

Learning AC Power Flow Solutions using a Data-Dependent Variational Quantum Circuit

Thinh Viet Le, Md Obaidur Rahman, and Vassilis Kekatos
Elmore Family School of Electrical and Computer Engineering, Purdue University
West Lafayette, IN 47906, USA
Emails: {le272,rahma160,kekatos}@purdue.edu

Abstract—Interconnection studies require solving numerous instances of the AC load or power flow (AC PF) problem to simulate diverse scenarios as power systems navigate the ongoing energy transition. To expedite such studies, this work leverages recent advances in quantum computing to find or predict AC PF solutions using a variational quantum circuit (VQC). VQCs are trainable models that run on modern-day noisy intermediate-scale quantum (NISQ) hardware to accomplish elaborate optimization and machine learning (ML) tasks. Our first contribution is to pose a single instance of the AC PF as a nonlinear least-squares fit over the VQC trainable parameters (weights) and solve it using a hybrid classical/quantum computing approach. The second contribution is to feed PF specifications as features into a data-embedded VQC and train the resultant quantum ML (QML) model to predict general PF solutions. The third contribution is to develop a novel protocol to efficiently measure AC-PF quantum observables by exploiting the graph structure of a power network. Preliminary numerical tests indicate that the proposed VQC models attain enhanced prediction performance over a deep neural network despite using much fewer weights. The proposed quantum AC-PF framework sets the foundations for addressing more elaborate grid tasks via quantum computing.

Index Terms—Variational quantum circuit, quantum machine learning, gradient descent, quantum observables, bus admittance matrix.

I. INTRODUCTION

Solving the AC power flow (AC PF) is an imperative task in power system simulation, operation, and planning. Given specified load conditions and generator settings, the AC PF seeks the complex voltage phasors at all buses [1]. Once voltage phasors are found, other quantities of interest, such as line power and currents, can be readily computed. Under the forthcoming energy transition, the increased uncertainty in generation, load demand, and prevailing technologies calls for conducting interconnection studies involving PF instances of ever-increasing volume, variability, and spatiotemporal resolution. This work explores recent advances in quantum machine learning (QML) to expedite bulk PF studies.

Traditional algorithms for solving the AC PF rely on linearization techniques or iterative nonlinear equation solvers [2]. Linearized variants of the PF fail to capture line flows and thermal losses. Iterative methods coping with the AC PF equations include the Gauss-Seidel, the impedance

matrix, and Newton-Raphson methods [3]. These tools have been studied extensively, and their mature implementations are widely used in industry practice. Reference [4] reformulates the AC PF task as a feasibility problem and relaxes it into a semidefinite program (SDP) to find AC-PF solutions even when Newton-Raphson iterates fail to converge.

Previous quantum computing attempts to cope with the AC PF task have focused on the classical Newton-Raphson iterations. The idea was to merely substitute the linear system solver with the Harrow–Hassidim–Lloyd (HHL) algorithm [5]–[7]. The HHL algorithm runs on fault-tolerant rather than noisy intermediate-scale quantum (NISQ) quantum computers and can solve systems of N linear equations in $\mathcal{O}(\log N)$. Nonetheless, its complexity scales unfavorably for linear systems involving the Jacobian of the AC-PF equations [8].

All previous solvers deal with a single PF instance at a time. The quasi-static time series (QSTS) method resorts to elaborate initializations and sampling schemes for solving the AC PF in distribution grids across successive instances [9]. Similar ideas may not scale equally well under the wider variability of planning studies in transmission systems. To address this challenge, a recent line of research advocates training machine learning (ML) models, including graph and physics-informed neural networks, to predict AC-PF solutions once presented with the problem specifications [10]–[13]. The motivation is to shift the computational burden to offline and perform the time-efficient prediction step during real-time operation. To sufficiently represent various AC-PF instances, the number of weights in classical ML models is massive, hindering both training and inference [13]. Moreover, previous ML models are typically trained in a supervised manner, requiring large datasets of labeled solved PF instances.

Quantum ML (QML) models have been recently proposed as an alternative to classical ones [14], [15]. However, available quantum hardware is confined in terms of width (number of qubits), depth (number of gate layers), and sensitivity to noise. Under these limitations, current QML models are designed to run on *variational quantum circuits* (VQC) [16, Ch. 5]. In a VQC, both inputs and trainable weights are entered as tunable parameters of controllable quantum gates, mathematically represented as parameterized unitary matrices. The VQC output is measured as the expectation over a *quantum observable*, described by a quadratic function of the VQC state

defined over a Hermitian matrix. Interestingly, recent results show that for the same number of trainable weights, a QNN can approximate smooth functions at higher accuracy than any classical NN [15], [17].

Measuring quantum observables is indispensable during both training and inference of a VQC. The VQC output can be measured using the technique of linear combination of unitaries (LCU) [16, p. 138]. Unfortunately, for problems featuring Hermitian matrices with dense diagonals, the number of unitaries scales exponentially with the number of qubits. Reference [7] suggests a variational quantum approach to implement the Newton-Raphson iteration for the AC PF. Because the method relies on the LCU decomposition of the PF Jacobian, it is challenged by the curse of dimensionality. Reference [18] proposes a novel quantum measurement protocol, wherein specific measurement matrices can be diagonalized using a few quantum-implementable unitaries. In [19], we adopted this protocol to quantum observables associated with the AC optimal power flow (AC OPF) problem. Despite its resemblance to the AC OPF, the AC PF problem requires special treatment in terms of measuring quantum observables.

