

Fair Sampling Error Analysis on NISQ Devices

JOHN GOLDEN, Information Sciences (CCS-3) and Computational Earth Sciences (EES-16), Los Alamos National Laboratory, USA

ANDREAS BÄRTSCHI, Information Sciences (CCS-3), Los Alamos National Laboratory, USA

DANIEL O'MALLEY, Computational Earth Sciences (EES-16), Los Alamos National Laboratory, USA

STEPHAN EIDENBENZ, Information Sciences (CCS-3), Los Alamos National Laboratory, USA

We study the status of *fair sampling* on **Noisy Intermediate Scale Quantum (NISQ)** devices, in particular the IBM Q family of backends. Using the recently introduced Grover Mixer-QAOA algorithm for discrete optimization, we generate fair sampling circuits to solve six problems of varying difficulty, each with several optimal solutions, which we then run on twenty backends across the IBM Q system. For a given circuit evaluated on a specific set of qubits, we evaluate: how frequently the qubits return an optimal solution to the problem, the fairness with which the qubits sample from all optimal solutions, and the reported hardware error rate of the qubits. To quantify fairness, we define a novel metric based on Pearson's χ^2 test. We find that fairness is relatively high for circuits with small and large error rates, but drops for circuits with medium error rates. This indicates that structured errors dominate in this regime, while unstructured errors, which are random and thus inherently fair, dominate in noisier qubits and longer circuits. Our results show that fairness can be a powerful tool for understanding the intricate web of errors affecting current NISQ hardware.

CCS Concepts: • **Hardware** → **Quantum error correction and fault tolerance**; • **Computer systems organization** → *Quantum computing*

Additional Key Words and Phrases: Quantum alternating operator ansatz, fair sampling

ACM Reference format:

John Golden, Andreas Bäertschi, Daniel O'Malley, and Stephan Eidenbenz. 2022. Fair Sampling Error Analysis on NISQ Devices. *ACM Trans. Quantum Comput.* 3, 2, Article 8 (May 2022), 23 pages.

<https://doi.org/10.1145/3510857>

1 INTRODUCTION

Combinatorial optimization problems, such as the Minimum Traveling Salesperson problem, the Maximum Satisfiability problem, or the Spin Glass Ising model, are widely viewed to be one of the most promising application domains for quantum computing. Such problems often have multiple optimal solutions, and depending on the algorithm being used to solve the problem, some optima may be much easier to find than others. This is problematic if one is interested not in finding a single solution, but instead determining the overall landscape of optimal solutions. Algorithms

Authors' addresses: J. Golden, Information Sciences (CCS-3) and Computational Earth Sciences (EES-16), Bldg. 1489, Casa Grande Drive, Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 87545; email: golden@lanl.gov; A. Bäertschi and S. Eidenbenz, Information Sciences (CCS-3), Bldg. 508, Diamond Drive, Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 87545; emails: {baertschi, eidenbenz}@lanl.gov; D. O'Malley, Computational Earth Sciences (EES-16), Bldg. 1400, Casa Grande Drive, Los Alamos National Laboratory, Los Alamos, New Mexico, USA, 87545; email: omalled@lanl.gov.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s).

2643-6817/2022/05-ART8

<https://doi.org/10.1145/3510857>

which fairly sample from the space of optimal solutions avoid this problem and have practical value in physics and engineering. For instance, it is often difficult in practice to encode all the engineering design goals in the function to be optimized. By allowing an engineer to explore an unbiased assortment of optimal solutions, fair sampling provides an opportunity to apply these other design goals to a more diverse set of options (especially useful when a customer adds one more requirement after the optimization has been performed). It can also be used to produce an ensemble of predictions from calibrated physical models that can be used to make a forecast including a range of possible outcomes (e.g., in the context of fluid flow in the earth's subsurface, which can be formulated as a discrete optimization problem [14, 24]).

In this paper, we determine the current status of fair sampling in the context of quantum computing. Specifically, we are interested both in algorithms which theoretically guarantee fair sampling (i.e., generate quantum states where all optimal solutions have the same amplitudes squared and thus are measured with equal probability), as well as actual hardware implementations of such algorithms. Previous studies have shown that quantum annealing cannot be used as a fair sampler for all problems [18]. We therefore focus our work on gate-based quantum computing algorithms and hardware.

The most widely-used quantum optimization framework is the **Quantum Alternating Operator Ansatz (QAOA)**, which alternates between a mixer operator that mixes the amplitudes of all feasible solutions and a phase separator operator that separates the phases of the feasible solutions [13]. Most mixers do not exhibit fair sampling, however, the recently discovered Grover Mixer [4] has the property that all feasible solutions with the same objective value have identical amplitudes. This novel Grover Mixer implementation of QAOA thus solves the theoretical side of fair sampling for gate-based quantum computing.

The majority of this work is dedicated to a practical evaluation of how fair sampling circuits are affected by the errors inherent in **Noisy Intermediate Scale Quantum (NISQ)** devices. In this study, a paradox quickly becomes apparent: while an error-corrected quantum device could obviously be used as a perfectly fair sampler, so could a quantum device completely dominated by random noise. Of course the noisy machine would be no different than a random number generator, and would sample from all feasible solutions rather than just the optimal solutions, but it would do so fairly. Therefore, our study seeks to clarify where current NISQ devices lie on the spectrum between these two extremes. Specifically, are the hardware errors predominantly random and unstructured, thus leading to relatively fair sampling of optimal solutions? Or do hardware biases and other structured errors play an important role, thus leading to unfair behavior?

Our methodology for answering these questions consists of the following steps:

- (1) We define six example problems, which are a slightly expanded set of examples studied in the context of fair sampling in quantum annealing [18].
- (2) We design quantum circuits for our problems and hand-tune them to minimize circuit depth and gate counts, taking into account different qubit connectivity topologies as encountered on IBM Q. Our circuits involve two to five qubits and our medium size examples result in circuits with gate counts in the 30–60 range, which is similar to the size of the (random) circuits used to establish the Quantum Volume metric on the IBM Q backends.
- (3) We execute our circuits in a large number of shots on twenty different IBM Q backends that have qubit counts of up to 65 qubits. In addition, we embed our circuits in all possible topologies on the circuits resulting in a total of 7,440 experiments at 40,960 shots each.
- (4) We define a novel metric for fairness in quantum computing based on Pearson's χ^2 test, a common statistical tool to evaluate how likely it is that a set of observational data came from a null hypothesis distribution. Since all NISQ hardware will necessarily involve some biases

and imperfections that inhibit fair sampling, we are interested in how much effort must be expended to prove that a piece of hardware is biased, i.e., reject the null hypothesis of fair sampling. We term this metric ‘number of shots to reject fair sampling,’ and it depends on both a specific set of qubits as well as a fair sampling circuit. This fairness metric is a separate contribution to the main results of the paper.

- (5) We analyze our experiments using our fairness metric and we calculate the aggregate error of each experiment from IBM’s reported gate-specific error rates.

We find that the NISQ devices under evaluation provide noticeably unfair sampling. However, fairness increases as aggregate error approaches both 0 and 1, matching with the theoretical extremes of error-corrected and noise-dominated devices discussed previously. In the case of small aggregate error, this is to be expected, as for 0 error we would retrieve the theoretical guarantees of Grover Mixer QAOA. Likewise, for large aggregate error, one should get samples uniformly at random from all computational basis states. Nonetheless, the result is perhaps counter-intuitive, as one might expect a linear interpolation between these two extremes. For example, the hardware-efficient random circuits in Google’s recent quantum supremacy experiment [2] give a sample distribution consisting of a convex combination of an ideal (error-free) Porter-Thomas distribution and a Uniform distribution (caused by errors). Our differing results indicate that we should not expect errors themselves to be evenly distributed (i.e., in a fair manner) across qubits for highly structured circuits such as ours. Indeed, a previous study found that structured errors (e.g., crosstalk, coherent noise) play a large role in the performance of structured circuits [25]. Our result holds for our selected set of circuits and across the twenty NISQ backends on IBM Q.

We propose that fair sampling be an addition in the emerging set of benchmarks for NISQ devices, which so far have largely focused on circuit correctness. As we show in this work, biased errors can artificially inflate circuit accuracy, and fair sampling on its own can be fooled by devices dominated by random noise. However, by studying both accuracy and fairness, we gain a more complete picture of the type and magnitude of errors affecting NISQ hardware.

