Wood Mackenzie

COLUMBIA | SIPA
School of International and Public Affairs

The Quantum Disruption of Energy

Advisors: Jeanne Fox, Ben Hertz-Shargel, Trent Yang
4/24/2026

Team Members

Akshay Kalyan MIA IFEP
E-mail ak5309@columbia.edu

Vincent Saputo, MPA CEE
E-mail vfs2113@columbia.edu

Ruby Wu MPA CEE
E-mail jw3912@columbia.edu

Coco Guo MPA CEE
E-mail zg2548@columbia.edu

Jialong W MPA CEE
E-mail jialong.w@columbia.edu

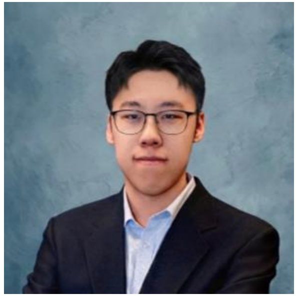

Terry Zhang MPA CEE
E-mail tz2645@columbia.edu

Jeanne Fox
E-mail jf2949@columbia.edu

Table of contents

01. Project Overview & Quantum Computing 101

02. Techno-economic Analysis And Matrix Evaluation

03. Applications & Analysis

04. Policy Recommendations

05. Commentary & Conclusion

Executive Summary – 1/3

- **Quantum computing will complement, not replace, classical high-performance computing.**
    - Use cases rely on hybrid quantum-classical routines to address specific classically intractable bottlenecks, in coordination with AI and HPC-enabled methods.
    - The quantum advantage will be dynamic. Improvement and integration of AI will evolve the best classical alternatives, raising the bar over time.
    - Notably, some quantum workflows may offer energy efficiency gains ahead of raw computational speedups.

- **Quantum computing will undergo a step-change adoption trajectory.**
    - Industry adoption of quantum computing will feature discrete step-change moments rather than steady uptake.
    - Qubit error rates and decoherence remain a key hurdle to overcome – robust, low-overhead error correction solutions will unlock commercial viability for many applications at once.
    - Fault-tolerant quantum computing will drive the broad commercial appeal of QC, most likely by 2035.
    - Use cases and investment momentum will mutually reinforce as these thresholds are sequentially crossed.

Executive Summary – 2/3

- **Quantum computing adoption will depend on the hardware modality race.**
  - There is no clear modality winner. Current leaders – annealing and semiconducting – may eventually cede ground to ascendant platforms, e.g. photonic or topological systems.
  - Industry convergence towards 2-3 dominant modalities will be driven by ancillary requirements (operating temperature, vacuum), fault tolerance progress, cost economics, application breadth, and geopolitical factors.
  - Progress on gate modalities will eventually erode quantum annealing's niche in combinatorial optimization.

- **Quantum computing has very compelling energy sector applications across time horizons.**
  - Near-term opportunities - Quantum sensing and grid optimization use cases address existing operational pain points with clear commercial outcomes at stake, without relying on fully-mature QC technology.
  - Long-term potential – Quantum computing is a research catalyst that can accelerate R&D in battery and solar technology, but success demands sustained policy support and investment. Commercial outcomes will also hinge on the economics of manufacturing novel materials.

Executive Summary – 3/3

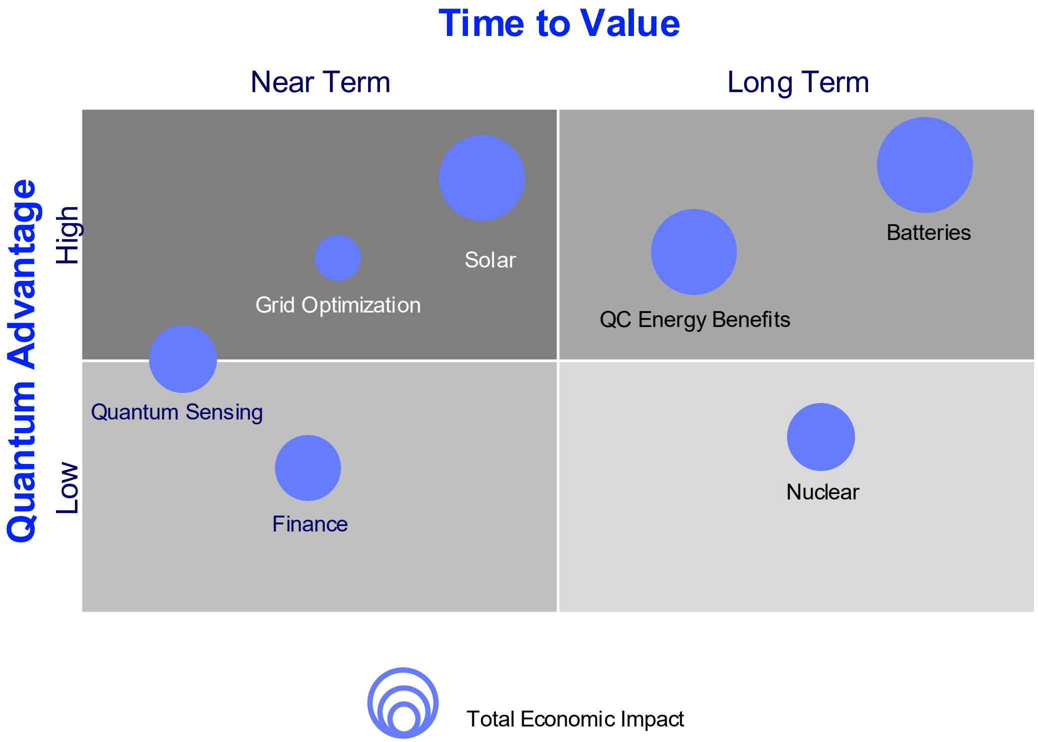

Wood Mackenzie

COLUMBIA | SIPA
School of International and Public Affairs

Project Overview

Project Overview

This capstone project conducted in partnership with Wood Mackenzie explores the transformative potential of quantum computing within the global energy landscape. The initiative focuses on comprehensive technoeconomic analysis to identify and evaluate the most viable quantum computing use cases across the energy value chain.

Key Objectives

This project outlines:

- What quantum computing can realistically achieve
- How it would "disrupt" the status quo
- Implications to industry, investor, and/or governmental stakeholders
- Current and expected R&D work
- A qualitative assessment of the likelihood of disruption being achieved, and the associated timelines for it
- Key overarching themes and strategic recommendations for industry leaders and policymakers

What is Quantum Computing?

Classical vs. Quantum	Regular computers process information as 0s and 1s. Quantum computers use *qubits* that can be both at the same time. • Ideal for problems with enormously large combinatorial search space (drug design, encryption, logistics, power grid optimization). • Google's 2024 Willow quantum chip performed a benchmark computation in **under 5 minutes** that would take the world's fastest classical supercomputer an estimated **10 septillion years** to complete.
Why it matters now?	Global electricity transmission and distribution losses waste an estimated **$130 billion annually.** Quantum optimization algorithms applied to grid management could reduce these inefficiencies at a new scale (IEA, 2023). • Estimates from McKinsey & Co. indicate quantum applications could generate **$400–$600 billion in economic value** in Finance industry by 2035. • The U.S. Department of Energy has committed **$625 million** to quantum energy applications through its National Quantum Initiative.
The Path Forward	The transition from current experimental devices to world-changing applications requires overcoming "noise" and scaling physical hardware. • **The Error Challenge**: Qubits are extremely fragile; environmental interference (heat, vibration) causes decoherence, leading to calculation errors. • **Quantum Error Correction (QEC)**: Moving beyond the NISQ era (Noisy Intermediate-Scale Quantum) by using thousands of physical qubits to create one stable logical qubit. • **A Hybrid Future**: Quantum units (QPUs) will function as specialized accelerators alongside classical CPUs/GPUs, much like how GPUs revolutionized AI training.

Quantum Computing 101

Qubits

The basic unit of quantum information

Qubits store information using quantum states rather than classical bits. Unlike classical bits (0 or 1), qubits can represent multiple possibilities simultaneously. These are the building blocks of quantum computing.

Example:
IBM and Google build superconducting qubit processors used in quantum computing research.

Quantum Gates

Operations that manipulate qubits

Quantum gates control and transform qubit states, similar to logic gates in classical computing. By applying sequences of gates, quantum computers perform calculations.

Example:
IBM's Qiskit platform allows developers to design circuits using quantum gates.

Quantum Algorithms

Problem-solving methods

Quantum algorithms organize many quantum gates into structured processes to solve specific computational problems.

Example:
JPMorgan and IonQ have explored quantum algorithms for financial portfolio optimization and option pricing.

Error Correction

Ensuring reliable computation

Quantum states are extremely sensitive to noise and hardware imperfections. Error correction techniques are used to detect and correct errors so that computations remain sufficiently reliable.

Example:
Google's quantum research team demonstrated improved quantum error correction on their Sycamore processor.

Wood Mackenzie

COLUMBIA | SIPA
School of International and Public Affairs

Qubits rewrite the rules of the game for Monte Carlo simulations

Monte Carlo - A method for estimating uncertain outcomes by running millions of random simulations, i.e. stochastic method. Used across finance (derivatives pricing, risk quantification), energy (weather forecasting, grid planning), drug discovery, and logistics.

