Application-Motivated, Holistic Benchmarking of a Full Quantum Computing Stack

Daniel Mills^{1,2}, Seyon Sivarajah², Travis L. Scholten³, and Ross Duncan^{2,4}

¹University of Edinburgh, 10 Crichton Street, Edinburgh EH8 9AB, UK

²Cambridge Quantum Computing Ltd, 9a Bridge Street, Cambridge, CB2 1UB, UK

³IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA

⁴University of Strathclyde, 26 Richmond Street, Glasgow, G1 1XH, UK

Quantum computing systems need to be benchmarked in terms of practical tasks they would be expected to do. Here, we propose 3 “application- motivated” circuit classes for benchmarking: deep (relevant for state preparation in the variational quantum eigensolver algorithm), shallow (inspired by IQP-type circuits that might be useful for near-term quantum machine learning), and square (inspired by the quantum volume benchmark). We quantify the performance of a quantum computing system in running circuits from these classes using several figures of merit, all of which require exponential classical computing resources and a polynomial number of classical samples (bitstrings) from the system. We study how performance varies with the compilation strategy used and the device on which the circuit is run. Using systems made available by IBM Quantum, we examine their performance, showing that noise-aware compilation strategies may be beneficial, and that device connectivity and noise levels play a crucial role in the performance of the system according to our benchmarks.

Contents

1	Introduction	2
2	Figures of Merit	3
2.1	Heavy Output Probability	4
2.2	Cross-Entropy Difference	5
2.3	ℓ_1 -Norm Distance	6
2.4	Metric Comparison	6
3	Circuit Classes	7
3.1	Shallow Circuits: IQP	7
3.2	Square Circuits: Random Circuit Sampling	8
3.3	Deep Circuits: Pauli Gadgets	10
4	Quantum Computing Stack	11
4.1	Software Development Kits	11
4.2	Compilers	11
4.3	Devices	12
5	Experimental Results	12
5.1	Full Stack Benchmarking	13
5.2	Application Motivated Benchmarks	14
5.3	Insights from Classical Simulation	18
6	Conclusion	19
	References	21
A	Exponential Distribution	25
A.1	Square Circuits	26
A.2	Deep Circuits	26
A.3	Shallow Circuits	26
B	Compilation Strategies	27
C	Device Data	29
C.1	Device Coupling Maps	29
C.2	Device Calibration Information	29
D	Empirical Relationship Between Heavy Output Generation Probability and ℓ_1-Norm Distance	30

Daniel Mills: daniel.mills@cambridgequantum.com

1 Introduction

The evolution of quantum computers from experimental devices comprising a handful of qubits, towards general-purpose, programmable, commercial-grade systems [1–5] necessitates new techniques for characterizing them. Quantum characterization, validation, and verification (QCVV) protocols to detect, diagnose, and quantify errors in quantum computers originally focused on properties of one or several qubits (e.g., T_1 and T_2 times, gate error rates, state preparation fidelity, etc). Such quantities measure the form and magnitude of noise present in the system, but are only *proxy* measures of device performance when implementing computations.

This can be addressed through “holistic” benchmarks, which stress test a quantum computing system in its *entirety*, and not just individual components. “Holistic benchmarking” of a quantum computing system has commonly referred to benchmarking the physical implementation of a collection of qubits, without referring to the computational task they would perform. This is most useful when testing *physical properties* of a collection of qubits.¹

A complementary view (taken in this work) is that holistic benchmarks test the *computational capabilities* of the complete system. Under this view, the entire compute *stack* – qubits, compilation strategy, classical control hardware, etc. – should be benchmarked collectively [6]. This captures the performance of the system as an integrated unit, giving a better indication of the system’s performance in practice.

Holistic benchmarks, while useful for comparing systems,² and tracking the performance of a given system over time, provide little information on how different combinations of the stack’s components could change system performance. For this reason, holistic benchmarking should, as much as possible, make explicit the variable components of the stack, and systematically vary them to identify how a particular component affects system-level performance. This ensures these benchmarks identify improvements due to changes at all layers of the quantum computing stack, towards which, for example, improvements in the compilation strategy can contribute significantly.

The benchmarks defined here have two parts: a *circuit class* and a *figure of merit*. The circuit class describes the type of circuit to be run by the system, and the figure of merit quantifies performance of the system when running circuits from that class.

Quantum computing systems are used for particu-

lar applications, so the circuit classes should test the performance of a system in those arenas [10]. Assessing the performance of quantum computing stacks when implementing circuits derived from applications ensures that the results are stronger indicators of whether those applications can be successfully implemented than they would otherwise be (for example, by numerical simulation of the application using a noise model).

At least two notions have been put forth for defining such classes. The first is benchmarks based on often-used quantum algorithmic primitives [11], the circuit classes being primitives of Grover iterations and Trotterized Hamiltonian simulation. This approach has the advantage that such computational primitives will be ubiquitous in quantum applications, but has the disadvantage that implementations on near-term quantum computers may yield poor performance.

The second is for the circuit class to be a particular instance of an application. Such benchmarks have been defined in the context of quantum simulation [12–16], quantum machine learning [10, 17–20], discrete optimisation [21–24], and quantum computational supremacy [3, 25–27]. This approach has the advantage that the definition of success when running a circuit from the class is fairly straightforward, but with the disadvantage that performance, as measured by one instance of an application, may not be predictive of performance for the application generically, or another application. Further, while the wide selection of circuits presented in the aforementioned literature covers an array of applications, deriving benchmarks independently of each other may result in a lack of coverage, or unnecessary repeated coverage, of circuit classes.

The “application-motivated” circuit classes defined in this work draw inspiration from [11] (looking at computational primitives) but also from the literature above (focusing on computational primitives of near-term quantum computing applications). This approach has the advantage that a system which does well on an application-motivated benchmark should do well in running the application the benchmark was derived from. Further, by considering application primitives, the result of the benchmarks apply to the application generally, and not only to a particular instance of the application.

We introduce 3 such “application-motivated” circuit classes, which cover varying depth regimes and are (somewhat) controllable in depth. The definitions and motivations for each class are given in Section 3 but, in brief, the classes – as labelled by their depth regimes – and the applications that motivate them, are:

Shallow: Inspired by hardware-efficient ansatzes [28, 29] which may be useful for near-term quantum machine learning and chemistry applications [30–32]. The width of shallow circuits grow with a minimal increase in depth, allowing the impact

¹A simple example is crosstalk detection, where the output of the benchmarking could be a table of coupling values between all connected qubits.

²Similar to the benchmarking of classical computers, with the LINPACK benchmarks [7, 8] being used to build the TOP500 ranking of supercomputers [9].

of including many qubits to be explored.

Square: Inspired by the circuits used to calculate a system’s quantum volume [33]. These circuits utilise gates sampled uniformly at random from all $SU(4)$ gates, making them a test of general-purpose, programmable quantum computers.

Deep: Inspired by product formula circuits, including state preparation circuits used in the variational quantum eigensolver (VQE) algorithm[34–36]. The depth of these circuits grows most quickly with width, giving a thorough coverage of depth regimes.

Importantly, as discussed in Section 3, these circuits are complementary; they cover a wide selection of applications and circuit types. By covering many applications with few circuit classes, we are able to concisely present far-reaching results, and minimise the number of computations to be performed in deriving those results.

How well a stack executes a circuit is assessed here via continuous figures of merit, rather than binary ones which may only verify correctness. This is because the outcomes from noisy devices will likely not be correct, while information about closeness to the correct answer is still highly valuable. Further, techniques for the verification of universal quantum computation requires many qubits, or qubit communication, or both, none of which are accessible using currently-available devices [37, 38].

We use three figures of merit, calculated using classical computers: heavy output generation probability [39], cross-entropy difference [27], and ℓ_1 -norm distance. Estimating each of these figures of merit requires knowledge about the *ideal* (noise-free) outcome (bitstring) probabilities.

In practice, calculating them requires direct simulation of the circuit under consideration. Consequently, scaling to tens or hundreds of qubits will be challenging in general, particularly if the ℓ_1 -norm distance is used as the figure of merit. However, by considering circuits with a few qubits, the circuits can be simulated classically, and insight into the behaviour of larger devices [3, 40] can be gained.

However, significant improvements in the time needed to perform benchmarks can be made if the circuits and figures of merit are developed jointly. Indeed, the calculation of heavy output generation and cross-entropy difference requires only a polynomial number of samples from a quantum device for some circuit classes. Importantly, we show that this is the case for deep circuits and square circuits. Therefore, deep circuits provide the first instance of chemistry-motivated circuits, which are likely not possible to classically simulate in general, but which can be benchmarked using polynomially many samples from the output distribution from a device implementing them.

We refer to a set of benchmarks as a *benchmarking suite*, each benchmark being defined by unique combinations of each circuit class and figure of merit. Using a benchmarking suite enables the derivation of broad insights about the behaviour and performance of a system across a wide variety of possible applications. Their varying demands on quantum computing resources (e.g., qubits, achievable depth) allows for the exploration of the best routes to extract the most utility from near-term quantum computers. In sum, in order to predict the systems’ performance in practice, our benchmarking approach is both application-motivated, holistic and full quantum computing stack. For the reasons stated above, we regard all of these properties to be necessary of a standard benchmark suite.

Here, we will benchmark systems made available by IBM Quantum using the benchmarking suite defined in this work, and investigate two components of the stack: the *compilation strategy* used to map an abstract circuit onto one that is executable on a quantum computer, and the *device* used to run the compiled circuit and return the results.³

The remainder of this paper is comprised as follows: Section 3 details the circuit classes, including algorithms for generating the circuits; Section 2 explains the figures of merit we use; Section 4 introduces the software stack, as well as hardware made available by IBM Quantum, that comprise the systems we’ll be benchmarking; and Section 5 shows the results of our benchmarking. We conclude in Section 6.

2 Figures of Merit

Suppose a quantum computer runs an n qubit circuit C , on the input $|0\rangle^n$. Repeated runs produce a set of classical bitstrings $x_1, \dots, x_k$ (with k being the total number of runs). Figures of merit compare $p_C(x_j) = |\langle x | C | 0^n \rangle|^2$, the ideal output probabilities of each x_j in C , and $\mathcal{D}_C(x_j)$, the probability that x_i is produced by a real implementation, which may be noisy.

In practice, direct access to $\mathcal{D}_C(x_j)$ is impossible, because the number of bitstrings generally grows exponentially with circuit width. Consequently, for a fixed k , there will be some bitstrings which are not observed. The k bitstrings which are observed can be post-processed to *estimate* a given figure of merit. Statistical uncertainty in the value of figures of merit can arise in this way. On the other hand, simulations can be used to calculate $p_C(x_j)$ directly, and the figures of merit we discuss will make use of this.

As such, focusing on figures of merit which can be reasonably approximated from samples is important. This is to say with a number of samples from $\mathcal{D}_C$

³While the particular systems used here have other components (such as pulse synthesisers), we do not look at the impact of those pieces on full-stack performance.

which is polynomial in the size of the circuit, as opposed to the exponential number which would be required for a full characterisation. This focus allows for a better scaling in the resource requirements as the size of the circuit grows.

The circuits we investigate are small enough that $\mathcal{D}_C$ can be well characterised by a reasonable absolute number of samples. As such we will also include a discussion and comparison with the ℓ_1 -norm distance, motivated by its relation to several theoretical results, notably Theorem 1.⁴

The remainder of this section outlines three figures of merit: Heavy output probability, Cross-entropy difference, and ℓ_1 -norm distance. We detail: their definition, the continuous range of values they can take, their dependence on noise, and the procedure for calculating their value from samples produced by an implementation. These figures of merit have been developed in other works, and so in this section we explore these properties in so far as they facilitate their use and justify our choosing them.

2.1 Heavy Output Probability

Heavy Output Generation [39] (HOG) is the problem of producing bitstrings that are predominantly those that are the most likely in the output distribution of C . That is to say, outputs with the highest probability in the ideal distribution should be produced most regularly.

If the ideal distribution is sufficiently far from uniform, HOG provides a means to distinguish between samples from the ideal distribution and a trivial attempt to mimic such a sampling procedure; namely by producing uniformly random bitstrings. Although a simply stated problem, which is verifiable by a classical device using a polynomial number of samples from $\mathcal{D}_C$, this task is conjectured to be hard for a classical computer to perform in general [39].

Definitions and Related Results An output $z \in \{0,1\}^n$ is *heavy* for a quantum circuit C , if $p_C(z)$ is greater than the median of the set $\{p_C(x) : x \in \{0,1\}^n\}$. We can define the probability that samples drawn from a distribution $\mathcal{D}_C$ will be heavy outputs in the distribution p_C , called the *heavy output generation probability* of $\mathcal{D}_C$, as follows:

$$\text{HOG}(\mathcal{D}_C, p_C) = \sum_{x \in \{0,1\}^n} \mathcal{D}_C(x) \delta_C(x). \quad (1)$$

Here $\delta_C(x) = 1$ if x is heavy for C , and 0 otherwise.

⁴In the case that sufficient samples can be obtained to characterise $\mathcal{D}_C$ well, several other figures of merit are accessible, such as the KL-divergence. We limit our investigations to the three discussed to remain thorough but concise, but invite further analysis of the data generated by our experiments. This data can be found under ‘Data Availability’ towards the end of the paper.

In the case $\text{HOG}(p_C, p_C) \approx 1/2$, outputs which are heavy have similar probabilities of occurring as outputs that are not. Such distributions are well approximated by the uniform distribution. Hence, for $\text{HOG}(\mathcal{D}_C, p_C)$ to help us distinguish between an ideal implementation of C and a trivial attempt to mimic it by generating random bitstrings, $\text{HOG}(p_C, p_C)$ should be greater than $1/2$. In fact, $\text{HOG}(p_C, p_C)$ is expected to be $(1 + \log 2)/2 \approx 0.846574$ [39] for circuit classes whose distribution of measurement probabilities, p , is of the exponential form $\Pr(p) = N e^{-Np}$, where $N = 2^n$.⁵ This is discussed at length in Appendix A where we will demonstrate that the deep circuits and square circuits classes have this property. When the output distributions of a class of circuits is shown to take this form it is meaningful to define the Heavy Output Generation problem.

Problem 1 (Heavy Output Generation [39]) *Given a measure μ over a class of circuits, the family of distributions $\{\mathcal{D}_C\}$ is said to satisfy HOG if the following is true.*

$$\mathbb{E}_{C \leftarrow \mu} [\text{HOG}(\mathcal{D}_C, p_C)] \geq \frac{2}{3} \quad (2)$$

The classical harness of HOG can be related to the difficulty of calculating the output probabilities of quantum circuits [39]. In particular, such calculations have been argued to be beyond classical computers using polynomial resources in the case of circuits similar to the square circuits of Section 3.2 [39].

