



A case study of variational quantum algorithms for a job shop scheduling problem

David Amaro¹ , Matthias Rosenkranz^{1*} , Nathan Fitzpatrick² , Koji Hirano³ and Mattia Fiorentini¹ 

*Correspondence:

matthias.rosenkranz@cambridgequantum.com

¹Cambridge Quantum Computing Limited, SW1P 1BX, London, UK
Full list of author information is available at the end of the article

Abstract

Combinatorial optimization models a vast range of industrial processes aiming at improving their efficiency. In general, solving this type of problem exactly is computationally intractable. Therefore, practitioners rely on heuristic solution approaches. Variational quantum algorithms are optimization heuristics that can be demonstrated with available quantum hardware. In this case study, we apply four variational quantum heuristics running on IBM's superconducting quantum processors to the job shop scheduling problem. Our problem optimizes a steel manufacturing process. A comparison on 5 qubits shows that the recent filtering variational quantum eigensolver (F-VQE) converges faster and samples the global optimum more frequently than the quantum approximate optimization algorithm (QAOA), the standard variational quantum eigensolver (VQE), and variational quantum imaginary time evolution (VarQITE). Furthermore, F-VQE readily solves problem sizes of up to 23 qubits on hardware without error mitigation post processing.

Keywords: Combinatorial optimization; Variational quantum algorithms; Heuristics; Quantum hardware

1 Introduction

One of the major drivers of industry's recent interest in quantum computing is the promise of improving combinatorial optimization. This could have a large impact across many sectors including manufacturing, finance, logistics and supply chain management. However, most combinatorial optimization problems are NP-hard making it unlikely that even quantum computers can solve them efficiently in the worst-case. Informally, NP-hardness means that finding exact solutions is not more efficient than going through all potential solutions—at a cost that grows exponentially with the problem size. Quantum algorithms such as Grover's perform exhaustive search with a quadratic speedup but require fault tolerant quantum hardware [1, 2]. Instead it is interesting to explore if quantum computers could speed up the average-case or special cases of practical interest or, indeed, improve approximate solutions in practice on non-fault-tolerant hardware.

A large body of research focuses on quantum-enhanced optimization heuristics for the noisy intermediate-scale quantum (NISQ) era [3–5]. Typically, these algorithms don't come equipped with convergence guarantees and instead solve the problem approximately within a given computational budget. While many fault-tolerant optimization al-

© The Author(s) 2022. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

gorithms can also be formulated as heuristics [6], our focus is on variational quantum algorithms (VQA). Typically VQA employ objective functions implemented with parameterized quantum circuits (PQCs) and update their parameters via a classical optimization routine. In our context, a common approach for combinatorial optimization encodes the optimal solution in the ground state of a classical multi-qubit Hamiltonian [7–9].

Studying the effectiveness of such heuristics relies on intuition and experimentation. However, today's quantum computers are noisy and fairly limited in size making such experimentation hard. Nevertheless it is important to gauge properties such as convergence speed, scalability and accuracy from the limited hardware we have available. To make the most of today's NISQ computers it is reasonable to compare different VQA on concrete problems.

We selected the popular quantum approximate optimization algorithm (QAOA) [10] and the variational quantum eigensolver (VQE) [11] as well as the less well studied variational quantum imaginary time evolution algorithm (VarQITE) [12] and the filtering variational quantum eigensolver (F-VQE) [13] recently introduced by some of the present authors. Despite its promising properties, such as supporting a form of quantum advantage, [14–16] and considerable progress with regards to its experimental realization [17], in general the QAOA ansatz requires circuit depths that are challenging for current quantum hardware. VQE, VarQITE and F-VQE employ more flexible, hardware-efficient ansätze tailored for the particular quantum processor. Those ansätze feature high expressibility and entangling capabilities [18], which suggests that they can lead to genuinely different heuristics compared to classical ones. On the other hand, they are prone to barren plateaus which could prevent the algorithms' convergence at larger problem sizes [19, 20]. In addition, the classical optimizer can significantly affect the performance of quantum heuristics on NISQ hardware, and the magnitude of this effect can vary between optimization problems [21–25]. Those effects have made it difficult in the past to scale common VQA beyond small-scale experiments. Here we compare VQA executed on IBM's superconducting quantum computers with a view towards scaling up a particular optimization problem of industrial relevance.

We compare the effectiveness of VQE, QAOA, VarQITE and F-VQE on the job shop scheduling problem (JSP). The JSP is a combinatorial optimization problem where jobs are assigned to time slots in a number of machines or processes in order to produce a final product at minimal cost. Typically costs are associated with delivery delays or reconfiguration of production processes between time slots. The JSP formulation considered herein was developed by Nippon Steel Corporation and applies to processes typical of steel manufacturing.

This article is structured as follows. Section 2 introduces the JSP formulation and the four VQA employed in this work, highlighting their similarities and differences. Section 3 analyses the performance of all VQA and shows results of scaling up F-VQE on hardware. We conclude in Sect. 4. Appendix A includes a derivation of the JSP formulation, App. B discusses the scaling of the JSP, App. C lists key properties of the quantum processors used for this work, and App. D provides several additional results from hardware experiments.

2 Methods

This section introduces the JSP and its mathematical formulation in Sects. 2.1–2.2 and introduces the VQE, QAOA, VarQITE and F-VQE with our choices for the various settings of these algorithms in Sect. 2.3.

2.1 Job shop scheduling in a steel manufacturing process

The general JSP is the problem of finding an assignment—also called a schedule—of J jobs to M machines, where each job needs to be processed in a certain order across the machines. Each job can carry additional data such as due time or processing time. A JSP is typically described by two further components: *processing characteristics and constraints* and an *objective*. The processing characteristics and constraints encode the specifics of an application such as setup times of machines and job families or production groups. Typical examples of objectives to minimise include makespan (total completion time) or mismatch of the jobs' completion and due times (for an overview of common scheduling formulations, see Ref. [26]).

The JSP formulation we consider applies to general manufacturing processes and was fine-tuned by Nippon Steel Corporation for steel manufacturing. We consider jobs $j = 1, \dots, J$ assigned to different machines or processes $m = 1, \dots, M$ at time slots $t_m = 1, \dots, T_m$. In this work, the processing times of all jobs for all processes are assumed to be equal. Accordingly, time slots can be common across the multiple processes and thus t_m is simplified as t throughout the paper. The processing times of all jobs are equal and each job is assigned a due time d_j . Each machine m is allowed to idle for a total number of time slots $e_m \geq 0$ at the beginning or end of the schedule. This number is an input of the problem. Hence, the maximum time slot for machine m is $T_m = J + e_m$.

The objective is to minimize the sum of *early delivery* and *late delivery* of jobs leaving the last machine, and the *production cost* associated with changing the processing conditions for subsequent jobs in each machine. Early (late) delivery is quantified by a constant c_e (c_l) multiplied by the number of time steps a job finishes before (after) its due date, summed over all jobs. To compute the production cost for each machine m each job j is assigned a *production group* P_{mj} . The production cost is quantified by a constant c_p multiplied by the total number of times consecutive jobs j_1, j_2 in a machine m switch productions groups, i.e. $P_{mj_1} \neq P_{mj_2}$. Figure 1 illustrates these costs for the largest (20-job) JSP instance we consider in this work.

We consider the following sets of constraints, which follow from the specifics of the manufacturing process.

1. *Job assignment constraints*. Each job is assigned to exactly one time slot in each machine.
2. *Time assignment constraints*. J jobs are distributed to J consecutive time slots in each machine.
3. *Process order constraints*. Each job must progress from machine 1 to M in non-descending order.
4. *Idle slot constraints*. Idle slots occur only at the beginning or end of a schedule.

2.2 Quadratic unconstrained binary optimization formulation of the JSP

We formulate the JSP defined in Sect. 2.1 as a Quadratic Unconstrained Binary Optimization (QUBO) problem. A feasible solution of the JSP is a set of two schedules (\mathbf{x}, \mathbf{y}) given by binary vectors $\mathbf{x} \in \mathbb{B}^{N_x}$ for the *real jobs* (those corresponding to jobs $1, \dots, J$) and $\mathbf{y} \in \mathbb{B}^{N_y}$ for the *dummy jobs* introduced to fill idle time slots at the beginning and end of each machine's schedule. Here $\mathbb{B} = \{0, 1\}$, $N_x = \sum_{m=1}^M J(J + e_m)$, and $N_y = \sum_{m=1}^M e_m$. N_y is independent of J because, owing to the idle slot constraints, the optimization only needs to decide on the number of consecutive dummy jobs at the beginning of the schedule per machine.

