Quantum methods for scientific optimisation and energy systems

Build an optimisation encoding that preserves the original engineering objective and feasibility requirements.

Undergraduate 22 min lesson Dr. Charith Jayasekara

By the end, you'll be able to

- Evaluate a constrained two-bit penalty model.
- Explain where phase and mixing change a candidate distribution.
- Benchmark feasible decisions rather than only average energy.

Begin with the engineering decision

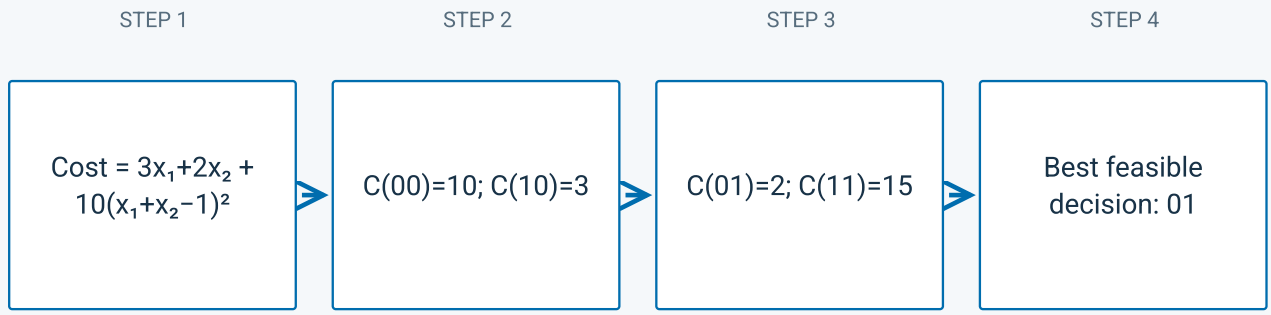
Scientific optimisation and energy-system planning contain many different problems: choosing equipment, scheduling generators, routing resources, fitting parameters or balancing a network. Some are discrete, some continuous, and some mix both. A quantum method should be selected only after the objective, constraints, uncertainty and required solution quality are explicit.

The same application can lead to several algorithmic routes. A binary selection model may be encoded as a QUBO and explored with QAOA or annealing. A structured linear subproblem may suggest a quantum linear-algebra primitive. Scenario evaluation may suggest amplitude estimation. These routes solve different mathematical tasks and have different access and output requirements.

Review oracle costs, complexity and hardware limitations. A useful classical prerequisite is distinguishing a feasible solution from an optimal one.

WORKED DIAGRAM · AA26

Penalty calibration changes the optimum



With penalty 10, 01 is the global minimum at cost 2. Reducing the penalty to 1 gives C(00)=1, so an infeasible decision becomes preferred.

State the optimisation setup

Specify decision variables, objective units, constraints and uncertainty model. Define whether the goal is exact optimality, a bounded approximation, a feasible solution within a deadline or a distribution of useful candidates. The output is usually a classical decision vector whose feasibility and cost can be checked. A low quantum expectation value is not itself that decision vector.

Also state the access model. Is the objective a short analytic expression, a large simulator or a database-dependent calculation? Can it be evaluated reversibly? Is there a natural quantum subproblem? A mathematical reformulation that hides an expensive simulator inside one oracle call does not establish an end-to-end improvement.

Work a binary selection model

Choose exactly one of two resources. Let x1 and x2 be binary variables, with operating costs three and two. Use the penalised objective

$$C(x_1, x_2) = 3x_1 + 2x_2 + 10(x_1 + x_2 - 1)^2.$$

For 00, the cost is ten because the selection constraint is violated. For 10, it is three. For 01, it is two. For 11, it is fifteen: operating cost five plus penalty ten. The best feasible decision is 01 with cost two. This complete table is the reference against which any circuit should be checked.

The penalty is not a physical law. It is a modelling choice large enough in this tiny example to make infeasible decisions unattractive. Too small a penalty can favour a constraint violation; excessively large penalties can create an awkward energy scale and make fine objective differences harder to resolve. A constrained mixer or another formulation may handle feasibility differently.

Enumerate the complete cost table before asking a solver to minimise it.

$$C = 10 - 7x_1 - 8x_2 + 20x_1x_2, \qquad \big(C(0,0), C(1,0), C(0,1), C(1,1)\big) = (10, 3, 2, 15).$$

This expands the existing penalised model using $x_i^2 = x_i$ for binary variables. The feasible optimum is 01 with cost two. The table checks both the operating-cost coefficients and the constraint penalty.

