www.q-solid.de

Quantum Computer in the Solid State

Solving constrained combinatorial optimization problems on quantum devices with linear penalty terms

Pim van den Heuvel[1,2], J. A. Montanez-Barrera[1], Dennis Willsch[1], Kristel Michielsen[1,2,3]

[1]Jülich Supercomputing Centre, Institute for Advanced Simulation, Forschungszentrum Jülich, 52425 Jülich, Germany
[2]RWTH Aachen University, 52056 Aachen, Germany
[3]AIDAS, 52425 Jülich, Germany

Constrained quantum optimization

Combinatorial optimization problems are usually written in quadratic unconstrained binary optimization (QUBO) form to run on quantum devices. (In)equality constraints are encoded into the cost function (C) as quadratic penalties, for example, for some budget B:

$$\max_{\mathbf{x}} \ C \ \text{s.t.} \ \sum_i x_i = B \longrightarrow \max_{\mathbf{x}} \ C - \lambda\left(\sum_i x_i - B\right)^2$$

For some problems, a linear instead of a quadratic penalty can suffice to correctly encode the optimization problem[1,2]:

$$\max_{\mathbf{x}} \ C \ \text{s.t.} \ \sum_i x_i = B \longrightarrow \max_{\mathbf{x}} \ C \pm \lambda\left(\sum_i x_i - B\right)$$

The linear penalty can only work for problems with a specific structure with respect to the Hamming weight ($\sum_i x_i$) of possible solutions. Namely, the difference in the maximum energy needs to be a monotonically decreasing function of the Hamming weight (for maximization problems). In other words, its 'derivative' is monotonically decreasing and the spectra of these problems seem to have a 'quadratic' shape with respect to the Hamming weight. The linear penalty is then able to shift the maxima of such spectra to the desired Hamming weight.

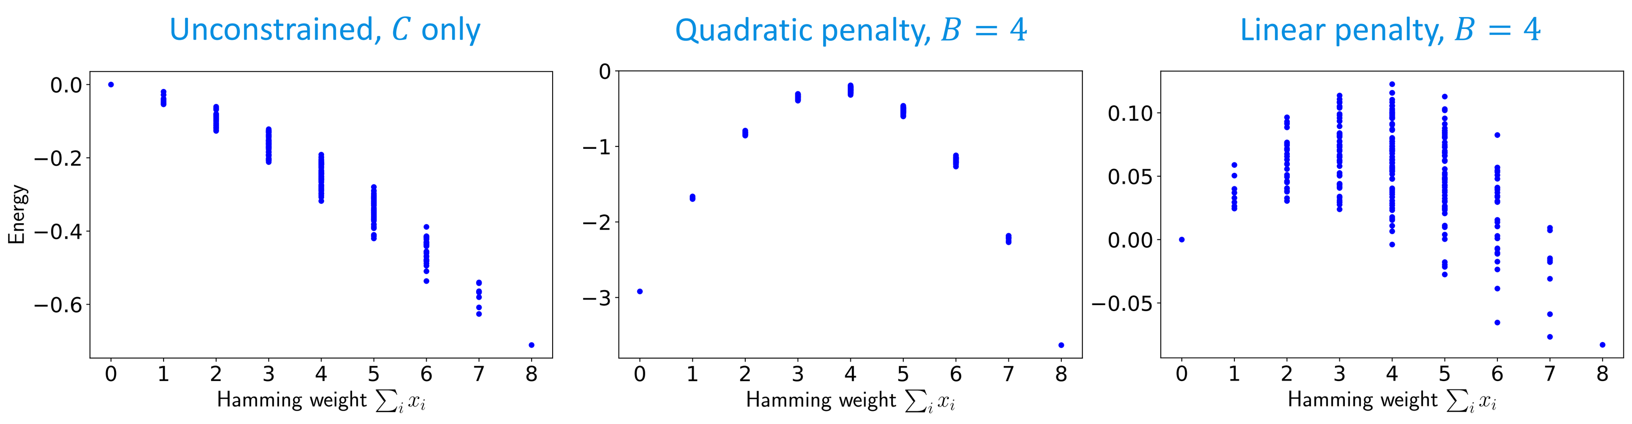

Advantages of linear penalties over quadratic ones:

- Fewer two-qubit terms in the cost function
- Smaller range of energy values, especially useful on devices with restrictions on this range (D-Wave devices)

QAOA for larger problem sizes

One specific Portfolio Optimization problem yielded noteworthy success probabilities at relatively modest depth (10 QAOA layers). It has a dominant quadratic term because of a relatively large q value and only positive covariances σ. We analyzed the scaling behavior of this problem for larger problem sizes (more qubits) using the Jülich Universal Quantum Computer Simulator (JUQCS)[3].