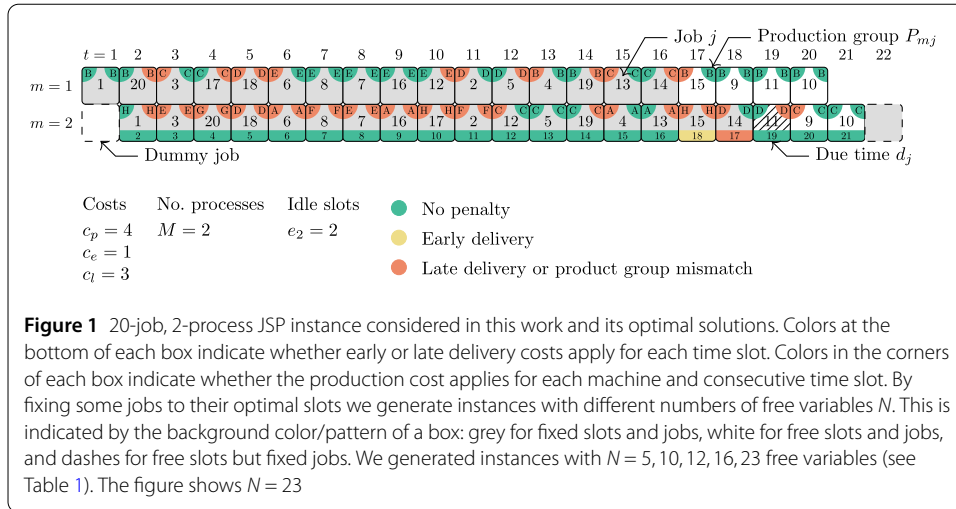


Table 1 Time slots and jobs needing assignment in each of the problem instances considered in this work

Free variables N	Machine m	Time slots	Jobs
5	2	1, 20–21	9–10
10	2	1, 19–21	9–11
12	2	19–22	9–11
16	2	18–21	9–11, 14
23	1; 2	17–20; 1, 19–21	9–11, 15; 9–10

A value $x_{mjt} = 1$ ($x_{mjt} = 0$) indicates that job j is assigned (is not assigned) to machine m at time t . Similarly, for dummy jobs, value $y_{mt} = 1$ ($y_{mt} = 0$) indicates that a dummy job is (is not) assigned to machine m at time slot t . With the cost and constraints of the JSP encoded in a quadratic form $Q: \mathbb{B}^{N_x} \times \mathbb{B}^{N_y} \rightarrow \mathbb{R}$ the JSP becomes

$$(\mathbf{x}^*, \mathbf{y}^*) = \arg \min_{(\mathbf{x}, \mathbf{y}) \in \mathbb{B}^{N_x} \times \mathbb{B}^{N_y}} Q(\mathbf{x}, \mathbf{y}). \tag{1}$$

The binary representation makes it straightforward to embed the problem on a quantum computer by mapping schedules to qubits.

The function Q for the JSP is

$$Q(\mathbf{x}, \mathbf{y}) = c(\mathbf{x}) + p \sum_{m=1}^M \sum_{j=1}^J (g_{mj}(\mathbf{x}) - 1)^2 + p \sum_{m=1}^M \sum_{t=1}^{T_m} (\ell_{mt}(\mathbf{x}, \mathbf{y}) - 1)^2 + p \sum_{m=1}^{M-1} \sum_{j=1}^J q_{mj}(\mathbf{x}) + p \sum_{m=2}^M \sum_{t=1}^{e_{m-1}} r_{mt}(\mathbf{y}). \tag{2}$$

All terms are derived in more detail in App. A. $c(\mathbf{x})$ is the cost of the schedule, Eq. (21), $g_{mj}(\mathbf{x})$ encodes the job assignment constraints, Eq. (22), $\ell_{mt}(\mathbf{x}, \mathbf{y})$ encodes the time assignment constraints, Eq. (23), $q_{mj}(\mathbf{x})$ encodes the process order constraints, Eq. (24), $r_{mt}(\mathbf{x})$ encodes the idle slot constraints, Eq. (25). The constraints are multiplied by a penalty p , which will be set to a sufficiently large value. To ensure non-negative penalties some constraints need to be squared. Note that Q is a quadratic form because all terms can be

written as polynomials of degree two in the binary variables \mathbf{x} and \mathbf{y} . To simplify notation we often denote the concatenation of the two sets of binary variables with $\mathbf{z} = (\mathbf{x}, \mathbf{y})$ and $Q(\mathbf{z}) = Q(\mathbf{x}, \mathbf{y})$. Figure 1 illustrates the largest JSP instance used in this work together with its optimal solution obtained via a classical solver, and Table 1 specifies all instances used. App. B derives the scaling of the total number of variables for this formulation.

Solving the JSP, Eq. (1), is equivalent to finding the ground state of the Hamiltonian

$$H = Q\left(\frac{I - \mathbf{Z}^{(\mathbf{x})}}{2}, \frac{I - \mathbf{Z}^{(\mathbf{y})}}{2}\right) = h_0 I + \sum_{n=1}^N h_n Z_n + \sum_{n,n'=1}^N h_{nn'} Z_n Z_{n'}, \tag{3}$$

where the vectors of Pauli Z operators $\mathbf{Z}^{(\mathbf{x})}, \mathbf{Z}^{(\mathbf{y})}$ correspond to the binary variables in \mathbf{x}, \mathbf{y} , respectively, \mathbf{Z} corresponds to \mathbf{z} , and $h_0, h_n, h_{nn'}$ are the coefficients of the corresponding operators. Note that this Hamiltonian is defined purely in terms of Pauli Z operators, which means that its eigenstates are separable and they are computational basis states.

2.3 Variational quantum algorithms for combinatorial optimization problems

VQA are the predominant paradigm for algorithm development on gate-based NISQ computers. They comprise several components that can be combined and adapted in many ways making them very flexible for the rapidly changing landscape of quantum hard- and software development. The main components are an *ansatz* for a PQC, a *measurement scheme*, an *objective function*, and a classical *optimizer*. The measurement scheme specifies the operators to be measured, the objective function combines measurement results in a classical function, and the optimizer proposes parameter updates for the PQC with the goal of minimising the objective function. As noted in Sect. 2.2, the JSP is equivalent to finding the ground state of the Hamiltonian Eq. (3). VQA are well suited to perform this search by minimising a suitable objective function. We focus on four VQA for solving the JSP: VQE, QAOA, VarQITE, and F-VQE.

We use conditional Value-at-Risk (CVaR) as the objective function for all VQA [27]. For a random variable X with quantile function F^{-1} the CVaR is defined as the conditional expectation over the left tail of the distribution of X up to a quantile $\alpha \in (0, 1]$:

$$\text{CVaR}_\alpha(X) = \mathbb{E}[X | X \leq F_X^{-1}(\alpha)]. \tag{4}$$

In practice we estimate the CVaR from measurement samples as follows. Prepare a state $|\psi\rangle$ and measure this state K times in the computational basis. Each measurement corresponds to a bitstring \mathbf{z}_k sampled from the distribution implied by the state $|\psi\rangle$ via the Born rule, $\mathbf{z}_k \sim |\langle \mathbf{z} | \psi \rangle|^2$. We interpret each bitstring as a potential solution to the JSP with energy (or cost) $E_k = Q(\mathbf{z}_k)$, $k = 1, \dots, K$. Given a sample of energies $\{E_1, \dots, E_K\}$ —without loss of generality assumed to be ordered from small to large—the CVaR estimator is

$$\widehat{\text{CVaR}}_\alpha(\{E_1, \dots, E_K\}) = \frac{1}{\lceil \alpha K \rceil} \sum_{k=1}^{\lceil \alpha K \rceil} E_k. \tag{5}$$

For $\alpha = 1$ the CVaR estimator is the sample mean of energies, which is the objective function often used in standard VQE. The CVaR estimator with $0 < \alpha < 1$ has shown advantages in applications that aim at finding ground states, such as combinatorial optimization problems [27] and some of our experiments confirmed this behaviour.

Table 2 VQA and settings used for the hardware experiments in Figs. 3–5. An initial parameter $|+\rangle^{\otimes N}$ means that the initial angles of all R_y in the first (second) layer of the ansatz are set to 0 ($\pi/2$). The last line highlights some key findings from our experiments

	N	VQE	QAOA	VarQITE	F-VQE
Ansatz	5	Fig. 2 ($p = 2$)	Eq. (7) ($p = 2$)	Fig. 2 ($p = 2$)	Fig. 2 ($p = 2$)
	>5	–	–	–	Fig. 2 ($p = 1$)
Initial param.		$ +\rangle^{\otimes N}$	uniform in $[0, \pi]$	$ +\rangle^{\otimes N}$	$ +\rangle^{\otimes N}$
Objective		CVaR Eq. (6) ($\alpha = 0.5$)	CVaR Eq. (8) ($\alpha = 0.5$)	Mean energy Eq. (9)	Custom Eq. (13)
Optimizer		COBYLA	COBYLA	Eq. (12)	Eq. (14)
No. shots	5	1000	1000	1000	1000
	10	–	–	–	500
	12	–	–	–	550
	16	–	–	–	650
	23	–	–	–	450
Quantum chip	5	multiple	multiple	multiple	multiple
	10	–	–	–	ibmq_toronto
	12	–	–	–	ibmq_guadalupe
	16	–	–	–	ibmq_manhattan
	23	–	–	–	ibmq_manhattan
Key findings		Flexible ansatz; converges slower than F-VQE	Ansatz fixed by problem topology; poor convergence likely due to noise	Flexible ansatz; strongly varying performance across runs; converges slower than F-VQE	Flexible ansatz; fastest, most consistent convergence

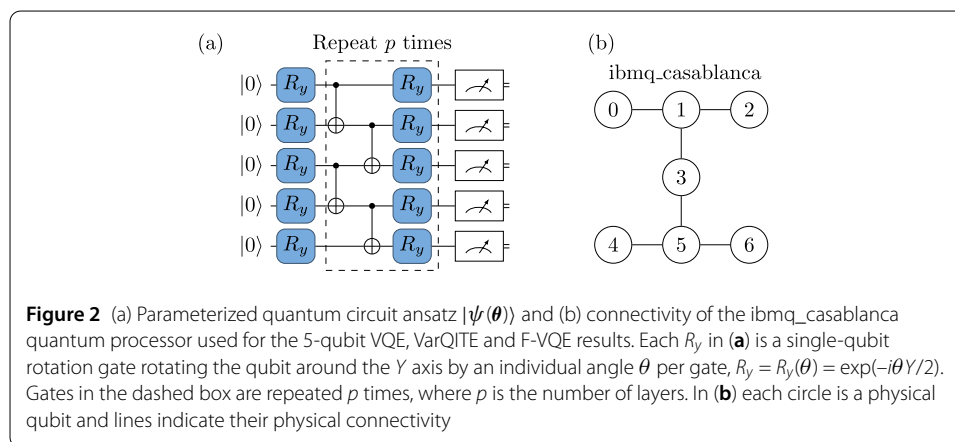


Figure 2 (a) Parameterized quantum circuit ansatz $|\psi(\theta)\rangle$ and (b) connectivity of the `ibmq_casablanca` quantum processor used for the 5-qubit VQE, VarQITE and F-VQE results. Each R_y in (a) is a single-qubit rotation gate rotating the qubit around the Y axis by an individual angle θ per gate, $R_y = R_y(\theta) = \exp(-i\theta Y/2)$. Gates in the dashed box are repeated p times, where p is the number of layers. In (b) each circle is a physical qubit and lines indicate their physical connectivity

The difference between the considered VQA boils down to different choices of the ansatz, measurement scheme, objective and optimizer. Table 2 compares the four algorithms and our concrete settings and Sects. 2.3.1–2.3.4 detail the algorithms. Appendix C lists the quantum processors used for the hardware execution.