In summary, the main contributions of our work are:

- (1) We propose a novel methodology for quantifying the fairness of sampling on NISQ hardware.
- (2) We elucidate two noise regimes (high noise and low noise) that result in fair sampling, and a third regime (intermediate noise) that is dominated by biased errors.
- (3) In practice, this approach provides a way to determine whether near-term improvements that reduce noise will improve or degrade the performance of fair sampling for a given problem.
- (4) As an example, we study how a practical error mitigation technique (namely, measurement error mitigation) impacts the quality of sampling on IBM Q hardware.

2 REVIEW

In this section we describe previous studies of fair sampling in the context of quantum computing. A note on nomenclature: since the optimization problems we study can all be phrased in terms of Ising models, we adopt the more physics-oriented language and refer to ‘optimal solutions’ as ground states. Problems with multiple optimal solutions are thus said to have degenerate ground states.

2.1 Fair Sampling in Quantum Annealing

After quantum annealing was proposed as a possibly useful optimization heuristic [16], it was discovered that naïve approaches to quantum annealing do not always sample degenerate ground states fairly [20]. In fact, certain Hamiltonians feature ground states that are sampled with

probability approaching zero as annealing time increases. This effect is inherent to the quantum annealing approaches, and is not the consequence of hardware noise or other technical limitations. This behavior was subsequently verified experimentally [19]. Since unfair sampling (especially when certain ground states are strongly suppressed) can have negative consequences for many applications, techniques were proposed to improve upon the naïve quantum annealing approach [17, 18, 27, 34]. None of these techniques have been tested on quantum annealing hardware, indicating some of the challenges with these approaches. It should also be noted that it has been proposed that there are some advantages of biased samplers, that is, they can be used effectively in an ensemble of samplers each of which has different biases [36].

One technique that has been proposed to improve the fairness of quantum annealing is the introduction of more complex Hamiltonians to drive the annealing process [19]. This generally improves fairness, but is hard to implement and does not resolve the underlying problem. Even for small problems, biases sometimes remain despite using fairly high-order Hamiltonians [18]. In studying these higher-order Hamiltonians, Könz et al. [18] introduced four Ising models that are suitable for our QAOA analyses. The problems are described in Table 1 as Problems (a)–(d). These problems all exhibit biased sampling with quantum annealing, so they present a non-trivial challenge for fair sampling algorithms. Two of them involve 4 qubits, one involves 5 qubits, and another involves 6 qubits. This small number of qubits make it plausible that these problems can be solved with reasonably high fidelity on current NISQ hardware.

Impact of noise: Recent work [29] achieved high-quality thermal Gibbs Sampling with the D-Wave 2000Q system for a restricted class of Ising Hamiltonians that are native to the hardware connectivity and have coupling strengths $J_{ij} \in \{-1, 0, 1\}$. This category includes Hamiltonians with degenerate ground states, but not the more general Hamiltonians studied in this paper (see Table 1). Fair Sampling of those ground states is achieved by re-scaling the input energy scale to a sweet spot that is large enough to be noise-resilient but small enough to overcome the biases from the transverse field. In the identified sweet spot energy scale, changing the annealing time influences the effective temperature of the Gibbs samples generated by the hardware, and thus the probability to find ground states.

2.2 Fair Sampling in the Quantum Alternating Operator Ansatz

The **Quantum Alternating Operator Ansatz** [13] (QAOA) is a promising heuristic quantum algorithm for (combinatorial) optimization problems. In its essence, for a problem instance I with feasible states F and cost Hamiltonian H_C on n qubits, a p -round QAOA prepares a parametrized state from which one would like to sample low-energy states with respect to H_C :

$$|\beta, \gamma\rangle := U_M(\beta_p)U_P(\gamma_p) \cdots U_M(\beta_1)U_P(\gamma_1)U_S |\uparrow^n\rangle. \quad (1)$$

The circuit consists of an initial *state preparation* unitary operator U_S that creates some superposition of all feasible solutions in I , followed by p applications of alternating parametrized *phase separating* and *mixing* unitaries $U_P(\gamma_k)$, $U_M(\beta_k)$ with real angle parameters $\gamma = (\gamma_1, \dots, \gamma_p)^T$ and $\beta = (\beta_1, \dots, \beta_p)^T$, and a final measurement in the computational basis, see e.g., Figure 1. The β, γ angles are traditionally found via a hybrid classical/quantum approach: given some initial parameters, take enough samples to reliably estimate the expectation value $\langle \beta, \gamma | H_C | \beta, \gamma \rangle$, use a classical optimizer to adjust the parameters, and repeat until one has a state with a good (low-energy) expectation value. For further description of, and theoretical justifications for this approach, see [6, 8, 31].

The role of the phase separating unitaries $U_P(\gamma)$ is to add multiplicative phase factors to the amplitudes of feasible computational basis states, with phases proportional to respective energies. We usually have (up to global phases) $U_P(\gamma) \cong e^{-i\gamma H_C}$. For many problems, such as MaxCut [8] or the problems considered in this paper, H_C is an Ising Hamiltonian with quadratic and sometimes linear




Fig. 1. Grover Mixer QAOA: The state preparation unitary U_S for an equal superposition of all feasible states $|F\rangle = \frac{1}{\sqrt{|F|}} \sum_{x \in F} |x\rangle$, its conjugate transpose U_S^\dagger , and a multi-controlled phase-shift gate $Z^{-\beta/\pi} = \begin{pmatrix} 1 & 0 \\ 0 & e^{-i\beta} \end{pmatrix}$ are used to implement the mixer $U_M(\beta) = e^{-i\beta|F\rangle\langle F|}$.

terms; but in other problems, H_C can also involve higher-order terms, such as for MaxE3Lin2 [9]. The interplay between the state preparation unitary U_S and the mixing unitary U_M is particularly important, and falls into three categories:

Original Approach. The original approach by Farhi et al. [8], termed Quantum Approximate Optimization Algorithm, was inspired by quantum annealing for unconstrained optimization problems, i.e., $U_S = H^{\otimes n}$ (preparing an equal superposition of all states) and $U_M(\beta) = e^{-i\beta \sum X_j}$ (a transverse field of Pauli-X operators acting on individual qubits).

General Ansatz. The general framework introduced in Equation (1) by Hadfield et al. [13] is a generalization also suitable for a wide variety of constrained optimization problems. They focus on designing involved mixing unitaries $U_M(\beta)$ and usually split U_S in a layer of Identity and Pauli-X gates preparing some feasible computational basis state, followed by an initial mixing unitary $U_M(\beta_0)$ with new angle β_0 .

Grover Mixer. If one can design efficient state preparation unitaries U_S preparing an equal superposition of all feasible basis states $|F\rangle = 1/\sqrt{|F|} \sum_{x \in F} |x\rangle$, then this gives rise to a mixing unitary resembling Grover's selective phase-shift operator [12, 22, 35], $U_M(\beta) = e^{-i\beta|F\rangle\langle F|} = Id - (1 - e^{-i\beta}) |F\rangle\langle F| = U_S(Id - (1 - e^{-i\beta}) |\uparrow\rangle\langle\uparrow|)U_S^\dagger$, see Figure 1. In some sense, this is the reverse of the method in the general ansatz, and it was shown to be applicable to many relevant optimization problems [4], but there also exists a range of constrained problems where it is ruled out under complexity-theoretic assumptions [5].

Although the Quantum Alternating Operator Ansatz is the most general framework of the three, both the Quantum Approximate Optimization Algorithm as well as Grover Mixer QAOA have their advantages. For example, the original approach by definition cannot do worse than a Trotterization of the adiabatic algorithm [10], thus converging to a ground state in the $p \rightarrow \infty$ limit [8]. Furthermore, its phase separating and mixing operators are periodic, restricting the angles to real values in $[-\pi, \pi)$. Grover Mixer QAOA retrieves these properties for constrained optimization problems with efficient state preparation unitaries U_S [5], but for unconstrained optimization problems the Grover Mixer $e^{-i\beta|F\rangle\langle F|}$ has a larger circuit depth than the Transverse Field Mixer $e^{-i\beta \sum X_j}$. In this paper we explore an advantage unique to Grover Mixer QAOA among the three categories: fair sampling.

