Why quantum?

Flood decisions are combinatorial selection problems: the number of options explodes as the network grows.

Which **k** sub-basins saturating **together** cause the largest downstream delay?

POSSIBLE SCENARIOS C(N, K) · LOG

C(20, 3) · Kelantan network · 20 sub-basins, k = 3 — **1,140** scenarios

Exact classical solution: milliseconds (ILP 14 ms). The same instance also ran on 3 IBM Heron processors: they found the best solution, but equal-budget random sampling found it too.

C(100, 10) · 100 sub-basins, k = 10 — **≈ 1.7×10^{13}** scenarios

C(200, 15) · 200 sub-basins, k = 15 — **≈ 1.46×10^{22}** scenarios

Exact value: 14,629,416,353,818,682,834,880

Same problem class: **sensor placement** · **coupled gate and pump operation**

A scenario count is not a solve time; classical solvers exploit structure and remain strong today.

Honest note: combinatorial optimisation of this kind is one of the problem classes where a benefit from quantum hardware is plausibly expected as hardware scales. That is an expectation, not a proof.

1. **Classical decides**
2. **Quantum runs in the shadow**
3. **We claim nothing unproven**

Measured, not promised

- **IBM Heron (ibm_kingston), 100 / 133 / 148 qubits:** in 60 of 60 pre-hashed instances, significantly better than random sampling; results were 0.93–0.97 of the pre-registered noise prediction.
- **Origin Wukong (WK_C180), 32 qubits:** better than random in 4 of 4 instances (with a pre-registered readout-error model).
- **9 quantum pre-registration documents** were sealed with sha256 before their hardware results were seen.
- **Failures published too:** the hardware reached the optimum in none of the 60 instances; our first QAOA circuit was worse than random (bug found and fixed).

Vendor-independent: IBM + Origin adapters, suitable for on-prem deployment

Cost to the customer: zero — shadow mode; fees are for the classical forecast

Quantum-ready. We will claim superiority only if a pre-registered test proves it.