For the circuit classes in Section 3 which have exponentially distributed output probabilities, we will often consider the largest n for which distributions $\{\mathcal{D}_{C_n}\}$ solve Problem 1. This approach is inspired by quantum volume and is useful as an indicator of the largest Hilbert space accessible to a quantum computing stack [33]. For circuit classes with output probabilities that are not distributed as such, we will explicitly calculate the ideal heavy output probability as a point of comparison.

Empirical Estimation From Samples We approximate $\text{HOG}(\mathcal{D}_C, p_C)$ in a number of operations which grows exponentially with the number of qubits, but using only a polynomial number of samples from the real distribution $\mathcal{D}_C$. To do so we evaluate the following expression, where the ideal probabilities $p_C(x)$ must be calculate in order to determine $\delta_C(x)$, and where $x_1, \dots, x_k$ are samples drawn from $\mathcal{D}_C$.

$$\frac{1}{k} \sum_{i=1, \dots, k} \delta_C(x_i) \quad (3)$$

By the law of large numbers, this converges to $\text{HOG}(\mathcal{D}_C, p_C)$ in the limit of increasing sample size.

⁵This is also commonly referred to as the Porter-Thomas distribution [41].

Evaluating Stack Performance In the case of extreme noise, and the convergence of the real distribution $\mathcal{D}_C$ to the uniform distribution $\mathcal{U}$, $\text{HOG}(\mathcal{D}_C, p_C) = 1/2$. This is compared to the case where $\mathcal{D}_C = p_C$, when we would expect to have $\text{HOG}(\mathcal{D}_C, p_C) = (1 + \log 2)/2$.

Accordingly, in the results of Section 5, we say a quantum computing stack has performed well if, on average over the class of circuits, $\text{HOG}(\mathcal{D}_C, p_C)$ is between $(1 + \log 2)/2$ and $2/3$, with $(1 + \log 2)/2$ being best of all. A poorer performance is indicated by average values between $2/3$ and $1/2$, with $1/2$ being worst of all. This gives a continuum of values, with an intuitive interpretation, with which to compare quantum computing stacks, which we call Heavy Output Generation Benchmarking.

2.2 Cross-Entropy Difference

Cross-entropy benchmarking [27] relates to the average probability, in the ideal distribution, p_C , of the outputs which are sampled from the empirical distribution, $\mathcal{D}_C$. For distributions which are far from uniform, this measure can be used to distinguish an ideal from a noisy implementation. Ideal implementations will regularly produce the higher probability outputs, obtaining a high benchmark value, while even a small shift in the distribution will lower the value.

The cross-entropy difference can be calculated using exponential classical resources, from a polynomial number of samples from a quantum computer, which allows for its utilisation in benchmarking smaller quantum devices [3, 27, 42, 43]. There are also well-developed means by which this quantity can be used as a means of extrapolating from the behaviour of smaller devices to that of larger devices, which might demonstrate quantum computational supremacy [3].

Definitions and Related Results The *cross-entropy* captures one's surprise when sampling from $\mathcal{D}$ when one is expecting $\mathcal{D}'$:

Definition 1 (Cross-Entropy) The cross-entropy between two probability distributions $\mathcal{D}$ and $\mathcal{D}'$ is

$$\text{CE}(\mathcal{D}, \mathcal{D}') = \sum_{x \in \{0,1\}^n} \mathcal{D}(x) \log \left(\frac{1}{\mathcal{D}'(x)} \right). \quad (4)$$

The *cross-entropy difference* $\text{CED}(\mathcal{D}, \mathcal{D}')$ is $\text{CE}(\mathcal{U}, \mathcal{D}') - \text{CE}(\mathcal{D}, \mathcal{D}')$, where $\mathcal{U}$ is the uniform distribution. Hence, the cross-entropy difference can be seen as answering “Is the distribution $\mathcal{D}'$ best predicted by $\mathcal{D}$ or by the uniform distribution?”.

The comparison to the uniform distribution which the cross-entropy difference provides is valuable as, if an honest attempt is being made to recreate a distribution, at worst $\mathcal{U}$ could be produced.

In some cases, the cross-entropy also gives an estimate for the average circuit fidelity [27], facilitating

the characterisation of noise levels in implementations of quantum circuits. While Cross-Entropy Benchmarking on its own cannot be used to distinguish error channels, in combination with the techniques introduced here, it can provide insight into this information.

Empirical Estimation From Samples By the law of large numbers, the following expression converges to $\text{CE}(\mathcal{D}_C, p_C)$, where $x_1, \dots, x_k$ are samples drawn from $\mathcal{D}_C$.

$$\frac{1}{k} \sum_{i=1, \dots, k} \log \left(\frac{1}{p_C(x_i)} \right) \quad (5)$$

This can be used by a classical computer to approximate the value for $\text{CED}(\mathcal{D}_C, p_C)$. While only a polynomial number of samples x_i are required, the best known algorithms for calculating $p_C(x_i)$ for an arbitrary circuit C require exponential time or space on a classical computer. It should be expected, therefore, that calculating the cross-entropy difference consumes exponential classical compute resources in general.

In our case, to avoid requiring the inverse of 0 in this approximation, we chose to use

$$\max \left\{ p_C(x), 2^{-n^2} \right\}. \quad (6)$$

This choice of an inverse exponential in the number of qubits is inspired by, although not directly derived from, the average case supremacy results related to random circuits [44, 45].

Evaluating Stack Performance $\text{CE}(\mathcal{D}_C, p_C)$, reduces to the entropy, $H(p_C)$, of p_C when there is no noise, as in this case $p_C = \mathcal{D}_C$. Further, when the probabilities $p_C(x)$ are approximately independent and identically distributed according to the exponential distribution, $H(p_C) = \log 2^n + \gamma - 1$ [27], where γ is Euler's constant.

If $\mathcal{D}_C(x)$ is uncorrelated with $p_C(x)$, the following result holds [27]:

$$\mathbb{E}_C [\text{CE}(\mathcal{D}_C, p_C)] = \log 2^n + \gamma \quad (7)$$

$\mathcal{D}_C(x)$ and $p_C(x)$ are uncorrelated if, for example, $\mathcal{D}_C$ is the uniform distribution, or, in the case of demonstrations of quantum computational supremacy, if $\mathcal{D}_C$ is the output of a polynomial cost classical algorithm [27].

These results allow us to identify the extreme values taken by the cross-entropy difference:

$$\mathcal{D}_C = p_C : \quad \text{CED}(\mathcal{D}_C, p_C) = 1 \quad (8)$$

$$\mathcal{D}_C = \mathcal{U} : \quad \text{CED}(\mathcal{D}_C, p_C) = 0 \quad (9)$$

As such, the cross-entropy difference gives a value between 0 and 1 which measures the accuracy of the

implementation of a circuit, the calculation of which is called Cross-Entropy Benchmarking.

Hence, in the results of Section 5 we will say that better performance is indicated by a value, averaged over the class of circuits, closer to 1, and worse by a value closer to 0. Notice that values above 1, or below 0, are possible for individual circuits if, for example, the ideal output distribution happens to be very skewed towards heavy outputs, or if unlikely outcomes occur very often, respectively. Unlike in the case of Heavy Output Generation Benchmarking, we do not give a value which a score above would indicate good performance.⁶

2.3 ℓ_1 -Norm Distance

The ℓ_1 -norm distance between two probability distributions measures the total difference between the probabilities the distributions assign to elements of their sample space. Such a metric is sufficiently strong that for several classes of quantum circuits it is known that classical simulation of all circuits in the class to within some ℓ_1 -norm distance of the ideal distribution would contradict commonly held computational complexity theoretic conjectures [44, 46, 47].

Unlike the previous two figures of merit, approximating the ℓ_1 -norm distance requires a full characterisation of the ideal output distribution. In the cases where few qubits are considered, as is so here, it is possible to perform such characterisations. For larger qubit counts, the cross-entropy benchmarking and heavy output generation are the preferred benchmarking schemes.

Definitions and Related Results In the case of distributions over the sample space $\{0,1\}^n$, the ℓ_1 -norm distance is defined as follows.

Definition 2 (ℓ_1 -norm distance) *For distributions $\mathcal{D}$ and $\mathcal{D}'$ over the sample space $\{0,1\}^n$ the ℓ_1 -norm distance between them is*

$$\ell_1(\mathcal{D}, \mathcal{D}') = \sum_{x \in \{0,1\}^n} |\mathcal{D}(x) - \mathcal{D}'(x)|. \quad (10)$$

Because of its independence from probability values themselves, it is regarded as reasonable to require quantum computers to produce samples from distributions within some ℓ_1 -norm distance of the ideal distribution. This is as opposed to measures of distance, such as multiplicative error, which require zero probability outcomes are preserved in the presence of noise, but for which very strong connections to quantum computational supremacy also exist [48].

⁶Where such a value has been defined, it is perhaps surprisingly close to 0, and so the uniform distribution [3].

Empirical Estimation From Samples In this work we will approximate the ℓ_1 -norm distance between the ideal and real distributions using samples from the real distribution. Given samples $s = \{x_1, \dots, x_m\}$ from $\mathcal{D}_C$, let s^x be the number of times x appears in s . Define $\widetilde{\mathcal{D}}_C$ by $\widetilde{\mathcal{D}}_C(x) = (s^x)/m$. Then the approximation we will use for $\ell_1(\mathcal{D}_C, p_C)$ is $\ell_1(\widetilde{\mathcal{D}}_C, p_C)$.

Evaluating Stack Performance An ideal implementation of a circuit would result in $\ell_1(\mathcal{D}_C, p_C) = 0$. As such, in the results of Section 5 we will say that a quantum computing stack performs better the closer this value is to 0. However, noise will likely make it incredibly difficult for even fault tolerant quantum computers to achieve a ℓ_1 -norm distance of 0. Hence bounds, such as that discussed in Theorem 1, are often put on the value instead. For the circuit classes for which such results are relevant, as we discuss in Section 3, we will use these results to define a values below which a performance can be considered as good. Once again, the ℓ_1 -norm distance takes a continuous range of values allowing for comparison between implementations of circuits.

2.4 Metric Comparison

Unfortunately, Cross-Entropy Benchmarking and Heavy Output Generation Benchmarking cannot be used to bound the ℓ_1 -norm distance [44], which, as noted in Section 2.3, provides strong guarantees of demonstrations of quantum computational supremacy.⁷ That is, the ℓ_1 -norm distance provides uniquely (amongst the metrics studied here) strong assurances about quantum computational supremacy. In general, such assurances come at the cost of requiring classical compute resources which grow exponentially with the circuit size.

In the case of Heavy Output Generation Benchmarking and Cross-Entropy Benchmarking only polynomially many single output probabilities are required, allowing the utilisation of *Feynman simulators* [39]. These compute output bitstring amplitudes by adding all Feynman path contributions. This extends the domain of classical simulation by overcoming the memory storage problem, establishing the frontier of what's possible on classical computers [3, 49, 50]. However, this method still requires exponential time to perform and so reaches its own limit for large numbers of qubits.

Importantly, since $\text{HOG}(\mathcal{D}_C, p_C)$ and $\text{CE}(\mathcal{D}_C, p_C)$ are expectations of different functions of ideal output probabilities, $\delta(p_C)$ and $-\log(p_C)$ respectively,

⁷Interestingly, empirical result show a slight negative correlation between ℓ_1 -norm distance and experimental heavy output probability (normalized to the heavy output probability of an ideal device). This relationship is discussed in Appendix D, which provides empirical details of this relationship.

over the experimental output distribution, they capture different features of the outputs [44]. In fact $\text{HOG}(\mathcal{D}_C, p_C)$ can also be used to approximate circuit fidelity, however the standard deviation of the estimator is larger than that for $\text{CE}(\mathcal{D}_C, p_C)$ [3].

3 Circuit Classes

This section presents the formal definitions of the circuits used in this work, while also identifying the motivations for their use in benchmarking. These motivations include both the class of applications they represent and the properties of the quantum computing stacks that they will probe. Collectively, this selection of circuit classes encompass an array of potential applications of quantum computing, covering circuits of varied depth, connectivity, and gate types.

3.1 Shallow Circuits: IQP

Instantaneous Quantum Polytime (IQP) circuits [51] consist of commuting gates and, as well as being simpler to implement than universal quantum circuits, are believed, even in the presence of noise, to be impossible to simulate efficiently using classical computers [47, 48, 52]. This has allowed for the application of noisy quantum technology in areas such as machine learning [30, 31, 53] and interactive two-player games [51, 54]. The *shallow* class of circuits introduced here, whose depth increases slowly with width, is a subclass of IQP circuits.

Definitions and Related Results An n -qubit IQP circuit consists of gates that are diagonal in the Pauli-X basis, acting on the $|0\rangle^n$ state, with measurement taking place in the computational basis. For this class of circuits, Theorem 1 applies.

Theorem 1 (Informal [47]) *Assuming either one of two conjectures, relating to the hardness of approximating the Ising partition function and the gap of degree 3 polynomials, and the stability of the Polynomial Hierarchy,⁸ it is impossible to classically sample from the output probability distribution of any IQP circuit in polynomial time, up to an ℓ_1 -norm distance of $1/192$.*

This class is called “instantaneous” because these gates commute with one another, which has the benefit of reducing the amount of time that the quantum state will need to be stored. Further, the impossibility of simulating IQP circuits is shown to hold when restricted by physically motivated constraints such as limited connectivity and constant error rates on each qubit [52]. This makes the class an exciting one to explore with near-term technology.

⁸The non-collapse of the Polynomial Hierarchy is widely conjectured to be true [55].

Algorithm 1 The pattern for building shallow circuits.

Input: Number of qubits, $n \in \mathbb{Z}$
Worst case depth: 7
Output: Circuit, C_n

```

1: Initialise  $n$  qubits, labelled  $q_1, \dots, q_n$ , in the state  $|0\rangle$ .
2:
3: for all  $i \in \{1, \dots, n\}$  do
4:   Act H on  $q_i$ 
5: end for
6:
7: Let  $G_n$  be the graph indicating between which qubits a CZ gate can act. Let  $G_n$  in this case be a random binomial graph,  $G_n$ , with  $n$  vertices and edge probability 0.5, post selecting on those that are connected and have degree less than 4.
8:
9: for all edges  $\{i, j\}$  in  $G_n$  do
10:   Act CZ between  $q_i$  and  $q_j$ 
11: end for
12:
13: for all  $i \in \{1, \dots, n\}$  do
14:   Generate  $\alpha_i \in [0, 2\pi]$  uniformly at random.
15:   Act RZ( $\alpha_i$ ) on  $q_i$ .
16: end for
17:
18: for all  $i \in \{1, \dots, n\}$  do
19:   Act H on  $q_i$ 
20: end for
21:
22: Measure  $q_1, \dots, q_n$  in the computational basis

```

An equivalent, commonly-considered definition is that IQP circuits consist of gates diagonal in the Pauli-Z basis, sandwiched between two layers of Hadamard gates acting on all qubits. Algorithm 1 is used to generate IQP circuits of this form. Note that Algorithm 1 limits the connectivity allowed between the qubits, so it does not generate *all* circuits in the IQP class.

