

LA-UR-21-20101

Accepted Manuscript

Sampling on NISQ Devices: "Who's the Fairest One of All?"

Golden, John Kimbell

Baertschi, Andreas

O'Malley, Daniel

Eidenbenz, Stephan Johannes

Provided by the author(s) and the Los Alamos National Laboratory (2023-07-31).

To be published in:

DOI to publisher's version: 10.1109/QCE52317.2021.00038

Permalink to record:

<https://permalink.lanl.gov/object/view?what=info:lanl-repo/lareport/LA-UR-21-20101>



Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Triad National Security, LLC for the National Nuclear Security Administration of U.S. Department of Energy under contract 89233218CNA000001. By approving this article, the publisher recognizes that the U.S. Government retains nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Sampling on NISQ Devices: “Who’s the Fairest One of All?”

Elijah Pelofske*, John Golden*, Andreas Bärtschi*, Daniel O’Malley† and Stephan Eidenbenz*

* CCS-3 Information Sciences, Los Alamos National Laboratory, Los Alamos, NM 87544, USA

Email: epelofske@lanl.gov, golden@lanl.gov, baertschi@lanl.gov, eidenben@lanl.gov

† EES-16 Earth Sciences, Los Alamos National Laboratory, Los Alamos, NM 87544, USA

Email: omalled@lanl.gov

Abstract—Modern NISQ devices are subject to a variety of biases and sources of noise that degrade the solution quality of computations carried out on these devices. A natural question that arises in the NISQ era, is how fairly do these devices sample ground state solutions. To this end, we run five fair sampling problems (each with at least three ground state solutions) that are based both on quantum annealing and on the Grover Mixer-QAOA algorithm for gate-based NISQ hardware. In particular, we use seven IBM Q devices, the Aspen-9 Rigetti device, the IonQ device, and three D-Wave quantum annealers.

For each of the fair sampling problems, we measure the ground state probability, the relative fairness of the frequency of each ground state solution with respect to the other ground state solutions, and the aggregate error as given by each hardware provider. Overall, our results show that NISQ devices do not achieve fair sampling yet. We also observe differences in the software stack with a particular focus on compilation techniques that illustrate what work will still need to be done to achieve a seamless integration of frontend (i.e. quantum circuit description) and backend compilation.

I. INTRODUCTION

While finding *any* optimal solution is usually the goal when solving optimization problems, sampling fairly from all optimal solutions is an essential component of many real-life optimization applications, such as satisfiability-based probabilistic membership filters [2], detecting equally likely fluid flow outcomes in subsurface modeling problems [10], [16], and generally in engineering/physics contexts where the objective function does not explicitly encode all design goals.

Fair sampling from all optimum solutions has been proposed [7] as a benchmark problem for Noisy Intermediate Scale Quantum (NISQ) devices, inspired by theoretical work on fairness in quantum annealing [13]. In this paper, we thoroughly examine the state of fair sampling on a wide range of NISQ platforms by considering both the probability that the device finds an optimum solution (which we call ground-state probability) and the fairness metric from [7] based on a statistical test on how many runs are needed to reject the hypothesis that device is actually a fair sampler. The NISQ platforms that we consider are seven backends from IBM Q ranging in Quantum Volume from 8 to 128, the Aspen-9 Rigetti device, an IonQ device and three different D-Wave

Research presented in this article was supported by the Laboratory Directed Research and Development program of Los Alamos National Laboratory under project number 20200671DI. LA-UR-21-25803

quantum annealing devices. We accessed IonQ, Rigetti, and two of the D-Wave devices through the Amazon Braket cloud service and had direct access to IBM Q and the third D-Wave device.

We choose five standard small optimization problems (from [13] and [7]) that are paradigmatic examples for fair sampling formulated as Ising problems, which can be solved by the Grover Mixer Quantum Alternating Operator Ansatz (GM-QAOA) [3] for general gate-level quantum devices (i.e., IBM Q, Rigetti, and IonQ) and by the standard annealing algorithm for Ising problems on D-Wave quantum annealers. The GM-QAOA algorithm theoretically guarantees fair sampling, while the quantum annealing algorithm only samples fairly at very short annealing times at the cost of lower ground state probabilities (see theoretical analysis in [13]). We present this background information in more detail in Section II. Our results give a good glimpse of how far practice and theory still diverge, due to hardware and software limits.

If we had error-corrected quantum computing, we could just code up our five example problems in a standard quantum circuit language, such as IBM’s QISKit, Amazon Braket, or Rigetti’s Quil and compile it into a standard intermediate representation (such as Microsoft’s QIR or IBM’s QASM) and then let a backend compiler turn it into machine code for the individual vendor-specific backends. The reality of 2021, however, is that such compilation can be done, but it results in lengthy circuits whose execution essentially returns just noise on current NISQ platforms. Until we have more advanced optimization passes for quantum compilers, we find that hand-optimization is a necessity. We describe our resulting circuits in Section III. While we will go into some detail regarding the compilation choices for IBM Q and D-Wave in separate Sections V-B and V-C, we report our comparative results along the two axes of probability of returning ground state and fairness in Section V. Our main findings are:

1) *IonQ achieves the highest fairness and ground state probability for Problems (a) and (b):* IonQ has the highest fairness and ground state probability among the gate model devices for Problems (a) and (b).

2) *IBM Q achieves the highest fairness for Problems (d) and (e) and highest ground state probability for Problems (d):* The aggregate results for IBM Q have the highest fairness for problems (d) and (e). IBM Q aggregate results also have

Problem with Diagram	Ising Hamiltonian H_C to be minimized	Ground States $ (lq_0 := \uparrow) \dots lq_{n-1}\rangle$	1-round GM-QAOA Optimum $(\beta, \gamma) \div \pi$	$\langle \beta, \gamma H_C \beta, \gamma \rangle$	GSP
(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\uparrow\rangle, \uparrow\uparrow\uparrow\downarrow\downarrow\rangle, \uparrow\uparrow\downarrow\downarrow\downarrow\rangle$	$(\frac{1}{2}, \frac{11}{12})$	-2.682 out of -4	0.498
(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\uparrow\rangle, \uparrow\uparrow\downarrow\uparrow\downarrow\rangle, \uparrow\uparrow\downarrow\downarrow\uparrow\rangle, \uparrow\downarrow\uparrow\uparrow\uparrow\rangle, \uparrow\downarrow\uparrow\uparrow\downarrow\rangle$	$(\frac{11}{15}, \frac{17}{60})$	-4.228 out of -5	0.846
(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\uparrow\uparrow\rangle, \uparrow\downarrow\uparrow\downarrow\downarrow\rangle$	$(\frac{23}{60}, -\frac{1}{15})$	-1.563 out of -4	0.215
(d)		$-[Z_0Z_1 + Z_1(-Z_2 - Z_3) - Z_2Z_3]$ $ \uparrow\uparrow\uparrow\downarrow\rangle, \uparrow\uparrow\downarrow\uparrow\rangle, \uparrow\uparrow\downarrow\downarrow\rangle$	$(\frac{5}{12}, -\frac{1}{10})$	-1.319 out of -2	0.702
(e)		$-[Z_0(-Z_1 - Z_2) - Z_1Z_2]$ $ \uparrow\uparrow\downarrow\rangle, \uparrow\downarrow\uparrow\rangle, \uparrow\downarrow\downarrow\rangle$	$(\frac{23}{60}, -\frac{3}{5})$	-0.999 out of -1	1.000

TABLE I: Ising models with degenerate ground states to be studied on NISQ hardware. **Problems** (a)–(d) are from [13]. 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. All **Ising Hamiltonians** $H_C = -\sum J_{ij}Z_iZ_j$ have only quadratic terms and no linear terms. Note that a symmetry-breaking fixed setting of the logical qubit $lq_0 := \uparrow$ results in a Hamiltonian H'_C on logical qubits lq_1, \dots, lq_{n-1} (without lq_0) with some linear terms. Only **Ground States** with $lq_0 = \uparrow$ are listed as all models are symmetric under simultaneous \uparrow / \downarrow -flips. The right column shows the **Optimum Energy** $\langle \beta, \gamma | H_C | \beta, \gamma \rangle$ of a 1-round Grover Mixer QAOA found with a fine grid search for angles (β, γ) with a grid resolution of $\frac{1}{60} \cdot \pi$ and the corresponding ground state probabilities.

the highest ground state probability for Problem (d), although IonQ has the highest ground state probability for Problem (e).