	Classical MC	Quantum MC (QAE)
How	Run each simulation sequentially, one at a time	Qubits encode all possible price paths in superposition simultaneously — then use interference to amplify the answer
Convergence rate	$O(M^{-1/2})$ - to halve the error, you need 4× more samples	$O(M^{-1})$ - quadratically fewer samples
Samples for 1% error on energy option pricing	~10,000 random paths	~100 quantum oracle calls
Speedup	—	~100× fewer samples
What changes	More samples = more time	Interference extracts answer from superposition

QC offers "all-to-all" connectivity instead of the classical "point-to-point" connectivity – e.g. QC can compare two time series all at once, instead of comparing individual points one pair at a time.

IBM showed an average speed-up between cubic and quartic for a quantum implementation of the Markov Chain Monte Carlo (MCMC) Metropolis Hastings algorithm.

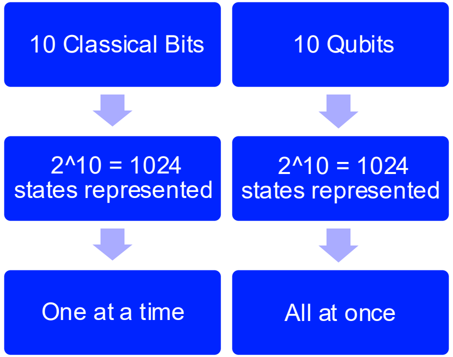

Source: Montanaro, Woerner and Egger

The quantum computing value chain comprises tech giants, startups, and international players

QC Infrastructure

- BLUEFORS
- QUANTUM MACHINES
- THORLABS
- elementsix DE BEERS GROUP

High-purity silicon, superconductors, dilution refrigerators

QC Hardware

- IBM Quantum
- XANADU
- Google Quantum AI
- IONQ
- rigetti
- D-Wave The Quantum Computing Company

Quantum Processing Units (QPUs)

QC Software

- Google Cloud
- Microsoft Azure
- FUJITSU
- riverlane
- QUANTINUUM

Quantum Software Developing Kits (SDKs), Quantum Error Correction Software

Quantum-as-a-Service (QaaS)

- 本源量子 Origin Quantum
- Terra Quantum
- BOSCH
- Q-CTRL
- exail TECHNOLOGIES

Industry and research applications in energy, finance, cryptography, and more

Quantum computing hardware modalities are in a race towards commercial viability

Several technologies are being tested and developed to build, stabilize, and manipulate qubits cost-effectively.

Technology	Description	Investment Outlook	Cooling Requirement
Superconducting	Encode quantum information within electromagnetic fields and manipulate qubits through high frequency microwaves.	Steady, market-leading funding	◆ (red)
Trapped Ion	Have ions suspended above a surface by radio-frequency electric fields and uses laser pulses to manipulate qubits.	Explosive growth (1B+ USD in 2025)	◆ (yellow)
Neutral Atom	Use neutral atoms (e.g., Rubidium) as qubits, manipulated by focused laser beams, i.e. optical tweezers in a vacuum.	Steady	◆ (yellow)
Photonic	Use light photons, which are challenging to disturb, but offer low decoherence and high speeds.	Explosive growth (1B+ USD in 2025)	◆ (green)
Silicon Spin	Encode quantum information within the intrinsic spins of electrons and compute by manipulating those spins.	Steady	◆ (red)
Topological	Encode quantum information in anyons (2-D quasiparticles) that intertwine to form braids.	Speculative/R&D, minimal public raise	◆ (red)
Quantum Annealer	Represents optimization problems as an energy landscape, with the solution being the lowest energy state.	Steady, D-Wave had a $400M raise in July	◆ (red)

◆ (red) Most intensive cooling infrastructure to achieve milliKelvin temperatures. Requires dilution refrigeration (DR), which is expensive.

◆ (yellow) Moderate cooling needs. Systems operate at room-temperature or inside mild cryogenics but require a vacuum and laser cooling of individual qubits for control.

◆ (green) Least intensive cooling infrastructure. Cryogenic qubit control is not needed, but detectors may require cooling.

Source: [QuEra](), [Quantum Insider]()

Wood Mackenzie

COLUMBIA | SIPA
School of International and Public Affairs

Quantum computing hardware modalities are in a race towards commercial viability

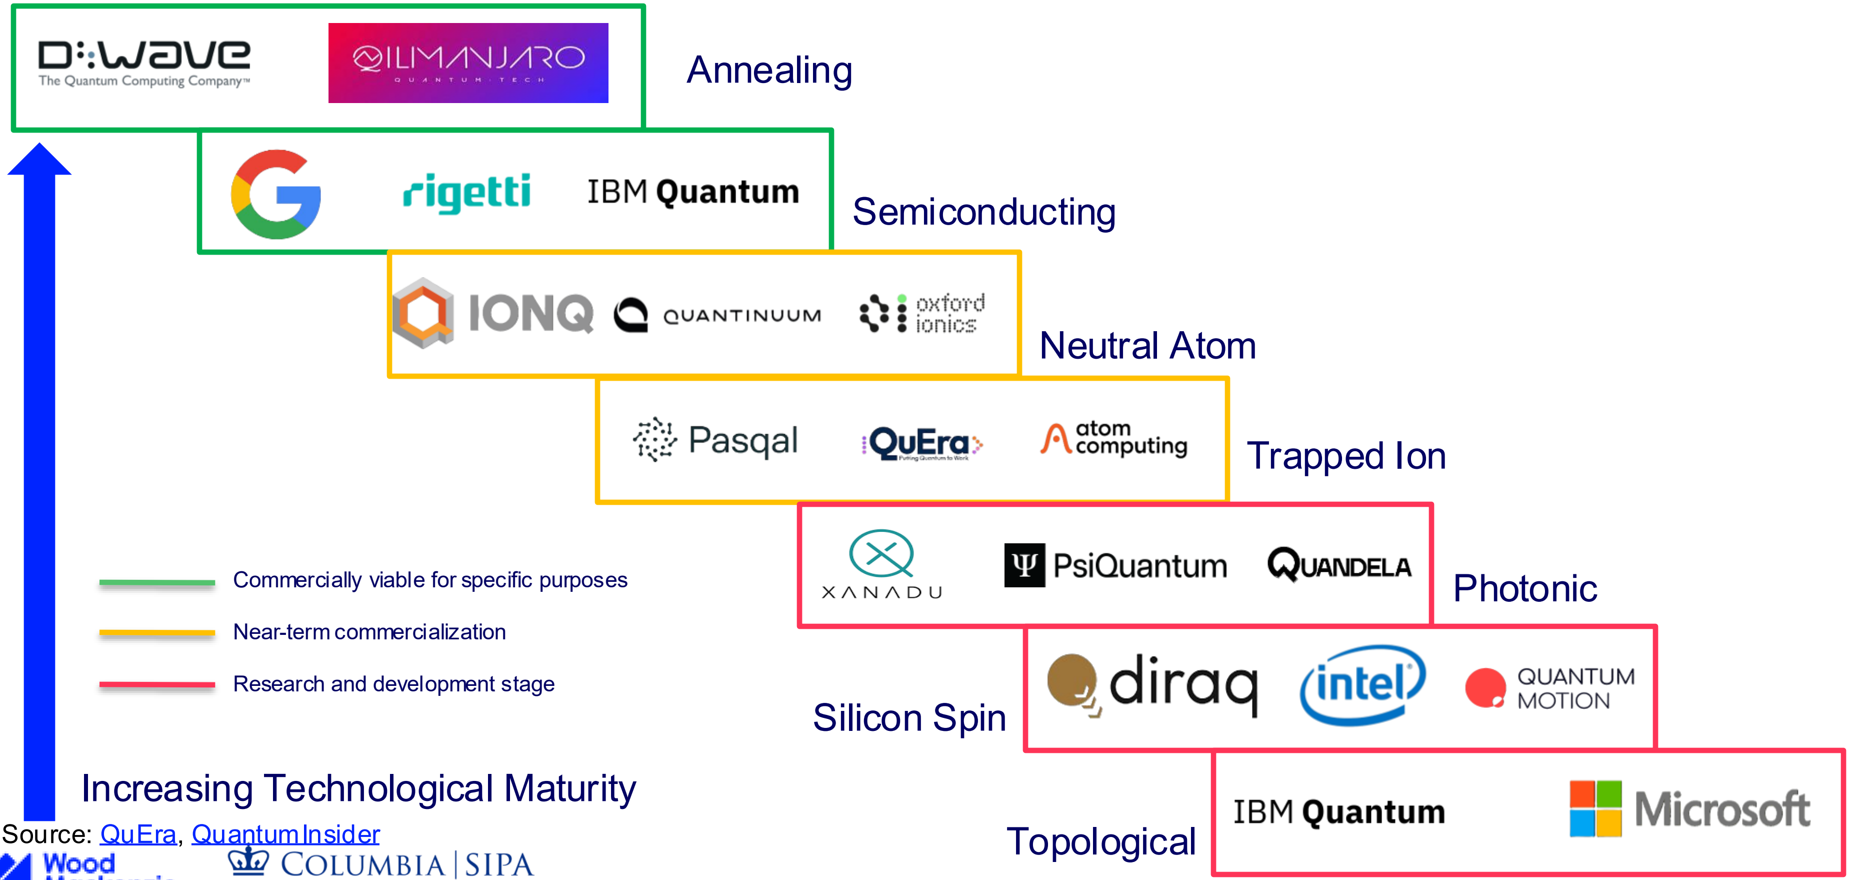

Source: QuEra, Quantum Insider

Wood Mackenzie

COLUMBIA | SIPA
School of International and Public Affairs

Quantum computing is not science fiction. This is what is already possible.

