

**Protocols for classically training quantum generative models on probability distributions**Sachin Kasture <sup>1</sup>, Oleksandr Kyriienko,<sup>2</sup> and Vincent E. Elfving<sup>1</sup><sup>1</sup>*PASQAL, 7 rue Léonard de Vinci, 91300 Massy, France*<sup>2</sup>*Department of Physics and Astronomy, University of Exeter, Stocker Road, Exeter EX4 4QL, United Kingdom*

(Received 7 July 2023; accepted 25 September 2023; published 5 October 2023)

Quantum generative modeling (QGM) relies on preparing quantum states and generating samples from these states as hidden—or known—probability distributions. As distributions from some classes of quantum states (circuits) are inherently hard to sample classically, QGM represents an excellent test bed for quantum supremacy experiments. Furthermore, generative tasks are increasingly relevant for industrial machine learning applications, and thus QGM is a strong candidate for demonstrating a practical quantum advantage. However, this requires that quantum circuits are trained to represent industrially relevant distributions, and the corresponding training stage has an extensive training cost for current quantum hardware in practice. In this work, we propose protocols for classical training of QGMs based on circuits of the specific type that admit an efficient gradient computation, while remaining hard to sample. In particular, we consider instantaneous quantum polynomial (IQP) circuits and their extensions. Showing their classical simulability in terms of the time complexity, sparsity, and anticoncentration properties, we develop a classically tractable way of simulating their output probability distributions, allowing classical training to a target probability distribution. The corresponding quantum sampling from IQPs can be performed efficiently, unlike when using classical sampling. We numerically demonstrate the end-to-end training of IQP circuits using probability distributions for up to 30 qubits on a regular desktop computer. When applied to industrially relevant distributions this combination of classical training with quantum sampling represents an avenue for reaching advantage in the noisy intermediate-scale quantum (NISQ) era.

DOI: [10.1103/PhysRevA.108.042406](https://doi.org/10.1103/PhysRevA.108.042406)**I. INTRODUCTION**

Recent breakthrough works in quantum computing demonstrated an improved scaling for sampling problems, leading to so-called quantum supremacy [1]. While an exact boundary of classical simulation is still to be established [2], for the carefully selected tasks of random circuit sampling and boson sampling [3,4] one achieves an exponential separation for the task of generating samples (bit strings from measurement readout). Generally, the task of generating samples from underlying probability distributions is the basis of generative modeling, and represents a highly important part of classical machine learning. Its quantum version—quantum generative modeling (QGM)—relies on training quantum circuits as adjustable probability distributions that can model sampling from some *particular* desired distribution [5–10]. QGM exploits the inherent superposition properties of a quantum state generated from a parametrized unitary, along with the probabilistic nature of quantum measurements, to efficiently sample from a trainable model. QGMs have potential applications in generating samples from solutions of stochastic differential equations (SDEs) for simulating financial and diffusive physical processes [11,12], scrambling data using a quantum embedding for anonymization [13–15], and generating solutions to graph-based problems like maximum independent set or maximum clique [16,17], among many others. Given the successes of quantum sampling this makes QGM a promising contender for achieving a quantum advantage. However, to date demonstrating QGM of practical significance has eluded the field as training

specific generative models is a complex time-intensive task.

The key asset of quantum generative modeling is that quantum measurements (collapsed to one of the eigenstates of a measurement operator) provide a new sample with each shot. Depending on the hardware platform, generating one sample can take on the order of a few hundreds of microseconds to several milliseconds [18–24]. For large probability distributions, represented by entangled 50+ qubit registers, performing the classical inversion is indeed much more costly. However, often the challenge comes from the inference side, when quantum circuits are required to match specific distributions. Typically, training of quantum generative models utilizes gradient-based parametric learning, similarly to training of deep neural networks [25]. Parametrized quantum circuits (also referred as quantum neural networks, QNNs) are trained by estimating the gradient of gate parameters  $\theta$ . Moreover, the gradient of a full QGM loss has to be estimated with respect to  $\theta$ . For quantum devices this can be done by the parameter-shift rule and its generalization [26–28], where number of circuits estimation increases linearly with the number of parameters. The overall training cost corresponds to the measurement of the loss function at each iteration step. In the case of quantum circuit Born machine the loss may correspond to Kullback-Leibler (KL) divergence [9], Sinkhorn divergence [29], or maximum mean discrepancy (MMD) [30], and may require extensive sampling for resolving the loss as an average. For quantum generative adversarial networks (QGAN) the loss minimization is substituted by the minimax

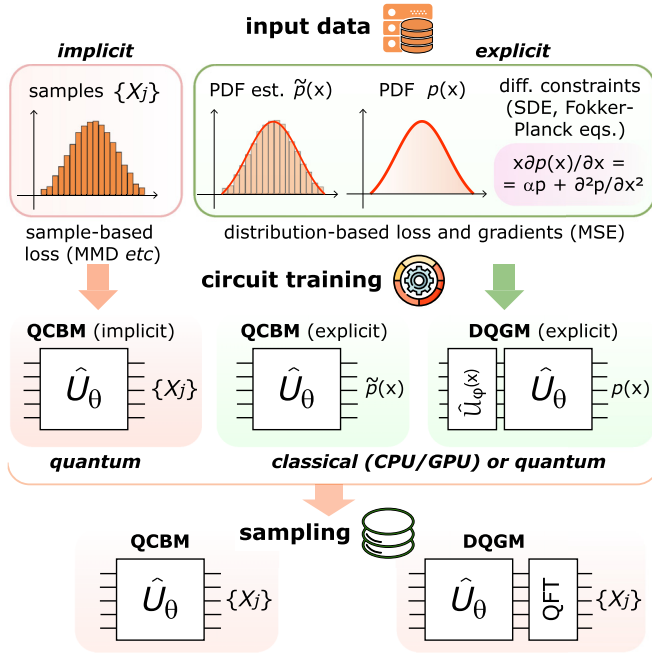


FIG. 1. Schematic of different steps in training a quantum generative model using QCBM and DQGM architectures. In a conventional QCBM setup, the loss and gradient estimation is done using input data samples directly, while in our work we focus on an explicit version of QCBM and DQGM, allowing for classical training.

game [31,32] requiring multicircuit estimation. In all cases the convergence is not guaranteed due to exponential reduction of gradients [33].

In this work we investigate the possibility of training the parameters of quantum generative models classically, while still retaining the quantum advantage in sampling [34]. For instance, the ability of classical training for a different paradigm was shown for Gaussian boson sampling (GBS) devices [35], but under certain conditions of fixing an initial set of samples and nonuniversal operation. For the digital quantum computing operation, results from previous works [36,37] motivate the possibility that estimating probability density classically can be feasible without losing the sampling complexity. We explore this possibility in more detail and use this further to develop methods to train circuits classically to output a desired distribution using a gradient-based approach. We show that our method is feasible using numerics for up to 30 qubits on a regular desktop computer. We explore different families of quantum circuits in detail and perform numerical studies to study their sampling complexity and expressivity. For expressivity studies in particular, we look at training a differentiable quantum generative model (DQGM) [38–40] architecture which allows training in the latent (or “frequency”) space, and sampling in the bit-basis. This presents a good testing ground for applying the proposed method to explicit quantum generative models. We also show that quantum circuit Born machines (QCBMs) can be trained classically for certain distributions, while still hard to sample. Our protocols contribute towards tools for achieving the practical advantage in sampling once the target distributions

are chosen carefully. We highlight the differences between well-known strategies for QGM and the method discussed in the paper in Fig. 1.