3) *As expected, D-Wave finds ground states most reliably compared to 1-round GM-QAOA:* Allowing for more QAOA rounds and thus longer circuits would increase the theoretical optimum achievable, but in practice would lead to a further decrease in ground state probabilities as the noise just dominates the algorithmic improvement of additional rounds.

4) *D-Wave achieves highest fairness for Problem (c):* While the gate model backends behave similarly for Problem (c), D-Wave achieves the highest fairness metric.

5) *Rigetti's circuits have the highest aggregate error rate:* Rigetti's relatively poor performance can be attributed to its high aggregate error rates. Using vendor provided qubit-level error rates, we calculate the aggregate error of each circuit from the examples. Rigetti's error rates are usually higher than 80 %, IonQ's is around 50 %, and IBM Q error rates span a wide range, largely corresponding to the different device generation and achieved Quantum volumes.

Details: The detailed scatter plots in Section V provide more insights. The Pareto front of backends in a fairness vs. ground-state probability plot is dominated by IBM backends with a few notable IonQ experiments interspersed. If we include D-Wave, D-Wave dominates the Pareto front on the ground state probability but only at a typically low fairness level. Recall that a data point is on the Pareto front if no other point exists that is better in both fairness and ground state probability.

We present more detailed results on the IBM Q compiler environment in connection with our examples in Section V-B. The compiler passes and flags that IBM Q QISKit offers produce at times unpredictable results. We look at the aggregate circuit error that we can compute without actually running the circuit. We find that selecting noise adaptive compilation option reduces errors often, albeit increasing it in other cases. A similar effect can be observed when manually selecting the circuit topology preserving initial qubit layout for the circuits. Since the changes in aggregate error rates for these circuits can achieve reductions from 80 % error to less than 30 %, our

Name	LNN	5T	5P	6A	7H	Clique
Hardware Sub-Topology	$pq0-pq1-pq2$ $pq2$ $pq3$	$pq0-pq1-pq4$ $pq2$ $pq4$	$pq0-pq3$ $pq1-pq2$ $pq4$	$pq0-pq3$ $pq1-pq2$ $pq4$ $pq5$	$pq0$ $pq4$ $pq1-pq2-pq3$ $pq5$ $pq6$	all-to-all
Devices	All IBMQ, Rigetti	All IBMQ, Rigetti	IBMQ Melbourne, Rigetti	IBMQ	IonQ	

TABLE II: Hardware topologies we compile our circuits to. The topologies correspond to graphs of physical qubits pq_0, \dots, pq_6 available as sub-topologies on IBM, Rigetti, and IonQ devices: LNN (Linear Nearest Neighbor), 5T, 5P, 6A, 7H (labeled by shape) and Clique (full connectivity). We compile the smaller problems (d)–(e) to LNN and the larger problems (a)–(c) to all topologies they fit on, using only the available gates of the corresponding device(s). For D-Wave embedding, see Section IV-A.

results suggest that users should always test a large number of compiler options. We also look at QISKit’s four different optimization levels for compilation (0,1,2,3,). These levels generally reduce aggregate error by a few percentage points as we increase the optimization level; however, optimization level 1 shows erratic behavior at times tripling error rates over optimization level 0, but at other times actually beating even level 3 by 10%.

Section V-C presents more detailed results on studying the effect of varying the annealing time parameter on D-Wave from 1 to 300 microseconds. In three of our examples, we see the theoretically predicted increase in ground state probability between 1 and 10 microseconds at a significant level. Overall we conjecture that even the minimum value of 1 microsecond for annealing time is already too large to still allow for fairness, thus real-life D-Wave is good at finding optimum solutions, but does not sample fairly among them.

Overall, our results characterize the state of fair sampling on a good subset of all available NISQ devices. While IBM and IonQ outperform Rigetti and D-Wave in our fairness metric, we want to be clear that none of the tested NISQ platforms are very fair in the first place. Much work remains to be done in quantum hardware design to truly achieve fair sampling.

II. BACKGROUND: SETTING THE STAGE

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. We follow the tradition of analyzing and assessing fair sampling in the quantum computing context as defined for quantum annealing [13] and general gate-level quantum computing [7]. Taking the paradigmatic examples suggested by these earlier works, we define five standard examples of optimization problems, whose graph representation is shown in the first column of Table I with the formal objective function expressed as Ising Hamiltonians in the second column, and the resulting optimal solutions or ground states (expressed in an \uparrow, \downarrow basis) in the third column. Problems (a)–(d) all exhibit theoretically understood biased sampling with quantum annealing, so they are a somewhat natural non-trivial challenge for fair sampling algorithms. We added Problem (e) as an even smaller test case for our studies. Since all problems are less than 7 qubits, it appears plausible that they could be solved with reasonably well on current NISQ hardware.

By way of giving background, basic quantum annealing [11] is known to not always sample degenerate ground states fairly from a theoretical [15] and experimental [14] perspective. Several approaches to improve fairness of quantum annealing have been made [12], [13], [17], [18]. An additional suggestion is to add to the transverse field driver Hamiltonian $\sum X_i$ all higher-order Pauli- X terms, resulting in $\sum X_i + \sum X_i X_j + \sum X_i X_j X_k + \dots = |\rightarrow^n\rangle \langle \rightarrow^n|$ [15]. The study of limited higher-order driver Hamiltonians (i.e., without going over the complete sum), led Könz et al [13] to introduce the paradigmatic Problems (a)–(d). As Ising problems consisting of sums of quadratic ZZ terms, the D-Wave quantum annealer takes these as native inputs. For gate-level NISQ devices, we need an appropriate algorithm to solve the problems. The most natural candidate for this is the Grover Mixer QAOA (GM-QAOA) algorithm [3], which in theory samples fairly.

GM-QAOA is a variation of the Quantum Alternating Operator Ansatz [9]. In its essence, for a problem instance I with feasible solution states F and cost Hamiltonian H_C on n qubits, a p -round QAOA prepares a parameterized state from which it samples 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 a superposition of all feasi-

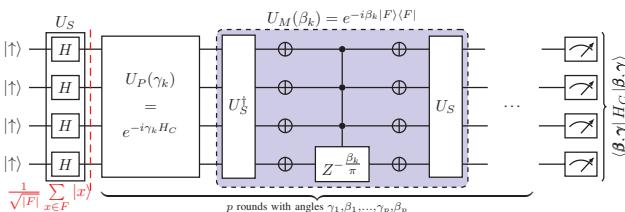


Fig. 1: GM-QAOA: State preparation U_S gives equal superposition of all feasible states $|F\rangle = |F|^{-1/2} \sum_{x \in F} |x\rangle$. U_S and U_S^\dagger are used to implement the Mixer $U_M(\beta) = e^{-i\beta |F\rangle\langle F|}$. For unconstrained problems, we have $U_S |\uparrow^n\rangle = H^{\otimes n} |\uparrow^n\rangle = |\rightarrow^n\rangle$ and $U_M(\beta) = e^{-i\beta |\rightarrow^n\rangle\langle \rightarrow^n|}$. The multi-control- $Z^{-\beta/\pi}$ gate is fully symmetric, thus we may swap controls and target.