Pasqal and EDF

Quantum computers help forecast renewable energy generation under different weather conditions and accordingly refine EV charging schedules.

D-Wave and Deloitte

Quantum-annealing optimizers allowed Deloitte to improve the scheduling of airport security shifts involving more 60,000 employees and 450 airports.

Fujitsu and University of Osaka

Fujitsu and the University of Osaka have developed an optimization technique that reduces the hardware requirements and calculation times for complex molecular simulations by orders of magnitude.

D-Wave and SLB

SLB demonstrated that quantum computing helped optimize the scheduling of equipment and labor resources to drill 88 oil wells in 33 days.

IBM and ExxonMobil

ExxonMobil has implemented quantum algorithms for maritime routing optimization of their LNG fleet involving over 500 ships. The optimization challenge involves 2^1,000,000 possible combinations of discrete decisions in their global LNG shipping network.

IBM and Cornell University

Cornell researchers encoded model predictive control (MPC) for building energy management as a quadratic unconstrained binary optimization (QUBO) challenge and proved a 6.6% energy efficiency gain compared to classical methods using IBM's 127-qubit QC.

Source: PV Magazine, HPC Wire

Approach to Quantum Opportunity Assessment & Prioritization

01
Technology Identification
Map quantum advantages to specific use cases

02
Baseline Benchmarking
Map quantum advantages to specific use cases

03
Quantum Lift Modeling
Quantify speed-up, accuracy gains & resource reduction

04
Scenario Analysis
Impact under near/mid/long term horizons

05
Comparative Assessment
Analysis on timeline to value and quantum advantage

Primary Goal: Score and rank quantum computing use cases across energy applications to prioritize further evaluation

Analysis of Quantum Use Cases: Matrix Methodology

Three pillars of evaluation

1

Time to Value (X-Axis)

Measures the feasibility of achieving quantum advantage relative to current hardware roadmaps.

Grading Logic:

5 20+ Year Horizon
4 15-20 Year Horizon
3 10-15 Year Horizon
2 5-10 Year Horizon
1 1-5 Year Horizon

Timeline scale is inverted: lower scores reflect nearer time to value

2

Quantum Advantage (Y-Axis)

Compares the marginal benefit of quantum applications vs. advanced classical computing.

Grading Logic:

5 Exponential Speedup
4 High-Order Polynomial
3 Substantial Marginal Gain
2 Minimal Classical Delta
1 No Quantum Advantage

3

Cost-Benefit (Bubble Size)

A composite estimate of hardware access, integration effort, and specialized application development.

Grading Logic:

5 Low barrier
4 Moderate costs
3 High R&D/custom dev
2 Prohibitive hardware CapEx
1 Theoretical

Scoring Process

Criteria responses are entered into a structured Excel model, which outputs matrix placements for each use case

Inputs

Literature & Research
Academic papers, hardware roadmaps, industry reports

Expert Interviews
Interviews with leaders in quantum research and industry

Use Case Data
Pilots, qubit benchmarks, cost estimates

Model

Excel Scoring Model

Rubric inputs score 1-5 based on criterion; directly maps to matrix logic

Sample Outputs

Grid Optimization
- Near-Term
- Moderate Quantum Advantage
- Low Cost Barrier

Battery Materials
- Long-Term
- High Quantum Advantage
- High Cost Barrier

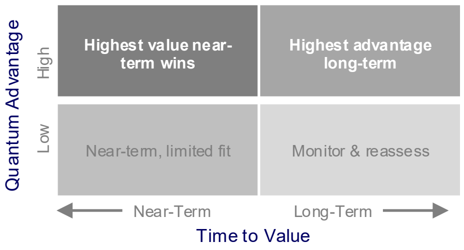

Upper-left is highest priority. Near time to value and high quantum advantage.

Upper-right represents transformational long-horizon opportunities. High quantum fit but requires fault tolerant hardware.

Lower quadrants reflect limited quantum advantage regardless of timeline. These should be monitored for future developments but not prioritized.

Quantum computing will save energy in AI data centers

Problem	• Classical High-Performance Computing (HPC) consumes lots of energy, e.g. the world's best supercomputer, Frontier, needs 21 MW, or 161 GWh/year • Data centres strain an overburdened energy grid, pushing prices and technical limits
Computational Constraint	• Efficiency improvements in classical computing hardware may be plateauing • AI use requires energy-hungry compute for training, tuning, and inference of LLMs • Inference comprises 60-90% of LLM energy needs due to high frequency of requests
Quantum Fit	• As a different paradigm, QC uses orders of magnitude less energy for computing • Algorithms and operations that generate more quantum entanglement will see the greatest relative energy savings; the green advantage of QC is problem-specific • QC energetic advantages could be realised well before computational time savings
Impact	• QC-based AI will improve all stages of the LLM lifecycle, potentially with highest impact in inference due to the repeated calculations required • This could significantly reduce data center loads, speeding up development and interconnection with much lower stress on energy systems and prices

Training GPT-4 consumed ~62000 MWh of electricity

~ 51 metric tons CO_2 saved per task by QC annealer vs. classical HPC

10 – 100 millikelvin temperatures needed for most QC

Source: Jaschke and Montangero, Pasqal, Qilimanjaro

Wood Mackenzie

COLUMBIA | SIPA
School of International and Public Affairs

Quantum computers hold great promise for more energy-efficient computing

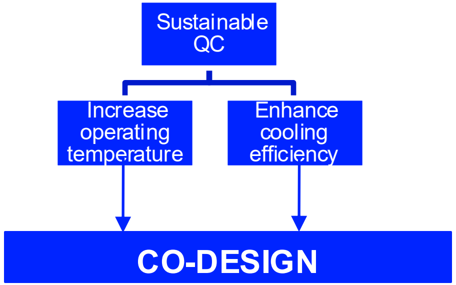

Enhancing the sustainability promise of QC requires integrated design. Today's small quantum devices already provide power savings compared to ultramodern classical machines.

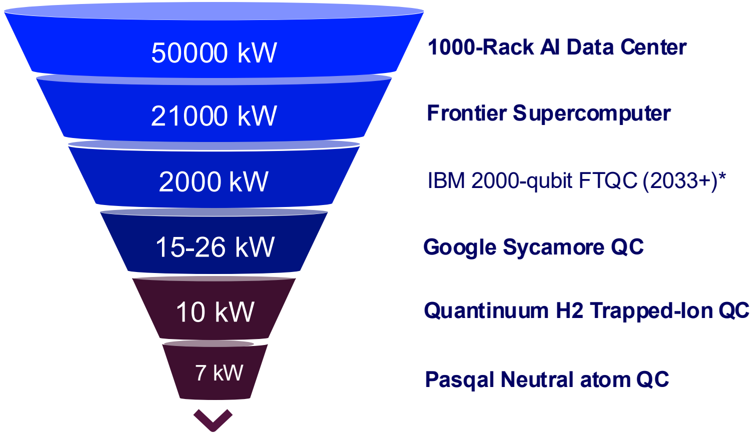

*Planned
The other quantum computers listed are built-for-purpose and have limited commercialization.

Theoretical heat transfer physics predicts computational capacity vs. energy usage relationships for QC.

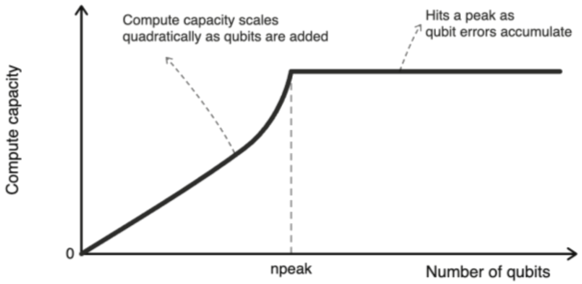

Quantum volume, a measure of computing power, will be scaled quadratically with qubits up to a limit.

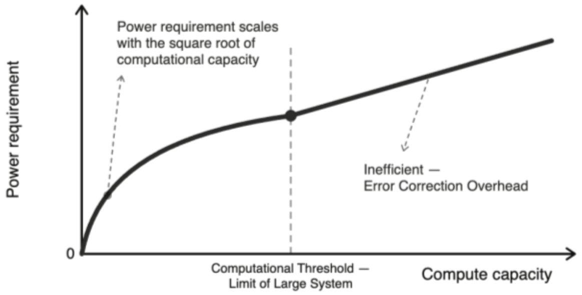

QC power requirement will scale to the square root of quantum volume until a threshold, beyond which, large errors cause inefficiency. Better qubit control will push the threshold outwards, making QC more useful.

Source: NREL, Pasqal

Quantum computing will save energy in AI data centers

Timeline and Risk
Fault-tolerant quantum hardware of practical scale for LLMs is likely 10-15 years away, but hybrid routines can offer energy advantages in the interim, even before computational speedup.

Quantum Fit
Classically intractable problems with high-entanglement algorithms show best fit. QC low-rank approximation and reinforcement learning are well-suited to model distillation for LLM training. Hybrid QC will save energy for inference problems using more than 46 qubits.

Economic Impact of Implementation
The potential energy savings could be enormous, leading to significant implications for the energy transition in the digital age.