The depth of this circuit may be arrived at by observing that finding a valid order in which to apply the CZ gates in the circuit is equivalent to finding an edge colouring of the graph G_n . By Vizing’s theorem, it is known that a colouring of undirected graphs using at most one more colour than the maximum vertex degree exists. Constructive proofs of Vizing’s theorem therefore demonstrate that a 4-colouring can be found in polynomial time [56].⁹ Algorithm 1 includes

⁹It is possible that the chromatic index of a particular graph is lower than 4, either because: the maximum degree is less than 3, as could be the case with these random graphs; or because a

discrete randomness over the graphs, G_n , and continuous randomness over the rotation angles, α_i .

The close connection, through Theorem 1, of quantum computational supremacy and shallow circuits,¹⁰ explicitly measured in ℓ_1 -norm distance, provides a measure of a quantum computing stack’s quality. This is namely, by analysing the closeness of the distributions it produces to the ideal ones, as measured by the ℓ_1 -norm distance, and comparing this value to $1/192$.

However, in the case of shallow circuits, the output probabilities are not exponentially distributed. As expanded on in Section 2.2 and Section 2.1, this property allows us to simplify calculations required when performing both Cross-Entropy Benchmarking and Heavy Output Generation Benchmarking. In particular the theoretical value of heavy output probability for circuits with exponentially distributed output probabilities cannot be used here.

Instead, we use the empirical value of the ideal heavy output probability, in the place of a theoretically derived one, as a point of comparison with the behaviour of the quantum computing stack being benchmarked. This approach requires calculation of all output probabilities and summation of the probabilities of those that are heavy, which is to say have greater than the median ideal probability. This can be done for the small circuits investigated here, but allows for the benchmarking of fewer qubits than would be accessible if a theoretical value was known.

Discussion The design of circuits in Algorithm 1 may be compared to other sparse IQP circuits [52], IQP circuits on 2D lattices [52, 57], and random 3-regular graphs used for benchmarking [21]. In our case we aim to avoid favouring particular architectures, and so avoid 2D lattices. In addition we avoid highly structured 3-regular graphs, in favour of allowing reduced vertex connectivity.¹¹ This avoids preferring very particular applications of low depth IQP circuits; instead exploring a variety of appli-

3-colouring exists, as is the lower bound on the chromatic index of an undirected graph of maximum vertex degree 3. However a 4-colouring certainly provide an upper bound and so it suffices for our discussion as we are concerned with upper bounding the depth of the circuit.

¹⁰While Theorem 1 is a worst case hardness result, and may not apply to shallow circuits, we regard their performance as indicative of that of those for which it does. Indeed, similar hardness results to Theorem 1 exist for other families of sparse, constant depth IQP circuits [57].

¹¹As n grows the probability that the maximum degree is at most 3 will tend to zero, and so Algorithm 1 may become infeasible to implement. While for the number of qubits considered in this paper Algorithm 1 works well, for larger n it may become necessary to alter the graph generation step, possibly to the use of 3-regular graphs, for which the fast random generation algorithms have been extensively researched and developed. However, regular graphs are too few in number for graphs with low numbers of qubits to provide a good benchmark set of many possible circuits, and so we avoid them here.

cations simultaneously. It is also of interest that there are sparse IQP circuits for which verification schemes exist [57, 58] although the connectivity is too architecture-specific for our purposes, with the verification scheme requiring limits to the measurement noise which we cannot guarantee.

Before compilation, shallow circuits have constant depth, allowing us to measure the impact of increasing circuit width independently of increasing circuit depth. Further, because Algorithm 1 limits the connectivity allowed between the qubits, the increase in circuit depth due to compilation onto limited-connectivity architectures is also minimised, while avoiding a choice of connectivity favouring one device in particular. By bounding connectivity, but allowing all connections in principle, we avoid biasing against architectures that allow all-to-all connectivity, which would still perform well.

In summary, inspired by a variety of near-term applications of IQP circuits [30, 31, 51, 53, 54], we introduce shallow circuits in Algorithm 1. Performance when implementing these circuits is indicative of the performance when implementing those applications, but also, more generally, of applications requiring circuits which grow slowly in depth. As a result, these circuits also probe the performance of a quantum computing stack in fine-grained detail by measuring the impact of including more qubits (quasi-) independently of increasing circuit depth. This is useful for understanding the performance of a device being utilised for applications whose qubit requirement grows more quickly than the circuit depth.

3.2 Square Circuits: Random Circuit Sampling

While circuits required for applications are typically not random, sampling from the output distributions of random circuits built from two-qubit gates has been suggested as a means to demonstrate quantum computational supremacy [27, 39, 44, 45]. By utilising uniformly random two-qubit unitaries and all-to-all connectivity, *square circuits*, introduced now, also provide a benchmark at all layers of the quantum computing stack. With a circuit depth that grows linearly with the number of qubits, square circuits also probes the performance of a quantum computing stack in ways not accessed by shallow circuits.

Definitions and Related Results A random circuit, for a fixed number of qubits n and graph G_n , is generated by applying $m = \text{poly}(n)$ uniformly random two-qubit $\text{SU}(4)$ gates between qubits connected by edges of G_n . Here, “uniformly random” means according to the Haar measure.

Random Circuit Sampling (RCS) is the task of producing samples from the output distribution of random circuits. To perform RCS approximately is to sample from a distribution close to that produced by

the random circuit. For certain classes of graphs G_n , this task has been shown to be hard in the average case [44, 45], as outlined in Theorem 2.

Theorem 2 (Informal [44]) *There exists a collection of graphs G_n , with one for each n , and procedure for generating random circuits respecting each G_n , for which there is no classical randomised algorithm that performs approximate RCS, to within inverse polynomial ℓ_1 -norm distance error, for a constant fraction of the random circuits.*

The conditions imposed on which graphs and circuit generation procedures are covered by this theorem are quite mild, but in particular this can be done using circuits with depth $\mathcal{O}(n)$ acting on a 2D square lattice [39, 44]. Note that in that case G_n are not generated at random, but systematically according to n , unlike the graphs used to generate shallow circuits in Algorithm 1. While this is relevant for devices built using superconducting technology [3], we wish to avoid biasing in favour of this technology in particular.

The circuits used here – which are almost identical to those used for the quantum volume benchmark [33] – are generated according to Algorithm 2. We refer to this class of circuits as square circuits, and note that they consist of n layers of two-qubit gates acting between a bipartition of the qubits. There is discrete randomness over the possible bipartition of the qubits, and continuous randomness over the matrix elements of two-qubit $SU(4)$ gates selected uniformly at random from the Haar measure.

In addition, square circuits fulfil the necessary conditions to apply HOG, as defined in Problem 1. Namely, the distribution p_C is sufficiently far from uniform in the required sense, as introduced in Section 2.1, which we demonstrate in Appendix A.1.

Discussion By utilising uniformly random two-qubit unitaries, square circuits provide a benchmark at all layers of the quantum computing stack. In particular this tests the ability of the device to implement a universal gate set, the diversity and quality of the gates available, and the compilation strategy’s ability to decompose these gates to the native architecture. Further, as quantum circuits can always be approximated up to arbitrary precision using two-qubit unitary gates [59], square circuits can help us understand the performance of quantum computing stacks when implementing computations requiring a universal gate set, which is less the case for shallow circuits and deep circuits.

By allowing two-qubit gates to act between any pair of qubits in the uncompiled circuit, square circuits avoid favouring any device in particular [3, 27, 39]. This choice adheres closely to our motivations of being hardware-agnostic. In addition, assuming all-to-all connectivity passes the burden of mapping

Algorithm 2 The pattern for building square circuits.

Input: Number of qubits, $n \in \mathbb{Z}$

Worst case depth: n

Output: Circuit, C_n

```

1: Initialise  $n$  qubits, labelled  $q_1, \dots, q_n$ , in the state  $|0\rangle$ 
2:
3: for each layer  $t$  up to depth  $n$  do
4:      $\triangleright$  The contents of this for loop constitutes a layer. The choice of the number of layers used here is discussed in Appendix A.1.
5:
6:     Divide the qubits into  $\lfloor \frac{n}{2} \rfloor$  pairs  $\{q_{i,1}, q_{i,2}\}$  at random.
7:     for all  $i \in \mathbb{Z}$ ,  $0 \leq i \leq \lfloor \frac{n}{2} \rfloor$  do
8:         Generate  $U_{i,t} \in SU(4)$  uniformly at random according to the Haar measure.
9:         Act  $U_{i,t}$  on qubits  $q_{i,1}$  and  $q_{i,2}$ .
10:    end for
11: end for
12:
13: Measure all qubits in the computational basis.
```

the circuit onto the device to the compilation strategy, which is in line with our wish to benchmark the full quantum computing stack. That said, any architecture whose coupling map closely mirrors the uncompiled circuit will be advantaged, as even a naive compilation strategy will perform well in that case.

In [33] similar circuits are used but with all-to-all connectivity restricted to nearest neighbour connectivity on a line, and the addition of permutation layers. As this disadvantages devices with a completely connected coupling map¹² [5], a property which would typically be an advantage, we choose not to make this restriction here. Notice, however, that naively compiling square circuits onto an architecture with nearest neighbour connectivity on a line would result in the circuits of [33]. This similarity makes a comparison between experiments involving these circuits relevant. As a result, compiling square circuits to superconducting devices (where connectivity is low) will generally result in a circuit similar to those used in the quantum volume benchmark, as many SWAP operations are required regardless.

In summary, by allowing for the most general gate set and connectivity, square circuits provide a means to thoroughly interrogate all layers of the quantum

¹²Note that some compilation strategies may identify that the SWAP gates in the permutation layer may be removed for devices with all-to-all connectivity. We avoid this dependence on the compilation strategy by fixing the connectivity in the uncompiled circuit to be all-to-all.

computing stack. This class of circuits is inspired by quantum volume circuits [33] but are somewhat generalised to avoid device bias. Because of this generality, and the linear increase in depth with number of qubits, square circuits are complementary to the other classes introduced in this work.

3.3 Deep Circuits: Pauli Gadgets

Pauli gadgets [60] are quantum circuits implementing an operation corresponding to exponentiating a Pauli tensor. Sequences of Pauli gadgets acting on qubits form *product formula* circuits, most commonly used in Hamiltonian simulation [34]. Many algorithms employing these circuits require fault-tolerant devices, but they are also the basis of trial state preparation circuits in many variational algorithms, which are the most promising applications of noisy quantum computers.

A notable example of this in quantum chemistry is the physically-motivated Unitary Coupled Cluster family of trial states used in the variational quantum eigensolver (VQE) [35, 61]. As near-term quantum computers hold promise as useful tools for studying quantum chemistry, we propose that the quality of an implementation of these gadgets is a useful benchmark, and use them to define the *deep* circuit class.

Definitions and Related Results These circuits are built as in Algorithm 3. They are constructed from several layers of Pauli Gadgets, each acting on a random subset of n qubits. In the worst case each Pauli Gadget will demand $4n+1$ gates: $2n$ Pauli gates, $2(n-1)$ CX gates, and one RZ gate. In the construction of deep circuits there is discrete randomness over the choice of Pauli string, s , and continuous randomness over a rotation angle α .

By establishing the exponential distribution of the output probabilities from deep circuits, as we do in Appendix A.2, we allow ourselves the capacity to use Heavy Output Generation Benchmarking and Cross-Entropy Benchmarking as introduced in Section 2. This constitutes a novel extension of those approaches to application motivated benchmarking, and the unique ability for us to benchmark application-motivated circuits, using polynomially many samples from a device. This provides a novel insight into the capacity of near-term hardware to implement quantum chemistry circuits.

Discussion Note that the circuits in this class differ from running the VQE end-to-end. Focusing on the state preparation portion of a VQE circuit, we might deduce performance of the quantum computing stack

Algorithm 3 The pattern for building deep circuits.

Input: Number of qubits, $n \in \mathbb{Z}$
Worst case depth: $(4n-1)(3n+1)$
Output: Circuit, C

```

1: function PHASEGADGET( $\alpha, \{\tilde{q}_1, \dots, \tilde{q}_p\}$ )
2:   if  $p = 1$  then
3:     Act RZ( $\alpha$ ) on  $\tilde{q}_1$ 
4:   else
5:     Act CX between  $\tilde{q}_1$  and  $\tilde{q}_2$ 
6:     PHASEGADGET( $\alpha, \{\tilde{q}_2, \dots, \tilde{q}_p\}$ )
7:     Act CX between  $\tilde{q}_1$  and  $\tilde{q}_2$ 
8:   end if
9: end function
10:
11: function PAULI( $\{\tilde{q}_1, \dots, \tilde{q}_p\}, s$ )
12:   if  $s_1 = X$  then
13:     Act RX( $\frac{\pi}{2}$ ) on  $\tilde{q}_1$ 
14:   else if  $s_1 = Y$  then
15:     Act H on the  $\tilde{q}_p$ 
16:   end if
17:   PAULI( $\{\tilde{q}_2, \dots, \tilde{q}_p\}, s$ )
18: end function
19:
20: function PAULIGADGET( $\alpha$ , qubits,  $s$ )
21:   PAULI(qubits,  $s$ )
22:   PHASEGADGET( $\alpha$ , qubits)
23:   Act the inverse of PAULI(qubits,  $s$ )
24: end function
25:
26: Initialise  $n$  qubits, labelled  $q_1, \dots, q_n$ , in the state  $|0\rangle$ .
27:
28: for each layer  $t$  up to depth  $3n+1$  do
29:   ▷ The contents of this for loop constitutes
   a layer. The choice of the number of layers used
   here is discussed in Appendix A.2.
30:
31:   Select a random string  $s \in \{I, X, Y, Z\}^n$ 
32:   Generate random angle  $\alpha \in [0, 2\pi]$ 
33:   PAULIGADGET( $\alpha, \{q_i : s_i \neq I\}, s$ )
34:
35: end for
36:
37: Measure all qubits in the computational basis

```

when running the VQE on a number of molecules.¹³ The intuition being that if the state preparation is accurate, then the error in the expectation values of measured observables will be due to errors in implementing those observables, or the readout process itself.

In summary, we have introduced deep circuits to probe the performance of a quantum computing stack when implementing common quantum chemistry applications. The deep circuits circuit class is complementary to the other classes featured in this work by having a depth which grows quickly with the number of qubits.