With Grover Mixer QAOA, all feasible basis states begin with amplitude $1/\sqrt{|F|}$ (after state preparation with U_S). The phase separating unitary $U_P(\gamma) = e^{-i\gamma H_C}$ then phases the amplitude of

every basis state proportional to its energy and γ , keeping the same phase for basis states of same energy. The mixing unitary $U_M(\beta) = Id - (1 - e^{-i\beta}) |F\rangle\langle F|$ then deducts from all amplitudes of their input state $|\psi\rangle$ the same weighted average of all of its amplitudes, $((1 - e^{-i\beta})/\sqrt{|F|})\langle F | \psi \rangle$. Therefore, basis states with the same energy are sampled with the same amplitude. This property holds for all feasible states (not only ground states) and is independent of the number of QAOA levels or the choice of angles for U_M, U_P . For a complete proof, see [4].

Impact of noise: While the original based Quantum Approximate Optimization Algorithm offers non-trivial provable performance guarantees already for a single for certain problems such as MAXCUT on $(d=3)$ -regular graphs [8], a higher number of rounds – $p \geq 6$ [33] or even $p \in \Omega(\log(n)/d)$ [3] – might be necessary to outperform the best-known classical approximation ratios.

However, simulations considering realistic noise models of current hardware show that performance in terms of approximation ratio may degrade even for $p = 2$ [1]. Noise leading to a decrease in performance has also been observed for growing problem sizes n , with the exception of problem graphs matching the hardware connectivity [15]. For the latter setting, experiments on actual hardware have shown an approximation ratio increase up to $p = 3$ (mean) and $p = 4$ (certain instances), but no further [15]. Furthermore, even with constantly improving hardware, complications may arise in the optimization of the variational angles β, γ due to noise-induced barren plateaus [30].

All of these results refer to transverse field based QAOA, which does not have the property of sampling fairly among ground states. For Grover Mixer QAOA, we have a theoretical guarantee of fair sampling for ideal noise-less devices. At the same time we have a larger circuit depth of the mixer, asymptotically matching the circuit depth of the phase separator [4]. Hence, for current NISQ devices, we restrict ourselves to a single-round Grover Mixer QAOA, analysing the impact of noise on both ground state probability and fair sampling among ground states.

2.3 Description of Test Circuits

In addition to the four problems from [18], we additionally studied two simpler Ising models with degenerate ground states but only involving two and three qubits. While these models do not feature ground state suppression on quantum annealing hardware, we include them in this study because they represent the simplest non-trivial models with degenerate ground states and thus serve as useful baselines as compared to the more complicated models of [18]. All of the models we studied are depicted in Table 1.

Our general procedure for generating test circuits was to begin with a 1-level Grover Mixer QAOA algorithm with q_0 fixed to \uparrow . We chose a 1-level implementation of Grover-QAOA in order to keep circuit depth low. Following [18], we fix $q_0 := \uparrow$ as all of the models in Table 1 are \uparrow / \downarrow symmetric. This allows us to reduce the problems from Ising Hamiltonians H_C with only quadratic terms acting on qubits q_0, \dots, q_{n-1} to new Hamiltonians H'_C with some linear terms acting only on qubits q_1, \dots, q_{n-1} . This also transforms q_0 into a classical control bit for the consecutive Grover Mixer, which can be removed. Thus, we can embed Ising problems on n qubits onto circuits with only $n - 1$ qubits. See Figure 9 in Appendix C for a visual representation of a generic 1-level Grover Mixer QAOA for an unconstrained optimization problem, along with a description of the simplifications resulting from fixing q_0 .

We then compiled the circuits to match connectivity graphs and gates available on the IBM Q backend. Instead of employing the tools available in IBM's qiskit software, we compiled the circuits by hand. This had two benefits. First, this reduced circuit depth by roughly a factor of two as compared to qiskit. Second, it allowed us to generate circuits for architectures that can be found as subgraphs of every hardware connectivity graph among the IBM Q backends, see

Table 1. Ising Models with Degenerate Ground States to Be Studied on NISQ Hardware

Problem	Diagram	Ising Hamiltonian H_C	Ground States $ (q_0 := \uparrow) \dots q_{n-1}\rangle$
(a)		$- [Z_0(Z_1 + Z_2 - Z_3) + Z_1(Z_2 - Z_4) + Z_2(Z_3 + Z_4) + Z_3Z_4]$	$ \uparrow\uparrow\uparrow\uparrow\rangle, \uparrow\uparrow\uparrow\downarrow\rangle, \uparrow\uparrow\downarrow\downarrow\rangle$
(b)		$- [Z_0(2Z_1 + Z_2 + 2Z_3 + Z_4) + Z_1(-2Z_2 - Z_3 + Z_4) + Z_2(Z_3 + 2Z_4) - 2Z_3Z_4]$	$ \uparrow\uparrow\uparrow\uparrow\rangle, \uparrow\uparrow\uparrow\downarrow\rangle, \uparrow\uparrow\downarrow\downarrow\rangle, \uparrow\uparrow\downarrow\uparrow\rangle, \uparrow\downarrow\uparrow\uparrow\rangle, \uparrow\downarrow\uparrow\downarrow\rangle$
(c)		$- [Z_0Z_2 + Z_1Z_3 + Z_2(-Z_3 + Z_4 - Z_5) + Z_3(Z_4 - Z_5) + Z_4Z_5]$	$ \uparrow\uparrow\uparrow\uparrow\downarrow\rangle, \uparrow\downarrow\uparrow\downarrow\uparrow\rangle, \uparrow\downarrow\uparrow\downarrow\downarrow\rangle$
(d)		$- [Z_0Z_1 + Z_1(-Z_2 - Z_3) - Z_2Z_3]$	$ \uparrow\uparrow\downarrow\rangle, \uparrow\uparrow\uparrow\rangle, \uparrow\uparrow\downarrow\rangle$
(e)		$- [Z_0(-Z_1 - Z_2) - Z_1Z_2]$	$ \uparrow\uparrow\downarrow\rangle, \uparrow\downarrow\uparrow\rangle, \uparrow\downarrow\downarrow\rangle$
(f)		$- [-Z_0Z_1]$	$ \uparrow\downarrow\rangle$

Problems (a)–(d) are from [18]. All of the Ising models $H_C = -\sum J_{ij}Z_iZ_j$ have only quadratic terms and no linear terms. The dark red edges indicate a ferromagnetic $J_{ij} = +2$ coupling and the light red edges a ferromagnetic $J_{ij} = +1$ coupling; the light blue edges represent an antiferromagnetic $J_{ij} = -1$ coupling and the dark blue edges an antiferromagnetic $J_{ij} = -2$ coupling. Only ground states with $q_0 = \uparrow$ are listed as all models are symmetric under simultaneous \uparrow/\downarrow -flips. Note that a symmetry-breaking fixed setting $q_0 := \uparrow$ results in a Hamiltonian H'_C on qubits q_1, \dots, q_n (without q_0) with some linear terms.

Figure 2. Thus, the same exact circuit could be evaluated across many backends, and in many instances across many different qubits on the same backend.

For Problems (a) and (b), which involve five variables (and with $q_0 := \uparrow$ require only 4 qubits) we generated three distinct circuits, each on a different architecture: 4T, 4L, and 5T (see Figure 2 for notation). The 5T circuits for these problems employ an ancilla qubit, which allows us to discard any sample in which the ancilla was not measured in the \uparrow state. The remaining Problems (c)–(f) each only have a single circuit associated with them. Note that Problem (f), with the Hamiltonian Z_0Z_1 and ground states $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$, is uniquely simple. In this case we do not fix $q_0 := \uparrow$, nor do we employ Grover-QAOA, as the problem can be “solved” with an almost-trivial circuit of a




Fig. 2. The hardware architectures we transpose our circuits to: Linear Nearest Neighbor connectivities on 2,3,4 qubits (2L,3L,4L), as well as T-shaped connectivities on 4,5 qubits (4T,5T). These architectures form connectivity subgraphs of all IBM Q devices. By fixing $q_0 := \uparrow$, each n -qubit problem in Table 1 can be mapped to a $(n - 1)$ -qubit architecture.

Hadamard on qubit q_0 , a X -gate on qubit q_1 , and a consecutive CNOT with control q_0 and target q_1 . Due to the small number of qubits and gates, we use this circuit to explore the highest reaches of accuracy and fairness capable with the hardware.