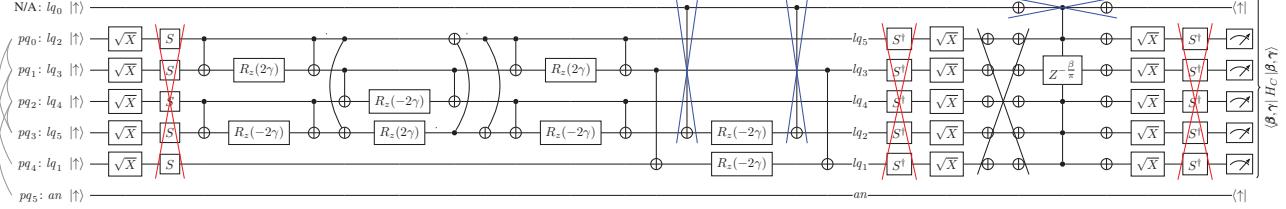


Fig. 2: Implementation of State Preparation U_S and the Phase Separator U_P for Problem (c) on a 6A-connectivity (pq_0, \dots, pq_5 , light gray) for IBM Q: The logical qubit $lq_0 := |\uparrow\rangle$ is fixed, effectively making all its gates and controls **redundant**. **(Phase separator)** The weighted Ising terms $w Z_i Z_j$ in the cost Hamiltonian H_C of the phase separator $U_P = e^{-i\gamma H_C}$ pairwise commute and can thus be implemented individually with 2 $CNOT$ s and 1 $R_z(-2\cdot w \cdot \gamma)$. IBM Q Melbourne does not contain a 4-clique in its topology, only a 4-cycle, hence implementing the 4-clique of logical qubits lq_2, lq_3, lq_4, lq_5 needs at least one $SWAP$. We incorporated it into the phase separator on logical qubits lq_2, lq_5 (drawn with curved $CNOT$ s), resulting in a permuted assignment to the physical qubits at the end. **(State preparation)** Given $H|\uparrow\rangle = |\rightarrow\rangle = S\sqrt{X}|\uparrow\rangle$, the state preparation unitary and its inverse have been rewritten as $U_S = (S\sqrt{X})^{\otimes n-1}$ and $U_S^\dagger = (\sqrt{X}S^\dagger)^{\otimes n-1} = (X\sqrt{X}S^\dagger)^{\otimes n-1}$. The phase shifts S and S^\dagger commute through the phase separator and cancel, and the action before measurement in the Z -basis only affects phase not probability; hence they can be **removed**. **(Grover Mixer)** Schematics only, for implementation see Fig. 3.

ble solutions F , 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 Fig. 1. The phase separating unitaries $U_P(\gamma)$ 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) $U_P(\gamma) \cong e^{-i\gamma H_C}$, where H_C is an Ising Hamiltonian with quadratic and linear terms as found in the second column of Table I for our example problems. The Grover mixer unitary [3] requires an efficient state preparation unitary U_S of an equal superposition of all feasible basis states $|F\rangle = 1/\sqrt{|F|} \sum_{x \in F} |x\rangle$. U_S can be used to design a mixing unitary U_M resembling Grover's selective phase-shift operator [1], [8], [19], where $U_M(\beta) = e^{-i\beta|F\rangle\langle F|} = Id - (1 - e^{-i\beta})|F\rangle\langle F| = U_S(Id - (1 - e^{-i\beta})|\uparrow\rangle\langle\uparrow|)U_S^\dagger$.

Grover Mixer QAOA samples fairly: 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 the same weighted average of the amplitudes, $((1 - e^{-i\beta})/\sqrt{|F|})\langle F|U_P(\gamma)U_S|\uparrow\rangle$. Therefore, basis states with the same energy are sampled with the same amplitude. For a complete proof, see [3].

III. DESIGNING CIRCUITS FOR NISQ DEVICES

Our general procedure for generating test circuits was to begin with a 1-round Grover Mixer QAOA algorithm in order to keep circuit depth low. Following [13], we fix $q_0 := \uparrow$ as all of the models in Table I are symmetric under simultaneous \uparrow / \downarrow swaps on all qubits. This reduces the problems from Ising Hamiltonians H_C with only quadratic terms acting on logical qubits lq_0, \dots, lq_{n-1} to a new Hamiltonians H'_C with some linear terms acting on fewer qubits lq_1, \dots, lq_{n-1} . This

also transforms lq_0 into a classical control bit for the consecutive Grover Mixer, which can be removed. Thus we can embed Ising problems without linear terms on n qubits onto circuits with only $n - 1$ qubits, see Fig. 2. The optimum energy $\langle \beta, \gamma | H_C | \beta, \gamma \rangle$ achievable with a 1-round Grover Mixer QAOA was found with a fine grid search for angles (β, γ) with a grid resolution of $\pi/60$; the values with corresponding ground state probabilities can be found in Table I.

We then compiled the circuits to match connectivity and gates available for each tested device individually. The available hardware topologies are listed in Table II. For IonQ and Rigetti through Braket, there is no available compilation to the available 1- and 2-qubit gates. Compared to IBM's qiskit compiler tools, we found that manual compilation reduced circuit depth by roughly a factor of two.

Hence we used compilation by hand down to the available 1- and 2-qubit gates for each individual device. To reduce circuit depth as much as possible, we made use of three techniques: (i) use the available 1- and 2-qubit gates wherever possible, (ii) when using $SWAP$ gates, classically track the permuted assignment of logical to physical qubits instead of restoring the original assignment in the circuit, (iii) use ancilla qubits whenever they help bring down the circuit depth of the Grover Mixer implementation. We demonstrate our techniques in detail for the circuit compilation of Problem (c) on topology 6A for IBM Q using gates $X, \sqrt{X}, CNOT, R_z(\theta)$ in Figures 2 and 3. We provide accompanying interactive Quirk circuits for these figures [here](#) (state preparation & phase separator) and [here](#) (multi-control- Z^{-t} phase shift unitary matrix illustrated by Quirk's amplitude display via state channel duality).

IV. METHODOLOGY / EXPERIMENTAL DESIGN

We briefly explain our experimental methodology and the experimental design. In general we send the vendor and backend-independent circuits from Section III to the compiler for each of the vendors, calculate the aggregate error on each circuit, and then execute a large number of runs on

