Solving Optimal Power Flow using a Variational Quantum Approach

Thinh Viet Le[1], Mark M. Wilde[2], and Vassilis Kekatos[1]

[1]Elmore Family School of Electrical and Computer Engineering, Purdue University, West Lafayette, Indiana 47906, USA
[2]School of Electrical and Computer Engineering, Cornell University, Ithaca, New York 14853, USA

The optimal power flow (OPF) is a large-scale optimization problem that is central in the operation of electric power systems. Although it can be posed as a nonconvex quadratically constrained quadratic program, the complexity of modern-day power grids raises scalability and optimality challenges. In this context, this work proposes a variational quantum paradigm for solving the OPF. We encode primal variables through the state of a parameterized quantum circuit (PQC), and dual variables through the probability mass function associated with a second PQC. The Lagrangian function can thus be expressed as scaled expectations of quantum observables. An OPF solution can be found by minimizing/maximizing the Lagrangian over the parameters of the first/second PQC. We pursue saddle points of the Lagrangian in a hybrid fashion. Gradients of the Lagrangian are estimated using the two PQCs, while PQC parameters are updated classically using a primal-dual method. We propose permuting primal variables so that OPF observables are expressed in a banded form, allowing them to be measured efficiently. Numerical tests on the IEEE 57-node power system using Pennylane's simulator corroborate that the proposed doubly variational quantum framework can find high-quality OPF solutions. Although showcased for the OPF, this framework features a broader scope, including conic programs with numerous variables and constraints, problems defined over sparse graphs, and training quantum machine learning models to satisfy constraints.

Contents

Thinh Viet Le: le272@purdue.edu
Mark M. Wilde: wilde@cornell.edu
Vassilis Kekatos: kekatos@purdue.edu

1 Introduction

Optimal power flow (OPF) is a fundamental problem in the operation of electric power systems. It aims to determine the most economical, reliable, and secure schedule for generators to meet electric power demand while

adhering to engineering limitations. Given the increasing load growth and amid a multi-faceted energy transition, grid operators must solve the OPF to optimally operate power systems of ever-increasing size and complexity at finer spatiotemporal scales.

OPF can be posed as an optimization problem over thousands of variables and constraints. If power flows need to be modeled precisely through the alternating-current (AC) power flow equations, the OPF can be formulated as a quadratically-constrained quadratic program (QCQP). Although this QCQP is not convex, the OPF can be relaxed to a convex semidefinite program (SDP) and solved to global optimality under certain conditions [1, 2]. Nonetheless, classical algorithms for solving SDPs [3], such as interior-point methods, may not scale well for large-scale OPF instances. Moreover, the SDP relaxation may not be always successful. In this context, and leveraging the success of variational quantum algorithms (VQAs) in other computational tasks, this work explores VQA alternatives to SDP solvers for addressing the OPF problem.

1.1 Contributions

The contribution of our paper is on three fronts:

c1) Modeling of the OPF using parameterized quantum circuits (PQC). The OPF can be posed as a non-convex QCQP. Both the number of variables and the number of constraints scale with the size of the power system. Given that modern power systems may consist of thousands of nodes, OPF is a large-scale optimization problem. To cope with the high dimensionality, we propose representing the OPF variables using two PQCs. Primal optimization variables are captured upon scaling the state of a primal PQC. Dual variables or Lagrange multipliers corresponding to constraints are captured by scaling the probability mass function (PMF) associated with a dual PQC. Due to this modeling, the Lagrangian function associated with the OPF can be expressed as an expectation over quantum observables operating on either or both PQCs. It is worth emphasizing that, although showcased for the OPF problem, the proposed modeling applies to other QCQP or more general problems that can be expressed in terms of expectations of quantum observables.

c2) Optimizing a doubly parameterized Lagrangian function. To find approximate primal/dual OPF solutions, we propose a VQA seeking a saddle point of the associated Lagrangian function. In particular, we aim to minimize the Lagrangian function with respect to the parameters of the primal PQC and maximize it with respect to the parameters of the dual PQC. This is accomplished in a hybrid fashion. The two PQCs are used to measure gradients of the Lagrangian with respect to PQC parameters. A classical computer subsequently uses the gradient information to update PQC parameters iteratively based on the primal-dual method. The proposed primal-dual updates apply directly to a broader class of optimization problems that can be solved via VQAs. Moreover, the proposed method unlocks potential towards training quantum machine learning (QML) models to comply with constraints. Such functionality is essential for endowing QML models with safety and stability features.

c3) Efficient measurement of OPF observables. Although OPF observables involve sparse Hermitian matrices, standard LCU-based decomposition methods do not yield qubit-efficient measurement protocols. To address this challenge, we utilize the structure of the OPF problem and the recently proposed method of extended Bell measurements (XBM) [4], to design a qubit-efficient protocol for measuring gradients of the Lagrangian function associated with the OPF. Specifically, the XBM method partitions the entries of an $N \times N$ Hermitian measurement matrix into groups, allowing entries within each group to be measured simultaneously using at most $\mathcal{O}(\log N)$ additional quantum gates. If a matrix can be decomposed into $\tilde{\mathcal{O}}(\log N)$ such groups, the XBM method is qubit-efficient. Two key points for successfully adopting the XBM method to the OPF are: *i)* The sparsity pattern (the positions of nonzero entries) of OPF observables is determined by the graph defined by the nodes and transmission lines of the power system, which is sparse; and *ii)* According to numerical tests, nodes in a power system can be permuted so that OPF observables can be partitioned in $\mathcal{O}((\log N)^3)$ groups. Finding the node permutation that yields the number of groups is an NP-hard problem [5]. Instead, we adopt the reverse Cuthill-McKee (RCM) algorithm [6], a greedy heuristic algorithm of linear complexity $\mathcal{O}(N)$, to permute nodes so that OPF observables are banded. This is effective because banded matrices are known to feature a small number of groups [4]. Beyond the OPF setting, the devised workflow can be instrumental in effectively measuring quantum observables associated with sparse connectivity graphs arising in ML and optimization applications over natural or engineered networked systems.

1.2 Literature review

Quantum computing for the OPF. A previous quantum computing effort to the OPF problem has focused on the classical primal-dual interior point method, wherein the Harrow–Hassidim–Lloyd (HHL) algorithm serves as a linear system solver in each Newton step [7]. The idea of utilizing the HHL algorithm as a linear system solver has also been explored for solving the nonlin-

ear AC power flow equations, which constitute a subset of constraints of the AC OPF problem [8, 9, 10]. Nonetheless, the computational complexity of the HHL algorithm does not scale well for linear systems involving the Jacobian matrix of the AC power flow equations, as established in [11]. Reference [12] integrates a PQC as an intermediate layer of a classical neural network, trained in a supervised fashion to predict OPF solutions. However, it is not clear which measurement outputs are encoded into the subsequent classical hidden layer, nor how the parameter-shift rule is applied under such a design.

Handling constrained optimization problems using VQAs. VQAs have emerged as a promising approach for solving optimization problems on current-day quantum computers. A well-studied exemplar of VQAs is the variational quantum eigensolver (VQE), which aims to find the smallest eigenvalue of a quantum observable by employing a PQC [13]. The quantum approximate optimization algorithm (QAOA) is a variant of VQE with a problem-dependent PQC and a diagonal observable targeting combinatorial problems [14]. While VQE and QAOA were originally devised for unconstrained problems, extensions to optimization problems with constraints have been proposed in [15, 16, 17, 18, 19, 20, 21, 22]. References [15, 19] convert constrained problems to unconstrained ones by penalizing constraint violations on the objective function. However, selecting a suitable penalty parameter is a nontrivial task. Unreasonably large values can cause ill-conditioning, while unreasonably small values can yield infeasible solutions. In [20], the QAOA's ansatz is adapted to ensure that the output state stays within the feasible subspace; yet, this strategy applies only to binary problems with linear constraints. References [21] propose a hybrid method for solving SDPs. The idea is to approximate the matrix variable of the original SDP by a semidefinite matrix of smaller dimension. The resultant SDP of reduced dimension is solved classically using standard interior point-based methods. Quantum measurements are used to translate the cost and constraint functions of the original to the compressed SDP formulation. Nonetheless, the solution of the compressed SDP may be infeasible or suboptimal for the original SDP.

An alternative approach to cope with constrained problems is via Lagrangian duality. Reference [16] seeks a saddle point of the Lagrangian function with respect to the PQC parameters and the dual variables. To tackle problems with numerous constraints, reference [18] encodes dual variables through a parameterized quantum state, in addition to the primal variables. Both [16] and [18] solve the associated dual problem using the iterative method of dual decomposition. Nevertheless, each iteration of dual decomposition involves solving a VQE-type problem to optimality, which is computationally overwhelming. In [17], this drawback is circumvented by using the primal-dual update method, which only requires estimating gradients of the Lagrangian function. Nonetheless, reference [17] applies only to binary and linear programs. In this work, we combine [17] and [18] to deal with a doubly variational VQA formulation of the OPF.

Measuring quantum observables. Measuring quantum observables is key to the effective implementation of VQAs. A prominent measurement method decomposes a Hermitian matrix into a linear combination of unitaries (LCU), and measures the resultant unitary observables via the destructive swap test [23, 24]. This method is practically relevant only if the resultant unitaries can be implemented efficiently using quantum gates. This is the case, for example, if the unitaries can be expressed as Pauli strings (tensor products of Pauli matrices), and the number of Pauli strings scales polynomially with the number of qubits, n; see [25, 26, 27]. Unfortunately, a general matrix may involve $\mathcal{O}(4^n)$ Pauli strings [28]. The number of Pauli strings grows exponentially with n, even for some structured sparse matrices, such as k-banded matrices that involve $\mathcal{O}(kn2^n)$ Pauli strings [28]. One possible strategy to reduce the number of measurements is to group Pauli strings into sets of terms that pairwise commute. If two Hermitian matrices commute, they can be jointly diagonalized, and thus, measured simultaneously [29, pg. 43]. To this end, various grouping methods have been proposed, including qubit-wise commutativity grouping [30, 31], general commutativity grouping [32], and unitary partitioning [33, 34]. Nevertheless, partitioning Pauli strings to achieve the minimum number of groups is an NP-hard problem [32].

Classical shadowing is an alternative technique for measuring observables [35]. It is particularly effective when the measurement matrix can be expressed as local Pauli strings, i.e., Pauli strings that act non-trivially on a small number of qubits. For observables that do not admit this form, however, the number of measurements scales linearly with the squared Frobenius norm of the matrix. Unfortunately, for connectivity matrices of sparse graphs, this quantity generally grows linearly with the network size.

Another recent method for measuring observables involves extended Bell measurements [4]. This method decomposes the measurement matrix into matrices whose associated observables are simultaneously measurable using qubit-efficient rotation unitaries. Nevertheless, the number of induced observables may reach 2^n for a general sparse measurement matrix [4]. Our key idea is to permute nodes once and beforehand, so that OPF quantum observables can be measured efficiently.

1.3 Notation

Column vectors are denoted by boldface lower-case letters. Matrices are denoted by boldface upper-case letters. The symbol $^\top$ stands for transposition, and † for conjugate transposition. The operator $\mathrm{dg}(\mathbf{x})$ returns a diagonal matrix with vector $\mathbf{x}$ on its main diagonal, while $\mathrm{dg}(\mathbf{X})$ returns a vector with the diagonal entries of matrix $\mathbf{X}$. The set of real-valued symmetric matrices of dimension M is denoted by $\mathbb{S}^M$. Logarithmic functions are of the binary base. The imaginary unit is denoted by $\imath := \sqrt{-1}$. The notation $m = 1 : M$ means that index m takes values from 1 to M. Symbol $\|\mathbf{H}\|$ denotes the spectral norm of a matrix $\mathbf{H}$.

2 Optimal power flow as a QCQP

The OPF is a large-scale, nonconvex optimization problem. Given forecasts of electric load demands for the next control period, the goal of the OPF is to schedule generators and flexible demand most economically while complying with engineering constraints imposed by the power transmission system and the generation units. This section states the OPF and poses it as a QCQP.

Let us first review a power system model; see [36] and references therein. An electric power system can be represented by an undirected connected graph $\mathcal{G} := \{\mathcal{N}, \mathcal{E}\}$. The vertex set $\mathcal{N} := \{1, \ldots, N\}$ comprises N nodes. Nodes correspond to points where electric power is produced by generators or consumed by loads. The edge set $\mathcal{E} := \{\ell = (n, m) : n, m \in \mathcal{N}\}$ consists of $L_e := |\mathcal{E}|$ transmission lines, each connecting two nodes. Because each node is connected only to a few other nodes, the number of lines, L_e, is typically a small multiple of N.

The power system can be described by a vector $\mathbf{v} \in \mathbb{C}^N$ of the complex AC voltages (voltage phasors) experienced at all nodes. Because many other quantities of interest can be expressed as functions of nodal voltages, the vector $\mathbf{v}$ is typically selected as the state of the power system. For example, the vector of AC currents, $\mathbf{i}$, injected into all nodes can be expressed as $\mathbf{i} = \mathbf{Y}\mathbf{v}$, where $\mathbf{Y} \in \mathbb{C}^{N\times N}$ is the node admittance matrix, which can be thought of as a complex-valued Laplacian matrix of the power system graph $\mathcal{G}$. Other grid quantities, such as the nodal power injections, line power flows, squared voltage magnitudes, and squared line magnitudes, can all be expressed as quadratic functions of $\mathbf{v}$, as $\mathbf{v}^\dagger\mathbf{M}_m\mathbf{v}$ for some Hermitian matrix $\mathbf{M}_m$. If v_n and i_n are the AC voltage and current at node n, respectively, the complex power injected into the power network through node n is defined as $s_n = v_n i_n^* = p_n + \imath q_n$, whose real part p_n is termed active power, and imaginary part q_n is reactive power.

The cost function of the OPF is usually a linear function of some of the active power injections. To comply with generation capacities and load demands, the OPF imposes upper and lower limits on active and reactive power injections in the form of

$$\underline{p}_n \le p_n \le \bar{p}_n \quad \text{and} \quad \underline{q}_n \le q_n \le \bar{q}_n \quad \text{for all } n \in \mathcal{N}.$$

Nodal voltage magnitudes are constrained within upper and lower limits as

$$\underline{v}_n^2 \le |v_n|^2 \le \bar{v}_n^2 \quad \text{for all } n \in \mathcal{N}.$$

Line current magnitudes are upper-limited by line capacities as

$$|i_\ell|^2 \le \bar{i}_\ell^2 \quad \text{for all } \ell \in \mathcal{E}.$$

Since all quantities mentioned above are quadratic in $\mathbf{v}$, the OPF can be posed as a quadratically constrained quadratic program (QCQP) over $\mathbf{v}$ as [2, 36]:

$$P^* := \min_{\mathbf{v}\in\mathbb{C}^N} \ \mathbf{v}^\dagger\mathbf{M}_0\mathbf{v} \tag{QCQP}$$
$$\text{subject to (s.to) } \mathbf{v}^\dagger\mathbf{M}_m\mathbf{v} \le b_m, \quad \text{for } \ m = 1 : M$$

where b_m, for all m, is a given real-valued parameter determined by lower/upper limits and each $\mathbf{M}_m$ is a known Hermitian matrix. For example, if constraint $\underline{p}_n \le p_n$ is indexed as the m-th constraint, then $b_m = -\underline{p}_n$. The OPF involves quadratic equality constraints, e.g., when $\underline{p}_n = \bar{p}_n$ for a subset of nodes $n \in \mathcal{N}$. Each of these equality constraints corresponds to two inequality constraints in (QCQP). Overall, the number of OPF constraints can be several times the network size, such as $M \simeq 8N$.

The precise form of $\mathbf{M}_0$ and $\{(\mathbf{M}_m, b_m)\}_{m=1}^M$ is delineated in Appendix A.1. It is worth stressing that every matrix $\mathbf{M}_m$ is highly sparse. This is because every node of a power system is directly connected to only a few other nodes. The number of nonzero entries per matrix $\mathbf{M}_m$ is upper bounded by a small constant much smaller than N. The (i, j)-th off-diagonal entry across every matrix $\mathbf{M}_m$ is nonzero only if nodes i and j are connected through a transmission line, that is, if $(i, j) \in \mathcal{E}$. Consequently, the sparsity pattern (location of nonzero entries) for each $\mathbf{M}_m$ is determined by a small subset of $\mathcal{E}$. Therefore, the union of the sparsity patterns across all matrices $\mathbf{M}_m$ coincides with the sparsity pattern of the nodal admittance matrix $\mathbf{Y}$. This feature will be crucial for efficiently measuring the quantum observables involved in the OPF, as expounded in Section 6.

3 Primal-dual iterations for the OPF

Although (QCQP) is nonconvex, it can be relaxed to an SDP and solved by standard interior-point solvers.

For many practical settings, this convex relaxation is successful in the sense that a globally optimal $\mathbf{v}$ can be readily recovered from the SDP solution. Nonetheless, the memory and time requirements imposed by SDP solvers may be challenging to meet as N increases. This motivates us to explore a VQA for the OPF.