2.3.1 Variational quantum eigensolver

VQE aims at finding the lowest energy state within a family of parameterized quantum states. It was introduced for estimating the ground state energies of molecules described by a Hamiltonian in the context of quantum chemistry. Exactly describing molecular ground states would require an exponential number of parameters. VQE offers a way to approximate their description using a polynomial number of parameters in a PQC ansatz. Since the JSP can be expressed as the problem of finding a ground state of the Hamiltonian Eq. (3), VQE can also be used for solving the JSP. This results in a heuristic optimization

algorithm for the JSP similar in spirit to classical heuristics, which aim at finding good approximate solutions.

Our VQE implementation employs the hardware-efficient ansatz in Fig. 2(a) for the PQC. Hardware-efficient ansätze are very flexible as they can be optimized for a native gate set and topology of a given quantum processor [28]. We denote the free parameters of the single-qubit rotation gates in the ansatz with the vector θ . The PQC implements the unitary operator $U(\theta)$ and $|\psi(\theta)\rangle = U(\theta)|0\rangle$ denotes the parameterized state after executing this PQC.

The measurement scheme for VQE is determined by the Hamiltonian we wish to minimize. In the case of JSP this reduces to measuring tensor products of Pauli Z operators given by Eq. (3). All terms commute so they can be computed from a single classical bitstring $\mathbf{z}_k \sim |\langle \mathbf{z} | \psi(\theta) \rangle|^2$ sampled from the PQC. Sampling K bitstrings and calculating their energies $E_k(\theta) = Q(\mathbf{z}_k(\theta))$ yields a sample of (ordered) energies $\{E_1(\theta), \dots, E_K(\theta)\}$ parameterized by θ . Plugging this sample into the CVaR estimator, Eq. (4), yields the objective function for VQE

$$O_{\text{VQE}}(\theta; \alpha) = \widehat{\text{CVaR}}_{\alpha}(\{E_1(\theta), \dots, E_K(\theta)\}). \tag{6}$$

We use the Constrained Optimization By Linear Approximation (COBYLA) optimizer to tune the parameters of the PQC [29]. This is a gradient-free optimizer with few hyperparameters making it a reasonable baseline choice for VQA [23].

2.3.2 Quantum approximate optimization algorithm

QAOA is a VQA which aims at finding approximate solutions to combinatorial optimization problems. In contrast to VQE, research on QAOA strongly focuses on combinatorial optimization rather than chemistry problems. QAOA can be thought of as a discretized approximation to quantum adiabatic computation [30].

The QAOA ansatz follows from applying the two unitary operators $U_M(\beta) = e^{-i\beta \sum_{n=1}^N X_n}$ and $U(\gamma) = e^{-i\gamma H}$ a number of p times to the N -qubit uniform superposition $|+\rangle = \frac{1}{\sqrt{2^N}} \sum_{n=0}^{2^N-1} |n\rangle$ in an alternating sequence. Here X_n is the Pauli X operator applied to qubit n and H is the JSP Hamiltonian, Eq. (3). The QAOA ansatz with $2p$ parameters (β, γ) is

$$|\psi(\beta, \gamma)\rangle = U_M(\beta_p)U(\gamma_p)U_M(\beta_{p-1})U(\gamma_{p-1}) \cdots U_M(\beta_1)U(\gamma_1)|+\rangle. \tag{7}$$

In contrast to our ansatz for VQE, in the QAOA ansatz the connectivity of the JSP Hamiltonian dictates the connectivity of the two-qubit gates. This means that implementing this ansatz on digital quantum processors with physical connectivity different from the JSP connectivity requires the introduction of additional gates for routing. This overhead can be partly compensated by clever circuit optimization during the compilation stage.

We use the same measurement scheme, objective function and optimizer for QAOA and VQE. Namely, we sample bitstrings $\mathbf{z}_k(\beta, \gamma)$ from the PQC and calculate their energies $E_k(\beta, \gamma) = Q(\mathbf{z}_k(\beta, \gamma))$. The objective function is the CVaR estimator

$$O_{\text{QAOA}}(\beta, \gamma; \alpha) = \widehat{\text{CVaR}}_{\alpha}(\{E_1(\beta, \gamma), \dots, E_K(\beta, \gamma)\}) \tag{8}$$

and the optimizer is COBYLA.

2.3.3 Variational quantum imaginary time evolution

Imaginary time evolution is a technique for finding ground states by evolving an initial state with the Schrödinger equation in imaginary time $\tau = it$. This technique has mainly been applied to study quantum many-body problems [31] and a variant of the algorithm shows promising results for combinatorial optimization [32]. Here we use a variational formulation of imaginary time evolution dubbed VarQITE [12] to find approximate solutions of the JSP.

Given an initial state $|\phi(0)\rangle$ the imaginary time evolution is defined by $|\phi(\tau)\rangle = e^{-H\tau}|\phi(0)\rangle/\sqrt{\mathcal{Z}(\tau)}$ with a normalization factor $\mathcal{Z}(\tau) = \langle\phi(0)|e^{-2H\tau}|\phi(0)\rangle$. The non-unitary operator $e^{-H\tau}$ cannot be mapped directly to a quantum circuit and is typically implemented via additional qubits and post-selection. To avoid additional qubits and post-selection, instead the VarQITE algorithm optimizes a PQC to approximate the action of the non-unitary operator. This is achieved by replacing the state $|\phi(\tau)\rangle$ with a state $|\psi(\theta)\rangle = |\psi(\theta(\tau))\rangle = U(\theta)|+\rangle$ and the parameters are assumed to be time-dependent $\theta = \theta(\tau)$. We use the PQC ansatz in Fig. 2(a) and set initial parameters such that the resulting initial state is $|+\rangle$.

We use the same measurement scheme as in VQE with the mean energy as the objective function, i.e. CVaR with $\alpha = 1$,

$$O_{\text{VarQITE}}(\theta) = \frac{1}{2} \widehat{\text{CVaR}}_1(\{E_1(\theta), \dots, E_K(\theta)\}). \tag{9}$$

VarQITE updates parameters with a gradient-based optimization scheme derived from McLachlan’s variational principle [31]. This lifts the imaginary time evolution of the state $|\phi(\tau)\rangle$ to an evolution of the parameters in the PQC via the differential equations

$$A(\theta) \frac{\partial \theta(\tau)}{\partial \tau} = -\nabla O_{\text{VarQITE}}(\theta), \tag{10}$$

where $A(\theta)$ is a matrix with entries

$$A_{ij} = \text{Re} \left(\left\langle \left. \frac{\partial \psi(\theta)}{\partial \theta_i} \right| \frac{\partial \psi(\theta)}{\partial \theta_j} \right\rangle \right). \tag{11}$$

We assume small time steps $\delta\tau$, denote $\tau_n = \tau_n + n\delta\tau$, $\theta_n = \theta(\tau_n)$ and approximate the parameter evolution Eq. (10) with the explicit Euler scheme

$$\theta_{n+1} = \theta_n - A^{-1}(\theta_n) \nabla O_{\text{VarQITE}}(\theta_n) \delta\tau. \tag{12}$$

We estimate the entries of A and $\nabla O_{\text{VarQITE}}$ with the Hadamard test. This requires an additional qubit and controlled operations.

2.3.4 Filtering variational quantum eigensolver

F-VQE is a generalization of VQE with faster and more reliable convergence to the optimal solution [13]. The algorithm uses *filtering operators* to modify the energy landscape at each optimization step. A filtering operator $f(H; \tau)$ for $\tau > 0$ is defined via a real-valued function $f(E; \tau)$ with the property that $f^2(E; \tau)$ is strictly decreasing on the spectrum of the Hamiltonian $E \in [E_{\min}, E_{\max}]$.