To bridge the identified gaps, this work proposes a data-based quantum power flow framework to solve the AC PF problem. Its technical contributions are on three fronts: *i*) Develop a model for solving the AC PF via a hybrid classical-quantum algorithm (Section III); *ii*) Leverage a data-embedded VQC and train a QML model to predict AC PF solutions in an unsupervised manner (Section IV); *iii*) Expedite the computation of VQC/QML by reformulating PF to involve fewer expectations over quantum observables (Section V). Numerical tests on the IEEE 14-bus system demonstrate that the novel QPF framework predicted AC PF solutions with smaller errors while using significantly fewer parameters than a deep neural network (Section VI). Section VII concludes our findings and suggests opening research directions.

II. QUANTUM COMPUTING PRELIMINARIES

A quantum computer operating on $\log N$ quantum bits (qubits) is a random sampler of $\log N$ classical bits. Every time this computer is run, we read a string of $\log N$ classical bits at its output. Quantum computing relies on manipulating the probability mass function (PMF) of the N bitstring outcomes, and is founded on four basic postulates [20]. According to the first, a quantum system operating on $\log N$ qubits is described by a *state vector* $\mathbf{x} \in \mathbb{C}^N$. The Dirac or *ket* notation $|\mathbf{x}\rangle$ highlights the fact that a quantum state vector is always of unit norm, so that $\|\mathbf{x}\|_2^2 = \mathbf{x}^H \mathbf{x} = \langle \mathbf{x} | \mathbf{x} \rangle = \sum_{k=0}^{N-1} |x_k|^2 = 1$, where the *bra* notation $\langle \mathbf{x} |$ denotes $\langle \mathbf{x} | = \mathbf{x}^H$. This vector captures the PMF related to the quantum system, in the sense that $|x_k|^2$ is the probability of sampling the bitstring corresponding to the binary representation of integer $k = 0, \dots, N-1$.

The second postulate predicates that the only operation to be applied on a quantum state is a unitary matrix $\mathbf{U} \in \mathbb{C}^{N \times N}$ to transform state $|\mathbf{x}\rangle$ to $|\mathbf{y}\rangle = \mathbf{U} |\mathbf{x}\rangle$. Not all unitaries can be efficiently implemented in quantum hardware. We can




Figure 1. A typical architecture of a VQC operating on an $\log N$ qubits. *Left*: An L -layered hardware-efficient VQC [21, Ch. 5]. *Middle*: A parameterized layer of the hardware-efficient VQC, where $\mathbf{R}_\sigma$ denotes single-qubit Pauli rotation gates ($\mathbf{R}_X, \mathbf{R}_Y, \mathbf{R}_Z$). *Right*: One cyclic entanglement layer comprises $\log N$ two-qubit CNOT gates.

typically implement unitaries operating on one or two qubits. Some examples of single-qubit unitaries are:

$$\mathbf{X} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{R}_X(\theta) = \begin{bmatrix} \cos(\theta) & -i \sin(\theta) \\ -i \sin(\theta) & \cos(\theta) \end{bmatrix}. \quad (1)$$

Matrix $\mathbf{X}$ implements the quantum NOT gate, and $\mathbf{R}_X(\theta)$ is an example of a parameterized unitary, where $\theta \in [0, 2\pi]$ is controlled externally.

According to the third postulate, the output of a quantum state can be associated with a Hermitian matrix $\mathbf{H}$ to provide the expectation over a *quantum observable* as the quadratic function

$$F = \langle \mathbf{x} | \mathbf{H} | \mathbf{x} \rangle = \mathbf{x}^H \mathbf{H} \mathbf{x}.$$

If $\mathbf{H}$ is diagonal, then obviously $F = \sum_{k=0}^{N-1} |x_k|^2 H_{kk}$ where H_{kk} are the diagonal entries of $\mathbf{H}$. If a random variable takes values H_{kk} with probability $|x_k|^2$ for all k , then F is essentially the mean of that random variable.

Per the fourth postulate, if a quantum system is described by $|\mathbf{x}_1\rangle$ and a second one by $|\mathbf{x}_2\rangle$, the joint system is characterized by state $|\mathbf{x}_1, \mathbf{x}_2\rangle = \mathbf{x}_1 \otimes \mathbf{x}_2$, where $\otimes$ is the Kronecker product.

The *variational quantum eigensolver* (VQE) is a computational task tailored to NISQ systems. Given an $N \times N$ Hermitian $\mathbf{H}_0$, VQE aims to find its minimum eigenvalue and associated eigenvector. Because N can be large, VQE captures the sought eigenvector as a parameterized quantum state $|\psi(\theta)\rangle$ using $\log N$ qubits, and solves the related *variational* eigenproblem:

$$\min_{\theta} F_0(\theta) := \langle \psi(\theta) | \mathbf{H}_0 | \psi(\theta) \rangle \quad (2)$$

over a vector of parameters $\theta \in \mathbb{R}^P$, where typically P scales polynomially with $\log N$. The state $|\psi(\theta)\rangle$ is prepared by feeding $|\mathbf{0}\rangle$ into a *variational quantum circuit* (VQC) as

$$|\psi(\theta)\rangle = \mathbf{V}(\theta) |\mathbf{0}\rangle$$

where $|\mathbf{0}\rangle$ is the first column of the identity $\mathbf{I}_N$, and $\mathbf{V}(\theta)$ is a parameterized unitary describing the VQC. The VQC consists of layers of parameterized single-qubit gates and fixed two-qubit gates, shown in green and orange, respectively, in Fig. 1.