To motivate a PQC-based solver for large-scale OPFs, let us discuss how (QCQP) could be solved on a classical computer. Among different approaches, we present the method of primal-dual (PD) decomposition as a representative example of a first-order optimization algorithm. The PD method relies on Lagrangian duality [3]. Each constraint indexed by m in (QCQP) is associated with a non-negative dual variable λ_m. Let a vector $\boldsymbol{\lambda} \in \mathbb{R}_+^M$ collect all dual variables. The related *Lagrangian function* is a function of both $\mathbf{v}$ and $\boldsymbol{\lambda}$:

$$\mathcal{L}(\mathbf{v};\boldsymbol{\lambda}) = \mathbf{v}^\dagger \mathbf{M}_0 \mathbf{v} + \sum_{m=1}^{M} \lambda_m (\mathbf{v}^\dagger \mathbf{M}_m \mathbf{v} - b_m). \tag{1}$$

The corresponding *dual problem* is posed as

$$D^* := \max_{\boldsymbol{\lambda} \succeq \mathbf{0}} \min_{\mathbf{v} \in \mathbb{C}^N} \mathcal{L}(\mathbf{v};\boldsymbol{\lambda}). \tag{2}$$

A *primal/dual solution* $(\mathbf{v}^*, \boldsymbol{\lambda}^*) \in \mathbb{C}^N \times \mathbb{R}_+^M$ of (2) satisfies

$$\mathcal{L}(\mathbf{v}^*;\boldsymbol{\lambda}) \leq \mathcal{L}(\mathbf{v}^*;\boldsymbol{\lambda}^*) \leq \mathcal{L}(\mathbf{v};\boldsymbol{\lambda}^*)$$

for all $(\mathbf{v}, \boldsymbol{\lambda}) \in \mathbb{C}^N \times \mathbb{R}_+^M$. Such a pair is called a *saddle point* of the Lagrangian function. In essence, $\mathbf{v}^*$ is a minimizer of $\mathcal{L}(\mathbf{v};\boldsymbol{\lambda}^*)$ and $\boldsymbol{\lambda}^*$ is a maximizer of $\mathcal{L}(\mathbf{v}^*;\boldsymbol{\lambda})$. For problems featuring strong duality (in which case $P^* = D^*$), a point is an *optimal primal/dual point* if and only if it is a saddle point of the Lagrangian function [3, p. 239]. Furthermore, at a primal/dual solution $(\mathbf{v}^*;\boldsymbol{\lambda}^*)$, the value of the Lagrangian satisfies $\mathcal{L}(\mathbf{v}^*;\boldsymbol{\lambda}^*) = P^* = D^*$. Finding an exact saddle point of the Lagrangian imposes stringent conditions on $\mathcal{L}(\mathbf{v};\boldsymbol{\lambda})$. However, approximate saddle points can be reached under specific settings using gradient-based methods, such as the PD method [37].

According to the PD method, during iteration t, the primal and dual variables are updated as:

$$\mathbf{v}^{t+1} := \mathbf{v}^t - \mu_v^t \nabla_{\mathbf{v}} \mathcal{L}(\mathbf{v}^t;\boldsymbol{\lambda}^t), \tag{3a}$$

$$\boldsymbol{\lambda}^{t+1} := \left[\boldsymbol{\lambda}^t + \mu_\lambda^t \nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{v}^{t+1};\boldsymbol{\lambda}^t)\right]_+, \tag{3b}$$

where μ_v^t and μ_λ^t are positive step sizes, and $[x]_+ := \max\{x, 0\}$. The update in (3a) is a gradient descent step on $\mathcal{L}$ in terms of the primal variable. The update in (3b) is a projected gradient ascent step on $\mathcal{L}$ in terms of the dual variable.

One may attempt to solve (QCQP) using the previous PD iterations. In this case, computing the Lagrangian gradients incurs complexity $\mathcal{O}(N)$. To see this, note that the partial derivative of $\mathcal{L}$ with respect to λ_m is

$$\frac{\partial \mathcal{L}}{\partial \lambda_m} = \mathbf{v}^\dagger \mathbf{M}_m \mathbf{v} - b_m.$$

This partial derivative can be computed in $\mathcal{O}(1)$ steps, simply because $\mathbf{M}_m$ has $\mathcal{O}(1)$ nonzero entries. Therefore, computing the entire gradient $\nabla_{\boldsymbol{\lambda}} \mathcal{L}$ entails complexity $\mathcal{O}(M) = \mathcal{O}(N)$ given that M is proportional to N. Similarly, computing the gradient

$$\nabla_{\mathbf{v}} \mathcal{L} = 2 \left(\mathbf{M}_0 + \sum_{m=1}^{M} \lambda_m \mathbf{M}_m \right) \mathbf{v}$$

incurs $\mathcal{O}(N)$ operations. Contemporary power systems consist of more than $N = 30,000$ nodes. Nonetheless, it is anticipated that the size of the OPF could increase dramatically due to the pressing need to schedule renewables at finer spatiotemporal scales. In this context, can a VQA attain reduced computational complexity? The next section proposes a variational model for the OPF.

4 A doubly variational OPF model

Variational quantum computing is an algorithmic tool for solving high-dimensional problems in a parameterized form using a parameterized quantum circuit (PQC). The well-known variational quantum eigensolver (VQE) solves a specific form of (QCQP) [13]:

$$\begin{aligned} \min_{\mathbf{v} \in \mathbb{C}^N} \quad & \mathbf{v}^\dagger \mathbf{M}_0 \mathbf{v} \\ \text{s.to} \quad & \mathbf{v}^\dagger \mathbf{v} = 1. \end{aligned} \tag{4}$$

The optimal cost in (4) is the smallest eigenvalue of $\mathbf{M}_0$. Moreover, the eigenvector corresponding to the smallest eigenvalue of $\mathbf{M}_0$ is the minimizer. The key idea of VQE is to model the original high-dimensional vector $\mathbf{v}$ as a quantum state $|\boldsymbol{\psi}\rangle$ using $\log N$ qubits. Suppose for now that N is a power of 2. The constraint in (4) can be omitted as the quantum state has unit Euclidean norm anyway. This quantum state is generated upon applying a PQC on an initial state $|\mathbf{0}\rangle_{\log N}$. The PQC is modeled by a unitary matrix $\mathbf{U}(\boldsymbol{\theta})$ parameterized by vector $\boldsymbol{\theta} \in \mathbb{R}^P$ with $P \ll N$. The PQC generates the state $|\boldsymbol{\psi}(\boldsymbol{\theta})\rangle = \mathbf{U}(\boldsymbol{\theta})\,|\mathbf{0}\rangle$, and then measures the expectation $F_0(\boldsymbol{\theta}) = \langle \boldsymbol{\psi}(\boldsymbol{\theta}) | \mathbf{M}_0 | \boldsymbol{\psi}(\boldsymbol{\theta}) \rangle$ and its gradient $\nabla_{\boldsymbol{\theta}} F_0(\boldsymbol{\theta})$. Subsequently, a classical computer receives the gradient information and updates $\boldsymbol{\theta}$ to minimize $F_0(\boldsymbol{\theta})$ using gradient descent or other classical optimization methods.

We propose solving the OPF using a similar workflow. Because the primal variable of the OPF does not have

a unit Euclidean norm, we parameterize it as [38]:

$$\mathbf{v}(\boldsymbol{\theta}) = \alpha\,|\boldsymbol{\psi}(\boldsymbol{\theta})\rangle \tag{5}$$

where $\alpha > 0$ is an auxiliary optimization variable. The PQC used to generate $|\boldsymbol{\psi}(\boldsymbol{\theta})\rangle$ will be henceforth termed the *primal PQC* or PQC$_p$.

Given the parameterization in (5), the problem (QCQP) can be expressed in the variational form:

$$\begin{aligned} P_\theta^* := \min_{\boldsymbol{\theta},\alpha>0}\ & \alpha^2 F_0(\boldsymbol{\theta}) \qquad (\text{QCQP}_\theta)\\ \text{s.to}\ & \alpha^2 F_m(\boldsymbol{\theta}) \le b_m, \quad m = 1:M \end{aligned}$$

where expectations are defined as

$$F_m(\boldsymbol{\theta}) := \langle \boldsymbol{\psi}(\boldsymbol{\theta})|\mathbf{M}_m|\boldsymbol{\psi}(\boldsymbol{\theta})\rangle\,, \quad m = 0:M. \tag{6}$$

As in VQE, the original (QCQP) over $\mathbf{v}$ has been replaced by the variational problem (QCQP$_\theta$) over the parameter vector $\boldsymbol{\theta}$. Attempting to solve (QCQP$_\theta$) using Lagrangian duality entails solving the dual problem:

$$D_\theta^* := \max_{\boldsymbol{\lambda}\ge\mathbf{0}}\ \min_{\boldsymbol{\theta},\alpha>0} \mathcal{L}(\boldsymbol{\theta},\alpha;\boldsymbol{\lambda})$$

over $\boldsymbol{\lambda} \in \mathbb{R}_+^M$. The related Lagrangian function is

$$\mathcal{L}(\boldsymbol{\theta},\alpha;\boldsymbol{\lambda}) = \alpha^2 F_0(\boldsymbol{\theta}) + \sum_{m=1}^{M} \lambda_m \left(\alpha^2 F_m(\boldsymbol{\theta}) - b_m\right). \tag{7}$$

Unfortunately, because the dual variable $\boldsymbol{\lambda}$ has length $\mathcal{O}(N)$, the curse of dimensionality remains.

To bypass this difficulty, dual variables can also be represented using a variational model. A similar idea was proposed in [18] for capturing the primal and dual variables of an SDP. To model $\boldsymbol{\lambda}$ variationally, we introduce a second PQC, henceforth termed the *dual PQC* or PQC$_d$. The dual PQC operates on $\log M$ qubits, supposing again that the number of constraints, M, is a power of two. The dual PQC is modeled by a unitary matrix $\mathbf{V}(\boldsymbol{\phi})$ parameterized by vector $\boldsymbol{\phi} \in \mathbb{R}^Q$ such that $Q \ll M$. The state of PQC$_d$ is $|\boldsymbol{\xi}(\boldsymbol{\phi})\rangle = \mathbf{V}(\boldsymbol{\phi})\,|\mathbf{0}\rangle$; see Fig. 1.

With the aid of PQC$_d$, the m-th entry of $\boldsymbol{\lambda}$ is parameterized as

$$\lambda_m(\boldsymbol{\phi}) = \beta^2 \cdot |\xi_m(\boldsymbol{\phi})|^2, \quad \text{for } m = 1:M \tag{8}$$

where $\beta > 0$ is an auxiliary variable and $\xi_m(\boldsymbol{\phi})$ is the m-th entry of the state $|\boldsymbol{\xi}(\boldsymbol{\phi})\rangle$. Like α in (5), variable β is introduced for scaling purposes. Unlike (5), however, dual variables are related to the *squared magnitudes* of quantum state entries. This ensures that $\lambda_m(\boldsymbol{\phi})$ takes real nonnegative values.

The Lagrangian function related to this doubly parameterized OPF can be expanded into three terms:

$$\mathcal{L}(\boldsymbol{\theta},\alpha;\boldsymbol{\phi},\beta) = \alpha^2 F_0(\boldsymbol{\theta}) + \alpha^2\beta^2 F(\boldsymbol{\theta},\boldsymbol{\phi}) - \beta^2 G(\boldsymbol{\phi}), \tag{9}$$

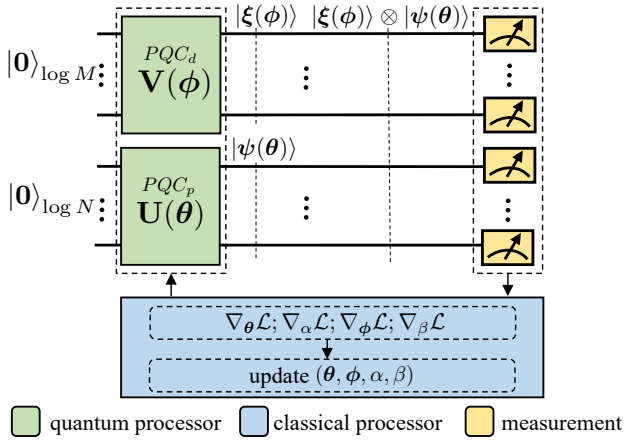

Figure 1: Workflow of approximately solving (QCQP) using a doubly variational approach. The primal PQC (bottom) parameterizes the primal variables. The dual PQC (top) parameterizes the dual variables.

where

$$F(\boldsymbol{\theta},\boldsymbol{\phi}) := \sum_{m=1}^{M} |\xi_m(\boldsymbol{\phi})|^2 F_m(\boldsymbol{\theta}) \tag{10}$$

and

$$G(\boldsymbol{\phi}) := \sum_{m=1}^{M} b_m |\xi_m(\boldsymbol{\phi})|^2. \tag{11}$$

The variational dual problem is defined as

$$D_{\theta,\phi}^* := \max_{\boldsymbol{\phi},\beta>0}\ \min_{\boldsymbol{\theta},\alpha>0} \mathcal{L}(\boldsymbol{\theta},\alpha;\boldsymbol{\phi},\beta). \tag{12}$$

One may wonder why we do not parameterize the primal and dual variables jointly using a single PQC operating on $\log N + \log M$ qubits. Unfortunately, under that design, we would not be able to update the PQC parameters in a meaningful way. Having two PQCs parameterized separately, as depicted in Figure 1, allows us to minimize the Lagrangian function in (9) over $(\boldsymbol{\theta},\alpha)$ and maximize it over $(\boldsymbol{\phi},\beta)$.

The standard Lagrangian dual function is known to be concave with respect to the dual variables, even if the primal problem is nonconvex. That is not the case for the dual function $\min_{\boldsymbol{\theta},\alpha>0} \mathcal{L}(\boldsymbol{\theta},\alpha;\boldsymbol{\phi},\beta)$ in (12). This is because the dual variables are now parameterized, and $\mathcal{L}(\boldsymbol{\theta},\alpha;\boldsymbol{\phi},\beta)$ is not concave in $(\boldsymbol{\phi},\beta)$.

5 Solving the doubly variational OPF

We propose seeking a saddle point of the Lagrangian function in (9) using a PD algorithm. The t-th iteration of the algorithm comprises four steps:

$$\boldsymbol{\theta}^{t+1} := \boldsymbol{\theta}^t - \mu_\theta^t \nabla_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}^t,\alpha^t;\boldsymbol{\phi}^t,\beta^t), \tag{13a}$$

$$\alpha^{t+1} := \left[\alpha^t - \mu_\alpha^t \nabla_\alpha \mathcal{L}(\boldsymbol{\theta}^t, \alpha^t; \boldsymbol{\phi}^t, \beta^t)\right]_+, \tag{13b}$$

$$\boldsymbol{\phi}^{t+1} := \boldsymbol{\phi}^t + \mu_\phi^t \nabla_{\boldsymbol{\phi}} \mathcal{L}(\boldsymbol{\theta}^t, \alpha^t; \boldsymbol{\phi}^t, \beta^t), \tag{13c}$$

$$\beta^{t+1} := \left[\beta^t + \mu_\beta^t \nabla_\beta \mathcal{L}(\boldsymbol{\theta}^t, \alpha^t; \boldsymbol{\phi}^t, \beta^t)\right]_+, \tag{13d}$$

where μ_θ^t, μ_α^t, μ_ϕ^t, and μ_β^t are positive step sizes. Steps (13a)–(13b) constitute gradient descent steps to update the primal variables, while (13c)–(13d) are gradient ascent steps to update the dual variables. The updates of (13) run on a classical computer. The suggested workflow is shown in Fig. 1.

The gradients in (13) can be estimated with the help of the two PQCs. The key observation is that each term of the Lagrangian function can be expressed as an expectation over an observable. More specifically, the first term can be computed using an observable operating on PQC_p through the expectation:

$$F_0(\boldsymbol{\theta}) = \langle \boldsymbol{\psi}(\boldsymbol{\theta}) | \mathbf{M}_0 | \boldsymbol{\psi}(\boldsymbol{\theta}) \rangle . \tag{14}$$

The third term can be computed using an observable operating on PQC_d through the expectation:

$$G(\boldsymbol{\phi}) = \langle \boldsymbol{\xi}(\boldsymbol{\phi}) | \mathbf{S} | \boldsymbol{\xi}(\boldsymbol{\phi}) \rangle , \tag{15}$$

where $\mathbf{S} := \text{dg}(\{b_m\}_{m=1}^M)$ is a diagonal matrix with the OPF parameters b_m on its main diagonal. Due to the parameterization in (8), the Lagrangian of the doubly parameterized problem is quadratic in $|\boldsymbol{\xi}(\boldsymbol{\phi})\rangle$, although the Lagrangian in (7) was linear in $\boldsymbol{\lambda}$.

The second term of the Lagrangian can be expressed as the expectation over an observable operating *jointly* on PQC_p and PQC_d, as shown in the ensuing lemma proved in Appendix A.2.

Lemma 1. *The second summand in* (10) *can be computed as the expectation*

$$F(\boldsymbol{\theta}, \boldsymbol{\phi}) = \langle \boldsymbol{\xi}(\boldsymbol{\phi}), \boldsymbol{\psi}(\boldsymbol{\theta}) | \mathbf{M} | \boldsymbol{\xi}(\boldsymbol{\phi}), \boldsymbol{\psi}(\boldsymbol{\theta}) \rangle \tag{16}$$

defined by the $MN \times MN$ Hermitian matrix

$$\mathbf{M} = \sum_{m=1}^{M} \mathbf{e}_m \mathbf{e}_m^\top \otimes \mathbf{M}_m \tag{17}$$

and $\mathbf{e}_m$ is the m-th column of $\mathbf{I}_M$.

