
Cambridge Quantum

A Case Study of Variational Quantum Algorithms for a Job Scheduling Program

*by David Amaro, Matthias Rosenkranz,
Nathan Fitzpatrick, Mattia Fiorentini, Koji Hirano*

Cambridge Quantum
DAVID AMARO
david.amaro@cambridgequantum.com
MATTHIAS ROSENKRANZ
matthias.rosenkranz@cambridgequantum.com
NATHAN FITZPATRICK
nathan.fitzpatrick@cambridgequantum.com
MATTIA FIORENTINI
mattia.fiorentini@cambridgequantum.com

Nippon Steel Corporation, Futtsu, Chiba, Japan
KOJI HIRANO
hirano.3c6.kohji@jp.nipponsteel.com

Cambridge Quantum
Terrington House
13-15 Hills Road
Cambridge CB2 1NL
United Kingdom

Published by Cambridge Quantum

09 September 2021

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

David Amaro,^{1,*} Matthias Rosenkranz,^{1,†} Nathan Fitzpatrick,² Koji Hirano,^{3,‡} and Mattia Fiorentini¹

¹*Cambridge Quantum Computing Limited, SW1E 6DR London, United Kingdom*

²*Cambridge Quantum Computing Limited, Terrington House 13-15,
Hills Road, Cambridge, CB2 1NL, United Kingdom*

³*Nippon Steel Corporation, 20-1 Shintomi, Futtsu, Chiba, 293-8511, Japan*

(Dated: September 9, 2021)

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. Combining F-VQE with error mitigation and causal cones could allow quantum optimization heuristics to scale to relevant problem sizes.

I. 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 algorithms 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

* david.amaro@cambridgequantum.com

† matthias.rosenkranz@cambridgequantum.com

‡ hirano.3c6.kohji@jp.nipponsteel.com

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. Sec. II introduces the JSP formulation and the four VQA employed in this work, highlighting their similarities and differences. Sec. III analyses the performance of all VQA and shows results of scaling up F-VQE on hardware. We conclude in Sec. IV. The appendix includes a derivation of the JSP formulation, discusses the scaling of the JSP and lists key properties of the quantum processors used for this work.

II. METHODS

This section introduces the mathematical formulation of the JSP in Sec. IIB and introduces the VQE, QAOA, VarQITE and F-VQE with our choices for the various settings of these algorithms in Sec. IIC.

A. Job shop scheduling in a steel manufacturing process

The JSP formulation we consider applies to general manufacturing processes and was fine-tuned by Nippon Steel Corporation for steel manufacturing. Job shop scheduling assigns a number of jobs $j = 1, \dots, J$ to different machines or processes $m = 1, \dots, M$ at time slots $t = 1, \dots, T_m$. The goal of the JSP is to find an assignment—also called a schedule—with minimal cost given a number of constraints. Each machine can idle for a number of time slots at the beginning or end of the schedule. The number of idle slots for each machine m is a nonnegative integer e_m provided as an input. Hence, the maximum time slot for machine m is $T_m = J + e_m$.

The cost of a schedule is determined by the cost of *early delivery* and *late delivery* of jobs from the last machine, and the *production cost* associated with changing the processing conditions for subsequent jobs in each machine. We consider the following sets of constraints, which follow from the specifics of this 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.

B. Quadratic unconstrained binary optimization formulation of the JSP

We formulate the JSP defined in subsection IIA 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 jobs 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. Sometimes we will call the jobs in \mathbf{x} *real jobs* to distinguish them from dummy jobs. A value $x_{mjt} = 1$ ($x_{mjt} = 0$) indicates that job j is assigned (is not assigned) to machine m at time t , and. 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

$$(\mathbf{x}^*, \mathbf{y}^*) = \underset{(\mathbf{x}, \mathbf{y}) \in \mathbb{B}^{N_x} \times \mathbb{B}^{N_y}}{\operatorname{arg\,min}} Q(\mathbf{x}, \mathbf{y}). \quad (1)$$

