

Quantum on the Dispatch Desk: Credible Grid-Optimisation Use Cases

A commercial-readiness framework for deciding where quantum and hybrid optimisation can support power-system planning and operations

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

Power-system operators solve demanding optimisation problems across planning, scheduling and real-time operations. Higher renewable penetration, distributed energy resources, storage, flexible demand and large new loads are increasing the number of discrete choices, uncertain states and network constraints. Quantum computing is frequently presented as a response to this complexity. The commercial question is narrower: which grid decisions have a mathematical structure that could benefit from quantum or hybrid methods, what evidence is required before operational reliance, and how should a utility or investor govern experimentation while classical algorithms continue to improve? This paper develops a Dispatch-Desk Readiness Framework for quantum grid optimisation. It separates the decision problem, mathematical formulation, classical baseline, quantum encoding, hardware execution, feasibility recovery, power-system validation and operational integration. It evaluates eight use-case families: unit commitment, security-constrained scheduling, transmission switching, storage and microgrid planning, restoration, distributed-energy coordination, maintenance scheduling and market portfolio optimisation. The framework requires each candidate to pass tests for decision value, formulation fit, qubit and interaction burden, constraint fidelity, solution quality, runtime, repeatability, cyber and data controls, and integration with existing operator processes. The available evidence supports disciplined experimentation rather than operational claims of general quantum advantage. Recent studies demonstrate workable decompositions and hybrid algorithms on small or simplified instances. Other benchmarks find that candidate methods do not outperform classical or random baselines under the tested conditions. United States Department of Energy programmes are therefore funding benchmark development and real-grid challenge problems that compare quantum and hybrid approaches with classical methods. Three figures and three tables in this paper use hypothetical author-developed assumptions to demonstrate the evaluation mechanics. They are not forecasts, market quotations, vendor performance claims or investment recommendations.

JEL Classification: C61, L94, O32, Q40, Q48

Keywords: quantum computing, power systems, grid optimisation, unit commitment, optimal power flow, energy storage, microgrids, hybrid quantum-classical algorithms, technology readiness, utility strategy

Introduction

Electricity systems are governed by decisions made at several timescales. Planners decide where generation, storage, transmission and flexible demand should be added. System operators schedule generating units and reserves. Control rooms adjust dispatch and respond to outages. Distribution utilities coordinate increasingly numerous devices at the grid edge. Market participants optimise bids, fuel, storage and risk. Many of these tasks combine discrete choices, continuous physical variables, uncertainty and strict security constraints.

Quantum computing attracts attention because several grid decisions contain difficult combinatorial components. A unit-commitment model chooses which generators should be on or off across time. Transmission-switching and restoration models choose network configurations. Storage and microgrid planning models choose locations, capacities and operating patterns. These choices can be represented in forms that quantum annealers or gate-based algorithms may address. A mathematical mapping alone does not establish operational value. The complete problem also includes continuous power flows, nonlinear physics, uncertainty, contingency requirements, market rules and tight time limits.

The central conclusion of this paper is that a credible quantum-grid programme begins with a decision and a benchmark, not with a machine. The utility should identify a costly or constrained decision, document the incumbent classical workflow, define operational acceptance criteria and test whether a quantum component improves the complete workflow. The useful architecture is likely to be hybrid for the foreseeable evaluation horizon: classical systems prepare data, decompose the problem, solve continuous subproblems, recover feasibility and validate security; a quantum processor addresses a carefully selected discrete subproblem.

1 Define the grid decision before selecting the technology

A technology experiment should begin with a named decision owner, operating cadence and economic consequence. “Optimise the grid” is too broad. “Produce a feasible day-ahead commitment schedule within the market deadline while reducing total production cost against the current solver” is testable. So is “identify a portfolio of storage and microgrid investments that improves resilience subject to a capital ceiling and network constraints.”

The decision definition should record inputs, output, horizon, update frequency, constraints, tolerance, fallback and approval rights. A scheduling problem may require a result every five minutes, every hour or once a day. A planning problem may allow hours of computation but require many scenarios. Restoration may demand a good feasible answer within seconds. These differences matter more than the abstract complexity class.

The owner should also define the value of improvement. A lower objective value has little commercial meaning if the schedule violates reserves, cannot be reproduced, arrives after the decision deadline or increases operational risk. Value may arise from lower production cost, reduced curtailment, better reserve allocation, faster study throughput, more robust restoration plans, deferred network investment or improved resilience. Each value claim should trace to an auditable operating or planning metric.


2 Understand the optimisation stack

Grid optimisation is a stack rather than one equation. Data preparation reconciles topology, equipment status, forecasts, offers, limits and contingencies. The model represents costs and physical constraints. Presolve reduces the problem. A solver searches for candidates. Feasibility checks test network and operational limits. Operators review exceptions and apply judgement. Results then enter market, control, planning or investment systems.

Quantum computing usually addresses one part of this stack. Quantum annealing and the Quantum Approximate Optimisation Algorithm are commonly associated with binary optimisation. Continuous dispatch, AC power flow and detailed security assessment remain classical in most proposed architectures. Decomposition methods can separate binary and continuous decisions, allowing the quantum processor to propose commitment or configuration choices while a classical solver calculates dispatch and checks feasibility.