Lemma 1 asserts that $F(\boldsymbol{\theta}, \boldsymbol{\phi})$ is an expectation applied to the composite state of the two PQCs. This is important as it bypasses the need to measure separately each one of the M constraint expectations $F_m(\boldsymbol{\theta})$ of the OPF. This powerful feature was first identified in [18] for solving doubly variational SDPs on a quantum computer. In this work, we adopt this idea to the OPF setting. We also propose a method for efficiently measuring the joint observable for the OPF associated with sparse graphs in general. We uniquely leverage the graph structure of a power system to efficiently measure the three observables using the two PQCs. We defer this discussion to Section 6. If the observables can be measured efficiently, their gradients can be measured efficiently as well via the so-called *parameter-shift rule* (PSR). We review this rule and adapt it to the PD iterations in (13).

Let us focus on the gradient $\nabla_{\boldsymbol{\theta}} F_0(\boldsymbol{\theta}^t)$. Suppose that the unitary matrix modeling the primal PQC takes the form

$$\mathbf{U}(\boldsymbol{\theta}) = \mathbf{W}_{P+1} \prod_{p=1}^{P} \exp(-j\theta_p \mathbf{G}_p) \mathbf{W}_p, \tag{18}$$

where each $\mathbf{G}_p$ is a single-qubit Hermitian generator with two distinct eigenvalues $\pm r$, and $\{\mathbf{W}_p\}_{p=1}^{P+1}$ is a set of fixed gates. Then, the partial derivative of $F_0(\boldsymbol{\theta}^t)$ with respect to the p-th entry of $\boldsymbol{\theta}$ can be computed by evaluating $F_0(\boldsymbol{\theta})$ at two values of $\boldsymbol{\theta}$ [39, 40]:

$$\frac{\partial F_0(\boldsymbol{\theta}^t)}{\partial \theta_p} = r\left(F_0(\boldsymbol{\theta}^t + \tfrac{\pi}{4r}\mathbf{e}_p) - F_0(\boldsymbol{\theta}^t - \tfrac{\pi}{4r}\mathbf{e}_p)\right), \tag{19}$$

where $\mathbf{e}_p$ is the p-th column of the identity matrix $\mathbf{I}_P$. If $\mathbf{G}_p$ is a Pauli rotation from the set $\frac{1}{2}\{\sigma_x, \sigma_y, \sigma_z\}$, we get $r = \frac{1}{2}$ and each θ_p^t is shifted by $\pm\frac{\pi}{2}$.

The expectation value $F_0(\boldsymbol{\theta})$ cannot be evaluated exactly; it can only be estimated via quantum measurements: To estimate $F_0(\boldsymbol{\theta})$ for a particular $\boldsymbol{\theta}$ with precision ϵ, PQC_p has to be sampled $S = \mathcal{O}(\epsilon^{-2})$ times. To obtain the complete gradient $\nabla_{\boldsymbol{\theta}} F_0(\boldsymbol{\theta}^t)$, the previous process is repeated $2P$ times and involves a total of $2PS$ measurement samples.

The gradient of the expectation value in (15) with respect to $\boldsymbol{\phi}$ can be computed similarly by running PQC_d by shifting the values of $\boldsymbol{\phi}^t$ for $2Q$ times or $2QS$ measurement samples. The gradient $\nabla_{\boldsymbol{\theta}} F(\boldsymbol{\theta}^t, \boldsymbol{\phi}^t)$ can be computed by running both PQCs while only shifting the entries of $\boldsymbol{\theta}^t$; and vice versa for computing the gradient $\nabla_{\boldsymbol{\phi}} F(\boldsymbol{\theta}^t, \boldsymbol{\phi}^t)$. It should be clear by now that, due to the PSR, evaluating gradients of expectations is equivalent to measuring expectations. Section 6 discusses how to measure the observables involved in the OPF.

Although the plain PD method has been extensively studied for convex/concave and nonconvex/concave saddle-point problems [41, 42, 43, 44], its convergence has not been fully understood in the nonconvex/nonconcave setting, like the one in (12). Recently, a variation of the PD method, termed the extragradient (EG) method [45], has been shown to converge under certain conditions for nonconvex/nonconcave saddle-point problems [46]. Compared to the plain PD method, each EG iteration doubles the number of gradient evaluations of the Lagrangian function. More specifically, let

vector $\mathbf{z}^t := [\boldsymbol{\theta}^t; \alpha^t; \boldsymbol{\phi}^t; \beta^t]$ collect the primal and dual variables at iterate t, and define the vector:

$$\mathbf{g}(\mathbf{z}^t) := [\nabla_{\boldsymbol{\theta}}\mathcal{L}(\mathbf{z}^t); \nabla_{\alpha}\mathcal{L}(\mathbf{z}^t); -\nabla_{\boldsymbol{\phi}}\mathcal{L}(\mathbf{z}^t); -\nabla_{\beta}\mathcal{L}(\mathbf{z}^t)]. \tag{20}$$

Then, the EG iterations can be expressed as

$$\bar{\mathbf{z}}^t = \mathbf{z}^t - 2\mu_z^t \mathbf{g}(\mathbf{z}^t), \tag{21a}$$

$$\mathbf{z}^{t+1} = \mathbf{z}^t - \mu_z^t \mathbf{g}(\bar{\mathbf{z}}^t), \tag{21b}$$

for a step size $\mu_z^t > 0$. In essence, the EG method first computes an intermediate point $\bar{\mathbf{z}}^t$ based on the plain PD method, and then updates $\mathbf{z}^t$ to $\mathbf{z}^{t+1}$ upon evaluating the gradient operator at the intermediate point. Despite doubling the number of gradient evaluations, EG iterates feature favorable convergence properties as established analytically in Section 7 and corroborated numerically in Section 8.

Before concluding this section, it is worth noting that, as with other VQAs, training the primal-dual PQC pair could be impeded by the so-called *barren plateaus* issue, according to which gradients vanish exponentially fast in terms of the number of qubits and depth of PQCs [47, 48]. This phenomenon can occur when designing VQAs for large-scale problems and/or using overly expressive PQCs. A related concern is that the PQC architectures that are both sufficiently expressive and free from barren plateaus are often efficiently simulated by classical algorithms [49]. Interestingly, recent results show that the so-called dynamic quantum phase gates (PQCs) [50], which incorporate early measurements followed by measurement-conditional unitaries, and quantum recurrent embedding neural networks [51] can provably avoid barren plateaus while retaining high expressivity. Moreover, dynamic PQCs inherit the worst-case classical hardness of the universal unitary circuits, precluding the possibility of classical simulatability. Investigating the performance of the doubly variational OPF on dynamic PQCs or quantum recurrent embedding neural networks for large-scale networks is an interesting future direction.

6 Measuring OPF observables

This section reviews the extended Bell measurement (XBM) method proposed in [4] and explains how it can be adapted to efficiently measure OPF observables.

6.1 Extended Bell measurement (XBM) method

Consider measuring the expectation value $\langle\boldsymbol{\psi}(\boldsymbol{\theta})|\,\mathbf{M}_m\,|\boldsymbol{\psi}(\boldsymbol{\theta})\rangle$ associated with the observable $\mathbf{M}_m$, which is a $N \times N$ Hermitian matrix. The matrix $\mathbf{M}_m$

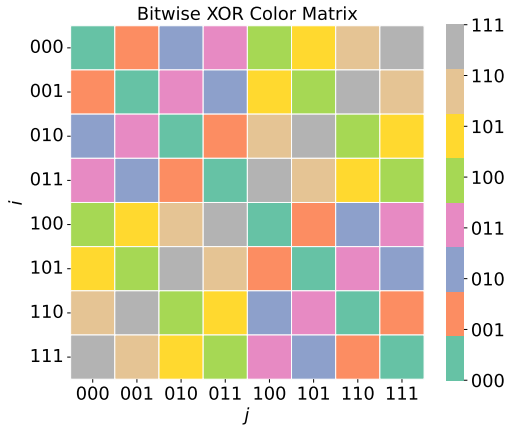

Figure 2: The XBM method of [4] groups the entries of an $N \times N$ Hermitian matrix $\mathbf{M}_m$ into N groups or colors. This XBM grouping or entry coloring is shown here for an 8×8 matrix associated with $\log 8 = 3$ qubits. Each color c is identified by the binary form of its index c. Color c includes all matrix entries with index pairs (i, j) for which $i \oplus j = c$. For example, the grey color $c = 7 = |111\rangle$ includes all entries lying on the main anti-diagonal of $\mathbf{M}_m$. Note that the color arrangement is symmetric.

is conventionally expressed as

$$\mathbf{M}_m = \sum_{i=0}^{N-1}\sum_{j=0}^{N-1} \langle i|\mathbf{M}_m|j\rangle\, |i\rangle\langle j|\,. \tag{22}$$

The XBM method proposes the alternative expansion:

$$\mathbf{M}_m = \sum_{c=0}^{N-1} \underbrace{\sum_{i=0}^{N-1} \langle i|\mathbf{M}_m|i\oplus c\rangle\, |i\rangle\langle i\oplus c|}_{:=\mathbf{M}_m^c} = \sum_{c=0}^{N-1} \mathbf{M}_m^c, \tag{23}$$

where $\oplus$ denotes the bitwise exclusive OR (XOR) operator if indices i and c are represented in binary form. To establish the equivalence between (22) and (23), we need to show that every pair (i, j) corresponds to a unique pair $(i, i \oplus c)$, which indeed holds for $c = i \oplus j$. Due to this alternative indexing, the matrix $\mathbf{M}_m$ can be decomposed into N matrices $\{\mathbf{M}_m^c\}_{c=0}^{N-1}$ defined as

$$\langle i|\mathbf{M}_m^c|j\rangle = \begin{cases} \langle i|\mathbf{M}_m|j\rangle\,, & \text{if } i \oplus j = c \\ 0, & \text{otherwise.} \end{cases}$$

For example, the matrix $\mathbf{M}_m^0$ carries the diagonal entries of $\mathbf{M}_m$ because $i \oplus i = 0$. Figure 2 uses color coding to identify the index pairs (i, j) corresponding to the same c for an observable over three qubits [4]. We will henceforth say that all index pairs (i, j) satisfying $i \oplus j = c$ for a specific c belong to the same *color* c.

Because $i \oplus j = j \oplus i$, every $\mathbf{M}_m^c$ is Hermitian. Therefore, the matrices

$$\hat{\mathbf{M}}_m^c := \mathrm{Re}\{\mathbf{M}_m^c\} \quad \text{and} \quad \check{\mathbf{M}}_m^c := \iota\, \mathrm{Im}\{\mathbf{M}_m^c\} \tag{24}$$

are also Hermitian for all c, and $\check{\mathbf{M}}_m^0 = \mathbf{0}$.

Based on (23), the expectation value associated with $\mathbf{M}_m$ can be expressed as a sum of $2N-1$ observables:

$$\langle \boldsymbol{\psi} | \mathbf{M}_m | \boldsymbol{\psi} \rangle = \sum_{c=0}^{N-1} \langle \boldsymbol{\psi} | \hat{\mathbf{M}}_m^c | \boldsymbol{\psi} \rangle + \sum_{c=1}^{N-1} \langle \boldsymbol{\psi} | \check{\mathbf{M}}_m^c | \boldsymbol{\psi} \rangle \,. \tag{25}$$

Because $\hat{\mathbf{M}}_m^0$ is diagonal, the expectation value $\langle \boldsymbol{\psi} | \hat{\mathbf{M}}_m^0 | \boldsymbol{\psi} \rangle$ can be measured in the computational basis. Upon measuring $|\boldsymbol{\psi}\rangle$, the binary vector outcome $|i\rangle$ is observed with probability $|\psi_i|^2$. Define a random variable taking the value $\langle i | \hat{\mathbf{M}}_m^0 | i \rangle$ when outcome $|i\rangle$ is observed. Then, the expectation $\langle \boldsymbol{\psi} | \hat{\mathbf{M}}_m^0 | \boldsymbol{\psi} \rangle$ in (25) is the expected value of this random variable.

The remaining expectation values in (25) are not as straightforward to measure because the matrices $(\hat{\mathbf{M}}_m^c, \check{\mathbf{M}}_m^c)$ are non-diagonal for $c = 1, \ldots, N-1$. The advantage of XBM's color decomposition is that these matrices can be diagonalized by unitaries corresponding to qubit-efficient circuits. This key result from [4] is summarized in the following lemma.

Lemma 2 ([4])**.** *Consider the color index $c \geq 1$. Let index k_c denote the most significant bit taking the value of one in the binary representation of $|c\rangle = |c_{\log N - 1} \cdots c_1 c_0\rangle$. Then, the matrices $\hat{\mathbf{M}}_m^c$ and $\check{\mathbf{M}}_m^c$ admit the eigenvalue decompositions:*

$$\hat{\mathbf{M}}_m^c = \mathbf{U}_c \hat{\boldsymbol{\Lambda}}_m^c \mathbf{U}_c^\top \quad \textit{and} \quad \check{\mathbf{M}}_m^c = \mathbf{S}_c^\dagger \mathbf{U}_c \check{\boldsymbol{\Lambda}}_m^c \mathbf{U}_c^\top \mathbf{S}_c$$

where

i) unitary matrix $\mathbf{U}_c$ can be implemented using one Hadamard gate and at most $(\log N - 1)$ CNOT gates;

ii) unitary matrix $\mathbf{S}_c$ applies a phase gate $\mathbf{S}$ on qubit k_c and acts trivially on all other qubits; and

iii) the diagonal eigenvalue matrices are

$$\hat{\boldsymbol{\Lambda}}_m^c := \frac{1}{2} \sum_{i=0}^{N-1} \langle i | \hat{\mathbf{M}}_m^c | i \oplus c \rangle \left(|i\rangle\langle i| - \mathbf{X}_{k_c} |i\rangle\langle i| \mathbf{X}_{k_c}^\top \right)$$

$$\check{\boldsymbol{\Lambda}}_m^c := \frac{-\iota}{2} \sum_{i=0}^{N-1} \langle i | \check{\mathbf{M}}_m^c | i \oplus c \rangle \left(|i\rangle\langle i| - \mathbf{X}_{k_c} |i\rangle\langle i| \mathbf{X}_{k_c}^\top \right)$$

with $\mathbf{X}_{k_c}$ applying a NOT gate on qubit k_c.

The practical value of Lemma 2 is that the expectation value

$$\langle \boldsymbol{\psi} | \hat{\mathbf{M}}_m^c | \boldsymbol{\psi} \rangle = \langle \boldsymbol{\psi} | \mathbf{U}_c \hat{\boldsymbol{\Lambda}}_m^c \mathbf{U}_c^\top | \boldsymbol{\psi} \rangle$$

can be measured by transforming $|\boldsymbol{\psi}\rangle$ to $\mathbf{U}_c^\top |\boldsymbol{\psi}\rangle$ and then measuring the latter in the computational basis using the diagonal matrix $\hat{\boldsymbol{\Lambda}}_m^c$. The transformation is qubit-efficient because the unitary $\mathbf{U}_c$ involves at most $\log N$ elementary gates. The expectation value $\langle \boldsymbol{\psi} | \check{\mathbf{M}}_m^c | \boldsymbol{\psi} \rangle$ can be measured similarly, with at most $\log N + 1$ elementary gates; $\log N$ for $\mathbf{U}_c$ plus one for $\mathbf{S}_c$. A critical feature is that the unitary $\mathbf{U}_c$ diagonalizing $(\hat{\mathbf{M}}_m^c, \check{\mathbf{M}}_m^c)$ depends solely on the color (locations of nonzero entries) of $(\hat{\mathbf{M}}_m^c, \check{\mathbf{M}}_m^c)$, and *not* on the values of those entries. Only the eigenvalue matrices $(\hat{\boldsymbol{\Lambda}}_m^c, \check{\boldsymbol{\Lambda}}_m^c)$ depend on those values, and can be computed without the need for an eigenvalue decomposition.

Although each expectation value in (25) can be estimated efficiently, there are $(2N-1)$ of them. Unfortunately, each one of them requires a different quantum circuit. Therefore, the expectation value $\langle \boldsymbol{\psi} | \mathbf{M}_m | \boldsymbol{\psi} \rangle$ can be measured efficiently only if all nonzero entries of $\mathbf{M}_m$ are covered by a few colors. Let $\mathcal{C}$ denote the set of colors for which either $\hat{\mathbf{M}}_m^c$ or $\check{\mathbf{M}}_m^c$ is nonzero, so that $\mathbf{M}_m$ is decomposed as $\mathbf{M}_m = \sum_{c \in \mathcal{C}} \mathbf{M}_m^c$. For the XBM protocol to be qubit-efficient, the number of colors $C = |\mathcal{C}|$ should scale polynomially in $\log N$. Reference [4] establishes that *banded* matrices occupy relatively few colors. Specifically, it is shown that a k-banded matrix can be decomposed using $C = \mathcal{O}(k \log N)$ rather than N colors. For the example of Figure 2, a matrix with bandwidth $k = 1$ requires 4 colors.

Based on Lemma 2, if the expectation value $F_m(\boldsymbol{\theta})$ occupies C colors, it can be decomposed into $2C-1$ expectations:

$$F_m(\boldsymbol{\theta}) = \sum_{c=1}^{2C-1} F_m^c(\boldsymbol{\theta}) = \sum_{c=1}^{2C-1} \langle \boldsymbol{\psi}_c(\boldsymbol{\theta}) | \boldsymbol{\Lambda}_m^c | \boldsymbol{\psi}_c(\boldsymbol{\theta}) \rangle \,, \tag{26}$$

where $|\boldsymbol{\psi}_c(\boldsymbol{\theta})\rangle = \mathbf{U}_c^\top |\boldsymbol{\psi}(\boldsymbol{\theta})\rangle$ and $\boldsymbol{\Lambda}_m^c = \hat{\boldsymbol{\Lambda}}_m^c$ for $c = 1, \ldots, C$; and $|\boldsymbol{\psi}_c(\boldsymbol{\theta})\rangle = \mathbf{U}_c^\top \mathbf{S}_c |\boldsymbol{\psi}(\boldsymbol{\theta})\rangle$ and $\boldsymbol{\Lambda}_m^c = \check{\boldsymbol{\Lambda}}_m^c$ for $c = C+1, \ldots, 2C-1$. Each one of the observables involved can be diagonalized by a different unitary matrix.