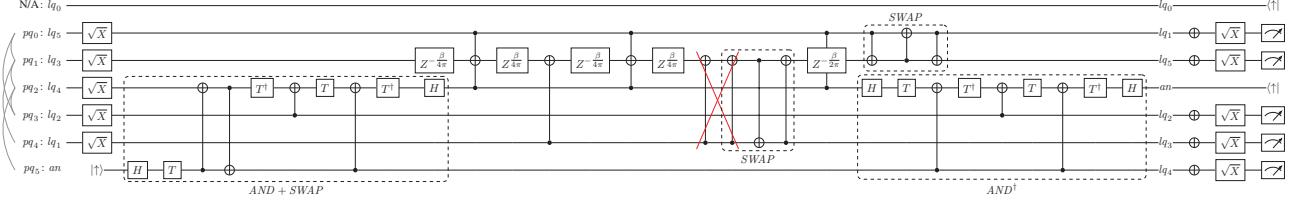


Fig. 3: Grover Mixer implementation details for logical qubits lq_1, \dots, lq_5 with fixed $lq_0 := |\uparrow\rangle$ on 6A-connectivity (pq_0, \dots, pq_5 , light gray) for IBM Q: **(AND gates)** To reduce the 4-control- $Z^{-\beta/\pi}$ from Fig. 2 to fewer controls, we first compute the logical $AND(lq_2, lq_4)$ into an ancilla qubit an , and swap it to pq_2 . The task is now to implement a 3-control- $Z^{-\beta/\pi}$ phase shift on central qubit lq_3 on pq_1 with controls on its neighbors lq_5, an, lq_1 on pq_0, pq_2, pq_4 , before uncomputing the AND . **(3-control phase shift)** We use the standard decomposition into single-qubit phase shifts on lq_3 interleaved with $CNOT$ and **TOFFOLI** gates, plus a recursively smaller 2-control- $Z^{-\beta/2\pi}$ on the control qubits lq_5, an, lq_1 . As lq_1 is not adjacent to lq_5, an , we first swap it to the central qubit pq_1 . Note that due to gate cancellations, this $SWAP$ only increases the $CNOT$ count by 1. **(Recursion)** We recursively decompose the 2-control- $Z^{-\beta/2\pi}$ (decomposition not shown). Similarly to before, we will need another $SWAP$ (shown) to implement the final base of a 1-control- $Z^{-\beta/4\pi}$ (not shown). **(TOFFOLI gates)** For the two **TOFFOLI**s, we can avoid such a further recursion (decompositions not shown): We will implement each **TOFFOLI** with 2 Hadamard gates on the target conjugating a 2-control- Z . If we interpret the second of these 2-control- Z as a 2-control- Z^{-1} (as $Z^{-1} = Z$), the resulting recursive 1-control- $Z^{1/2}$ and 1-control- $Z^{-1/2}$ of the two **TOFFOLI** implementations cancel. Finally, we can replace all (non-controlled) phase-shift gates by (up to global phase) equivalent R_z rotations, native to IBM Q.

the 12 backends in Table III with different circuit topologies described in Table II, collecting data to calculate our metrics of ground state probability (GSP) and our fairness metric called “number of shots to reject the fair sampling hypothesis”. We describe each of these individual steps briefly, particularly highlighting differences among the four vendors.

	Backend	Type	QV	Qubits
IBM	melbourne	Gate model	8	15
IBM	casablanca	Gate model	32	7
IBM	guadalupe	Gate model	32	16
IBM	toronto	Gate model	32	27
IBM	sydney	Gate model	32	27
IBM	manhattan	Gate model	32	65
IBM	montreal	Gate model	128	27
IonQ	ionq	Gate model		11
Rigetti	Aspen-9	Gate model		31
DWave	LANL 2000Q	Annealer		2032
DWave	2000Q-6	Annealer		2048
DWave	Advantage	Annealer		5760

TABLE III: Device properties for the NISQ backends we studied. Only IBM Q uses the Quantum Volume metric (QV) to characterize their systems.

A. Backend Compilation

IBM Q: We selected seven IBM Q backends shown in Table III, which all allowed us to embed all five fair sampling problems onto the hardware connectivity, that is without having to resort to extensive swapping. We implemented the circuits of the five test problems in IBM’s QISKit (version 0.25.0). QISKit has a method that allows you to compile the logical circuit onto the specified hardware using the native

gateset. We run each circuit using three different compiler arguments: (i) simply leaving everything to default, except supplying the specific *backend*; (ii) setting the *layout method* as *noise-adaptive*, setting the *optimization level* to 3, default values otherwise; (iii) supplying an argument for *initial-layout*, and set *optimization level* to 3. The initial layout we supply is a random subgraph of the hardware connectivity graph which is **isomorphic** to the original circuit graph structure (e.g. 5T, 7H, etc). We used the circuit topologies as described in Table II.

Rigetti: We accessed the IonQ system through the AWS Braket cloud service. AWS Braket defines its own Python-based quantum circuit language. In order to successfully run on the Aspen-9 device, we found it necessary to hand-compile some aspects of the backend independent circuits as AWS Braket did not transform the circuits supplied into the native gateset used by Rigetti. Rigetti has a compilation step from the provided circuit into their own Quil language. For each circuit we submit through AWS Braket, we get that compiled Quil code as part of the task metadata. The compilation step involves re-mapping the logical circuit onto a set of physical qubits. We found that this compilation step significantly changes the structure of the original circuit. The circuit topologies that were run on the Aspen-9 device are given in Table II. For problem (b), both 5T and 5P were embeddable onto the Aspen-9 hardware, so both topologies were used. We used the 6A circuit topology because 7H is not embeddable onto the Aspen-9 connectivity.

IonQ: In order to access the IonQ system, we used AWS Braket. IonQ accepts a different gate set than Rigetti, so we had to again hand-compile the circuits. IonQ’s all-to-all topology made this relatively easy.

DWave: The circuits we run on D-Wave are in the form of Ising models. Recall that Table I column 3 **Ising Hamiltonian** H_C shows the exact Ising models we use for these

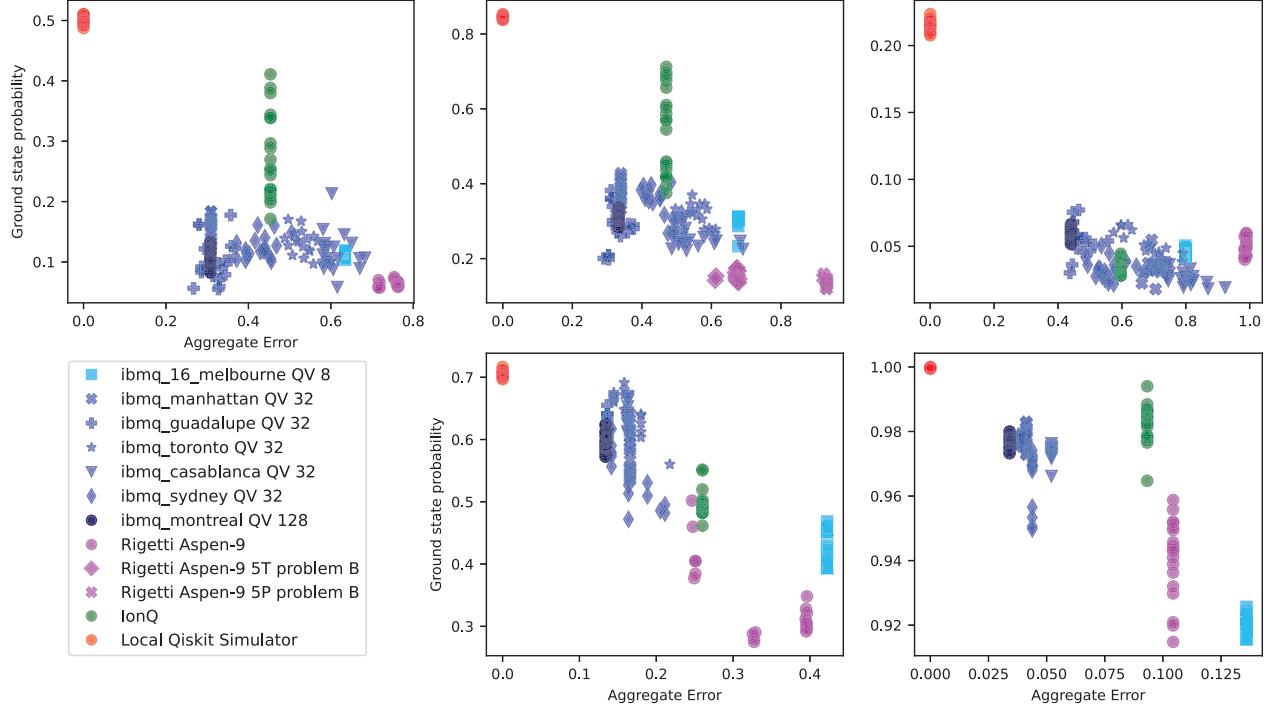


Fig. 4: Aggregate error vs GSP for Problems (a) (top left), (b) (top center), (c) (top right), (d) (bottom center), (e) (bottom right).

experiments.