## II. METHODS

### A. Preliminaries: QCBM and DQGM as implicit vs explicit generative models

Generally, there are two types of generative models. *Explicit* generative models assume a direct access (or “inference”) to probability density functions (PDF). At the same time, *implicit* generative models are described by hidden parametric distributions, where samples are produced by transforming randomness via inversion procedure. These two types of models have crucial differences. For example, training explicit models involves loss functions measuring distances between a given PDF,  $p_{\text{target}}(x)$ , and a model PDF,  $p_{\text{model}}(x)$ , for example, with a mean square error (MSE) loss which is defined as

$$\mathcal{L}_{\text{MSE}} = \sum_x |p_{\text{model}}(x) - p_{\text{target}}(x)|^2, \quad (1)$$

where explicit knowledge of the  $p_{\text{target}}(x)$  is used. On the other hand, training implicit models involves comparing the samples generated by the model with given data samples (e.g., with a MMD loss [30]). The MMD loss is defined as

$$\mathcal{L}_{\text{MMD}} = \mathbb{E}_{x \sim p_{\text{model}}, y \sim p_{\text{model}}} K(x, y) - 2 \mathbb{E}_{x \sim p_{\text{model}}, y \sim p_{\text{target}}} K(x, y) + \mathbb{E}_{x \sim p_{\text{target}}, y \sim p_{\text{target}}} K(x, y), \quad (2)$$

where  $K(x, y)$  is an appropriate kernel function. The MMD loss measures the distance between two probability distributions using samples drawn from the respective distributions as shown in the above equation. In the context of QGM, QCBM is an excellent example of implicit training where typically a MMD like loss-function is used. On the other hand, recent work showcases how *explicit* quantum models such as DQGM [38] and quantum quantile mechanics [11] benefit from a *functional* access to the model probability distributions, allowing input-differentiable quantum models [39,40] to solve stochastic differential equations or to model distributions with differential constraints.

Let us consider a quantum state  $|\Psi\rangle$  created by applying a quantum circuit  $\hat{U}$  (which can be parametrized) to a zero basis state. For a general  $\hat{U}$ , simulating the output PDF values that follow the Born rule  $p_{\text{model}}(x) = |\langle x|\Psi\rangle|^2$  and producing samples from  $|\Psi\rangle$  are both classically hard. But what if *estimating* the PDF for certain  $\hat{U}$  to sufficient accuracy is classically tractable? In this case one can use an explicit training, and at the inference stage have access not only to probabilities but also the capacity to sample efficiently via quantum measurements. This scenario describes a potential for classical training of quantum generative models.

To enable the classical training, we propose a strategy described in a schematic shown in Fig. 2. Here, our goal is finding circuits satisfying “classical training + quantum sampling” conditions.

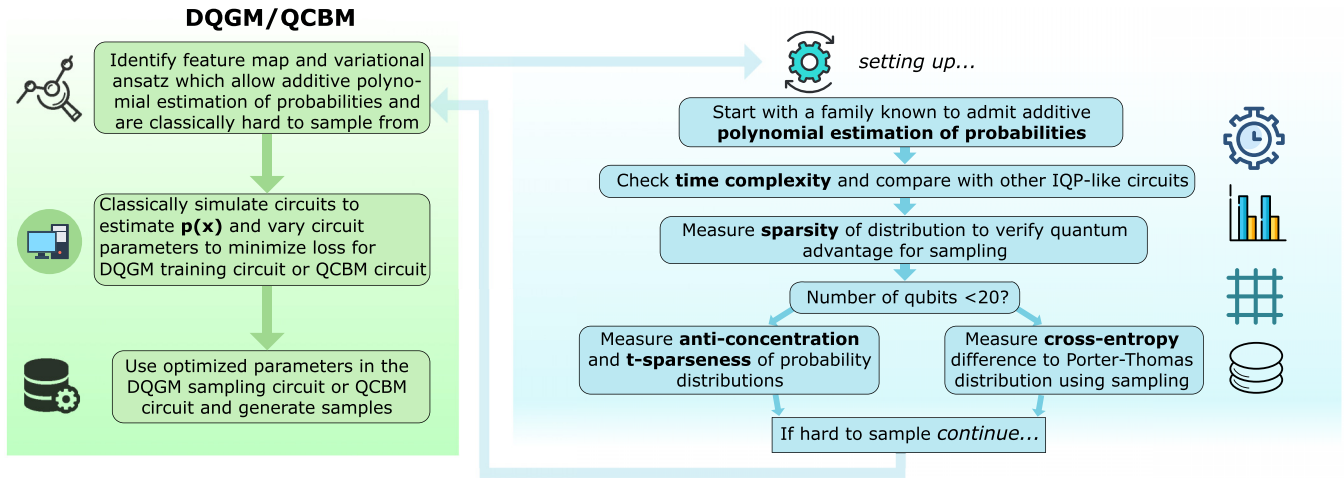


FIG. 2. Workflow used for classical training of quantum samplers, both in the DQGM and the QCBM setting. First, we identify circuits  $\hat{U}$  suitable for classically tractable probability estimation and hard sampling (see the right chart). For this we (1) check that a circuit admits additive polynomial estimation of probability; (2) compare the time complexity for exact probability calculation with known circuit families; (3) verify the sparseness of the probability distributions, depending on the number of qubits  $n$ ; and (4) measure anti-concentration/ $t$ -sparseness or cross entropy difference using sampling. Analyzing these properties, we confirm whether  $\hat{U}$  admits classical training and computationally hard sampling. After the classical training is performed variationally by minimizing DQGM/QCBM loss, we use optimized parameters for quantum sampling circuits.

### B. $\pi$ -simulable circuits that are hard to sample

We note that certain families of quantum circuits, including Clifford [41,42] and match-gate sequences [43–45], admit a classical tractable *estimation* of probabilities  $p_{\text{model}}(x)$ . At the same time, for these circuits the generation of samples is also classically “easy,” thus limiting the potential for achieving a quantum advantage. However, there exist families of quantum circuits that allow for complexity separation between the two tasks. For instance, in Ref. [36] the authors show that one can estimate probabilities for IQP (instantaneous quantum polynomial) circuits [46,47] up to an additive polynomial error, while retaining a classical hardness for sampling. For a typical IQP circuit with input  $|0^{\otimes n}\rangle$ , the amplitude to obtain a certain bit string  $x$  at output is given by

$$\Psi(x) = \langle x | \hat{H} \hat{U} \hat{H} | 0^{\otimes n} \rangle, \quad (3)$$

where  $\hat{U}$  consists of single and two-qubit Z-basis rotations and  $\hat{H}$  represents the Hadamard gate applied to all the qubits. Moreover, it is known to be classically hard to even *approximately* sample from IQP circuits in an average case [48]. Therefore, such circuits offer an opportunity for explicit training of a quantum generative model, and a potential for quantum advantage in sampling. IQP circuits by its structure are also strongly related to the so-called forrelation problem [49–51], they only differ by an additional  $\hat{H}$  layer in the middle. This corresponds to calculating an overlap between states defined as

$$\Phi = \langle 0^{\otimes n} | \hat{H} \hat{U}_2 \hat{H} \hat{U}_1 \hat{H} | 0^{\otimes n} \rangle \quad (4)$$

after the action of quantum circuit  $\hat{U}_F := \hat{H} \hat{U}_2 \hat{H} \hat{U}_1 \hat{H}$ , where  $\hat{U}_1, \hat{U}_2$  are circuits that consist of  $z$ -basis rotations  $\hat{R}_z$  and Ising-type propagators  $\hat{R}_{zz}$ . Hereafter,  $\hat{H}$  is a layer of Hadamard gates applied to all  $n$  qubits. It was shown that  $\Phi$  can be calculated efficiently classically. Therefore, from the squared forrelation  $|\Phi|^2$  one can estimate the probability to

obtain  $|0^{\otimes n}\rangle$  at the output after the action of  $\hat{U}_F$ , provided the input is  $|0^{\otimes n}\rangle$ . At the same time, in the following we show that the variant of forrelation can be used for achieving the sampling advantage.

