification, & HPC integration |

| AI & Accelerated Compute | | NVIDIA, Anthropic | | GPU-accelerated decoding, simulation, & AI-assisted codesign |

Why Conventional Quantum Error Correction Scales Slowly

To protect fragile quantum states from bit-flip and phase-flip errors, quantum error correction encodes a single logical qubit across many physical qubits arranged in a redundant entangled state -- such as a two-dimensional topological surface code or a quantum low-density parity-check (qLDPC) code. Repeated syndrome-extraction circuits measure stabilizer operators on ancilla qubits without collapsing the encoded logical information, allowing a classical decoder to infer which physical errors occurred and apply corrective Pauli frames.

In standard "application-agnostic" fault-tolerant architectures, engineers design QEC for worst-case arbitrary quantum circuits, demanding ultra-low logical error rates per gate (often 10 12 or lower for Shor's factoring algorithm). Achieving those error floors with conventional surface codes can require hundreds or thousands of physical qubits per logical qubit, plus slow lattice-surgery or magic-state distillation factories that consume the vast majority of hardware space and clock cycles.

Application-Aware Codesign: Tailoring Fault Tolerance to Physics

Scientific quantum simulations differ fundamentally from cryptographic factoring. Many condensed-matter, quantum-chemistry, and lattice-gauge-theory workloads possess intrinsic symmetries, local interaction structures, and natural tolerance to small statistical fluctuations in expectation values -- meaning not every circuit layer requires identical worst-case code distance.

Application-aware codesign exploits that structure across three coupled axes:

First, symmetry-preserving encoding and compilation aligns logical gate sequences with the native physical symmetries of the target Hamiltonian -- such as particle-number conservation or spin-rotation invariance -- so that unphysical symmetry-violating errors are flagged or suppressed naturally at the algorithmic level. Second, heterogeneous code-distance allocation assigns deeper error protection only to sensitive non-Clifford operations or long-lived memory registers, while keeping short-depth state-preparation or measurement stages lean. Third, hardware-native connectivity mapping restructures the quantum circuit to exploit the

exact native entangling gates and physical qubit geometry of the target processor rather than forcing generic all-to-all decompositions.

Reconfigurable Neutral-Atom Logical Processors and Real-Time AI Decoders

Neutral-atom arrays -- pioneered by Mikhail Lukin's Harvard laboratory and commercialized by QuEra Computing -- offer an unusually flexible hardware substrate for application-aware QEC. In a landmark 2024 Nature demonstration (DOI: 10.1038/s41586-023-06927-3), Lukin and colleagues showed that individual rubidium atoms trapped in optical tweezers can be coherently shuttled across a two-dimensional processor while preserving quantum superposition, enabling parallel transversal entangling gates between entire logical qubit blocks.

Because optical-tweezer shuttling provides dynamic, non-local connectivity, neutral-atom processors are not restricted to nearest-neighbor planar surface codes; they can implement higher-rate qLDPC codes and transversal logical operations that dramatically reduce circuit depth. However, non-local connectivity and correlated shuttling noise also make syndrome decoding far more complex than decoding a static grid. Traditional minimum-weight perfect matching algorithms struggle with hypergraph error correlations across transversal logical gates -- creating an ideal opening for AI-driven syndrome decoders trained on realistic physical noise models and executed on GPU hardware to keep pace with microsecond-to-millisecond experimental cycle times.

Machine-Learning Quantum Compilers and Hybrid Materials Simulation

At Cornell University, Eun-Ah Kim's research group brings deep expertise in applying machine learning to quantum many-body physics and quantum information. Within ASQC, the Cornell team focuses on machine-learning-guided quantum compilers and hybrid quantum-classical simulation workflows tailored to strongly correlated quantum materials -- systems such as unconventional superconductors, frustrated magnets, and topological phases of matter where classical tensor-network and quantum Monte Carlo methods hit exponential entanglement or sign-problem walls.

An AI-enhanced quantum compiler does far more than translate abstract gates into hardware pulses. By learning how specific hardware noise channels propagate through variational, Trotterized, or phase-estimation circuits, neural compilers can automatically synthesize shorter, noise-resilient pulse schedules and error-mitigated measurement protocols -- extracting accurate ground-state energies and dynamical correlation functions from hybrid quantum-classical runs before full-scale, million-qubit fault tolerance arrives.

Bridging Classical Supercomputers and Early Fault-Tolerant Hardware

Across DOE national laboratories -- Los Alamos, Lawrence Berkeley, and Sandia -- high-performance computing centers already operate some of the world's fastest GPU-accelerated supercomputers. Integrating NVIDIA's accelerated computing stack and Anthropic's AI systems into ASQC reflects a broader architectural shift: future scientific quantum computers will not operate as standalone replacements for classical supercomputers, but as tightly coupled quantum co-processors orchestrated

by classical AI.

During a hybrid simulation, classical GPUs and AI agents handle real-time QEC decoding, circuit optimization, variational parameter updates, and classical shadow or tensor-network post-processing, reserving scarce logical quantum operations strictly for the classically intractable entanglement steps. By co-designing the physics application, the error-correcting code, and the AI decoder simultaneously, ASQC aims to pull practical scientific quantum advantage years closer -- turning early logical-qubit processors into usable instruments for materials discovery and fundamental physics.

References & Primary Sources

* [1] Bluvstein, D., Evered, S. J., ... & Lukin, M. D. (2024). Logical quantum processor based on reconfigurable atom arrays. Nature, 626, 58-65. -- https://doi.org/10.1038/s41586-023-06927-3

* [2] Harvard University (2026). Application-Aware Error Correction Codesign for Scientific Quantum Computing (ASQC) -- DOE Genesis Mission Award. -- https://news.harvard.edu/

* [3] Cornell University (2026). Cornell Arts and Sciences Researchers Join DOE Quantum AI Project. -- https://news.cornell.edu/

* [4] Quantum Zeitgeist (2026). Harvard & Cornell Arts And Sciences Join $75M Energy Department Quantum AI Project. -- https://quantumzeitgeist.com/