We next describe how to effect an OPF measurement matrix $\mathbf{M}$ that uses a few colors.

6.2 Node permutation for qubit-efficient measurements

To solve the doubly parameterized OPF, we need to evaluate the Lagrangian function in (9) by measuring the observable in Lemma 1 so that we can compute its gradients over $(\boldsymbol{\theta}, \alpha; \boldsymbol{\phi}, \beta)$. We next adopt the XBM protocol to efficiently measure OPF observables.

According to (17), observable matrix $\mathbf{M}$ has the $N \times N$ matrices $\{\mathbf{M}_m\}_{m=1}^M$ as its diagonal blocks. Despite being block-diagonal, matrix $\mathbf{M}$ may still occupy up to

N colors. Critically, we do not necessarily have to work with $\mathbf{M}$. Instead, we can permute variables to effect a new measurement matrix that is amenable to more efficient measurements. Suppose we permute vector $\mathbf{v}$ by a permutation matrix $\mathbf{P}$ to get $\mathbf{Pv}$. This entails reordering the nodes of the power system, which is a trivial task that can be performed before solving the OPF. The observable in Lemma 1 can be expressed as

$$\sum_{m=1}^{M} \lambda_m \mathbf{v}^\dagger \mathbf{M}_m \mathbf{v} = \sum_{m=1}^{M} \lambda_m \mathbf{v}^\dagger \mathbf{P}^\top (\mathbf{P}\mathbf{M}_m\mathbf{P}^\top)\mathbf{P}\mathbf{v}. \quad (27)$$

Ideally, we would like to design a *single* $\mathbf{P}$ to minimize the maximum bandwidth across all $\mathbf{P}\mathbf{M}_m\mathbf{P}^\top$. To this end, we suggest designing $\mathbf{P}$ to minimize the bandwidth of $\mathbf{P}\mathbf{Y}\mathbf{P}^\top$, where $\mathbf{Y}$ is the node admittance matrix of the power system. This is because the sparsity pattern of $\mathbf{Y}$ captures the *union* of sparsity patterns of all $\mathbf{M}_m$. Therefore, minimizing the bandwidth of $\mathbf{P}\mathbf{Y}\mathbf{P}^\top$ is equivalent to minimizing the maximum bandwidth across all $\mathbf{P}\mathbf{M}_m\mathbf{P}^\top$.

Unfortunately, minimizing the bandwidth of a matrix is an NP-hard problem even if the matrix is sparse [5]. Nonetheless, heuristic approaches, such as the reverse Cuthill-McKee (RCM) algorithm [6], provide reasonably good solutions in linear complexity $\mathcal{O}(L_e) = \mathcal{O}(N)$, where L_e is the number of edges in the power system graph. The RCM algorithm has been used to reduce the bandwidth of the node admittance matrix $\mathbf{Y}$ in classical computing [52, 53, 54]. Albeit permuting $\mathbf{M}_m$'s using the RCM algorithm takes $\mathcal{O}(N)$ operations, the crucial point is that this permutation can be performed once. This is because the power system topology changes infrequently, and such changes only affect a few lines. As long as the network topology remains unchanged, an instance of the OPF is characterized by parameters b_m, which vary frequently. In contrast to the permutation step, OPF instances over the same topology must be solved repeatedly, every few minutes during real-time operation and potentially millions of times in long-term planning.

We numerically validated the effect of node permutation on the bandwidth of $\mathbf{M}$. We ran the RCM algorithm on the 66 benchmark power system models of the `pglib` dataset [55]. Figure 3 shows the bandwidth of $\mathbf{M}$ before and after node permutation. Because the RCM solution depends on the initial node, we conducted 200 Monte Carlo runs per power system, with the initial node drawn uniformly at random per run. Out of the 200 runs, we retained the permutation that yielded the smallest bandwidth. Evidently, permuting nodes reduces the bandwidth of $\mathbf{M}$ dramatically. Upon data fitting and cross-validation, the reduced bandwidth was numerically found to scale as $(\log N)^3$. We also empirically tested the effect of node permutation on the number of colors in the OPF measurement matrix. Figure 4 shows that node permutation significantly reduces the number of colors, too. If the bandwidth of $\mathbf{P}\mathbf{M}\mathbf{P}^\top$ is $k = \mathcal{O}((\log N)^3)$ per data fitting in Figure 3, the analysis of [4] upper bounds the number of its colors as $C = \mathcal{O}(k \log N) = \mathcal{O}((\log N)^4)$. However, the number of colors in $\mathbf{P}\mathbf{M}\mathbf{P}^\top$ was found to scale as $(\log N)^3$.

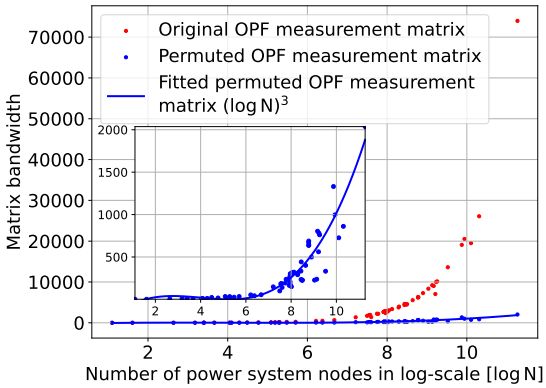

Figure 3: The bandwidth of the original and permuted OPF measurement matrices as a function of the logarithmic network size $(\log N)$ for the power system graphs in `pglib`. Nodes were permuted by the RCM algorithm. Permuted matrices feature patently smaller bandwidths. Upon data fitting and 5-fold cross-validation over different polynomial and exponential functions, the reduced bandwidth was numerically found to scale as $(\log N)^3$. The run time of the RCM algorithm (implemented by function `scipy.sparse.csgraph.reverse_cuthill_mckee` in Python 3.11.9) on the largest power system graph of $N = 78,784$ is a few seconds using a MacBook laptop equipped with an M3 Pro processor and 36 GB of RAM.

We have so far focused on measuring the second observable of the Lagrangian function in (9). The third observable is diagonal and is thus straightforward to measure. The first observable $F_0(\boldsymbol{\theta})$, however, is defined over the non-diagonal matrix $\mathbf{M}_0$. In general, among all $\mathbf{M}_m$'s, matrix $\mathbf{M}_0$ could have the most nonzero off-diagonal entries. These entries correspond to power injections from nodes hosting generators and can be $\mathcal{O}(N)$ of those entries. Fortunately, the sparsity pattern of $\mathbf{M}_0$ is a subset of the sparsity pattern of $\mathbf{Y}$. Therefore, if nodes are permuted so that $\mathbf{M}$ occupies a few colors C, the same permutation works well also for $\mathbf{M}_0$. In other words, matrix $\mathbf{P}\mathbf{M}_0\mathbf{P}^\top$ would occupy C colors or fewer.

Since node permutation yields OPF measurement matrices with much fewer colors, we suggest the following workflow: *i)* Given a power system model, feed the sparsity pattern of $\mathbf{Y}$ into the RCM algorithm to find a near-optimal node permutation matrix $\mathbf{P}$; *ii)* Formulate the OPF based on the permuted nodes; *iii)* Solve the OPF using the VQA of Section 4; and *iv)* If needed, apply the reverse permutation to the OPF primal solution

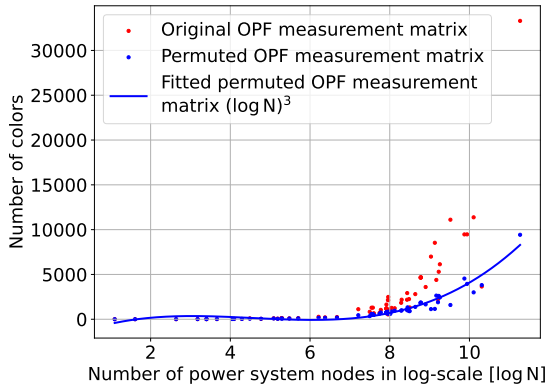

Figure 4: The number of colors in the original and permuted OPF measurement matrices as a function of the logarithmic network size $(\log N)$ for the power system graphs in `pglib`. Permuting nodes using the RCM algorithm was numerically found to significantly reduce the number of colors in use. Data fitting and 5-fold cross validation showed that the number of colors scales as $(\log N)^3$.

to recover the original node ordering. It is henceforth assumed that primal variables and measurement matrices have been permuted already. Therefore, we will henceforth use the lighter notation $\mathbf{M}$ and $\mathbf{M}_m$ instead of $\mathbf{PMP}^\top$ and $\mathbf{PM}_m\mathbf{P}^\top$.

6.3 Implementation details

We next explain how to estimate the Lagrangian gradients needed in (13) using the primal/dual PQC pair. We start with $\nabla_\alpha \mathcal{L}$ in (13b), which entails estimating $F_0(\boldsymbol{\theta})$ and $F(\boldsymbol{\theta},\boldsymbol{\phi})$. To simplify the presentation of measuring $F_0(\boldsymbol{\theta})$, suppose $\mathbf{M}_0$ occupies all C colors used in $\mathbf{Y}$. According to (26), measuring $F_0(\boldsymbol{\theta})$ requires $2C-1$ circuits. Each rotated primal circuit c involves the same primal PQC $\mathbf{U}(\boldsymbol{\theta})$ followed by a different XBM unitary $\mathbf{U}_c$. These $2C-1$ circuits can be executed sequentially or in parallel, as shown in Figure 5. The sequential implementation operates on $\log N$ qubits and requires compiling $2C-1$ circuits in a sequential fashion. The parallel implementation is $2C-1$ times faster but operates on $\log N(2C-1)$ qubits and requires replicating the primal PQC $2C-1$ times.

Measuring observable $F(\boldsymbol{\theta},\boldsymbol{\phi})$ is more complicated as it involves estimating the double summation

$$F(\boldsymbol{\theta},\boldsymbol{\phi}) = \sum_{m=1}^{M}\sum_{c=1}^{2C-1} |\xi_m(\boldsymbol{\phi})|^2 F_m^c(\boldsymbol{\theta}) = \sum_{c=1}^{2C-1}\sum_{m=1}^{M} |\xi_m(\boldsymbol{\phi})|^2 F_m^c(\boldsymbol{\theta}). \quad (28)$$

Interestingly, for a specific index c, the corresponding sum $\sum_{m=1}^{M} |\xi_m(\boldsymbol{\phi})|^2 F_m^c(\boldsymbol{\theta})$ can be evaluated in two

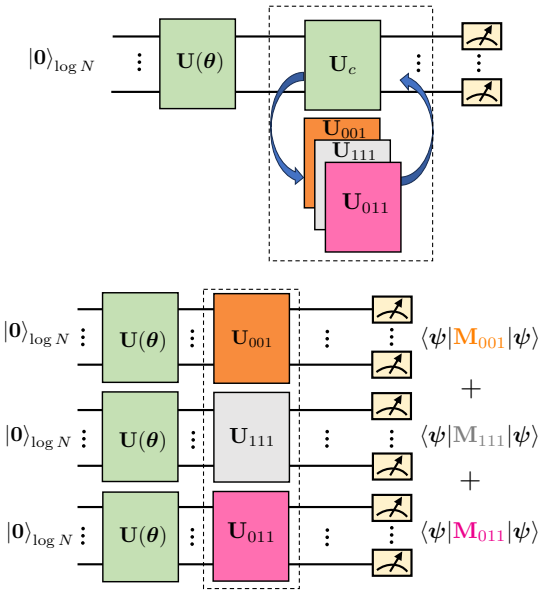

Figure 5: Implementation of rotation unitaries. *Top:* A unitary $\mathbf{U}_c$ is recompiled $2C-1$ times to consider for $2C-1$ rotation unitaries running in sequence. *Bottom:* $2C-1$ rotation unitaries running in parallel.

steps. We first measure the dual PQC and observe $|m\rangle$ with probability $|\xi_m(\boldsymbol{\phi})|^2$. Given the observed $|m\rangle$, we can now easily estimate the expectation $F_m^c(\boldsymbol{\theta})$ corresponding to a diagonal observable.

Measuring $\nabla_\beta \mathcal{L}$ in (13d) is trivial as it depends on $F(\boldsymbol{\theta},\boldsymbol{\phi})$ and $G(\boldsymbol{\phi})$.

The gradient $\nabla_{\boldsymbol{\theta}}\mathcal{L}(\boldsymbol{\theta})$ can be estimated based on the PSR. For example, the partial derivative of $\mathcal{L}$ with respect to θ_p can be found as

$$\frac{\partial \mathcal{L}}{\partial \theta_p} = \frac{\alpha^2}{2}\left(F_0(\boldsymbol{\theta}_p^+) - F_0(\boldsymbol{\theta}_p^-)\right) + \frac{\alpha^2\beta^2}{2}\left(F(\boldsymbol{\theta}_p^+,\boldsymbol{\phi}) - F(\boldsymbol{\theta}_p^-,\boldsymbol{\phi})\right) \quad (29)$$

where $\boldsymbol{\theta}_p^\pm := \boldsymbol{\theta} \pm \frac{\pi}{2}\mathbf{e}_p$. According to (29), estimating $\nabla_{\boldsymbol{\theta}}\mathcal{L}$ requires running $2C-1$ rotated primal PQCs, $2P$ times each for $\{\boldsymbol{\theta}_p^+, \boldsymbol{\theta}_p^-\}_{p=1}^P$, while the dual PQC is set at $\boldsymbol{\phi}$.

Similarly, evaluating $\nabla_{\boldsymbol{\phi}}\mathcal{L}$ in (13c) requires running $2C-1$ rotated primal circuits set at $\boldsymbol{\theta}$, and the dual PQC for $2Q$ values for $\{\boldsymbol{\phi}_q^+, \boldsymbol{\phi}_q^-\}_{q=1}^Q$ as

$$\frac{\partial \mathcal{L}}{\partial \phi_q} = \frac{\alpha^2\beta^2}{2}\left(F(\boldsymbol{\theta},\boldsymbol{\phi}_q^+) - F(\boldsymbol{\theta},\boldsymbol{\phi}_q^-)\right) - \frac{\beta^2}{2}\left(G(\boldsymbol{\phi}_q^+) - G(\boldsymbol{\phi}_q^-)\right) \quad (30)$$

where $\boldsymbol{\phi}_q^\pm := \boldsymbol{\phi} \pm \frac{\pi}{2}\mathbf{e}_q$.

Altogether, each PD iteration requires running $(2C-1)(2P+1)$ rotated primal circuits and $2Q+1$ dual circuits. Each circuit is run for S samples. According to [13], the number of PQC parameters (P, Q) is expected to be $\tilde{\mathcal{O}}(\log N)$. If the number of colors grows approximately as $C = \mathcal{O}((\log N)^3)$, each PD iteration in (13) incurs a computational complexity of $\tilde{\mathcal{O}}(\log N)$, which is more efficient compared to $\mathcal{O}(N)$ of the PD iteration in (3) running on a classical computer.

7 Convergence analysis

This section studies the convergence of the EG iterates in (21). For a general optimization problem, if the Lagrangian function is convex in the primal variables, concave in the dual variables, and has a Lipschitz gradient in primal and dual variables, the EG iterates converge to a saddle point of the Lagrangian function; see [45]. For nonconvex/nonconcave saddle-point problems, however, establishing convergence of the EG iterates is challenging. Reference [46] has recently shown that, under certain conditions, the EG iterates converge to an approximately stationary point of the Lagrangian function. Of course, a stationary point may not necessarily be a saddle point of the Lagrangian. Similar convergence claims have also been established in the stochastic setting, where EG updates rely on noisy yet unbiased gradient estimates in lieu of the actual Lagrangian gradients. Given that the Lagrangian function in (12) is nonconvex in $(\boldsymbol{\theta}, \alpha)$ and nonconcave in $(\boldsymbol{\phi}, \beta)$, we adopt the convergence proof from [46] for the EG iterates in (21) and provide an upper bound on the number of measurement samples across all EG iterations. To this end, the following lemma shows that the operator $\mathbf{g}(\mathbf{z})$ is L-Lipschitz continuous; see Appendix A.3 for a proof.

Lemma 3. *The vector-valued function $\mathbf{g}(\mathbf{z})$ defined in (20) is Lipschitz continuous with the Lipschitz constant*

$$L = \left(P\bar{\alpha}^2\bar{\beta}^2 + Q\bar{\alpha}^2\bar{\beta}^2 + 2\bar{\alpha}\bar{\beta}^2 + 2\bar{\alpha}^2\bar{\beta}\right)\max_m \|\mathbf{M}_m\| + \max\{(P\bar{\alpha}^2 + 2\bar{\alpha})\|\mathbf{M}_0\|, (Q\bar{\beta}^2 + 2\bar{\beta})\max_m |b_m|\}, \tag{31}$$

where $\bar{\alpha}$ and $\bar{\beta}$ are upper bounds on the optimal values of α and β, respectively.

The convergence analysis relies on the ensuing assumption [46].