Classical estimation of probability using the forrelation has been shown to be possible [52] under the condition that the circuit’s entangling properties are constrained. To understand this, we use the concept of connectivity in graph theory. Consider a graph  $G$  with  $n$  nodes, where each node represents a qubit in a quantum circuit. Two nodes are connected by an edge if there is a two-qubit entangling gate between the corresponding qubits in the circuit. Typically IQP circuits have all-to-all connectivity, usually by using a two-qubit entangling gate. However, if we restrict the connectivity such that the resulting connectivity graph is bipartite, we can obtain probabilities up to an additive polynomial error classically efficiently for these “extended-IQP” circuits. More concretely, whenever the connectivity graph can be partitioned into two disjoint subsets such that the tree-decomposition of each of the subsets has a small tree width, then a classical algorithm is possible with a run time of  $O(n4^w \epsilon^{-2})$ , where  $w$  is the maximum tree width of the decomposition [52],  $n$  is the number of qubits, and  $\epsilon$  is error in the estimated probability. We show examples of a bipartite graph with four nodes and a complete graph in Appendix, as well as the corresponding circuits. In the rest of the text, we use the term extended-IQP circuits to mean these quantum circuits which have a bipartite connectivity graph and an additional Hadamard layer between the set of commuting gates.

### C. Training a DQGM efficiently classically

We focus on DQGM circuits as an explicitly trained generative methodology, although the ideas hold also for architectures typically considered as implicit, like QCBM [29,30] or QGAN [53]. DQGMs allow for separation of

training and the sampling stages and allow leveraging of frequency taming techniques like feature map sparsification, qubit-wise learning, and Fourier initialization for improving and simplifying the training. DQGM also naturally allows for generative modeling for sampling from solutions to stochastic differential equations, inspired by physics-inspired neural networks (PINN) [54,55] and derivative quantum circuit (DQC)-like [38] approaches for finding solutions to the time-dependent probability density function and sampling from that. The training part of the DQGM consists of a kernel  $\hat{U}_\phi(x)$  followed by a variational circuit  $\hat{U}_\theta$ . Following Ref. [38], we define  $\hat{U}_\phi(x)$  as

$$\hat{U}_\phi(x) = \prod_{i=1}^n \left[ \hat{R}_i^z \left( \frac{2\pi x}{2^i} \right) \right] \hat{H}, \quad (5)$$

where  $\hat{R}_i^z$  are single-qubit rotation gates around the  $z$  axis, which are preceded by  $\hat{H} = \prod_{i=1}^n \hat{H}_i$  as the layer of single-qubit Hadamard gates. The operator  $\hat{U}_\phi(x)$  maps an initial state  $|\phi\rangle$  (taken as a state of  $|0^{\otimes n}\rangle$  for all qubits) to a product state  $|\tilde{x}\rangle$ , which is a latent space representation of the variable  $x$ . The transform  $\hat{U}_{T\phi}$  transforms this to a binary state  $|x\rangle$  as a bijection. This circuit is dependent on the feature map and for the above described map [Eq. (5)] corresponds to the inverse quantum Fourier transform (iQFT). The training stages can be described as a sequence of steps,

$$|\phi\rangle \xrightarrow{\hat{U}_\phi(x)} |\tilde{x}\rangle \xrightarrow{\hat{U}_{T\phi}} |x\rangle \xrightarrow{\hat{U}_\theta \hat{U}_{T\phi}^\dagger} P_{\text{train}}(0^{\otimes n}), \quad (6)$$

where  $P_{\text{train}}(0^{\otimes n})$  denotes finding classically the probability of observing  $0^{\otimes n}$  bit string.

Similarly, for the sampling stage we have

$$|\phi\rangle \xrightarrow{\hat{U}_{T\phi} \hat{U}_\theta^\dagger} P_{\text{sampling}}(|x\rangle\langle x|), \quad (7)$$

where at the end we perform quantum measurements in the computational basis. It can be shown that for a given  $x$ ,  $P_{\text{train}}(0^{\otimes n}) = P_{\text{sampling}}(|x\rangle\langle x|)$ . We therefore train  $\hat{U}_\theta$  so that  $P_{\text{train}}(0^{\otimes n}) = p_{\text{target}}(x)$  for all values of  $x$ , where  $p_{\text{target}}(x)$  is the target probability distribution we want to sample from.

Next, to be able to train the DQGM efficiently, we now rewrite the corresponding circuits  $\hat{U}_\phi(x)\hat{U}_\theta$  as an extended-IQP circuit. The extended IQP has the form  $\hat{H}\hat{U}_1\hat{H}\hat{U}_2\hat{H}$  where the depth is doubled. Thus, we can assign  $\hat{U}_\phi(x) = \hat{H}\hat{U}_1(x)$  and  $\hat{U}_\theta = \hat{U}_1\hat{H}\hat{U}_2\hat{H}$ . Here we have split up  $\hat{U}_1$  so that part of it can be used as a feature map (consisting of single-qubit  $x$ -dependent  $\hat{R}_z$  rotations) and the remaining consisting of  $\hat{R}_z$  and  $\hat{R}_{zz}$  with bipartite connectivity as a part of the variational ansatz  $\hat{U}_\theta$ . Similarly, the DQGM training circuit can be written as an IQP circuit by setting  $\hat{U}_\phi(x) = \hat{H}\hat{U}_1$  and  $\hat{U}_\theta = \hat{U}_2\hat{H}$ . Note that in the case of IQP  $\hat{U}_1$  and  $\hat{U}_2$  do not have to be limited to bipartite connectivity.

We now show that gradients with respect to circuit parameters  $\theta$  can also be classically estimated efficiently for an extended-IQP circuit. Here we assume a DQGM setting, although the conclusions hold for a QCBM as well (which we show in the next section). We have seen that for an extended-IQP circuit, we can estimate the probability of having output  $|0^{\otimes n}\rangle$ , i.e.,  $p_{\text{model}}(x) = p(0) = \text{tr}\{|0^{\otimes n}\rangle\langle 0^{\otimes n}| \hat{\rho}_{\text{out}}\}$  can be estimated efficiently, where  $\hat{\rho}_{\text{out}}$  is the output density matrix.

Suppose we have a gate  $\hat{U}_{2,j} = e^{i\theta_j \hat{P}_j/2}$  in  $\hat{U}_2$  where  $\hat{P}_j$  are the Pauli operators. Then the gradients using parameter-shift rule [26] can be written as

$$\begin{aligned} \frac{\partial p_{\text{model}}(x)}{\partial \theta_j} &= \text{tr}\{|0^{\otimes n}\rangle\langle 0^{\otimes n}| \hat{H} \hat{U}_{2,l:j+1} \hat{U}_{2,j}(\pi/2) \hat{\rho}_j \\ &\quad \times \hat{U}_{2,j}^\dagger(\pi/2) \hat{U}_{2,l:j+1}^\dagger \hat{H}\} \\ &\quad - \text{tr}\{|0^{\otimes n}\rangle\langle 0^{\otimes n}| \hat{H} \hat{U}_{2,l:j+1} \\ &\quad \times \hat{U}_{2,j}(-\pi/2) \hat{\rho}_j \hat{U}_{2,j}^\dagger(-\pi/2) \hat{U}_{2,l:j+1}^\dagger \hat{H}\}, \quad (8) \end{aligned}$$

where  $\hat{\rho}_j = \hat{U}_{2,1:j} \hat{H} \hat{U}_1 \hat{H} |0^{\otimes n}\rangle\langle 0^{\otimes n}| \hat{H} \hat{U}_1^\dagger \hat{H} \hat{U}_{2,1:j}^\dagger$ . Since both the terms in the above equation are probabilities of obtaining  $|0^{\otimes n}\rangle$  from the extended-IQP circuit, they can be estimated efficiently. A similar approach also works for an IQP circuit.