For F-VQE we used the ansatz in Fig. 2(a). In contrast to our VQE implementation, F-VQE uses a gradient-based optimizer. At each optimization step n the objective function is

$$O_{\text{F-VQE}}^{(n)}(\boldsymbol{\theta}; \tau) = \frac{1}{2} \left\| |\psi(\boldsymbol{\theta})\rangle - |F_n \psi_{n-1}\rangle \right\|^2, \tag{13}$$

where $|\psi_{n-1}\rangle = |\psi(\boldsymbol{\theta}_{n-1})\rangle$ and $|F_n \psi_{n-1}\rangle = F_n |\psi_{n-1}\rangle / \sqrt{\langle F_n^2 \rangle_{\psi_{n-1}}}$ with $F_n = f(H; \tau_n)$. We use the *inverse filter* $f(H; \tau) = H^{-\tau}$. It can be shown that the algorithm minimises the mean energy of the system, i.e. CVaR with $\alpha = 1$. The update rule of the optimizer at step n is

$$\boldsymbol{\theta}_{n+1} = \boldsymbol{\theta}_n - \eta \nabla O_{\text{F-VQE}}^{(n)}(\boldsymbol{\theta}_n; \tau), \tag{14}$$

where η is a learning rate. The gradient in Eq. (14) is computed with the parameter shift rule [33, 34]. This leads to terms of the form $\langle F \rangle_{\psi}$ and $\langle F^2 \rangle_{\psi}$ for states $|\psi\rangle$. They can be estimated from bitstrings $\mathbf{z}_k^{\psi}(\boldsymbol{\theta}) \sim |\langle \mathbf{z} | \psi(\boldsymbol{\theta}) \rangle|^2$ sampled from the PQC. A sample of K bitstrings yields a sample of filtered energies $\{f_1^{\psi}(\boldsymbol{\theta}; \tau), \dots, f_K^{\psi}(\boldsymbol{\theta}; \tau)\}$ with $f_k^{\psi}(\boldsymbol{\theta}; \tau) = f(Q(\mathbf{z}_k^{\psi}(\boldsymbol{\theta}); \tau))$. Then all $\langle F \rangle_{\psi}$ are estimated from such samples via

$$\langle F \rangle_{\psi}(\boldsymbol{\theta}; \tau) \approx \widehat{\text{CVaR}}_1(\{f_1^{\psi}(\boldsymbol{\theta}; \tau), \dots, f_K^{\psi}(\boldsymbol{\theta}; \tau)\}) \tag{15}$$

and equivalently for $\langle F^2 \rangle_{\psi}$. Our implementation of F-VQE adapts the parameter τ dynamically at each optimization step to keep the gradient norm of the objective close to some large, fixed value (see [13] for details).

3 Results and discussion

We have tested the algorithms in Sect. 2.3 on instances of the JSP on IBM quantum processors. First we compared all algorithms on a 5-qubit instance to evaluate their convergence. Then, based on its fast convergence, we selected F-VQE to study the scaling to larger problem sizes. A comparison against classical solvers is not in scope of this work (in fact, all instances can be easily solved exactly). Instead we focus on convergence and scaling the VQA for this particular optimization problem of industrial relevance. All quantum processors were accessed via tket [35]. Hardware experiments benefitted from tket’s out-of-the-box noise-aware qubit placement and routing, but we did not use any other error mitigation techniques involving additional post-processing.

All problem instances for the experiments have been obtained as sub-schedules of the 20-job 2-machine problem whose solution is illustrated in Fig. 1. Table 1 provides information on which machine, time slot and job needed to be assigned a schedule in each of the problem instances.

Throughout this section we plot average energies scaled to the range $[0, 1]$:

$$\epsilon_{\psi} = \frac{\langle H \rangle_{\psi} - E_{\min}}{E_{\max} - E_{\min}} \in [0, 1], \tag{16}$$

where E_{\min}, E_{\max} are the minimum and maximum energy of the Hamiltonian, respectively, and $\langle H \rangle_{\psi} = \langle \psi | H | \psi \rangle$ for a given state $|\psi\rangle$. We calculated E_{\min}, E_{\max} exactly. A value $\epsilon_{\psi} = 0$ corresponds to the optimal solution of the problem. To assess the convergence speed

to good approximation ratios we would like an algorithm to approach values $\epsilon_\psi \approx 0$ in few iterations. We also plot the frequency of sampling the ground state of the problem Hamiltonian $|\psi_{\text{gs}}\rangle$:

$$P_\psi(\text{gs}) = |\langle \psi | \psi_{\text{gs}} \rangle|^2. \quad (17)$$

Ideally, we would like an algorithm to return the ground state with a frequency $P_\psi(\text{gs}) \approx 1$, which implies small average energy $\epsilon_\psi \approx 0$. The converse is not true because a superposition of low-energy excited states $|\psi\rangle$ can exhibit a small average energy $\epsilon_\psi \approx 0$ but small overlap with the ground state $P_\psi(\text{gs}) \approx 0$ [13].

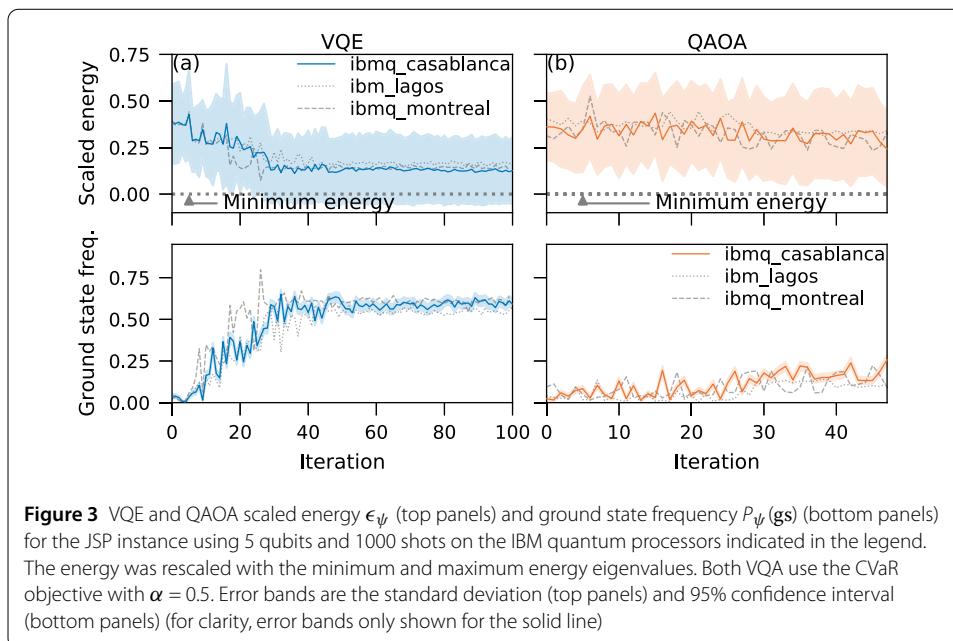
3.1 Performance on 5-variable JSP

We analyzed all algorithms on a JSP instance with 5 free variables requiring 5 qubits. This is sufficiently small to run, essentially, on all available quantum processors. We performed experiments for all VQA on a range of IBM quantum processors. To make the results more comparable, all experiments in this section use the same quantum processors, number of shots, ansatz (VQE, VarQITE, F-VQE) and number of layers for each of the VQA (see Table 2 for all settings). We chose to highlight the results from the `ibmq_casablanca` device in the following plots since it showed the best final ground state frequency for QAOA and good overall performance for VQE and VarQITE. Appendix D presents additional hardware experiments for VQE, QAOA and F-VQE and also VQE and QAOA results for CVaR quantile $\alpha = 0.2$. The goal of these experiments is to analyse the general convergence of the algorithms without much fine-tuning and to select candidate algorithms for the larger experiments in Sect. 3.2.

First, we analyzed VQE. Due to its simplicity it is ideal for initial experimentation. We compared the CVaR objective with $\alpha < 1$ against the standard VQE mean energy objective ($\alpha = 1$). We observed that the CVaR mainly leads to lower variance in the measurement outcomes.

Figure 3(a) shows the results for VQE using CVaR with $\alpha = 0.5$ and 1000 shots and $p = 2$ layers of the ansatz Fig. 2(a). VQE on `ibmq_casablanca` converged after around 40 iterations with a frequency of sampling the ground state of approximately 59%. The frequency of sampling the ground state is approximately bounded by the value α of the CVaR. This is because CVaR optimises the left tail of the empirical distribution up to quantile α . If all the probability mass of the distribution up to quantile α is on the ground state, the cost function achieves its optimal value: the conditional expectation is the ground state energy. At the same time, on average a fraction $1 - \alpha$ of the distribution sits in the right tail of excited states. Results for CVaR with $\alpha = 0.2$ in Fig. 6(b) of App. D are consistent with this observation. All quantum processors showed similar final energies and ground state frequencies for VQE (cf. Fig. 3(a)) with a moderate amount of variance across devices during the initial iterations. Different choices of optimizers could potentially improve convergence rate of VQE [22, 36] but their fine-tuning was not in scope of this study.

QAOA with $p = 2$ showed very slow convergence across all tested quantum processors. The optimizer COBYLA terminated after 47, 50, 48 iterations for `ibmq_casablanca`, `ibmq_lagos` and `ibmq_montreal`, respectively, when it was unable to improve results further. Figure 3(b) shows the scaled energy and ground state frequency with 1000 shots and CVaR $\alpha = 0.5$ (same as VQE). In contrast to VQE, QAOA did not saturate the ground state

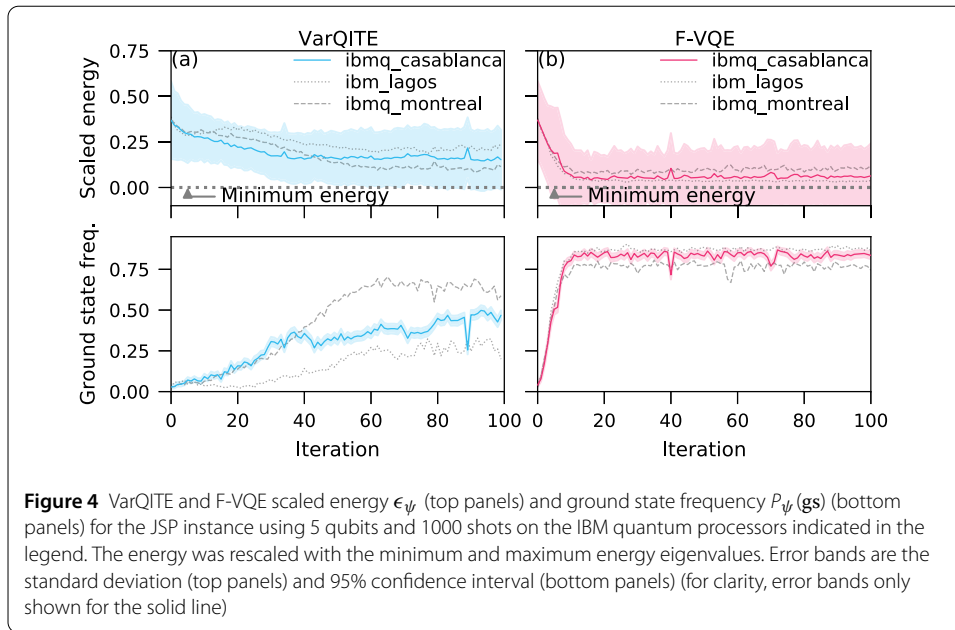


frequency bound at α . We repeated QAOA experiments with CVaR $\alpha = 0.2$ on several quantum processors (see Fig. 7(b)). In this case the ground state frequencies saturated at around $\alpha = 0.2$ but final average energies showed similar performance as the $\alpha = 0.5$ case.