This division should be explicit in any business case. Claims about a quantum subroutine should not be presented as claims about the end-to-end process. Data-transfer latency, queue time, repeated sampling, embedding, parameter tuning and feasibility recovery can dominate the elapsed time. The benchmark must therefore measure the full workflow from accepted input to validated output.

Figure 1. The Dispatch-Desk Readiness Framework



The framework is author-developed. Each gate requires documented evidence and accountable approval before progression.

3 Establish a serious classical baseline

The baseline is the most important control in a quantum evaluation. It should represent the best method that the organisation could reasonably deploy, including commercial solvers, decomposition, heuristics, high-performance computing and domain-specific reformulations. A weak baseline can make an experimental method appear valuable. A baseline from an old academic example may bear little resemblance to the tuned system used by a market operator or utility.

The United States Department of Energy Grid Optimization Competition illustrates the level of detail in a realistic benchmark. Its formulations include AC network physics, multi-period commitment, contingencies, ramping and solution evaluation. The competition framework makes objective value, feasibility and computational performance measurable. A quantum programme can use such open formulations to develop capability, while proprietary tests should preserve the utility's actual operational distributions and constraints.

Benchmark records should include hardware, software version, preprocessing, solver settings, stopping criteria, warm starts, random seeds and elapsed time. Solution quality should be reported with a valid lower or upper bound where available. Repeated runs matter because sampling-based methods can return different results. The evaluation should also distinguish time to first feasible solution from time to best solution.

4 Select problems with a plausible formulation fit

A candidate problem needs a material discrete structure, repeated value and a tractable interface with the rest of the workflow. Binary choices such as unit on or off, line open or closed, project selected or rejected, and crew assigned or not assigned can map to quadratic unconstrained binary optimisation or Ising forms. The mapping often introduces penalty terms and auxiliary variables. Poor penalty calibration can produce low-energy solutions that violate the original constraints.

Interaction density is as important as the number of variables. A model with many pairwise relationships can require additional physical qubits or complex embedding on an annealer. A gate-based circuit can become too deep for current noise levels. Time coupling, minimum up and down times, reserves, network limits and scenario relationships can expand the encoded problem rapidly.

Decomposition offers a practical route. The classical master can coordinate scenarios or continuous variables while a quantum routine searches a binary neighbourhood. Benders decomposition, Lagrangian relaxation, alternating-direction methods and local branching can define smaller subproblems. The commercial test is whether the decomposition improves the total process after communication, sampling and recovery costs.

Table 1. Screening grid-optimisation use cases for quantum experimentation

Use-case family	Discrete core	Classical component retained	Credible near-term experiment	Principal acceptance risk
Unit commitment	On, off, start and configuration choices	Economic dispatch, network checks and reserves	Quantum search inside decomposition or neighbourhood improvement	Feasibility across time and security constraints
Transmission switching	Line-status and remedial-action choices	AC or DC power flow and contingency analysis	Candidate topology generation followed by classical validation	Network-security violation and combinatorial embedding burden
Storage and microgrid planning	Site, size band and project selection	Chronological dispatch, reliability and finance	Portfolio selection across scenario-reduced cases	Loss of chronology and weak resilience representation
Restoration	Switch, crew and sequence decisions	Electrical feasibility and state estimation	Generate feasible candidate sequences for a bounded zone	Response time and safety-critical validation
Distributed-energy coordination	Device or cluster participation choices	Local controls, forecasting and feeder power flow	Cluster-level scheduling with classical device control	Scale, privacy, communications and real-time requirements
Maintenance scheduling	Outage windows, crews and task assignment	Reliability, inventory and network adequacy	Combinatorial schedule generation with resource constraints	Incomplete asset and failure data

Use-case family	Discrete core	Classical component retained	Credible near-term experiment	Principal acceptance risk
Market portfolio optimisation	Bid blocks, commitment and hedge choices	Price scenarios, risk and settlement rules	Scenario-based portfolio selection	Model risk and changing market rules
Expansion planning	Project and corridor selection	Production-cost simulation and power flow	Candidate-plan search over a reduced portfolio	Long-horizon uncertainty and model-form error

The assessment is an author-developed screening framework. Readiness depends on the actual formulation, data and hardware available.

5 Unit commitment is the leading test case

Unit commitment is a natural research target because it contains binary generator-status decisions and continuous dispatch. It is operationally important and supported by standard test systems. Renewable variability, storage, demand response and shorter scheduling intervals increase the number of states and scenarios. Security-constrained variants add network and contingency requirements.

Published approaches include QAOA, quantum annealing, surrogate Lagrangian relaxation and hybrid decomposition. Some studies report feasible or near-optimal results on simplified instances and identify methods for reducing the required qubits. The 2025 work by Barrass, Nagarajan and Coffrin uses Benders-style cut generation and neighbourhood recovery to reduce quantum interactions and improve solution quality on a modified IEEE RTS-96 test system. Other work decomposes the binary commitment subproblem from continuous dispatch and coordinates them iteratively.