Note that in the context of solving SDEs [11,38,56], it can be shown that differentials with respect to  $x$  such as  $d p_{\text{model}}(x)/dx$  and higher-order derivatives can also be estimated efficiently classically.

We now show how gradients of probabilities with respect to circuit parameters  $\theta$  can be estimated classically efficiently. Working in the DQGM setting and using the extended-IQP architecture, we know that

$$p(x) = P(0) = |\Phi|^2, \quad (9)$$

where we use the definition of correlation  $\Phi$  from Eq. (4) and write  $\hat{U}_1 = \hat{U}_1(\theta_1)$  and  $\hat{U}_2 = \hat{U}_2(\theta_2)$  as parametrized trainable layers. Following Ref. [52],  $\Phi$  can be written as

$$\Phi = \sum_y P(y) R(y), \quad (10)$$

where  $R(y) = \frac{\langle \beta|y\rangle}{\langle \alpha|y\rangle}$ ,  $|\alpha\rangle = (\hat{H}_A \otimes \hat{I}_B) \hat{U}_1 \hat{H} |0^{\otimes n}\rangle$ ,  $|\beta\rangle = (\hat{H}_B \otimes \hat{I}_A) \hat{U}_2^\dagger \hat{H} |0^{\otimes n}\rangle$ , and  $P(y) = |\langle y|\alpha\rangle|^2$ . Thus,  $\Phi$  can be estimated by sampling from  $P(y)$ . One way of calculating the gradients with respect to  $\theta$  would be to use the parameter shift rule, which has been described in Eq. (8). However, this would involve resampling from  $P(y)$  with shifted circuit parameters. For a large number of parameters this becomes inefficient. Instead, we now use the following approach: The derivative of  $|\Phi|^2$  can be written as

$$\frac{d|\Phi|^2}{d\theta} = \Phi^* \frac{d\Phi}{d\theta} + \frac{d\Phi^*}{d\theta} \Phi. \quad (11)$$

Using Eq. (10), we differentiate  $\Phi$  with respect to  $\theta$  and get

$$\frac{d\Phi}{d\theta} = \sum_y \frac{dP(y)}{d\theta} R(y) + \sum_y P(y) \frac{dR(y)}{d\theta}. \quad (12)$$

The second term can be estimated using the same samples used to estimate  $\Phi$ . For the first term we can write

$$\sum_y \frac{dP(y)}{d\theta} R(y) = \sum_y P(y) \frac{1}{P(y)} \frac{dP(y)}{d\theta} R(y). \quad (13)$$

Now, using samples drawn from  $P(y)$  we can estimate the value of  $\frac{1}{P(y)} \frac{dP(y)}{d\theta} R(y)$ . Thus, this method avoids the need for repeated resampling from  $P(y)$  to estimate gradients. Equations (11)–(13) can then be used to estimate the gradients for  $|\Phi|^2$ .

Although we have focused on using estimating probability densities classically for training, for solving general QGM problems we would also like to be able to estimate more general observables or “cost functions.” These can involve different sets of operators other than the zero state overlap. We now show that for an extended-IQP circuit, also more general expectation values can be calculated classically efficiently. For example, let us consider an expectation value of the operator  $\Gamma = \sum_{i,j} Z_i Z_j$ , where  $i, j$  index the qubits. These terms occur in an Ising Hamiltonian which is used in the formulation of binary optimization problems using digital or analog quantum devices. For the extended-IQP circuit we can write

$$\langle \Gamma \rangle = \sum_{i,j} \langle 0^{\otimes n} | \hat{H} \hat{U}_p^\dagger \hat{H} \hat{U}_q^\dagger \hat{H} Z_i Z_j \hat{H} \hat{U}_q \hat{H} \hat{U}_p \hat{H} | 0^{\otimes n} \rangle. \quad (14)$$

We now consider a single term in the summation and consider the term  $Z_1 Z_2$ , and writing  $\hat{I}_{12} = \sum_{a,b \in \{0,1\}} |z_a z_b\rangle \langle z_a z_b|$ , we get

$$\begin{aligned} \langle Z_1 Z_2 \rangle &= \sum_{a,b,c,d \in \{0,1\}} \langle 0^{\otimes n} | \hat{H} \hat{U}_p^\dagger \hat{H} | z_a z_b \rangle \\ &\quad \times \langle z_a z_b | \hat{U}_q^\dagger \hat{H} Z_1 Z_2 \hat{H} \hat{U}_q | z_c z_d \rangle \langle z_c z_d | \hat{H} \hat{U}_p \hat{H} | 0^{\otimes n} \rangle. \end{aligned} \quad (15)$$

A single term in the summation can be written as

$$\begin{aligned} \langle \hat{Z}_1 \hat{Z}_2 \rangle_{abcd} &= \langle z_a z_b | \hat{U}_{12} | z_c z_d \rangle \langle 0^{\otimes n} | \hat{H} \hat{U}_p^\dagger \hat{H} | z_a z_b \rangle \\ &\quad \times \langle z_c z_d | \hat{H} \hat{U}_p \hat{H} | 0^{\otimes n} \rangle, \end{aligned} \quad (16)$$

where  $\hat{U}_{12} = \hat{U}_{q,12}^\dagger \hat{H}_1 \hat{H}_2 Z_1 Z_2 \hat{H}_1 \hat{H}_2 \hat{U}_{q,12}$ , with  $\hat{H}_i$  being the Hadamard gate acting on qubit  $i$ , and the terms in  $\hat{H}$ ,  $\hat{U}_q$  which do not contain terms for qubits 1,2 commute through  $Z_1, Z_2$  and meet their conjugates and are converted to identity. This term can be calculated classically efficiently since it is only a two-qubit overlap integral. The second term in the above product can be written as

$$\begin{aligned} \langle 0^{\otimes n} | \hat{H} \hat{U}_p^\dagger \hat{H} | z_a z_b \rangle \langle z_c z_d | \hat{H} \hat{U}_p \hat{H} | 0^{\otimes n} \rangle &= \\ = \langle 0^{\otimes n} | \hat{H} \hat{U}_p^\dagger \hat{\delta} \hat{U}_p \hat{H} | 0^{\otimes n} \rangle, \end{aligned} \quad (17)$$

where  $\hat{\delta} = \hat{H} | z_a z_b \rangle \langle z_c z_d | \hat{H}$ , which is a tensor product operator. The authors of Ref. [52] prove that this term can also be calculated classically efficiently up to an additive polynomial error. Hence,  $\langle Z_1 Z_2 \rangle$  and consequently  $\langle \Gamma \rangle$  can be calculated classically efficiently. It can be similarly shown that expectation values for operators like  $\sum Z_i$  can be calculated classically efficiently. This means various expectation values and thus different loss/cost functions can be estimated classically efficiently up to additive polynomial error. Therefore, the extended-IQP circuits can be classically trained using not just probability densities, but also a variety of cost functions based on measuring expectation values of different observables. This may be useful, for example, when one had to sample bit strings which minimized a certain Hamiltonian.