Apart from the optimizer, a contributing factor of this poor performance is likely that the QAOA ansatz is not hardware-efficient, i.e. the compiler needs to add SWAP gates for routing. On `ibmq_casablanca` the compiler embedded the problem on qubits 1-3-4-5-6 (see Fig. 2(b) for the device’s connectivity). In our instance each layer p requires six 2-qubit operations of the form $e^{-i\theta Z_i Z_j}$ each requiring 2 CNOTs. For $p = 2$ layers this is a total of 24 CNOTs to implement the unitaries $U(\gamma_1), U(\gamma_2)$. Routing requires an additional 6 SWAPs, which are implemented with 3 CNOTs each, for a total of 18 CNOTs for routing. In total QAOA required 42 CNOTs. In contrast, the hardware-efficient ansatz Fig. 2(a) for the other VQA can be embedded on a linear chain such as 0-1-3-5-6. This requires no SWAPs and results in a total of 8 CNOTs for our VQE and F-VQE runs. The challenge of scaling QAOA on quantum processors with restricted qubit connectivity was also highlighted in [17] and our results appear to confirm that QAOA running on NISQ hardware requires fine-tuned optimizers even for small-scale instances [23, 24].

VarQITE converged somewhat more gradually compared to VQE but reached similar final mean energies as VQE. Figure 4(a) shows its performance on different quantum processors with 1000 shots and $p = 2$ layers of the ansatz Fig. 2(a). In contrast to VQE, VarQITE exhibited a higher variance of the final mean energy and ground state frequency across different quantum processors. One of the issues of VarQITE is inversion of the matrix A in Eq. (12), which is estimated from measurement shots. This can lead to unstable evolutions. Compared to QAOA, for our problem instance VarQITE converged much faster and smoother across all quantum processors.

F-VQE converged fastest on all quantum processors. Moreover, Fig. 4(b) shows that its convergence is very consistent across devices and the final mean energies are closest to the minimum compared to the other VQA. F-VQE also showed high probability of sampling the optimal solution after just 10-15 iterations, and high final probabilities of 84%, 87%



and 75% after 100 iterations on `ibmq_casablanca`, `ibmq_lagos` and `ibmq_montreal`, respectively. We repeated the F-VQE experiment with a single layer of an ansatz using a linear chain of CNOTs instead of the CNOT pattern in Fig. 2(a) with, essentially, identical results (not shown). This confirms the fast convergence of this algorithm first observed for the weighted MaxCut problem in Ref. [13]. Another advantage of F-VQE compared to VarQITE is that F-VQE does not require inversion of the—typically ill-conditioned—matrix A in Eq. (10), which is estimated from measurement samples. Based on these results we chose to focus on F-VQE for scaling up to larger JSP instances.

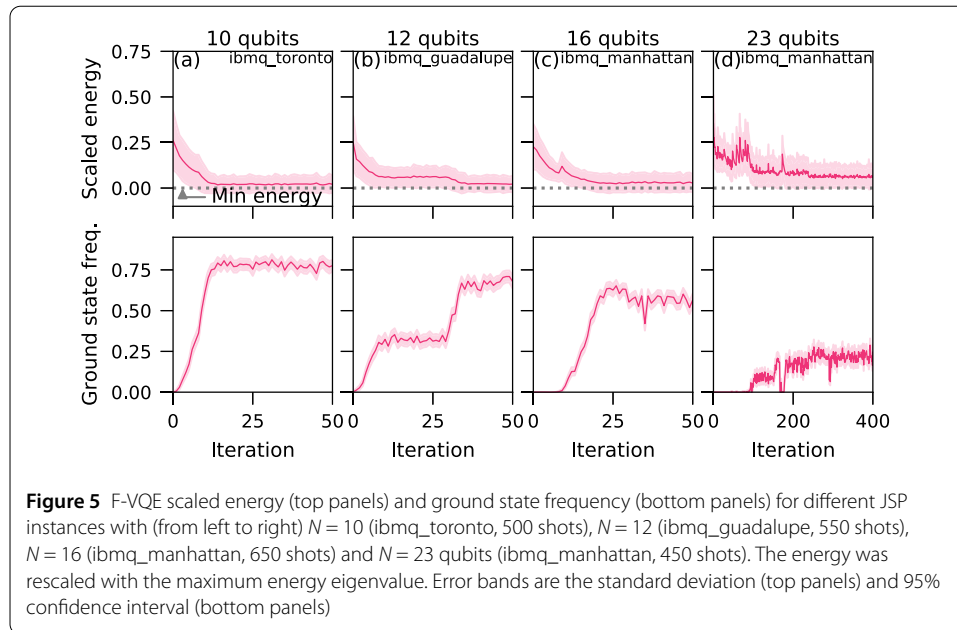
3.2 Performance on larger instances

This section analyzes the effectiveness of F-VQE on larger JSP instances executed on NISQ hardware. Figure 5 summarises the results for up to 23 qubits executed on several IBM quantum processors. For practical reasons (availability, queuing times on the largest device) we ran those experiments on different processors. However, based on the results in Sect. 3 we expect similar performance across different quantum processors. F-VQE converges quickly in all cases. All experiments reach a significant nonzero frequency of sampling the ground state: $P_\psi(\text{gs}) \approx 80\%$ for 10 qubits, $P_\psi(\text{gs}) \approx 70\%$ for 12 qubits, $P_\psi(\text{gs}) \approx 60\%$ for 16 qubits, and $P_\psi(\text{gs}) \approx 25\%$ for 23 qubits.

An interesting case is $N = 12$ (Fig. 5(b)). From iteration 10-30 F-VQE sampled the ground state and one particular excited state with roughly equal probability. However, the algorithm was able to recover the ground state with high probability from iteration 30.

The $N = 23$ results show convergence in terms of the scaled energy and ground state frequency. F-VQE sampled the ground state for the first time after 45 iterations and gradually builds up the probability of sampling it afterwards. This means F-VQE is able to move to a parameter region with high probability of sampling the optimal solution in a computational space of size 2^{23} despite device errors and shot noise.

To our knowledge, the 23-qubit experiment is one of the largest experimental demonstrations of VQA for combinatorial optimization. Otterbach *et al.* [37] demonstrated



QAOA with $p = 1$ on Rigetti's 19-qubit transmon quantum processor. Pagano *et al.* [38] demonstrated the convergence of QAOA ($p = 1$) for up to 20 qubits on a trapped-ion quantum processor. In addition, they present QAOA performance close to optimal parameters with up to 40 qubits without performing the variational parameter optimization. Harrigan *et al.* [17] demonstrated QAOA on Google's superconducting quantum processor Sycamore for up to 23 qubits when the problem and hardware topologies match ($p = 1, \dots, 5$) and up to 22 qubits when the problem and hardware topologies differ ($p = 1, \dots, 3$).

4 Conclusions

In this case study, we solved a combinatorial optimization problem of wide industrial relevance—job shop scheduling—on IBM's superconducting, gate-based quantum processors. Our focus was on the performance of four variational algorithms: the popular VQE and QAOA, as well as the more recent VarQITE and F-VQE. Performance metrics were convergence speed in terms of the number of iterations and the frequency of sampling the optimal solution. We tested these genuinely quantum heuristics using up to 23 physical qubits.

In a first set of experiments we compared all algorithms on a JSP instance with 5 variables (qubits). F-VQE outperformed the other algorithms by all metrics. VarQITE converged slower than F-VQE but was able to sample optimal solutions with comparably high frequency. VQE converged slowly and sampled optimal solutions less frequently. Lastly, QAOA struggled to converge owing to a combination of deeper, more complex circuits and the optimizer choice. QAOA convergence can possibly be improved with a fine-tuned optimizer [24]. In the subsequent set of experiments, we focused on F-VQE as the most promising algorithm and studied its performance on increasingly large problem instances up to 23 variables (qubits). To the best of our knowledge, this is amongst the largest combinatorial optimization problems solved successfully by a variational algorithm on a gate-based quantum processor.