Evidence remains instance-specific. A 2026 benchmark of hybrid variational methods on small AC optimal-power-flow unit-commitment instances reported no advantage over uniform sampling for the method tested. Earlier work found classical solvers effective on simulated instances below a stated scale and identified fault-tolerant hardware as a possible requirement at larger scale. These findings support a disciplined test plan: preserve the best classical baseline, publish the complete instance, report feasibility and quality, and avoid extrapolating from a small example to an independent-system-operator workflow.

6 Security-constrained scheduling raises the evidentiary bar

Security-constrained unit commitment and economic dispatch require the schedule to remain acceptable after credible contingencies. The model may include thousands of monitored limits and corrective actions. Operators also use reserve requirements, stability constraints and remedial-action schemes. A candidate schedule that is cheap but insecure has no operational value.

Quantum experiments can focus on candidate generation while classical contingency analysis acts as the safety gate. The workflow may add violated constraints iteratively, similar to constraint generation. A quantum subroutine could search commitment or topology choices within a bounded neighbourhood. Every returned candidate would then pass an independent deterministic validation using the accepted network model.

The success measure should include how many candidates survive validation, how quickly the first secure candidate appears, the objective gap, and performance under stressed cases. The test set should cover peak demand, low inertia, renewable ramps, outages and forecast error. A result on a normal case does not establish resilience under the conditions that make security-constrained optimisation valuable.

7 Transmission switching and remedial action

Opening or closing selected transmission elements can relieve congestion, improve resilience or support restoration. The discrete topology choices interact with continuous power flow and protection requirements. Quantum optimisation may assist in searching a large set of candidate configurations, especially when the allowed switching set is limited and every candidate is subsequently checked with established power-flow and contingency tools.

The model must represent switching restrictions, transient concerns, protection settings, voltage limits and operator procedures. A DC approximation can be useful for screening while remaining insufficient for final approval. An implementation should therefore use a staged process: generate candidates, run AC power flow, conduct contingency and stability checks, apply operating rules, and present the feasible set with explanations.

Commercial value depends on congestion cost, avoided curtailment, restoration time or deferred investment. The business case should use observed congestion patterns and realistic switching opportunities. It should also price the engineering and control-system changes needed to execute a recommendation. A computational improvement without an executable field action has limited value.

8 Storage and microgrid planning

Storage and microgrid planning combines project selection, capacity, location and operating strategy. The decision is influenced by load growth, renewable output, outages, market prices, network limits and resilience objectives. The United States Department of Energy's 2026 Quantum Grid Planning Competition used storage and microgrid siting and sizing as a challenge problem. The programme required quantum and hybrid approaches to be compared with classical methods, creating a useful model for evidence-led evaluation.

The experiment should preserve chronology where storage energy balance matters. Representative days can reduce size, yet poor treatment of transitions can overstate value. Resilience analysis should represent the timing and duration of outages, critical loads, islanding constraints, fuel and restoration. A portfolio selected on simplified scenarios must be validated through detailed chronological simulation.

Quantum search may be most useful at the portfolio-selection layer. A classical model can calculate project attributes or operating outcomes across scenarios. The quantum component selects combinations subject to budget, geography, network and policy constraints. This separation creates a clear interface and makes it easier to compare the candidate portfolio with classical integer programming and heuristics.

9 Grid restoration and emergency planning

Restoration involves switch sequences, energisation paths, black-start resources, crews, communications and safety constraints. The problem can be combinatorial and time-sensitive. A quantum programme should begin in offline planning, where many credible outage states can be analysed and candidate playbooks prepared. Operational use requires a substantially higher standard of validation, availability and human control.

The model should respect electrical feasibility at every step rather than only at the final state. It should include damaged equipment, uncertain status, communication loss, crew travel and resource constraints where relevant. Candidate sequences should be tested in the utility's network-analysis and training environments. Operators should understand why a sequence was proposed and retain authority to reject or modify it.

The first commercial metric may be study throughput rather than control-room automation. If a hybrid method allows planners to evaluate more outage combinations and produce more robust playbooks within the same planning cycle, it can create value before any online deployment. This pathway reduces the risk of relying on immature hardware during an emergency.

10 Distributed-energy and virtual-power-plant coordination

Distribution systems can contain very large numbers of solar inverters, batteries, electric vehicles, flexible loads and other controllable devices. Centralised optimisation can become burdensome, while communications, ownership and privacy limit the available data. NREL has demonstrated hierarchical and distributed classical control methods for more than one million distributed resources. That achievement is a demanding baseline for any quantum proposal.

A plausible hybrid design aggregates devices into clusters. Local classical controllers preserve device constraints and privacy. A higher-level optimiser allocates targets or market participation across clusters. A quantum component could search discrete participation, commitment or routing choices for a subset of the coordination problem. Feeder power flow and device feasibility remain classical validation gates.

The evaluation should measure communications, failed-device behaviour, forecast error and rebound effects. A schedule that performs well in simulation may fail when devices respond differently or communications are delayed. The fallback should return control to a proven local policy. Cybersecurity and customer-consent requirements are part of the operating design, not a later compliance exercise.

11 Maintenance and outage scheduling

Utilities and generators schedule maintenance across assets, crews, contractors, spares and permitted outage windows. The combinatorial structure can fit quantum experimentation, particularly when tasks have precedence, mutual exclusion or resource constraints. The objective can combine reliability, cost, lost production and network adequacy.

