Iterative warm-started QAOA for optimising QUBO problems

Ryszard Kukulski

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

Quantum Computing Lab at IT4Innovations National Supercomputing Center

26 October 2026

QUBO: one normal form for many hard problems

QUBO Problem

$$\min_{x\in\{0,1\}^n} \; x^\top Q x \;=\; \sum_i Q_{ii}\, x_i + \sum_{i<j} Q_{ij}\, x_i x_j$$

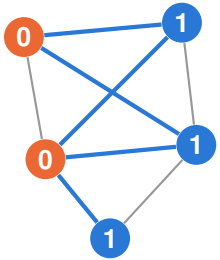
Ising model $x_i = \frac{1-s_i}{2},\; s_i \in \{\pm 1\}$

$$E(s) = \sum_{i<j} J_{ij}\, s_i s_j + \sum_i h_i\, s_i \;(+\text{const})$$

Qubits $s_i \mapsto Z_i$

$$H_C = \sum_{i<j} J_{ij} Z_i Z_j + \sum_i h_i Z_i$$

Constraints $Ax = b \;\Rightarrow\; +\lambda \|Ax - b\|^2$



Max-Cut: maximize the **cut edges**

Max-Cut $\subset$ QUBO $\Rightarrow$ NP-hard

Where QUBO shows up

Lucas (2014): Ising forms for **all of Karp's 21** NP-complete problems

Max-Cut



$$\max \sum_{(ij)\in E} w_{ij}\,(x_i + x_j - 2x_i x_j)$$

Number partitioning



$$H = \Big(\sum_i a_i\, s_i\Big)^2$$

$= 0 \iff$ two equal halves

Graph colouring



$$\sum_v \big(1 - \sum_c x_{v,c}\big)^2$$
$$+ \sum_{(uv)\in E} \sum_c x_{u,c}\, x_{v,c}$$

Travelling salesman



one-hot $x_{v,t}$ (city v at step t)

$$+ \sum_t \sum_{u,v} d_{uv}\, x_{u,t}\, x_{v,t+1}$$

Knapsack



$$-\sum_i c_i x_i$$
$$+\lambda\big(\sum_i w_i x_i + \text{slack} - W\big)^2$$

Portfolio selection

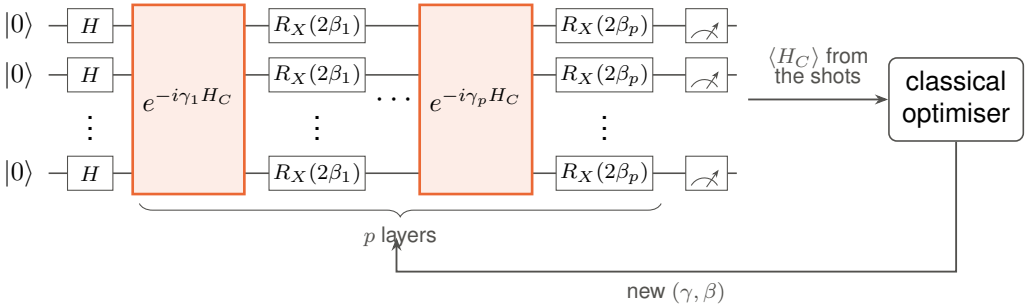


$$q\, x^\top \Sigma x - \mu^\top x$$
$$+\lambda\big(\sum_i x_i - B\big)^2$$

A. Lucas, Front. Phys. 2, 5 (2014).

Standard QAOA

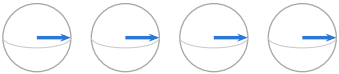
$$|\gamma,\beta\rangle = \prod_{k=1}^{p} e^{-i\beta_k H_M}\, e^{-i\gamma_k H_C}\, |+\rangle^{\otimes n}, \qquad H_M = \sum_i X_i$$



E. Farhi, J. Goldstone, S. Gutmann, arXiv:1411.4028 (2014);

Warm-start QAOA

standard QAOA



every qubit on the equator: $|+\rangle$

warm start



tilted towards a classical solution

1. relaxed solution $c^* \in [0,1]^n$, clipped to $[\varepsilon, 1-\varepsilon]$

2. start state

$$|c_i\rangle = R_Y(\theta_i)\,|0\rangle = \sqrt{1-c_i}\;|0\rangle + \sqrt{c_i}\;|1\rangle$$