See Table 2 for a brief description of all of the circuits, including their architectures, gate count, and β , γ parameters. We describe our compilation technique in detail in Appendix C, and give QASM representations of our circuits in the ancillary file associated to this note.

3 METHODOLOGY

We evaluated the previously described set of circuits across a wide swath of hardware on the IBM Q platform. Furthermore, we collected data in two large sets, one in October 2020 and the other in August 2021, allowing for an analysis of the change in fairness of the IBM Q platform over time. Since our interest was not in identifying the best performing qubits, but rather observing trends in circuit correctness vs. fairness, we evaluated the same circuit many times per chip, each time using a different subset of qubits. For example, testing the circuit for Problem (e), which has only 2 qubits, on a chip with 5 qubits in an LNN architecture would involve four separate evaluations (on the pairs of qubits $\{q_i, q_{i+1}\}$ for $i = 1 \dots 4$). The machines we tested, along with the number of qubit subsets for each architecture, are listed in Table 3. It is important to emphasize that there was no additional transpilation or optimization of these circuits beyond that discussed in the previous section.

We collected 40,960 shots for each circuit on each set of applicable qubits, which gave us sufficient data to achieve repeatable statistics. None of the hardware consistently reported fair results, i.e., hardware noise introduced significant biases for even the simplest circuits. In order to characterize how fair a given set of qubits performed with a given circuit, we adopted the metric: how many shots does one need of a given circuit on a given set of qubits in order to *reject* the hypothesis that these qubits are fair samplers at 95% significance level? This is equivalent to the question of: how many flips (on average) does it take to show that a coin is unfair? A 90/10 coin will take relatively few flips to reveal its bias, whereas a 51/49 coin will take many flips. We call this metric “number of shots to reject fair sampling.”

To calculate the number of shots to reject fair sampling, we utilized Pearson’s χ^2 test, which evaluates how likely a set of observations $\{O_i\}$ is given some expected null hypothesis values $\{E_i\}$. If a given circuit has d degenerate ground states, and $n_{g.s.}$ is the number of ground state observations (out of the 40,960 total shots), our null hypothesis is that

$$E_i = E = \frac{n_{g.s.}}{d}. \quad (2)$$

The χ^2 value is determined via

$$\chi^2 = \sum_i \frac{(O_i - E)^2}{E}. \quad (3)$$

Table 2. Details for 1-round Grover-QAOA Fair Sampling Circuits Studied on NISQ Hardware

Prob.	Arch.	# Rot.	# CNOT	(β, γ)	$\langle \beta, \gamma H_C \beta, \gamma \rangle$	GSP
(a)	4L	57	35	$(-1/2, -11/12) \cdot \pi$	-2.682 out of -4	0.498
(a)	4T	51	29	$(-1/2, -11/12) \cdot \pi$	-2.682 out of -4	0.498
(a)	5T	47	25	$(-1/2, -11/12) \cdot \pi$	-2.682 out of -4	0.498
(b)	4L	59	39	$(-11/15, -17/60) \cdot \pi$	-4.228 out of -5	0.846
(b)	4T	53	32	$(-11/15, -17/60) \cdot \pi$	-4.228 out of -5	0.846
(b)	5T	49	29	$(-11/15, -17/60) \cdot \pi$	-4.228 out of -5	0.846
(c)	5T	90	62	$(-23/60, 1/15) \cdot \pi$	-1.563 out of -4	0.215
(d)	3L	26	14	$(-5/12, 1/10) \cdot \pi$	-1.319 out of -2	0.702
(e)	2L	16	4	$(-23/60, 3/5) \cdot \pi$	-0.999 out of -1	1.000
(f)	2L	2	1	N/A	N/A	1.000

For each circuit, we list the problem from Table 1 being solved, the qubit connectivity from Figure 2 employed, the total number of single-qubit (i.e. rotation) gates, the total number of CNOT gates, the β, γ parameters used, the expectation value $\langle \beta, \gamma | H_C | \beta, \gamma \rangle$, and the probability of selecting a ground state (GSP).

The null hypothesis is rejected at a significance level based on the value of χ^2 as compared against the upper-tail critical values of the χ^2 distribution

$$f(x, k) = \frac{x^{k/2-1} e^{-x/2}}{2^{k/2} \Gamma(k/2)}, \quad (4)$$

where k is the number of degrees of freedom in the observed data. For example, with five degrees of freedom a χ^2 value of 11.070 indicates a rejection of the null hypothesis with 95% significance level.

In principle, the value of χ^2 is itself a measure of fairness for a set of observational data, however we found it a poor metric for our application. This is because the actual significance of a given χ^2 value is dependent on the size of the data set as well as the number of degrees of freedom. Since our circuits reported widely varying numbers of ground state observations, and had different numbers of degenerate ground states, comparing χ^2 values across circuits did not compare apples to apples. One could instead report the significance level at which the observational data rejects the fair sampling hypothesis, as that takes sample size and degrees of freedom into account. However, current hardware rejects the fair sampling hypothesis at significance levels so close to 100% that numerical differences between rejection levels were difficult to meaningfully interpret. Instead, determining how many samples from a distribution are necessary (on average) to reject fair sampling at 95% significance can be done with observational data of any size and degrees of freedom, and provides an intuitive measure of the fairness of the underlying qubits.

A formal description of our algorithm for determining the number of shots to reject fair sampling can be found in Appendix B, here we describe the method in general terms. Based on our observational data of a given circuit C on a set of qubits Q , we determined the relative frequency of observing each ground state. We then generated 1,000 synthetic samples, each containing n ground state observations matching the relative frequency in the observed data. For each sample we calculated χ^2 , then took the median value for the set of 1,000 samples.¹ We repeated for increasing values of n until we arrived at a median value of χ^2 matching the level necessary to reject the

¹We found that taking fewer than 1,000 samples resulted in inconsistent median values, and more than 1,000 samples resulted in slow performance.

fair sampling hypothesis with 95% statistical significance given the number of degenerate ground states. As an example, with this methodology a coin biased 60/40 requires an average of 74 flips to reject fair sampling.

4 RESULTS

As mentioned in the introduction, none of the qubits or circuits consistently reported fair sampling. Instead, we sought to understand the correlation between qubit fidelity and fair sampling. Our method for quantifying the degree of fair sampling for a circuit on a piece of hardware was described in the previous section. We used two complementary methods of quantifying qubit fidelity. The first is simply the frequency with which a given set of qubits accurately identify an optimal solution for the problem at hand. In this context, higher fidelity qubits are those that obtain a ground state more frequently. The benefit of this method is that it relies on direct observation of qubits solving a real problem, while the downside is that it can't be used to directly compare results from different problems. Our second method for quantifying qubit fidelity is to use the hardware error rates as reported by `qiskit`. This method, which we will discuss in more detail later in this section, has the benefit of more accurately comparing results from different circuits. However, the practical relevance of this reported error data for our application is unclear.

A broad issue for experimental tests of existing NISQ hardware is the highly stochastic nature of the results. We observed significant variance in all of our observed data, and therefore used a very large experiment count (with 7,440 distinct evaluations of the IBM Q hardware) to observe consistent trends. To be precise, each data point in the following plots is generated by specifying: a circuit to be evaluated, an IBM Q backend, and a subset of qubits on the backend which match the connectivity of the given circuit. We then collected 40,960 shots of the circuit on those specific qubits, and the results of those shots correspond to a single data point in the plots below. This large number of data points allowed us to identify repeatable positive or negative correlations, and we include lines of best fit to indicate these trends.

4.1 Ground State Probability vs. Fairness

With that proviso out of the way, let us now discuss the results of directly comparing quantity (as measured by the number of ground state observations) vs. quality (as measured by the number of shots to reject fair sampling). We observed a positive correlation between ground state probability and fair sampling for three circuits, those solving the “easy” problems: (d), (e), (f), see Figure 3. Note that these circuits all have ≤ 3 qubits and circuit depths 40, 20, and 3, respectively. We call this “soft” unfairness, in that improving ground state probability and fairness of sampling are both accomplished by evaluating the circuit on the highest performing qubits possible.