4 Quantum Computing Stack

Each component of a quantum computing stack exerts an influence on overall performance, and identifying the distinct impact of a particular component is often hard. To disentangle these factors, we must clearly identify the components used during benchmarking. Here we detail the components used to build the quantum computing stacks explored in Section 5. The diverse selection of components allows us to investigate a variety of ways of building a quantum computing stack.

4.1 Software Development Kits

We use a combination of tools available via `pytket` [62, 63] and `Qiskit` [1, 64]. `pytket` is a Python module which provides an environment for constructing and implementing quantum circuits, as well as for interfacing with `t|ket`, a retargetable compiler for near term quantum devices, featuring hardware-agnostic optimisation. `Qiskit` is an open-source quantum computing software development framework for programming, simulating, and interacting with quantum processors, which also provides a compiler. Details of the versions of the software used are seen in Table 2 of Appendix B.

We use three parts of `Qiskit` in this work. First is the *transpiler architecture*, which enables users to define a custom compilation strategy by executing a series of *passes* on the input circuit, as discussed in Section 4.2. The second part of `Qiskit` we use is the library of predefined passes. Finally, a *provider* is used to access hardware made available over the cloud by IBM Quantum. The provider enables users to send circuits to hardware, retrieve results, and query the

¹³We do not explore the relationship between the performance of a quantum computing stack when implementing deep circuits and when implementing VQE, but regard it to be important future work.

hardware for its properties.¹⁴

Similarly, we use `pytket` to generate and manipulate circuits in several ways. Firstly we use the `t|ket` compiler to construct compilation strategies which optimise the input circuit for the target hardware, utilising predefined passes available in `t|ket`. Secondly we use `pytket` to define abstract circuits and to convert `t|ket`'s native representation of the circuit into a `Qiskit QuantumCircuit` object which is then dispatched to IBM Quantum's systems for execution.

4.2 Compilers

Compilers provide tools to construct executable quantum circuits from abstract circuit models. This is done by defining *passes* which may manipulate a representation of a quantum circuit, often by taking account of limited connectivity architectures, or minimising quantities such as gate depth, but need not perform any manipulation.¹⁵ These passes are composed to form *compilation strategies* which should output executable quantum circuits. Quantum compiling is an active area of research [65–71], and there are many pieces of software available for quantum compiling. As noted above, in this work we use two: `t|ket` and the compiler available in `Qiskit`.

For the purposes of this work, the problem of quantum compilation is divided into three tasks.

Placement: Determine onto which *physical* qubits of a given device the *virtual* qubits in the circuit's representation should be initially mapped.

Routing: Modify a circuit to conform to the qubit layout of a specific architecture, for example, by inserting `SWAP` gates to allow non-adjacent qubits to interact [72]. Circuits are often designed without the device's coupling map in mind, so this step is important [21].¹⁶

Optimisation: Optimise the circuit according to some cost metric, usually gate count or circuit depth.

Each of these tasks could consider such things as the trade-offs between the connectivity of a particular subgraph of the device and the amount of crosstalk present in that subgraph [73].

Both `pytket` and `Qiskit` have multiple placement, optimisation, and routing passes. We compare the performance of 5 compilation strategies built from these

¹⁴These properties include the graph connectivity, single- and two-qubit error rates, and qubit T_1 and T_2 times. Some of the *noise-aware* compilation strategies we use require knowledge of these properties.

¹⁵An example of this is a pass which counts the gates in the circuit.

¹⁶While hardware-efficient *ansätze* can prove advantageous in task such as machine learning and quantum chemistry in the near term, in general they may be restrictive and so we consider the more general problem.

passes. Two of them, noise-unaware pytket and noise-unaware Qiskit, compile the circuit *without* knowledge of the device’s noise properties. Another two, noise-aware pytket and noise-aware Qiskit, do take noise properties into account. As a base line, we consider a simple compilation strategy from pytket using only routing, without optimisation or noise-awareness; we refer to this compilation strategy as *only pytket routing*. We detail these schemes in Appendix B. The main difference between the noise-aware schemes is that noise-aware pytket prioritises the minimisation of gate errors during placement,¹⁷ whereas noise-aware Qiskit prioritises readout and CX errors [69].

4.3 Devices

We benchmark some of the devices made available over the cloud by IBM Quantum. The devices we use are referred to by the unique names `ibmqx2`, `ibmq_16_melbourne`, `ibmq_singapore` and `ibmq_ourense`. Each device has a set of *native gates* which all gates in a given circuit must be decomposed to. For all the devices considered here, the native gates are: an identity operation, `I`; 3 “*u-gates*” [74], as defined in equation (11); and a controlled-NOT (CX) gate.

$$\begin{aligned} U_3(\theta, \phi, \lambda) &= \begin{pmatrix} \cos(\frac{\theta}{2}) & -e^{i\lambda} \sin(\frac{\theta}{2}) \\ e^{i\phi} \sin(\frac{\theta}{2}) & e^{i(\lambda+\phi)} \cos(\frac{\theta}{2}) \end{pmatrix} \\ U_2(\phi, \lambda) &= U_3(\frac{\pi}{2}, \phi, \lambda) \\ U_1(\lambda) &= U_3(0, 0, \lambda) \end{aligned} \quad (11)$$

Two of the device properties used by the noise-aware compilation strategies are their *connectivity* and *calibration data*. Information about the connectivity of a device is contained in its *coupling map* which, in the cases of the devices studied here, are shown in Appendix C.1 and summarised in Table 1. A coupling map is a graph with a vertex for each qubit, and edges between vertices when there exists coupling between the corresponding qubits.

We expect the coupling map of a device will influence the device’s performance according to our benchmarks. Coupling maps with a high average *degree* (mean number of edges incident on each vertex) have more qubits directly connected to one another, reducing the requirement for SWAP operations that would increase the circuit depth. Coupling maps with low *radius* (minimax distance over all pairs of qubits) have qubits which are closer to one another, again reducing the need for SWAP gates. Coupling maps with a high number of *vertices* (i.e., qubits), enables higher-width circuits to be run. Finally, coupling maps where the number of vertices equals the *minimum cycle length* (smallest number of edges per cycle over all cycles)

¹⁷This is true of pytket 0.3.0; as a result of this work, later versions of pytket take into account readout error.

have the advantage that any routing operation has more paths available to it, potentially allowing for some parallelisation.

Device calibration data includes information about single- and two-qubit error rates, readout error, and qubit frequency, T_1 , and T_2 times. The noise-aware compilation strategies we investigate use the gate error rates and readout error rates. Full details of noise levels can be found in Appendix C.2 with average values given in Figure 1. This information is updated twice daily, with the data in Figure 1 averaged over the period 2020-01-29 to 2020-02-10 during which time our experiments were conducted.

The results of Section 5 depend heavily on the noise levels of the device at the time at which the computation is implemented. This is doubly true in the case of the noise-aware optimisation schemes as a circuit optimised at one time may not perform as well over time as the noise levels of the devices change. To reduce this effect we endeavoured to compile and run circuits within as short a time interval as possible.

5 Experimental Results

This section presents experimental results of running the circuit families defined in Section 3 on various quantum computing systems. We present a sub-sample of all our results in the figures which follow. We study these results in 3 contexts:

Section 5.1: Full Stack Benchmarking – Incorporating and thoroughly investigating the compilation strategy helps develop an understanding of how circuit compilation influences the performance of the quantum computing stack. For noise-aware compilation strategies, our results show that the assumptions made by the strategy about the importance of different kinds of noise impacts performance.

Section 5.2: Application Motivated Benchmarks – By including three quite different circuit classes in our benchmark suite, we explore how a quantum computing stack may perform when implementing a wide array of applications.

Section 5.3: Insights from Classical Simulation – We explore how benchmarks assist in developing new noise models. By identifying when benchmark values for real implementations and those we expect from simulations using noise models differ, noise channels which should be added to the noise models to achieve greater agreement with real devices can be identified. This is of particular importance as noise-aware compilation strategies often utilise noise properties.

For each circuit class and fixed number of qubits, 200 circuits were generated according to the algorithms of

Device	#Vertices	Average Degree	Radius	Minimum Cycle Length
ibmqx2	5	2.4	1	3
ibmq_16_melbourne	15	$2\frac{2}{3}$	4	4
ibmq_ourense	5	1.6	2	N/A
ibmq_singapore	20	2.3	4	6

Table 1: **Selected graph properties of the coupling maps of devices studied in this work.** See Appendix C.1 for full details of the coupling maps of the devices explored here.

Figure 1: **Average error rates across devices used in this work.** Bars show the mean error rates across the whole device, while error bars give the standard deviation. Devices shown here are: ibmqx2 [■], ibmq_ourense [■], ibmq_singapore [■], ibmq_16_melbourne [■]. Further details can be found in Appendix C.2

Section 3. Each circuit is compiled by a given compilation strategy onto a particular device. The compiled circuits were then run on the device, using 8192 repetitions (samples) from each compiled circuit, which generates 8192 bitstrings. The compiled circuits are also classically simulated using a noise model built from the device calibration information at the time of the device run. See *Data Availability* for access to the full experimental data set.

The resulting bitstrings are then processed according to the figures of merit given in Section 2. The distribution of the figures of merit are compared by their mean and distribution via a box-and-whisker plot. In particular boxes show quartiles of the dataset, whiskers extend to 1.5 times the interquartile range past the low and high quartiles, and white circles give the mean. Uncompiled circuits were also perfectly simulated without noise in order to calculate the ideal heavy output probability. These points are referred to as *Noise-Free* in the figures below.

5.1 Full Stack Benchmarking

Impact of the Compilation Strategy

We study the compilation strategy and the device on which the compiled circuit is run. Using a fixed device and comparing multiple strategies allows us to determine which strategy tends to perform well. Fur-

ther, aggregating performance over all compilation strategies provides an estimate of the performance of a “generic” strategy. Similarly, fixing the strategy, and comparing its performance on multiple devices, shows how (possibly-erroneous) assumptions made by the strategy about the devices impact performance.

Figure 2 displays experimental results when implementing square circuits on ibmq_16_melbourne, using heavy output generation probability as the figure of merit. The noise-aware pytket compilation strategy performs somewhat better, on average, than a generic strategy. Because the aggregated information (“All Strategies” in Figure 2) includes aggregation over noise-aware pytket, these results indicate that other compilation strategies perform a bit worse, since the performance of the aggregate is generally lower than that of noise-aware pytket.¹⁸ This reveals the potential for compilation strategy driven improvements in performance.

Aggregation over compilation strategies provides a way of identifying devices which perform well, by “washing out” the effect of the compilation strategy on performance. Figure 3 shows that by considering performance with a fixed compilation strategy

¹⁸Note that due to the fact we aggregate over all 5 compilation strategies, the distribution of heavy output probabilities amongst the “All Strategies” category contains five times as many points as compared to those for noise-aware pytket.

Figure 2: **Comparison of fixed compilation strategy to average of all strategies using, the heavy outputs probability metric.** Here we have run square circuits using the real `ibmq_16_melbourne` device. Values above $2/3$ (dotted blue line), and closer to Noise Free, indicate better performance.

`ibmq_singapore` would be considered to perform similarly, if not slightly better than `ibmq_ourense`, as measured by ℓ_1 -norm distance. However, aggregating over all strategies, (Figure 4) shows `ibmq_ourense` to perform better, suggesting that `ibmq_ourense` might be a better device for a “generic” compilation strategy to compile to.

An instance-by-instance comparison of different compilation strategies also reveals their limitations and advantages. For example, Figure 5 shows `noise-aware pytket` works best at reproducing the ideal distribution of heavy output probabilities of square circuits on `ibmq_16_melbourne`. When compared to the strong performance of only `pytket` routing, this suggests these results are due in part to the routing scheme.

Similarly, Figure 6 shows that `noise-aware pytket` is amongst the worst-performing compilation strategies for lower numbers of qubits, while it is amongst the best-performing for higher numbers. This could be a result of the way in which `noise-aware pytket` prioritises noise in its routing scheme, with gate errors taking precedence.¹⁹

Noise Level, Connectivity Trade Off

More highly-connected architectures typically allow for shallower implementations of a given circuit as compared to less-connected ones. However, the noise levels may be higher due to crosstalk [75], resulting in a trade-off between connectivity and the total amount of noise incurred when running a computation. Noise affects the accuracy of the computation, so this trade-off has practical implications for the performance of a device. Reducing the connectivity between superconducting qubits has been used to reduce noise levels [75]. In superconducting qubits, this can also be counteracted using tunable couplers [3], but this is

¹⁹For larger numbers of qubits and deeper circuits, gate errors becomes more impactful on the total noise, and materialises as giving `noise-aware pytket` an advantage for larger numbers of qubits.

not utilised in the devices studied here.²⁰

Figure 7 shows that devices with lower noise levels (`ibmq_singapore` and `ibmq_ourense`) typically outperform devices with higher noise levels (`ibmqx2` and `ibmq_16_melbourne`) despite the latter’s higher connectivity. An interesting exception to this is for 4 qubits, where `ibmq_16_melbourne` performs best, likely because of the 4-qubit cycles in its connectivity graph, as mentioned in Table 1. This reduces the SWAP operations necessary for implementing the circuit, reducing the overall circuit depth. This reveals the increase in performance that can be expected when the connectivity of the device and the problem instance are similar [21]. Similar results hold for Cross-Entropy Benchmarking, as shown in Figure 8.

In general, we expect that circuits whose structure can naturally be mapped to the connectivity of the device will generally perform well, whereas those which cannot, will not. In general though, lower-noise devices will tend to perform best.

5.2 Application Motivated Benchmarks

The same quantum computing stack will perform differently when running different applications, as the structure of the circuits they require will generally be different. Differences in performance are seen in the context of our application-motivated benchmarks. For example, consider Figure 9, which shows performance when implementing sparsely connected circuits, and Figure 10, which shows performance when implementing chemistry-motivated circuits. In the case of Figure 9, the `ibmqx2` device outperforms `ibmq_singapore`, while in the case of Figure 10 the reverse is true.

²⁰While we focus on the connectivity of superconducting architectures here, more generally the comparison between the limited connectivity of superconducting devices, and the completely connected coupling maps of ion trap devices is of interest [10, 18, 20].

Figure 3: **Comparison of devices using the ℓ_1 -norm distance metric, when running shallow circuits compiled using noise-aware pytket.** Values close to 0 indicate better performance. Values below $1/192$ (dotted blue line) can be regarded as performing very well. Both simulations using Qiskit noise models, and implementations on real devices, are included.

Figure 4: **Comparison of real devices, using the ℓ_1 -norm distance metric, when running shallow circuits compiled using all compilation strategies.** Here we compile onto each device using all compilation strategies, including all compiled circuits in this plot. Values close to 0 indicate better performance. Values below $1/192$ (dotted blue line) can be regarded as performing very well.