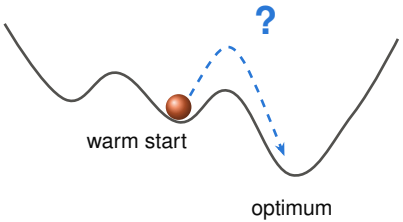
3. aligned mixer ($|c\rangle$ = its ground state)

$$H_M^{\mathrm{WS}} = -\sum_i \left(\sin\theta_i\, X_i + \cos\theta_i\, Z_i\right)$$

$c_i = \frac{1}{2}$ gives back standard QAOA.

D. J. Egger, J. Mareček, S. Woerner, Quantum 5, 479 (2021); R. Tate et al., Quantum 7, 1121 (2023).

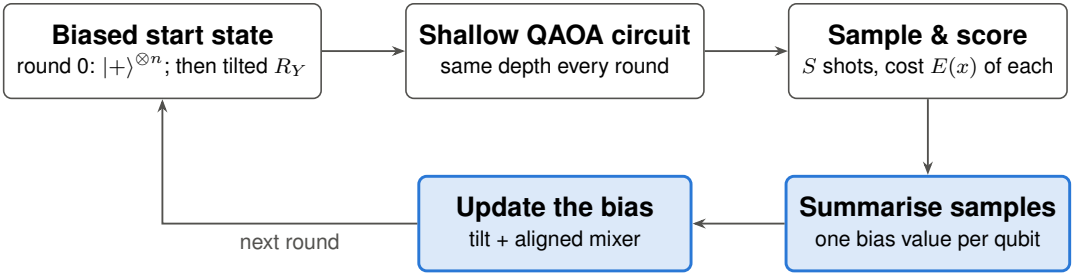
The catch



- Start in one good string
  $\Rightarrow$ QAOA at any accessible depth barely improves it
- Bias close to the poles
  $\Rightarrow$ very slow convergence

M. Cain et al., arXiv:2207.05089 (2022); R. Tate, S. Eidenbenz, arXiv:2410.00027 (2024).

Looping QAOA's warm start



Circuit depth never grows — only single-qubit rotations change between rounds.

One update

$$m_i = \sum_{x \in \text{samples}} w(x)\,(1 - 2x_i) \;\longrightarrow\; p_i = \mathrm{clip}\Big(\tfrac{1-\eta\, m_i}{2},\, \varepsilon,\, 1-\varepsilon\Big) \;\longrightarrow\; \theta_i = 2\arcsin\sqrt{p_i}$$

which samples count?

- best string only
- Boltzmann: $w \propto e^{-E(x)/T}$
- plain mean $\langle Z_i \rangle$

how strong a tilt?

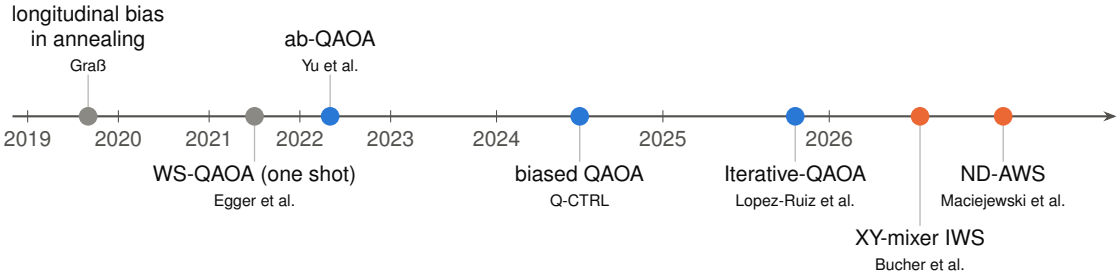
- learning rate η
- clip ε, tightened as quality improves
- e.g. $0.5 \to 0.1 \to 0.05 \to 0.025$

which angles?