The quality of asset data determines the value of optimisation. Failure probabilities, task durations and resource availability are often uncertain. A sophisticated solver cannot repair an incomplete work-order history or an unrecorded spare-parts constraint. The pilot should begin with a bounded fleet where task definitions, durations and dependencies can be audited.

Candidate schedules should be tested through reliability and production simulations. The commercial measure can include expected unserved energy, lost margin, overtime, contractor mobilisation and schedule stability. Operators may prefer a slightly more expensive schedule that remains feasible when a task overruns. Robustness and explainability therefore belong in the acceptance criteria.

12 Market bidding and portfolio decisions

Generators, storage owners, retailers and traders make discrete and continuous decisions across bids, commitments, fuel, contracts and risk limits. Quantum optimisation may be tested for portfolio selection or scenario-based bidding when the organisation can define a bounded decision set and validate settlement consequences.

Market models change and price distributions are unstable. Backtests should use walk-forward periods, include transaction and imbalance costs, and avoid information unavailable at the decision time. A favourable historical result should be separated from evidence that the optimisation method caused the improvement. The method should be compared with mixed-integer optimisation, stochastic programming, robust optimisation and practical heuristics.

Governance should mirror other material trading models. Independent validation, access controls, versioning, limits, monitoring and kill switches are required. A quantum vendor should not receive sensitive bids or network information without appropriate security, contractual and jurisdictional controls.

13 Encode constraints without losing the original problem

Quantum algorithms often require a quadratic unconstrained binary form. Grid problems are constrained and frequently include continuous variables. Converting them can add binary encodings, slack variables and penalty coefficients. The encoded objective may behave differently from the original model if penalties are too weak, too strong or poorly scaled.

The team should maintain a formulation ledger. It records each original variable and constraint, its encoded representation, scaling, approximation and validation test. It should show which physics were simplified, which limits are hard or soft, and how infeasible samples are treated. This ledger enables independent review and prevents a low encoded energy from being mistaken for a good grid decision.

Feasibility recovery is part of the algorithm. A classical repair step may adjust commitment, dispatch or topology. The benchmark should report quality before and after repair, the share of samples that require repair, and the computational cost. If the repair routine creates most of the value, the organisation should test the routine with classical candidate generators as well.

14 Measure hardware and workflow economics

A commercial assessment includes more than solver runtime. It should measure data preparation, network transfer, queue time, compilation or embedding, repeated shots, calibration, post-processing, feasibility checks and operator review. Cloud access may introduce availability, sovereignty and security considerations. Dedicated hardware may require utilisation sufficient to justify capital and specialist staff.

Energy consumption claims require careful boundaries. The power used by a quantum processing unit is only part of the system. Cooling, control electronics, classical orchestration, network and repeated experiments contribute to total consumption. Comparisons should use the same problem, solution-quality target and accounting boundary.

Pricing can be uncertain while the technology matures. The business case should use scenario ranges for access cost, support, integration, internal capability and switching. It should also assign value to learning only when the programme produces reusable formulations, benchmarks, data pipelines or staff capability. General exposure to a new technology is too vague to support a large commitment.

Table 2. Evidence required at each commercial-readiness gate

Gate	Required evidence	Minimum decision	Reasons to stop or redesign
Problem qualification	Named owner, baseline cost, cadence, constraints and value metric	Approve a bounded experiment	No material pain point or no measurable outcome
Formulation review	Original and encoded models, approximation ledger and test instances	Approve benchmark implementation	Critical physics or constraints cannot be represented or recovered
Baseline validation	Tuned classical solver, reproducible environment and accepted metrics	Freeze comparison protocol	Weak baseline, undisclosed tuning or incomparable hardware
Technical benchmark	Repeated quality, feasibility, runtime and scaling results	Continue to operational pilot	No reproducible improvement on relevant instances
Operational validation	Security checks, integration, fallback, operator review and stressed cases	Permit shadow operation	Unacceptable failure modes, latency or control weakness
Commercial case	Total cost, benefit attribution, vendor risk and exit plan	Contract, defer or stop	Benefits depend on unsupported scale or price assumptions

The gates are author-developed. Thresholds should be set by the utility or investor for the specific decision.

15 Build a benchmark portfolio rather than one showcase

A single instance can favour a method by chance or design. The benchmark portfolio should include standard public cases, synthetic scaled cases and proprietary cases drawn from the intended operating distribution. It should contain normal, stressed and pathological conditions. Instance generation and selection should be documented before results are examined.

Metrics should include objective gap, feasibility, constraint violation, runtime, time to first feasible result, success probability, variance, memory, interactions with the quantum service and total cost. Scaling should be tested across variables, interactions, time periods, scenarios and contingencies. Results should show where performance degrades.

The test protocol should reserve hold-out instances. Algorithm and penalty tuning on the full portfolio can overfit the benchmark. Independent replication by an internal validation team or external research partner improves confidence. Source code, seeds and configuration should be preserved to the extent permitted by security and intellectual-property constraints.


16 Use technology-readiness levels tied to operating evidence

Technology-readiness labels can conceal important differences. A method may run on real quantum hardware while remaining far from a utility pilot. The readiness ladder should therefore be tied to the decision workflow.