Figure 5: **Comparison of compilation strategies, using the heavy outputs probability metric, when running square circuits on the real ibmq_16_melbourne device.** Values above $2/3$ (dotted blue line), and closer to Noise Free, indicate better performance.

Figure 6: **Comparison of compilation strategies, using the ℓ_1 -norm distance metric, when running shallow circuits on the real ibmq_ourense device.** Values close to 0 indicate better performance. Values below $1/192$ (dotted blue line) can be regarded as performing very well.

Figure 7: **Comparison of devices using the heavy outputs probability metric, when running square circuits compiled using noise-aware pytket.** Values above $2/3$ (dotted blue line), and closer to Noise Free, indicate better performance. Both simulations using Qiskit noise models, and implementations on real devices, are included.

Figure 8: **Comparison of devices using the cross entropy difference metric, when running square circuits compiled using noise-aware pytket.** Values close to 1 indicate better performance. Both simulations using Qiskit noise models, and implementations on real devices, are included.

Figure 9: **Comparison of real devices, using the cross entropy difference metric, when running shallow circuits compiled using noise-aware Qiskit.** Values close to 1 indicate better performance.

Figure 10: **Comparison of devices using the cross entropy difference metric, when running deep circuits compiled using noise-aware Qiskit.** Values close to 1 indicate better performance. Both simulations using Qiskit noise models, and implementations on real devices, are included.

Quantum Chemistry

Figure 10 suggests `ibmq_ourense` is best for quantum chemistry applications, because it performs well when running deep circuits.²¹ In particular Figure 10 indicates that the average circuit fidelity is highest for implementations on `ibmq_ourense`.

In Figure 10, all devices converge to the minimum value of cross-entropy difference at 4 qubits. To extend an investigation of this sort to more qubits would require lower noise levels or chemistry motivated circuits which generate exponentially distributed output probabilities at lower depth.

Shallow Circuits as a Benchmark

Figure 11 demonstrates that shallow circuits allow us to benchmark the behaviour of a quantum computing stack for applications involving circuits with many qubits but low circuit depth [30, 51, 54].

The results show `ibmq_singapore` outperforms the comparably sized `ibmq_16_melbourne` and has comparable performance to `ibmq_ourense` for smaller numbers of qubits. `ibmq_singapore` outperforms `ibmq_ourense` by having more qubits available. This superior performance of `ibmq_singapore` is in comparison to the results of Figure 7, where `ibmq_ourense` was shown to perform well. This justifies our sugges-

²¹This comes with the caveat, as mentioned in Section 3.3, that the connection between the quality of an implementation of these computational primitives, as measured by this benchmark, and accurate ground state energy calculations in VQE has not been demonstrated experimentally.

tion that shallow circuits should be included in benchmarking suites. Doing so enables exploring higher-width circuits, and in this setting devices that perform poorly when implementing square circuits or deep circuits may perform well.

5.3 Insights from Classical Simulation

Noise in a non-fault-tolerant quantum computer results in discrepancies between results obtained from running on real hardware and those that would be obtained from an ideal quantum computer. Noise models are utilised to help identify why these discrepancies occur [76]. However, a perfect model of the noise – which could reproduce the results of real hardware (up to statistical error) – could require many parameters to completely specify it. Therefore, most noise models consider only a small handful of physically-motivated noise sources. Consequently discrepancies between the results of noisy simulation and running experiments on real hardware always remain.

Historically, closing this gap required developing noise models of increasing sophistication. Doing so typically requires a great deal of physics expertise to identify new noise channels. Further, new experiments would have to be designed in order to estimate the parameters in the noise model.

Here, we show how application-motivated benchmarks can be helpful in identifying whether new noise channels should be incorporated into a noise model. By isolating the circuit types and coupling maps for which the discrepancies are greatest, we gain intuition about the possible causes of the mismatch.

Figure 11: **Comparison of real devices, using the heavy outputs probability metric, when running shallow circuits compiled using noise-aware pytket.** Values close to Noise Free indicate better performance.

For the devices explored here, the noise models are built using Qiskit. They are derived from a device’s properties, and include one- and two-qubit gate errors²² and single-qubit readout errors. We find these noise models are inadequate to explain some of the discrepancies observed in the data.

Noise Does Not Just Flatten Distributions

One discrepancy between experiments and noisy simulations is the spread of the data. For example, Figure 7 shows that only in the experimental case do the whiskers of the plot fall below the value 0.5, indicating the heavy outputs are less likely than they would be in the uniform distribution. Some noise type, in particular one which shifts the probability density, rather than uniformly flattening it, is not considered, or is under appreciated, by the noise models used. Identifying that noise channel is left to future work, though we speculate it may be related to a kind of thermal relaxation error.

Noise Models Under Represent Some Noise Channels

The classical simulations in Figure 7 suggest ibmqx2 should perform similarly to ibmq_ourense in most cases. In fact, it quite consistently performs worse. This is isolated in Figure 8, with the same phenomenon being observed in Figure 3 and Figure 10, showing the behavior is consistent across all circuit types and figures of merit.

This difference between simulated and experimental results is pronounced in the case of Figure 10, where deep circuits are used, suggesting the noise models may be underestimating the error from time-dependent noises such as depolarising and dephasing, or from two-qubit gates which are more prevalent in deep circuits.

Another such example of a two-qubit noise channel, which is explicitly not accounted for in our noise models, is crosstalk. The results in Figure 8 are consistent with the expectation that cross-talk should have

²²These are modelled to consist of a depolarising errors followed by a thermal relaxation errors.

the greatest impact on more highly connected devices [75]. As such crosstalk may be the origin of the discrepancy. Of note is the fact this benchmark wasn’t explicitly designed to capture the effects of crosstalk, and yet those effects manifest themselves in its results. We anticipate that including crosstalk-aware passes in compilation strategies [73] would reduce the discrepancy.

6 Conclusion

The performance of quantum computing devices is highly dependent on several factors. Amongst them are the noise levels of the device, the software used to construct and manipulate the circuits implemented, and the applications for which the device is used. The impact of these factors on the performance of a quantum computing stack are intertwined, making the task of predicting its holistic performance from knowledge of the performance of each component impossible. In order to understand and measure the performance of quantum computing stacks, benchmarks must take this into consideration.

In this work we have addressed this problem by introducing a methodology for performing application-motivated, holistic benchmarking of the full quantum computing stack. To do so we provide a benchmark suite utilising differing circuit classes and figures of merit to access a variety of properties of the device. This includes the use of three circuit classes: deep circuits and shallow circuits, which are novel to this paper; and square circuits, which resemble random circuits used in other benchmarking experiments [33]. In addition we make use of a diverse selection of figures of merit to measure the performance of the quantum computing stacks considered, namely: Heavy Output Generation Benchmarking, Cross-Entropy Benchmarking, and the ℓ_1 -norm distance.

In particular, in the form of deep circuits we present an alternative to previous approaches to application-motivated benchmarking. This is by considering circuits inspired by one of the primitives utilised in VQE, namely Pauli gadgets employed for state preparation,

rather than VQE itself. Further, while we have found that the performances of quantum computing stacks are indistinguishable when using square circuits and Heavy Output Generation Benchmarking for a large number of qubits, shallow circuits extend the number of qubits for which detail can be observed, while also being consistent with philosophy of volumetric benchmarking.

We demonstrate this benchmark suite by employing it on `ibmqx2`, `ibmq_16_melbourne`, `ibmq_ourense`, and `ibmq_singapore`. In doing so we justified our thesis that the accuracy of a computation depends on several levels of the quantum computing stack, and that each layer should not be considered in isolation. For example, identifying that the increased connectivity of a device does not compensate for the increased noise, as we do in Section 5.1, shows the impact of this layer of the stack, and justifies investigating devices with a variety of coupling maps and noise levels. By showing the differing performance between five compilation strategies, we are able to identify, in Section 5.1, the dependence of the best compilation strategy to use on the device and the dimension of the circuit. This illustrates the dependence of the performance of the quantum computing stack on the compilation layer, and the interdependence between the compilation strategy, device and application on the overall performance of the quantum computing stack. In particular, noise-aware compilation strategies often perform well, when the noise model used by the strategy is accurate, as discussed in Section 5.3.

In Section 5.2, the wide selection of circuits within the proposed benchmark suite reveals that the same device, evaluated according to a fixed figure of merit, will perform differently when running different applications, whose circuits are compiled by the same compilation strategy. Indeed the comparative performance of (compilation strategy, device) pairs is shown to vary between our circuit classes. This justifies our inclusion of circuit classes which collectively cover a wide selection of applications in the benchmark suite proposed here, and our full quantum computing stack approach.

We foresee the benchmarks conducted in this work providing a means to select the best quantum computing stack, of those explored here, for a particular task, and vice versa. As such we also anticipate that a variety of new quantum computing stacks could be benchmarked in the way described in this work, empowering the user with knowledge about the performance of current quantum technologies for particular tasks.

These benchmarks may, in time, come to complement noise models and calibration information as a means to disseminate information about a device’s performance. This parallels the use of the LINPACK benchmarks [7] alongside FLOPS to compare diverse classical computers. Recently, quantum volume, as

defined in [33], has started to be adopted as one such metric [77], and we hope the benchmark suite developed here will be incorporated similarly. Further, our benchmarks may facilitate an understanding of how new, or hard-to-characterise, noise affects the practical performance of quantum computers, as implied by the classical simulations of Section 5.3.

The work presented here could be extended in several directions. The first is to examine the impact of incorporating these benchmarks into a compilation strategy. While noise-aware compilation strategies currently use properties of qubits to decide how to compile a circuit, it would be interesting to explore if instead optimising for these benchmarks would change the compilation. The trade off between the benefits of doing so against the increased compilation time resulting from the time taken to perform the benchmarks should then also be assessed. This information would help in the understanding of the interplay between the amount of classical circuit optimisation performed and the amount by which the performance of a quantum system can be increased.

Second, the philosophy of application-motivated benchmarking could be extended to circuits which are more easily classically simulable. Because of their reliance on classical simulation, the benchmarks introduced here may be used up to, but not after, the point of demonstrating quantum computational supremacy. Hence new circuit classes will need to be introduced which can be classically simulated in this regime. Alternatively, application-motivated benchmarks that are derived from combining benchmarks of smaller devices [3] could be developed.

Third, we envision a need to systematically study how properties of hardware, such as noise levels or connectivity, influence a given device’s performance. In this work, we were limited to the particular devices made available by IBM Quantum, which limits our ability to perform such a systematic inquiry. It is nevertheless vital to do so, as the results of Section 5.1 show that changing the hardware can dramatically influence performance. Indeed, this would allow us to understand if the observations made in Section 5.1 are typical, and to explore the existence of other relationships. This could be achieved by implementing this benchmark suite on more devices, or synthetic devices with tunable coupling maps and noise information.

Finally, there is a need to study the correlation between the results of an application-motivated benchmark and the performance of a quantum computing stack at running the application which motivated it. This would show that benchmarking application sub-routines provides reliable predictors of performance when running the application itself. While similar work has explored the correlation between the classification accuracy and circuit properties of parametrised quantum circuits [32], comparing the performance of the benchmarks defined here with their applications

is a subject for future work. For example, comparing the performance of a stack at implementing deep circuits and running the VQE algorithm would show the extent to which quantum computing stacks that perform well at a particular kind of state preparation circuit also perform well in estimating properties of a wide range of molecules.

Data Availability

The data gathered during the experiments conducted for this work, as presented in Section 5, are available at [78] - <http://doi.org/10.5281/zenodo.3832121>. QASM files representing the circuits executed, the per sample bitstring outputs from device and simulator executions, and device calibration data gathered throughout the course of the experiments, are provided.

Acknowledgements

DM acknowledges support from the Engineering and Physical Sciences Research Council (grant EP/L01503X/1). This work was made possible, in part, by systems built by IBM as part of the IBM Quantum program, and made accessible via membership in the IBM Q Network.

All authors would like to thank the anonymous reviewers for their careful consideration of this work, and for their insightful comments and suggestions.

IBM, IBM Q, and Qiskit are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product or service names may be trademarks or service marks of IBM or other companies.

References

- [1] Gadi Aleksandrowicz et al. *Qiskit: An Open-source Framework for Quantum Computing*. Version 0.7.2. Jan. 2019. DOI: [10.5281/zenodo.2562111](https://doi.org/10.5281/zenodo.2562111).
- [2] Peter J Karalekas et al. “A quantum-classical cloud platform optimized for variational hybrid algorithms”. In: *Quantum Science and Technology* 5.2 (Apr. 2020), p. 024003. DOI: [10.1088/2058-9565/ab7559](https://doi.org/10.1088/2058-9565/ab7559).
- [3] Frank Arute et al. “Quantum supremacy using a programmable superconducting processor”. In: *Nature* 574.7779 (2019), pp. 505–510. ISSN: 1476-4687. DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5).
- [4] Simon J. Devitt. “Performing quantum computing experiments in the cloud”. In: *Phys. Rev. A* 94 (3 2016), p. 032329. DOI: [10.1103/PhysRevA.94.032329](https://doi.org/10.1103/PhysRevA.94.032329).
- [5] S. Debnath et al. “Demonstration of a small programmable quantum computer with atomic qubits”. In: *Nature* 536.7614 (2016), pp. 63–66. ISSN: 1476-4687. DOI: [10.1038/nature18648](https://doi.org/10.1038/nature18648).
- [6] Prakash Murali et al. “Full-Stack, Real-System Quantum Computer Studies: Architectural Comparisons and Design Insights”. In: *Proceedings of the 46th International Symposium on Computer Architecture*. ISCA '19. Phoenix, Arizona: Association for Computing Machinery, 2019, 527–540. ISBN: 9781450366694. DOI: [10.1145/3307650.3322273](https://doi.org/10.1145/3307650.3322273).
- [7] A Petitet, R C Whaley, J Dongarra, and A Cleary. *HPL - A Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers*. <http://www.netlib.org/benchmark/hpl/>.
- [8] Jack J. Dongarra, Piotr Luszczek, and Antoine Petitet. “The LINPACK Benchmark: past, present and future”. In: *Concurrency and Computation: Practice and Experience* 15.9 (2003), pp. 803–820. DOI: [10.1002/cpe.728](https://doi.org/10.1002/cpe.728).
- [9] Jack Dongarra and Piotr Luszczek. “TOP500”. In: *Encyclopedia of Parallel Computing*. Ed. by David Padua. Boston, MA: Springer US, 2011, pp. 2055–2057. ISBN: 978-0-387-09766-4. DOI: [10.1007/978-0-387-09766-4_157](https://doi.org/10.1007/978-0-387-09766-4_157).
- [10] Norbert M. Linke et al. “Experimental comparison of two quantum computing architectures”. In: *Proceedings of the National Academy of Sciences* 114.13 (2017), pp. 3305–3310. ISSN: 0027-8424. DOI: [10.1073/pnas.1618020114](https://doi.org/10.1073/pnas.1618020114).
- [11] Robin Blume-Kohout and Kevin C. Young. *A volumetric framework for quantum computer benchmarks*. Nov. 2020. DOI: [10.22331/q-2020-11-15-362](https://doi.org/10.22331/q-2020-11-15-362).
- [12] Sam McArdle et al. “Quantum computational chemistry”. In: *Rev. Mod. Phys.* 92 (1 Mar. 2020), p. 015003. DOI: [10.1103/RevModPhys.92.015003](https://doi.org/10.1103/RevModPhys.92.015003).
- [13] Alexander J. McCaskey et al. “Quantum chemistry as a benchmark for near-term quantum computers”. In: *npj Quantum Information* 5.1 (2019), p. 99. ISSN: 2056-6387. DOI: [10.1038/s41534-019-0209-0](https://doi.org/10.1038/s41534-019-0209-0).
- [14] Pierre-Luc Dallaire-Demers et al. “An application benchmark for fermionic quantum simulations”. In: (2020). arXiv: [2003 . 01862 \[quant-ph\]](https://arxiv.org/abs/2003.01862).
- [15] E. F. Dumitrescu et al. “Cloud Quantum Computing of an Atomic Nucleus”. In: *Phys. Rev. Lett.* 120 (21 May 2018), p. 210501. DOI: [10.1103/PhysRevLett.120.210501](https://doi.org/10.1103/PhysRevLett.120.210501).