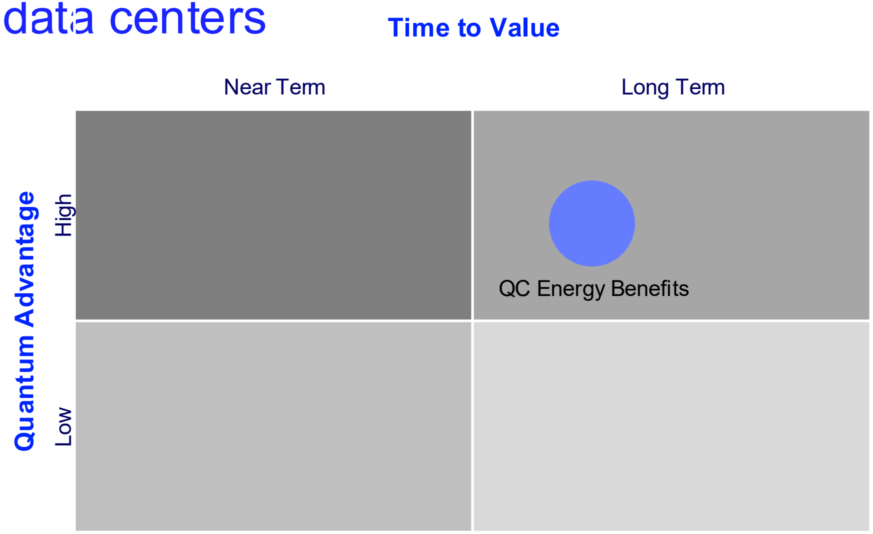

Scenario	Quantum Milestone	Energy Savings	Economic Impact
Conservative	Quantum development stalls prior to achieving energy advantage.	~10%	Incremental energy use reductions; primarily in tailor-made facilities for research applications.
Base	Limited use of QC in hybrid algorithms, co-located in larger data centers.	~20-30%	Significant energy use reductions by integrating QC into classically difficult subroutines of LLM training.
Optimistic	Fault-tolerant QCs integrated into core data centers in key markets.	~50-80% for LLM development and operation	Complete overhaul of the digital infra-energy nexus.

Quantum Computing has transformative potential in energy grid optimization

Problem

- Rapid integration of intermittent renewable energy sources has turned the grid into a highly decentralized network.
- Balancing supply and demand in real-time and coordinating thousands of unpredictable generation nodes makes this extremely difficult.

Computational Constraint

- Energy grid optimization involve massive, high-dimensional problems such as Optimal Power Flow (OPF) and Unit Commitment (UC).
- The sheer combinatorics of these problems can be overwhelming. Current classical computing methods struggle to process these calculations quickly.

Quantum Fit

- Current research is heavily focused on hybrid frameworks that pair classical processors with Noisy Intermediate-Scale Quantum (NISQ) devices.
- QC can evaluate vast combinatorial solution spaces simultaneously.

Impact

- QC can achieve further efficiency gains, allowing electricity markets to clear with better efficiency, and better integrating renewable energy.
- Quantum-embedded robust optimization models can identify "worst-case" damage scenarios during extreme weather events, which allow for preventive response plans that balance operating costs while minimizing load-shedding losses.

Using FERC's framework, we project ~**$10b annual savings with 5% efficiency gain in dispatch** from AC-OPF algorithms

Between 2022 and 2027, distributed fuel-based generation and distributed storage will **grow 240% and 460%** respectively

Source: FERC, Wood Mackenzie

Real-time operations (UC & OPF) present the highest value-add for near-term QC use

QC formulations of power grid problems use well-known classical problem types. Unit Commitment is often coded as Mixed-Integer Linear Programming (MILP), while Optimal Power Flow is typically coded as Non-linear Programming (NLP).

	Unit Commitment and Optimal Power Flow	Grid Partitioning and Expansion Planning
Frequency of Use	High; every 5-15 minutes	Low; quarterly or annually
QC Value-add	Operational efficiency and OPEX savings	Strategic advantages and CAPEX savings
Computational Complexity	Large-scale on-linear problems with time and physical constraints	Extensive variables, but can run for hours on classical computers
QC Fit	High; QC delivers valuable, real-time solutions	Low-medium; relaxed time constraint makes classical computing more appealing

Source: IonQ, QuantumInsider

OAK RIDGE National Laboratory | IONQ

- Successfully utilized IonQ's 36-qubit system to solve a multi-variable generation scheduling problem with **24 time periods and 26 generators.**
- Grid-scale deployment is actively targeting the 100-to-200 qubit systems arriving in 2026.

NREL NATIONAL RENEWABLE ENERGY LABORATORY | atom computing

- Designed a **Quantum-in-the-loop (QIL)** framework which allows a quantum computer to interact with real-time grid simulations and actual hardware.
- Researchers demonstrated the interface by coordinating electric vehicle charging and exploring how to quickly switch power sources during emergencies, such as a city evacuation for a hurricane, to improve resilience.

QC achieves grid optimization improvements in experimental settings

Timeline and Risk
Utilities operate in a high-stakes environments when delivering power, and demand high optimality rates that current QC struggles to meet. Still, market-relevant QC application is envisioned in the next 5-10 years.

Quantum Fit
QC excels at solving combinatorial problems such as unit commitment and optimal power flow that have to be solved very frequently. Other problems, such as capacity and transmission planning, are less suitable for QC due to less frequency. The speed advantages of QC may not outweigh the costs.

Economic Impact of Implementation
In high frequency calculations such as unit commitment, the benefits and cost savings of QC are magnified. However, utilities have traditionally hesitated to use novel methods for efficiency gains due to the risks of errors.

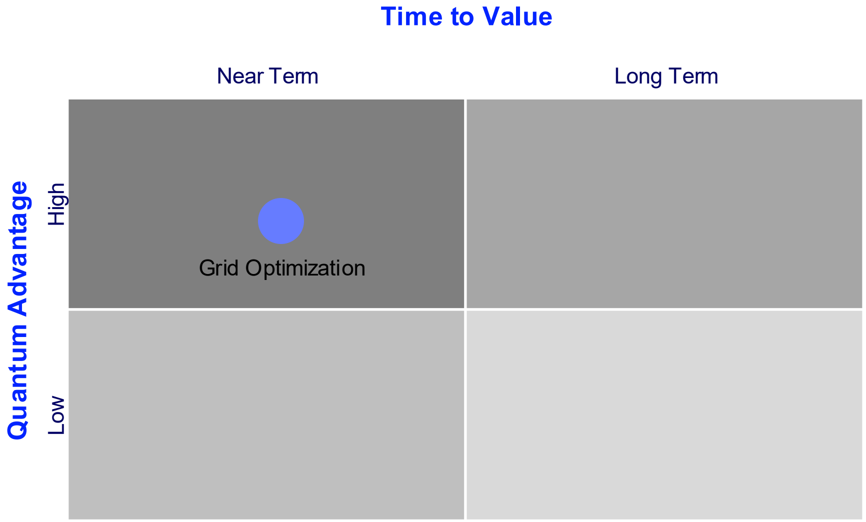

Scenario	Quantum Milestone	Efficiency Achieved	Economic Impact
Conservative	Optimality rates hovering around 95% (the current state)	Inconsistent performance reduces any efficiency gains	Limited: utilities unlikely to adopt, or only greenlight pilot projects
Base	Optimality rates reaching 98-99%	Consistent but modest performance gains	Moderate: Some regional or national utility adoption with tangible improvements
Optimistic	Optimality rates reaching 99.5%+	Quadratic or exponential improvements	Significant: Better grid management lowers grid OpEx.

Wood Mackenzie
COLUMBIA | SIPA
School of International and Public Affairs

Quantum computing can help design better batteries

Problem

- Battery limitations - energy density vs duration trade-off, cyclic degradation, low efficiency, and safety (flammability) prevent scale-up of storage technologies.
- Better battery chemistry is critical to increasing energy density and durability.

Computational Constraint

- Exact solutions to Schrödinger's equation for reactions are impossible, so Density Functional Theorem (DFT) is used to make approximations
- Electrolyte and electrode chemistry research pushes beyond the limits of DFT and the resultant errors prevent meaningful research progress

Quantum Fit

- QC runs on the same principles as quantum mechanical systems or reactions
- Algorithms allow direct mapping of molecular Hamiltonians to quantum hardware
- QC offers near-exact solutions with comparable computational costs.

Impact

- QC could expand the search space for new materials and enhance predictive accuracy about how those molecular systems will behave
- Battery cost, capacity, and durability improvements will enhance energy storage economics, transforming transportation decarbonization and storage for the grid.