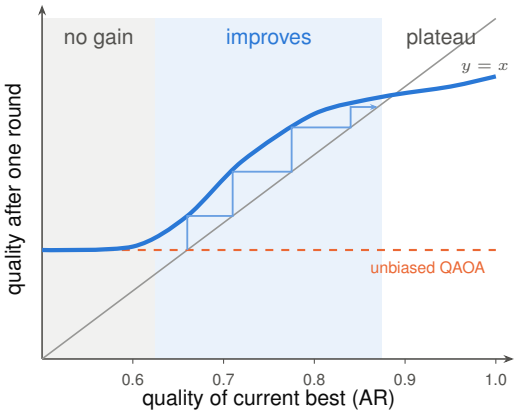
- re-optimise every round
- optimise once, reuse
- fixed schedule (no optimiser at all)

Yu et al., PRR 4, 023249 (2022); Romero et al., Commun. Phys. 8, 348 (2025); Lopez-Ruiz et al., arXiv:2510.26859; Maciejewski et al., arXiv:2607.09368.

A growing family, 2019 – 2026



Why chaining works



- Warm start pays off once the candidate is good enough
  AR $\gtrsim 0.6$–0.65
- Each round improves it, up to a plateau
  AR ≈ 0.85–0.9

Schematic (shape only) after the $p = 1$, 500-qubit analysis in F. B. Maciejewski et al., arXiv:2607.09368 (2026).

Variational or not?

re-optimise every round

ab-QAOA,
biased QAOA

− many circuit calls
− barren plateaus

optimise once, reuse

XY-mixer:
good angles roughly fixed

+ cheap
− worse quality

fixed schedule, no optimiser

Iterative-QAOA (linear ramp)

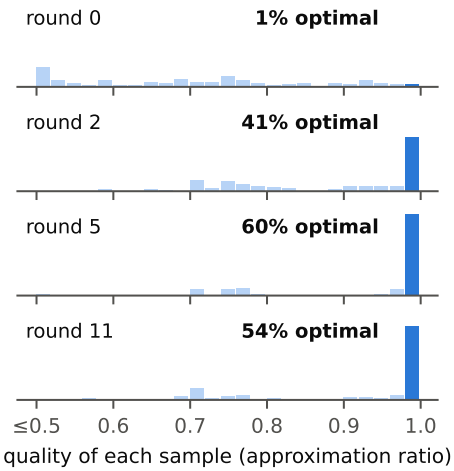
+ cheapest
− must suit the instance

less classical optimisation

The loop can replace the expensive outer optimiser.

J. A. Montañez-Barrera, K. Michielsen, npj Quantum Inf. 11 (2025) (linear-ramp QAOA); D. Bucher et al., arXiv:2604.02083 (2026).

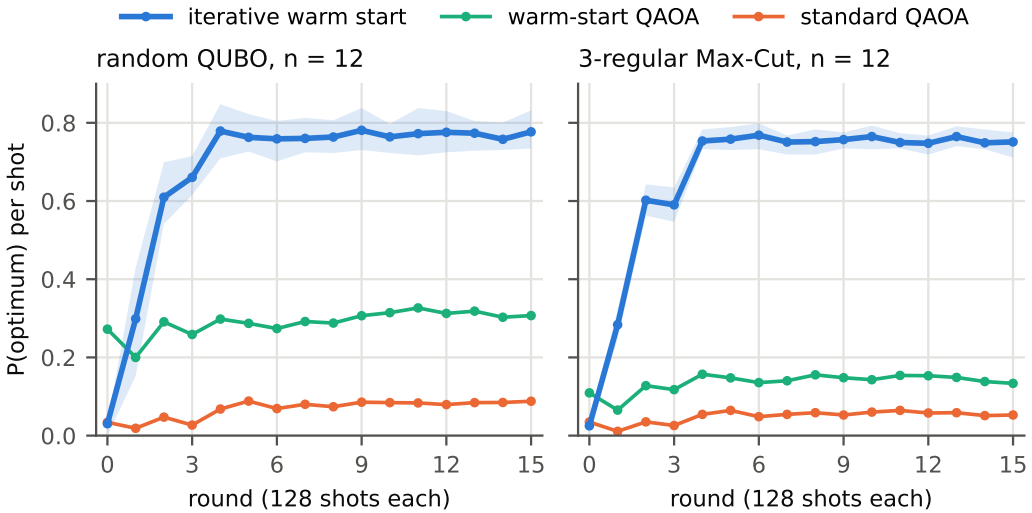
Example 1: what the loop does to the samples