When using quantum annealing, if the Ising model connectivity does not embed onto the hardware connectivity, then we must minor-embed the Ising model onto the hardware graph. We use a D-Wave’s standard embedding method called *minorminer* as well as a hand-tuned embeddings.

When we use minor-embedding, we also need to use a parameter called **chain strength** in order to ensure that physical qubits representing a logical variable (we refer to these physical qubits as a chain) take the same value, (meaning they do not disagree on the logical variable state). If a chain of physical qubits disagrees, then this is called a **broken chain**. In order to resolve a broken chain, in our experiments we use the *majority vote* method. Additionally, we compute chain strength using the default method in D-Wave’s Ocean SDK.

For the standard embedding experiments, we use a random minor-embedding found by D-Wave’s *minorminer* [4]–[6], with potentially unequal chain lengths. Using the full Isings means that the results we get from D-Wave include not only the ground states shown in the 3rd column of Table I, but also their complement, thus there are two ways of quantifying ground state bias: (i) treat the complement of each ground state as a separate ground-state; (ii). treat the complement of each ground state as being the same as the original. Call the first option *separate ground-state complements* and the second option *combined ground-state complements*. We ran standard embedding experiments only on the LANL D-Wave 2000Q.

For manual embedding experiments, we designed *balanced embeddings* across all devices and problems. In order to achieve this, we first fixed the value of qubit 0 on all problems to be 1 (note that for the gate model problems, we also fix the state of qubit 0). In this process, the quadratic terms ($q_0, q_i : n$) associated with qubit 0 became linear terms on the neighboring qubit ($q_i : n$). Fixing qubit 0 to be 1 allowed us to then have chain length 2 embeddings for both 2000Q backends, and chain length 1 (i.e. native embedding) embeddings for Advantage-system1.1. Fixing qubit 0 to be 1 also causes the ground state behavior to change from above. Specifically, now there are no complement ground states; only the original groundstates shown in column 4 of Table I (except missing qubit 0 because we fixed its state). We ran this experiment on all three D-Wave backends.

We also varied the *annealing time* parameter in all our experiments. Specifically, for the LANL 2000Q we varied annealing time between 1 microsecond and 300 microseconds in step size of 1 microsecond. The hardware does not allow for annealing times less than 1 microsecond.

For the backends Advantage-system1.1 and 2000Q-6, we used AWS Braket to execute the tasks, and therefore getting large numbers of results was more limited. As such, we tried an annealing time of 2 microseconds, and an annealing time of 100 microseconds on both of these backends.

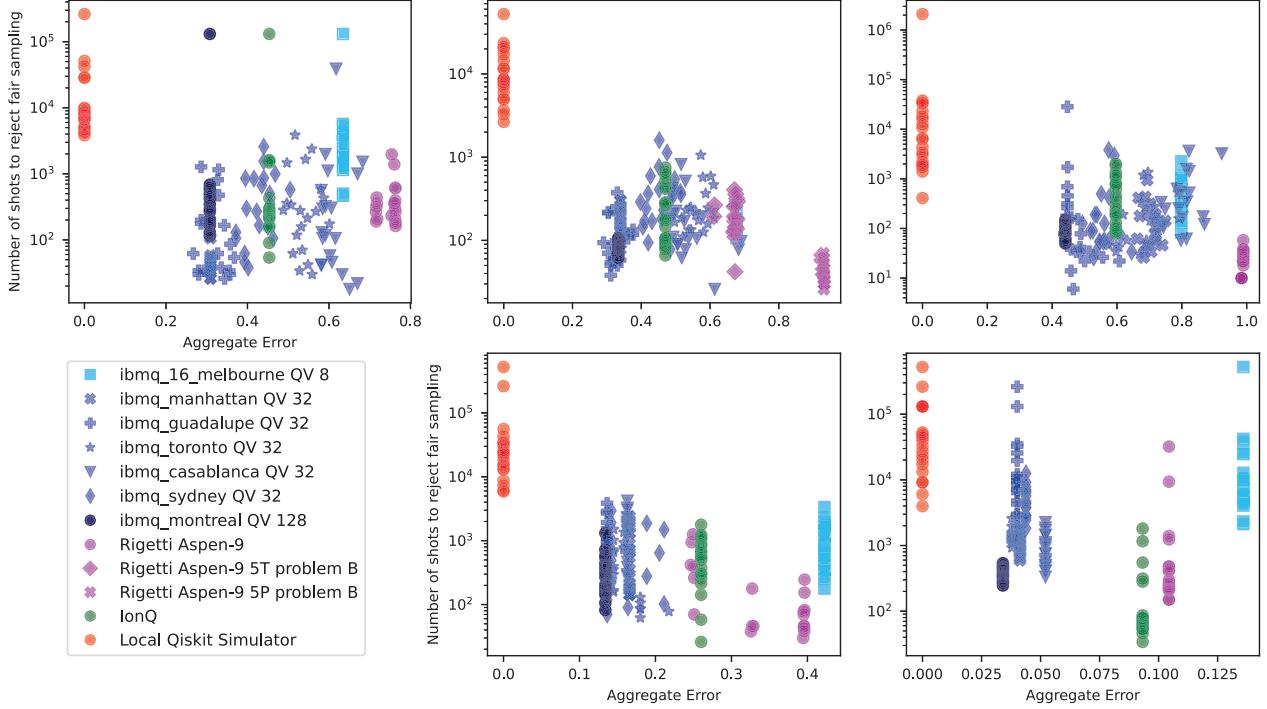


Fig. 5: Aggregate error vs Fairness for Problems (a) (top left), (b) (top center), (c) (top right), (d) (bottom center), (e) (bottom right).

B. Ground state probability metric

Our first metric is simply *ground state probability* (abbreviated as GSP), which is the fraction of shots which return a ground state solution for that circuit.

C. Fairness: Number of shots to reject fair sampling metric

We define our main metric, as the *Number of shots to reject fair sampling* (for short simply called Fairness). We want to quantify how biased or fair the distribution of these found ground states is. Importantly, each of the five test problems has at least three ground state solutions, meaning that we can calculate a distribution of how frequently each ground state solution was found. The exact description of this method is outlined in [7]. The goal of this method is to provide a reasonable metric for fairness which translates p-values from the χ^2 test into a larger (and more human-readable) metric. This metric specifically reports the number of shots (drawn from the ground state distribution) one would need in order to reject the fair sampling hypothesis with 95% statistical significance. Thus, the smaller the metric is, the less fair the ground state distribution is, and the larger it is the more fair the ground state distribution is.

The only adjustable parameter for this method, other than the input distribution, is the number of inner loops (called n_i). For all gate model results, we use $n_i = 100,000$ when applying the method, and for all D-Wave quantum annealing

results we use $n_i = 1,000$ in order to reduce computation time.

D. Aggregate Error metric

In order to quantify how much noise a particular circuit was subject to during its execution, we used the *aggregate error* metric introduced in [7]. To be specific, for a sequence of instructions (gates) g_1, \dots, g_n that make up a circuit C , the *aggregate error* is defined as $E_C = 1 - \prod_{i=1}^n (1 - e_i)$ where e_i is the associated error for the gate g_i . In this computation, we include every gate including the readout operation.

In the case of IBM Q, every task includes a full description of error rates for each native gate operation for each qubit and edge on the hardware. The calibration data for each backend is updated consistently (on the order of hours to days). Before submitting the task to each IBM Q backend, we compiled the circuit into the native gateset for all IBM Q devices using the QISKit transpiler, see more details in Section IV-A. Therefore, we used these device compiled circuits in order to compute the *aggregate error* for each circuit ran on the IBM Q devices.