- [16] Frank Arute et al. *Hartree-Fock on a superconducting qubit quantum computer*. 2020. DOI: [10.1126/science.abb9811](https://doi.org/10.1126/science.abb9811).
- [17] Vedran Dunjko and Hans J Briegel. “Machine learning & artificial intelligence in the quantum domain: a review of recent progress”. In: *Reports on Progress in Physics* 81.7 (June 2018), p. 074001. DOI: [10.1088/1361-6633/aab406](https://doi.org/10.1088/1361-6633/aab406).
- [18] Marcello Benedetti et al. “A generative modeling approach for benchmarking and training shallow quantum circuits”. In: *npj Quantum Information* 5.1 (2019), p. 45. ISSN: 2056-6387. DOI: [10.1038/s41534-019-0157-8](https://doi.org/10.1038/s41534-019-0157-8).
- [19] Kathleen E. Hamilton and Raphael C. Pooser. “Error-mitigated data-driven circuit learning on noisy quantum hardware”. In: *Quantum Machine Intelligence* 2.1 (2020), p. 10. ISSN: 2524-4914. DOI: [10.1007/s42484-020-00021-x](https://doi.org/10.1007/s42484-020-00021-x).
- [20] Kathleen E. Hamilton, Eugene F. Dumitrescu, and Raphael C. Pooser. “Generative model benchmarks for superconducting qubits”. In: *Phys. Rev. A* 99 (6 June 2019), p. 062323. DOI: [10.1103/PhysRevA.99.062323](https://doi.org/10.1103/PhysRevA.99.062323).
- [21] Matthew P. Harrigan et al. *Quantum approximate optimization of non-planar graph problems on a planar superconducting processor*. 2021. DOI: [10.1038/s41567-020-01105-y](https://doi.org/10.1038/s41567-020-01105-y).
- [22] Madita Willsch et al. *Benchmarking the quantum approximate optimization algorithm*. 2020. DOI: [10.1007/s11128-020-02692-8](https://doi.org/10.1007/s11128-020-02692-8).
- [23] Andreas Bengtsson et al. *Improved Success Probability with Greater Circuit Depth for the Quantum Approximate Optimization Algorithm*. 2020. DOI: [10.1103/PhysRevApplied.14.034010](https://doi.org/10.1103/PhysRevApplied.14.034010).
- [24] Guido Pagano et al. *Quantum approximate optimization of the long-range Ising model with a trapped-ion quantum simulator*. 2020. DOI: [10.1073/pnas.2006373117](https://doi.org/10.1073/pnas.2006373117).
- [25] John Preskill. *Quantum computing and the entanglement frontier*. 2012. arXiv: [1203.5813](https://arxiv.org/abs/1203.5813) [quant-ph].
- [26] Aram W. Harrow and Ashley Montanaro. “Quantum computational supremacy”. In: *Nature* 549.7671 (Sept. 2017), 203–209. ISSN: 1476-4687. DOI: [10.1038/nature23458](https://doi.org/10.1038/nature23458).
- [27] Sergio Boixo et al. “Characterizing quantum supremacy in near-term devices”. In: *Nature Physics* 14.6 (June 2018), pp. 595–600. ISSN: 1745-2481. DOI: [10.1038/s41567-018-0124-x](https://doi.org/10.1038/s41567-018-0124-x).
- [28] Sukin Sim, Peter D. Johnson, and Alán Aspuru-Guzik. “Expressibility and Entangling Capability of Parameterized Quantum Circuits for Hybrid Quantum-Classical Algorithms”. In: *Advanced Quantum Technologies* 2.12 (2019), p. 1900070. DOI: [10.1002/qute.201900070](https://doi.org/10.1002/qute.201900070).
- [29] Abhinav Kandala et al. “Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets”. In: *Nature* 549.7671 (Sept. 2017), pp. 242–246. ISSN: 1476-4687. DOI: [10.1038/nature23879](https://doi.org/10.1038/nature23879).
- [30] Brian Coyle, Daniel Mills, Vincent Danos, and Elham Kashefi. “The Born supremacy: quantum advantage and training of an Ising Born machine”. In: *npj Quantum Information* 6.1 (2020), p. 60. ISSN: 2056-6387. DOI: [10.1038/s41534-020-00288-9](https://doi.org/10.1038/s41534-020-00288-9).
- [31] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao. *Expressive power of parametrized quantum circuits*. 2020. DOI: [10.1103/PhysRevResearch.2.033125](https://doi.org/10.1103/PhysRevResearch.2.033125).
- [32] Thomas Hubregtsen, Josef Pichlmeier, and Koen Bertels. “Evaluation of Parameterized Quantum Circuits: on the design, and the relation between classification accuracy, expressibility and entangling capability”. In: (2020). arXiv: [2003.09887](https://arxiv.org/abs/2003.09887) [quant-ph].
- [33] Andrew W. Cross et al. “Validating quantum computers using randomized model circuits”. In: *Phys. Rev. A* 100 (3 Sept. 2019), p. 032328. DOI: [10.1103/PhysRevA.100.032328](https://doi.org/10.1103/PhysRevA.100.032328).
- [34] Dominic W. Berry, Graeme Ahokas, Richard Cleve, and Barry C. Sanders. “Efficient Quantum Algorithms for Simulating Sparse Hamiltonians”. In: *Communications in Mathematical Physics* 270.2 (2007), pp. 359–371. ISSN: 1432-0916. DOI: [10.1007/s00220-006-0150-x](https://doi.org/10.1007/s00220-006-0150-x).
- [35] Alberto Peruzzo et al. “A variational eigenvalue solver on a photonic quantum processor”. In: *Nature Communications* 5.1 (July 2014), p. 4213. ISSN: 2041-1723. DOI: [10.1038/ncomms5213](https://doi.org/10.1038/ncomms5213).
- [36] Jonathan Romero et al. “Strategies for quantum computing molecular energies using the unitary coupled cluster ansatz”. In: *Quantum Science and Technology* 4.1 (2018), p. 014008. DOI: [10.1088/2058-9565/aad3e4](https://doi.org/10.1088/2058-9565/aad3e4).
- [37] Alexandru Gheorghiu, Theodoros Kapourniotis, and Elham Kashefi. “Verification of Quantum Computation: An Overview of Existing Approaches”. In: *Theor. Comp. Sys.* 63.4 (May 2019), pp. 715–808. ISSN: 1432-4350. DOI: [10.1007/s00224-018-9872-3](https://doi.org/10.1007/s00224-018-9872-3).
- [38] U. Mahadev. “Classical Verification of Quantum Computations”. In: *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. 2018, pp. 259–267. DOI: [10.1109/FOCS.2018.00033](https://doi.org/10.1109/FOCS.2018.00033).
- [39] Scott Aaronson and Lijie Chen. “Complexity-Theoretic Foundations of Quantum Supremacy Experiments”. In: (2016). arXiv: [1612.05903](https://arxiv.org/abs/1612.05903) [quant-ph].

- [40] Nathan Wiebe, Christopher Granade, and D G Cory. “Quantum bootstrapping via compressed quantum Hamiltonian learning”. In: *New Journal of Physics* 17.2 (Feb. 2015), p. 022005. DOI: [10.1088/1367-2630/17/2/022005](https://doi.org/10.1088/1367-2630/17/2/022005).
- [41] C. E. Porter and R. G. Thomas. “Fluctuations of Nuclear Reaction Widths”. In: *Phys. Rev.* 104 (2 Oct. 1956), pp. 483–491. DOI: [10.1103/PhysRev.104.483](https://doi.org/10.1103/PhysRev.104.483).
- [42] C. Neill et al. “A blueprint for demonstrating quantum supremacy with superconducting qubits”. In: *Science* 360.6385 (2018), pp. 195–199. ISSN: 0036-8075. DOI: [10.1126/science.aao4309](https://doi.org/10.1126/science.aao4309).
- [43] Sergio Boixo, Vadim N. Smelyanskiy, and Hartmut Neven. “Fourier analysis of sampling from noisy chaotic quantum circuits”. In: (2017). arXiv: [1708.01875](https://arxiv.org/abs/1708.01875) [quant-ph].
- [44] Adam Bouland, Bill Fefferman, Chinmay Nirkhe, and Umesh Vazirani. “On the complexity and verification of quantum random circuit sampling”. In: *Nature Physics* 15.2 (2019), pp. 159–163. ISSN: 1745-2481. DOI: [10.1038/s41567-018-0318-2](https://doi.org/10.1038/s41567-018-0318-2).
- [45] Ramis Movassagh. “Efficient unitary paths and quantum computational supremacy: A proof of average-case hardness of Random Circuit Sampling”. In: (2018). arXiv: [1810.04681](https://arxiv.org/abs/1810.04681) [quant-ph].
- [46] Scott Aaronson and Alex Arkhipov. “The Computational Complexity of Linear Optics”. In: *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*. STOC ’11. San Jose, California, USA: Association for Computing Machinery, 2011, 333–342. ISBN: 9781450306911. DOI: [10.1145/1993636.1993682](https://doi.org/10.1145/1993636.1993682).
- [47] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. “Average-Case Complexity Versus Approximate Simulation of Commuting Quantum Computations”. In: *Phys. Rev. Lett.* 117 (8 Aug. 2016), p. 080501. DOI: [10.1103/PhysRevLett.117.080501](https://doi.org/10.1103/PhysRevLett.117.080501).
- [48] Michael J. Bremner, Richard Jozsa, and Dan J. Shepherd. “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 467.2126 (2011), pp. 459–472. DOI: [10.1098/rspa.2010.0301](https://doi.org/10.1098/rspa.2010.0301).
- [49] Benjamin Villalonga et al. “A flexible high-performance simulator for verifying and benchmarking quantum circuits implemented on real hardware”. In: *npj Quantum Information* 5.1 (2019), p. 86. ISSN: 2056-6387. DOI: [10.1038/s41534-019-0196-1](https://doi.org/10.1038/s41534-019-0196-1).
- [50] Benjamin Villalonga et al. “Establishing the quantum supremacy frontier with a 281 Pflop/s simulation”. In: *Quantum Science and Technology* 5.3 (2020), p. 034003. DOI: [10.1088/2058-9565/ab7eeb](https://doi.org/10.1088/2058-9565/ab7eeb).
- [51] Dan Shepherd and Michael J. Bremner. “Temporally unstructured quantum computation”. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 465.2105 (2009), pp. 1413–1439. DOI: [10.1098/rspa.2008.0443](https://doi.org/10.1098/rspa.2008.0443).
- [52] Michael J. Bremner, Ashley Montanaro, and Dan J. Shepherd. “Achieving quantum supremacy with sparse and noisy commuting quantum computations”. In: *Quantum* 1 (Apr. 2017), p. 8. ISSN: 2521-327X. DOI: [10.22331/q-2017-04-25-8](https://doi.org/10.22331/q-2017-04-25-8).
- [53] Vojtěch Havlíček et al. “Supervised learning with quantum-enhanced feature spaces”. In: *Nature* 567.7747 (2019), pp. 209–212. ISSN: 1476-4687. DOI: [10.1038/s41586-019-0980-2](https://doi.org/10.1038/s41586-019-0980-2).
- [54] Daniel Mills, Anna Pappa, Theodoros Kapourniotis, and Elham Kashefi. “Information Theoretically Secure Hypothesis Test for Temporally Unstructured Quantum Computation (Extended Abstract)”. In: *Electronic Proceedings in Theoretical Computer Science* 266 (Feb. 2018), 209–221. ISSN: 2075-2180. DOI: [10.4204/eptcs.266.14](https://doi.org/10.4204/eptcs.266.14).
- [55] Richard M. Karp and Richard J. Lipton. “Some Connections between Nonuniform and Uniform Complexity Classes”. In: *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing*. STOC ’80. Los Angeles, California, USA: Association for Computing Machinery, 1980, 302–309. ISBN: 0897910176. DOI: [10.1145/800141.804678](https://doi.org/10.1145/800141.804678).
- [56] J. Misra and David Gries. “A constructive proof of Vizing’s theorem”. In: *Information Processing Letters* 41.3 (1992), pp. 131–133. ISSN: 0020-0190. DOI: [10.1016/0020-0190\(92\)90041-S](https://doi.org/10.1016/0020-0190(92)90041-S).
- [57] Juan Bermejo-Vega et al. “Architectures for Quantum Simulation Showing a Quantum Speedup”. In: *Phys. Rev. X* 8 (2 Apr. 2018), p. 021010. DOI: [10.1103/PhysRevX.8.021010](https://doi.org/10.1103/PhysRevX.8.021010).
- [58] D Hangleiter, M Kliesch, M Schwarz, and J Eisert. “Direct certification of a class of quantum simulations”. In: *Quantum Science and Technology* 2.1 (Feb. 2017), p. 015004. DOI: [10.1088/2058-9565/2/1/015004](https://doi.org/10.1088/2058-9565/2/1/015004).
- [59] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. 2010. DOI: [10.1017/CBO9780511976667](https://doi.org/10.1017/CBO9780511976667).