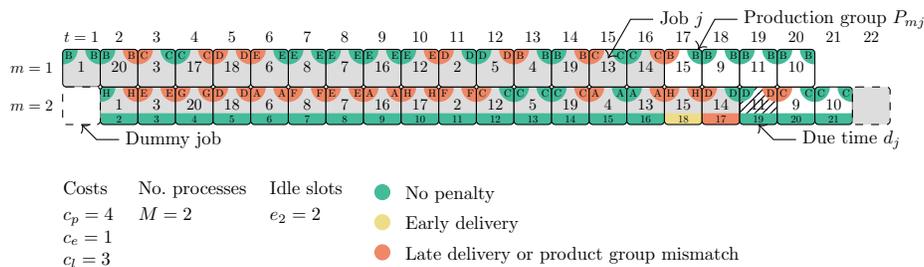


FIG. 1. 20-job, 2-process JSP instance considered in this work and its optimal solutions. Colors at the bottom and in the corners of each box indicate which penalty applies for each time slot and machine. By fixing some jobs to their optimal slots we generate instances with different numbers of free variables N . This is indicated by the background color/pattern of a box: grey for fixed slots and jobs, white for free slots and jobs, and dashes for free slots but fixed jobs. We generated instances with $N = 5, 10, 12, 16, 23$ free variables (see Tab. I). The figure shows $N = 23$.

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

TABLE I. Time slots and jobs needing assignment in each of the problem instances considered in this work.

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

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

Here, $c(\mathbf{x})$ is the cost of the schedule and the last four terms correspond to the job assignment constraints, time assignment constraints, process order constraints and idle slot constraints, respectively. All terms are derived in App. A. 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 $Q(\mathbf{z}) = Q(\mathbf{x}, \mathbf{y})$. Fig. 1 illustrates the largest JSP instance used in this work together with its optimal solution obtained via a classical solver, and Tab. I specifies all instances used.

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, and \mathbf{Z} corresponds to \mathbf{z} and $h_0, h_n, h_{nn'}$ are the coefficients of the corresponding operators. Note that this is a classical Hamiltonian, which means that its eigenstates are separable and they are eigenstates of \mathbf{Z} .

C. 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

Algorithm	Ansatz	Initial param.	Objective	Optimizer
VQE	Fig. 2(a)	$ +\rangle$	Eq. (6)	COBYLA
QAOA	Eq. 7(a) ($p = 2$)	uniform in $[0, \pi]$	Eq. (8)	COBYLA
VarQITE	Fig. 2(a)	$ +\rangle$	Eq. (9)	Eq. (12)
F-VQE	Fig. 2(a) (5 qubits), (b) (> 5 qubits)	$ +\rangle$	Eq. (13)	Eq. (14)

TABLE II. VQA and settings used in this work. Initial param. $|+\rangle$ means that the initial angles of all R_y in the first (second) layer are set to 0 ($\pi/2$).

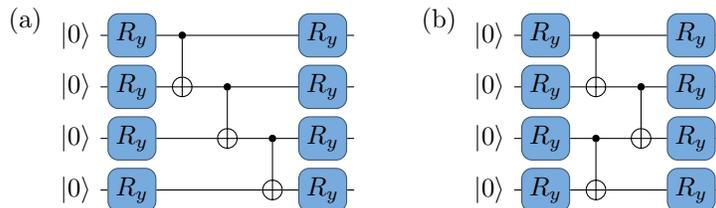


FIG. 2. Parameterized quantum circuit ansatz $|\psi(\boldsymbol{\theta})\rangle$ used for the (a) 5-qubit VQE, VarQITE and F-VQE results and (b) larger F-VQE instances. Each rotation gate is parameterized by an angle θ_n .

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 subsection II B, the JSP is equivalent to finding the ground state of a Hamiltonian. 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 [26]. 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 $\alpha \in (0, 1]$:

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

In practice we estimate the CVaR from measurement samples as follows. Consider a state $|\psi\rangle$ and sample K classical bitstrings $\mathbf{z}_k \sim \|\langle \mathbf{z} | \psi \rangle\|^2$ by measuring this state repeatedly in the computational basis. 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

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

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 [26] and some of our experiments confirmed this behaviour.

The difference between the considered VQA boils down to different choices of the ansatz, measurement scheme, objective and optimizer. Table II summarizes and subsections II C 1–II C 4 detail our choices. Appendix C lists the quantum processors used for the hardware execution.