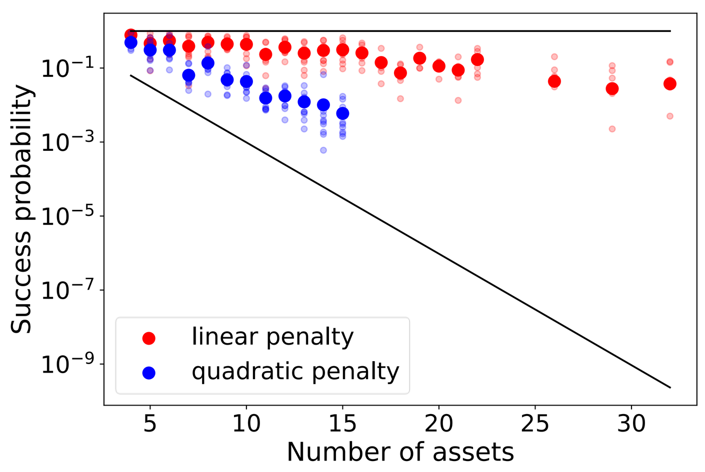

Although the success probability does drop off when the number of qubits is increased, the success probability for this problem can still exceed 10^{-1} even for > 30 qubits. For a QAOA circuit with only 10 layers, these are promising values. Since these calculations are more computationally expensive, no optimization is performed.

Neglecting some two-qubit terms

The entries of the σ covariance matrices in the Portfolio Optimization problems are random and some of them might be so small that they have negligible impact on the problems. Removing the smallest entries means that fewer two-qubit terms have to be implemented on the hardware with the linear penalty method.

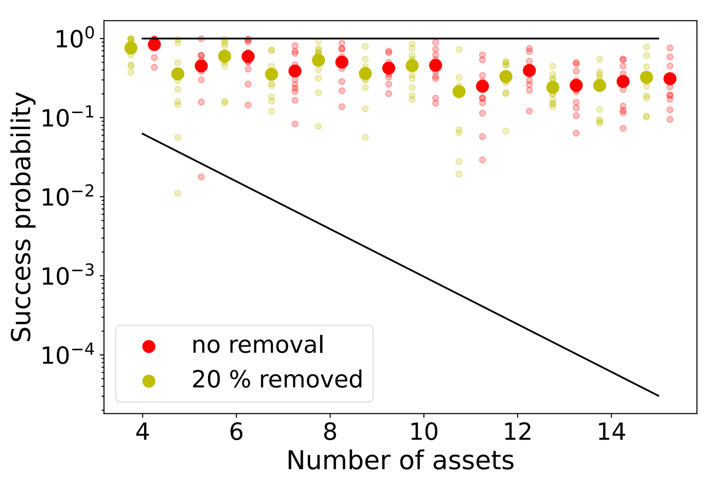

If 20 % of the smallest σ matrix entries are removed for the well-performing Portfolio Optimization problem, the success probability changes minimally. So at least in theory it seems fine to neglect some of the two-qubit terms. With current noisy devices, it seems likely that removing two-qubit terms provides an advantage in practice.

QAOA comparison of linear vs quadratic penalty

To test the two penalty methods, we ran QAOA simulations over a range of problem sizes (n), always with ten layers and for ten randomly generated problems. Multipliers λ are selected from a range (of 20 values) that depends on the problem. The range is the same for the linear and quadratic penalty method. Only the result from the best multiplier is kept and plotted.

Quadratic Knapsack Problem (QKP)

$$\max_{\mathbf{x}} \sum_{i,j} \sigma_{ij} x_i x_j$$
$$\text{subject to} \ \sum_i x_i = B$$
$$\sigma_{ij} \sim \mathcal{U}(0, b)$$
$$B = \lfloor n/2 \rfloor$$

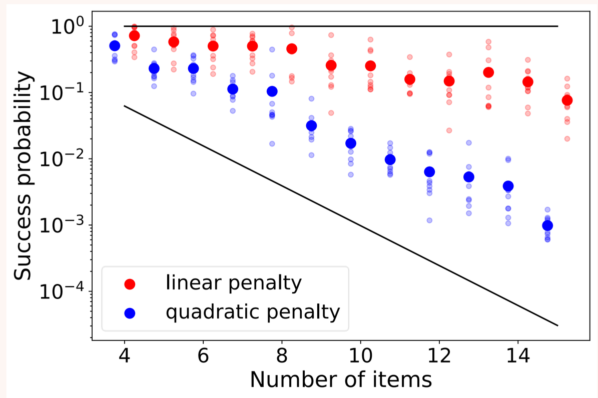

This QKP formulation doesn't have the correct spectrum shape for the linear penalty to work. Instead, the problem can be inverted by the transformation $\sigma_{ij} \to \max(\sigma) - \sigma_{ij}$ and changed to a minimization problem. This is not the same problem anymore, but the valid ground state is still encoded correctly.

