Potential for Quantum Computing Applied to Power System Applications

ESIG Advanced Grid Solutions Workshop
Session 4A: What is the Next Set of Advanced Grid Solutions?

Hongyu Wu | Kansas State University
Wednesday, June 17, 2026 • Crystal B/C

The practical question is workflow value — not quantum hype

Grid studies are becoming more repeated, coupled, and scenario-driven.

- Planning and operations increasingly evaluate many futures, not one base case.
- Classical solvers remain the backbone, but repeated high-fidelity runs can become the bottleneck.
- The question is where a small quantum device can improve a much larger classical workflow.

Pressure points

More uncertainty

weather, large-load growth, DERs, inverter behavior, contingencies

More discrete choices

commitment, switching, restoration, protection settings, corrective actions

More evidence needed

trustworthy feasibility, reliability, cost, and runtime comparisons

Bottom line: quantum is interesting only if it improves end-to-end decision quality.

What is quantum computing? A mechanism view

1 Encode amplitudes

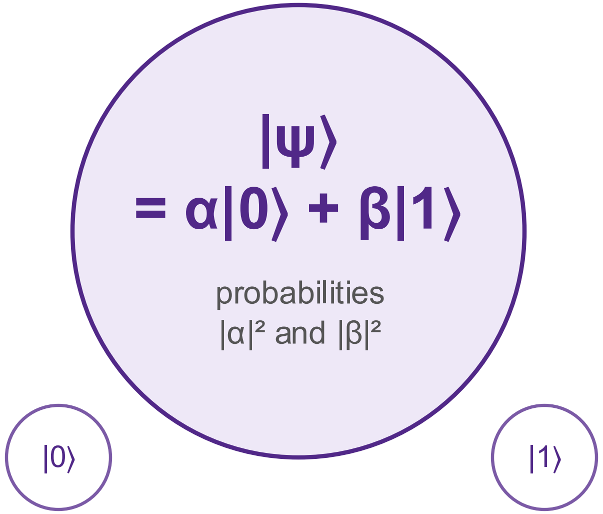

$|\psi\rangle$

$= \alpha|0\rangle + \beta|1\rangle$

probabilities

$|\alpha|^2$ and $|\beta|^2$

$|0\rangle$

$|1\rangle$

A qubit carries amplitudes, not a fixed zero-or-one value.

2 Evolve with gates

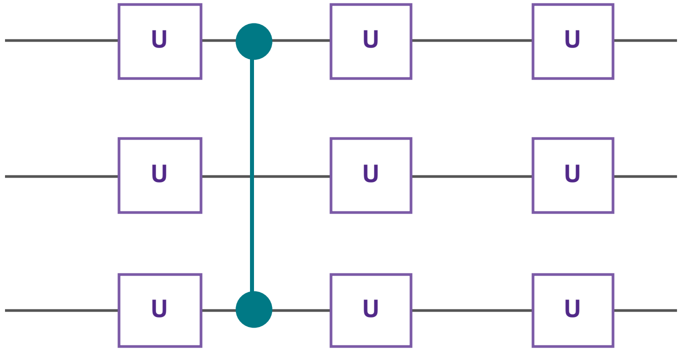

Gates rotate amplitudes; entanglement couples variables; interference changes which outcomes are likely.

3 Measure outcomes

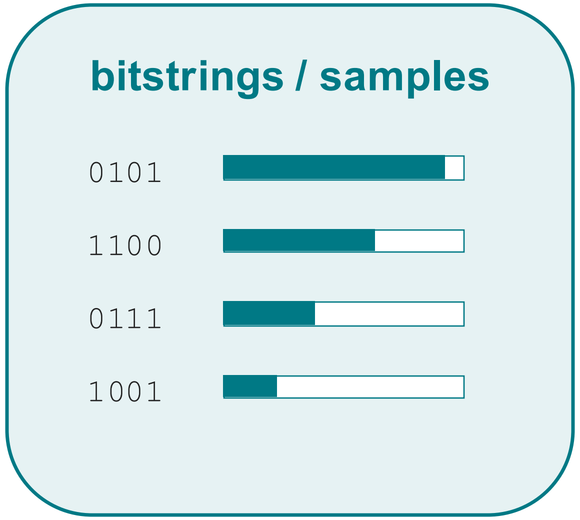

Many shots estimate a distribution or an expectation value.

Mechanism takeaway

Quantum does not return a certified grid decision by itself. It returns samples or expectations that a classical workflow must decode, repair, and verify.