#### D. Training a QCBM efficiently classically

In the QCBM setting, we estimate  $p_{\text{model}}(x)$  directly using classically simulated output bit strings for fixed input  $|0^{\otimes n}\rangle$ . The amplitude to obtain a certain bit string  $x$  at the output of

an extended-IQP circuit can be written as

$$\Phi_x = \langle x | \hat{H} \hat{U}_2 \hat{H} \hat{U}_1 \hat{H} | 0^{\otimes n} \rangle. \quad (18)$$

Writing  $|x\rangle = (\prod_{i=1}^k X_i) |0^{\otimes n}\rangle$ , where  $i$  is indexed over locations where 1 occurs in state  $|x\rangle$ . Thus, we can write

$$\begin{aligned} \Phi_x &= \langle 0^{\otimes n} | \left( \prod_{i=1}^k \hat{X}_i \right) \hat{H} \hat{U}_2 \hat{H} \hat{U}_1 \hat{H} | 0^{\otimes n} \rangle \\ &= \langle 0^{\otimes n} | \hat{H} \left( \prod_{i=1}^k \hat{Z}_i \right) \hat{U}_2 \hat{H} \hat{U}_1 \hat{H} | 0^{\otimes n} \rangle, \end{aligned} \quad (19)$$

where we used the fact that  $\hat{H}_i \hat{X}_i = \hat{Z}_i \hat{H}_i$ . Absorbing the  $\hat{Z}$  gates into  $\hat{U}_2$ , the above equation can be written as a correlation and thus can be computed classically efficiently. Using Eq. (19),  $p_{\text{model}}(x)$  can be computed. Similarly, just like Eq. (8), the gradients with respect to  $\theta$  can be written as

$$\begin{aligned} \frac{\partial p_{\text{model}}(x)}{\partial \theta_j} &= \text{tr}\{|x\rangle \langle x| \hat{H} \hat{U}_{2,l:j+1} \hat{U}_{2,j}(\pi/2) \hat{\rho}_j \\ &\quad \times \hat{U}_{2,j}^\dagger(\pi/2) \hat{U}_{2,l:j+1}^\dagger \hat{H}\} - \text{tr}\{|x\rangle \langle x| \hat{H} \\ &\quad \times \hat{U}_{2,l:j+1} \hat{U}_{2,j}(-\pi/2) \hat{\rho}_j \hat{U}_{2,j}^\dagger(-\pi/2) \hat{U}_{2,l:j+1}^\dagger \hat{H}\}, \end{aligned} \quad (20)$$

where  $\hat{\rho}_j = \hat{U}_{2,1:j} \hat{H} \hat{U}_1 \hat{H} | 0 \rangle \langle 0 | \hat{H} \hat{U}_1^\dagger \hat{H} \hat{U}_{2,1:j}^\dagger$ . Since both the terms in the above equation are probabilities of obtaining a certain bit string  $x$  at the output of an extended-IQP circuit, using Eqs. (18) and (19), they can be estimated classically. However, as described in the previous section, estimating gradients using the parameter shift rule will require re-sampling from  $P(y)$  for shifted parameters. Similar to the DQGM setting, we can estimate the gradients without the need for resampling for each parameter by replacing  $\Phi$  with  $\Phi_x$  and using Eqs. (11)–(13).

#### E. Complexity of classical simulability: Probabilities and sampling

As discussed before, to enable classical training and hard sampling we need to check the properties of quantum circuits that we train. For this, we develop a workflow used for studying different properties of the chosen circuits (see the chart in Fig. 2, right). We start selecting a family of circuits that allows additive polynomial estimation of probabilities. To show that it is still hard to sample from, we show that probabilities generated from these circuits are not polysparse [36,42]. We use two different approaches to show this. One approach involves numerical random sampling of these circuits and looking at their anticoncentration properties [57,58]. An output distribution of a unitary  $\hat{U}$  for some settings of its parameters is said to anticoncentrate when

$$\Pr_{\hat{U} \sim \mu} \left( |\langle x | \hat{U} | 0 \rangle|^2 \geq \frac{\alpha}{N} \right) \geq \beta \quad (21)$$

for constants  $\alpha, \beta$ , where  $N = 2^n$  and  $\hat{U}$  is drawn from a certain measure  $\mu$ . For example, Ref. [58] shows a class of families for which  $\beta = 1/e$ . The probability distributions of these circuits along with families discussed in Ref. [59] converge to the Porter-Thomas distribution. In Ref. [36], the

authors prove that the anticoncentration and polysparsity cannot coexist. Thus, if we show that probability distributions from a family anticoncentrate, then we can conclude that the probability distributions are not polysparse and hence are hard to sample from. We use the approach discussed above for studying systems with up to 20 qubits. For a larger number of qubits we use the fact that the probability distributions converge to the Porter-Thomas distribution and use the cross-entropy difference to approximately measure the distance with the Porter-Thomas distribution. The cross-entropy difference is defined as

$$\Delta H(p_{\text{samp}}) \approx H_0 - \frac{1}{m} \sum_{j=1}^m \ln \frac{1}{p_U(x_j^{\text{samp}})}, \quad (22)$$

where  $H_0 = \ln(N) + \gamma$ , and  $\gamma \approx 0.577$  is Euler's constant.  $p_U(x_j^{\text{samp}})$  corresponds to the probability computed classically for the generated samples only. The error in  $\Delta H(p_{\text{samp}})$  is given by  $\kappa/\sqrt{m}$ , where  $\kappa \approx m$ . Thus, if we have a way of generating a finite number of samples, we can approximately characterize the distribution without the need of calculating all the probabilities. This is especially useful for larger registers ( $n \geq 25$ ) where state-vector calculations for all the probabilities (needed to measure anticoncentration or sparsity) rapidly becomes unfeasible. This approach has been used for classical benchmarking of data from random quantum circuits for  $\sim 50$  qubits [1,59].

To study resource requirements of various circuit families we use tensor networks to represent our quantum circuits, and we analyze their properties with classical simulation. Tensor networks use a tree-based decomposition to estimate the time complexity of calculating probabilities. This is done by estimating the size of the largest tensor during the contraction process [60]. The maximum size depends on the contraction order, and various algorithms are used to find the contraction order which gives the smallest tensor size [2,61,62].

Apart from using anticoncentration, we can also measure whether a probability is polysparse or not. This is done by measuring the number of terms needed to  $\epsilon$ -approximate it with a sparse distribution. A  $t$ -sparse distribution, with only  $t$  nonzero terms, can  $\epsilon$ -approximate a probability distribution  $P(X)$  if and only if  $\sum_x |P(x) - P_t(x)| \leq \epsilon$  [37]. Here,  $P_t(x)$  is the probability distribution containing only the highest  $t$  terms from  $P_t(x)$  as the nonzero terms. We know that for  $\epsilon = 0$ ,  $t = N$ , where  $N = 2^n$  ( $n$  is the number of qubits). Therefore, we can approximate the behavior of  $t$  as  $t(\epsilon) = N[1 - f(1/\epsilon)]$ , where function  $f$  shows a polynomial behavior if the distribution is polysparse. For an exponential behavior  $f \sim e^{-1/\epsilon}$ , the distribution is dense. Therefore, after calculating  $t$  for different values of  $\epsilon$ , we calculate  $f(1/\epsilon) = 1 - t/N$  and plot this as a function of  $1/\epsilon$ .

### III. RESULTS

We proceed to implement the proposed strategies in practice. For enabling the classical training, we choose different quantum circuit families that include extended-IQP circuits compared with product, Hadamard, IQP, and IQP one-dimensional (1D)-chain circuits (see corresponding diagrams in the Appendix). First, we compare the time complexity for




FIG. 3. Time complexity for different families of circuits shown on a logarithmic scale. We observe that for IQP and extended-IQP the time complexity scales exponentially. Simple Hadamard circuits has a constant complexity. For product circuits and IQP circuits in a 1D chain, we see that the complexity has a saturation behavior and thus only a polynomial increase (also see inset).

different families shown in Fig. 3. These plots have been generated using Julia libraries YAOTOEINSUM for tensor network representation of quantum circuits built in YAO, which is based on the generic tensor contraction tool OMEINSUM [61–64]. As expected, for product, Hadamard, and 1D chain, the maximum size of the tensor during the contraction grows and quickly saturates (see inset in Fig. 3). For IQP and extended IQP, the time complexity grows linearly in the logarithmic scale. This implies that the classical computational complexity for calculating exact probabilities of the extended-IQP circuits, just like for IQPs, is exponential in the number of qubits.