The VQE optimization in (2) is solved in a hybrid classical/quantum fashion. The VQC evaluates the objective function $F_0(\theta)$ and its gradient, vector θ is updated using gradient

descent (GD) updates on a classical computer. Critically, for a broad family of VQC architectures, the gradient $\nabla_{\theta} F_0(\theta)$ can be computed by evaluating F_0 at $2P$ additional points. In detail, according to the *parameter shift rule* (PSR), a partial derivative can be evaluated exactly as the finite difference

$$\frac{\partial F_0(\theta)}{\partial \theta_p} = \frac{1}{2} (F_0(\theta + \frac{\pi}{2} \mathbf{e}_p) - F_0(\theta - \frac{\pi}{2} \mathbf{e}_p)) \quad (3)$$

where $\mathbf{e}_p$ is the p -th column of the identity $\mathbf{I}_P$. Building on VQE, we next explain how a single instance of the AC PF task can be solved using a variational quantum approach.

III. SINGLE-INSTANCE QUANTUM POWER FLOW

After briefly reviewing the AC-PF task, we reformulate it as an interesting generalization of the VQE problem. Consider a power system consisting of N buses indexed by $n \in \{1, \dots, N\}$. Let $p_n + jq_n$ and v_n denote the complex power injection and the complex voltage phasor at bus n , respectively. Vector $\mathbf{v} \in \mathbb{C}^N$ collects all voltage phasors. In its standard form, the AC power flow problem entails solving the AC PF equations to find $\mathbf{v}$ given two specifications per bus. Depending on its specifications, a bus can belong to one of the following three types: *i*) One generator bus is selected as the reference or slack bus, for which we specify $|v_n|$ and set its phase angle to zero; *ii*) For the remaining generator buses (also known as PV buses), we specify $(p_n, |v_n|)$ or, equivalently, $(p_n, |v_n|^2)$. We will henceforth assume the latter for convenience. *iii*) For load buses (also known as PQ buses), we specify (p_n, q_n) . Because the AC PF equations are invariant with respect to a phase shift in $\mathbf{v}$, nulling the voltage phase angle of the reference bus is unnecessary. We are thus left with $S = 2N - 1$ real-valued specifications. Let b_s denote the value of specification s , and collect specifications in $\mathbf{b} \in \mathbb{R}^S$.

Solving the AC-PF problem is equivalent to solving the following set of S quadratic equations over $\mathbf{v}$ [4], [22]

$$\mathbf{v}^H \mathbf{H}_s \mathbf{v} = b_s, \quad s = 1, \dots, S \quad (4)$$

where $\mathbf{H}_s$ are known sparse Hermitian matrices that depend on the bus admittance matrix $\mathbf{Y}$. Grid operators solve (4) routinely under different loading conditions and generator settings. As long as $\mathbf{Y}$ remains unaltered, an instance of the AC-PF is characterized by its specification vector $\mathbf{b}$.

For a particular PF instance $\mathbf{b}$, a solution is typically found via iterative methods, such as the Gauss-Seidel, the impedance matrix, or the Newton-Raphson method. Alternatively, recent approaches train ML models to predict AC-PF solutions once presented with a $\mathbf{s}$ at their input. The idea is to offload computational cost from real-time to offline, given that the ML model should be trained over a sufficiently representative dataset of AC-PF instances $\{\mathbf{b}_t\}_{t=1}^T$. In lieu of a classical ML model, we suggest predicting AC-PF solutions using a VQC. Before dealing with multiple PF instances, let us first explain how to deal with a single PF instance.

To this end, first note that rather than solving the quadratic equations in (4), a PF solution can be found as a minimizer of the unconstrained optimization

$$\min_{\mathbf{v} \in \mathbb{C}^N} \sum_{s=1}^S (\mathbf{v}^H \mathbf{H}_s \mathbf{v} - b_s)^2. \quad (5)$$

This formulation is also relevant for solving the power system state estimation (PSSE) task, in which case the specifications b_s correspond to noisy measurements of grid quantities in addition to $(p_n, q_n, |v_n|)$, such as apparent currents and line power flows, so that $S \geq 2N - 1$. The objective in (5) is quartic in $\mathbf{v}$ and can be easily shown as non-convex; convex reformulations of (5) for handling the PF and the PSSE tasks have been studied in [4], [22], [23].