Quantum-generated information trains classical AI models

Simulations run in **weeks not years**

750+ miles theoretical range for future solid-state battery EVs

Source: IonQ

Quantum computing can help design better batteries

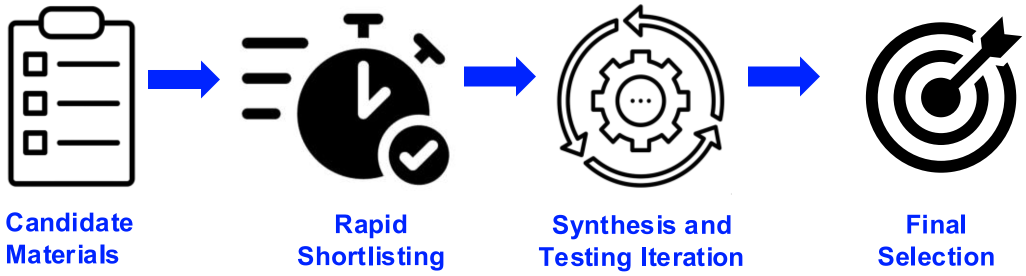

IBM x Daimler (2020)	IonQ x 1Qbit x Dow (2021)	Microsoft x Pacific Northwest National Lab (PNNL) (2024)	IBM x University Consortium (2026)
Simulated ground state energies for intermediate products of Li-S batteries and calculated the dipole moment of Li-H using 4 qubits of a QC.	Used problem decomposition approach – split complex problems into subparts that are designed for QC supremacy and implemented hybrid routines.	Used AI models to screen 32 million new candidate materials for batteries in a week-long process, avoiding 20 years' worth of calculations – basis for new Quantum-AI approaches.	Engineered a new molecule, $C_{13}Cl_2$, with a novel "cork-screw" electronic structure and unique chemical properties.
Li-S batteries could have up to twice the energy density of existing Li-Ion chemistries.	Achieved chemical accuracy in calculating electron structure for a ring of 10 hydrogen atoms.	Found a novel material for batteries to test - NaxLi3−xYCl6.	Decoded the half-Mobius electronic topology through quantum diagonalization on a 100-qubit QC.

Source: IBM, IonQ, IBM, IEEE Spectrum

Quantum computing can help design better batteries

Timeline and Risk
10 to 20-year timeline for industry-relevant quantum computing for chemistry research acceleration to be commercially useful.

Quantum Fit
High fit. Complex chemical processes are best embodied by quantum mechanical systems. Quantum advantage will be seen for small-to-medium sized chemical systems.

Economic Impact of Implementation
The potential energy savings could be enormous, leading to significant implications for the energy transition in the digital age.

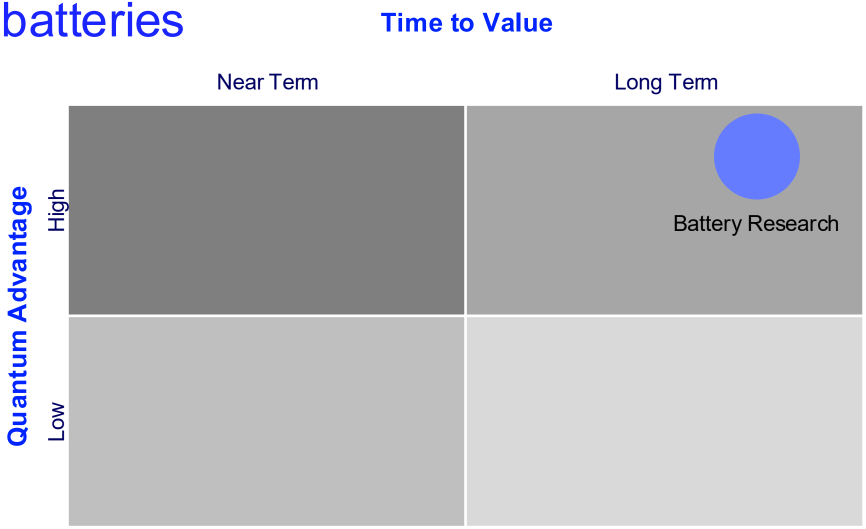

Scenario	Quantum Milestone	Efficiency Achieved	Economic Impact
Conservative	Error-prone QC persists, with applications restricted to academic research	~5%	Better scientific understanding of chemical systems and processes
Base	Limited commercial use of quantum annealing and QPE for specific questions yield moderate efficiency gains	~20%	New computing protocols adopted in industrial research efforts.
Optimistic	Fault-tolerant QC and AI-integrated research efforts rapidly produce new electrolytes and electrode materials	Orders of magnitude speedup of research.	New battery chemistries with 50% higher energy density and cycle endurance revolutionize EVs, solar + storage, and the energy mix

Quantum computing could disrupt solar energy economics

Problem	• Current solar PV extracts one electron from one photon, and efficiency is running up against the Shockley-Queisser limit: 29% • R&D endeavours are now at the tail end of diminishing marginal returns. In current form, solar is near the maximum possible efficiency by area.
Computational Constraint	• Singlet Fission (SF) promises to extract two electrons from one photon, significantly increasing electricity output from solar PV. • Classical techniques struggle to accurately model SF stability in materials suited for PV design, and the Shockley-Queisser limit holds.
Quantum Fit	• Quantum computers can simulate correlated electron behavior and multiple exciton quantum states natively. • Quantum algorithms can model necessary wavefunctions efficiently, enabling screening of SF materials at a scale impossible for classical computers to reach.
Impact	• SF-enabled solar PV stands to dramatically increase output per square meter of newly installed capacity, disrupting solar project economics. • Addition of new materials would also reduce degradation and extend lifetime of solar PV.

$718B
Market Size by 2035

2x
Output Gains

3-5yr
Panel Life Extension

Source: Omega Silicon

Singlet Fission: The Quantum Computing-Enabled Breakthrough

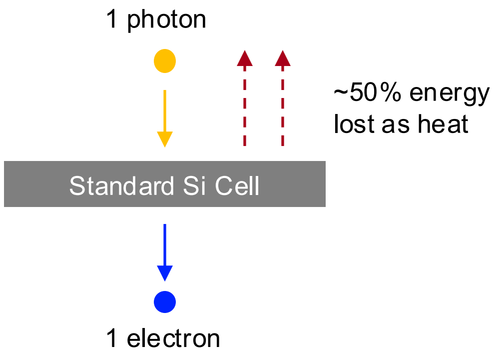

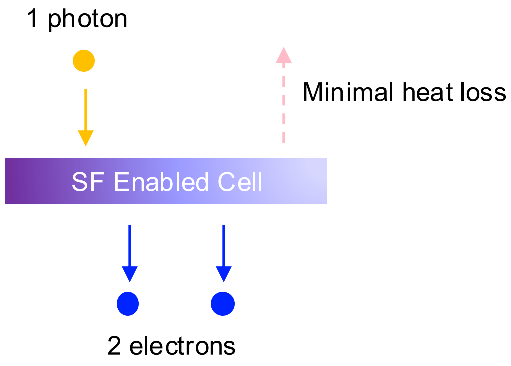

Traditional Solar PV

- Current PV extracts exactly one electron from one photon; any excess energy is lost
- High-energy photons generate excess heat rather than additional current
- Current R&D is hitting the Shockley-Queisser Limit of 29% because classical modeling cannot easily find ways around this 1:1 ratio

Quantum Enabled Solar PV

- One high-energy singlet exciton is split into two lower-energy triplet excitons via Singlet Fission
- Modeling these correlated electron behaviors and multiple exciton states is a native quantum task that classical computers struggle to perform accurately
- Quantum algorithms can efficiently model the wavefunctions needed to find stable SF materials at a scale impossible for classical techniques

Source: Claudino et al. Baldo et al.

Wood Mackenzie

COLUMBIA | SIPA School of International and Public Affairs

QC is positioned to unlock 2x solar PV efficiency gains

Timeline and Risk
Quantum hardware capable of simulating SF-enabled materials is likely 7-10 years out, but manufacturing scale up would take additional time, placing this in the medium- to long-term bucket.

Quantum Fit
Quantum algorithms unlock R&D of materials that enable significant efficiency advantages.

Economic Impact of Implementation
Marginal efficiency gains in solar translate to enormous economic value at grid scale.

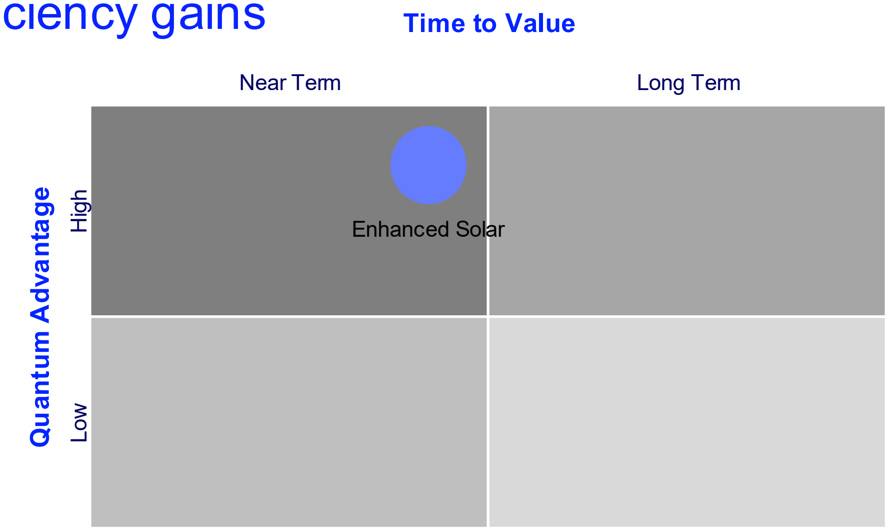

Scenario	Quantum Milestone	Efficiency Achieved	Economic Impact
Conservative	Limited quantum simulation gains	~30-32%	Incremental LCOE reduction, modest market differentiation
Base	Fault-tolerant simulation of SF materials	~38-42%	Significant LCOE reduction, new market entrants
Optimistic	Full SF-enabled PV commercialization	44%+	Structural shifts in global solar economics

Quantum computing can unlock a faster path to next-gen nuclear

Problem	• Classical computing struggles with reactor physics simulations that model many-body quantum systems for neutron transport, isotopic evolution, or radiation damage. • These design problems require approximations or conservative assumptions which consume massive compute budgets, inflate costs, and delay certifications by years.
Computational Constraint	• Neutron transport, uncertainty quantification (UQ), and heavy-element chemistry are exponentially hard problems that quantum algorithms are purpose-built to solve. • Classical methods face curse of dimensionality - memory bottlenecks, inconvenient scaling of variance, under-sampling of rare events – causing fidelity vs. cost trade-offs.
Quantum Fit	• QC directly encodes many-body nuclear Hamiltonians, quantum solvers handle classically intractable thermodynamics, QC random walks mitigate under-sampling. • Hybrid QC protocols can target the most difficult subroutines: constrained core-design optimization, transport/UQ sampling, and molecular simulation of fuel chemistry.
Impact	• Engineers can run more safety scenarios faster, giving regulators higher-confidence answers and reducing design iterations that slow down the licencing process. • QC workflows would speed up R&D cycles for next-gen reactors in the pipeline, enabling modelling of advanced fuels, cladding materials, and grid integration.