One of the many challenges for variational quantum optimization heuristics is solving larger and more realistic problem instances. It will be crucial to improve convergence of heuristics using more qubits as commercial providers plan a 2- to 4-fold increase of the qubit number on their flagship hardware in the coming years.¹ Our experiments suggest that F-VQE is a step in this direction as it converged quickly even on the larger problems we employed. Another challenge on superconducting quantum processors with hundreds of qubits is sparse connectivity and cross-talk noise. F-VQE can address this concern with ansätze that are independent of the problem's connectivity and that can be embedded in a quantum processor's topology with lower or even zero SWAP gate overhead from routing. In addition, error mitigation post processing can be utilized [39], although recent results show that this requires careful analysis as these techniques can either improve or hinder trainability of VQA [40]. Trapped-ion quantum hardware may be soon equipped with dozens of qubits. Their low noise levels and all-to-all qubit connectivity should be more suitable for deeper and more complex ansätze. Hence, trapped-ion quantum processors may benefit from the combination of F-VQE with *causal cones* [13]. Causal cones can split the evaluation of the cost function into batches of circuits with fewer qubits [41]. This allows quantum computers to tackle combinatorial optimization problems with more variables than their physical qubits and parallelize the workload.

The combination of the results of this case study together with the aforementioned algorithmic and hardware improvements paint the optimistic picture that near term quantum computers may be able to tackle combinatorial optimization problems with hundreds of variables in the coming years.

Appendix A: Derivation of the QUBO formulation of the JSP

This appendix describes the derivation of the QUBO formulation of the JSP in Eq. (2).

The cost of a schedule comprises three parts: the early delivery cost, late delivery cost and production cost. The early and late delivery costs are a penalty added when a job j passes the last machine M before or after its *due time* d_j , respectively:

$$u_j(\mathbf{x}) = c_e \sum_{t=1}^{d_j} (d_j - t)x_{Mjt} + c_l \sum_{t=d_j+1}^{T_M} (t - d_j)x_{Mjt} \quad \forall j = 1, \dots, J. \quad (18)$$

The constants c_e and c_l determine the magnitude of the early and late delivery cost, respectively. Figure 1 illustrates the 20-job instance used in our results together with its optimal schedule.

The production cost is a penalty added for production group switches of two jobs entering a machine at subsequent time slots. The production group of job j for machine m is determined by a matrix with entries P_{mj} . For each machine m we define a matrix $G^{(m)}$ with entries

$$G_{j_1 j_2}^{(m)} = \begin{cases} 0 & \text{if } P_{mj_1} = P_{mj_2}, \\ 1 & \text{otherwise.} \end{cases} \quad (19)$$

¹See development roadmaps by IBM <https://research.ibm.com/blog/ibm-quantum-roadmap> and Quantinuum (formerly Honeywell Quantum Solutions) <https://www.honeywell.com/us/en/news/2020/10/get-to-know-honeywell-s-latest-quantum-computer-system-model-h1>, for instance (accessed on 2022-02-04).

Hence, the production cost for machine m is given by

$$s_m(\mathbf{x}) = c_p \sum_{j_1, j_2=1}^J \sum_{t=1}^{T_m-1} G_{j_1 j_2}^{(m)} x_{mj_1 t} x_{mj_2(t+1)} \quad \forall m = 1, \dots, M. \tag{20}$$

The constant c_p determines the magnitude of the production cost.

The total cost of a schedule \mathbf{x} is

$$c(\mathbf{x}) = \sum_{j=1}^J u_j(\mathbf{x}) + \sum_{m=1}^M s_m(\mathbf{x}). \tag{21}$$

We model the constraints of the JSP as additional cost functions. The *job assignment constraints* enforces that each real job is assigned to exactly one time slot in each machine

$$g_{mj}(\mathbf{x}) \equiv \sum_{t=1}^{T_m} x_{mjt} = 1, \quad \forall m = 1, \dots, M \forall j = 1, \dots, J. \tag{22}$$

The *time assignment constraints* ensure that each time slot in each machine is occupied by exactly one job:

$$\ell_{mt}(\mathbf{x}, \mathbf{y}) = \left\{ \begin{array}{ll} y_{mt} + \sum_{i=1}^J x_{mjt} & \text{for } 1 \leq t \leq e_m \\ \sum_{j=1}^J x_{mjt} & \text{for } e_m < t \leq J \\ 1 - y_{m(t-J)} + \sum_{j=1}^J x_{mjt} & \text{for } J < t \leq T_m \end{array} \right\} = 1 \quad \begin{array}{l} \forall t = 1, \dots, T_m, \\ \forall m = 1, \dots, M. \end{array} \tag{23}$$

The *process order constraints* ensure that the processing time of a real job does not decrease from one machine to the next:

$$q_{mj}(\mathbf{x}) = \sum_{t=2}^{T_m} \sum_{t'=1}^{t-1} x_{mjt} x_{(m+1)jt'} = 0 \quad \forall m = 1, \dots, M-1 \forall j = 1, \dots, J. \tag{24}$$

The *idle slot constraints* ensure that dummy jobs are placed before all real jobs in each machine. Due to constraints ℓ_{mt} in Eq. (23) we only need to enforce that the transition from a real job to a dummy job is prohibited at the beginning of a schedule:

$$r_{mt}(\mathbf{y}) = (1 - y_{mt})y_{m(t+1)} = 0 \quad \forall t = 1, \dots, e_m - 1 \forall m = 2, \dots, M. \tag{25}$$

Note that constraints of this form are not required for machines with $e_m = 1$.

Appendix B: Worst-case scaling of the JSP

The total number of variables in the JSP formulation of Sect. 2.2 is

$$N = \sum_{m=1}^M J(J + e_m) + e_m. \tag{26}$$

The best-case scaling $\mathcal{O}(J^2M)$ is achieved for fixed e_m . In the worst case the number of dummy jobs needs to grow by $J - 1$ per machine to allow for a complete reordering of all

Table 3 Hardware devices used in this study

Device	No. of qubits	Quantum volume	Connectivity
ibmq_belem	5	16	T-shaped
ibmq_casablanca	7	32	Fig. 2(b)
ibmq_guadalupe	16	32	Heavy-hexagon [42]
ibmq_jakarta	7	16	Fig. 2(b)
ibmq_lagos	7	32	Fig. 2(b)
ibmq_manhattan	65	32	Heavy-hexagon
ibmq_montreal	27	128	Heavy-hexagon
ibmq_sydney	27	32	Heavy-hexagon
ibmq_toronto	27	32	Heavy-hexagon

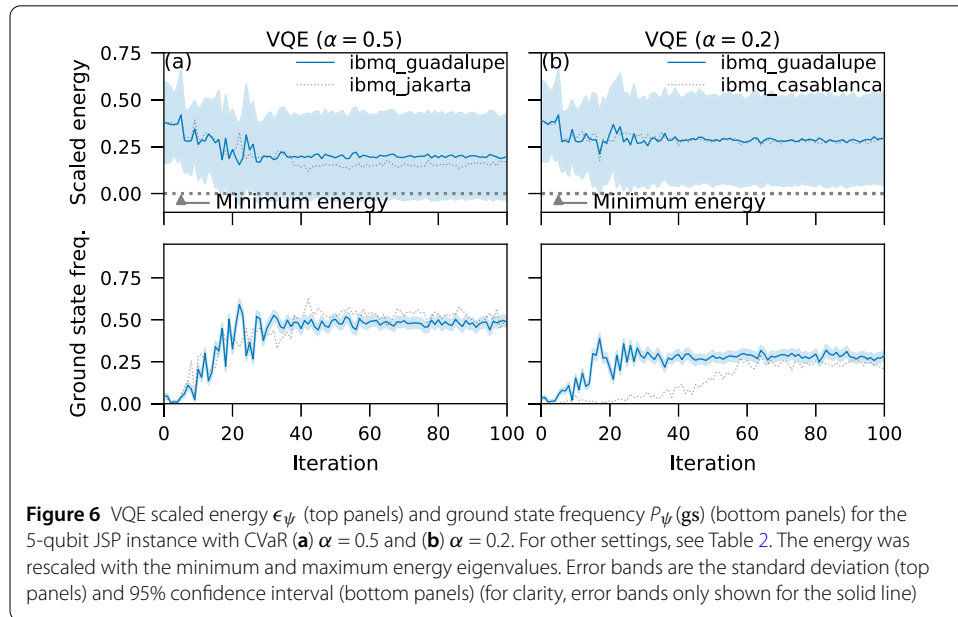
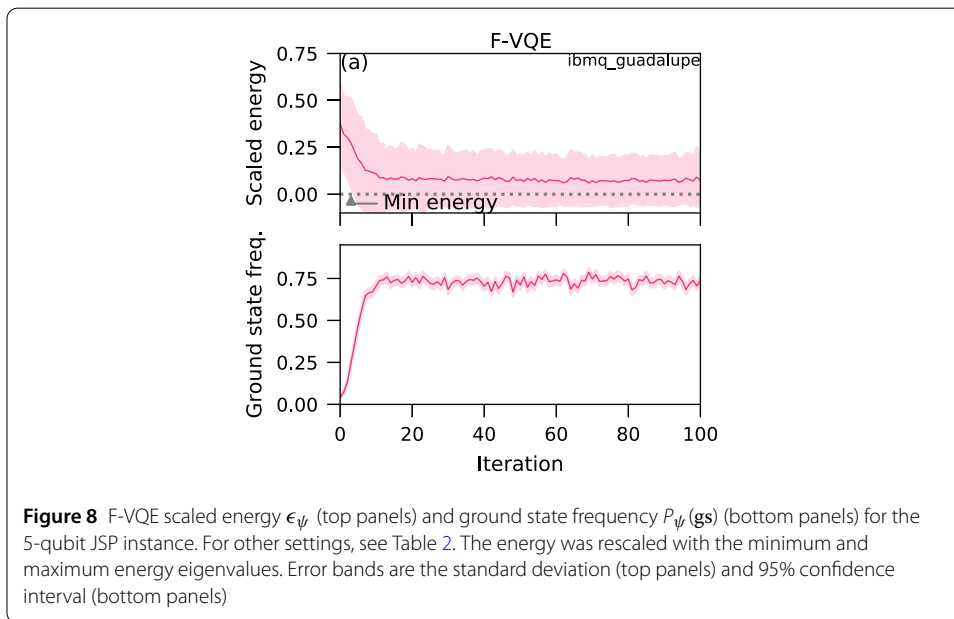
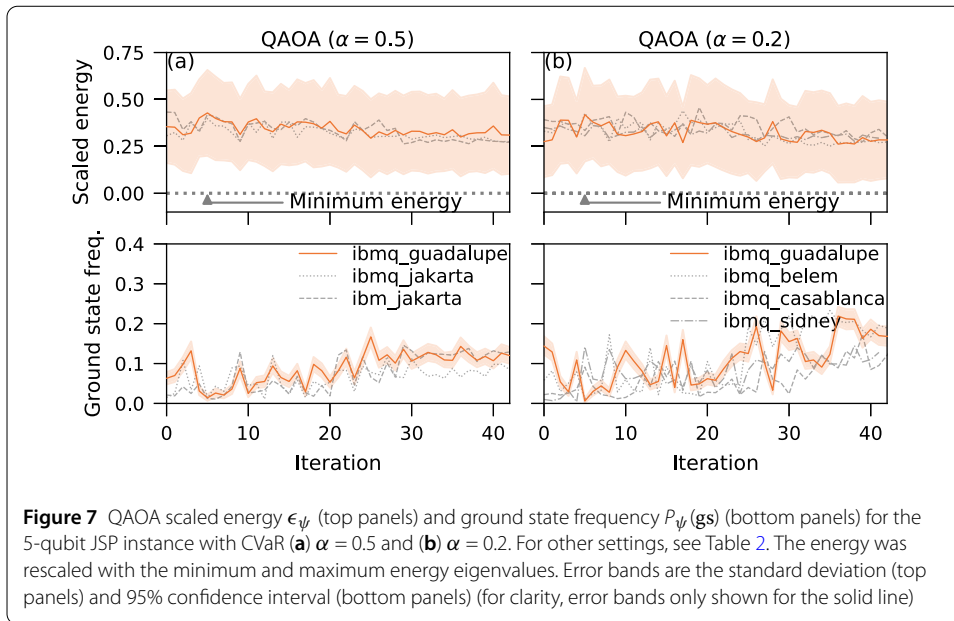