In the case of Rigetti, which we accessed through AWS Braket, we queried the latest calibration data at the same time we were executing each task. This calibration data seems to be updated regularly (on the order of hours or day). The Rigetti calibration information is similar to the IBM Q data, except that they provided single qubit gate fidelities as an entire class

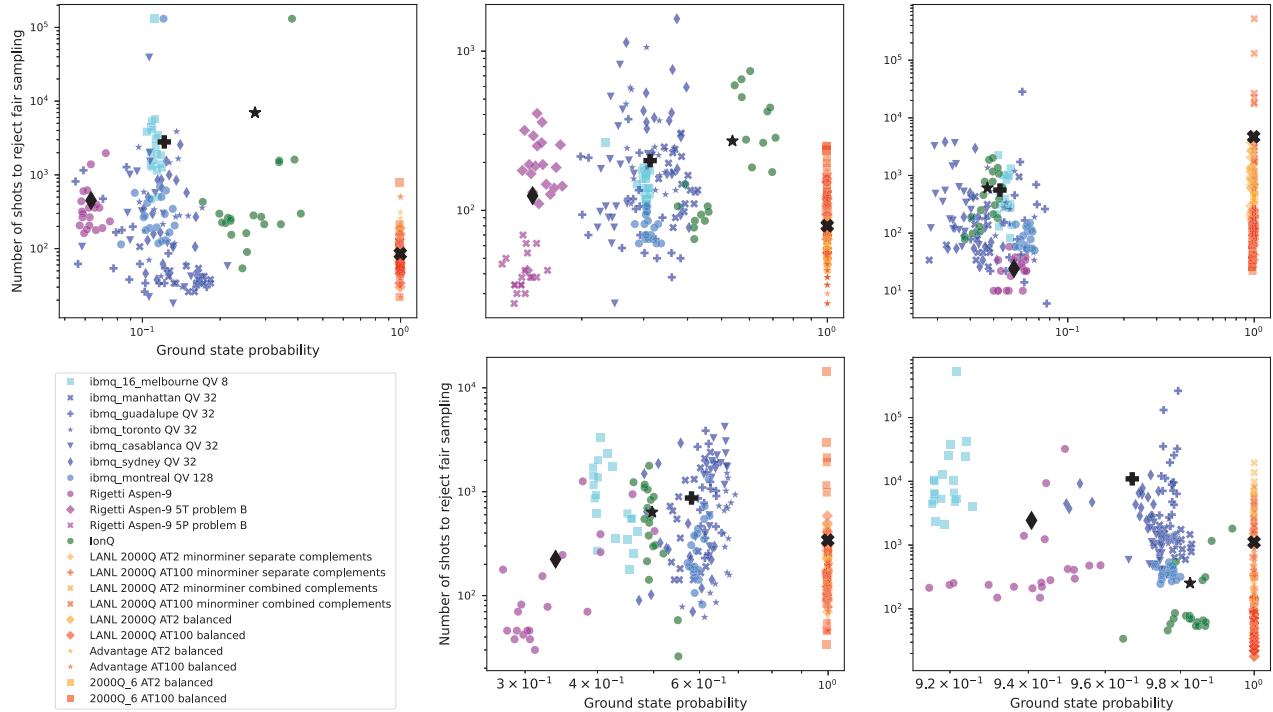


Fig. 6: Log Ground state probability vs log Fairness for problems (a) (top left), (b) (top center), (c) (top right), (d) (bottom center), (e) (bottom right) across all 12 NISQ backends. Black markers denote centroids of the point clouds for each of the four classes of NISQ device (IBM Q is \blacktriangleleft , Rigetti is \blacklozenge , IonQ is \star , and D-Wave is \blacktimes)

(instead of differentiating between different gate’s) and the data was given as gate fidelities instead of error rates. For each task submitted, we also get the device compiled Quil code for that circuit (see Section IV-A). We then used that Quil circuit (and the queried calibration data) in order to compute aggregate error.

IonQ provides less detailed information, but does give mean two qubit gate fidelity, a mean single qubit fidelity, and a mean readout fidelity, which we used to calculate aggregate error for the IonQ circuits.

For the D-Wave quantum annealers we test, there is not a reported error metric. Therefore, we only use the aggregate error metric for the gate model devices.

V. RESULTS

A. Fair sampling comparison across IBM Q, Rigetti, and IonQ

Because error is the main reason for falling short of theoretical performance, we first study how the two metrics GSP and Fairness change as a function of *aggregate error* for the different gate model backends. Figure 4 shows how GSP changes for the five different test problems on each of the nine gate model backends. Similarly, Figure 5 shows how Fairness changes across the nine gate model backends. We include points obtained from a classical Local QISKit Simulator, which shows roughly what to expect from a zero-error quantum device. Across all figures in this subsection, we plot

the metric values for each of the 20 individual calls (each call having 8192 shots) for each backend. Thus each plot in Figures 4, 5 and 6 shows at least 200 individual points plotted. For the IBM Q devices, the results plotted in Figures 4, 5 and 6 all used the noise adaptive transpiler options outlined in Section IV-A.

In theory, increasing aggregate error should correlate with decreasing values for both the GSP and Fairness. While both Figures 4 and 5 show this general trend, there are fairly notable exceptions. The correct interpretation of these plots looks at the 20 points of each backend as a cloud of points; while we did not discard outlier points, we should take care to not interpret results solely based on outliers, but rather on the position of the cloud. With respect to GSP in Figure 5, IonQ reaches the best values in Problems (a), (b), and (e), despite having higher aggregate error than many of the IBM Q systems. Among the IBM Q system, there is substantially smaller aggregate error as we go from the older Quantum Volume (QV) 8 systems to the newer higher QV systems (light blue to dark blue). However, this reduced error does not consistently lead to improved GSP values, in fact several of the QV32 systems outperform the QV128 in most examples. IBM Q does however perform best for Problems (c) and (d), coming very close to the theoretical optimum (QISKit simulator) in Problems (d) and (e). Rigetti’s performance falls behind, largely due to high aggregate error. With respect to

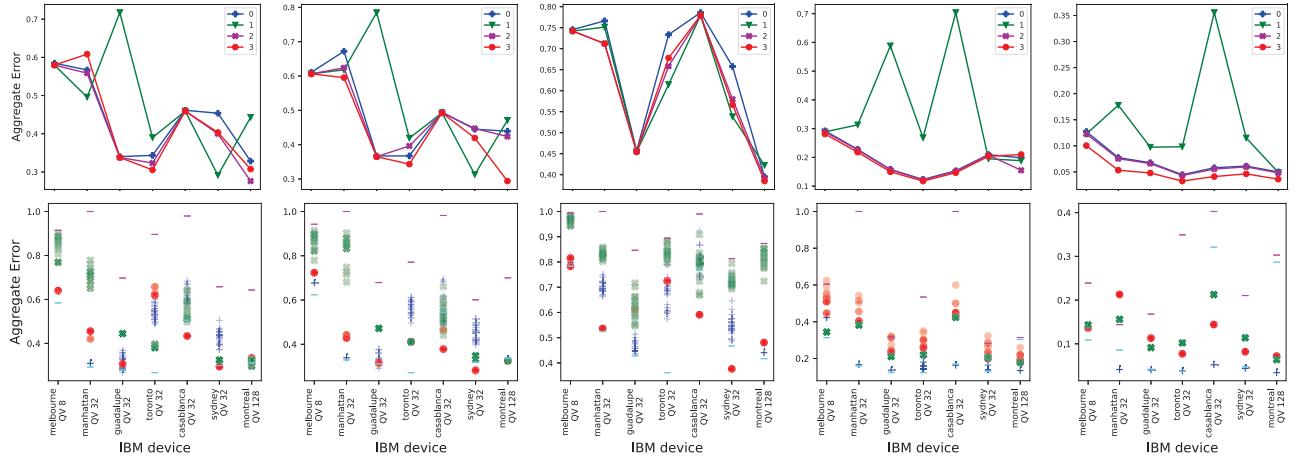


Fig. 7: IBM Q device-specific aggregate errors for Problems (a), (b), (c), (d), (e) (from left to right) depending on transpilation: **(top)** We compare varying optimization levels 0, 1, 2, 3 for transpilation with noise adaptive layout for the respective devices. **(bottom)** We plot 3 different transpilation methods: Default (\times), Initial layout (\circ) and Noise adaptive ($+$), as well as minimum ($-$) and maximum ($-$) aggregate error over *all possible* IBM transpiler flag combinations. Default is calling `transpile` with no additional flags. Initial layout 3 specifies a randomly chosen layout isomorphic to the original circuit (with level 3 optimization). Noise adaptive 3 uses a noise adaptive layout (with level 3 optimization).