Spurred by the VQE idea, it is tempting to capture directly $\mathbf{v}$ by the state vector $|\psi(\theta)\rangle$ of a VQC on $\log N$ qubits, presuming for now that the number of buses N is a power of 2. This mapping does not work as $\mathbf{v}$ is not of unit norm. In fact, because voltage magnitudes are typically in the range of one per unit (pu), we expect $\|\mathbf{v}\|_2$ to be in the order of $\sqrt{N}$. As in [24], we can introduce a classical optimization variable $\alpha > 0$ and model the power system state variationally as

$$\mathbf{v}(\theta) = \sqrt{\alpha} |\psi(\theta)\rangle \quad (6)$$

so that $\alpha = \|\mathbf{v}(\theta)\|_2^2$. If we define the expectation $F_s(\theta) := \langle \psi(\theta) | \mathbf{H}_s | \psi(\theta) \rangle$, a PF specification can be expressed in its variational form as

$$(\mathbf{v}(\theta))^H \mathbf{H}_s \mathbf{v}(\theta) = \alpha \langle \psi(\theta) | \mathbf{H}_s | \psi(\theta) \rangle = \alpha F_s(\theta). \quad (7)$$

Based on (7), we can attempt solving the AC-PF variationally over (θ, α) instead of $\mathbf{v}$ as

$$\min_{\theta, \alpha > 0} f(\theta, \alpha) = \sum_{s=1}^S (\alpha F_s(\theta) - b_s)^2. \quad (\text{QPF})$$

However, unlike VQE, the objective in (QPF) is not a linear function of a single expectation over a quantum observable. It is rather the sum of a large number S of *squared* expectation values. These two issues render measuring the objective and its gradient on a VQC non-straightforward. To address these challenges, we leverage the special structure of (QPF) and introduce a novel quantum measuring protocol.

The presentation of this measurement protocol is postponed to Section V. For now, it suffices to assume that we can readily evaluate the gradients of $f(\theta, \alpha)$ on a VQC, so that we can find a stationary point of (QPF) using GD updates

$$\begin{aligned} \theta^{k+1} &= \theta^k - \mu \nabla_{\theta} f(\theta^k, \alpha^k) \\ \alpha^{k+1} &= \max \left\{ \alpha^k - \mu \frac{\partial f(\theta^k, \alpha^k)}{\partial \alpha}, 0 \right\} \end{aligned}$$

for a step size $\mu > 0$.

We close this section with a note on zero-padding the grid state. If the size of the power system N is a power of 2, the voltage vector $\mathbf{v}$ can be modeled using $\log N$ qubits. Otherwise, we need $\lceil \log N \rceil$ qubits and zero-pad matrices $\mathbf{H}_s$




Figure 2. A data-embedding VQC includes three main blocks: the data embedding block, the trainable block, and the measurement block. Data $\mathbf{b} \in \mathbb{R}^S$ can be embedded as parameters of S parameterized single-qubit gates (colored in blue), arranged into $d = S/\log N$ layers.

so they become of size $2^{\lceil \log N \rceil} \times 2^{\lceil \log N \rceil}$. In the latter case, the $2^{\lceil \log N \rceil}$ -dimensional quantum state $|\psi(\theta)\rangle$ is partitioned into the first N entries comprising subvector $\psi_1(\theta)$, and the remaining ones comprising $\psi_2(\theta)$. We now have that

$$\mathbf{v}(\theta) = \alpha \psi_1(\theta) \quad \text{and} \quad F_s(\theta) = \alpha (\psi_1(\theta))^H \mathbf{H}_s \psi_1(\theta).$$

However, scaling is not as straightforward as before because

$$\|\mathbf{v}(\theta)\|_2^2 = \alpha \|\psi_1(\theta)\|_2^2 = \alpha(1 - \|\psi_2(\theta)\|_2^2)$$

given that $|\psi(\theta)\rangle$ has unit norm. Subvector $\psi_2(\theta)$ does not contribute to any expectation as it is multiplied by the zero-padded part of the augmented $\mathbf{H}_s$, yet its norm may still be considerable. If $\|\psi_2(\theta)\|_2$ is large, then $\|\psi_1(\theta)\|_2$ gets small and challenges the quantum measurement process. To avoid solutions where a small $\|\psi_1(\theta)\|_2$ is compensated by a large value of α , we add the constraint $\alpha \leq 1.1^2 N$. This is a safe upper bound on α as voltages are typically less than 1.1 pu.

IV. DATA-BASED QUANTUM POWER FLOW

So far, we have explained how the VQC parameters θ can be optimized to solve a single PF instance. In Section III, the vector $\mathbf{b}$ of PF specifications appeared only in the objective of (QPF). This section embeds PF specifications into the VQC state and trains the VQC parameters θ based on multiple PF instances $\{\mathbf{b}_t\}_{t=1}^T$. In classical ML models trained to predict PF solutions, vector $\mathbf{b}$ is treated as the input or feature vector. Analogously, our goal is to make the VQC state dependent on data and parameters alike as $|\psi(\mathbf{b}, \theta)\rangle$.

To accomplish this goal, we propose using a *data-embedded VQC*; see [16, Ch. 5]. This type of VQC comprises three blocks, as shown in Fig. 2: a data-embedding block, a trainable block, and the measurement layers. The last two are identical to those described in Section III. In the QML context, the trainable parameters $\mathbf{b}$ are analogous to the weights of classical ML models. The data embedding block inputs data into the QML model. Similarly to VQC weights θ , VQC data $\mathbf{b}$ are loaded into the VQC as parameters of another unitary $\mathbf{W}(\mathbf{b})$.