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 combinatorial optimization is equivalent to finding a ground state of a Hamiltonian, 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 [27]. We denote the free parameters of the single-qubit rotation gates in the ansatz with the vector $\boldsymbol{\theta}$. The PQC implements the unitary operator $U(\boldsymbol{\theta})$ and $|\psi(\boldsymbol{\theta})\rangle = U(\boldsymbol{\theta})|0\rangle$ 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 $\mathbf{z}_k(\boldsymbol{\theta}) \sim \|\langle \mathbf{z} | \psi(\boldsymbol{\theta}) \rangle\|^2$ sampled from the PQC. Sampling K bitstrings and calculating their energies $E_k(\boldsymbol{\theta}) = Q(\mathbf{z}_k(\boldsymbol{\theta}))$ yields a sample of (ordered) energies $\{E_1(\boldsymbol{\theta}), \dots, E_K(\boldsymbol{\theta})\}$ parameterized by $\boldsymbol{\theta}$. Plugging this sample into the CVaR estimator, Eq. (4), yields the objective function for VQE

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

We use the Constrained Optimization By Linear Approximation (COBYLA) optimizer to tune the parameters of the PQC [28]. This is a gradient-free optimizer with few hyperparameters making it a reasonable baseline choice for VQA [23].

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 [29].

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 $(\boldsymbol{\beta}, \boldsymbol{\gamma})$ is

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

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 $|\mathbf{z}_k(\boldsymbol{\beta}, \boldsymbol{\gamma})\rangle$ from the PQC and calculate their energies $E_k(\boldsymbol{\beta}, \boldsymbol{\gamma}) = Q(\mathbf{z}_k(\boldsymbol{\beta}, \boldsymbol{\gamma}))$. The objective function is the CVaR estimator

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

and the optimizer is COBYLA.

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 [30] and a variant of the algorithm shows promising results for combinatorial optimization [31]. 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 $|\psi(\boldsymbol{\theta})\rangle = |\psi(\boldsymbol{\theta}(\tau))\rangle = U(\boldsymbol{\theta})|+\rangle$ and the parameters are assumed to be time-dependent $\boldsymbol{\theta} = \boldsymbol{\theta}(\tau)$. 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 $\alpha = 1$,

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

VarQITE updates parameters with a gradient-based optimization scheme derived from McLachlan's variational principle [30]. 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

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

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

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

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

$$\boldsymbol{\theta}_{n+1} = \boldsymbol{\theta}_n - A^{-1}(\boldsymbol{\theta}_n) \nabla O_{\text{VarQITE}}(\boldsymbol{\theta}_n) \delta\tau. \quad (12)$$

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

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 experimented with both ansätze in Fig. 2. In contrast to our VQE implementation, F-VQE uses a gradient-based optimizer. At each optimization step n the objective function is

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

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

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

where η is a learning rate. The gradient in Eq. (14) is computed with the parameter shift rule [32, 33]. 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

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

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).

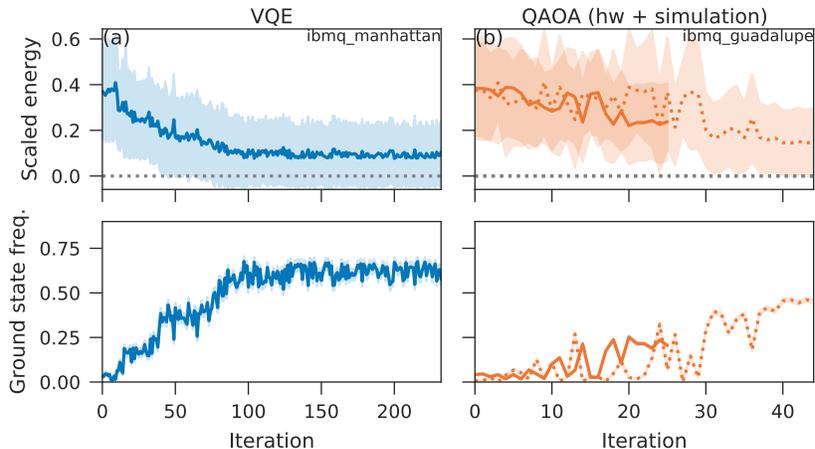


FIG. 3. VQE and QAOA scaled energy (top panels) and ground state frequency (bottom panels) for the JSP instance using 5 qubits. The energy was rescaled with the maximum energy eigenvalue. VQE with the CVaR objective ($\alpha = 0.5$) was executed on `ibmq_manhattan` (1,024 shots) and includes basic error mitigation techniques (measurement error and optimized qubit allocation). QAOA with $p = 2$ was executed on `ibmq_guadalupe` (solid line, 1,024 shots) and a noiseless simulator (dotted, 8,000 shots). Error bands are the standard deviation (top panels) and 95% confidence interval (bottom panels).