Next, we study the anticoncentration properties of quantum circuits. Figure 4 shows the anticoncentration as a fraction of nonuniform probabilities compared to random circuits with




FIG. 4. Figure shows the anticoncentration properties of extended IQP, which has a bipartite connectivity graph, and IQP with full connectivity. The fraction of probabilities  $\geq 1/2^n$  is very close to  $1/e$ , which is shown by a dotted line. The error bars show the standard deviation over 100 instances of unitary matrices for each family.




FIG. 5. Figure shows the total variation distance with respect to the Porter-Thomas distribution. The distance rapidly approaches zero as we increase the number of qubits. The plot shows the mean and the variance for the total variation distance over 100 distributions for each qubit number.

bipartite connectivity. The randomness is chosen as follows. The first layer, the middle layer, and the end layer are all composed of Hadamard gates. The  $\hat{R}_{zz}(\theta)$  and the  $\hat{R}_z(\theta)$  gates are chosen such that  $\theta = k\pi/8$ , with  $k$  uniformly randomly chosen from  $[0, 1, \dots, 7]$  [48]. We observe that this set of gates approximates  $\hat{U}$  drawn uniformly randomly from the Haar measure [33]. Specifically, we observe from Fig. 4 that the fraction of probabilities  $> 1/2^n$  is very close to  $1/e$  (dashed line), which is a good indicator that the probabilities do indeed anticentralize. The averaging is performed over 100 random circuits for each qubit number.

We can also show that probabilities for these circuits converge towards Porter-Thomas distribution for a greater

number of qubits, with the variance reducing as well. Figure 5 shows the total variation distance measured for numerically obtained probability distributions for the extended-IQP circuit for 100 random configurations with respect to the Porter-Thomas distribution. The corresponding PDF is  $P_{\text{PT}}(p) = Ne^{-Np}$ . Following Ref. [58], the variation distance is defined as

$$\|P - Q\|_{\text{TV}} := \frac{1}{2} \sum_{X \in \Omega} |P(X) - 1/m|, \quad (23)$$

where we divide the set of probabilities into  $m$  equally weighted bins  $[p_0, \dots, p_m]$  and take  $\int_{p_i}^{p_{i+1}} P_{\text{PT}} dp = 1/m$ . The set  $\Omega$  is the set of probabilities in the interval  $[p_i, p_{i+1}]$ , where  $i$  goes from 0 to  $m$ .  $Q$  is the set of probabilities observed numerically over the set  $\Omega$ . We observe the distance rapidly approaching zero as the number of qubits is increased. We continue to study the sparsity. In Figs. 6(a)–6(d) we show the log-log plots of  $f(1/\epsilon)$  vs  $1/\epsilon$  for different numbers of qubits. The results are for a single random distribution. The downward curvature shows a superpolynomial decay rate, which indicates that the probability distribution is not polysparse.

Figure 7 shows the results of training a quantum generative model as QNN based on the extended-IQP architecture for six qubits for a Gaussian probability density function. The circuit consists of an initial phase feature map as a part of the extended-IQP architecture. Using the training stage as described in Sec. II [Eq. (6)], we try to maximize  $P_{\text{train}}(0^{\otimes n})$  for different values of  $x$  by training  $\hat{U}_\theta$ . For the cost function we use the mean square error,  $\mathcal{L} = \sum_x |p_{\text{model}}(x) - p_{\text{target}}(x)|^2$ . We see from Fig. 7(a) that the trained QGM is able to closely follow the curve over the entire domain. The training has been performed using 128 points and a phase-feature map defined in Eq. (5). We then use the trained circuit to generate samples using the sampling stage [Eq. (7)]. The results for 20 000 shots is shown in Fig. 7(b).




FIG. 6. Plot of  $f(\epsilon^{-1})$  vs  $\epsilon^{-1}$  for a random probability distribution for different number of qubits (10, 12, 14, and 16) as a log-log plot. The downward curvature is an indication of superpolynomial behavior of  $f(\epsilon^{-1})$ .




FIG. 7. (a) Results for training QNN based on an extended-IQP circuit to generate a Gaussian probability density for six qubits. The figure shows an excellent fit between the trained and the target distribution. (b) Result of using the trained circuit in the training stage to generate samples. The plot shows count density for 20 000 shots. (c) How the MSE loss goes down for 100 training steps.

We also performed training for 30 qubits using the classical algorithms to estimate probabilities and their gradients based on Eqs. (10)–(13). To avoid issues related to barren plateau for training a large number of qubits, we choose a particular probability distribution for training. Specifically, we choose  $p(x)$  to be a Gaussian probability density function from 0 to 63 and thus can be generated by effectively training qubits 1–6. For qubits from 7 to 30, we apply the identity transformation as an initial setting. To do this, while still using the code for training 30 qubits, we apply the setup as shown in Fig. 8(a). The feature map is applied to qubits 1–6. To apply an identity transformation for the rest of the qubits we use the following fixed settings:

1. The  $\theta_s$  for all the two-qubit gates involving qubits 7–30 are set to 0.
2. The  $\theta_s$  for all the single-qubit gates are set to  $\pi/2$  (which effectively sets the angle to  $\pi/4$ ). These gates, along with the three  $\hat{H}$  layers in the extended-IQP architecture, effectively implement identity transformation on these qubits.

We now allow all the parameters to be trained (including for qubits 7–30). Starting with an initial identity transformation for qubits 7–30 ensures that the nonzero probability

density largely remains confined to the events involving qubits 1–6 during the entire course of the training. Figure 8(b) shows the results of training for 32 points and 100 training steps. The loss function value at convergence is  $6.22 \times 10^{-6}$ . This simulation took approximately 42 h on a regular desktop computer. While effectively the distribution is defined over six qubits, we stress that calculation of quantities like  $R(y)$ , sampling from  $P(y)$ , and calculation of  $\Phi$  as defined in Eq. (10) and calculation of gradients involved all 30 qubits.

#### IV. CONCLUSIONS AND OUTLOOK

Our results show that certain circuit families, which we here call extended IQP, can be trained classically by estimating probabilities up to an additive polynomial error, using the explicit generative modeling paradigm. We show that these circuits can be trained by estimating gradients classically in QCBM and DQGM settings. Using these techniques, we train a probability distribution for 30 qubits on a regular desktop computer. At the same time we show that these circuits still retain quantum advantage in terms of sampling. This we did by looking at the anticoncentration as well as the t-sparseness