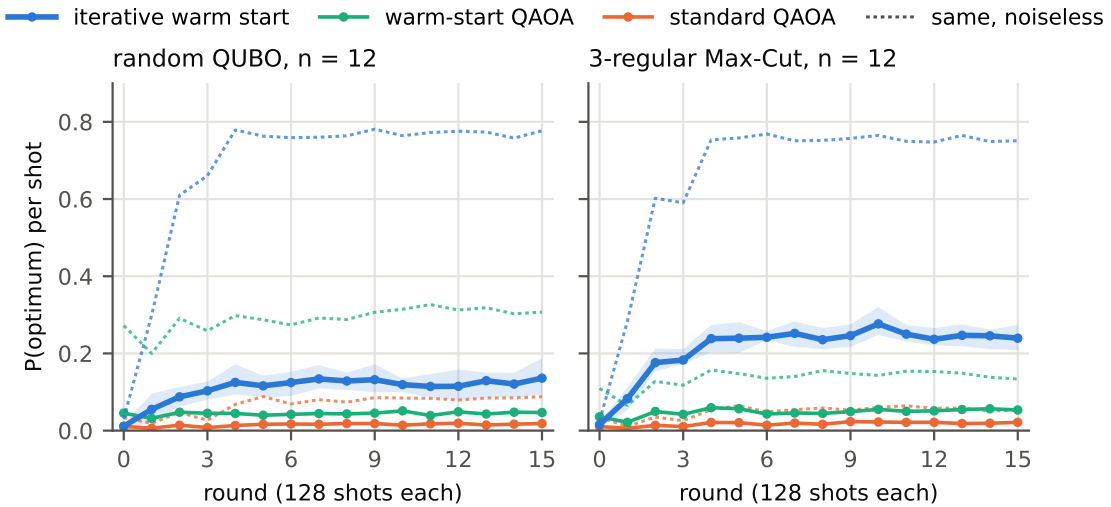
One random QUBO, $n = 12$
depth $p = 2$, noiseless, 128 shots per round

Round 0 is plain QAOA.

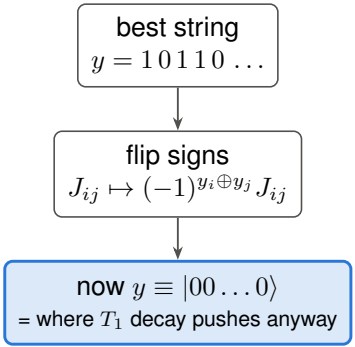
Example 2: 40 instances, noiseless



Example 3: noisy simulation with Qiskit Aer



Noise-directed adaptive warm-starting



Backend	`ibm_boston`
30	random Ising instances
0.969–1.0	approximation ratio
9–23	iterations

F. B. Maciejewski, S. Hadfield, O. Wallis, G. Pennington, S. Brandhofer, S. Woerner, D. J. Egger, D. Venturelli, arXiv:2607.09368 (2026).

Open problems and directions

- **Theory**
  when does the loop converge, and to what?
- **Self-reinforcing bias**
  escaping a wrong pool: restarts, several candidates
- **Bias schedules**
  annealing ε, adaptive per qubit
- **Noise as a resource**
  noise-directed variants, sampling-friendly mitigation
- **Fair benchmarks**
  null hypothesis, end-to-end time, strong baselines

Thank you for your attention!

Simulation details

Instances (exact optimum by enumeration)
A: 20 random QUBOs, $n = 12$, $Q_{ij} \sim U[-1, 1]$, 50 % dense
B: 20 weighted Max-Cut, random 3-regular, $n = 12$, $w \sim U[0, 1]$

Methods (all $p = 2$, 12×128 shots)
M1 standard QAOA, angles optimised once
M2 one-shot WS from one greedy local search, $\varepsilon = 0.25$
M3 iterative WS, round-0 angles reused
M4 same loop, $\gamma = \beta = 0$ (null hypothesis)

Noise (quantum trajectories, one per shot)
heavy-hex 27-qubit layout, greedy SWAP routing
CZ error 0.3 %, 84 ns; sx error 0.025 %, 32 ns
$T_1 = 200\,\mu$s, $T_2 = 150\,\mu$s; readout $1\,\% / 2\,\%$