III. RESULTS AND DISCUSSION

We have tested the algorithms in Sec. II C on instances of the JSP on IBM quantum processors and simulators. 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 with exhaustive search). Instead we focus on convergence and scaling the VQA for this particular optimization problem of industrial relevance. All quantum processors and simulators were accessed via `tket` [34]. Hardware experiments benefitted from `tket`'s out-of-the-box noise-aware qubit placement and routing, and IBM's error measurement reduction, 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 I 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 energies scaled to the range $[0, 1]$:

$$\epsilon = \frac{\langle H \rangle - E_{\min}}{E_{\max} - E_{\min}} \in [0, 1], \quad (16)$$

where E_{\min} , E_{\max} are the minimum and maximum energy of the Hamiltonian, respectively. We calculated E_{\min} , E_{\max} exactly, but for larger problems it is possible to compute bounds via semidefinite programming relaxation of the problem [35]. A value $\epsilon = 0$ 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 $\epsilon \approx 0$ in few iterations. For a given state $|\psi\rangle$ we also plot the frequency of sampling the ground state of the problem Hamiltonian $|\psi_{\text{gs}}\rangle$:

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

Ideally, we would like an algorithm to return the ground state with a frequency $P_{\text{gs}} \approx 1$. In general, a high frequency of sampling the ground state implies a low scaled energy but the converse is not true.

A. 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 and simulators. Table II summarises the settings of all algorithms used in the following results. 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 Sec. III B.

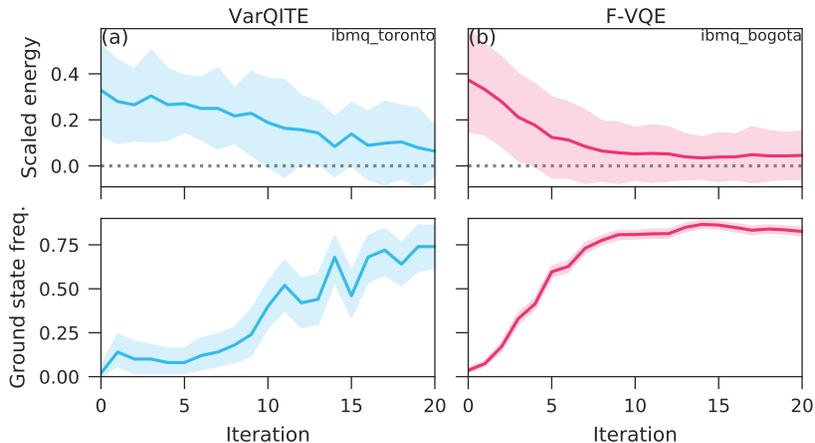


FIG. 4. VarQITE and F-VQE scaled energy (top panels) and ground state frequency (bottom panels) for the JSP instance using 5 qubits. The energy was rescaled with the maximum energy eigenvalue. VarQITE was executed on `ibmq_toronto` (50 shots), F-VQE was executed on `ibmq_bogota` (1,024 shots). Error bands are the standard deviation (top panels) and 95% confidence interval (bottom panels).

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 ($\alpha = 1$) and also investigated the influence of simple error mitigation techniques. We observed that the CVaR mainly leads to lower variance in the measurement outcomes and slightly increases the probability of sampling the ground state. Measurement error mitigation and noise-aware qubit placement had a more noticeable impact on the probability of sampling the ground state. Fig. 3(a) shows the results for VQE using CVaR with both error mitigation techniques and 1,024 shots. VQE converged after around 120 iterations with a frequency of sampling the ground state of approximately 60%. While the slow convergence and ground state frequency could likely be improved by a different optimizer [22, 36], our results appear to confirm that standard VQE without tuning is unlikely to yield competitive results for combinatorial optimization on NISQ computers [13, 23].