With two blocks of parameterized unitaries, the VQC state is now described as

$$|\psi(\theta, \mathbf{b})\rangle = \mathbf{V}(\theta) \mathbf{W}(\mathbf{b}) |0\rangle$$

so the expectation $F_s(\theta, \mathbf{b}) = \langle \psi(\theta, \mathbf{b}) | \mathbf{H}_s | \psi(\theta, \mathbf{b}) \rangle$ becomes a nonlinear function of θ and $\mathbf{b}$ as with DNNs [25]. Since this function is trigonometric, PF data are normalized within $[0, 2\pi]$. Both $\mathbf{V}(\theta)$ and $\mathbf{W}(\mathbf{b})$ operate on $\log N$ qubits. To achieve this, we partition $\mathbf{b}$ into $\log N$ subvectors, each of dimension $d = \lceil S/\log N \rceil$. Each subvector is embedded into d gates operating on the same qubit, as shown in Fig. 2. The entire data-embedding block can be repeated several times to enhance the representation capability of the VQC [25]. If S is not divisible by $\log N$, vector $\mathbf{b}$ is padded with zeros. Note that parameterized single-qubit gates, such as $\mathbf{R}_X(\theta)$ in (1), when fed with 0, become identity operators and thus act trivially on qubits.

Having embedded PF specifications into the VQC state, the VQC can now be trained in an unsupervised fashion to learn the solutions of multiple PF instances. More specifically, the optimal VQC parameters (θ, α) can be found by minimizing the ensuing loss function using gradient descent:

$$\min_{\theta, \alpha > 0} \sum_{t=1}^T \sum_{s=1}^S (\alpha F_s(\theta, \mathbf{b}_t) - b_t^s)^2 \quad (8)$$

where b_t^s is the s -th entry of $\mathbf{b}_t$. Mini-batch gradients can be implemented by sampling a few instances $\mathbf{b}_t$ per iteration [26].

V. EFFICIENT QUANTUM POWER FLOW MEASUREMENTS

In this section, we propose a novel approach for efficiently measuring the PF objective and its gradient using a VQC. The proposed protocol applies to QPF on a per-instance basis, so we drop dependence on $\mathbf{b}_t$. Upon ignoring inconsequential constants, the cost in (QPF) can be expanded as

$$f(\theta, \alpha) = \alpha^2 \sum_{s=1}^S F_s^2(\theta) - 2\alpha \sum_{s=1}^S b_s F_s(\theta). \quad (9)$$

Leaving the α factor aside, the second summand can be expressed as a single expectation

$$G(\theta) = \langle \psi(\theta) | \mathbf{G} | \psi(\theta) \rangle \quad \text{where} \quad \mathbf{G} = \sum_{s=1}^S b_s \mathbf{H}_s. \quad (10)$$

Can we efficiently measure $G(\theta)$? One option would be to express $\mathbf{G}$ as a linear combination of unitaries (LCU) [16, p. 138]. An LCU-based approach is practical only if the unitaries can be efficiently implemented on a quantum computer and their number scales polynomially in $\log N$. We conjecture that neither of these conditions is met.

Another idea would be trying the eigenvalue decomposition of $\mathbf{G} = \mathbf{U}_g \mathbf{\Lambda}_g \mathbf{U}_g^H$. We can apply the unitary $\mathbf{U}_g^H$ on the VQC state and measure the derived state $\mathbf{U}_g^H |\psi(\theta)\rangle$ using the diagonal matrix $\mathbf{\Lambda}_g$. This approach requires computing an eigenvalue decomposition per PF instance, and matrix $\mathbf{U}_g^H$ is unlikely to be efficiently realizable.

In our previous work [19], we adapted the so-called extended Bell measurement (XBM) protocol of [18] to efficiently measure observables to solve the OPF on a VQC. We briefly review the key ideas and extend them to the PF setting.

Lemma 1 ([18]). *Given the admittance matrix $\mathbf{Y}$ of a power system, there exists a set of unitary matrices $\{\mathbf{U}_i\}_{i=1}^C$ so that the specification matrices can be decomposed as*

$$\mathbf{H}_s = \sum_{i=1}^C \mathbf{U}_i \mathbf{\Lambda}_i^s \mathbf{U}_i^H \text{ for } s = 1, \dots, S$$

where

- i) Matrices $\mathbf{\Lambda}_i^s$ are diagonal, real-valued, and can be easily computed. The number of non-zero entries scales with the maximum degree of the power system graph, that is $\mathcal{O}(1)$; and
- ii) The unitary matrices $\mathbf{U}_i^H$ are qubit efficient, i.e., they can be realized using at most $\log N + 1$ single- and two-qubit gates.

The critical feature of Lemma 1 is that all $S = 2N - 1$ specification matrices can be decomposed using the same C unitaries. Numerical tests using all power system benchmarks in the `pglib` database indicate that the number of unitaries C scales polynomially in $\log N$.

Based on Lemma 1, matrix $\mathbf{G}$ can be decomposed as

$$\begin{aligned} \mathbf{G} &= \sum_{s=1}^S b_s \mathbf{H}_s = \sum_{s=1}^S b_s \left(\sum_{i=1}^C \mathbf{U}_i \mathbf{\Lambda}_i^s \mathbf{U}_i^H \right) \\ &= \sum_{i=1}^C \mathbf{U}_i \left(\underbrace{\sum_{s=1}^S b_s \mathbf{\Lambda}_i^s}_{:=\mathbf{\Lambda}_i(\mathbf{b})} \right) \mathbf{U}_i^H = \sum_{i=1}^C \mathbf{U}_i \mathbf{\Lambda}_i(\mathbf{b}) \mathbf{U}_i^H. \end{aligned}$$