Assumption 1. *There exists a stationary point $\mathbf{z}^*$ of $\mathcal{L}(\mathbf{z})$ such that:*

$$\mathbf{g}(\mathbf{z})^\top(\mathbf{z} - \mathbf{z}^*) \geq -\frac{\rho}{2}\|\mathbf{g}(\mathbf{z})\|_2^2, \tag{32}$$

for all $\mathbf{z} \in \mathbb{R}^{P+Q+2}$ and some parameter $\rho \in \left[0, \frac{1}{4\sqrt{2}L}\right)$, where L is defined in (31).

Another challenge for our problem is that the EG updates in (21) rely on noisy estimates of $\mathbf{g}(\mathbf{z})$ because observables are measured only based on samples. Specifically, as discussed earlier, each PD iteration in (13) utilizes S measurement samples per rotated primal/dual circuit to evaluate the Lagrangian gradients. Similarly, each step (21a) and (21b) of the EG updates requires S measurement samples per circuit to estimate $\mathbf{g}(\mathbf{z})$ as:

$$\hat{\mathbf{g}}(\mathbf{z}) \coloneqq \frac{1}{S}\sum_{s=1}^{S}\hat{\mathbf{g}}_s(\mathbf{z}), \tag{33}$$

where $\hat{\mathbf{g}}_s(\mathbf{z})$ denotes the estimate of $\mathbf{g}(\mathbf{z})$ based on the single sample s. Under this stochastic setting, we adapt the result from [46] to bound the number of iterations T and measurement samples S such that

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\left[\left\|\hat{\mathbf{g}}(\bar{\mathbf{z}}^t)\right\|_2\right] \leq \epsilon. \tag{34}$$

The expectation in (34) is taken over the randomness of $\hat{\mathbf{g}}(\mathbf{z})$. Suppose now that the final output $\hat{\mathbf{z}}$ of the EG method is drawn by uniformly sampling at random from the EG iterates $\{\bar{\mathbf{z}}^t\}_{t=0}^{T-1}$. Then, it is easy to see that $\hat{\mathbf{z}}$ satisfies

$$\mathbb{E}[\|\hat{\mathbf{g}}(\hat{\mathbf{z}})\|_2] \leq \epsilon, \tag{35}$$

where the expectation is now taken over the noisy gradients as well as the sampling process to draw the final output $\hat{\mathbf{z}}$. The convergence result is formalized in the following theorem from [46].

Theorem 1 ([46])**.** *Let $\mathbf{g}(\mathbf{z})$ be an L-Lipschitz operator that satisfies Assumption 1. Let $\hat{\mathbf{g}}(\mathbf{z})$ be an unbiased estimator of $\mathbf{g}(\mathbf{z})$ such that the variance of $\hat{\mathbf{g}}(\mathbf{z})$ is bounded as*

$$\mathbb{E}\left[\|\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})\|_2^2\right] \leq \frac{\sigma^2}{S}, \tag{36}$$

where S is the number of measurement samples per rotated primal/dual circuit. Given an arbitrary initial point $\mathbf{z}^0$, run the EG iterations in (21) for T times with the step size $\mu_z^t = \frac{1}{2\sqrt{2}L}$. Accordingly, let $\{\mathbf{z}^t\}_{t=1}^{T}$ and $\{\bar{\mathbf{z}}^t\}_{t=0}^{T-1}$ denote the sequences of points generated by the EG iteration. Select a point $\hat{\mathbf{z}}$ uniformly at random from $\{\bar{\mathbf{z}}^t\}_{t=0}^{T-1}$. If the number of iterations is selected as

$$T = \left\lceil \frac{32L^2\left\|\mathbf{z}^* - \mathbf{z}^0\right\|_2^2}{\epsilon^2(1 - 4\sqrt{2}L\rho)} \right\rceil \tag{37}$$

and each step (21a) and (21b) uses

$$S = \left\lceil \frac{8\sigma^2(8 + \sqrt{2}L\rho)}{\epsilon^2(1 - 4\sqrt{2}L\rho)} \right\rceil \tag{38}$$

samples per iteration, then $\mathbb{E}[\|\hat{\mathbf{g}}(\hat{\mathbf{z}})\|_2] \leq \epsilon$.

We next expound on whether the Lagrangian function $\mathcal{L}(\mathbf{z})$ satisfies the requirements of Theorem 1. First, we were unable to verify Assumption 1, so it is taken as given. Second, the L-Lipschitz continuity of the operator $\mathbf{g}(\mathbf{z})$ has been shown in Lemma 3. Third, the estimator $\hat{\mathbf{g}}(\mathbf{z})$ is unbiased because Lagrangian gradients $\mathcal{L}(\mathbf{z})$ are computed via the PSR, and estimates of quantum observables based on sample averages are known to be unbiased estimators [56]. Lastly, the bound on the variance in (36) is established in the ensuing lemma shown in Appendix A.4.

Lemma 4. *The parameter* σ^2 *upper bounding the variance of the gradient estimator in* (36) *can be found as:*

$$\sigma^2 = \frac{Q\bar{\beta}^4 + 8\bar{\beta}^2}{2} \max_m b_m^2 + \frac{8\bar{\alpha}^2 + P\bar{\alpha}^4}{2} \sum_{c=1}^{2C-1} \|\mathbf{M}_0^c\|^2 + \frac{8\bar{\alpha}^2\bar{\beta}^4 + 8\bar{\alpha}^4\bar{\beta}^2 + (P+Q)\bar{\alpha}^4\bar{\beta}^4}{2} \sum_{c=1}^{2C-1} \max_m \|\mathbf{M}_m^c\|^2. \tag{39}$$

Having determined the Lipschitz constant L for $\mathbf{g}(\mathbf{z})$ and the variance bound σ^2 for $\hat{\mathbf{g}}(\mathbf{z})$, the following proposition provides an upper bound on the total number of measurement samples running on all circuits.

Proposition 1. *Given an arbitrary point* $\mathbf{z}^0$*, run the EG iteration in* (21) *for* T *times with the step size* $\mu_z^t = \frac{1}{2\sqrt{2}L}$*. Accordingly, let* $\{\mathbf{z}^t\}_{t=1}^T$ *and* $\{\bar{\mathbf{z}}^t\}_{t=0}^{T-1}$ *denote the sequences of points generated by the EG iteration in* (21)*. Choose a point* $\hat{\mathbf{z}}$ *uniformly from* $\{\bar{\mathbf{z}}^t\}_{t=0}^{T-1}$*. If the operator* $\mathbf{g}(\mathbf{z})$ *satisfies Assumption 1, then after*

$$((2P+1)(2C-1)+2Q+1)\left\lceil \frac{4,224L^2\sigma^2\|\mathbf{z}^* - \mathbf{z}^0\|_2^2}{\epsilon^4(1-4\sqrt{2}L\rho)^2} \right\rceil \tag{40}$$

measurement samples across all iterations and for all circuits, it holds that

$$\mathbb{E}[\|\hat{\mathbf{g}}(\hat{\mathbf{z}})\|_2] \leq \epsilon.$$

The proposition can be established by multiplying the number of primal/dual circuits $((2P+1)(2C-1)+2Q+1)$, the number of iterations in (37), and twice the number of samples in (38). Here, we upper bound the factor $(8+\sqrt{2}L\rho)$ appearing in the numerator of (38) by 8.25 for $\rho \in \left[0, \frac{1}{4\sqrt{2}L}\right)$.

Proposition 1 shows that the total number of measurement samples scales polynomially with the number of colors C, the number of PQC parameters (P, Q), the Lipschitz constant L, and the variance bound σ^2.

As discussed in Section 6.3, the number of colors C and the number of PQC parameters (P, Q) grow polynomially in $\log N$. Nonetheless, the dependence of L in (31) and σ^2 in (39) on $\bar{\alpha}$ and $\bar{\beta}$ is unfavorable. This is because the magnitudes of the entries in $\mathbf{v}$ in (QCQP) are constrained to lie in the range of $[0.9, 1.1]$, so that $\bar{\alpha} = 1.1\sqrt{N}$. Furthermore, as OPF entails equality constraints for nodes connected to loads, each of these equality constraints is converted to two inequality constraints upon (QCQP). Since these inequalities should be active at optimality, the complementary slackness of the Karush–Kuhn–Tucker conditions written for (QCQP) implies that the associated dual variables are generally nonzero [3, Ch. 5]. Because the number of load nodes in real-world power systems is typically proportional to N, it follows that $\bar{\beta}$ may scale with $\sqrt{N}$ too. Overall, in practice, parameters L in (31) and σ^2 in (39) could grow as N and N^2, respectively, which indicates the total number of measurement samples in (40) could scale with N^4. This large number of measurement samples ultimately presents a challenge to attaining quantum advantage. A similar undesirable dependence of the measurement sample complexity on the scaling prefactors has also been encountered in [57], where the objective function entails expectations multiplied by prefactors exponentially in $\log N$, like the one in (9).

8 Numerical tests

The performance of the proposed doubly variational quantum OPF (Q-OPF) approach was validated using the IEEE 57-node power system benchmark [55]. For this benchmark, we have $N = 57$ primal variables and $M = 422$ dual variables. Therefore, the number of qubits needed to capture $\mathbf{v}(\boldsymbol{\theta})$ and $\boldsymbol{\lambda}(\boldsymbol{\phi})$ is 6 and 9, respectively. To simplify the implementation, nodes hosting generators were assumed to have no load, i.e., $p_n^d = q_n^d = 0$ for all $n \in \mathcal{N}_g$. Reactive power demands were set so that load nodes have $q_n^d = 0.33 \times p_n^d$ for all $n \in \mathcal{N}_l$. We generated 15 different instances of the OPF by scaling original load demands $\{(p_n^d, q_n^d)\}_{n\in\mathcal{N}_\ell}$ by factors drawn uniformly distributed within $[0.90, 1.05]$, independently per node and problem instance. Each of the 15 OPF instances was solved to global optimality using MATPOWER, a power system toolbox in MATLAB [58].

Selecting PQC architecture. To select the architecture of PQC_p, we numerically tested different architectures. Each candidate architecture was validated over all 15 problem instances. For each problem instance k, we found the OPF solution $\mathbf{v}_k^*$ using MATPOWER, and then solved the variational problem

$$\min_{\boldsymbol{\theta}_k} \left(1 - \frac{\langle \boldsymbol{\psi}(\boldsymbol{\theta}_k) | \mathbf{v}_k^* \rangle}{\|\mathbf{v}_k^*\|_2}\right)$$

Model	$\mathbf{x}$ [%]	$\boldsymbol{\lambda}$ [%]	(a)	(b)	(c)
QCQP-PD	14.95	19.30	25.87	13.73	0.49
QCQP-EG	13.67	19.20	24.53	12.29	0.39
QCQP_θ-PD	11.82	12.87	11.73	13.03	0.23
QCQP_θ-EG	7.62	12.17	11.53	11.86	0.21

Table 1: Average relative errors and constraint violation statistics across problem instances: (a) Number of constraint violations per instance; (b) maximum constraint violation [%]; (c) average constraint violation [%].

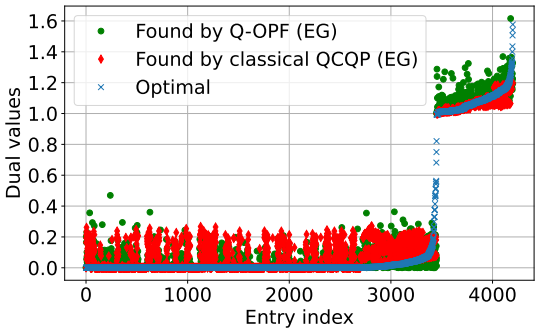

Figure 6: Comparing dual entries obtained by the Q-OPF and the classical QCQP, both solved by the EG iteration, to the optimal values found by MATPOWER. Dual vectors of 15 instances were concatenated and sorted in increasing order. Entries smaller than 10^{-6} were assigned to 0.

with $|\psi(\boldsymbol{\theta}_k)\rangle$ generated by the specified PQC_p architecture. Assuming the previous cost is optimized to global optimality, it measures the alignment between the generated quantum state and the actual OPF solution. The cost is inspired by [59] and is lower-bounded by zero. We tested architectures with different gate types and number of layers L, as detailed in Appendix A.5. We selected the PQC that attained the smallest average cost across all 15 OPF problem instances. A similar procedure was followed to determine the PQC_d architecture. The optimal PQC_p architecture consists of 10 layers, with each layer comprising the sequence $\mathbf{R}_y(\theta) - \text{CX} - \mathbf{R}_z(\theta') - \text{CX}$ gates, resulting in $P = 2 \times 10 \times 6 = 120$ parameters. The optimal PQC_d architecture is formed by 35 layers of $\mathbf{R}_y(\phi) - \text{CX}$ gates, yielding $Q = 1 \times 35 \times 9 = 315$ parameters. The average fitting errors obtained were 2×10^{-4} and 2×10^{-5} for PQC_p and PQC_d, respectively. This demonstrates that the selected PQC pairs are sufficiently expressive.

Algorithm details. PQC parameters $\boldsymbol{\theta}$ and $\boldsymbol{\phi}$ were initialized uniformly at random within the range $[0, 2\pi]$. Because voltages are normally around 1 per unit, the scaling variable α was initialized to $\sqrt{N}$; the variable β was initialized to $2\,|\mathcal{N}_\ell|$. Step sizes were set according to the exponentially decaying rule as $\mu_\theta^t = 0.015 \times 0.99985^t$, $\mu_\phi^t = 0.01 \times 0.99985^t$, and $\mu_\alpha^t = \mu_\beta^t = 10^{-5} \times 0.999^t$. Both PD and EG iterations were terminated when both conditions $\left\|\boldsymbol{\theta}^t - \boldsymbol{\theta}^{t-1}\right\|_2 \leq 10^{-6}$ and $\left\|\boldsymbol{\phi}^t - \boldsymbol{\phi}^{t-1}\right\|_2 \leq 10^{-6}$ were satisfied. Simulation scripts were written in Python and run on Pennylane's exact simulator [60].

We first explored the performance of Q-OPF in finding AC OPF solutions. Rather than evaluating errors on OPF solutions in terms of voltage phasors $\mathbf{v}$, we focused on finding the active power injection and voltage magnitude at all generator nodes. This is because power system operators are primarily interested in generator setpoints. Define the vector of generator setpoints $\mathbf{x} = [\mathbf{p}_g; \mathbf{v}_g]^\top$ with $\mathbf{p}_g$ and $\mathbf{v}_g$ collecting generator active power injections and voltage magnitudes, respectively. Generator active power injections can be computed from (44b) given $\mathbf{v}$ and the corresponding load demand p_n^d. In terms of dual variables, we focus on Lagrange multipliers λ_m associated with power balance and line limit constraints because these are used in electricity markets to compute the so-called *locational marginal prices*.

We compared solving (QCQP) and (QCQP_θ) in terms of the relative errors $\|\mathbf{x} - \mathbf{x}^*\|_2 / \|\mathbf{x}^*\|_2$ and $\|\boldsymbol{\lambda} - \boldsymbol{\lambda}^*\|_2 / \|\boldsymbol{\lambda}^*\|_2$ of the found generator setpoints $\mathbf{x}$ and dual variables $\boldsymbol{\lambda}$ from their global optimal solutions found using MATPOWER. The entries of the dual vector $\boldsymbol{\lambda}$ correspond to the Lagrangian multipliers associated with 280 power balance and line limit constraints. The non-convex (QCQP) was solved classically using the PD iteration in (3) and the EG iteration in (21) over voltage phasors $\mathbf{v}$ and dual variables $\boldsymbol{\lambda}$. Voltages were initialized at the flat voltage profile as $\mathbf{v}^0 = \mathbf{1}$, while the entries of $\boldsymbol{\lambda}$ were initialized independently at random based on the standard normal distribution and scaled by $2\,|\mathcal{N}_\ell|$. The related step sizes were set as $\mu_v = \mu_\lambda = 10^{-3} \times 0.9999^t$, and PD iterates were terminated when $\left\|\mathbf{v}^t - \mathbf{v}^{t-1}\right\|_2 \leq 10^{-6}$ and $\left\|\boldsymbol{\lambda}^t - \boldsymbol{\lambda}^{t-1}\right\|_2 \leq 10^{-6}$ were met.

Table 1 shows the relative errors for generator setpoints and dual variables attained by solving (QCQP_θ) and (QCQP) via the PD and EG iterations. Among the four models, solving (QCQP_θ) using the EG iteration achieved the smallest relative errors on $\mathbf{x}$ and $\boldsymbol{\lambda}$. Comparing algorithms when solving either (QCQP_θ) or (QCQP), the EG method consistently outperforms the PD method. Similarly, fixing the algorithm to be either EG or PD, solving (QCQP_θ) using a hybrid quantum-classical computer offers smaller errors compared to solving (QCQP) using the PD method.

Given that the EG iteration outperforms the PD iteration, Fig. 6 demonstrates the dual variables obtained by solving (QCQP_θ) and (QCQP) by the EG method.

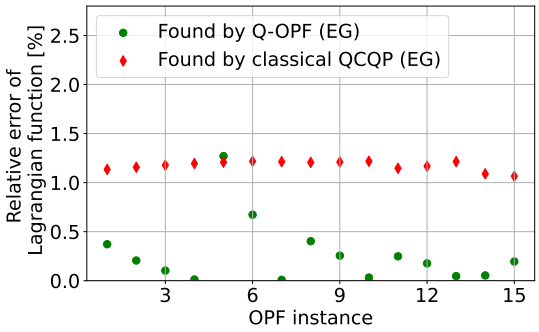

Figure 7: Comparing the relative error of the Lagrangian function for Q-OPF and the classical non-convex QCQP solved by the EG iteration across 15 OPF instances. The globally optimal value of the Lagrangian was computed using MATPOWER.