QAOA with $p = 2$ executed on a quantum processor showed very slow convergence. In fact, the optimizer COBYLA stopped after around 25 iterations leading to poor performance on the JSP as shown in Fig. 3(b). Apart from the optimizer another contributing factor is likely that the QAOA ansatz is not hardware-efficient, i.e. the compiler needs to add 2-qubit gates for routing. The challenge of scaling QAOA on quantum processors with restricted qubit connectivity was also highlighted in [17]. To extend our results slightly we also include a noiseless simulation with 8,000 shots as a dotted line in Fig. 3(b). In this simulation, COBYLA was able to improve average energy and ground state frequency slightly further up to 44 iterations. In summary, our results appear to confirm that QAOA running on NISQ hardware requires fine-tuned optimizers even for small-scale instances [23, 24]. Instead of fine-tuning QAOA, however, we chose to focus on other VQA showing better out-of-the-box performance for our problem.

Both VarQITE and F-VQE converged faster on quantum processors than VQE or QAOA (see Fig. 4). Both implementations used a specific gradient-based optimizer (see Sec. II C 3, II C 4). VarQITE reaches lower average energies after 20 iterations than VQE after over 200 iterations even though VarQITE was not fully converged yet (see Fig. 4(a)). It also reaches high probability of sampling the ground state showing that a sample of bitstrings is very likely to include the optimal solution. Overall, F-VQE achieved the fastest convergence for the JSP. Fig. 4(b) shows it produces mean energies close to the minimum and high probability (around 80%) of the sampling optimal solution after just 10-15 iterations. We repeated the F-VQE experiment with the ansatz in Fig. 2(b) 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 estimated from measurement samples (cf. Eq. (10)). Based on these results we chose to focus on F-VQE for scaling up to larger JSP instances.

B. 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. F-VQE converges quickly in all cases. All experiments reach a significant nonzero frequency of sampling the ground state: $P_{\text{gs}} \approx 80\%$ for 10

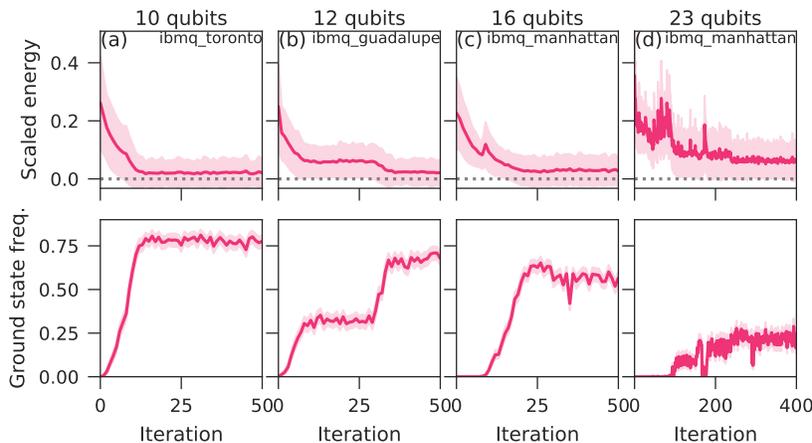


FIG. 5. F-VQE scaled energy (top panels) and ground state frequency (bottom panels) for different JSP instances with (from left to right) $N = 5$ (ibmq_toronto, 500 shots), $N = 12$ (ibmq_guadalupe, 550 shots), $N = 16$ (ibmq_manhattan, 650 shots) and $N = 23$ qubits (ibmq_manhattan, 450 shots). The energy was rescaled with the maximum energy eigenvalue. Error bands are the standard deviation (top panels) and 95% confidence interval (bottom panels).

qubits, $P_{\text{gs}} \approx 70\%$ for 12 qubits, $P_{\text{gs}} \approx 60\%$ for 16 qubits, and $P_{\text{gs}} \approx 25\%$ 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 ($p = 1, \dots, 3$).

IV. 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.[39] 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 100s of qubits is sparse connectivity and cross-talk

noise. F-VQE can address this concerns 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 overheads from routing. In addition, error mitigation post processing can be readily utilized [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.

V. 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

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:

$$u_j(\mathbf{x}) = c_e \sum_{t=1}^{d_j} (d_j - t)x_{Mjt} + c_l \sum_{t=d_j+1}^{T_M} (t - d_j)x_{Mjt} \quad \forall j = 1, \dots, J. \quad (\text{A1})$$

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

$$G_{j_1 j_2}^{(m)} = \begin{cases} 0 & \text{if } P_{mj_1} = P_{mj_2}, \\ 1 & \text{otherwise.} \end{cases} \quad (\text{A2})$$

Hence, the production cost for machine m is given by

$$s_m(\mathbf{x}) = c_p \sum_{j_1, j_2=1}^J \sum_{t=1}^{T_m-1} G_{j_1 j_2}^{(m)} x_{mj_1 t} x_{mj_2(t+1)} \quad \forall m = 1, \dots, M. \quad (\text{A3})$$

The constant c_p determines the magnitude of the production cost.