- [60] Alexander Cowtan et al. “Phase Gadget Synthesis for Shallow Circuits”. In: *Electronic Proceedings in Theoretical Computer Science* 318 (Apr. 2020), 214–229. ISSN: 2075-2180. DOI: [10.4204/eptcs.318.13](https://doi.org/10.4204/eptcs.318.13).
- [61] Panagiotis Kl. Barkoutsos et al. “Quantum algorithms for electronic structure calculations: Particle-hole Hamiltonian and optimized wavefunction expansions”. In: *Phys. Rev. A* 98 (2 Aug. 2018), p. 022322. DOI: [10.1103/PhysRevA.98.022322](https://doi.org/10.1103/PhysRevA.98.022322).
- [62] Seyon Sivarajah et al. “t|ket): A retargetable compiler for NISQ devices”. In: *Quantum Science and Technology* (2020). DOI: [10.1088/2058-9565/ab8e92](https://doi.org/10.1088/2058-9565/ab8e92).
- [63] *pytket documentation*. <https://cqcl.github.io/pytket/build/html/index.html>.
- [64] *qiskit documentation*. <https://qiskit.org/documentation/>.
- [65] Sumeet Khatri et al. “Quantum-assisted quantum compiling”. In: *Quantum* 3 (May 2019), p. 140. ISSN: 2521-327X. DOI: [10.22331/q-2019-05-13-140](https://doi.org/10.22331/q-2019-05-13-140).
- [66] Aram W. Harrow, Benjamin Recht, and Isaac L. Chuang. “Efficient discrete approximations of quantum gates”. In: *Journal of Mathematical Physics* 43.9 (2002), pp. 4445–4451. DOI: [10.1063/1.1495899](https://doi.org/10.1063/1.1495899).
- [67] Joel J. Wallman and Joseph Emerson. “Noise tailoring for scalable quantum computation via randomized compiling”. In: *Phys. Rev. A* 94 (5 Nov. 2016), p. 052325. DOI: [10.1103/PhysRevA.94.052325](https://doi.org/10.1103/PhysRevA.94.052325).
- [68] Jeff Heckey et al. “Compiler Management of Communication and Parallelism for Quantum Computation”. In: *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’15. Istanbul, Turkey: Association for Computing Machinery, 2015, 445–456. ISBN: 9781450328357. DOI: [10.1145/2694344.2694357](https://doi.org/10.1145/2694344.2694357).
- [69] Prakash Murali et al. “Noise-Adaptive Compiler Mappings for Noisy Intermediate-Scale Quantum Computers”. In: *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’19. Providence, RI, USA: Association for Computing Machinery, 2019, 1015–1029. ISBN: 9781450362405. DOI: [10.1145/3297858.3304075](https://doi.org/10.1145/3297858.3304075).
- [70] Ali JavadiAbhari et al. *ScaffCC: Scalable compilation and analysis of quantum programs*. Tech. rep. Computing Frontiers 2014: Best Papers. 2015, pp. 2–17. DOI: [10.1016/j.parco.2014.12.001](https://doi.org/10.1016/j.parco.2014.12.001).
- [71] Alexander J McCaskey et al. “XACC: a system-level software infrastructure for heterogeneous quantum–classical computing”. In: *Quantum Science and Technology* 5.2 (Feb. 2020), p. 024002. DOI: [10.1088/2058-9565/ab6bf6](https://doi.org/10.1088/2058-9565/ab6bf6).
- [72] Alexander Cowtan et al. “On the Qubit Routing Problem”. In: *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019)*. Ed. by Wim van Dam and Laura Mancinska. Vol. 135. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019, 5:1–5:32. ISBN: 978-3-95977-112-2. DOI: [10.4230/LIPIcs.TQC.2019.5](https://doi.org/10.4230/LIPIcs.TQC.2019.5).
- [73] Prakash Murali, David C. McKay, Margaret Martonosi, and Ali Javadi-Abhari. “Software Mitigation of Crosstalk on Noisy Intermediate-Scale Quantum Computers”. In: *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*. ASPLOS ’20. Lausanne, Switzerland: Association for Computing Machinery, 2020, 1001–1016. ISBN: 9781450371025. DOI: [10.1145/3373376.3378477](https://doi.org/10.1145/3373376.3378477).
- [74] IBM Quantum Experience User Guide. *Advanced Single-Qubit Gates*. <https://quantum-computing.ibm.com/support/guides/user-guide>. 2019.
- [75] Christopher Chamberland et al. “Topological and Subsystem Codes on Low-Degree Graphs with Flag Qubits”. In: *Phys. Rev. X* 10 (1 Jan. 2020), p. 011022. DOI: [10.1103/PhysRevX.10.011022](https://doi.org/10.1103/PhysRevX.10.011022).
- [76] Iskren Vankov, Daniel Mills, Petros Wallden, and Elham Kashefi. “Methods for classically simulating noisy networked quantum architectures”. In: *Quantum Science and Technology* 5.1 (Nov. 2019), p. 014001. DOI: [10.1088/2058-9565/ab54a4](https://doi.org/10.1088/2058-9565/ab54a4).
- [77] J. M. Pino et al. *Demonstration of the QCCD trapped-ion quantum computer architecture*. 2020. arXiv: [2003.01293](https://arxiv.org/abs/2003.01293) [quant-ph].
- [78] Daniel Mills, Seyon Sivarajah, Travis L. Scholten, and Ross Duncan. *Application-Motivated, Holistic Benchmarking of a Full Quantum Computing Stack: Experimental Data*. Version 1.0. May 2020. DOI: [10.5281/zenodo.3832121](https://doi.org/10.5281/zenodo.3832121).
- [79] Joseph Emerson et al. “Pseudo-Random Unitary Operators for Quantum Information Processing”. In: *Science* 302.5653 (2003), pp. 2098–2100. ISSN: 0036-8075. DOI: [10.1126/science.1090790](https://doi.org/10.1126/science.1090790).

- [80] Joseph Emerson, Etera Livine, and Seth Lloyd. “Convergence conditions for random quantum circuits”. In: *Phys. Rev. A* 72 (6 Dec. 2005), p. 060302. DOI: [10.1103/PhysRevA.72.060302](https://doi.org/10.1103/PhysRevA.72.060302).
- [81] Fernando G. S. L. Brandão, Aram W. Harrow, and Michał Horodecki. “Local Random Quantum Circuits are Approximate Polynomial-Designs”. In: *Communications in Mathematical Physics* 346.2 (Sept. 2016), pp. 397–434. ISSN: 1432-0916. DOI: [10.1007/s00220-016-2706-8](https://doi.org/10.1007/s00220-016-2706-8).
- [82] Andrew Fagan and Ross Duncan. “Optimising Clifford Circuits with Quantomatic”. In: *Electronic Proceedings in Theoretical Computer Science* 287 (Jan. 2019), 85–105. ISSN: 2075-2180. DOI: [10.4204/eptcs.287.5](https://doi.org/10.4204/eptcs.287.5).
- [83] M Blaauboer and R L de Visser. “An analytical decomposition protocol for optimal implementation of two-qubit entangling gates”. In: *Journal of Physics A: Mathematical and Theoretical* 41.39 (2008), p. 395307. DOI: [10.1088/1751-8113/41/39/395307](https://doi.org/10.1088/1751-8113/41/39/395307).
- [84] Easwar Magesan and Jay M. Gambetta. “Effective Hamiltonian models of the cross-resonance gate”. In: *Phys. Rev. A* 101 (5 2020), p. 052308. DOI: [10.1103/PhysRevA.101.052308](https://doi.org/10.1103/PhysRevA.101.052308).
- [85] Timothy Proctor et al. “What Randomized Benchmarking Actually Measures”. In: *Phys. Rev. Lett.* 119 (13 Sept. 2017), p. 130502. DOI: [10.1103/PhysRevLett.119.130502](https://doi.org/10.1103/PhysRevLett.119.130502).
- [86] E. Knill et al. “Randomized benchmarking of quantum gates”. In: *Phys. Rev. A* 77 (1 Jan. 2008), p. 012307. DOI: [10.1103/PhysRevA.77.012307](https://doi.org/10.1103/PhysRevA.77.012307).
- [87] Arnaud Carignan-Dugas, Kristine Boone, Joel J Wallman, and Joseph Emerson. “From randomized benchmarking experiments to gate-set circuit fidelity: how to interpret randomized benchmarking decay parameters”. In: *New Journal of Physics* 20.9 (Sept. 2018), p. 092001. DOI: [10.1088/1367-2630/aadcc7](https://doi.org/10.1088/1367-2630/aadcc7).
- [88] Timothy J. Proctor et al. “Direct Randomized Benchmarking for Multiqubit Devices”. In: *Phys. Rev. Lett.* 123 (3 July 2019), p. 030503. DOI: [10.1103/PhysRevLett.123.030503](https://doi.org/10.1103/PhysRevLett.123.030503).
- [89] Jay M. Gambetta et al. “Characterization of Addressability by Simultaneous Randomized Benchmarking”. In: *Phys. Rev. Lett.* 109 (24 Dec. 2012), p. 240504. DOI: [10.1103/PhysRevLett.109.240504](https://doi.org/10.1103/PhysRevLett.109.240504).
- [90] David C. McKay, Andrew W. Cross, Christopher J. Wood, and Jay M. Gambetta. *Correlated Randomized Benchmarking*. 2020. arXiv: [2003.02354](https://arxiv.org/abs/2003.02354) [quant-ph].
- [91] Easwar Magesan, J. M. Gambetta, and Joseph Emerson. “Scalable and Robust Randomized Benchmarking of Quantum Processes”. In: *Phys. Rev. Lett.* 106 (18 May 2011), p. 180504. DOI: [10.1103/PhysRevLett.106.180504](https://doi.org/10.1103/PhysRevLett.106.180504).
- [92] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. “Characterizing quantum gates via randomized benchmarking”. In: *Phys. Rev. A* 85 (4 Apr. 2012), p. 042311. DOI: [10.1103/PhysRevA.85.042311](https://doi.org/10.1103/PhysRevA.85.042311).
- [93] A. D. Córcoles et al. “Process verification of two-qubit quantum gates by randomized benchmarking”. In: *Phys. Rev. A* 87 (3 Mar. 2013), p. 030301. DOI: [10.1103/PhysRevA.87.030301](https://doi.org/10.1103/PhysRevA.87.030301).
- [94] David C. McKay et al. “Three-Qubit Randomized Benchmarking”. In: *Phys. Rev. Lett.* 122 (20 May 2019), p. 200502. DOI: [10.1103/PhysRevLett.122.200502](https://doi.org/10.1103/PhysRevLett.122.200502).

A Exponential Distribution

The *exponential distribution*, with rate λ , is a probability distribution with the probability density function

$$\Pr(x) = \lambda e^{-\lambda x}.$$

This is the distribution of waiting times between events in a Poisson process. We are concerned with showing that output probabilities of the circuit classes considered here are exponentially distributed. Such a property is a signature of quantum chaos, and that a class of circuits is approximately Haar random [27, 79, 80]. It also allows for a simplified calculation of both the ideal value of the cross-entropy discussed in Section 2.2, and the ideal heavy output probability as discussed in Section 2.1, as discussed in those sections. This in turn allows us to fully exploit Cross-Entropy Benchmarking and Heavy Output Generation Benchmarking. Here we will argue numerically which of the circuits we introduce in Section 3 generate output probabilities of this form,²³ and discuss the implications when they do not.

We also demonstrate why the circuit depths used in Section 3 are necessary to generate output probabilities of this form. To do this we generate 100 circuits of each type and number of layers, where a layer is as defined in the respective Algorithms of Section 3. We then calculate the ideal output probabilities using classical simulation and compare this distribution of output probabilities to the exponential distribution. In the case of square circuits and deep circuits, we notice a better approximation of the exponential distribution by the distribution of output probabilities,

²³This numerical approach to demonstrating properties of distributions of output probabilities from particular circuit classes parallels that taken in other work on benchmarking [27, 39, 46, 57].

measured by the ℓ_1 -norm distance between the two, as the number of layers increases. We can use this to isolate the number of layers at which the difference approaches its minimum.

A.1 Square Circuits

The exponential form of the distribution of the output probabilities from random circuits similar to square circuits has been established [27, 39]. As the procedure we use to generate square circuits, seen in Algorithm 2, differs slightly from that used for other similar random circuits [27, 33, 39], we explore the distribution of its output probabilities here.

The relevant results are seen in Figure 12. In particular, it can be seen from Figure 12b that the minimum value of ℓ_1 -norm distance between the distribution of output probabilities and the exponential distribution is approached at a number of layers equal to the number of qubits, justifying our choice of layer numbers in Algorithm 2. It may be that asymptotically the number of layers required is sub-linear [27], although for the circuit sizes used here a linear growth in depth is appropriate. Figure 12a illustrates the closeness of fit of the two distributions.

A.2 Deep Circuits

Unlike with square circuits, there is no precedent for utilising deep circuits to generate exponentially distributed output probabilities, as we do here. This allows us to use deep circuits as a uniquely insightful benchmark of the performance of quantum computing stacks, grounded both in the theoretical results of Section 2, and in pertinent applications.

The relevant results are seen in Figure 13. In particular, it can be seen from Figure 13b that the minimum value of ℓ_1 -norm distance between the distribution of output probabilities and the exponential distribution is approached at a number of layers equal to three times the number of qubits, plus one, justifying our choice of layer numbers in Algorithm 3. Figure 13a illustrates the closeness of fit of the two distributions.

The depth required to achieve an exponential distribution of outcome probabilities with deep circuits is greater than is the case for square circuits. Indeed, random circuits were initially introduced as the shallowest circuits required to generate such output probabilities [27]. This sacrifice in depth is made to achieve a benchmark which is uniquely application motivated, as discussed in Section 3.

A.3 Shallow Circuits

Unlike in the case of square circuits and deep circuits, the output probabilities of shallow circuits are not exponentially distributed. This is unsurprising since random circuits with this limited connectivity are thought to require at least depth $\mathcal{O}(\sqrt{n})$ to create

(a) The distribution of output probabilities from a circuit C , where C is a 5 qubit circuit, from the square circuits class as defined in Algorithm 2.

(b) The ℓ_1 -norm distance between the distribution of output probabilities of square circuits and the exponential distribution $2^n e^{-2^n x}$, where n is the number of qubits. A layer is defined as in Algorithm 2. Colours correspond to numbers of qubits in the following way: 2 [yellow], 3 [blue], 4 [green], 5 [magenta].

Figure 12: Exponential distribution fitting data for square circuits.

(a) The distribution of output probabilities from a circuit C , where C is a 5 qubit circuit, from the deep circuits class as defined in Algorithm 3.

(b) The ℓ_1 -norm distance between the distribution of output probabilities of deep circuits and the exponential distribution $2^n e^{-2^n x}$, where n is the number of qubits. A layer is defined as in Algorithm 3. Colours correspond to numbers of qubits in the following way: 2 [yellow], 3 [blue], 4 [green], 5 [red].