Level one demonstrates the mathematical mapping on a toy case. Level two reproduces results on standard test systems. Level three compares with strong classical baselines across a portfolio. Level four runs on relevant hardware and preserves original constraints through validation. Level five operates in shadow mode using real data and deadlines. Level six supports an operator decision under controlled conditions with fallback. Higher levels require sustained performance, cybersecurity, availability, auditability and change control.

Movement between levels should require evidence. A larger qubit count or a vendor roadmap does not advance an application automatically. Improvements in error correction, connectivity, fidelity or algorithm design matter only when they change the end-to-end benchmark for the chosen problem.

Figure 2. Hypothetical readiness scores across candidate use cases



Scores are hypothetical author-developed assumptions for illustrating the framework. They are not assessments of any vendor, utility or quantum platform.

17 Design the pilot around shadow operation

Shadow operation uses live or representative data without controlling the grid or market. The quantum and classical workflows run in parallel. Results are compared after independent feasibility and security checks. Operators can review candidate decisions and record whether they would have accepted them.

The pilot should specify duration, cases, service levels and exit criteria. It should include hardware unavailability, network failure, stale data and unusual system states. The organisation should test the fallback by design rather than assume it works. Model and hardware changes should be frozen or versioned during comparison periods.

Operator feedback should be structured. A candidate may be mathematically feasible but operationally unattractive because it creates excessive switching, unstable schedules or hard-to-explain changes. Acceptance metrics can include schedule churn, intervention count and reason codes. This evidence helps refine the objective and constraints.

18 Protect critical-infrastructure data and operations

Grid models can reveal topology, asset limits, operating practices and vulnerabilities. Quantum-cloud experimentation can involve external processing and multiple software layers. The utility should classify data, minimise transfers, segment environments and apply access, logging, encryption and retention controls. Contract terms should address location, subcontractors, incident notification, model use and deletion.

Operational technology should remain isolated from experimental environments. A pilot can use controlled extracts or synthetic data until security approval permits broader access. Software supply-chain review should cover classical orchestration, quantum libraries and dependencies. Reproducible builds and signed artefacts support audit and incident response.

Post-quantum cryptography is a separate issue from quantum optimisation. Utilities may need to migrate vulnerable cryptographic systems as fault-tolerant quantum capabilities develop. That programme should not be conflated with testing quantum solvers. Each has different risk, timing, ownership and standards.

19 Manage vendor and platform dependence

Quantum platforms differ in architecture, connectivity, noise, programming model and commercial access. A formulation optimised for one platform may not transfer easily. The organisation should preserve a platform-neutral problem definition and benchmark harness, then maintain adapters for specific hardware or simulators.

Contracts should address data, code, results, portability, service availability, security, intellectual property and termination. Vendor performance claims should be reproduced on the buyer's test portfolio where possible. Roadmap milestones can inform options, while procurement decisions should rely on demonstrated current capability and clearly priced commitments.

A multi-platform benchmark can reveal whether an observed improvement comes from the algorithm, hardware or surrounding classical stack. It can also reduce lock-in. The additional engineering cost should be justified by the strategic value of optionality.

20 Link the experiment to investment and transaction decisions


Utilities and investors may encounter quantum-grid capability through internal development, vendor procurement, venture investment, joint research or acquisition. Each route requires a different diligence lens. A procurement decision focuses on service levels, integration and cost. A venture investment also considers intellectual property, talent, funding needs, competitive differentiation and route to market. An acquisition requires technical, commercial, cyber and people diligence.

Revenue projections should distinguish research contracts, pilots and repeatable production deployments. A large announced partnership does not necessarily establish recurring product demand. Backlog should be tested for cancellation, milestones and gross margin. Customer concentration and dependence on public research funding may be material.

Valuation should use scenario-based milestones. Near-term value can attach to classical optimisation capability, domain data, benchmark infrastructure and customer relationships. Quantum-specific upside can be linked to demonstrated performance thresholds. This structure avoids assigning full value to an unverified future advantage while recognising assets that create value today.

21 Construct a stage-gated economics case

Figure 3. Hypothetical annual value versus total programme cost



All values are hypothetical author-developed assumptions for analytical illustration. They are not vendor prices, savings forecasts or client results.

The hypothetical figure illustrates an evidence rule rather than a forecast. A research prototype may have learning value while its direct operating benefit is small. A benchmark programme creates reusable test infrastructure.

Shadow operation can produce credible evidence about workflow performance. Controlled production can support a larger value case only after security, reliability and operator acceptance are demonstrated.

Every gate should cap spend and define the evidence required for the next commitment. Sunk research cost should not justify progression. The programme can stop successfully after establishing that classical methods remain superior for the intended use. That finding protects capital and preserves the benchmark for future hardware or algorithm improvements.

22 Allocate decision rights and independent challenge

The programme needs accountable owners across grid operations, planning, markets, optimisation, quantum technology, cybersecurity, procurement, legal and model risk. The business owner defines the decision and value. Power-system specialists own physical validation. The quantum team owns implementation evidence. An independent team challenges the benchmark and claims.

Decision rights should cover problem selection, data release, model approval, pilot progression, production use and incident response. A vendor should not approve its own benchmark. Operators should retain authority over operational adoption. Material model changes should trigger revalidation.