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

$$c(\mathbf{x}) = \sum_{j=1}^J u_j(\mathbf{x}) + \sum_{m=1}^M s_m(\mathbf{x}). \quad (\text{A4})$$

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

$$g_{mj}(\mathbf{x}) = \sum_{t=1}^{T_m} x_{mjt} = 1, \quad \forall m = 1, \dots, M \quad \forall j = 1, \dots, J. \quad (\text{A5})$$

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

$$\ell_{mt}(\mathbf{x}, \mathbf{y}) = \left\{ \begin{array}{ll} y_{mt} + \sum_{i=1}^J x_{mjt} & \text{for } 1 \leq t \leq e_m \\ \sum_{j=1}^J x_{mjt} & \text{for } e_m < t \leq J \\ 1 - y_{m(t-J)} + \sum_{j=1}^J x_{mjt} & \text{for } J < t \leq T_m \end{array} \right\} = 1 \quad \begin{array}{l} \forall t = 1, \dots, T_m \\ \forall m = 1, \dots, M. \end{array} \quad (\text{A6})$$

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

$$q_{mj}(\mathbf{x}) = \sum_{t=2}^{T_m} \sum_{t'=1}^{t-1} x_{mjt} x_{(m+1)jt'} = 0 \quad \forall m = 1, \dots, M-1 \quad \forall j = 1, \dots, J. \quad (\text{A7})$$

The *idle slot constraints* ensure that dummy jobs are placed before all real jobs in each machine. Due to constraints ℓ_{mt} in Eq. (A6) we only need to enforce that the transition from a real job to a dummy job is prohibited at the beginning of a schedule:

$$r_{mt}(\mathbf{y}) = (1 - y_{mt})y_{m(t+1)} = 0 \quad \forall t = 1, \dots, e_m - 1 \quad \forall m = 2, \dots, M. \quad (\text{A8})$$

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 Sec. II B is

$$N = \sum_{m=1}^M J(J + e_m) + e_m. \quad (\text{B1})$$

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

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 $O(J^2M)$ 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 III lists the devices 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?systems=all>.

Device	No. of qubits	Quantum Volume	Connectivity
ibmq_manhattan	65	32	Heavy-hexagon [42]
ibmq_toronto	27	32	Heavy-hexagon
ibmq_guadalupe	16	32	Heavy-hexagon
ibmq_bogota	5	32	Linear

TABLE III. Hardware devices used in this study.

AVAILABILITY OF DATA AND MATERIALS

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

COMPETING INTERESTS

The authors declare that they have no competing interests.

AUTHOR'S 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.

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.

-
- [1] L. K. Grover, A fast quantum mechanical algorithm for database search, in *Proceedings of the twenty-eighth annual ACM symposium on Theory of Computing*, STOC '96 (Association for Computing Machinery, New York, NY, USA, 1996) pp. 212—219, [arXiv:quant-ph/9605043](#).
 - [2] C. Durr and P. Hoyer, A Quantum Algorithm for Finding the Minimum (1999), [arXiv:quant-ph/9607014](#).
 - [3] J. Preskill, Quantum Computing in the NISQ era and beyond, [Quantum](#) **2**, 79 (2018).
 - [4] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, Variational Quantum Algorithms (2020), [arXiv:2012.09265](#).
 - [5] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, Noisy intermediate-scale quantum (NISQ) algorithms (2021), [arXiv:2101.08448](#).
 - [6] Y. R. Sanders, D. W. Berry, P. C. Costa, L. W. Tessler, N. Wiebe, C. Gidney, H. Neven, and R. Babbush, Compilation of Fault-Tolerant Quantum Heuristics for Combinatorial Optimization, [PRX Quantum](#) **1**, 020312 (2020).
 - [7] G. Kochenberger, J.-K. Hao, F. Glover, M. Lewis, Z. Lü, H. Wang, and Y. Wang, The unconstrained binary quadratic programming problem: a survey, [J. Comb. Optim.](#) **28**, 58 (2014).
 - [8] A. Lucas, Ising formulations of many NP problems, [Front. Phys.](#) **2**, 5 (2014), [arXiv:1302.5843](#).
 - [9] F. Glover, G. Kochenberger, and Y. Du, Quantum Bridge Analytics I: a tutorial on formulating and using QUBO models, [4OR-Q J Oper Res](#) **17**, 335 (2019).
 - [10] E. Farhi, J. Goldstone, and S. Gutmann, A Quantum Approximate Optimization Algorithm, ArXiv e-prints (2014), [1411.4028](#).
 - [11] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, [Nat. Commun.](#) **5**, 4213 (2014).
 - [12] S. McArdle, T. Jones, S. Endo, Y. Li, S. C. Benjamin, and X. Yuan, Variational ansatz-based quantum simulation of imaginary time evolution, [npj Quantum Inf.](#) **5**, 75 (2019).
 - [13] D. Amaro, C. Modica, M. Rosenkranz, M. Fiorentini, M. Benedetti, and M. Lubasch, Filtering variational quantum algorithms for combinatorial optimization, [arXiv:2106.10055 \[quant-ph\]](#) (2021), [arXiv:2106.10055 \[quant-ph\]](#).