With the remaining circuits we observed more complicated behavior. For Problem (a), the 4T circuit exhibits “soft” unfairness, the 4L circuit shows no variation in fairness with respect to ground state probability, and the 5T circuit has a *negative* correlation between ground state probability and fairness, see Figure 4. In other words, qubits that more frequently obtained ground states did so with increasing bias. All circuits for Problems (b) and (c) share this negative correlation. We call this “hard” unfairness, as improving ground state probability results in less fair results, and vice versa.

We believe this “hard” unfairness comes from the fact that long circuits on noisy qubits produce increasingly random results. This means that the likelihood of selecting a ground state decreases, but the distribution across ground states is fairly even. To test this conjecture, we compared the fairness and ground state probabilities of both a “soft” and “hard” circuit as functions of the error rate of the qubits.




Fig. 3. Each data point represents 40,960 shots of the specified circuit on a fixed set of qubits on an IBM Q backend. The **Quantum Volume (QV)** of the backend is indicated by yellow for QV128, blue for QV32, green for QV16, and red for QV8. In these shorter circuits we observe a positive correlation (as indicated by the black lines of best fit) between ground state probability and number of shots to reject fair sampling.

4.2 Aggregate Error Rate vs. Fairness

The IBM Q hardware reports the latest calibration data for a given chip at the time of evaluating a circuit. Specifically, `qiskit` reports the error rates for their native gateset (CX, SX, RX, and RZ) as well as individual qubit measurement error. We used this data to generate what we call an aggregate error of the chip and circuit in combination. For a circuit C acting on qubits q_1, \dots, q_l , composed of a sequence of gates g_1, \dots, g_k evaluated on a chip H , this aggregate error $E_{C,H}$ is calculated via

$$E_{C,H} = 1 - \left(\prod_{i=1}^k (1 - e_i) \right) \left(\prod_{i=1}^l (1 - m_i) \right), \quad (5)$$

where e_i is the error for gate g_i , and m_i is the measurement error for qubit q_i , as reported by `qiskit` at the time of evaluation. We believe this aggregate error metric evaluates roughly the likelihood of success of a given circuit on a specific piece of hardware. However, the usefulness of the error rates as reported by `qiskit` has been drawn into question [25, 32]. These experiments therefore serve as a test of the relevance of this reported error value in the context of fair sampling.

In Figure 5 we see that both the soft and hard circuits report lower ground state probabilities as aggregate error rates increase, as expected.² For Problem (e), we see that increased error rates results in less fair results. In other words, the errors result in systemic bias in favor of some ground states over others. For Problem (b) on the 5T architecture, we see the opposite behavior: increased error rates result in fairer solutions. We can see in Figure 5 that the likelihood of obtaining a ground state is ≈ 0.4 in the high error rate case. There are six degenerate ground states for this problem, out of 16 total states. Therefore a circuit that selected a state at random would select a ground state with $6/16 = 0.375$ probability.

By combining all of the data in to a single plot we can get a bigger picture view of how structured vs. unstructured errors impact our results. In Figure 6 we see that fairness is highest for low aggregate error, decreases as the error approaches approximately 0.5, and then begins increasing again. This suggests that structured, i.e., biased, errors dominate in the middle of our error spectrum. As the overall error rate increases, either due to longer circuits or worse performing qubits, the net result is more unstructured error and thus fairer results.

²In fact, we observed a decrease in ground state probability as aggregate error increased for all circuits, which serves as a sanity check on our methodology.




Fig. 4. In these longer circuits, involving four or five qubits and between 72 and 152 gates, we mostly observe a negative correlation between ground state probability and number of shots to reject fair sampling, with Problem (a) on the 4L and 4T architectures as the exceptions.

4.3 Comments on the IBM Q System Performance

We believe these results constitute a serious evaluation of the IBM Q system. Our two most prominent observations based on this data are:

Correlation between Ground State Probability and Fairness shifted. By specifically comparing the data generated in October 2020 against the data generated in August 2021, we can see how the overall fairness of the IBM Q system has evolved. Interestingly, if one only looks at the 2020 data, all circuits from Problem (a) exhibit “hard” unfairness. Running the experiments again in 2021, the 4T circuit shifted into the “soft” regime, see Figure 7. Furthermore, all circuits saw the slope of the fairness vs. ground state probability trend lines increase. Meanwhile, the average ground state probability across all problems decreased slightly, from 44.7% to 43.4%. The average fairness increased significantly in the new data,




Fig. 5. Results for a hard circuit, Problem (b) on the 5T architecture, as compared to results for a soft circuit, Problem (e). All circuits studied featured a negative correlation between ground state probability and aggregate error, as pictured here for these two circuits in the leftmost plots. For hard circuits, increase in aggregate error leads to behavior reminiscent of a random number generator, decreasing ground state probability but increasing fairness, as indicated in the top right plot.

however this is due solely to a few data points in the two-qubit problems with extremely high ($O(10^8)$) number of shots to reject fair sampling. The median number of shots to reject fair sampling actually decreased when looking at the new data, from 190 to 142. From this data we conclude that the newer machines are more dominated by structured error as opposed to random error.

Quantum Volume only loosely correlates with fairness. The IBM metric of Quantum Volume [7] did broadly correlate with decreased aggregate error, which led to improved fairness in short circuits (see Figure 5). However, the Quantum Volume metric is only based off of the highly tuned performance of the best performing qubits available on a machine, instead of the overall hardware performance. This study, which ranges over all available qubits on a wide range of available chips, gives a more comprehensive view of performance. For example, we have plotted the results for Problem (d) evaluated on all 190 connected 3-qubit subsets on the manhattan backend (which has QV32 and 65 qubits), see Figure 8. In these plots one can see a wide range of results, and the best performing qubits for a single




Fig. 6. Fairness as a function of aggregate error rate across all experiments. Fairness is highest for small and large error rates, with a decrease in the middle. This indicates that structured errors dominate in the middle regime, while unstructured errors play an increasingly important role as error increases. The log-quadratic line of best fit (in black) is included to graphically indicate these overall trends.

metric (ground state probability, number of samples to reject fair sampling, and aggregate error rate) can perform quite poorly when evaluated by a different metric. Thus, this study emphasizes the need for more comprehensive benchmarking that does not only focus on random circuits evaluated on a small subset of qubits, as also discussed in [25].

5 MEASUREMENT ERROR MITIGATION

In this section we analyze how measurement error affects fairness. The measurement error included in our aggregate error metric (see Equation (5)) is for each individual qubit. This tensored approach paints an incomplete picture of measurement error, as there are correlations between qubits leading to correlated errors [28]. Therefore, we collected more thorough measurement error data immediately before evaluating our fair sampling circuits. Specifically, for a given n -qubit circuit we explicitly prepared each of the 2^n basis states and then measured the likelihood that the all of the other states would be observed. Using this data, we created a measurement error calibration matrix for each set of qubits and used the inverse of that matrix to correct for measurement error [26]. Due to the small number of qubits in our problems under consideration, this simple approach is feasible; we relied on it rather than more advanced measurement error mitigation schemes [11, 23, 28].

Our complete results are visualized in Figure 13, here we describe the results in general. When going from the raw to the mitigated results, we tracked the change in ground state probability,




Fig. 7. Results for Problem (a) on the 4T architecture. Data collected across the IBM Q system in October 2020 showed a negative correlation between ground state probability and number of shots to reject fair sampling, whereas data collected in August 2021 showed a positive correlation. This indicates that the newer dates are more dominated by structured errors.

number of shots to reject fair sampling, and the slope of the line of best fit correlating the two. We found mixed results in terms of ground state probability and fairness. Averaged over all problems, ground state probability improved by an average of 2.5%. Meanwhile, the mean number of shots to reject fair sampling decreased by an average of 19.1%. Finally, we observed that the slope of the line of best fit increased on 9 out of 10 problems, with an average improvement of 79.8%. These results show that accounting for measurement error slightly improves ground state probability at the expense of fairness. The increase in slope suggests that the mitigated results are more dominated by structured error than the raw data.