The Q-OPF approach can approximate dual variables reasonably well. This test corroborates that Q-OPF is capable of finding near-optimal AC OPF solutions.

Feasibility. Does the power system safely operate under the generator setpoints computed by Q-OPFs? To answer this question, for each instance, given the input load demands and the setpoints $\mathbf{x}$ obtained by Q-OPF, we found an AC power flow solution via MATPOWER. The obtained voltages were used to evaluate all 222 inequality constraints in (44). Violations in line current and generator constraints were normalized by their maximum limits. We employed three constraint violation metrics: *a)* the average number of constraint violations surpassing a normalized magnitude of 10^{-6} across OPF instances; *b)* the maximum constraint violation incurred across instances; *c)* the violations averaged over all constraints and instances. As shown in Table 1, solving (QCQP$_\theta$) using the EG iteration attains the best constraint violation metrics among the tested models.

Figure 7 shows the relative errors in terms of the Lagrangian function $|\mathcal{L} - P_\theta^*|/P_\theta^*$ for the last EG iteration. The relative errors of the Q-OPF are consistently below 1.5% and smaller than those of QCQP in 14 out of 15 instances. This test indicates that under the Q-OPF's generator setpoints, all physical constraints of the network are nearly satisfied.

9 Conclusions

This work has proposed a doubly variational quantum approach to solve the OPF, which is a large-scale non-convex QCQP. Primal/dual OPF variables are modeled by scaling the state and the PMF corresponding to two PQCs. These two PQCs are trained using exclusively the gradient estimates of the Lagrangian function. The training procedure aims to seek an approximate stationary point of the Lagrangian function. By adopting the XBM method and leveraging the sparsity patterns of power systems, the OPF observables are judiciously permuted, resulting in a banded form so that the Lagrangian function can be measured efficiently. Numerical tests on the IEEE 57-node benchmark power system using PennyLane's quantum simulator have demonstrated that: *i)* the EG iteration performs markedly better than the PD iteration; *ii)* solving (QCQP$_\theta$) on PQCs offers smaller errors in terms of the generator setpoint and dual variables, as well as smaller constraint violations compared to solving (QCQP) on a classical computer using PD; and *iii)* the Q-OPF found generator setpoints and dual variables with small errors, and under the generator setpoints obtained by the Q-OPF, the network constraints are nearly satisfied.

The proposed doubly variational quantum framework sets a solid foundation for exploring exciting directions: *d1)* The Q-OPF can be a building block to investigate other computational tasks, such as multi-period OPF, which entails the time-coupled ramp constraints on generators; *d2)* Thus far, the PQC parameters $(\boldsymbol{\theta}, \boldsymbol{\phi})$ have been optimized to solve one OPF instance, and parameters b_m have solely appeared on the diagonal of the observable $\mathbf{S}$. Nevertheless, vector $\mathbf{b}$ can be encoded as parameters of another unitary $\mathbf{T}(\mathbf{b})$, under which the corresponding primal and dual states are $|\boldsymbol{\psi}(\boldsymbol{\theta}, \mathbf{b})\rangle = \mathbf{T}(\mathbf{b})\mathbf{U}(\boldsymbol{\theta})|\mathbf{0}\rangle$ and $|\boldsymbol{\xi}(\boldsymbol{\phi}, \mathbf{b})\rangle = \mathbf{T}(\mathbf{b})\mathbf{V}(\boldsymbol{\phi})|\mathbf{0}\rangle$. Measuring the states $|\boldsymbol{\psi}(\boldsymbol{\theta}, \mathbf{b})\rangle$ and $|\boldsymbol{\xi}(\boldsymbol{\phi}, \mathbf{b})\rangle$ and utilizing the proposed Q-OPF for the Lagrangian in (9) evaluated across instances can solve multiple instances of the OPF; *d3)* Prudently construct an alternative objective function for (QCQP) to escape the problem of measurement samples scaling with $\bar{\alpha}$ and $\bar{\beta}$; and *d4)* Designing network-informed PQCs and/or testing the Q-OPF using dynamic PQCs and quantum recurrent embedding neural networks will be imperative for dealing with large-scale networks.

10 Acknowledgments

VK and TL acknowledge support from the US National Science Foundation under grant no. 2412947, and the US Office of Naval Research under grant N000142412614. MMW acknowledges support from the National Science Foundation under grant no. 2329662 and from the Cornell School of Electrical and Computer Engineering.

References

[1] Xiaoqing Bai, Hua Wei, Katsuki Fujisawa, and Yong Wang. "Semidefinite programming for opti-

mal power flow problems". *Intl. Journal of Electrical Power & Energy Syst.*, 30, pp. 383–392 (2008).

[2] Javad Lavaei and Steven H. Low. "Zero duality gap in optimal power flow problem". *IEEE Trans. Power Syst.*, 27, pp. 92–107 (2011).

[3] Stephen Boyd and Lieven Vandenberghe. "Convex Optimization". New York, NY: Cambridge University Press, 2004.

[4] Ruho Kondo, Yuki Sato, Satoshi Koide, Seiji Kajita, and Hideki Takamatsu. "Computationally Efficient Quantum Expectation with Extended Bell Measurements". *Quantum*, 6, p. 688 (2022).

[5] Uriel Feige. "Coping with the NP-hardness of the graph bandwidth problem". *Scandinavian Workshop on Algorithm Theory*, pp. 10–19 (2000).

[6] Elizabeth Cuthill and James McKee. "Reducing the bandwidth of sparse symmetric matrices". *Proc. National Conference*, pp. 157–172 (1969).

[7] Farshad Amani and Amin Kargarian. "Quantum-inspired optimal power flow". *Proc. IEEE Texas Power and Energy Conference*, pp. 1–6 (2024).

[8] Brynjar Sævarsson, Spyros Chatzivasileiadis, Hjörtur Jóhannsson, and Jacob Østergaard. "Quantum computing for power flow algorithms: Testing on real quantum computers", arXiv: 2204.14028 (2022).

[9] Fei Feng, Yifan Zhou, and Peng Zhang. "Quantum power flow". *IEEE Trans. Power Syst.*, 36, pp. 3810–3812 (2021).

[10] Junyu Liu, Han Zheng, Masanori Hanada, Kanav Setia, and Dan Wu. "Quantum power flows: From theory to practice". *Quantum Machine Intelligence*, 6, p. 55 (2024).

[11] Parikshit Pareek, Abhijith Jayakumar, Carleton Coffrin, and Sidhant Misra. "Demystifying quantum power flow: Unveiling the limits of practical quantum advantage", arXiv: 2402.08617 (2024).

[12] Ze Hu, Ziqing Zhu, Linghua Zhu, Xiang Wei, Siqi Bu, and Ka W. Chan. "Advancing Hybrid Quantum Neural Network for Alternative Current Optimal Power Flow", arXiv: 2410.20275 (2024).

[13] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'brien. "A variational eigenvalue solver on a photonic quantum processor". *Nature Communications*, 5, p. 4213 (2014).

[14] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. "A quantum approximate optimization algorithm", arXiv: 1411.4028 (2014).

[15] Andrew Lucas. "Ising formulations of many NP problems". *Front. in Phys.*, 2, pp. 1–15 (2014).

[16] Sarthak Gupta and Vassilis Kekatos. "A Quantum Approach for Stochastic Constrained Binary Optimization". *Proc. IEEE Intl. Conf. on Acoustics, Speech, and Signal Process.* (2023).

[17] Thinh V. Le and Vassilis Kekatos. "Solving Constrained Optimization Problems via the Variational Quantum Eigensolver with Constraints". *Physical Review A*, 110, pp. 1–19 (2024).

[18] Dhrumil Patel, Patrick J. Coles, and Mark M. Wilde. "Variational Quantum Algorithms for Semidefinite Programming". *Quantum*, 8, p. 1374 (2024).

[19] Jingxuan Chen, Hanna Westerheim, Zoë Holmes, Ivy Luo, Theshani Nuradha, Dhrumil Patel, Soorya Rethinasamy, Kathie Wang, and Mark M. Wilde. "Slack-variable approach for variational quantum semidefinite programming". *Physical Review A*, 112, p. 022607 (2025).

[20] Stuart Hadfield, Zhihu Wang, Bryan O'Gorman, Eleanor G. Rieffel, Davide Venturelli, and Rupak Biswas. "From the quantum approximate optimization algorithm to a quantum alternating operator ansatz". *Algorithms*, 12, p. 34 (2019).

[21] Kishor Bharti, Tobias Haug, Vlatko Vedral, and Leong-Chuan Kwek. "Noisy intermediate-scale quantum algorithm for semidefinite programming". *Physical Review A*, 105, p. 052445 (2022).

[22] Kishor Bharti and Tobias Haug. "Iterative quantum-assisted eigensolver". *Physical Review A*, 104, p. L050401 (2021).

[23] Maria Schuld and Francesco Petruccione. "Machine learning with quantum computers". Cham, Switzerland: Springer, 2021.

[24] Juan Carlos Garcia-Escartin and Pedro Chamorro-Posada. "SWAP test and Hong-Ou-Mandel effect are equivalent". *Physical Review A*, 87, p. 052330 (2013).

[25] Jarrod R. McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. "The theory of variational hybrid quantum-classical algorithms". *New Journal of Physics*, 18, p. 023023 (2016).

[26] Ophelia Crawford, Barnaby V. Straaten, Daochen Wang, Thomas Parks, Earl Campbell, and Stephen Brierley. "Efficient quantum measurement of Pauli operators in the presence of finite sampling error". *Quantum*, 5, p. 385 (2021).

[27] Michael A. Nielsen and Isaac L. Chuang. "Quantum Computation and Quantum Information". New York, NY: Cambridge University Press, 2000.

[28] Océane Koska, Marc Baboulin, and Arnaud Gazda. "A tree-approach Pauli decomposition algorithm with application to quantum computing". *Proc. ISC High Performance 2024 Research Paper*, pp. 1–11 (2024).

[29] Ramamurti Shankar. "Principles of Quantum Mechanics". Baden-Wurttemberg, Germany: Springer Science & Business Media, 2012.

[30] Vladyslav Verteletskyi, Tzu-Ching Yen, and Artur F. Izmaylov. "Measurement optimization in the variational quantum eigensolver using a minimum clique cover". *The Journal of Chemical Physics*, 152 (2020).

[31] Tzu-Ching Yen, Vladyslav Verteletskyi, and Artur F. Izmaylov. "Measuring all compatible operators in one series of single-qubit measurements using unitary transformations". *Journal of Chemical Theory and Computation*, 16, pp. 2400–2409 (2020).

[32] Pranav Gokhale, Olivia Angiuli, Yongshan Ding, Kaiwen Gui, Teague Tomesh, Martin Suchara, Margaret Martonosi, and Frederic T. Chong. "$O(N^3)$ Measurement Cost for Variational Quantum Eigensolver on Molecular Hamiltonians". *IEEE Transactions on Quantum Engineering*, 1, pp. 1–24 (2020).

[33] Artur F. Izmaylov, Tzu-Ching Yen, Robert A. Lang, and Vladyslav Verteletskyi. "Unitary partitioning approach to the measurement problem in the variational quantum eigensolver method". *Journal of Chemical Theory and Computation*, 16, pp. 190–195 (2019).

[34] Andrew Zhao, Andrew Tranter, William M. Kirby, Shu Fay Ung, Akimasa Miyake, and Peter J. Love. "Measurement reduction in variational quantum algorithms". *Physical Review A*, 101, p. 062322 (2020).

[35] Hsin-Yuan Huang, Richard Kueng, and John Preskill. "Predicting many properties of a quantum system from very few measurements". *Nature Physics*, 16, pp. 1050–1057 (2020).

[36] Vassilis Kekatos, Gang Wang, Hao Zhu, and Georgios B. Giannakis. "PSSE Redux: Convex Relaxation, Decentralized, Robust, and Dynamic Approaches". *Advances in Power System State Estimation*. Wiley, 2021, pp. 173–202. DOI: https://doi.org/10.1002/9781119480402.ch6.

[37] Angelia Nedić and Asuman Ozdaglar. "Subgradient methods for saddle-point problems". *Journal of Optimization Theory and Applications*, 142, pp. 205–228 (2009).

[38] Yuki Sato, Ruho Kondo, Satoshi Koide, Hideki Takamatsu, and Nobuyuki Imoto. "Variational quantum algorithm based on the minimum potential energy for solving the Poisson equation". *Physical Review A*, 104, p. 052409 (2021).

[39] Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. "Quantum circuit learning". *Physical Review A*, 98, p. 032309 (2018).

[40] Maria Schuld, Ville Bergholm, Christian Gogolin, Josh Izaac, and Nathan Killoran. "Evaluating analytic gradients on quantum hardware". *Physical Review A*, 99, p. 032331 (2019).

[41] Markku Kallio and Andrzej Ruszczyński. *Perturbation methods for saddle point computation*. Tech. rep. International Institute for Applied Systems Analysis, Laxenburg, Austria: WP-94-038, 1994. URL: `https://pure.iiasa.ac.at/id/eprint/4171/`.

[42] Hirofumi Uzawa. "Iterative methods for concave programming". *Studies in Linear and Nonlinear Programming*, 6, pp. 154–165 (1958).

[43] Jayash Koshal, Angelia Nedić, and Uday V. Shanbhag. "Multiuser optimization: Distributed algorithms and error analysis". *SIAM Journal on Optimization*, 21, pp. 1046–1081 (2011).

[44] Tianyi Lin, Chi Jin, and Michael Jordan. "On gradient descent ascent for nonconvex-concave minimax problems". *Proc. Intl. Conf. on Machine Learning*, pp. 6083–6093 (2020).

[45] Galina M. Korpelevich. "The extragradient method for finding saddle points and other problems". *Matecon*, 12, pp. 747–756 (1976).

[46] Jelena Diakonikolas, Constantinos Daskalakis, and Michael I. Jordan. "Efficient methods for structured nonconvex-nonconcave min-max optimization". *Proc. Intl. Conf. on Artificial Intelligence and Statistics*, pp. 2746–2754 (2021).

[47] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. "Barren plateaus in quantum neural network training landscapes". *Nature Communications*, 9, p. 4812 (2018).

[48] Martin Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo. "Barren plateaus in variational quantum computing". *Nature Reviews Physics*, 7, pp. 174–189 (2025).

[49] Marco Cerezo, Martin Larocca, Diego García-Martín, Nelson L. Diaz, Paolo Braccia, Enrico Fontana, Manuel S. Rudolph, Pablo Bermejo, Aroosa Ijaz, Supanut Thanasilp, et al. "Does provable absence of barren plateaus imply classical simulability? Or, why we need to rethink variational quantum computing". *Nature Communications*, 16, p. 7907 (2025).

[50] Abhinav Deshpande, Marcel Hinsche, Sona Najafi, Kunal Sharma, Ryan Sweke, and Christa Zoufal. "Dynamic parameterized quantum circuits: expressive and barren-plateau free", arXiv: 2411.05760 (2024).

[51] Mingrui Jing, Erdong Huang, Xiao Shi, Shengyu Zhang, and Xin Wang. "Quantum Recurrent Embedding Neural Network", arXiv: 2506.13185 (2025).

[52] Krishna M. Sambarapu and S. Mark Halpin. "Sparse matrix techniques in power systems". *Proc. Southeastern Symposium on System Theory*, pp. 194–198 (2007).

[53] Alexander J. Flueck and Hsiao-Dong Chiang. "Solving the nonlinear power flow equations with an inexact Newton method using GMRES". *IEEE Trans. on Power Syst.*, 13, pp. 267–273 (1998).

[54] Ganqu Cui, Wenjian Yu, Xin Li, Zhiyu Zeng, and Ben Gu. "Machine-learning-driven matrix ordering for power grid analysis". *Proc. Design, Automation & Test in Europe Conference & Exhibition*, pp. 984–987 (2019).

[55] Sogol Babaeinejadsarookolaee, Adam Birchfield, Richard D. Christie, Carleton Coffrin, Christopher DeMarco, Ruisheng Diao, Michael Ferris, Stephane Fliscounakis, Scott Greene, Renke Huang, et al. "The power grid library for benchmarking AC optimal power flow algorithms", arXiv: 1908.02788 (2019).

[56] Aram W. Harrow and John C. Napp. "Low-depth gradient measurements can improve convergence in variational hybrid quantum-classical algorithms". *Phys. Rev. Letters*, 126, p. 140502 (2021).

[57] Abhijat Sarma, Thomas W. Watts, Mudassir Moosa, Yilian Liu, and Peter L. McMahon. "Quantum variational solving of nonlinear and multidimensional partial differential equations". *Physical Review A*, 109, p. 062616 (2024).

[58] Ray D. Zimmerman, Carlos E. Murillo-Sánchez, and Robert J. Thomas. "MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education". *IEEE Trans. Power Syst.*, 26, pp. 12–19 (2010).

[59] Carlos Bravo-Prieto, Ryan LaRose, M. Cerezo, Yigit Subasi, Lukasz Cincio, and Patrick J. Coles. "Variational Quantum Linear Solver". *Quantum*, 7, p. 1188 (2023).

[60] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, M. Sohaib Alam, Shahnawaz Ahmed, Juan Miguel Arrazola, Carsten Blank, Alain Delgado, Soran Jahangiri, Keri McKiernan, Johannes Jakob Meyer, Zeyue Niu, Antal Szava, and Nathan Killoran. "PennyLane: Automatic differentiation of hybrid quantum-classical computations", arXiv: 1811.04968 (2022).