Figure 6 VQE scaled energy ϵ_ψ (top panels) and ground state frequency $P_\psi(\text{gs})$ (bottom panels) for the 5-qubit JSP instance with CVaR **(a)** $\alpha = 0.5$ and **(b)** $\alpha = 0.2$. For other settings, see Table 2. The energy was rescaled with the minimum and maximum energy eigenvalues. Error bands are the standard deviation (top panels) and 95% confidence interval (bottom panels) (for clarity, error bands only shown for the solid line)

jobs. With the convention that $e_1 = 0$ this leads to $e_m = (m - 1)(J - 1)$ and the worst-case scaling $\mathcal{O}(J^2M^2)$.

Note that the dummy variables y_{m1} can be dropped from the problem for every machine with $e_m = 1$. For $e_m = 1$ the constraints $\ell_{m1}(\mathbf{x}, \mathbf{y})$ and $\ell_{m(J+1)}(\mathbf{x}, \mathbf{y})$ are automatically satisfied given the rest of constraints for $e_m = 1$. First, from the rest of constraints $\ell_{mt}(\mathbf{x}, \mathbf{y})$ the $J - 1$ time slots $t = 2, \dots, J$ contain one job. Second, from the constraints $g_{mj}(\mathbf{x})$ there are J jobs. Therefore, exactly one job is placed either in the time $t = 1$ or in the time $t = J + 1$ without the need of forcing the constraints $\ell_{m1}(\mathbf{x}, \mathbf{y})$ and $\ell_{m(J+1)}(\mathbf{x}, \mathbf{y})$.

It is possible to cut down the worst-case scaling to $\mathcal{O}(J^2M)$ with an alternative formulation of the JSP. This alternative uses a binary encoding for the e_m dummy jobs. However, in this work we focused on fixed e_m for all instances, which leads to the same scaling. Furthermore, we fix most of the time slots to the optimal solution and only leave the positions of a few jobs free. This way we can systematically increase problem sizes and analyse scaling of the algorithms.



Appendix C: Quantum hardware

Table 3 lists the quantum processors used in this work and some of their basic properties at the time of execution. More information is available at <https://quantum-computing.ibm.com/services>.

Appendix D: Additional experiments

Figure 6 shows results of additional hardware experiments for VQE with CVaR quantiles $\alpha = 0.5$ (Fig. 6(a)) and $\alpha = 0.2$ (Fig. 6(b)) for the 5-qubit JSP instance discussed in Sect. 3.1. For all other parameters see Table 2. In both cases VQE reaches a ground state frequency of approximately α indicating that the CVaR objective was achieved. Generally, the $\alpha = 0.2$

case converged to a mean energy considerably further from the optimal value than for $\alpha = 0.5$.

Figure 7 shows results of additional hardware experiments for QAOA with CVaR quantiles $\alpha = 0.5$ (Fig. 7(a)) and $\alpha = 0.2$ (Fig. 7(b)) for the 5-qubit JSP instance discussed in Sect. 3.1. For all other parameters see Table 2. QAOA with $\alpha = 0.2$ reaches a ground state frequency of approximately α indicating that the CVaR objective was achieved in this case.

Figure 8 shows results of one additional hardware experiment for F-VQE on ibmq_guadalupe for the 5-qubit JSP instance discussed in Sect. 3.1. For all other parameters see Table 2. The overall performance is comparable to its performance on other quantum processors in Sect. 3.1.

Acknowledgements

The authors would like to thank Carlo Modica for helping with the execution of some experiments on quantum hardware, and Michael Lubasch and Marcello Benedetti for the helpful discussions.

Funding

Not applicable.

Abbreviations

COBYLA, Constrained Optimization By Linear Approximation; CVaR, Conditional Value-at-Risk; F-VQE, Filtering Variational Quantum Eigensolver; JSP, Job Shop Scheduling problem; NISQ, Noisy Intermediate-Scale Quantum; PQC, Parameterized Quantum Circuit; QAOA, Quantum Approximate Optimization Algorithm; QUBO, Quadratic Unconstrained Binary Optimization; VarQITE, Variational Quantum Imaginary Time Evolution; VQA, Variational Quantum Algorithms; VQE, Variational Quantum Eigensolver.

Availability of data and materials

The datasets used and/or analysed during the current study are available from the corresponding author on reasonable request.

Declarations

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

All authors contributed to the drafting of the manuscript. DA, MF and MR designed the work and experiments. KH designed the JSP instance analyzed in this work. DA acquired the data. DA, MF and MR interpreted and analysed the data. DA and NF created the software for this work. All authors read and approved the final manuscript.

Author details

¹Cambridge Quantum Computing Limited, SW1P 1BX, London, UK. ²Cambridge Quantum Computing Limited, Terrington House 13-15 Hills Road, CB2 1NL, Cambridge, UK. ³Nippon Steel Corporation, 20-1 Shintomi, 293-8511, Futtsu, Chiba, Japan.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 13 August 2021 Accepted: 28 January 2022 Published online: 10 February 2022

References

1. Grover LK. A fast quantum mechanical algorithm for database search. In: Proceedings of the twenty-eighth annual ACM symposium on theory of computing. STOC'96. New York: Association for Computing Machinery; 1996. p. 212–9. [quant-ph/9605043](https://doi.org/10.1145/237814.237866). <https://doi.org/10.1145/237814.237866>.
2. Durr C, Hoyer P. A Quantum Algorithm for Finding the Minimum. 1999. [quant-ph/9607014](https://arxiv.org/abs/quant-ph/9607014).
3. Preskill J. Quantum Computing in the NISQ era and beyond. *Quantum*. 2018;2:79. <https://doi.org/10.22331/q-2018-08-06-79>.
4. Cerezo M, Arrasmith A, Babbush R, Benjamin SC, Endo S, Fujii K, McClean JR, Mitarai K, Yuan X, Cincio L, Coles PJ. Variational quantum algorithms. *Nat Rev Phys*. 2021;3(9):625–44. <https://doi.org/10.1038/s42254-021-00348-9>.
5. Bharti K, Cervera-Lierta A, Kyaw TH, Haug T, Alperin-Lea S, Anand A, Degroote M, Heimonen H, Kottmann JS, Menke T, Mok W-K, Sim S, Kwok L-C, Aspuru-Guzik A. Noisy intermediate-scale quantum (NISQ) algorithms. 2021. [2101.08448](https://arxiv.org/abs/2101.08448).
6. Sanders YR, Berry DW, Costa PCS, Tessler LW, Wiebe N, Gidney C, Neven H, Babbush R. Compilation of fault-tolerant quantum heuristics for combinatorial optimization. *PRX Quantum*. 2020;1(2):020312. <https://doi.org/10.1103/PRXQuantum.1.020312>.