This result may appear somewhat surprising, as measurement error on these devices is itself a structured error: qubits prepared in the 0 state are far less likely to be measured as 1 than vice versa. However, in the raw data we observe that ground states that are composed mostly of 0 qubits are disfavored by an average of 2% (i.e., they appear 2% less frequently than expected from a perfectly fair sampler). Meanwhile, states that are composed of a majority of 1 qubits are favored by an average of 3.4%. It is unclear what errors are occurring during the evaluation of these circuits which lead to this bias in favor of 1, however it is likely measurement error is in fact helping results appear more fair by creating a slight counterbalancing bias in favor of 0. Indeed, after imposing our measurement error mitigation scheme we see that ground states that are composed mostly of 0 qubits are disfavored by an average of 5.9%, and those with mostly 1 qubits are favored by an average of 8%. By accounting for the measurement bias in favor of 0, we have highlighted the fact that at least some of the purported fairness of these devices was in fact due to multiple biases partially cancelling each other out.

We see from this simple example that fair results can be produced by three different quantum computers: those dominated by random noise, those with biased errors that nearly cancel each other out, and those with low error rates. When determining the quality of these and future quantum computers, detailed analysis such as this is helpful in understanding which of the above scenarios best describes the device at hand.




Fig. 8. Detailed results for Problem (d) evaluated on all 3-qubit subsets of the manhattan backend. Each plot represents the same underlying data, and the set of qubits with the maximum ground state probability (max GSP), maximum number of shots to reject fair sampling (max NSRFS), and minimum aggregate error (min AR) are highlighted. This plot shows the high variance in results from a single backend.

6 CONCLUSION

We have studied the extent to which today’s gate-based quantum computing hardware can fairly sample from the space of optimal solutions for discrete optimization problems. First, we introduced a novel metric for evaluating the fairness of current hardware: the number of shots to reject fair sampling. Using this metric, we found that hardware errors had a significant impact on fairness. For short circuits with ≤ 3 qubits, higher fidelity qubits resulted in increased fairness. We call these circuits “soft.” However, for longer circuits with ≥ 4 qubits, higher fidelity qubits actually resulted decreased fairness. We call these circuits “hard.” This is because low fidelity hardware begins behaving like a random number generator for long circuits, which contains no information but is relatively fair.

It would be interesting to test multiple quantum computing hardware technologies (e.g., optical, trapped ions, etc.) to see if they have different fairness profiles. Our studies show that individual machines on the IBM Q backend can have dramatically different fairness and ground state probabilities when different qubits are used. As quantum computing hardware improves, we hope and expect to see an increase in qubit fairness and consistency.

In this study we have also briefly explored the ways in which mitigating errors can affect fairness. Specifically, we observed that measurement error was partially correcting other structured errors. Measurement error mitigation therefore decreased fairness while modestly increasing ground state probability. We believe that studying fairness in addition to ground state probability can give a more nuanced picture of the errors affecting NISQ hardware. This is important for two reasons. First, random benchmark circuits are less susceptible to structured errors than circuits designed to solve realistic problems. And second, reported calibration and error rates reported are known to present an incomplete picture of how ordered circuits will perform [25]. This can be slightly improved upon by employing just-in-time noise data [32]. However, a broader theory of how different types of errors (read-out, crosstalk, etc.) affect not only ground state probability but also fairness is worth future study. This may involve utilizing different circuits or more refined statistical tools to highlight individual sources of error.

APPENDICES

A IBM Q BACKEND INFORMATION

Table 3. Details of IBM Q Backends to Be Studied

	Backend	QV	Qubits	2L	3L	4L	4T	5T
October 2020	manhattan	32	65	72	190	232	96	92
	montreal	32	27	28	74	80	48	36
	toronto	32	27	28	74	80	48	36
	rome	32	5	4	6	4	0	0
	santiago	32	5	4	6	4	0	0
	bogota	32	5	4	6	4	0	0
	valencia	16	5	4	8	4	6	2
	vigo	16	5	4	8	4	6	2
	melbourne	8	15	20	72	130	66	120
	ourense	8	5	4	8	4	6	2
August 2021	montreal	128	27	28	74	80	48	36
	mumbai	128	27	28	74	80	48	36
	manhattan	32	65	72	190	232	96	92
	brooklyn	32	65	72	190	232	96	92
	toronto	32	27	28	74	80	48	36
	sydney	32	27	28	74	80	48	36
	guadalupe	32	16	16	40	40	24	16
	bogota	32	5	4	6	4	0	0
	lagos	32	7	6	14	8	12	4
	jakarta	16	7	6	14	8	12	4
	lima	8	5	4	8	4	6	2

The name of the backend, the Quantum Volume (QV), the number of qubits, and the number of subsets of qubits on each device matching the relevant circuit architectures (2L, 3L, 4L, 4T, and 5T) are listed.

B ALGORITHM FOR NUMBER OF SHOTS TO REJECT FAIR SAMPLING

For the calculations in this paper we used $n_s = 40960$ and $n_i = 1000$. As explained in the main text, these values were chosen via a ‘goldilocks’ principle: much smaller, and the stochastic nature of the experiments resulted in inconsistent results, much larger and the computational burden was considerable. It is also worth noting that this algorithm becomes less useful as the samples become more evenly distributed, e.g., calculating the number of shots to reject fair sampling on a coin biased 50.001/49.999 takes $O(1\text{hr})$.

ALGORITHM 1: Number of Shots to Reject Fair Sampling

Input: Qubits Q , quantum circuit C for a problem P with k degenerate ground states labeled s_1, \dots, s_k ; n_s number of shots, n_i number of inner loops

Output: Number N of shots to reject fair sampling for C on Q .

```

1: Execute  $n_s$  shots of circuit  $C$  on qubits  $Q$ 
2:  $o_i :=$  the number of times ground state  $s_i$  observed
3:  $w_i := o_i / \sum_{i=1}^k o_i$  ▷ relative frequency of ground state  $s_i$ 
4:  $\text{chiSqToReject} :=$  the  $\chi^2$  value necessary to reject the fair sampling hypothesis with 95% statistical significance for  $k - 1$  degrees of freedom
5: function MEDIANCHISQ( $N$ )
6:    $\text{chiSqArray} := \{\}$ 
7:   while  $\text{length}(\text{chiSqArray}) < n_i$  do
8:      $\text{testSamples} := N$  samples drawn from the set  $\{s_i\}$  with weights  $\{w_i\}$ 
9:      $\text{testChiSq} := \chi^2$  for  $\text{testSamples}$ 
10:     $\text{append}(\text{chiSqArray}, \text{testChiSq})$ 
11:   end while
12:   return  $\text{median}(\text{chiSqArray})$ 
13: end function
14:  $N := 2$  ▷  $\chi^2$  ill-defined on sets of length 1
15: while  $\text{MEDIANCHISQ}(N) < \text{chiSqToReject}$  do
16:    $N := 2N$ 
17: end while
18:  $\text{upperBound} := N$ 
19:  $\text{lowerBound} := N/2$ 
20: while  $\text{upperBound} - \text{lowerBound} > 2$  do
21:    $N := \lfloor (\text{upperBound} + \text{lowerBound})/2 \rfloor$ 
22:   if  $\text{MEDIANCHISQ}(N) < \text{chiSqToReject}$  then
23:      $\text{lowerBound} := N$ 
24:   else
25:      $\text{upperBound} := N$ 
26:   end if
27: end while
28: return  $N$ 

```

C DETAILS ON CIRCUIT COMPILATION

In this appendix we provide supplementary information detailing the specifics of our compilation techniques. Full details of every hand compilation procedure would be too lengthy, instead we illustrate the compilation of Problem (d) as an example. Before we delve into details, we mention some broad ideas which aided in the compilation.

- **Circuit parameters:** The values for β and γ in Table 2 were found by conducting a grid search over the angles $(\beta, \gamma) \in [-\pi, 0) \times [-\pi, \pi)$ with a resolution of $\pi/60$ and calculating the expectation value $\langle \beta, \gamma | H_C | \beta, \gamma \rangle$ classically (where the search space for β was cut in half due to a $(\beta, \gamma) \cong (-\beta, -\gamma)$ symmetry of the expectation value). This was efficient due to the relatively low circuit depths involved; for more complex circuits the hybrid approach detailed in Section 2.2 is more appropriate.




Fig. 9. (left) A 1-level Grover Mixer QAOA for an unconstrained optimization problem, $U_S |\uparrow^n\rangle = |\rightarrow^n\rangle$, $U_M(\beta) = e^{-i\beta} |\rightarrow^n\rangle\langle\rightarrow^n|$, and the Hamiltonian $|\rightarrow^n\rangle\langle\rightarrow^n| \cong Id + \sum_i X_i + \sum_{\langle i,j \rangle} X_i X_j + \sum_{\langle i,j,k \rangle} X_i X_j X_k + \dots$ which has also been proposed to achieve fair sampling in quantum annealing [21], going beyond limited higher-order terms.