Thanks to this decomposition, the expectation in (10) can be measured as the sum of C expectation values:

$$\begin{aligned} G(\boldsymbol{\theta}) &= \sum_{i=1}^C \langle \psi(\boldsymbol{\theta}) | \mathbf{U}_i \mathbf{\Lambda}_i(\mathbf{b}) \mathbf{U}_i^H | \psi(\boldsymbol{\theta}) \rangle \\ &= \sum_{i=1}^C \underbrace{\langle \psi_i(\boldsymbol{\theta}) | \mathbf{\Lambda}_i(\mathbf{b}) | \psi_i(\boldsymbol{\theta}) \rangle}_{:=G_i(\boldsymbol{\theta}; \mathbf{b})} = \sum_{i=1}^C G_i(\boldsymbol{\theta}; \mathbf{b}). \end{aligned}$$

To measure $G_i(\boldsymbol{\theta}; \mathbf{b})$, apply unitary $\mathbf{U}_i^H$ on the VQC state to generate $|\psi_i(\boldsymbol{\theta})\rangle = \mathbf{U}_i^H |\psi(\boldsymbol{\theta})\rangle$, and measure the latter state on $\mathbf{\Lambda}_i(\boldsymbol{\theta})$. The expectations $G_i(\boldsymbol{\theta}; \mathbf{b})$ can be measured using a serial or parallel architecture. The parallel one consists of C replicas of the original VQC, and applies a different unitary at the end of each replica. The serial architecture operates on a single VQC replica and appends a different unitary each time, iterating over C . The parallel architecture is C times faster but requires C times more quantum hardware.

Let us focus on the first summand in (9). Although this term does not depend on PF specifications, it is a quartic function of $|\psi(\boldsymbol{\theta})\rangle$. To address this, we propose using two replicas of the VQC to generate the joint quantum state

$$|\tilde{\psi}(\boldsymbol{\theta})\rangle = |\psi(\boldsymbol{\theta})\rangle \otimes |\psi(\boldsymbol{\theta})\rangle.$$

Note that although the two VQCs have identical states, when sampled, they produce different binary strings due to the random nature of the quantum circuit. The two states are

Algorithm 1 Measure expectation $\tilde{G}(\boldsymbol{\theta})$

Initialize: $\tilde{G}(\boldsymbol{\theta}) = 0$.

- 1: **for** $i = 1$ to C **do**
 - 2: **for** $j = 1$ to C **do**
 - 3: Apply $\mathbf{U}_i^H$ on first VQC to generate $|\psi_i(\boldsymbol{\theta})\rangle = \mathbf{U}_i^H |\psi(\boldsymbol{\theta})\rangle$.
 - 4: Apply $\mathbf{U}_j^H$ on second VQC to generate $|\psi_j(\boldsymbol{\theta})\rangle = \mathbf{U}_j^H |\psi(\boldsymbol{\theta})\rangle$.
 - 5: Measure $|\psi_i(\boldsymbol{\theta})\rangle \otimes |\psi_j(\boldsymbol{\theta})\rangle$ on diagonal $\mathbf{\Lambda}_{ij}$.
 - 6: Add measurement to $\tilde{G}(\boldsymbol{\theta})$.
 - 7: **end for**
 - 8: **end for**
-

identical because the two VQCs have the same topologies and are fed with the same parameters $\boldsymbol{\theta}$.

Dropping the α^2 factor, the dependence on $\boldsymbol{\theta}$, and the ket notation for simplicity, we write the first summand in (9) as

$$\begin{aligned} \tilde{G} &= \sum_{s=1}^S F_s^2 = \sum_{s=1}^S (\psi^H \mathbf{H}_s \psi) (\psi^H \mathbf{H}_s \psi) \\ &= \sum_{s=1}^S (\psi \otimes \psi)^H (\mathbf{H}_s \otimes \mathbf{H}_s) (\psi \otimes \psi) \\ &= \tilde{\psi}^H \left(\underbrace{\sum_{s=1}^S \mathbf{H}_s \otimes \mathbf{H}_s}_{:=\tilde{\mathbf{G}}} \right) \tilde{\psi} = \tilde{\psi}^H \tilde{\mathbf{G}} \tilde{\psi}. \end{aligned}$$

Therefore, the term $\tilde{G}(\boldsymbol{\theta})$ is an expectation over the joint state $|\tilde{\psi}(\boldsymbol{\theta})\rangle$ with associated Hermitian matrix $\tilde{\mathbf{G}}$.

Leveraging Lemma 1, matrix $\tilde{\mathbf{G}}$ can be decomposed as

$$\begin{aligned} \tilde{\mathbf{G}} &= \sum_{s=1}^S \left(\sum_{i=1}^C \mathbf{U}_i \mathbf{\Lambda}_i^s \mathbf{U}_i^H \right) \otimes \left(\sum_{j=1}^C \mathbf{U}_j \mathbf{\Lambda}_j^s \mathbf{U}_j^H \right) \\ &= \sum_{i=1}^C \sum_{j=1}^C (\mathbf{U}_i \otimes \mathbf{U}_j)^H \left(\underbrace{\sum_{s=1}^S \mathbf{\Lambda}_i^s \otimes \mathbf{\Lambda}_j^s}_{:=\mathbf{\Lambda}_{ij}} \right) (\mathbf{U}_i \otimes \mathbf{U}_j) \\ &= \sum_{i=1}^C \sum_{j=1}^C (\mathbf{U}_i \otimes \mathbf{U}_j)^H \mathbf{\Lambda}_{ij} (\mathbf{U}_i \otimes \mathbf{U}_j). \end{aligned}$$

Based on this decomposition, the term $\tilde{G}(\boldsymbol{\theta})$ can be measured using Algorithm 1. Because we need to measure all unique pairs of unitaries, the two VQCs are recompiled $C(C+1)/2$ times. It is important to note that measuring pairs of unitaries is required only during training.