fairness (Figure 5), we note the large variability in the Fairness metric that even extends to the simulator data. The high-error low QV IBM Q Melbourne system performs surprisingly well, most likely due to its unique topology, beating out its higher QV siblings, Rigetti and IonQ in Problems (a) and (e) and matching their performance in the other problems. The Rigetti system manages a similar performance coming close to IBM and IonQ in most problems despite high error. The QV128 IBM system again does not really stand out despite its generally low error rates. IonQ's fairness solidly matches IBM Q's in four problems, but falls well short in Problem (e).

Figure 6 directly shows a scatter plot of our two main metrics GSP vs. Fairness for all 12 NISQ backends using a logarithmic scale. There are several ways in which the D-Wave experiments were performed (see Section IV-A), resulting in several different results being plotted for this particular comparison. Note that for the gate model devices, we used 1 round QAOA. This means that we can not achieve the GSP on the gate model circuits (see I) that we can on the quantum annealers. Thus, in Figure 6 the GSP for all D-Wave backends is larger than the gate model devices.

Ignoring individual outlier points, Figure 6 lets us visualize the Pareto front in this bi-criteria optimization, where a point lies on the Pareto front if it does not get outperformed in both metrics by other points. For both Problems (a) and (b), IonQ has the highest Fairness metric, followed by IBM Q devices, then Rigetti, and lastly DWave. For Problems (a) and (b) on gate model devices, GSP is highest for IonQ, followed by IBM Q devices, and then Rigetti. This result is to be expected because problems (a) and (b) are the densest of the five Problems tested, and they take full advantage of the all to

all connectivity of IonQ, whereas the circuit implementations of (a) and (b) on IBM Q and Rigetti devices used an ancilla. Problem (c) shows that DWave ranks the highest for Fairness, followed by IonQ and IBM Q being nearly equal, and lastly Rigetti. For Problem (c) on gate model devices Rigetti has the highest GSP followed by IBM and IonQ, which again are nearly equal. Problems (d) Fairness goes from IBM Q devices, to IonQ, to DWave, and lastly Rigetti. For gate model devices on Problem (d) GSP ranking goes from IBM Q to IonQ to Rigetti. Problem (e) Fairness ranking goes from IBM Q devices, to Rigetti, to DWave, and lastly to IonQ. Problem (e) on gate model devices results in GSP going from IonQ, to IBM Q, and lastly Rigetti.

B. Compilation options for IBM Q backends

The particular QISKit transpilation method used in order to compile each of the five circuits onto the IBM Q backends greatly affects the results. In this Section, we present an analysis of how different transpilation options change the *aggregate error*.

Figure 7 (bottom row) shows the *aggregate error* across all IBM Q backends for each of the five test problems. In particular, we plot the *aggregate error* for the three transpiler options outlined in Section IV-A. Additionally, we perform a grid search across the transpiler options *layout-method*, *routing-method*, *translation-method*, and *optimization-level*, and we plot the minimum and maximum aggregate error found in the grid search.

This figure shows that there is no clear winning transpiler parameter combination across all backends and circuits, however the resulting aggregate varies by as much as 60 %. In practice the figures suggest that taking the better performing

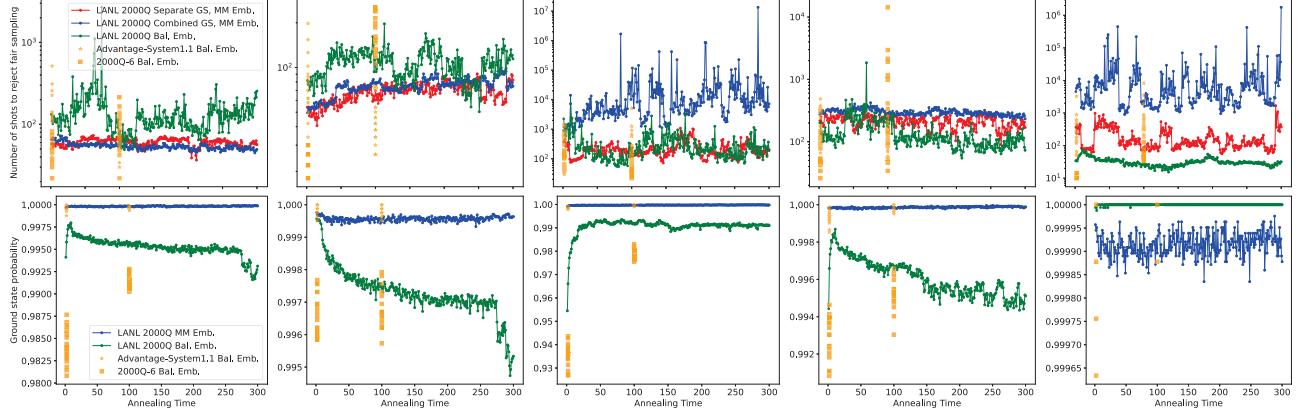


Fig. 8: Annealing-time dependent results on three D-Wave devices for Problems (a), (b), (c), (d), (e) (from left to right) for **(top)** log Fairness and **(bottom)** Ground state probability.

of the noise adaptive compilation and giving initial layout compilation should result in near minimum aggregate error in most cases.

Figure 7 (top row) plots *aggregate error* as a function of IBM Q backend for the 4 different possible QISKit transpiler optimization levels. Besides the optimization flag, the *layout-method* option was set to *noise adaptive*. The figure shows that optimization level 3 performs best with the exception of the Sydney backend. Variance in the performance of different backends (such as Sydney) is likely due a mismatch between the latest calibration data for that system and the noise experienced by the circuit when executed on the backend.

Generally, increasing optimization levels decrease aggregate error as expected, with the exception of optimization level 1. Optimization level 1 seems to work better on average in the case of a *noise adaptive layout method*; the reason being that optimization level 1 reduces overall circuit depth (thus decreasing *aggregate error*), but does not introduce a large number of CNOT gates as a result of SWAP operations attempting to adapt to the noise levels of different qubits (as higher optimization levels do).

C. D-Wave

We investigate how our metrics change as a function of *annealing time* for three D-Wave quantum annealing backends. For the LANL 2000Q backend, we plot the mean metric out of the 20 device calls (each call with 8192 shots) as a function of *annealing time*. For the other two backends (Advantage-System1.1 and 2000Q-6), we plot all 20 points for the two annealing times they were run on (2 microseconds and 100 microseconds).

In Figure 8 (top row), we observe a high variance in the fairness metric across all problems. Problems (a), (b) consistently show that the LANL 2000Q balanced embedding has the highest fairness over annealing time. Problems (c), (d), (e) on the other hand show that the LANL 2000Q minorminer embedding has the highest fairness over annealing time.