(right) For an unconstrained Ising problem with *only quadratic terms* in the Hamiltonian H_C on q_0, \dots, q_{n-1} , we break symmetries by setting $q_0 = |\uparrow\rangle$ and $U_S |\uparrow^n\rangle = Id \otimes H^{\otimes n-1} |\uparrow^n\rangle = |\uparrow \rightarrow^{n-1}\rangle$. The quadratic terms of H_C involving q_0 turn into linear terms of a Hamiltonian H'_C on q_1, \dots, q_{n-1} , thereby reducing q_0 to a classical control bit. The multi-control- $Z^{-\beta/\pi}$ gate is fully symmetric, thus we may swap controls and target.




Fig. 10. Two compilations of a large multi-control phase-shift gate (left) to gates with fewer controls: (middle) Direct decomposition into single-qubit phase-shift gates, large Toffolis, CNOTs, and a smaller multi-control phase-shift gate. (right) Decomposition using a $|\uparrow\rangle$ -initialized ancilla into which we compute an AND (i.e., a Toffoli with mismatched phases), followed by a smaller multi-control phase-shift gate and an uncomputation of the AND.

- **Multi-controlled phase-shifts:** First, we note that since terms in an Ising Hamiltonian pairwise commute, implementing the phase separating unitary can be done for each term individually, with 2 CNOTs and one phase-shift gate $Z^{\mp 2\gamma/\pi}$. Care has to be taken, however, to address the connectivity constraints in the architectures. More difficult is the compilation of the multi-control phase-shift gate at the heart of the Grover Mixer. Figure 10 gives two recursive compilations into gates with fewer controls, one without ancillas, and one with a $|\uparrow\rangle$ -initialized ancilla. We note that careful recursive calls to these compilations and using symmetries can result in some of the smaller multi-control gates to cancel each other.
- **SWAP operations:** We allow SWAP operations for qubits to change location in the circuit without having to go back to the original location. This is possible because the Grover Mixer is fully symmetric, i.e., we can arbitrarily exchange controls with the phase-shift gate, thus the input order to the Grover Mixer does not matter. We track qubit SWAPs with a permutation array from which in the end we read a read-out map for the measurements. This gives us a significant advantage over transpilation tools which are not aware of such symmetries. Introducing SWAPs (which usually cost 3 CNOTs) at the right place in a circuit (next to a CNOT) will cost only 1 CNOT. Similarly, moving SWAPs through to the beginning or end of the circuit results in complete removal of the SWAP cost, as the SWAP gate can be replaced with a change in the permutation array and corresponding readout.

Example: Problem (d)

Let us now describe in detail the compilation of Problem (d) to the 3L architecture. Starting with the general form of the 1-level Grover Mixer QAOA with fixed qubit $q_0 := \uparrow$, we implement each quadratic term of the Ising Hamiltonian with 2 CNOTs and a single-qubit phase-shift gate $Z^{\mp 2\gamma/\pi} = \begin{pmatrix} 1 & 0 \\ 0 & e^{\mp i\gamma} \end{pmatrix}$.




Fig. 11. Phase Separator Compilation, Problem (d): Implementing the state preparation U_S and the phase separator $U_P(\gamma) = e^{-i\gamma(-Z_0Z_1+Z_1Z_2+Z_1Z_3+Z_2Z_3)}$ on a Linear Nearest Neighbor architecture. Quadratic terms commute pairwise, so they can be implemented individually with 2 CNOTs and one phase-shift gate each. Circuit size is reduced in three ways: (i) Setting $q_0 := \uparrow$ results in q_0 acting as a classical control bit, which can be removed. (ii) Directly combining a single quadratic-term phase separator with a SWAP gate increases the CNOT count by 1 (instead of the usual 3). (iii) We keep track of the resulting permutation of the qubits instead of undoing the SWAPs.




Fig. 12. Grover Mixer Compilation, Problem (d): Implementing the grover mixer unitary $U_M(\beta)$ according to the high-level approach for multi-control phase-shift gates in Figure 10. Circuit size is reduced in two ways: (i) Moving SWAP gates through the circuit such that they appear at the beginning and the end, and thus can be tracked in the qubit permutation, rather than implemented. (ii) Moving (multi-)control gates through the circuit such that they can be cancelled with adjoint gates (not applicable in this example).

The CNOTs with control q_0 can be removed, leaving q_0 in the role of a classical bit, see Figure 11. We notice that the other three quadratic terms form a triangle in the problem graph, hence we cannot map them to a 3L Linear Nearest Neighbor architecture, but rather need at least one SWAP. Inserting the SWAP as 3 CNOTs right to an existing CNOT of the phase separating unitary $e^{-i\gamma(Z_2Z_3)}$ gives a cancellation of 2 CNOTs, thus the addition of the SWAP results in only 1 additional CNOT. Rather than reversing this SWAP at the end of $U_P(\gamma)$, we keep track of the permutation of the qubits (q_2 and q_3 have swapped places).

This end configuration (swapped qubits) now builds the start of the Grover Mixer unitary. The X-gates and the control on qubit q_0 can be dropped, as explained in Figure 9 (right). Using the recursive decomposition from Figure 10, we again see that we need to add a SWAP, see Figure 12. The remaining single-controlled phase-shift gate between the middle two circuit wires can be implemented with 2 CNOTs and 3 phase-shift gates as shown. Keeping track of the SWAP throughout the end of the circuit, we see that qubits q_2 and q_3 end up back in their original position—in general, this need not be the case.

Comment on 5-qubit Ising Problems on 5T Using an Ancilla. Since an ancilla offers no benefit for the implementation of the phase separating unitary, and due to limited connectivity in the considered architectures, we found that decomposing multi-controlled phase-shift gates using an ancilla as shown in Figure 10 (right) can only be used advantageously for the 5-qubit Ising Problems (a) and (b).

Following the approach outlined for Problem (d), both Problems (a) and (b) result in a multi-controlled phase-shift gate acting on 4 qubits. Using a remaining qubit on the 5T architecture as a $|\uparrow\rangle$ -initialized ancilla (which has to be swapped around), we can implement the mentioned decomposition, which should return the ancilla to the $|\uparrow\rangle$ -state. As mentioned in the primary text, we discarded all samples where the ancilla qubit was not measured in the $|\uparrow\rangle$ state.

D MEASUREMENT ERROR MITIGATION RESULTS




Fig. 13. Results of measurement error mitigation across all problems. Ground state probability improved by an average of 2.5%, number of shots to reject fair sampling decreased by an average of 19.1%, and the slope of the line correlating the two increased by an average of 79.8%.