- [14] E. Farhi and A. W. Harrow, Quantum supremacy through the quantum approximate optimization algorithm (2016), [arXiv:1602.07674](https://arxiv.org/abs/1602.07674).
- [15] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, Quantum approximate optimization algorithm: Performance, mechanism, and implementation on near-term devices, *Phys. Rev. X* **10**, 021067 (2020).
- [16] C. Moussa, H. Calandra, and V. Dunjko, To quantum or not to quantum: towards algorithm selection in near-term quantum optimization, *Quantum Sci. Technol.* **5**, 044009 (2020), [arXiv:2001.08271](https://arxiv.org/abs/2001.08271).
- [17] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo, M. Broughton, B. B. Buckley, D. A. Buell, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, Ben Chiaro, R. Collins, W. Courtney, S. Demura, A. Dunsworth, D. Eppens, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, S. Habegger, A. Ho, S. Hong, T. Huang, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, C. Jones, D. Kafri, K. Kechedzhi, J. Kelly, S. Kim, P. V. Klimov, A. N. Korotkov, F. Kostritsa, D. Landhuis, P. Laptev, M. Lindmark, M. Leib, O. Martin, J. M. Martinis, J. R. McClean, M. McEwen, A. Megrant, X. Mi, M. Mohseni, W. Mruczkiewicz, J. Mutus, O. Naaman, C. Neill, F. Neukart, M. Y. Niu, T. E. O'Brien, B. O'Gorman, E. Ostby, A. Petukhov, H. Putterman, C. Quintana, P. Roushan, N. C. Rubin, D. Sank, A. Skolik, V. Smelyanskiy, D. Strain, M. Streif, M. Szalay, A. Vainsencher, T. White, Z. J. Yao, P. Yeh, A. Zalcman, L. Zhou, H. Neven, D. Bacon, E. Lucero, E. Farhi, and R. Babbush, Quantum approximate optimization of non-planar graph problems on a planar superconducting processor, *Nature Physics* **17**, 332 (2021).
- [18] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, Expressibility and Entangling Capability of Parameterized Quantum Circuits for Hybrid Quantum-Classical Algorithms, *Advanced Quantum Technologies* **2**, 1900070 (2019), [arXiv:1905.10876](https://arxiv.org/abs/1905.10876).
- [19] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nature Communications* **9**, 4812 (2018).
- [20] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nature Communications* **12**, 1791 (2021).
- [21] G. G. Guerreschi and M. Smelyanskiy, Practical optimization for hybrid quantum-classical algorithms (2017), [arXiv:1701.01450](https://arxiv.org/abs/1701.01450).
- [22] G. Nannicini, Performance of hybrid quantum-classical variational heuristics for combinatorial optimization, *Physical Review E* **99**, 013304 (2019), [arXiv:1805.12037](https://arxiv.org/abs/1805.12037).
- [23] W. Lavrijsen, A. Tudor, J. Müller, C. Iancu, and W. de Jong, Classical Optimizers for Noisy Intermediate-Scale Quantum Devices, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (2020) pp. 267–277, [arXiv:2004.03004](https://arxiv.org/abs/2004.03004).
- [24] K. J. Sung, J. Yao, M. P. Harrigan, N. C. Rubin, Z. Jiang, L. Lin, R. Babbush, and J. R. McClean, Using models to improve optimizers for variational quantum algorithms, *Quantum Sci. Technol.* **5**, 044008 (2020), [arXiv:2005.11011](https://arxiv.org/abs/2005.11011).
- [25] A. Pellow-Jarman, I. Sinayskiy, A. Pillay, and F. Petruccione, A Comparison of Various Classical Optimizers for a Variational Quantum Linear Solver, *Quantum Information Processing* **20**, 202 (2021), [arXiv:2106.08682](https://arxiv.org/abs/2106.08682).