The steering committee should receive a concise evidence pack: baseline, instances, metrics, results, exceptions, cost, security status, operator feedback and next-gate request. Demonstrations and qubit counts can supplement the pack; they do not replace it.

Table 3. Ninety-day quantum grid-optimisation evaluation plan

Period	Core work	Evidence produced	Gate decision
Days 1 to 15	Name decision owner, select one use case, map workflow and define value	Decision charter, risk classification and baseline protocol	Confirm material problem and bounded scope
Days 16 to 30	Freeze public and proprietary instances; reproduce incumbent solver	Reproducible classical baseline and accepted metrics	Approve quantum formulation work
Days 31 to 50	Build encoded model, decomposition and feasibility recovery	Formulation ledger, unit tests and simulated results	Confirm constraint fidelity and test readiness
Days 51 to 70	Run repeated hardware and simulator benchmarks	Quality, runtime, scaling, variance and cost results	Select, redesign or stop candidate method
Days 71 to 85	Integrate with validation tools and conduct stressed cases	End-to-end results, cyber review, fallback and operator assessment	Approve or reject shadow pilot
Days 86 to 90	Review economics, vendor terms and capability plan	Investment memo, stage budget, milestones and exit plan	Fund next gate, defer or close programme

The plan is an author-developed example. Timing and activities should be adapted to the organisation, use case and procurement requirements.

23 Identify credible success conditions

A pilot can be considered technically credible when it uses relevant instances, a strong reproducible baseline, identical acceptance criteria and repeated end-to-end measurements. The output must survive independent power-system validation. The result should remain within the decision deadline and meet predefined repeatability and availability thresholds.

Commercial credibility requires attributable value after integration, support and risk costs. The use case should have an owner willing to adopt it, a secure operating model and a fallback. Contractual terms should preserve data rights and exit options. The organisation should be able to state which evidence would cause it to stop.

Strategic credibility comes from reusable assets. A good programme produces a clean problem definition, benchmark set, data pipeline, validation harness, skilled team and vendor knowledge. These assets allow future methods to be tested quickly. They can create value even when the first quantum approach does not proceed.

24 Recognise failure modes early

The most common failure is starting with hardware and searching for a problem. The result is a showcase with no accountable decision. A second failure is comparing with a weak or untuned baseline. A third is reporting encoded objective value without original-problem feasibility. Others include using only small favourable instances, excluding queue and post-processing time, extrapolating from simulation to hardware, and relying on vendor roadmaps as evidence.

Operational failure can arise from insecure data transfers, unavailable services, unexplained schedule changes or poor integration. Commercial failure can arise from expensive specialist dependence, platform lock-in or benefits too small to justify process change. Governance failure occurs when the innovation team owns both the experiment and its approval.

These failure modes are manageable through the framework in this paper. A named decision, frozen benchmark, formulation ledger, independent validation, shadow operation and stage-gated economics make claims testable.

25 Implications for utilities, investors and advisers

Utilities should treat quantum optimisation as an applied research and capability programme tied to hard decisions. The best starting points are repeated problems with a material combinatorial core, strong data and clear classical baselines. Unit commitment and planning are credible research domains. Production adoption should follow evidence from the complete workflow.

Investors should separate platform potential from application evidence. Technical diligence should examine encoded problem scale, interaction density, sampling, error, classical assistance, feasibility recovery and benchmark quality. Commercial diligence should test recurring demand, customer conversion, funding dependence, talent concentration and integration burden.

Advisers can create value by translating between engineering, optimisation, technology and finance. The mandate is to define the decision, structure the benchmark, challenge claims, quantify value, design governance and connect capability to procurement, partnership, investment or transaction structure. The quality of this translation determines whether a quantum initiative becomes a durable option or a costly demonstration.

Conclusion

Quantum computing may contribute to power-system optimisation where a material discrete problem can be isolated, encoded and recombined with classical physics, validation and operations. Unit commitment, topology selection, storage and microgrid planning, restoration, distributed-resource coordination, maintenance scheduling and portfolio optimisation provide credible experiment domains. Their readiness varies by formulation, hardware, data and operating requirement.

Current evidence supports benchmark-led experimentation. Hybrid methods have produced useful research results on bounded problems, while other studies show no advantage for the tested methods. DOE and ARPA-E programmes are investing in rigorous comparisons and real energy use cases. This is the appropriate commercial posture: measure capability against relevant classical baselines and preserve the conditions under which a result is valid.

The Dispatch-Desk Readiness Framework makes the evaluation accountable. It begins with the decision, freezes the baseline, controls the formulation, measures the entire workflow, validates physical feasibility, protects critical-infrastructure data and requires a stage-gated economics case. A utility can then invest, defer or stop on the basis of evidence. That discipline creates strategic value regardless of whether the first quantum algorithm reaches the control room.

Frequently Asked Questions

Which power-system problem is the strongest first candidate for quantum optimisation?

Unit commitment is a leading research candidate because it contains material binary decisions, has established formulations and test systems, and can be decomposed from continuous dispatch. A particular utility should still select the problem with the clearest value, best data and strongest baseline.