REFERENCES

- [1] Mahabubul Alam, Abdullah Ash-Saki, and Swaroop Ghosh. 2020. Design-space exploration of quantum approximate optimization algorithm under noise. In *2020 IEEE Custom Integrated Circuits Conference (CICC)*. 1–4. <https://doi.org/10.1109/CICC48029.2020.9075903>
- [2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A. Buell, et al. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574, 7779 (2019), 505–510. <https://doi.org/10.1038/s41586-019-1666-5>
- [3] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. 2020. Obstacles to variational quantum optimization from symmetry protection. *Physical Review Letters* 125 (2020), 260505. Issue 26. arXiv:1910.08980, <https://doi.org/10.1103/PhysRevLett.125.260505>
- [4] Andreas Bärttschi and Stephan Eidenbenz. 2020. Grover mixers for QAOA: Shifting complexity from mixer design to state preparation. In *IEEE International Conference on Quantum Computing & Engineering (QCE'20)*. 72–82. arXiv:2006.00354, <https://doi.org/10.1109/QCE49297.2020.00020>
- [5] Andreas Bärttschi and Stephan Eidenbenz. 2021. Grover Mixer QAOA: Feasible State Superpositions, Convergence to Optimum, and Limitations. (2021). in preparation.
- [6] Jeremy Cook, Stephan Eidenbenz, and Andreas Bärttschi. 2020. The quantum alternating operator ansatz on maximum k-vertex cover. In *IEEE International Conference on Quantum Computing & Engineering (QCE'20)*. 83–92. arXiv:1910.13483, <https://doi.org/10.1109/QCE49297.2020.00021>
- [7] Andrew W. Cross, Lev S. Bishop, Sarah Sheldon, Paul D. Nation, and Jay M. Gambetta. 2019. Validating quantum computers using randomized model circuits. *Physical Review A* 100, 3 (2019), 032328. arXiv:1811.12926, <https://doi.org/10.1103/PhysRevA.100.032328>
- [8] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A quantum approximate optimization algorithm. *arXiv e-prints* (2014). arXiv:1411.4028.
- [9] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A quantum approximate optimization algorithm applied to a bounded occurrence constraint problem. *arXiv e-prints* (2014). arXiv:1412.6062.
- [10] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. 2000. Quantum computation by adiabatic evolution. *arXiv e-prints* (2000). arXiv:quant-ph/0001106.
- [11] Lena Funcke, Tobias Hartung, Karl Jansen, Stefan Kühn, Paolo Stornati, and Xiaoyang Wang. 2020. Measurement error mitigation in quantum computers through classical bit-flip correction. (7 2020). arXiv:quant-ph/2007.03663.
- [12] Lov K. Grover. 2005. Fixed-point quantum search. *Physical Review Letters* 95, 15 (2005), 150501. <https://doi.org/10.1103/PhysRevLett.95.150501>
- [13] Stuart Hadfield, Zhihui Wang, Bryan O’Gorman, Eleanor G. Rieffel, Davide Venturelli, and Rupak Biswas. 2019. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms* 12, 2 (2019), 34. arXiv:1709.03489, <https://doi.org/10.3390/a12020034>
- [14] Dylan R. Harp, Zhenxue Dai, Andrew V. Wolfsberg, Jasper A. Vrugt, Bruce A. Robinson, and Velimir V. Vesselinov. 2008. Aquifer structure identification using stochastic inversion. *Geophysical Research Letters* 35, 8 (2008), L08404. <https://doi.org/10.1029/2008GL033585>
- [15] Matthew P. Harrigan et al. 2021. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics* 17, 3 (2021), 332–336. arXiv:2004.04197, <https://doi.org/10.1038/s41567-020-01105-y>
- [16] Tadashi Kadowaki and Hidetoshi Nishimori. 1998. Quantum annealing in the transverse Ising model. *Physical Review E* 58, 5 (1998), 5355. arXiv:cond-mat/9804280, <https://doi.org/10.1103/PhysRevE.58.5355>
- [17] Vaibhaw Kumar, Casey Tomlin, Curt Nehr Korn, Daniel O’Malley, and Joseph Duiny III. 2020. Achieving fair sampling in quantum annealing. *arXiv e-prints* (2020). arXiv:2007.08487.
- [18] Mario S. Könz, Guglielmo Mazzola, Andrew J. Ochoa, Helmut G. Katzgraber, and Matthias Troyer. 2019. Uncertain fate of fair sampling in quantum annealing. *Physical Review A* 100, 3 (2019), 030303. arXiv:1806.06081, <https://doi.org/10.1103/PhysRevA.100.030303>
- [19] Salvatore Mandra, Zheng Zhu, and Helmut G. Katzgraber. 2017. Exponentially biased ground-state sampling of quantum annealing machines with transverse-field driving Hamiltonians. *Physical Review Letters* 118, 7 (2017), 070502. arXiv:1606.07146, <https://doi.org/10.1103/PhysRevLett.118.070502>
- [20] Yoshiki Matsuda, Hidetoshi Nishimori, and Helmut G. Katzgraber. 2009. Ground-state statistics from annealing algorithms: Quantum versus classical approaches. *New Journal of Physics* 11, 7 (2009), 073021. <https://doi.org/10.1088/1367-2630/11/7/073021> arXiv:0808.0365
- [21] Yoshiki Matsuda, Hidetoshi Nishimori, and Helmut G. Katzgraber. 2009. Quantum annealing for problems with ground-state degeneracy. *Journal of Physics: Conference Series* 143 (2009), 012003. <https://doi.org/10.1088/1742-6596/143/1/012003>

- [22] Mauro E. S. Morales, Timur Tlyachev, and Jacob Biamonte. 2018. Variational learning of Grover’s quantum search algorithm. *Physical Review A* 98, 6 (2018), 062333. arXiv:1805.09337, <https://doi.org/10.1103/PhysRevA.98.062333>
- [23] Paul D. Nation, Hwajung Kang, Neereja Sundaresan, and Jay M. Gambetta. 2021. Scalable mitigation of measurement errors on quantum computers. *arXiv e-prints* (2021). arXiv:2108.12518.
- [24] Daniel O’Malley. 2018. An approach to quantum-computational hydrologic inverse analysis. *Scientific Reports* 8, 1 (2018), 1–9. <https://doi.org/10.1038/s41598-018-25206-0>
- [25] Timothy Proctor, Kenneth Rudinger, Kevin Young, Erik Nielsen, and Robin Blume-Kohout. 2020. Measuring the capabilities of quantum computers. *arXiv e-prints* (2020). arXiv:2008.11294.
- [26] Qiskit Development Team. Qiskit Documentation: Measurement Error Mitigation. (n.d.). https://qiskit.org/documentation/tutorials/noise/3_measurement_error_mitigation.html.
- [27] Lukas M. Sieberer and Wolfgang Lechner. 2018. Programmable superpositions of Ising configurations. *Physical Review A* 97, 5 (2018), 052329. arXiv:1708.02533, <https://doi.org/10.1103/PhysRevA.97.052329>
- [28] Swamit S. Tannu and Moinuddin K. Qureshi. 2019. Mitigating measurement errors in quantum computers by exploiting state-dependent bias. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO’52)*. Association for Computing Machinery, New York, NY, USA, 279–290. <https://doi.org/10.1145/3352460.3358265>
- [29] Marc Vuffray, Carleton Coffrin, Yaroslav A. Kharkov, and Andrey Y. Lokhov. 2020. Programmable quantum annealers as noisy Gibbs Samplers. (2020). arXiv:quant-ph/2012.08827.
- [30] Samson Wang, Enrico Fontana, Marco Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J. Coles. 2020. Noise-induced barren plateaus in variational quantum algorithms. *arXiv e-prints* (2020). arXiv:2007.14384.
- [31] Zhihui Wang, Nicholas C. Rubin, Jason M. Dominy, and Eleanor G. Rieffel. 2020. XY mixers: Analytical and numerical results for the Quantum Alternating Operator Ansatz. *Physical Review A* 101, 1 (2020), 012320. arXiv:1904.09314, <https://doi.org/10.1103/PhysRevA.101.012320>
- [32] Ellis Wilson, Sudhakar Singh, and Frank Mueller. 2020. Just-in-time quantum circuit transpilation reduces noise. In *IEEE International Conference on Quantum Computing & Engineering (QCE’20)*. 345–355. arXiv:2005.12820, <https://doi.org/10.1109/QCE49297.2020.00050>
- [33] Jonathan Wurtz and Peter Love. 2021. MaxCut quantum approximate optimization algorithm performance guarantees for $p > 1$. *Physical Review A* 103 (2021), 042612. Issue 4. arXiv:2010.11209, <https://doi.org/10.1103/PhysRevA.103.042612>
- [34] Masayuki Yamamoto, Masayuki Ohzeki, and Kazuyuki Tanaka. 2020. Fair sampling by simulated annealing on quantum annealer. *Journal of the Physical Society of Japan* 89, 2 (2020), 025002. arXiv:1912.10701, <https://doi.org/10.7566/JPSJ.89.025002>
- [35] Theodore J. Yoder, Guang Hao Low, and Isaac L. Chuang. 2014. Fixed-point quantum search with an optimal number of queries. *Physical Review Letters* 113, 21 (2014), 210501. arXiv:1409.3305, <https://doi.org/10.1103/PhysRevLett.113.210501>
- [36] Brian Hu Zhang, Gene Wagenbreth, Victor Martin-Mayor, and Itay Hen. 2017. Advantages of unfair quantum ground-state sampling. *Scientific Reports* 7 (2017), 1044. arXiv:1701.01524, <https://doi.org/10.1038/s41598-017-01096-6>

Received January 2021; revised September 2021; accepted January 2022