- [26] P. K. Barkoutsos, G. Nannicini, A. Robert, I. Tavernelli, and S. Woerner, Improving Variational Quantum Optimization using CVaR, *Quantum* **4**, 256 (2020).
- [27] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549**, 242 (2017).
- [28] M. J. D. Powell, A Direct Search Optimization Method That Models the Objective and Constraint Functions by Linear Interpolation, in *Advances in Optimization and Numerical Analysis*, edited by S. Gomez and J.-P. Hennart (Springer Netherlands, Dordrecht, 1994) pp. 51–67.
- [29] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, Quantum Computation by Adiabatic Evolution (2000), [arXiv:quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106).
- [30] X. Yuan, S. Endo, Q. Zhao, Y. Li, and S. C. Benjamin, Theory of variational quantum simulation, *Quantum* **3**, 191 (2019).
- [31] M. Motta, C. Sun, A. T. K. Tan, M. J. O'Rourke, E. Ye, A. J. Minnich, F. G. S. L. Brandão, and G. K.-L. Chan, Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution, *Nat. Phys.* **16**, 205 (2020).
- [32] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, *Phys. Rev. A* **99**, 032331 (2019), [arXiv:1811.11184](https://arxiv.org/abs/1811.11184).
- [33] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum Circuit Learning, *Phys. Rev. A* **98**, 032309 (2018), [arXiv:1803.00745](https://arxiv.org/abs/1803.00745).
- [34] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, t|ket>: A Retargetable Compiler for NISQ Devices, *Quantum Sci. Technol.* **6**, 014003 (2020), [arXiv:2003.10611](https://arxiv.org/abs/2003.10611).
- [35] C. Helmberg, Semidefinite programming for combinatorial optimization (2000).
- [36] J. Stokes, J. Izaac, N. Killoran, and G. Carleo, Quantum Natural Gradient, *Quantum* **4**, 269 (2020).
- [37] J. S. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, P. Karalekas, C. B. Osborn, A. Papageorge, E. C. Peterson, G. Prawiroatmodjo, N. Rubin, C. A. Ryan, D. Scarabelli, M. Scheer, E. A. Sete, P. Sivarajah, R. S. Smith, A. Staley, N. Tezak, W. J. Zeng, A. Hudson, B. R. Johnson, M. Reagor, M. P. da Silva, and C. Rigetti, Unsupervised Machine Learning on a Hybrid Quantum Computer (2017), [arXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [38] G. Pagano, A. Bapat, P. Becker, K. S. Collins, A. De, P. W. Hess, H. B. Kaplan, A. Kyprianidis, W. L. Tan, C. Baldwin, L. T. Brady, A. Deshpande, F. Liu, S. Jordan, A. V. Gorshkov, and C. Monroe, Quantum Approximate Optimization of the Long-Range Ising Model with a Trapped-Ion Quantum Simulator, *Proceedings of the National Academy of Sciences* **117**, 25396 (2020), [arXiv:1906.02700](https://arxiv.org/abs/1906.02700).
- [39] See development roadmaps by IBM <https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap/> and 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 2021-07-01).
- [40] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, Hybrid Quantum-Classical Algorithms and Quantum Error Mitigation, *Journal of the Physical Society of Japan* **90**, 032001 (2021), [arXiv:2011.01382](#).
 - [41] M. Benedetti, M. Fiorentini, and M. Lubasch, Hardware-efficient variational quantum algorithms for time evolution (2020), [arXiv:2009.12361](#).
 - [42] C. Chamberland, G. Zhu, T. J. Yoder, J. B. Hertzberg, and A. W. Cross, Topological and Subsystem Codes on Low-Degree Graphs with Flag Qubits, *Physical Review X* **10**, 011022 (2020).