VI. NUMERICAL TESTS

The proposed QPF framework was evaluated using the IEEE 14-bus system with $S = 27$ specifications. To generate multiple PF instances, benchmark PF specifications were perturbed by zero-mean, white Gaussian noise of standard deviation of 0.05 pu for voltage magnitude specifications, and 20% of the nominal values for power injections. VQCs were coded in




Figure 3. VQC architecture used in numerical tests. Blocks in orange represent cyclic entanglement with CNOT gates. For the single-instance and multi-instance QPFs, we used $L = 3$ and $L = 6$ layers, respectively.




Figure 4. Convergence of the normalized mean absolute error (NMAE) for the single-instance QPF across gradient descent iterations. The plot shows the confidence intervals within one standard deviation around the mean, computed over 10 PF problem instances.

Python using PennyLane’s exact quantum simulator [27]. The VQC operated on 4 qubits with matrices zero-padded to reach the dimension of 16. All tests used the VQC architecture of Fig. 3. VQC weights θ were initialized so the related quantum state corresponds to the flat voltage profile, using the method of [28]. Variable α was initialized to $\sqrt{N}$. Gradient descent iterations were terminated when the Euclidean norm of the gradient was smaller than 0.01.

We first tested the performance of the single-instance OPF in terms of the normalized mean absolute error (NMAE) $\|\hat{\mathbf{b}} - \mathbf{b}\|_1 / \|\mathbf{b}\|_1$, where vector $\hat{\mathbf{b}}$ evaluates the PF specifications based on the obtained voltage vector. We ran the single-instance QPF on 10 instances and $P = 28$ trainable parameters. Step sizes were set as μ_θ and μ_α were set to 5×10^{-5} across iterations. Figure 4 shows that the relative error $\|\hat{\mathbf{b}} - \mathbf{b}\|_1 / \|\mathbf{b}\|_1$ converges to a point around 5% with small deviation. This demonstrates that QPF found near-optimal PF solutions across all 10 instances.

We subsequently trained a QML model to learn multiple PF solutions and compared it against a classical DNN model. The DNN architecture was selected based on four-fold cross-validation on 80 PF problem instances upon testing different step sizes. The number of input layers was set to the length of $\mathbf{b}_t$. We used two fully connected hidden layers of dimension 10 each. The DNN outputs the real and imaginary parts of voltages of a total dimension of 32. Rectified linear units (ReLU) were used for all layers. Overall, the chosen DNN architecture used 690 weights. Regarding QML, we used the VQC architecture of Fig. 3 with $L = 6$ layers, resulting in only $P = 52$ trainable weights. Gradient descent step sizes for DNN and VQC/QML were adjusted per the exponentially




Figure 5. Convergence of NMAE attained by VQC/QML and DNN across gradient descent iterations.




Figure 6. Comparing the VQC/QML and DNN regarding their NMAE over 20 testing AC-PF instances.

decaying rule as 0.9999^k and 0.9995^k at iteration k , respectively. Figure 5 shows that the VQC/QML achieved faster convergence and smaller NMAE than DNN. Correspondingly, the VQC/QML model attained smaller prediction errors than DNN in 17 out of 20 instances, as shown in Fig. 6. This test corroborates that VQC/QML attained better prediction errors and faster training than DNN, while using an order of magnitude fewer trainable parameters.

VII. CONCLUSIONS

To cope with the challenges of solving multiple PFs promptly, this work has explored training a VQC/QML model to predict AC PF solutions. PF specifications are embedded as VQC parameters, allowing data-based VQC/QML to be trained unsupervised. To expedite gradient computations of VQC/QML, the PF problem has been judiciously reformulated via a few expectations, whose gradients can be measured efficiently on VQCs. Numerical tests on the IEEE 14-bus system have shown that the proposed data-based QPF achieved improved prediction error while using fewer trainable parameters (weights) than a classical DNN. Capitalizing on these results, several exciting directions are opened for future research: d1) skip the unitary pairs in training VQC/QML; d2) design grid-informed VQC topologies; and d3) explore other VQC data encoding strategies.