Has quantum advantage been demonstrated for operational grid optimisation?

The cited evidence does not establish general operational quantum advantage. Published studies report promising methods on bounded or simplified cases, while other benchmarks report no improvement for the tested method. Claims should be limited to the specific instances, hardware, workflow and baseline evaluated.

Why are hybrid quantum-classical methods important?

Grid problems combine binary choices, continuous variables, nonlinear physics and security constraints. A hybrid method can assign a bounded discrete search to a quantum processor while classical systems handle decomposition, dispatch, power flow, feasibility recovery and validation.

What should a utility measure in a benchmark?

It should measure end-to-end runtime, time to first feasible answer, objective quality, original-problem feasibility, constraint violations, repeatability, scaling, service availability, data-transfer burden and total cost against a tuned classical baseline.

Can a quantum result be used directly by a control room?

Operational use requires independent feasibility and security validation, tested integration, cybersecurity controls, service availability, operator acceptance and a proven fallback. Shadow operation is an appropriate intermediate stage.

How should a buyer value a quantum-grid company?

The buyer should value current assets such as classical software, domain data, customer contracts, benchmark infrastructure and talent separately from quantum-specific upside. Future upside can be linked to demonstrable technical and commercial milestones.

Does quantum optimisation address post-quantum cybersecurity?

They are separate programmes. Quantum optimisation explores computation for decisions. Post-quantum cybersecurity replaces cryptography that may become vulnerable to fault-tolerant quantum computers. They require different owners, standards and investment cases.

When should a quantum grid-optimisation programme stop?

It should stop or redesign when the problem lacks material value, the encoded model loses critical constraints, the method fails a relevant classical benchmark, integration or cyber risk is unacceptable, or the evidence does not justify the next stage of spending.