Portfolio Optimization (PO)

v_i are values of items / assets
$\sigma_{i,j}$ are values / covariances of pairs of items
q and A are constants; q denotes risk
B is the available budget

$$\max_{\mathbf{x}} \sum_i v_i x_i - q \sum_{i,j} \sigma_{ij} x_i x_j$$
$$\text{subject to} \ \sum_i x_i = B$$
$$v_i \sim \mathcal{U}(0, b)$$
$$\sigma \text{ is PSD}, \sigma_{ij} \sim \mathcal{N}(0, c)$$
$$B = \lfloor n/2 \rfloor$$
$$q = 0.5$$

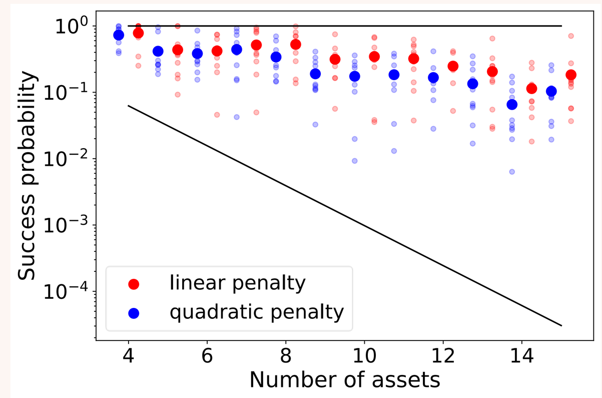

Weighted Independent Set (WIS)

$$\max_{\mathbf{x}} \sum_i v_i x_i - A \sum_{i,j \in E} x_i x_j$$
$$\text{subject to} \ \sum_i x_i = B$$
$$v_i \sim \mathcal{U}(0, b)$$
$$A = \max_i v_i$$
$$B = \lfloor n/4 \rfloor$$
$$i, j \in E \text{ with } P = 0.3$$

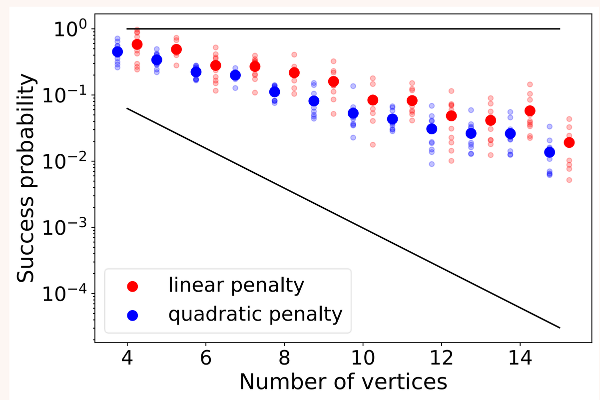

The initial parameters for the QAOA optimization procedure

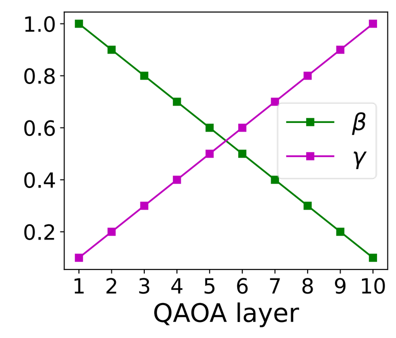

- On average, the linear penalty method performs better than the quadratic one across all problems and problem sizes
- There are significant differences between problems, and the linear penalty might do better (or the quadratic worse) for problems with relatively large quadratic coefficients.

Analysis of the difference between penalty methods

One-layer QAOA landscapes for the well-performing Portfolio Optimization problem ($n = 9$) show that the linear penalty method has only one clear peak and trough, which could explain its good performance. However, it's not certain that such qualitative observations extend to deeper circuits too.