State of the art: quantum-for-grid is broad, but still early

What has been explored / demonstrated

1	**Operation & control**	unit commitment, switching, restoration, contingency-corrective actions	QUBO/Ising, QAOA, annealing
2	**Steady-state security**	power flow, optimal power flow, contingency analysis	quantum linear-algebra and combinatorial reformulations
3	**Monitoring & dynamics**	state estimation, PMU/SCADA fusion, EMT/DAE, transient and voltage stability	quantum ML and circuit-based simulators
4	**Risk & uncertainty**	probabilistic power flow, rare-event and scenario-heavy studies	amplitude estimation, quantum Monte Carlo, stochastic workflows

What remains missing

1. Kernel demos are not enough

The full workflow includes encoding, data loading, measurement, decoding, repair, and validation.

2. Grid semantics must survive

Limits, reserves, contingencies, uncertainty, and time coupling must remain auditable.

3. Evidence must be operator-grade

Benchmarks need strong classical baselines, realistic scenarios, and metrics that operators recognize.

Where quantum could plausibly enter grid applications

1 **Search over discrete decisions**

Unit commitment, switching, restoration sequencing, corrective actions

2 **Sampling and risk exploration**

Scenario selection, rare events, probabilistic assessments, uncertainty quantification

3 **Linear algebra and simulation kernels**

Longer-term paths for state estimation, power flow, stability, and EMT-style models

Near-term recipe

- Do not send the full grid problem to a quantum computer.

- Instead, isolate a structured subproblem, use quantum resources to generate candidate decisions, and let classical grid tools enforce physics and operational constraints.

That is where today's hardware limitations and operator requirements can meet.

A concrete stress test: unit commitment

What the operator needs

Choose which generators are on or off over a planning horizon, then dispatch them to serve load at least cost while respecting ramping, limits, and reserves.

Why it is a good quantum test case

- Large binary search space with physics constraints.
- Strong classical baselines exist, so comparisons are meaningful.
- A candidate schedule can be checked by a classical dispatch solver.

Commitment schedule

	t1	t2	t3	t4
G1	ON	ON	OFF	ON
G2	OFF	ON	ON	ON
G3	ON	OFF	ON	ON

↓

Dispatch and feasibility

MW output + load balance
ramp limits + reserve margins
minimum / maximum generation

The quantum part only proposes. The grid math still has to verify.

KANSAS STATE
UNIVERSITY

Example prototype: quantum proposes, classical dispatch validates

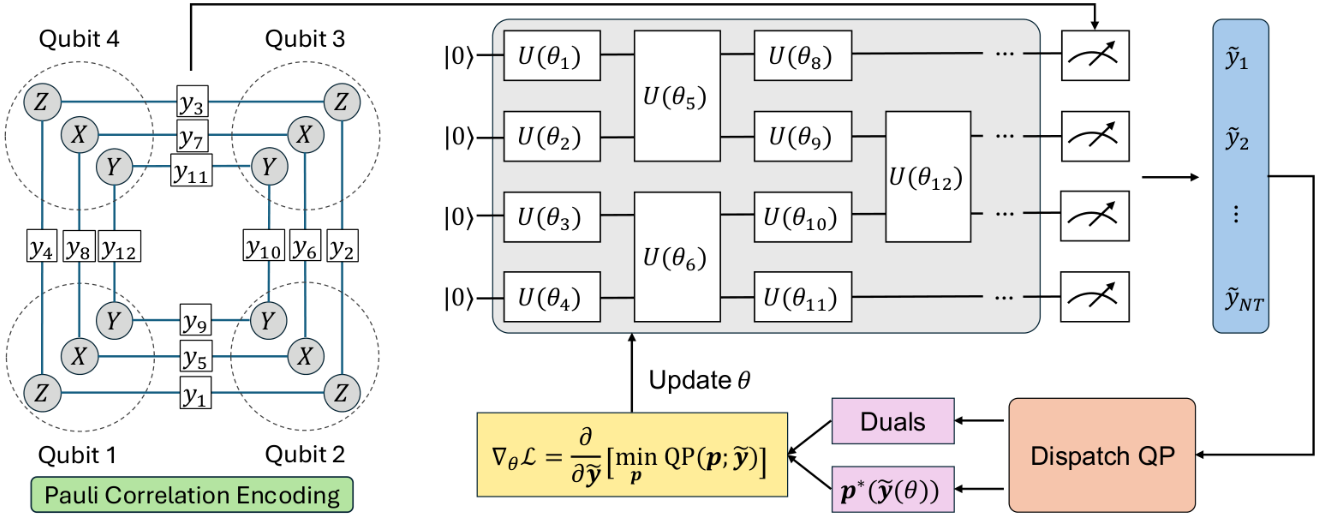

Hybrid leader–follower loop

- Quantum circuit encodes many commitment variables using Pauli correlations.
- Classical QP computes dispatch and dual sensitivities.
- Training updates the circuit, then thresholding produces a hard schedule.

Industry interpretation

This is not a standalone quantum UC solver. It is a candidate-generation layer wrapped by classical feasibility and cost evaluation.