811M core-hours for HPC nuclear research (FY2022)

30 GW SMR interests and agreements for data centers

Commercial fission works, but advanced reactors still depend on compute-heavy analysis

Fission & Fusion

Fission
Scientists discovered that splitting certain heavy atoms releases a great deal of heat and more neutrons. Today's plants use that controlled chain reaction.

Fusion
Fusion joins light atoms into a heavier one and releases energy. It powers the sun and remains a separate, longer-term research path.

How a reactor works, and where labs spend compute time

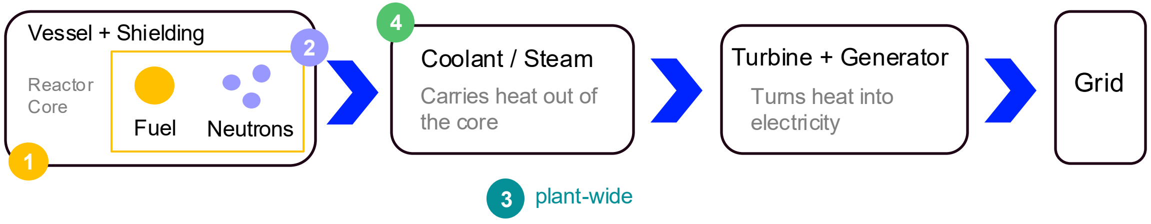

1 Fuel and core layout
Fuel pellets are packed into rods and fuel assemblies. Their arrangement changes power, fuel use, and safety margins.

2 Neutrons and shielding
Fission releases neutrons that keep the chain reaction going. Analysts calculate where they travel and how much shielding is needed.

3 Safety and uncertainty
Regulators need confidence ranges, not one answer. Engineers rerun many cases to prove the plant stays within safety limits.

4 Materials and chemistry
Fuel, coolant, salts, and metals must survive heat, radiation, and corrosion. Heavy-element chemistry is hard for common approximations.

What advanced reactors are trying to improve

- Generation III and III+ are today's large commercial reactor families, mostly water-cooled.
- SMRs are smaller units aimed at modular construction, flexible siting, and lower upfront capital.
- Generation IV designs use new coolants or higher temperatures to improve sustainability, economics, safety and reliability, and proliferation resistance.

Some Small Modular Reactor projects are already moving toward early commercialization, and the same research gains can carry over there too.

Why these tasks are a fit for QC

- These tasks match the clearest known quantum strengths: constrained optimization, repeated transport and uncertainty sampling, and hard chemistry calculations.
- They already consume large high-performance-computing budgets, so even a narrow gain in runtime, uncertainty, or chemistry accuracy could matter.
- Near-term path is hybrid: classical high-performance computing runs the overall reactor workflow, while quantum is tested on one hard subroutine.

Nuclear × Energy: Hybrid QC can accelerate advanced nuclear R&D—supporting faster SMR / Gen-IV deployment for firm AI-era power

Timeline and Risk
Near-term value is in hybrid pilots for optimization and transport/UQ—not full reactor simulation. System-level impact depends on benchmarked gains in neutron-transport UQ or fuel optimization, likely in the long-term.

Quantum Fit
QC fits narrow nuclear workflow better than whole-plant models: optimization, transport/UQ, and newer materials. LLMs help with coding and surrogates, but not validated transport results, shielding margins, or licensing-grade uncertainty.

Economic Impact of Implementation
Nuclear R&D already needs top-of-the-line HPC. QC could accelerate design loops, expand simulation scenarios, and tighten uncertainty as demand for firm low-carbon power rises.

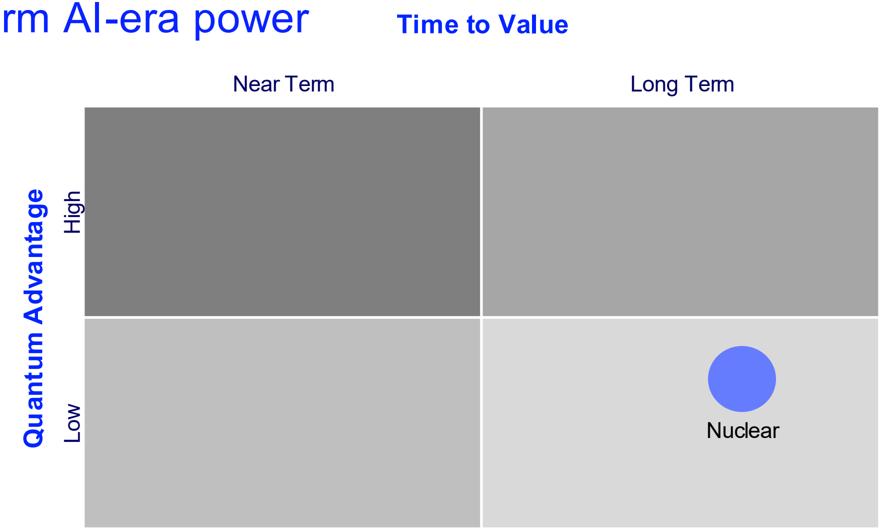

Scenario	Quantum Milestone	Efficiency Achieved	Economic Impact
Conservative	Benchmark-first hybrid pilots in fuel optimization or neutron-transport UQ	~5–10% targeted kernel acceleration	More analyses per budget; incremental engineering throughput
Base	QC integrated into selected HPC/ML workflows for SMR / Gen-IV studies	~10–25% better cost-to-accuracy on chosen kernels	Faster design iteration; tighter safety / shielding uncertainty.
Optimistic	Early fault-tolerant QC for transport or materials subroutines	Potential order-of-magnitude gains on narrow workloads	Shorter R&D / licensing cycles; stronger nuclear cost and time competitiveness as firm power.

Quantum sensing can unlock a new layer of physical intelligence

Unlike in the other computational quantum applications, quantum sensing requires dedicated hardware devices that leverage quantum measurement principles.

Problem	• Subsurface resources, reservoir behavior, and emissions are difficult to measure directly, forcing reliance on noisy geophysical proxies and costly exploration. • Key geophysical signals (e.g., micro-gravity or magnetic anomalies) are extremely weak and often buried in environmental noise or survey variability.
Computational Constraint	• Existing sensors lack sensitivity to detect small gravitational, magnetic, or chemical signals indicating mineral deposits, methane leaks or reservoir changes. • Extracting useful information requires complex modelling and large sensor networks such as repeated 4D seismic surveys, increasing monitoring cost and complexity.
Quantum Fit	• Quantum gravimeters (e.g., Exail Absolute Quantum Gravimeter) detect minute subsurface density changes linked to fluid movement or reservoir depletion. • Quantum photonic and diamond sensors enable ultra-sensitive detection of methane leaks and subtle magnetic anomalies.
Impact	• Higher-resolution subsurface monitoring enables improved reservoir management, mineral discovery, and geothermal resource identification. • Energy operators gain earlier methane leak detection and improved carbon storage monitoring while reducing the need for costly exploration and seismic campaigns.

Emissions Detection
ppm-level methane detection (≤100m range)

Subsurface Imaging
Detect micro-gravity field changes 100x smaller than classical limits

Infrastructure Monitoring
~1–3 kg/hr methane detection threshold (industry benchmark)

Quantum sensing has applications across energy & infrastructure

0 - 2 years

Oil & Gas Exploration
Detect subsurface density changes for precise reservoir mapping, avoiding dry wells and reducing exploration risk.

Emissions Detection
Identify methane leaks and fugitive emissions with parts-per-trillion sensitivity across wide areas.

3 - 5 years

Mineral Exploration
Map subsurface ore deposits and geological structures with precision that classical surveys cannot match.

Groundwater Monitoring
Track aquifer depletion, recharge, and subsurface water movement for operations in water stressed areas.

Geothermal Resources
Map geothermal heat sources and fluid pathways, cutting power plant siting survey costs.

10+ years

Grid Infrastructure
Real-time monitoring of grid stability, structural integrity, and early warning of system-scale failures.

Quantum sensors detect gravity, magnetic, chemical, and motion signals, with orders of magnitude greater precision, unlocking new visibility into subsurface and infrastructure systems.

Quantum sensing is market-ready, offering near-term precision gains

Timeline and Risk
Closer to commercialization than QC, with early deployments emerging in geophysical surveying and navigation. Quantum gravimeters developed by Exail are already used in geophysics and volcano monitoring.

Quantum Fit
Quantum sensors can detect extremely weak physical signals such as gravity changes, magnetic fields, and gas emissions, offering capabilities beyond the reach of traditional sensing devices.