Linear penalty

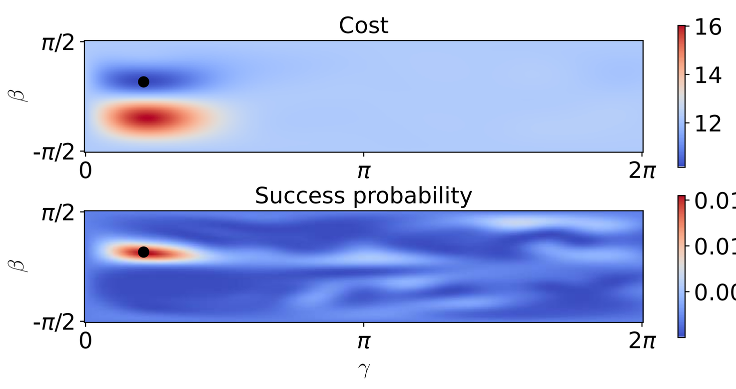

Quadratic penalty

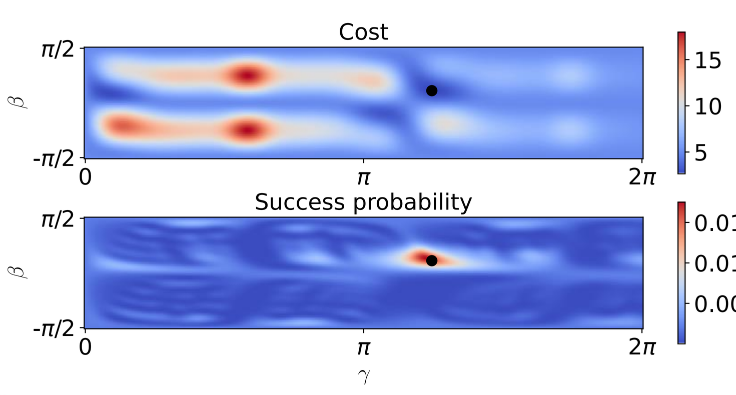

The linear penalty method reaches its highest success probability for the multiplier that encodes the solution in the ground state, as expected. For the quadratic penalty method, the highest success probability is sometimes reached when the problem is incorrectly encoded, and the success probability is more unpredictable.

Linear penalty

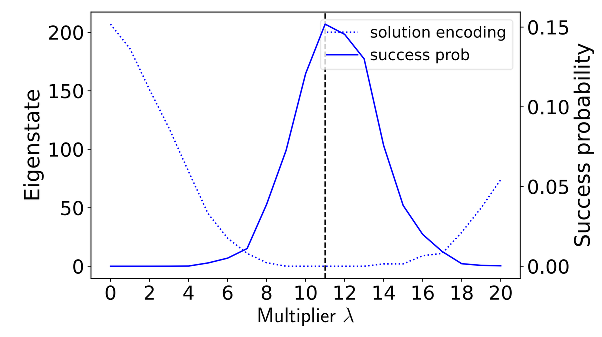

Quadratic penalty

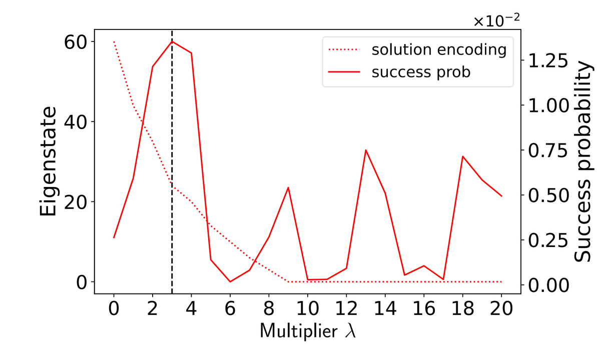

SPONSORED BY THE

Federal Ministry of Education and Research

The QSolid project acknowledges the support of the Federal Ministry of Education and Research (BMBF) within the framework programme "Quantum technologies – from basic research to market" (Grant No. 13N16149).

[1] Puya Mirkarimi, David C. Hoyle, Ross Williams, and Nicholas Chancellor. *Experimental demonstration of improved quantum optimization with linear Ising penalties*. 2024. arXiv: 2404.05476
[2] Puya Mirkarimi, Ishaan Shukla, David C. Hoyle, Ross Williams, and Nicholas Chancellor. *Quantum optimization with linear Ising penalty functions for customer data science*. 2024. arXiv: 2404.05467
[3] Hans De Raedt, Fengping Jin, Dennis Willsch, Madita Willsch, Naoki Yoshioka, Nobuyasu Ito, Shengjun Yuan, and Kristel Michielsen. "Massively parallel quantum computer simulator, eleven years later". In: Computer Physics Communications 237 (2019), pp. 47–61