K. X. Nguyen, I. Safro, and X. Liu, “Scaling quantum optimization for unit commitment via Pauli correlation encoding,” *arXiv preprint arXiv:2605.17145*, 2026. doi: 10.48550/arXiv.2605.17145.

Evidence so far: promising compression with clear caveats

A larger-scale UC test case

312 → 15

commitment variables → qubits with PCE

0%

constraint violations; 13.61% cost gap in one UC 26a run, as compared with CPLEX solution

16.7%

CPLEX warm-start speedup on UC 26a

What the result means

- Compression is the headline: the quantum layer can represent a much larger logical schedule.
- Cost gaps remain material in benchmark runs, but still behind CPLEX benchmark.
- Feasibility is not yet consistent, especially under tighter ramp and reserve constraints.

Honest takeaway

The near-term value is likely better warm starts, candidate schedules, and decomposition support — not a black-box quantum advantage claim.

Caveat: operator acceptance requires repeatable feasibility and end-to-end accounting of all overheads.

From prototype to practice: an evaluation framework

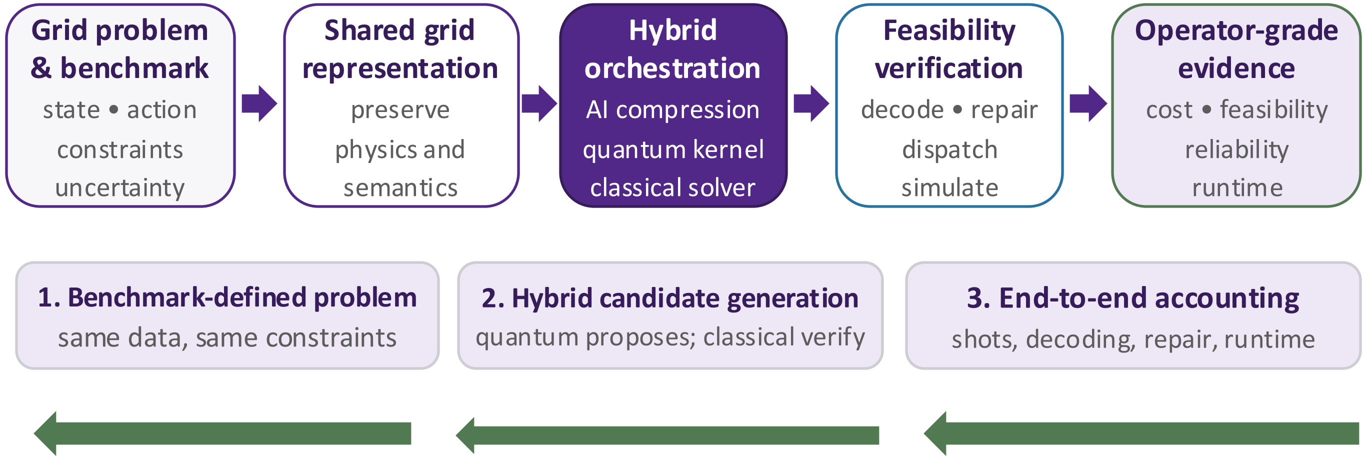

Grid problem & benchmark
state • action
constraints
uncertainty

Shared grid representation
preserve
physics and
semantics

Hybrid orchestration
AI compression
quantum kernel
classical solver

Feasibility verification
decode • repair
dispatch
simulate

Operator-grade evidence
cost • feasibility
reliability
runtime

1. Benchmark-defined problem
same data, same constraints

2. Hybrid candidate generation
quantum proposes; classical verify

3. End-to-end accounting
shots, decoding, repair, runtime

Feedback: evidence and failure modes refine representation, decomposition, and validation metrics

What has to be built around the algorithm

Shared grid representation

Preserve physics, uncertainty, and feasibility semantics across solvers.

Hybrid orchestration

AI compression, classical solvers, and quantum kernels each do the part they are suited for.

Operator-grade validation

Benchmark against strong classical baselines with partner-defined scenarios and metrics.

Barriers to bring quantum research into industry practice

1. Preserve operational meaning

Encodings must keep grid constraints auditable: limits, ramping, reserve, uncertainty, contingency logic, and time coupling.

2. Count end-to-end overheads

Data loading, measurement, post-processing, and classical wrapper time all count. Kernel speed alone is not enough.

3. Enforce feasibility reliably

Low cost is not useful if ramping, reserve, security, or stability constraints fail. Feasibility has to be repeatable.

4. Integrate with existing toolchains

The easiest first use cases are warm starts, candidate screening, and decomposition support — not solver replacement.

Adoption criterion

A quantum-enabled method should earn its place by improving a measurable planning or operations workflow under the same data, constraints, and baseline comparisons used by practitioners.

KANSAS STATE
UNIVERSITY