Economic Impact of Implementation
Higher-precision sensing enables earlier detection of infrastructure risks and more accurate subsurface mapping. In energy systems, this could improve fossil fuel reservoir monitoring, carbon storage verification, and methane leak detection.

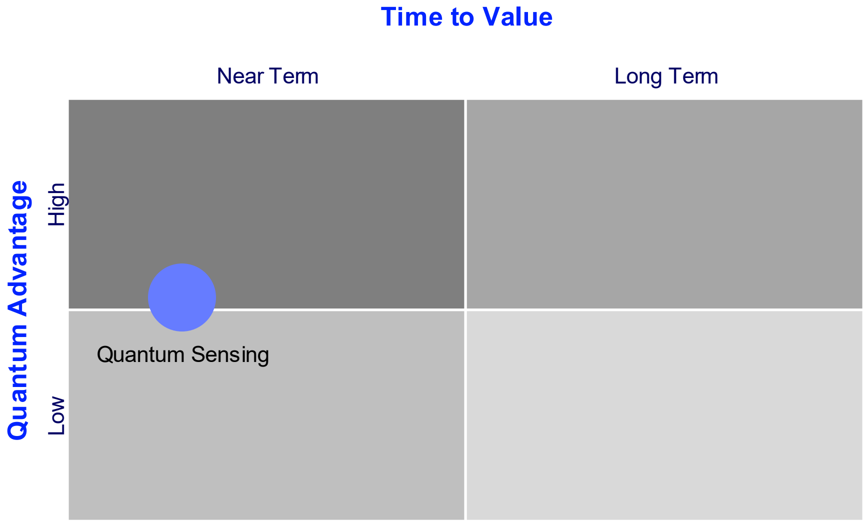

Scenario	Quantum Milestone	Efficiency Achieved	Economic Impact
Conservative	Early deployment of quantum magnetometers.	~5–10× sensitivity improvement	Incremental gains in grid and infrastructure monitoring.
Base	Commercial quantum gravimeters and inertial sensors.	~10–100× precision improvement	Improved underground mapping for geothermal and carbon storage
Optimistic	Integrated quantum sensing networks.	Real-time ultra-precision monitoring	Large cost and time savings in risk management, discovery of new oil and gas fields, safer energy systems.

Quantum computing speeds up pricing of energy derivatives

Problem	• Energy derivatives require constant re-pricing under volatile, weather-dependent market conditions, often running millions of Monte Carlo simulations daily. • Carbon credit and renewable energy certificate markets add new layers of pricing complexity that classical models struggle to handle at scale.
Computational Constraint	• Classical **Monte Carlo** converges at $O(M^{-1/2})$, meaning that accurate pricing of complex multi-asset energy options demands significant compute time and cost. • Portfolio-level risk (VaR/CVaR) on commodity books also requires nested **constrained optimization and high-dimensional pattern recognition**.
Quantum Fit*	• Quantum Amplitude Estimation (QAE) provides a near-quadratic speedup to $O(M^{-1})$, achieving equivalent pricing accuracy with 1,000 samples instead of 1,000,000. • QUBO/QAOA formulations map directly to commodity portfolio optimization solving problems that classical solvers handle poorly beyond 50 assets.
Impact	• E.ON and IBM developed a quantum algorithm for weather-risk energy derivatives hedging; Goldman Sachs targets production quantum pricing by 2028. • HSBC × IBM achieved 34% improvement in bond trade prediction on real production data (1.1M trades, 5,000+ bonds).

$97B
TAM by 2035

1,000x
Performance Uplift on Pricing

16x
Cost Reduction

*Note: QAE's quadratic Monte-Carlo speedup will have multi-industry benefits. Energy derivatives have particularly intensive workloads due to weather-driven volatility and continuous repricing.

Source: McKinsey, 2025; Goldman Sachs & IBM; HSBC;

Additional Applications in Energy Finance & Quantum Advantage Outlook

E.ON × IBM Collaboration

Quantum Monte Carlo for Energy Derivatives Pricing: Application to weather-risk hedging.

- Utilizes QAE-based Monte Carlo acceleration on IBM's 27-qubit processor.
- IBM anticipates quantum advantage **by 2029** for specific applications.

Goldman Sachs × IBM × Microsoft Initiative

Pioneering **resource estimation** for quantum advantage in derivatives pricing.

- Quantum Subspace Expansion (QSP) achieves T-gate reduction of **~16x** and logical qubit reduction of **~4x** for complex models.
- Potential for **1,000x faster** derivatives pricing on fully capable quantum hardware.

Quantum Carbon Finance Innovations

Leveraging **Quantum Amplitude Estimation for Carbon Emission Rights Option pricing** and portfolio optimization.

- Validated using data from the Beijing Green Exchange to ensure real-world applicability.
- Carbon credit portfolio optimization demonstrates a **31.6% improvement** over classical benchmarks in risk-adjusted returns.

Timeline to Quantum Advantage in Energy Finance

Achieving quantum advantage involves distinct phases, each defined by critical milestones and technological advancements:

Near-term (Now–2028)

Focus on **hybrid quantum-classical pilots** and the development of quantum feature extraction techniques. Early applications include optimizing carbon credit portfolios and initial explorations in energy derivatives with current noisy intermediate-scale quantum (NISQ) devices.

Medium-term (2028–2033)

Anticipated arrival of **early fault-tolerant quantum computers**. Key milestones include Goldman Sachs targeting QC pricing algorithms by 2028 and IBM's roadmap aiming for 200 logical qubits by 2029. This phase will see more robust applications in complex financial modeling.

Long-term (2033+)

Expected realization of **full-scale quantum advantage**, enabling real-time Monte Carlo simulations for previously intractable problems. This phase promises up to a 1,000x speedup in derivatives pricing and transformative impacts across energy finance, allowing for highly accurate risk assessments and portfolio optimizations.

Source: IBM, Liu et al, Stamatopoulos & Zeng (Quantum, 2024), Goldman Sachs & IBM, Ribeiro et al., Goldman Sachs via *The Quantum Record*.

Potential to Achieve 1000x Speedup in Derivatives Pricing

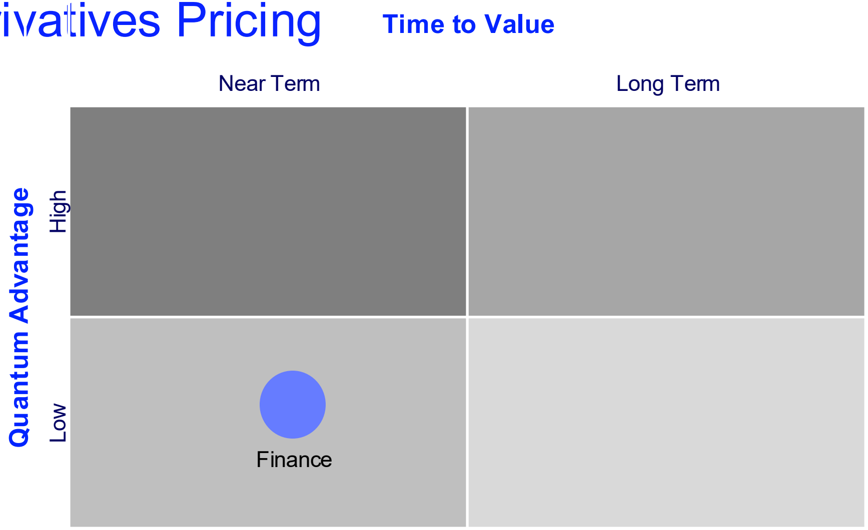

Timeline and Risk

Early quantum advantage targeted for 2028–2029. Near-term hybrid models face performance caveats in mid-range probabilities, advancements like Quantum Signal Processing (QSP) are bringing viable timelines closer.

Quantum Fit

Quantum algorithms are distinctly suited for running complex Monte Carlo simulations and solving high-dimensional combinatorial problems.

Economic Impact of Implementation

Achieving up to a 1,000x speedup in derivatives pricing or a 30%+ improvement in portfolio optimization translates to a massive competitive edge in volatile energy and carbon markets.

Scenario	Quantum Milestone	Efficiency Achieved	Economic Impact
Conservative	Hybrid algorithms and feature extraction.	Mainly model validation	Incremental validation of quantum models for utilities.
Base	Early fault-tolerant advantage (2028–2029) via QSP.	~4x fewer logical qubits required	Accelerated pricing of energy derivatives.
Optimistic	Full-scale Quantum Amplitude Estimation (QAE)	1,000x faster Monte Carlo sampling	Structural shift in energy trading; real-time pricing of complex, multi-asset energy options at scale.

Quantum Application Matrix: Technologies that stand to benefit most

Near Term Wins

- Grid optimization and quantum sensing applications offer the highest near-term ROI, making them ideal targets for first-mover investors.
- Quantum annealing-based applications have already yielded commercial benefits across industries.
- Financial institutions have shown the scope for quantum applications in simulation and trading.

Risk-Return Tradeoff

- High quantum advantage technologies (batteries, enhanced solar) require longer development horizons but represent transformational opportunities for energy systems at large.
- While quantum computing may unlock material discovery, total economic benefit is contingent upon cost constraints and manufacturing capabilities.