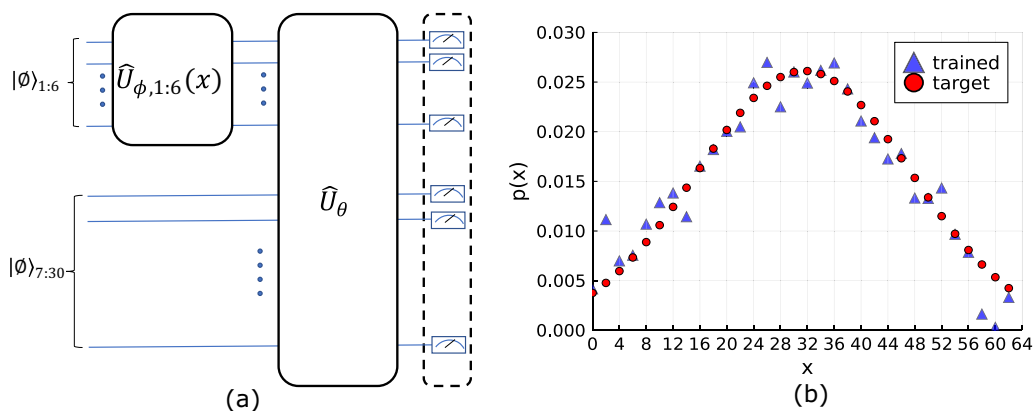


FIG. 8. (a) Setup used to train a probability distribution for 30 qubits. Only the parameters involving qubits 1–6 are updated during training, while the remaining qubits are kept unmodified. (b) Results of training the DQGM circuit with an extended-IQP architecture to output a Gaussian probability distribution for  $x$  between 1 and 64 with 32 equally spaced training points (integers label consecutive bit strings). The estimates to the model probability density after training are obtained using Eq. (9) with the obtained trained parameters. Loss values of  $6.22 \times 10^{-6}$  are reached at convergence.



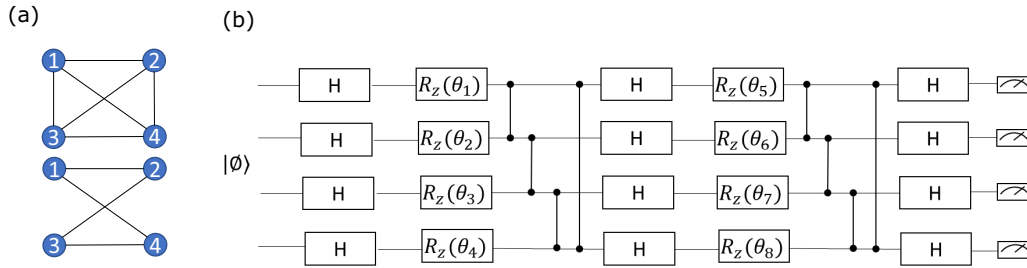


FIG. 9. (a) Qubit connectivity graph for all-to-all and bipartite qubit connectivity. (b) Extended-IQP circuit for qubits with bipartite connectivity.

properties of the probability distributions up to 16 qubits. For higher numbers of qubits, cross-entropy benchmarking using samples based on tensor networks will be studied.

Recent work [65] has highlighted the difficulty of training quantum generative models in the worst case, provided one has access only to estimates of quantities related to the target probability distribution. While in the case of QCBM training using MMD loss this could be very important, in our case this issue does not arise since we are assuming knowledge of the target probability distribution.

So far we have focused on a single layer of Hadamards in the middle of commuting gates. But it may be possible to also extend these results to other single-qubit operators. In addition it has been shown [52] that depth = 2 QAOA circuits also allow additive polynomial estimation of the  $\langle \psi | \hat{H}_{\text{prob}} | \psi \rangle$ , thus allowing classical training of these circuits while still showing quantum advantage in sampling. It could be an interesting possibility to classically simulate quantum annealing schedules to optimize annealing parameters. It would further-

more be interesting to study applications for this architecture for optimization problems [66,67].

**ACKNOWLEDGMENTS**

We thank QuiX Quantum for fruitful discussions.

A patent application for the method described in this manuscript has been submitted by PASQAL [34].

**APPENDIX: CIRCUITS**

Figure 9(a) shows an example of a qubit connectivity graph for an all-to-all (top) and bipartite (bottom) qubit connectivity. Our extended-IQP circuits are limited to bipartite connectivity to ensure classical training. Figure 9(b) show the extended-IQP circuit corresponding to the bipartite graph in Fig. 9(a).

Figure 10 shows different architectures studied for comparison of time-complexity with extended-IQP circuits in Fig. 3.

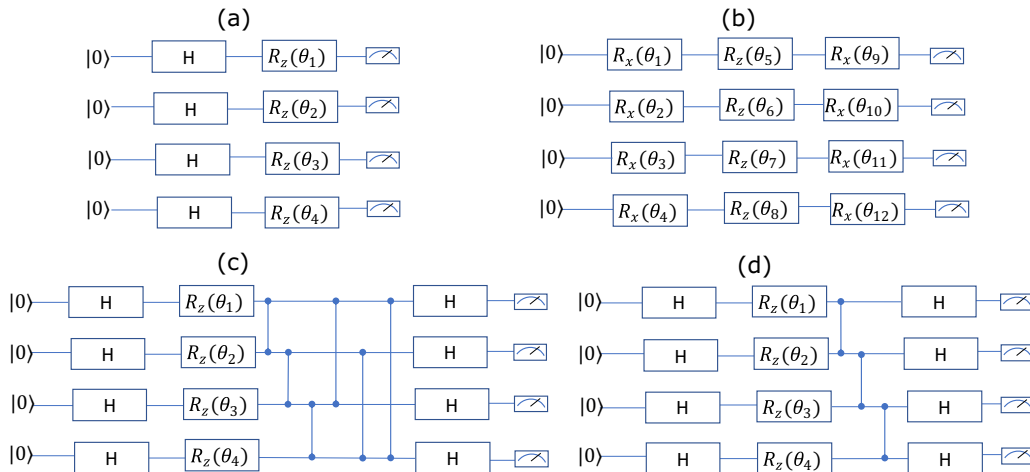


FIG. 10. Different architectures studied: (a) Hadamard, (b) product, (c) IQP, and (d) IQP 1D chain.

[1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell *et al.*, *Nature (London)* **574**, 505 (2019).  
 [2] F. Pan, K. Chen, and P. Zhang, *Phys. Rev. Lett.* **129**, 090502 (2022).

[3] H. S. Zhong, H. Wang, Y. H. Deng, M. C. Chen, L. C. Peng, Y. H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X. Y. Yang, W. J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang *et al.*, *Science* **370**, 1460 (2020).