REFERENCES

- [1] W. F. Tinney and C. E. Hart, "Power flow solution by Newton's method," *IEEE Trans. Power App. Syst.*, vol. PAS-86, no. 11, pp. 1449–1460, Nov 1967.
- [2] A. Gómez-Expósito, A. J. Conejo, and C. Canizares, Eds., *Electric Energy Systems, Analysis and Operation*. Boca Raton, FL: CRC Press, 2009.
- [3] M. Bazrafshan and N. Gatsis, "Convergence of the Z-bus method for three-phase distribution load-flow with ZIP loads," *IEEE Trans. Power Syst.*, vol. 33, no. 1, pp. 153–165, Jan. 2018.
- [4] R. Madani, J. Lavaei, and R. Baldick, "Convexification of power flow equations in the presence of noisy measurements," *IEEE Trans. Autom. Contr.*, vol. 64, no. 8, pp. 3101–3116, Feb 2019.
- [5] B. Sævarsson, S. Chatzivasileiadis, H. Jóhannsson, and J. Østergaard, "Quantum computing for power flow algorithms: Testing on real quantum computers," *arXiv preprint arXiv:2204.14028*, Apr 2022.
- [6] F. Feng, Y. Zhou, and P. Zhang, "Quantum power flow," *IEEE Trans. Power Syst.*, vol. 36, no. 4, pp. 3810–3812, Jul 2021.
- [7] J. Liu, H. Zheng, M. Hanada, K. Setia, and D. Wu, "Quantum power flows: From theory to practice," *Quantum Machine Intelligence*, vol. 6, no. 2, p. 55, Dec 2024.
- [8] P. Pareek, A. Jayakumar, C. Coffrin, and S. Misra, "Demystifying quantum power flow: Unveiling the limits of practical quantum advantage," *arXiv preprint arXiv:2402.08617*, Feb 2024.
- [9] M. J. Reno, J. Deboever, and B. Mather, "Motivation and requirements for quasi-static time series (qsts) for distribution system analysis," in *Proc. IEEE Power & Energy Society General Meeting*, Chicago, IL, Jul 2017, pp. 1–5.
- [10] X. Hu, H. Hu, S. Verma, and Z.-L. Zhang, "Physics-guided deep neural networks for power flow analysis," *IEEE Trans. Power Syst.*, vol. 36, no. 3, pp. 2082–2092, May 2021.
- [11] L. Guo, Y. Zhang, X. Li, Z. Wang, Y. Liu, L. Bai, and C. Wang, "Data-driven power flow calculation method: A lifting dimension linear regression approach," *IEEE Trans. Power Syst.*, vol. 37, no. 3, pp. 1798–1808, May 2022.
- [12] M. Gao, J. Yu, Z. Yang, and J. Zhao, "Physics embedded graph convolution neural network for power flow calculation considering uncertain injections and topology," *IEEE Trans. Neural Netw.*, vol. 35, no. 11, pp. 15 467–15 478, Nov 2024.
- [13] J. Jalving, M. Eydenberg, L. Blakely, A. Castillo, Z. Kilwein, J. K. Skolfield, F. Boukouvala, and C. Laird, "Physics-informed machine learning with optimization-based guarantees: Applications to ac power flow," *International Journal of Electrical Power & Energy Systems*, vol. 157, p. 109741, June 2024.
- [14] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019.
- [15] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, "The power of quantum neural networks," *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, Jun 2021.
- [16] M. Schuld and F. Petruccione, *Machine learning with quantum computers*. Cham, Switzerland: Springer, 2021.
- [17] Z. Yu, Q. Chen, Y. Jiao, Y. Li, X. Lu, X. Wang, and J. Yang, "Non-asymptotic approximation error bounds of parameterized quantum circuits," *Advances in Neural Information Processing Systems*, vol. 37, pp. 99 089–99 127, Sep 2024.
- [18] R. Kondo, Y. Sato, S. Koide, S. Kajita, and H. Takamatsu, "Computationally efficient quantum expectation with extended Bell measurements," *Quantum*, vol. 6, p. 688, Apr 2022.
- [19] T. V. Le, M. M. Wilde, and V. Kekatos, "Solving optimal power flow using a variational quantum computer," Jun. 2025, in preparation: <https://engineering.purdue.edu/kekatos/QOPF.pdf>.
- [20] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge, UK: Cambridge University Press, 2000.
- [21] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, "Evaluating analytic gradients on quantum hardware," *Phys. Rev. A*, vol. 99, no. 3, p. 032331, Nov 2019.
- [22] H. Zhu and G. B. Giannakis, "Power system nonlinear state estimation using distributed semidefinite programming," *IEEE J. Sel. Topics Signal Process.*, vol. 8, no. 6, pp. 1039–1050, Dec. 2014.
- [23] V. Kekatos, G. Wang, H. Zhu, and G. B. Giannakis, "PSSE redux: Convex relaxation, decentralized, robust, and dynamic approaches," in *Advances in Power System State Estimation*, M. El-Hawary, Ed. Wiley, 2021, pp. 173–202.
- [24] Y. Sato, R. Kondo, S. Koide, H. Takamatsu, and N. Imoto, "Variational quantum algorithm based on the minimum potential energy for solving the Poisson equation," *Physical Review A*, vol. 104, no. 5, p. 052409, Nov 2021.
- [25] M. Schuld, R. Sweke, and J. J. Meyer, "Effect of data encoding on the expressive power of variational quantum-machine-learning models," *Physical Review A*, vol. 103, no. 3, p. 032430, Mar 2021.
- [26] R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. Fährmann, B. Meynard-Piganeau, and J. Eisert, "Stochastic gradient descent for hybrid quantum-classical optimization," *Quantum*, vol. 4, p. 314, Aug 2020.
- [27] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, M. S. Alam, S. Ahmed, J. M. Arrazola, C. Blank, A. Delgado, S. Jahangiri, K. McKiernan, J. J. Meyer, Z. Niu, A. Szava, and N. Killoran, "PennyLane: Automatic differentiation of hybrid quantum-classical computations," 2022. [Online]. Available: <https://arxiv.org/abs/1811.04968>
- [28] C. Bravo-Prieto, R. LaRose, M. Cerezo, Y. Subasi, L. Cincio, and P. J. Coles, "Variational Quantum Linear Solver," *Quantum*, vol. 7, p. 1188, Nov 2023.