[61] Guillaume Garrigos and Robert M. Gower. "Handbook of convergence theorems for (stochastic) gradient methods", arXiv: 2301.11235 (2023).

[62] Andrea Mari, Thomas R. Bromley, and Nathan Killoran. "Estimating the gradient and higher-order derivatives on quantum hardware". *Physical Review A*, 103, p. 012405 (2021).

A Appendix

A.1 Power system modeling

A transmission line connecting node n to node m is represented by its series conductance G_{nm} and series susceptance B_{nm}. Given G_{nm} and B_{nm}, power injections at node n are computed through nodal voltages based upon the following quadratic power flow equations:

$$p_n = \sum_{m=1}^{N} v_n^r(v_m^r G_{nm} - v_m^i B_{nm}) + v_n^i(v_m^i G_{nm} + v_m^r B_{nm}),$$

$$q_n = \sum_{m=1}^{N} v_n^i(v_m^r G_{nm} - v_m^i B_{nm}) - v_n^r(v_m^i G_{nm} + v_m^r B_{nm}),$$

where (v_n^r, v_n^i) and (v_m^r, v_m^i) are the real and imaginary components of the voltages at nodes n and m, respectively. Define the $N \times N$ node admittance matrix as $\mathbf{Y} \coloneqq \mathbf{G} + \iota\mathbf{B}$, where $\mathbf{G}$ collects G_{nm} and $\mathbf{B}$ collects

B_{nm} of the network lines. Then the power flow equations can be compactly represented as

$$p_n = \mathbf{v}^\dagger \mathbf{M}_{p_n} \mathbf{v}, \tag{41}$$

$$q_n = \mathbf{v}^\dagger \mathbf{M}_{q_n} \mathbf{v}, \tag{42}$$

where $\mathbf{M}_{p_n}$ and $\mathbf{M}_{q_n}$ are $N \times N$ Hermitian matrices taken the forms as

$$\mathbf{M}_{p_n} \coloneqq \frac{1}{2}(\mathbf{Y}^\dagger \mathbf{e}_n \mathbf{e}_n^\top + \mathbf{e}_n \mathbf{e}_n^\top \mathbf{Y}),$$

$$\mathbf{M}_{q_n} \coloneqq \frac{1}{2\iota}(\mathbf{Y}^\dagger \mathbf{e}_n \mathbf{e}_n^\top - \mathbf{e}_n \mathbf{e}_n^\top \mathbf{Y}).$$

Similarly, the squared voltage magnitude at node n is represented as

$$(v_n^r)^2 + (v_n^i)^2 = \mathbf{v}^\dagger \mathbf{M}_{v_n} \mathbf{v} \tag{43}$$

where $\mathbf{M}_{v_n} \coloneqq \mathbf{e}_n \mathbf{e}_n^\top$. Given a line $(n, m) \in \mathcal{E}$ with the series admittance Y_{nm}, the complex current flowing on from node n to node m is provided as $i_{nm} = Y_{nm}(v_n - v_m)$. This poses $|i_{nm}| = \mathbf{v}^\dagger \mathbf{M}_{i_{nm}} \mathbf{v}$ where $\mathbf{M}_{i_{nm}} \coloneqq |Y_{nm}| \, (\mathbf{e}_n - \mathbf{e}_m)(\mathbf{e}_n - \mathbf{e}_m)^\top$.

The power injection (p_n, q_n) consists of a dispatchable constituent (p_n^g, q_n^g) and an inflexible constituent (p_n^d, q_n^d) such that $p_n = p_n^g - p_n^d$ and $q_n = q_n^g - q_n^d$. The former represents the power dispatch of a generator or a flexible load placed at node n, and the latter expresses the inelastic load supplied at node n. Let nodes with dispatchable power injections belong to a subset $\mathcal{N}_g \subseteq \mathcal{N}$. The remaining nodes hosted inelastic loads form the subset $\mathcal{N}_\ell = \mathcal{N} \backslash \mathcal{N}_g$. Node $n = 1$ is referred to as the angle reference node, under which it follows $\mathbf{v}^\dagger \mathbf{e}_1 \mathbf{e}_1^\top \mathbf{v} = 1$, where $\mathbf{e}_1$ is the first column of the identity. To simplify the exposition, each node in $\mathcal{N}_g$ is assumed to host only one dispatchable unit. Given the inflexible loads $\{p_n^d, q_n^d\}_{n \in \mathcal{N}}$, the OPF's objective is to minimize the dispatch cost of generators and flexible load while respecting the resource and network limits. The OPF is formulated as the ensuing quadratically constrained quadratic program (QCQP) over nodal voltages and power of generators:

$$\min \quad \sum_{n \in \mathcal{N}_g} c_n p_n^g \tag{44a}$$

$$\text{over} \quad \mathbf{v} \in \mathbb{C}^N, \{p_n^g, q_n^g\}_{n \in \mathcal{N}_g}$$

$$\text{s.to} \quad \mathbf{v}^\dagger \mathbf{M}_{p_n} \mathbf{v} = p_n^g - p_n^d \qquad \forall n \in \mathcal{N}_g \tag{44b}$$

$$\mathbf{v}^\dagger \mathbf{M}_{q_n} \mathbf{v} = q_n^g - q_n^d \qquad \forall n \in \mathcal{N}_g \tag{44c}$$

$$\mathbf{v}^\dagger \mathbf{M}_{p_n} \mathbf{v} = -p_n^d \qquad \forall n \in \mathcal{N}_\ell \tag{44d}$$

$$\mathbf{v}^\dagger \mathbf{M}_{q_n} \mathbf{v} = -q_n^d \qquad \forall n \in \mathcal{N}_\ell \tag{44e}$$

$$\underline{p}_n^g \leq \mathbf{v}^\dagger \mathbf{M}_{p_n} \mathbf{v} + p_n^d \leq \bar{p}_n^g \qquad \forall n \in \mathcal{N}_g \tag{44f}$$

$$\underline{q}_n^g \leq \mathbf{v}^\dagger \mathbf{M}_{q_n} \mathbf{v} + q_n^d \leq \bar{q}_n^g \qquad \forall n \in \mathcal{N}_g \tag{44g}$$

$$\underline{v}_n \leq \mathbf{v}^\dagger \mathbf{M}_{v_n} \mathbf{v} \leq \bar{v}_n \qquad \forall n \in \mathcal{N} \tag{44h}$$

$$\mathbf{v}^\dagger \mathbf{e}_1 \mathbf{e}_1^\top \mathbf{v} = 1 \tag{44i}$$

$$\mathbf{v}^\dagger \mathbf{M}_{i_{nm}} \mathbf{v} \leq \bar{i}_{nm} \qquad \forall (n, m) \in \mathcal{E} \tag{44j}$$

Constraints (44b)–(44e) capture the power flow equations at generator and load nodes. Constraints (44f)–(44g) confine limits of generators and flexible loads. Constraint (44h) limits the squared voltage magnitudes within the specified range, and constraint (44i) represents the squared magnitude of the angle reference node. Constraint (44j) enforces the line current magnitude within the line ratings. By substituting $\{p_n^g, q_n^g\}$ in (44b)–(44c) into the objective of (44), the OPF model is represented as a QCQP in terms of the nodal voltage vector $\mathbf{v}$.

A.2 Proof of Lemma 1

The claim follows readily by substituting $\mathbf{M}$ into the observable as

$$\begin{aligned}
\langle \boldsymbol{\xi}, \boldsymbol{\psi} | \mathbf{M} | \boldsymbol{\xi}, \boldsymbol{\psi} \rangle &= (\boldsymbol{\xi} \otimes \boldsymbol{\psi})^\dagger \left(\sum_{m=1}^{M} \mathbf{e}_m \mathbf{e}_m^\top \otimes \mathbf{M}_m \right) (\boldsymbol{\xi} \otimes \boldsymbol{\psi}) \\
&= \sum_{m=1}^{M} (\boldsymbol{\xi} \otimes \boldsymbol{\psi})^\dagger \left(\mathbf{e}_m \mathbf{e}_m^\top \otimes \mathbf{M}_m \right) (\boldsymbol{\xi} \otimes \boldsymbol{\psi}) \\
&= \sum_{m=1}^{M} \left(\boldsymbol{\xi}^\dagger \mathbf{e}_m \mathbf{e}_m^\top \boldsymbol{\xi} \right) \otimes \left(\boldsymbol{\psi}^\dagger \mathbf{M}_m \boldsymbol{\psi} \right) \\
&= \sum_{m=1}^{M} |\xi_m(\boldsymbol{\phi})|^2 F_m(\boldsymbol{\theta}).
\end{aligned}$$

A.3 Proof of Lemma 3

Quantum observables are trigonometric functions of $(\boldsymbol{\theta}, \boldsymbol{\phi})$. Moreover, the Lagrangian function $\mathcal{L}(\mathbf{z})$ is quadratic in (α, β). Therefore, the operator $\mathbf{g}(\mathbf{z})$ is continuously differentiable with respect to $\mathbf{z}$. Let $\mathbf{J}(\mathbf{z}) \coloneqq \nabla_{\mathbf{z}} \mathbf{g}(\mathbf{z})$ denote the Jacobian matrix of $\mathbf{g}(\mathbf{z})$. If the spectral norm of $\mathbf{J}(\mathbf{z})$ is upper bounded by a constant L, then $\mathbf{g}(\mathbf{z})$ is L-Lipschitz continuous [61, Lemma 2.6]. Based on matrix norm inequalities, the spectral norm of $\mathbf{J}(\mathbf{z})$ can be bounded as:

$$\|\mathbf{J}(\mathbf{z})\| \leq \sqrt{\|\mathbf{J}(\mathbf{z})\|_1 \, \|\mathbf{J}(\mathbf{z})\|_\infty} \tag{45}$$

where $\|\mathbf{J}(\mathbf{z})\|_1$ is the maximum absolute column sum of $\mathbf{J}(\mathbf{z})$, and $\|\mathbf{J}(\mathbf{z})\|_\infty$ is the maximum absolute row sum. Because $\mathcal{L}(\mathbf{z})$ has continuous second-order partial derivatives, Schwarz's theorem predicates that $\frac{\partial^2 \mathcal{L}}{\partial z_i \partial z_j} = \frac{\partial^2 \mathcal{L}}{\partial z_j \partial z_i}$, so that $\|\mathbf{J}(\mathbf{z})\|_1 = \|\mathbf{J}(\mathbf{z})\|_\infty$ and (45) simplifies

as

$$\|\mathbf{J}(\mathbf{z})\| \leq \|\mathbf{J}(\mathbf{z})\|_\infty = \max_{j\in\{1,\dots,P+Q+2\}} \left\|\mathbf{e}_j^\top \mathbf{J}(\mathbf{z})\right\|_1 \quad (46)$$

where $\mathbf{e}_j$ is the j-th column of the identity.

Consider first the first P rows of $\mathbf{J}(\mathbf{z})$ corresponding to the partial Hessian matrix $\nabla^2_{\boldsymbol{\theta},\mathbf{z}}\mathcal{L}(\mathbf{z})$. More specifically, the absolute row sum for row $j \in \{1,\dots,P\}$ is

$$\begin{aligned}\left\|\mathbf{e}_j^\top \mathbf{J}(\mathbf{z})\right\|_1 = &\left(\sum_{p=1}^{P}\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\theta_p}\right|\right) + \left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\alpha}\right| \\ &+ \left(\sum_{q=1}^{Q}\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\phi_q}\right|\right) + \left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\beta}\right|. \quad (47)\end{aligned}$$

Applying the triangle inequality on each summand of the first term in (47) yields

$$\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\theta_p}\right| \leq \alpha^2\left|\frac{\partial^2 F_0(\boldsymbol{\theta})}{\partial\theta_i\partial\theta_p}\right| + \alpha^2\beta^2\sum_{m=1}^{M}|\xi_m(\boldsymbol{\phi})|^2\left|\frac{\partial^2 F_m(\boldsymbol{\theta})}{\partial\theta_i\partial\theta_p}\right|$$

The second-order partial derivatives of quantum expectations can be computed according to the PSR as [62]:

$$\begin{aligned}\left|\frac{\partial^2 F_m(\boldsymbol{\theta})}{\partial\theta_j\partial\theta_p}\right| &= \frac{1}{4}\Big|F_m\left(\boldsymbol{\theta}_j^+ + \frac{\pi}{2}\mathbf{e}_p\right) - F_m\left(\boldsymbol{\theta}_j^+ - \frac{\pi}{2}\mathbf{e}_p\right) \\ &\quad - F_m\left(\boldsymbol{\theta}_j^- + \frac{\pi}{2}\mathbf{e}_p\right) + F_m\left(\boldsymbol{\theta}_j^- - \frac{\pi}{2}\mathbf{e}_p\right)\Big| \\ &\leq \frac{1}{4}\left|F_m\left(\boldsymbol{\theta}_j^+ + \frac{\pi}{2}\mathbf{e}_p\right)\right| + \frac{1}{4}\left|F_m\left(\boldsymbol{\theta}_j^+ - \frac{\pi}{2}\mathbf{e}_p\right)\right| \\ &\quad + \frac{1}{4}\left|F_m\left(\boldsymbol{\theta}_j^- + \frac{\pi}{2}\mathbf{e}_p\right)\right| + \frac{1}{4}\left|F_m\left(\boldsymbol{\theta}_j^- - \frac{\pi}{2}\mathbf{e}_p\right)\right| \\ &\leq \|\mathbf{M}_m\|\end{aligned}$$

for all m, p, and j.

Based on the latter bounds, we get that:

$$\begin{aligned}\sum_{p=1}^{P}\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\theta_p}\right| &\leq P\alpha^2\left(\|\mathbf{M}_0\| + \beta^2\sum_{m=1}^{M}|\xi_m(\boldsymbol{\phi})|^2\|\mathbf{M}_m\|\right) \\ &\leq P\alpha^2\left(\|\mathbf{M}_0\| + \beta^2\langle\boldsymbol{\xi}(\boldsymbol{\phi})|\boldsymbol{\xi}(\boldsymbol{\phi})\rangle\max_m\|\mathbf{M}_m\|\right) \\ &= P\alpha^2\left(\|\mathbf{M}_0\| + \beta^2\max_m\|\mathbf{M}_m\|\right). \quad (48)\end{aligned}$$

The second term of (47) can be upper bounded similarly using the PSR for first-order partial derivatives:

$$\begin{aligned}\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\alpha}\right| &\leq 2\alpha\left|\frac{\partial F_0(\boldsymbol{\theta})}{\partial\theta_j}\right| + 2\alpha\beta^2\sum_{m=1}^{M}|\xi_m(\boldsymbol{\phi})|^2\left|\frac{\partial F_m(\boldsymbol{\theta})}{\partial\theta_j}\right| \\ &\leq 2\alpha\|\mathbf{M}_0\| + 2\alpha\beta^2\sum_{m=1}^{M}|\xi_m(\boldsymbol{\phi})|^2\|\mathbf{M}_m\| \\ &\leq 2\alpha\|\mathbf{M}_0\| + 2\alpha\beta^2\langle\boldsymbol{\xi}(\boldsymbol{\phi})|\boldsymbol{\xi}(\boldsymbol{\phi})\rangle\max_m\|\mathbf{M}_m\| \\ &\leq 2\alpha\|\mathbf{M}_0\| + 2\alpha\beta^2\max_m\|\mathbf{M}_m\|. \quad (49)\end{aligned}$$

The summands in the third term of (47) can be upper bounded as:

$$\begin{aligned}\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\phi_q}\right| &\leq \alpha^2\beta^2\sum_{m=1}^{M}\left|\frac{\partial|\xi_m(\boldsymbol{\phi})|^2}{\partial\phi_q}\right|\left|\frac{\partial F_m(\boldsymbol{\theta})}{\partial\theta_j}\right| \\ &\leq \frac{\alpha^2\beta^2}{2}\sum_{m=1}^{M}\left||\xi_m(\boldsymbol{\phi}_q^+)|^2 - |\xi_m(\boldsymbol{\phi}_q^-)|^2\right|\cdot\|\mathbf{M}_m\| \\ &\leq \frac{\alpha^2\beta^2}{2}\left(\sum_{m=1}^{M}|\xi_m(\boldsymbol{\phi}_q^+)|^2 + \sum_{m=1}^{M}|\xi_m(\boldsymbol{\phi}_q^-)|^2\right)\|\mathbf{M}_m\| \\ &\leq \alpha^2\beta^2\max_m\|\mathbf{M}_m\|\end{aligned}$$

for all q. Then, the third term of (47) is bounded as

$$\sum_{q=1}^{Q}\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\phi_q}\right| \leq Q\alpha^2\beta^2\max_m\|\mathbf{M}_m\|. \quad (50)$$

The fourth term in (47) can be bounded as

$$\begin{aligned}\left|\frac{\partial^2\mathcal{L}(\mathbf{z})}{\partial\theta_j\partial\beta}\right| &\leq 2\alpha^2\beta\sum_{m=1}^{M}|\xi_m(\boldsymbol{\phi})|^2\left|\frac{\partial F_m(\boldsymbol{\theta})}{\partial\theta_j}\right| \\ &\leq 2\alpha^2\beta\max_m\|\mathbf{M}_m\|. \quad (51)\end{aligned}$$

Combining (48)–(51) yields

$$\begin{aligned}\max_{j\in\{1,\dots,P\}}\left\|\mathbf{e}_j^\top\mathbf{J}(\mathbf{z})\right\|_1 &\leq (P\alpha^2 + 2\alpha)\|\mathbf{M}_0\| \\ &+ \left(P\alpha^2\beta^2 + Q\alpha^2\beta^2 + 2\alpha\beta^2 + 2\alpha^2\beta\right)\max_m\|\mathbf{M}_m\|. \quad (52)\end{aligned}$$