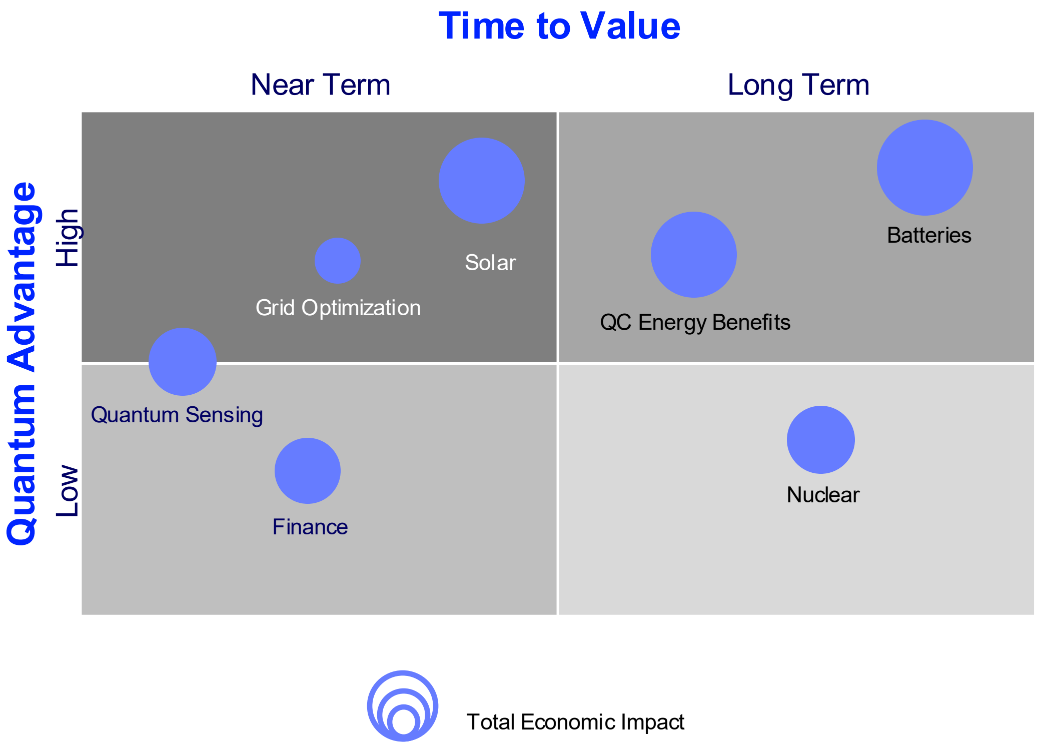

Global policy is moving quantum computing beyond lab research towards wider access, demonstrations, and security planning

North America

- U.S.: Dept. Of Energy (DOE) renewed the five National Quantum Information Science (QIS) Research Centers with $625M in 2025.
- Canada: added $245M over 5 years and launched Canadian Quantum Champions Program Phase 1 (up to $67M) plus a Benchmarking Quantum Platform.

Europe

- UK: $3.4B / 10-year quantum strategy, plus $60M for testbeds and catalyst funding; government explicitly cites energy-grid use cases.
- EU: adopted the Quantum Europe Strategy in 2025; a Quantum Act is planned for 2026.
- Germany: DLR Quantum Computing Initiative (QCI) is already running 60+ software/application projects

China

- The 2026 Two Sessions / 15th Five-Year Plan cycle identified quantum technology as a future-industry priority.
- Separately, China's 2025 State Council guideline defines "scenarios" as a bridge from R&D to market adoption — a deployment model that is relevant to quantum as well.

Asia-Pacific

- Japan calls 2025 the "first year of quantum industrialization."
- Singapore is backing its National Quantum Strategy with $240M in investments.
- Australia is funding challenge-led programs, including optimizing energy networks and quantum timing for distributed energy systems.

The strongest national strategies now do more than fund research: they open test environments, support real demonstrations, and prepare critical systems for the energy and security transition.

Source: Canada, DLR, QuantumInsider

Stakeholders have distinct roles in moving quantum from pilots to energy sector deployment

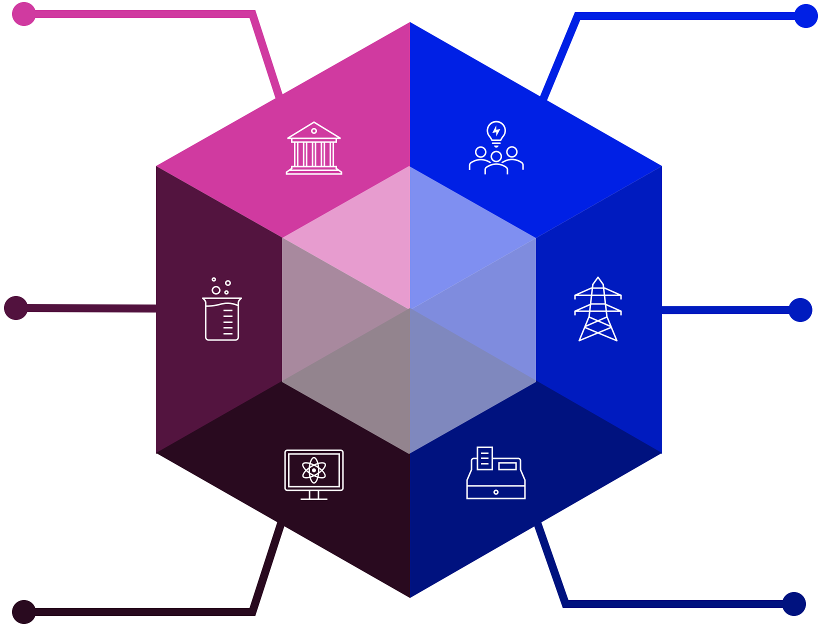

Government

- Regulation support for quantum startups
- Funding for the quantum pilot programs
- Establish coordinating entity for inter-departmental collaboration on applications
- Cooperate with other countries on standards and technology sharing

Academia / National Labs

- Increase academic focused on QC, offering students new career pathways
- Open test-beds and QC-centric labs
- Undertake QC industry partnerships to enhance research access to QC hardware

Cybersecurity / Standards Bodies

- Invest in post-quantum cryptography (PQC) and migrate systems before fault-tolerant QC
- Integrate QC into procurement risk reviews by vetting vendors and protocols

Private Investors

- Number of quantum startups with VC-backing has skyrocketed
- Short-term: QC should be a strategic focus within a broader HPC exposure, not a core pillar of the strategy
- Diversify along the QC value chain

Utilities / Grid Operators

- Undertake QC pilots for NP-hard problems such as UC and power flow
- Upgrade data collection and analytical frameworks for future quantum preparedness
- Develop internal QC literacy and protocols

Corporations / Energy Buyers

- Explore QaaS opportunities to streamline R&D and business operations
- Integrate hybrid QC in energy-intensive computational workflows
- Plan for long-term QC-driven energy shocks

The organizations that build literacy, partnerships, and quantum-ready infrastructure today will define the competitive landscape of tomorrow.

Glossary

QC — Quantum Computing
QPU — Quantum Processing Unit
CPU — Central Processing Unit
GPU — Graphics Processing Unit
HPC — High-Performance Computing
QEC — Quantum Error Correction
NISQ — Noisy Intermediate-Scale Quantum
FTQC — Fault-Tolerant Quantum Computing
QV — Quantum Volume
QAE — Quantum Amplitude Estimation
QAOA — Quantum Approximate Optimization Algorithm
QPE — Quantum Phase Estimation
QSP — Quantum Signal Processing
VQE — Variational Quantum Eigensolver
MCMC — Markov Chain Monte Carlo
MC — Monte Carlo
QUBO — Quadratic Unconstrained Binary Optimization
MPC — Model Predictive Control

UC — Unit Commitment
SDK — Software Development Kit
QaaS — Quantum-as-a-Service
DFT — Density Functional Theory
SF — Singlet Fission
SMR — Small Modular Reactor
UQ — Uncertainty Quantification
VaR — Value at Risk
CVaR — Conditional Value at Risk
TAM — Total Addressable Market
CAPEX — Capital Expenditure
OPEX — Operational Expenditure
LCOE — Levelized Cost of Energy
LLM — Large Language Model
NQI — National Quantum Initiative
QIS — Quantum Information Science
PQC — Post-Quantum Cryptography
OPF — Optimal Power Flow

Statement on AI Use

Generative AI was used in this Capstone project primarily as a research aid. We individually used Large-Language Models (LLMs) such as Claude, Perplexity, and ChatGPT for general research, for e.g. to ask for a list of websites or sources that offer information on quantum computing. All information within this presentation came directly from sources that we have read and cited within the respective slides. No text on these slides was generated by LLMs.

An LLM was only used to generate the conceptual graphs on page 19 that describe symbolically the proposed relationship between qubit count, quantum volume (computational capacity), and power requirements of a quantum computer. These images were generated using Claude (2026) - Anthropic. The scientific information being conveyed through these images is from the following source -

Martin, M. J., Hughes, C., Moreno, G., Jones, E. B., Sickinger, D., Narumanchi, S., & Grout, R. (2022). Energy use in quantum data centers: Scaling the impact of computer architecture, qubit performance, size, and thermal parameters. *IEEE Transactions on Sustainable Computing*, *7*(4), 864-874. https://ieeexplore.ieee.org/ielaam/7274860/9976243/9827605-aam.pdf

The specific chat with Claude that was used to generate and improve those images iteratively can be accessed here - https://claude.ai/share/dbbafa35-0bbe-4205-849b-e3ccdca6f57d