- [4] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, A. E. Lita, T. Gerrits, S. W. Nam, V. D. Vaidya, M. Menotti, I. Dhand, Z. Vernon, N. Quesada, and J. Lavoie, *Nature (London)* **606**, 75 (2022).
- [5] I. Georgescu, *Nat. Rev. Phys.* **4**, 362 (2022).
- [6] K. M. Abeywickrama, S. Ganguly, L. Gerardo Ayala Bertel, and S. Mohanty, *Quantum and Blockchain for Modern Computing Systems: Vision and Advancements*, edited by A. Kumar, S. S. Gill, and A. Abraham (Springer International Publishing, Cham, 2022), Vol. 133, pp. 205–222.
- [7] P.-L. Dallaire-Demers and N. Killoran, *Phys. Rev. A* **98**, 012324 (2018).
- [8] L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou, and L. Sun, *Sci. Adv.* **5**, eaav2761 (2019).
- [9] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, *npj Quantum Inf.* **5**, 45 (2019).
- [10] G. Verdon, J. Marks, S. Nanda, S. Leichenauer, and J. Hidary, [arXiv:1910.02071](https://arxiv.org/abs/1910.02071).
- [11] A. E. Paine, V. E. Elfving, and O. Kyriienko, *Adv Quantum Technol.* 2300065 (2023).
- [12] K. Kubo, Y. O. Nakagawa, S. Endo, and S. Nagayama, *Phys. Rev. A* **103**, 052425 (2021).
- [13] K. A. Landsman, C. Figgatt, T. Schuster, N. M. Linke, B. Yoshida, N. Y. Yao, and C. Monroe, *Nature (London)* **567**, 61 (2019).
- [14] X. Li, G. Zhu, M. Han, and X. Wang, *Phys. Rev. A* **100**, 032309 (2019).
- [15] J. Harris, B. Yan, and N. A. Sinitsyn, *Phys. Rev. Lett.* **129**, 050602 (2022).
- [16] L. Banchi, M. Fingerhuth, T. Babej, C. Ing, and J. M. Arrazola, *Sci. Adv.* **6**, eaax1950 (2020).
- [17] S. Ebadi, A. Keesling, M. Cain, T. T. Wang, H. Levine, D. Bluvstein, G. Semeghini, A. Omran, J.-G. Liu, R. Samajdar, X.-Z. Luo, B. Nash, X. Gao, B. Barak, E. Farhi, S. Sachdev, N. Gemelke, L. Zhou, S. Choi, H. Pichler *et al.*, *Science* **376**, 1209 (2022).
- [18] C. Gross and I. Bloch, *Science* **357**, 995 (2017).
- [19] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, M. Kwon, M. Ebert, J. Cherek, M. T. Lichtman, M. Gillette, J. Gilbert, D. Bowman, T. Ballance, C. Campbell, E. D. Dahl *et al.*, *Nature (London)* **604**, 457 (2022).
- [20] L. Henriot, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Raymond, and C. Jurczak, *Quantum* **4**, 327 (2020).
- [21] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, *Nature (London)* **460**, 240 (2009).
- [22] M. P. Harrigan, K. J. Sung, M. Neeley, K. J. Satzinger, F. Arute, K. Arya, J. Atalaya, J. C. Bardin, R. Barends, S. Boixo *et al.*, *Nat. Phys.* **17**, 332 (2021).
- [23] E. A. Martinez, C. A. Muschik, P. Schindler, D. Nigg, A. Erhard, M. Heyl, P. Hauke, M. Dalmonte, T. Monz, P. Zoller, and R. Blatt, *Nature (London)* **534**, 516 (2016).
- [24] C. Figgatt, D. Maslov, K. A. Landsman, N. M. Linke, S. Debnath, and C. Monroe, *Nat. Commun.* **8**, 1918 (2017).
- [25] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Nature (London)* **323**, 533 (1986).
- [26] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, *Phys. Rev. A* **98**, 032309 (2018).
- [27] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, *Phys. Rev. A* **99**, 032331 (2019).
- [28] O. Kyriienko and V. E. Elfving, *Phys. Rev. A* **104**, 052417 (2021).
- [29] B. Coyle, D. Mills, V. Danos, and E. Kashefi, *npj Quantum Inf.* **6**, 60 (2020).
- [30] J.-G. Liu and L. Wang, *Phys. Rev. A* **98**, 062324 (2018).
- [31] C. Zoufal, A. Lucchi, and S. Woerner, *npj Quantum Inf.* **5**, 103 (2019).
- [32] H.-L. Huang, Y. Du, M. Gong, Y. Zhao, Y. Wu, C. Wang, S. Li, F. Liang, J. Lin, Y. Xu, R. Yang, T. Liu, M.-H. Hsieh, H. Deng, H. Rong, C.-Z. Peng, C.-Y. Lu, Y.-A. Chen, D. Tao, X. Zhu *et al.*, *Phys. Rev. Appl.* **16**, 024051 (2021).
- [33] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Nat. Commun.* **9**, 4812 (2018).
- [34] A patent application for the method described in this manuscript has been submitted by *PASQAL* with Sachin Kasture and Vincent E. Elfving as inventors.
- [35] L. Banchi, N. Quesada, and J. M. Arrazola, *Phys. Rev. A* **102**, 012417 (2020).
- [36] H. Pashayan, S. D. Bartlett, and D. Gross, *Quantum* **4**, 223 (2020).
- [37] M. van den Nest, *Quantum Inf. Comput.* **11**, 784 (2011).
- [38] O. Kyriienko, A. E. Paine, and V. E. Elfving, [arXiv:2202.08253](https://arxiv.org/abs/2202.08253).
- [39] O. Kyriienko, A. E. Paine, and V. E. Elfving, *Phys. Rev. A* **103**, 052416 (2021).
- [40] A. E. Paine, V. E. Elfving, and O. Kyriienko, *Phys. Rev. A* **107**, 032428 (2023).
- [41] D. Gottesman, in *Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, edited by S. P. Corney, R. Delbourgo, and P. D. Jarvis (International Press, Cambridge, 1999), pp. 32–43.
- [42] M. van Den Nest, *Quantum Inf. Comput.* **10**, 0258 (2010).
- [43] R. Jozsa and A. Miyake, *Proc. R. Soc. A* **464**, 3089 (2008).
- [44] B. M. Terhal and D. P. DiVincenzo, *Phys. Rev. A* **65**, 032325 (2002).
- [45] D. J. Brod, *Phys. Rev. A* **93**, 062332 (2016).
- [46] D. Shepherd, [arXiv:1005.1744](https://arxiv.org/abs/1005.1744).
- [47] D. J. Shepherd, [arXiv:1005.1425](https://arxiv.org/abs/1005.1425).
- [48] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Phys. Rev. Lett.* **117**, 080501 (2016).
- [49] S. Aaronson and A. Ambainis, in *Proceedings of the 47th Annual ACM Symposium on Theory of Computing* (ACM, New York, 2015), pp. 307–316.
- [50] N. Bansal and M. Sinha, in *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing* (ACM, New York, 2021), pp. 1303–1316.
- [51] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, *Nature (London)* **567**, 209 (2019).
- [52] S. Bravyi, D. Gosset, D. Grier, and L. Schaeffer, [arXiv:2102.06963](https://arxiv.org/abs/2102.06963).
- [53] S. Lloyd and C. Weedbrook, *Phys. Rev. Lett.* **121**, 040502 (2018).
- [54] M. Raissi, P. Perdikaris, and G. E. Karniadakis, *J. Comput. Phys.* **378**, 686 (2019).
- [55] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, *Nat. Rev. Phys.* **3**, 422 (2021).

- [56] H. Alghassi, A. Deshmukh, N. Ibrahim, N. Robles, S. Woerner, and C. Zoufal, [Quantum](#) **6**, 730 (2022).
- [57] D. Hangleiter, J. Bermejo-Vega, M. Schwarz, and J. Eisert, [Quantum](#) **2**, 65 (2018).
- [58] J. Bermejo-Vega, D. Hangleiter, M. Schwarz, R. Raussendorf, and J. Eisert, [Phys. Rev. X](#) **8**, 021010 (2018).
- [59] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, [Nat. Phys.](#) **14**, 595 (2018).
- [60] I. L. Markov and Y. Shi, [SIAM J. Comput.](#) **38**, 963 (2008).
- [61] G. Kalachev, P. Pantelev, and M.-H. Yung, [arXiv:2108.05665](#).
- [62] J. Gray and S. Kourtis, [Quantum](#) **5**, 410 (2021).
- [63] X.-Z. Luo, J.-G. Liu, P. Zhang, and L. Wang, [Quantum](#) **4**, 341 (2020).
- [64] S. Kourtis, C. Chamon, E. Mucciolo, and A. Ruckenstein, [SciPost Phys.](#) **7**, 060 (2019).
- [65] M. Hinsche, M. Ioannou, A. Nietner, J. Haferkamp, Y. Quek, D. Hangleiter, J.-P. Seifert, J. Eisert, and R. Sweke, [arXiv:2110.05517](#).
- [66] E. Farhi, J. Goldstone, and S. Gutmann, [arXiv:1411.4028](#).
- [67] S. Varsamopoulos, E. Philip, H. W. T. van Vlijmen, S. Menon, A. Vos, N. Dyubankova, B. Torfs, A. Rowe, and V. E. Elfving, [arXiv:2205.02807](#).