In Figure 8 (bottom row), we observe different *annealing time* dependent trends across each of the problems, largely from the LANL 2000Q balanced embedding data (the LANL 2000Q minorminer embedding results are relatively constant over annealing time). For Problems (a), (c), and (d), we observe a sharp increase in GSP in the first few microseconds of annealing time. Problems (a) and (d) then decrease in GSP after that initial increase. Problem (c) continues to increase GSP up until it plateaus at roughly 100 microseconds. Problem (b) shows no initial increase in GSP, instead it starts out at a high GSP and decreases over all annealing times. Problem (e) shows effectively no difference in GSP over time. It is hard to tell whether these D-Wave results are in agreement with the theoretical predictions of [13], which call for initial fairness with low annealing times and then unfairness but higher GSP as we increase annealing times, because all modern D-Wave backends have a minimum annealing time of 1 microsecond.

VI. CONCLUSION

In this article, we compare IBM Q’s, Rigetti’s, IonQ’s and D-Wave NISQ devices with respect to their performance in finding ground states and fairness across a set of five paradigmatic problems. While no platform turned out to be dominant, we noticed performance differences in both metrics. It is encouraging to see that leading NISQ vendors are relatively close in performance and we hope that further technological advances will increase the fairness of sampling as well as bring advances to various compilation techniques.

As for open problems, we plan to compare D-Wave performance against QAOA-based circuits that use the transverse field mixer (as opposed to the Grover mixer), we would like to relate our D-Wave results more cleanly to theoretical predictions. We also plan to continue our studies for novel NISQ backends.

REFERENCES

[1] V Akshay, H Philathong, Mauro ES Morales, and Jacob D Biamonte. Reachability deficits in quantum approximate optimization. *Physical*

Review Letters, 124(9):090504, 2020. [arXiv:1906.11259](https://arxiv.org/abs/1906.11259), doi: [10.1103/PhysRevLett.124.090504](https://doi.org/10.1103/PhysRevLett.124.090504).

[2] Marlon Azinović, Daniel Herr, Bettina Heim, Ethan Brown, and Matthias Troyer. Assessment of quantum annealing for the construction of satisfiability filters. *SciPost Physics*, 2(2):013, 2017. [arXiv:1607.03329](https://arxiv.org/abs/1607.03329), doi: [10.21468/SciPostPhys.2.2.013](https://doi.org/10.21468/SciPostPhys.2.2.013).

[3] Andreas Bärtschi and Stephan Eidenbenz. Grover Mixers for QAOA: Shifting Complexity from Mixer Design to State Preparation. In *IEEE International Conference on Quantum Computing & Engineering QCE'20*, pages 72–82, 2020. [arXiv:2006.00354](https://arxiv.org/abs/2006.00354), doi: [10.1109/QCE49297.2020.00020](https://doi.org/10.1109/QCE49297.2020.00020).

[4] Jun Cai, William G. Macready, and Aidan Roy. A practical heuristic for finding graph minors. *arXiv e-prints*, 2014. [arXiv:1406.2741](https://arxiv.org/abs/1406.2741).

[5] D-Wave Systems. minorminer. <https://github.com/dwavesystems/minorminer>, 2017. A heuristic tool for minor embedding.

[6] D-Wave Systems. *D-Wave Ocean: minorminer*, 2021.

[7] John Golden, Andreas Bärtschi, Daniel O’Malley, and Stephan Eidenbenz. QAOA-based Fair Sampling on NISQ Devices. *arXiv e-prints*, 2021. [arXiv:2101.03258](https://arxiv.org/abs/2101.03258).

[8] Lov K. Grover. Fixed-Point Quantum Search. *Physical Review Letters*, 95(15):150501, 2005. doi: [10.1103/PhysRevLett.95.150501](https://doi.org/10.1103/PhysRevLett.95.150501).

[9] Stuart Hadfield, Zhihui Wang, Bryan O’Gorman, Eleanor G Rieffel, Davida Venturelli, and Rupak Biswas. From the quantum approximate optimization algorithm to a quantum alternating operator ansatz. *Algorithms*, 12(2):34, 2019. [arXiv:1709.03489](https://arxiv.org/abs/1709.03489), doi: [10.3390/algorithms12020034](https://doi.org/10.3390/algorithms12020034).

[10] Dylan R Harp, Zhenxue Dai, Andrew V Wolfsberg, Jasper A Vrugt, Bruce A Robinson, and Velimir V Vesselinov. Aquifer structure identification using stochastic inversion. *Geophysical Research Letters*, 35(8):L08404, 2008. doi: [10.1029/2008GL033585](https://doi.org/10.1029/2008GL033585).

[11] Tadashi Kadokami and Hidetoshi Nishimori. Quantum annealing in the transverse ising model. *Physical Review E*, 58(5):5355, 1998. [arXiv:cond-mat/9804280](https://arxiv.org/abs/cond-mat/9804280), doi: [10.1103/PhysRevE.58.5355](https://doi.org/10.1103/PhysRevE.58.5355).

[12] Vaibhaw Kumar, Casey Tomlin, Curt Nehrkorn, Daniel O’Malley, and Joseph Dulny III. Achieving fair sampling in quantum annealing. *arXiv e-prints*, 2020. [arXiv:2007.08487](https://arxiv.org/abs/2007.08487).

[13] Mario S. Könz, Guglielmo Mazzola, Andrew J. Ochoa, Helmut G. Katzgraber, and Matthias Troyer. Uncertain fate of fair sampling in quantum annealing. *Physical Review A*, 100(3):030303, 2019. [arXiv:1806.06081](https://arxiv.org/abs/1806.06081), doi: [10.1103/PhysRevA.100.030303](https://doi.org/10.1103/PhysRevA.100.030303).

[14] Salvatore Mandra, Zheng Zhu, and Helmut G Katzgraber. Exponentially Biased Ground-State Sampling of Quantum Annealing Machines with Transverse-Field Driving Hamiltonians. *Physical Review Letters*, 118(7):070502, 2017. [arXiv:1606.07146](https://arxiv.org/abs/1606.07146), doi: [10.1103/PhysRevLett.118.070502](https://doi.org/10.1103/PhysRevLett.118.070502).

[15] Yoshiki Matsuda, Hidetoshi Nishimori, and Helmut G Katzgraber. Ground-state statistics from annealing algorithms: quantum versus classical approaches. *New Journal of Physics*, 11(7):073021, 2009. [arXiv:0808.0365](https://arxiv.org/abs/0808.0365), doi: [10.1088/1367-2630/11/7/073021](https://doi.org/10.1088/1367-2630/11/7/073021).

[16] Daniel O’Malley. An approach to quantum-computational hydrologic inverse analysis. *Scientific reports*, 8(1):1–9, 2018. doi: [10.1038/s41598-018-25206-0](https://doi.org/10.1038/s41598-018-25206-0).

[17] Lukas M Sieberer and Wolfgang Lechner. Programmable superpositions of Ising configurations. *Physical Review A*, 97(5):052329, 2018. [arXiv:1708.02533](https://arxiv.org/abs/1708.02533), doi: [10.1103/PhysRevA.97.052329](https://doi.org/10.1103/PhysRevA.97.052329).

[18] Masayuki Yamamoto, Masayuki Ohzeki, and Kazuyuki Tanaka. Fair Sampling by Simulated Annealing on Quantum Annealer. *Journal of the Physical Society of Japan*, 89(2):025002, 2020. [arXiv:1912.10701](https://arxiv.org/abs/1912.10701), doi: [10.7566/JPSJ.89.025002](https://doi.org/10.7566/JPSJ.89.025002).

[19] Theodore J Yoder, Guang Hao Low, and Isaac L Chuang. Fixed-Point Quantum Search with an Optimal Number of Queries. *Physical Review Letters*, 113(21):210501, 2014. [arXiv:1409.3305](https://arxiv.org/abs/1409.3305), doi: [10.1103/PhysRevLett.113.210501](https://doi.org/10.1103/PhysRevLett.113.210501).