7. Kochenberger G, Hao J-K, Glover F, Lewis M, Lü Z, Wang H, Wang Y. The unconstrained binary quadratic programming problem: a survey. *J Comb Optim.* 2014;28(1):58–81. <https://doi.org/10.1007/s10878-014-9734-0>.
8. Lucas A. Ising formulations of many NP problems. *Front Phys.* 2014;2:5. <https://doi.org/10.3389/fphy.2014.00005.1302.5843>.
9. Glover F, Kochenberger G, Du Y. Quantum Bridge Analytics I: a tutorial on formulating and using QUBO models. *4OR-Q. J Oper Res.* 2019;17(4):335–71. <https://doi.org/10.1007/s10288-019-00424-y>.
10. Farhi E, Goldstone J, Gutmann S. A Quantum Approximate Optimization Algorithm. 2014. 1411.4028.
11. Peruzzo A, McClean J, Shadbolt P, Yung M-H, Zhou X-Q, Love PJ, Aspuru-Guzik A, O'Brien JL. A variational eigenvalue solver on a photonic quantum processor. *Nat Commun.* 2014;5(1):4213. <https://doi.org/10.1038/ncomms5213>.
12. McArdle S, Jones T, Endo S, Li Y, Benjamin SC, Yuan X. Variational ansatz-based quantum simulation of imaginary time evolution. *npj Quantum Inf.* 2019;5(1):75. <https://doi.org/10.1038/s41534-019-0187-2>.
13. Amaro D, Modica C, Rosenkranz M, Fiorentini M, Benedetti M, Lubasch M. Filtering variational quantum algorithms for combinatorial optimization. *Quantum Sci Technol.* 2022, to appear. 2106.10055. <https://doi.org/10.1088/2058-9565/ac3e54>.
14. Farhi E, Harrow AW. Quantum supremacy through the quantum approximate optimization algorithm. 2016. 1602.07674.
15. Zhou L, Wang S-T, Choi S, Pichler H, Lukin MD. Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices. *Phys Rev X.* 2020;10:021067. <https://doi.org/10.1103/PhysRevX.10.021067>.
16. Moussa C, Calandra H, Dunjko V. To quantum or not to quantum: towards algorithm selection in near-term quantum optimization. *Quantum Sci Technol.* 2020;5(4):044009. <https://doi.org/10.1088/2058-9565/abb8e5.2001.08271>.
17. Harrigan MP, Sung KJ, Neeley M, Satzinger KJ, Arute F, Arya K, Atalaya J, Bardin JC, Barends R, Boixo S, Broughton M, Buckley BB, Buell DA, Burkett B, Bushnell N, Chen Y, Chen Z, Chiaro Collins RB, Courtney W, Demura S, Dunswoth A, Eppens D, Fowler A, Foxen B, Gidney C, Giustina M, Graff R, Habegger S, Ho A, Hong S, Huang T, Ioffe LB, Isakov SV, Jeffrey E, Jiang Z, Jones C, Kafri D, Kechedzhi K, Kelly J, Kim S, Klimov PV, Korotkov AN, Kostritsa F, Landhuis D, Laptev P, Lindmark M, Leib M, Martin O, Martinis JM, McClean JR, McEwen M, Megrant A, Mi X, Mohseni M, Mruzckiewicz W, Mutus J, Naaman O, Neill C, Neukart F, Niu MY, O'Brien TE, O'Gorman B, Ostby E, Petukhov A, Putterman H, Quintana C, Roushan P, Rubin NC, Sank D, Skolik A, Smelyanskiy V, Strain D, Streif M, Szalay M, Vainsencher A, White T, Yao ZJ, Yeh P, Zalcman A, Zhou L, Neven H, Bacon D, Lucero E, Farhi E, Babbush R. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nat Phys.* 2021;17(3):332–6. <https://doi.org/10.1038/s41567-020-01105-y>.
18. Sim S, Johnson PD, Aspuru-Guzik A. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Adv Quantum Technol.* 2019;2(12):1900070. <https://doi.org/10.1002/qute.201900070>. 1905.10876.
19. McClean JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H. Barren plateaus in quantum neural network training landscapes. *Nat Commun.* 2018;9(1):4812. <https://doi.org/10.1038/s41467-018-07090-4>.
20. Cerezo M, Sone A, Volkoff T, Cincio L, Coles PJ. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat Commun.* 2021;12(1):1791. <https://doi.org/10.1038/s41467-021-21728-w>.
21. Guerreschi GG, Smelyanskiy M. Practical optimization for hybrid quantum-classical algorithms. 2017. 1701.01450.
22. Nannicini G. Performance of hybrid quantum-classical variational heuristics for combinatorial optimization. *Phys Rev E.* 2019;99(1):013304. <https://doi.org/10.1103/PhysRevE.99.013304>. 1805.12037.
23. Lavrijsen W, Tudor A, Müller J, Iancu C, de Jong W. Classical optimizers for noisy intermediate-scale quantum devices. In: 2020 IEEE international conference on quantum computing and engineering (QCE). 2020. p. 267–77. <https://doi.org/10.1109/QCE49297.2020.00041>. 2004.03004.
24. Sung KJ, Yao J, Harrigan MP, Rubin NC, Jiang Z, Lin L, Babbush R, McClean JR. Using models to improve optimizers for variational quantum algorithms. *Quantum Sci Technol.* 2020;5(4):044008. <https://doi.org/10.1088/2058-9565/abb6d9.2005.11011>.
25. Pellow-Jarman A, Sinayskiy I, Pillay A, Petruccione F. A comparison of various classical optimizers for a variational quantum linear solver. *Quantum Inf Process.* 2021;20(6):202. <https://doi.org/10.1007/s11128-021-03140-x>. 2106.08682.
26. Pinedo ML. *Scheduling: theory, algorithms, and systems.* Boston: Springer; 2012.
27. Barkoutsos PK, Nannicini G, Robert A, Tavernelli I, Woerner S. Improving Variational Quantum Optimization using CVaR. *Quantum.* 2020;4:256. <https://doi.org/10.22331/q-2020-04-20-256>.
28. Kandala A, Mezzacapo A, Temme K, Takita M, Brink M, Chow JM, Gambetta JM. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature.* 2017;549(7671):242–6. <https://doi.org/10.1038/nature23879>.
29. Powell MJD. A direct search optimization method that models the objective and constraint functions by linear interpolation. In: Gomez S, Hennart J-P, editors. *Advances in optimization and numerical analysis.* Netherlands: Springer; 1994. p. 51–67. https://doi.org/10.1007/978-94-015-8330-5_4.
30. Farhi E, Goldstone J, Gutmann S, Sipser M. Quantum Computation by Adiabatic Evolution. 2000. [quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
31. Yuan X, Endo S, Zhao Q, Li Y, Benjamin SC. Theory of variational quantum simulation. *Quantum.* 2019;3:191. <https://doi.org/10.22331/q-2019-10-07-191>.
32. Motta M, Sun C, Tan ATK, O'Rourke MJ, Ye E, Minnich AJ, Brandão FGSL, Chan GK-L. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nat Phys.* 2020;16(2):205–10. <https://doi.org/10.1038/s41567-019-0704-4>.
33. Schuld M, Bergholm V, Gogolin C, Izaac J, Killoran N. Evaluating analytic gradients on quantum hardware. *Phys Rev A.* 2019;99(3):032331. <https://doi.org/10.1103/PhysRevA.99.032331>. 1811.11184.
34. Mitarai K, Negoro M, Kitagawa M, Fujii K. Quantum circuit learning. *Phys Rev A.* 2018;98(3):032309. <https://doi.org/10.1103/PhysRevA.98.032309>. 1803.00745.
35. Sivarajah S, Dilkes S, Cowtan A, Simmons W, Edgington A, Duncan R. t|ket>: a retargetable compiler for NISQ devices. *Quantum Sci Technol.* 2020;6(1):014003. <https://doi.org/10.1088/2058-9565/ab8e92.2003.10611>.
36. Stokes J, Izaac J, Killoran N, Carleo G. Quantum natural gradient. *Quantum.* 2020;4:269. <https://doi.org/10.22331/q-2020-05-25-269>.

37. Otterbach JS, Manenti R, Alidoust N, Bestwick A, Block M, Bloom B, Caldwell S, Didier N, Fried ES, Hong S, Karalekas P, Osborn CB, Papageorge A, Peterson EC, Prawiroatmodjo G, Rubin N, Ryan CA, Scarabelli D, Scheer M, Sete EA, Sivarajah P, Smith RS, Staley A, Tezak N, Zeng WJ, Hudson A, Johnson BR, Reagor M, da Silva MP, Rigetti C. Unsupervised machine learning on a hybrid quantum computer. 2017. [1712.05771](https://doi.org/10.1073/pnas.2006373117).
38. Pagano G, Bapat A, Becker P, Collins KS, De A, Hess PW, Kaplan HB, Kyprianidis A, Tan WL, Baldwin C, Brady LT, Deshpande A, Liu F, Jordan S, Gorshkov AV, Monroe C. Quantum approximate optimization of the long-range Ising model with a trapped-ion quantum simulator. *Proc Natl Acad Sci*. 2020;117(41):25396–401. <https://doi.org/10.1073/pnas.2006373117>. [1906.02700](https://doi.org/10.1073/pnas.2006373117).
39. Endo S, Cai Z, Benjamin SC, Yuan X. Hybrid quantum-classical algorithms and quantum error mitigation. *J Phys Soc Jpn*. 2021;90(3):032001. <https://doi.org/10.7566/JPSJ.90.032001>. [2011.01382](https://doi.org/10.7566/JPSJ.90.032001).
40. Wang S, Czarnik P, Arrasmith A, Cerezo M, Cincio L, Coles PJ. Can Error Mitigation Improve Trainability of Noisy Variational Quantum Algorithms? 2021. [2109.01051](https://doi.org/10.1103/PhysRevResearch.3.033083).
41. Benedetti M, Fiorentini M, Lubasch M. Hardware-efficient variational quantum algorithms for time evolution. *Phys Rev Res*. 2021;3(3):033083. <https://doi.org/10.1103/PhysRevResearch.3.033083>.
42. Chamberland C, Zhu G, Yoder TJ, Hertzberg JB, Cross AW. Topological and subsystem codes on low-degree graphs with flag qubits. *Phys Rev X*. 2020;10(1):011022. <https://doi.org/10.1103/PhysRevX.10.011022>.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ [springeropen.com](https://www.springeropen.com)