Figure 13: **Exponential distribution fitting data for deep circuits.**

such a feature [44, 57, 81]. This has the unfortunate side effect that the results of Section 2.2 do not apply, and so the simplifications to the calculations performed when conducting Cross-Entropy Benchmarking cannot be used.

While it is also true that the predictions made about the ideal heavy output probability, as discussed in Section 2.1, also do not apply, a study of the heavy output probability is still of interest. In particular, while we cannot connect the benchmark to the HOG problem of Problem 1, we can compare the probability of generating heavy outputs to the ideal probability of producing heavy outputs, as calculated by classical simulation.

B Compilation Strategies

This section details the compilation strategies explored in each of our experiments. For the circuit families and figures of merit investigated here, the compilation strategies we used were designed and empirically confirmed to perform well at the compilation tasks at hand. The version of each package used are listed in Table 2.

noise-unaware pytket and noise-aware pytket The noise-unaware pytket and noise-aware pytket compilation strategies are generated using Algorithm 4. noise-unaware pytket is generated by passing `False` as input to Algorithm 4, and noise-aware pytket by passing `True`.

Of particular interest are the following functions:

OptimiseCliffors: Simplifies Clifford gate sequences [82].

KAKDecomposition: Identifies two-qubit sub-circuits with more than 3 CXs and reduces them via the KAK/Cartan decomposition [83].

route: Modifies the circuit to satisfy the architectural constraints [72]. This will introduce SWAP gates.

noise_aware_placement: Selects initial qubit placement taking into account reported device gate error rates [62].

line_placement: Attempts to place qubits next to those they interact with in the first few time slices. This does not take device error rates into account.

Package	Version
Qiskit [1, 64]	0.12.0
pytket [62, 63]	0.3.0

Table 2: **Packages used in this work, and their corresponding versions.**

Algorithm 4 pytket compilation strategies. The passes listed here are named as in the documentation for pytket [63], where additional detail on their actions can be found.

Input: `noise_aware` $\in \{\text{True}, \text{False}\}$

```

1: OptimiseCliffords
2: KAKDecomposition
3:
4: RebaseToRzRx          ▷ Convert to IBM gate set
5: CommuteRzRxThroughCX
6:
7: if noise_aware then
8:   noise_aware_placement
9: else
10:  line_placement
11: end if
12:
13: route
14: decompose_SWAP_to_CX
15: redirect_CX_gates ▷ Orientate CX to coupling
    map
16:
17: OptimisePostRouting    ▷ Optimisation
    preserving placement and orientation

```

noise-unaware Qiskit and noise-aware Qiskit The noise-unaware Qiskit and noise-aware Qiskit compilation strategies, as defined in Algorithm 5, are heavily inspired by `level_3_passmanager`, a preconfigured compilation strategy made available in Qiskit. `noise-unaware Qiskit` is generated by passing `noise_aware` as `False` in Algorithm 4, and `noise-aware Qiskit` by passing `True`.

Where possible we passed `stochastic` as `True` in order to use `StochasticSwap` instead of `BasicSwap` during the swap mapping pass. In general, `StochasticSwap` generates circuits with lower depth; however, for the versions listed in Table 2, it proved faulty for some circuit sizes and device coupling maps used in this work. `StochasticSwap` may also result in repeated measurement of the same qubit, which cannot be implemented. Repeated compilation attempts may therefore be necessary, and if this fails the circuit is not included in the plots of Section 5.

Of particular note are the following functions:

NoiseAdaptiveLayout: Selects initial qubit placement based on minimising readout error rates [69].

DenseLayout: Chooses placement by finding the most connected subset of qubits.

Unroller: Decomposes unitary operation to desired gate set.

StochasticSwap: Adds SWAP gates to adhere to coupling map using a randomised algorithm.

BasicSwap: Produces a circuit adhering to coupling map using a simple rule: CX gates in the circuit which are not supported by the hardware are preceded with necessary SWAP gates.

Algorithm 5 Qiskit compilation strategies. The passes listed here are named as in the documentation for Qiskit [64], where additional detail on their actions can be found.

Input:
`noise_aware` $\in \{\text{True}, \text{False}\}$
`stochastic` $\in \{\text{True}, \text{False}\}$

```

1: Unroller
2:
3: if noise_aware then
4:   NoiseAdaptiveLayout
5: else
6:   DenseLayout
7: end if
8: AncillaAllocation    ▷ Assign idle qubits as
    ancillas
9:
10: if stochastic then
11:   StochasticSwap
12: else
13:   BasicSwap
14: end if
15:
16: Decompose(SwapGate)  ▷ Decompose SWAP to
    CX
17: CXDirection          ▷ Orientate CX to coupling map
18:
19:                      ▷ Gather 2 qubit blocks
20: Collect2qBlocks
21: ConsolidateBlocks
22:
23: Unroller              ▷ Unroll two-qubit blocks
24: Optimize1qGates ▷ Combine chains of one-qubit
    gates
25: CXDirection

```

only pytket routing In this case we perform, in the order as listed, the pytket operations: `route`, `decompose_SWAP_to_CX`, and `redirect_CX_gates`. We then account for the architecture gate set, without any further optimisation.

C Device Data

Two device properties leveraged by our compilation strategies are the *coupling maps*, describing the connectivity of the qubits and in which directions CX gates can be performed, and the *calibration information*, describing the noise levels of the device. These properties, and devices noise levels in particular, are considered valuable benchmarks of the performance of the device in their own right.

These properties are collectively influential in noise-aware compiling, as detailed in Appendix B. There circuits are compiled to adhere to the device’s coupling map, while also aiming to minimise some function of the calibration information. Because full quantum computing stack holistic benchmarking encompasses the circuit compilation strategies, it provides a novel way of using device information to benchmark an entire system, instead of simply the physical qubits which comprise it.

C.1 Device Coupling Maps

A *coupling map* of a device is a graphical representation of how two-qubit gates can be applied across the device. In this representation, each qubit is represented by a vertex, with directed edges joining qubits between which a two-qubit gate can be applied. For the devices considered here, this two-qubit gate is a CX gate, implemented using the cross-resonance interaction of transmon qubits [84]. The direction of the edge is from the control to the target qubit of the CX gate, with bi-directional edges indicating that both qubits can be used as either the control or target. The coupling maps of the devices investigated in this work are shown in Figure 14. For those devices all edges are bi-directional, although this is not typical when the asymmetric CX is employed.

As discussed in Section 4, a trade-off exists between the connectivity of the device and the number of two-qubit gates necessary to implement a given circuit. More highly connected coupling maps typically require fewer two-qubit gates to implement a fixed unitary than less connected ones, owing to the reduced need for SWAP gates to account for discrepancies between the coupling maps of the uncompiled circuit and the device. While this reduced depth can reduce the impact of time based noise channels, this is counterbalanced by the higher levels of cross-talk experienced by qubits corresponding to vertices with high degree in the device’s coupling map [75].

C.2 Device Calibration Information

The noise-aware tools employed by the compilation strategies explored in this work consider three kinds of errors which can occur, namely: readout error, single-qubit gate error, and two-qubit gate error. For the devices provided through IBM Quantum, this information is contained in calibration data which is accessible using tools in the Qiskit library, and is updated twice daily. The experiments in this paper were conducted between 2020-01-29 and 2020-02-10 with the calibration data in Figure 15 and Figure 16 aggregated over this time period.

An assignment or readout error corresponds to an incorrect reading of the state of the qubit; for example, returning “0” when the proper label is “1”, or vice-versa. The probability of incorrectly labelling the qubit is called the *readout error*, denoted, ϵ^a , and is calculated as

$$\epsilon^a = \frac{\Pr(\text{“0”}||1\rangle) + \Pr(\text{“1”}||0\rangle)}{2}. \quad (12)$$

ϵ^a is estimated by repeatedly preparing a qubit in a known state, immediately measuring it, and then counting the number of times the measurement returns the wrong label. This value, for the devices explored in this paper, is reported in Figure 15a.

Errors affecting the gates of the device correspond to an incorrect operation applied by the device. There are many ways to quantify the effect of this error, with IBM Quantum’s devices reporting randomized benchmarking (RB) numbers [85, 86]. The RB number, ϵ^C , is estimated by running many self-inverting Clifford circuits, consisting of m layers of gates drawn from the n -qubit Clifford group, inverted at layer $m + 1$. The *survival probability*, which is the probability the input state is unchanged, can then be estimated. Under a broad set of noise models and assumptions [85, 87], this survival probability can be shown to decay exponentially with m . Consequently, it can be estimated by fitting a decay curve of the form $Ap^m + B$. The RB number is related to $p \in [0, 1]$, called the *depolarisation/decay rate*, by

$$\epsilon^C = (1 - p)(1 - 1/D), \quad (13)$$

where $D = 2^n$, and n is the number of qubits acted on by the Clifford gates. ϵ^C , which is also referred to as the *error per Clifford* of the device, is minimised at $p = 1$, in which case the survival probability is constant and set by the state preparation and measurement errors.

The Clifford gates necessary for RB must be compiled to the native gate set of the device. Using an estimate of ϵ^C , an estimate of the *error per gate*, ϵ_G^g , for a gate G , can be obtained by multiplying ϵ^C by a factor related to the average number of uses of G when implementing a random Clifford operation:

$$\epsilon_G^g \sim \epsilon^C \times \# \text{ uses of } G \text{ per Clifford}. \quad (14)$$

Figure 14: **Coupling maps of the devices studied in this work.** Vertices, represented by blue circles, correspond to qubits, while edges are directed from the control to the target qubits of permitted two-qubit gates.

Values for $\epsilon_{U_2}^g$, the error per gate for U_2 gates, can be found in Figure 15b, and ϵ_{CX}^g , that for CX gates, in Figure 16. The commonly reported average fidelity for U_3 gates is $1 - (1 - \epsilon_{U_2}^g)^2$.

There are many variants of randomized benchmarking, such as direct RB [88], simultaneous RB [89], and correlated RB [90]. For details on the randomized benchmarking protocol used by IBM Quantum, see [89, 91–94].

The experiments necessary for cross-entropy benchmarking may themselves also be used to estimate a depolarisation rate in a similar way to RB [3]. Instead of using random Clifford circuits, however, the random circuits are run. Under the assumption that the action of a random circuit can be described using a depolarising error model (with equal-probability Pauli errors), then the Pauli error, ϵ^P , can be estimated as

$$\epsilon^P = (1 - p)(1 - 1/D^2). \quad (15)$$

Here, p is the depolarisation rate of the survival probability under the action of random circuits, estimated as above. Interestingly, ϵ^P can be estimated using single and two-qubit RB information.

Several important noise channels, most notably cross-talk, are not included in the device calibration data. As shown in Section 5, the effects of this noise can be inferred through the application-motivated benchmarks we introduce in this work, by showing the trade-off between connectivity of the device and cross-talk [75].

D Empirical Relationship Between Heavy Output Generation Probability and ℓ_1 -Norm Distance

As discussed in Section 2.4, the theoretical foundations for believing that implementing shallow circuits to within a fixed ℓ_1 -norm distance constitutes a demonstration of quantum computational supremacy are stronger than for implementations with high heavy output generation probability. That said, Figure 3 and Figure 11 contain similar features. For example, `ibmq_16_melbourne` consistently performs the worst, with `ibmq_singapore` and `ibmq_ourense` performing the best by both figures of merit. How then

do these two figures of merit generally relate to one another.

If the ℓ_1 -norm distance was 0, the experimental outcome frequencies would equal the ideal outcome probabilities. Consequently, the heavy output probabilities would be the same between the device and an ideal quantum computer. Because the heavy output probability depends on the circuit in question, it is useful to normalise the latter by the heavy output probability of an ideally-implemented circuit. We define the *normalised heavy output generation probability* as the ratio of the heavy output probability of the device and the heavy output probability from an ideal quantum computer. Hence if the ℓ_1 -norm distance was 0, the normalised heavy output generation probability would be 1.

As the ℓ_1 -norm distance increases, the experimental frequencies increasingly differ from the ideal outcome probabilities. Two things then may happen: heavy outputs are produced more regularly, in which case the normalised heavy output generation probability will grow above 1; or less regularly, in which case the normalised heavy output generation probability will fall below 1. In practice, we expect the distribution produced by the device to converge to the uniform one as the noise increases, so we expect the normalised heavy output generation probability to fall with increasing ℓ_1 -norm distance.

The empirical relationship between the normalised heavy output generation probability and ℓ_1 -norm distance is shown in Figure 17. For each circuit, Figure 17 plots the ℓ_1 -norm distance of the distribution produced by a real device against the normalised heavy output generation probability. As expected, a negative correlation exists between these two figures of merit. For the deepest circuits, and in particular the widest circuits from the deep circuits class, the cluster of points indicates that the the normalised heavy output generation probability falls more slowly as the ℓ_1 -norm distance becomes larger. This is because the minimum value of heavy output generation probability is being reached, which is to say that the output distribution from the real device has converged to the uniform one, while more detail can be extracted by considering the ℓ_1 -norm distance.

This correlation is encouraging as, in the regime

(a) **Average readout error.** The readout error is the probability the state of a given qubit is incorrectly labelled.

(b) **Average error per U_2 gate.** The error per gate is a measure of how accurately the U_2 gate is applied.

Figure 15: **Error per single qubit operations on the devices used in this work.** Bars indicate the average error rates; error bars are one standard deviation. Data aggregated based on calibration data collected over the course of our experiments. Devices shown here are: ibmqx2 [■], ibmq_ourense [■], ibmq_singapore [■], ibmq_16_melbourne [■]. A logarithmic scale is used.

Figure 16: **Average error per CX operation on the devices used in this work.** The error per CX gate is a measure of how accurately the CX gate is applied. Bars indicate the average error rates; error bars are one standard deviation. Data aggregated based on calibration data collected over the course of our experiments. Devices shown here are: ibmqx2 [■], ibmq_ourense [■], ibmq_singapore [■], ibmq_16_melbourne [■]. A logarithmic scale is used.

Figure 17: **Scatter plot and linear regression line comparing the normalised heavy output generation probability and ℓ_1 -norm distance.** Each point corresponds to one circuit of the class and width as labelled. Colours correspond to numbers of qubits in the following way: 2 [■], 3 [■], 4 [■], 5 [■], 6 [■], 7 [■].

where it becomes impossible to calculate the ℓ_1 -norm distance, we can be justified in believing the correlation between the features present in Section 5 persist. This is to say that we may be able to make loose predictions about the performance of a quantum computing stack as measured by ℓ_1 -norm distance based on its performance as measured by Heavy Output Generation Benchmarking. This reasoning is similar to that used when Cross-Entropy Benchmarking is used to predict demonstrations of quantum computational supremacy in the regime when it becomes impossible to calculate [3].