Convert the model to a circuit objective

Replace each binary variable by one half of one minus its Z operator. The resulting diagonal Hamiltonian has the same objective values on computational-basis labels. A QAOA-style circuit alternates a cost-dependent phase operation with a mixer, often built from X rotations. Parameters are chosen to improve an objective estimated from sampled bitstrings.

The cost phase alone does not increase the probability of a cheap decision. It changes phases, and the mixer converts those relationships into a new distribution. This is the same amplitude-to-probability distinction introduced in the foundation interference lessons. A correct explanation should show where interference enters rather than claiming that low energies are automatically selected.

Farhi, Goldstone and Gutmann introduce QAOA as a family of approximate optimisation algorithms and analyse specified settings. Its general practical performance depends on circuit depth, parameter optimisation and the problem family. The original proposal is not a theorem that every industrial QUBO is solved efficiently. Read the QAOA paper.

Follow an application workflow

Start with a small instance whose feasible solutions can be enumerated. Verify the objective and constraints classically. Construct the quantum encoding and check every basis-state energy. Select an ansatz or annealing schedule, define the optimiser and shot budget, and collect candidate decisions. Evaluate feasibility and original-unit cost after measurement, including any repair procedure.

Then scale a family of instances while tracking approximation quality, probability of a feasible sample and total time. If a repair heuristic supplies much of the improvement, compare that heuristic with classical samples too. A quantum pipeline should not claim the benefit of a strong classical postprocessor without testing what the postprocessor achieves independently.

Energy applications add physical constraints

Real energy systems can include network flows, ramp limits, reliability requirements and continuous operating quantities. A simple binary penalty model is only a teaching abstraction. Discretisation can introduce many variables and approximation error. Coupling a binary commitment decision to a continuous dispatch or power-flow calculation may require a hybrid decomposition.

A faster optimiser is useful only if the resulting decision satisfies the original engineering model. Report constraint violations in physical units, not merely a low penalised objective. Some penalty formulations can trade a serious violation for a small objective improvement unless the scales are calibrated carefully. Verification should therefore return to the original constraints.

Compare with strong baselines

Mixed-integer optimisation, dynamic programming, decomposition, local search and domain heuristics can be powerful. Match objective accuracy, feasibility, runtime limits and preprocessing. Compare distributions across random seeds rather than one best quantum sample against an average classical run. Include parameter-search time, failed runs and repeated circuit executions.

For fault-tolerant search or estimation proposals, count reversible objective evaluation and workspace. For near-term variational proposals, count measurement and optimisation overhead and distinguish empirical results from asymptotic guarantees. A small proof of concept can validate an encoding without demonstrating a production advantage.

Resource and uncertainty questions

Record qubit count, circuit depth, two-qubit operations, connectivity, shots and classical optimisation cost. If uncertain scenarios are sampled, separate statistical uncertainty from optimisation error. If the model changes with new data, include the cost of rebuilding the encoding or retraining parameters. Reusing a preprocessed model across many decisions can change the comparison, but that reuse should be explicit.

The scientific lesson is to keep the original decision visible throughout the algorithm. Qubit encodings and cost Hamiltonians are tools for solving a specified problem. They should not replace the question of whether the final decision is feasible, reliable and useful in the system being modelled.

Practice with answers

Exercise 1. What happens if the penalty coefficient in the example is reduced from ten to one?

Check the worked answer

The costs become one, three, two and six. The infeasible 00 decision now has the smallest penalised cost. The encoding no longer makes the constrained optimum the global minimum of the penalised objective.

Exercise 2. Does a low average penalised energy guarantee that every sampled decision is feasible?

Check the worked answer

No. An average can mix good and invalid samples, and the penalty may be miscalibrated. Report the feasible-sample probability and evaluate each returned decision against the original constraints. The application needs useful decisions, not only an expectation value.

Check your understanding

PAUSE & THINK

What goes wrong when the penalty coefficient becomes 1?

- [ ] The feasible optimum changes to 10
- [ ] Infeasible 00 becomes the lowest penalised cost
- [ ] Every state has the same cost

Answer: Infeasible 00 becomes the lowest penalised cost. The costs become 1,3,2,6 in order 00,10,01,11. The penalty no longer enforces the intended constraint through the global minimum.

Sources

- Read the QAOA paper