References

1. United States Department of Energy, Office of Electricity, *Grid Modeling*. <https://www.energy.gov/oe/grid-modeling>
2. United States Department of Energy, *Grid Modernization Multi-Year Program Plan*, 2025. <https://www.energy.gov/downloads/grid-modernization-multi-year-program-plan-myp>
3. United States Department of Energy, *Grid Modernization Initiative Lab Call 2023*, 2023. https://www.energy.gov/sites/default/files/2023-05/2023-gmi-lab-call_0523.pdf
4. United States Department of Energy, Office of Technology Commercialization, *DOE Announces Winning Teams in Quantum Grid Planning Competition*, 2026. <https://www.energy.gov/technologycommercialization/articles/doe-announces-winning-teams-quantum-grid-planning-competition>
5. Advanced Research Projects Agency-Energy, *Enhancing Neutral-atom Computers for Optimizing Delivery of Energy*, 2026. <https://arpa-e.energy.gov/programs-and-initiatives/search-all-projects/enhancing-neutral-atom-computers-optimizing-delivery-energy-encode>
6. United States Department of Energy, Office of Electricity, *DOE Announces GridFM 2.0 Project*, 2026. <https://www.energy.gov/oe/articles/does-office-electricity-announces-115m-genesis-mission-project-meet-growing-electricity>
7. Grid Optimization Competition, *Challenge 3 Problem Formulation*, 2024. https://gocompetition.energy.gov/sites/default/files/Challenge3_Problem_Formulation_20240122.pdf
8. Grid Optimization Competition, *Challenge 2 Problem Formulation*, 2021. https://gocompetition.energy.gov/sites/default/files/Challenge2_Problem_Formulation_20210531.pdf
9. National Renewable Energy Laboratory, *Planning for Reliable Operations*. <https://www.nrel.gov/grid/planning-for-reliable-operations.html>
10. National Renewable Energy Laboratory, *Large-Scale Experiments Demonstrate Advanced Controls for Autonomous Energy Systems*, 2021. <https://www.nrel.gov/news/program/2021/large-scale-experiments-demonstrate-advanced-controls-for-autonomous-energy-systems.html>
11. National Renewable Energy Laboratory, *OptGrid: Optimizing Distributed Energy Resources for Stability, Resilience, and Efficiency*. <https://www.nrel.gov/news/video/optgrid-optimizing-distributed-energy-resources-for-stability-resilience-and-efficiency-text.html>
12. National Renewable Energy Laboratory, *REopt: Energy Integration and Optimisation*. <https://www.nrel.gov/reopt/>
13. Watson, J.-P., *Rigorous Assessment and Development of Quantum and Hybrid Quantum-Classical Computing Approaches for Solving the Unit Commitment Problem*, Lawrence Livermore National Laboratory, 2024. <https://www.osti.gov/servlets/purl/2281795>
14. Barrass, R., Nagarajan, H. and Coffrin, C., *Leveraging Quantum Computing for Accelerated Classical Algorithms in Power Systems Optimization*, CPAIOR, 2025. <https://www.osti.gov/servlets/purl/3009672>
15. Eskandarpour, R. et al., *Quantum Computing for Enhancing Grid Security*, IEEE, 2020. <https://www.osti.gov/servlets/purl/1674938>
16. Zhou, Z. et al., *Quantum Computing in Power Systems*, iEnergy, 2022. <https://doi.org/10.23919/IEN.2022.0021>
17. Feng, C. et al., *Novel Resolution of Unit Commitment Problems Through Quantum Surrogate Lagrangian Relaxation*, IEEE Transactions on Power Systems, 2023. <https://doi.org/10.1109/TPWRS.2022.3181221>
18. Mahroo, R. and Kargarian, A., *Hybrid Quantum-Classical Unit Commitment*, 2022. <https://arxiv.org/abs/2201.03701>
19. Salgado, B., Sequeira, A. and Santos, L.P., *A Hybrid Classical-Quantum Approach to Highly Constrained Unit Commitment Problems*, 2024. <https://arxiv.org/abs/2412.11312>
20. Gaidai, I. and Mukherjee, R., *Benchmarking Hybrid Quantum-Classical Algorithms for Power Grid Optimization Problems*, 2026. <https://arxiv.org/abs/2607.15543>
21. Carrillo-Muñoz, M. et al., *Towards Scalable Quantum Unit Commitment in Power Systems: Encoding and Complexity Reduction Techniques*, 2026. <https://arxiv.org/abs/2609.36999>
22. Ajagekar, A. and You, F., *Quantum Computing for Energy Systems Optimization: Challenges and Opportunities*, Energy, 2019. <https://doi.org/10.1016/j.energy.2019.04.186>
23. Nikmehr, N. and Zhang, P., *Quantum Approximation for Unit Commitment*, IEEE, 2021. <https://doi.org/10.1109/QCE52317.2021.00035>
24. Padhy, N.P., *Unit Commitment: A Bibliographical Survey*, IEEE Transactions on Power Systems, 2004. <https://doi.org/10.1109/TPWRS.2003.821611>
25. Morales-España, G., Ramos, A. and García-González, J., *An MIP Formulation for Joint Market-Clearing of Energy and Reserves Based on Ramp Scheduling*, IEEE Transactions on Power Systems, 2014. <https://doi.org/10.1109/TPWRS.2013.2259601>
26. Morales-España, G., Latorre, J.M. and Ramos, A., *Tight and Compact MILP Formulation for the Thermal Unit Commitment Problem*, IEEE Transactions on Power Systems, 2013. <https://doi.org/10.1109/TPWRS.2012.2222938>
27. Knueven, B., Ostrowski, J. and Watson, J.-P., *On Mixed-Integer Programming Formulations for the Unit Commitment Problem*, INFORMS Journal on Computing, 2020. <https://doi.org/10.1287/ijoc.2019.0944>
28. Federal Energy Regulatory Commission, *Energy Primer: A Handbook of Energy Market Basics*, 2020. <https://www.ferc.gov/energy-primer>
29. North American Electric Reliability Corporation, *Reliability Standards for the Bulk Electric Systems of North America*. <https://www.nerc.com/pa/stand/Pages/ReliabilityStandards.aspx>
30. North American Electric Reliability Corporation, *Technical Reference Document: Considerations for Performing an Energy Reliability Assessment*, 2024. https://www.nerc.com/comm/RSTC_Reliability_Guidelines/Technical%20Reference%20Document%20Considerations%20for%20Performing%20an%20ERA%20V2.pdf
31. International Energy Agency, *Electricity Grids and Secure Energy Transitions*, 2023. <https://www.iea.org/reports/electricity-grids-and-secure-energy-transitions>
32. International Energy Agency, *Modernising Grids in the Age of Electricity*, 2026. <https://www.iea.org/reports/modernising-grids-in-the-age-of-electricity>
33. International Renewable Energy Agency, *Innovation Landscape Brief: Market Integration of Distributed Energy Resources*, 2019. <https://www.irena.org/publications>
34. International Renewable Energy Agency, *Electricity Storage Valuation Framework*, 2020. <https://www.irena.org/publications/2020/mar/electricity-storage-valuation-framework-2020>
35. National Institute of Standards and Technology, *Cybersecurity Framework 2.0*, 2024. <https://www.nist.gov/cyberframework>
36. National Institute of Standards and Technology, *Post-Quantum Cryptography Standards*, 2024. <https://www.nist.gov/pqcrypto>
37. Cybersecurity and Infrastructure Security Agency, *Cybersecurity Performance Goals*. <https://www.cisa.gov/cybersecurity-performance-goals>

38. Cybersecurity and Infrastructure Security Agency, *Principles of Operational Technology Cybersecurity*, 2024. <https://www.cisa.gov/news-events/alerts/2024/10/01/asds-acsc-cisa-fbi-nsa-and-international-partners-release-guidance-principles-ot-cybersecurity>
39. European Union Agency for Cybersecurity, *Cybersecurity for Smart Infrastructures*. <https://www.enisa.europa.eu/topics/critical-information-infrastructures-and-services>
40. International Organization for Standardization, *ISO 31000 Risk Management*. <https://www.iso.org/iso-31000-risk-management.html>
41. Basel Committee on Banking Supervision, *Principles for Operational Resilience*, 2021. <https://www.bis.org/bcbs/publ/d516.htm>
42. IFRS Foundation, *IAS 38 Intangible Assets*. <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
43. IFRS Foundation, *IFRS 3 Business Combinations*. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
44. IFRS Foundation, *IFRS 13 Fair Value Measurement*. <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>

About the Author

Authored by Chennakeshav Adya

Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.