For the remaining rows of $\mathbf{J}(\mathbf{z})$ corresponding to the partial Hessian matrices $\nabla^2_{\alpha,\mathbf{z}}\mathcal{L}(\mathbf{z})$, $-\nabla^2_{\boldsymbol{\phi},\mathbf{z}}\mathcal{L}(\mathbf{z})$, and $-\nabla^2_{\beta,\mathbf{z}}\mathcal{L}(\mathbf{z})$, the norms $\left\|\mathbf{e}_j^\top\mathbf{J}(\mathbf{z})\right\|_1$ can be upper bounded using similar arguments. A key point here is that $\|\mathbf{M}\| = \max_m\|\mathbf{M}_m\|$ because $\mathbf{M}$ in (17) is block diagonal. To avoid repetition and tedious algebra, we omit the detailed derivations and present the final bounds.

$$\begin{aligned}\|\mathbf{e}_{P+1}^\top\mathbf{J}(\mathbf{z})\|_1 &\leq (2P\alpha + 2)\|\mathbf{M}_0\| \\ &+ (2P\alpha\beta^2 + 4Q\alpha\beta^2 + 2\beta^2 + 4\alpha\beta)\max_m\|\mathbf{M}_m\|, \quad (53)\end{aligned}$$

$$\begin{aligned}\max_{j\in\{P+2:P+Q+1\}}\left\|\mathbf{e}_j^\top\mathbf{J}(\mathbf{z})\right\|_1 &\leq (Q\beta^2 + 2\beta)\max_m|b_m| \\ &+ \left(P\alpha^2\beta^2 + Q\alpha^2\beta^2 + 2\alpha\beta^2 + 2\alpha^2\beta\right)\max_m\|\mathbf{M}_m\|, \quad (54)\end{aligned}$$

$$\|\mathbf{e}_{P+Q+2}^\top\mathbf{J}(\mathbf{z})\|_1 \leq (2Q\beta + 2)\max_m|b_m|$$

$$+ \left(P\alpha^2\beta + 2Q\alpha^2\beta + 4\alpha\beta + 2\alpha^2\right) \max_m \|\mathbf{M}_m\| . \quad (55)$$

According to (46), the Lipschitz constant L of $\mathbf{g}(\mathbf{z})$ is the maximum of the right-hand sides of (52)–(55). If $\alpha \geq 4$, the right-hand side (RHS) of (52) is larger than the RHS of (53). The condition $\alpha \geq 4$ holds trivially if the number of power system nodes is $N > 32$. Moreover, if $\beta \geq 2$, the RHS of (54) is larger than that of (55). The condition $\beta \geq 2$ is expected to occur for problems with many constraints. Therefore, we are left with the RHS of (52) and (54), whose maximum can be compactly expressed as

$$\left(P\alpha^2\beta^2 + Q\alpha^2\beta^2 + 2\alpha\beta^2 + 2\alpha^2\beta\right) \max_m \|\mathbf{M}_m\| + \max\{\left(P\alpha^2 + 2\alpha\right) \|\mathbf{M}_0\| , \left(Q\beta^2 + 2\beta\right) \max_m |b_m|\}.$$

Then, the claim of Lemma 3 follows readily.

A.4 Proof of Lemma 4

We first upper bound the variance of the sample-average estimators of the three terms in (9). Note that the three terms in (9) can be decomposed into quantum expectations across c. This is obvious for $F_0(\boldsymbol{\theta})$ and $G(\boldsymbol{\phi})$. Regarding $F(\boldsymbol{\theta}, \boldsymbol{\phi})$, we obtain from (28):

$$F(\boldsymbol{\theta}, \boldsymbol{\phi}) = \sum_{c=1}^{2C-1} \sum_{m=1}^{M} |\xi_m(\boldsymbol{\phi})|^2 \langle \boldsymbol{\psi}_c(\boldsymbol{\theta}) | \boldsymbol{\Lambda}_m^c | \boldsymbol{\psi}_c(\boldsymbol{\theta}) \rangle$$
$$= \sum_{c=1}^{2C-1} \langle \boldsymbol{\xi}(\boldsymbol{\phi}), \boldsymbol{\psi}_c(\boldsymbol{\theta}) | \left(\sum_{m=1}^{M} \mathbf{e}_m \mathbf{e}_m^\top \otimes \boldsymbol{\Lambda}_m^c \right) | \boldsymbol{\xi}(\boldsymbol{\phi}), \boldsymbol{\psi}_c(\boldsymbol{\theta}) \rangle .$$

Because the three expectations in (9) are defined over different states and observables, we consider a general expectation $E \coloneqq \langle \boldsymbol{\zeta} | \mathbf{N} | \boldsymbol{\zeta} \rangle$, where $|\boldsymbol{\zeta}\rangle$ and $\mathbf{N}$ are dimensionally compatible and $\mathbf{N}$ is a Hermitian matrix with the same color decomposition as OPF matrices $\mathbf{M}_m$'s. The expectation E can be color-decomposed as

$$E = \sum_{c=1}^{2C-1} \langle \boldsymbol{\zeta} | \mathbf{N}^c | \boldsymbol{\zeta} \rangle = \sum_{c=1}^{2C-1} \langle \boldsymbol{\zeta}_c | \boldsymbol{\Lambda}^c | \boldsymbol{\zeta}_c \rangle , \quad (56)$$

where $\mathbf{N}^c = \mathbf{U}_c \boldsymbol{\Lambda}^c \mathbf{U}_c^\dagger$ is the eigendecomposition of $\mathbf{N}^c$ and $|\boldsymbol{\zeta}_c\rangle = \mathbf{U}_c^\dagger |\boldsymbol{\zeta}\rangle$.

Upon measuring $|\boldsymbol{\zeta}_c\rangle$ in the computational basis using S measurement samples, the binary outcome $|i\rangle$ is observed with probability $|\zeta_c^i|^2 \coloneqq |\langle i | \boldsymbol{\zeta}_c \rangle|^2$. For each sample s, define a random variable $\hat{E}_s^c$, taking the value $\langle i | \boldsymbol{\Lambda}^c | i \rangle$ when outcome $|i\rangle$ is observed while sampling $|\boldsymbol{\zeta}_c\rangle$. Subsequently, the sample-average estimator of E can be expressed as

$$\hat{E} = \frac{1}{S} \sum_{c=1}^{2C-1} \sum_{s=1}^{S} \hat{E}_s^c. \quad (57)$$

The variance of $\hat{E}$ is bounded in the next lemma.

Lemma 5. *The variance of the estimator in (57) can be computed as:*

$$\mathrm{Var}(\hat{E}) = \frac{1}{S} \sum_{c=1}^{2C-1} \left(\langle \boldsymbol{\zeta}_c | (\boldsymbol{\Lambda}^c)^2 | \boldsymbol{\zeta}_c \rangle - \langle \boldsymbol{\zeta}_c | \boldsymbol{\Lambda}^c | \boldsymbol{\zeta}_c \rangle^2 \right), \quad (58)$$

and can be upper bounded as

$$\mathrm{Var}(\hat{E}) \leq \frac{1}{S} \sum_{c=1}^{2C-1} \|\mathbf{N}^c\|^2 . \quad (59)$$

Proof. The variance of each $\hat{E}_s^c$ can be computed and upper bounded as follows:

$$\begin{aligned}
&\mathrm{Var}(\hat{E}_s^c) = \mathbb{E}[(\hat{E}_s^c)^2] - (\mathbb{E}[\hat{E}_s^c])^2 \\
&= \sum_{i=0}^{N-1} |\zeta_c^i|^2 \langle i | \boldsymbol{\Lambda}^c | i \rangle^2 - \left(\sum_{i=0}^{N-1} |\zeta_c^i|^2 \langle i | \boldsymbol{\Lambda}^c | i \rangle \right)^2 \\
&= \langle \boldsymbol{\zeta}_c | (\boldsymbol{\Lambda}^c)^2 | \boldsymbol{\zeta}_c \rangle - \langle \boldsymbol{\zeta}_c | \boldsymbol{\Lambda}^c | \boldsymbol{\zeta}_c \rangle^2 \\
&\leq \langle \boldsymbol{\zeta}_c | (\boldsymbol{\Lambda}^c)^2 | \boldsymbol{\zeta}_c \rangle \\
&\leq \|\mathbf{N}^c\|^2.
\end{aligned}$$

Because the random variables $\hat{E}_s^c$ are independent across c and s, we get that

$$\mathrm{Var}(\hat{E}) = \frac{1}{S^2} \sum_{c=1}^{2C-1} \sum_{s=1}^{S} \mathrm{Var}(\hat{E}_s^c) = \frac{1}{S} \sum_{c=1}^{2C-1} \mathrm{Var}(\hat{E}_s^c)$$

for any s. Substituting $\mathrm{Var}(\hat{E}_s^c)$ above yields the expression in (58). The final bound in (59) follows readily. □

We are now ready to upper bound the variance of the estimate $\hat{\mathbf{g}}(\mathbf{z})$ per Theorem 1.

Proof of Lemma 4. The RHS of (36) can be expanded as

$$\begin{aligned}
&\mathbb{E}\left[\|\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})\|_2^2 \right] = \\
&\underbrace{\mathbb{E}[|\hat{\nabla}_\alpha \mathcal{L}(\mathbf{z}) - \nabla_\alpha \mathcal{L}(\mathbf{z})|^2]}_{\coloneqq \mathrm{Var}(\hat{\nabla}_\alpha \mathcal{L}(\mathbf{z}))} + \mathbb{E}[\|\hat{\nabla}_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{z}) - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{z})\|_2^2] \\
&+ \mathbb{E}[\|\hat{\nabla}_{\boldsymbol{\phi}} \mathcal{L}(\mathbf{z}) - \nabla_{\boldsymbol{\phi}} \mathcal{L}(\mathbf{z})\|_2^2] + \underbrace{\mathbb{E}[|\hat{\nabla}_\beta \mathcal{L}(\mathbf{z}) - \nabla_\beta \mathcal{L}(\mathbf{z})|^2]}_{\coloneqq \mathrm{Var}(\hat{\nabla}_\beta \mathcal{L}(\mathbf{z}))} .
\end{aligned} \quad (60)$$

Among four terms in (60), the terms $\mathrm{Var}(\hat{\nabla}_\alpha \mathcal{L}(\mathbf{z}))$ and $\mathrm{Var}(\hat{\nabla}_\beta \mathcal{L}(\mathbf{z}))$ are simpler to bound. Let $\hat{F}_0(\boldsymbol{\theta})$ and $\hat{F}(\boldsymbol{\theta}, \boldsymbol{\phi})$ denote the sample-based estimates of $F_0(\boldsymbol{\theta})$ and $F(\boldsymbol{\theta}, \boldsymbol{\phi})$, respectively. Then, we can express the first term in (60) as:

$$\mathrm{Var}(\hat{\nabla}_\alpha \mathcal{L}(\mathbf{z})) = \mathrm{Var}\left(2\alpha \hat{F}_0(\boldsymbol{\theta}) + 2\alpha\beta^2 \hat{F}(\boldsymbol{\theta}, \boldsymbol{\phi}) \right)$$

No.	architecture of 1 layer	L (primal)	P	error (primal)	L (dual)	Q	error (dual)
1	$\mathbf{R}_x$−CX	20	120	9×10^{-1}	35	315	8×10^{-5}
2	$\mathbf{R}_y$−CX	20	120	5×10^{-3}	35	315	2×10^{-5}
3	$\mathbf{R}_z$−CX	20	120	1	35	315	2×10^{-3}
4	$\mathbf{R}_x$−CX−$\mathbf{R}_y$−CX	10	120	3×10^{-4}	18	324	7×10^{-5}
5	$\mathbf{R}_x$−CX−$\mathbf{R}_z$−CX	10	120	3×10^{-4}	18	324	7×10^{-5}
6	$\mathbf{R}_y$−CX−$\mathbf{R}_z$−CX	10	120	2×10^{-4}	18	324	7×10^{-5}
7	$\mathbf{R}_x$−CX−$\mathbf{R}_y$−CX−$\mathbf{R}_z$−CX	6	108	3×10^{-3}	12	324	7×10^{-5}
8	$\mathbf{R}_x - \mathbf{R}_y - \mathbf{R}_z$−CX	7	126	10^{-3}	12	324	4×10^{-6}

Table 2: Tested PQC architectures.

$$= 4\alpha^2 \operatorname{Var}(\hat{F}_0(\boldsymbol{\theta})) + 4\alpha^2\beta^4 \operatorname{Var}(\hat{F}(\boldsymbol{\theta}, \boldsymbol{\phi})). \tag{61}$$

It follows from (59) that

$$\operatorname{Var}(\hat{F}_0(\boldsymbol{\theta})) \leq \frac{1}{S} \sum_{c=1}^{2C-1} \|\mathbf{M}_0^c\|^2. \tag{62}$$

Regarding $\operatorname{Var}(\hat{F}(\boldsymbol{\theta}, \boldsymbol{\phi}))$, because $\mathbf{M}$ is a block-diagonal matrix with blocks $\mathbf{M}_m$, it follows that

$$\operatorname{Var}(\hat{F}(\boldsymbol{\theta}, \boldsymbol{\phi})) \leq \frac{1}{S} \sum_{c=1}^{2C-1} \max_m \|\mathbf{M}_m^c\|^2. \tag{63}$$

Substituting (62)–(63) into (61) provides

$$\operatorname{Var}(\hat{\nabla}_\alpha \mathcal{L}(\mathbf{z})) \leq \frac{4\alpha^2}{S} \sum_{c=1}^{2C-1} \|\mathbf{M}_0^c\|^2 + \frac{4\alpha^2\beta^4}{S} \sum_{c=1}^{2C-1} \max_m \|\mathbf{M}_m^c\|^2. \tag{64}$$

Let $\hat{\partial}_{\theta_p} F_0(\boldsymbol{\theta})$ and $\hat{\partial}_{\theta_p} F(\boldsymbol{\theta}, \boldsymbol{\phi})$ denote the sample-based estimates of $\partial_{\theta_p} F_0(\boldsymbol{\theta})$ and $\partial_{\theta_p} F(\boldsymbol{\theta}, \boldsymbol{\phi})$, respectively. The second expectation in (60) entails computing

$$\begin{aligned}
&\mathbb{E}[\|\hat{\nabla}_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{z}) - \nabla_{\boldsymbol{\theta}} \mathcal{L}(\mathbf{z})\|_2^2] = \sum_{p=1}^{P} \operatorname{Var}(\hat{\partial}_{\theta_p} \mathcal{L}(\mathbf{z})) \\
&= \sum_{p=1}^{P} \operatorname{Var}\left(\alpha^2 \hat{\partial}_{\theta_p} F_0(\boldsymbol{\theta}) + \alpha^2\beta^2 \hat{\partial}_{\theta_p} F(\boldsymbol{\theta}, \boldsymbol{\phi})\right) \\
&= \alpha^4 \sum_{p=1}^{P} \left(\operatorname{Var}(\hat{\partial}_{\theta_p} F_0(\boldsymbol{\theta})) + \beta^4 \operatorname{Var}(\hat{\partial}_{\theta_p} F(\boldsymbol{\theta}, \boldsymbol{\phi}))\right) \\
&= \frac{\alpha^4}{4} \sum_{p=1}^{P} \left(\operatorname{Var}(\hat{F}_0(\boldsymbol{\theta}_p^+)) + \operatorname{Var}(\hat{F}_0(\boldsymbol{\theta}_p^-))\right) \\
&+ \frac{\alpha^4\beta^4}{4} \sum_{p=1}^{P} \left(\operatorname{Var}(\hat{F}(\boldsymbol{\theta}_p^+, \boldsymbol{\phi})) + \operatorname{Var}(\hat{F}(\boldsymbol{\theta}_p^-, \boldsymbol{\phi})\right)
\end{aligned}$$

$$\leq \frac{P\alpha^4}{2S} \sum_{c=1}^{2C-1} \|\mathbf{M}_0^c\|^2 + \frac{P\alpha^4\beta^4}{2S} \sum_{c=1}^{2C-1} \max_m \|\mathbf{M}_m^c\|^2. \tag{65}$$

Analogously, the third and fourth expectations in (60) can be upper bounded as

$$\mathbb{E}[\|\hat{\nabla}_{\boldsymbol{\phi}} \mathcal{L}(\mathbf{z}) - \nabla_{\boldsymbol{\phi}} \mathcal{L}(\mathbf{z})\|_2^2] \leq \frac{Q\beta^4}{2S} \max_m b_m^2 + \frac{Q\alpha^4\beta^4}{2S} \sum_{c=1}^{2C-1} \max_m \|\mathbf{M}_m^c\|^2, \tag{66}$$

$$\operatorname{Var}(\hat{\nabla}_\beta \mathcal{L}(\mathbf{z})) \leq \frac{4\beta^2}{S} \max_m b_m^2 + \frac{4\alpha^4\beta^2}{S} \sum_{c=1}^{2C-1} \max_m \|\mathbf{M}_m^c\|^2. \tag{67}$$

Summing (64)–(67), the RHS of (36) is bounded as:

$$\mathbb{E}\left[\|\hat{\mathbf{g}}(\mathbf{z}) - \mathbf{g}(\mathbf{z})\|_2^2\right] \leq \frac{Q\bar{\beta}^4 + 8\bar{\beta}^2}{2S} \max_m b_m^2 + \frac{8\bar{\alpha}^2 + P\bar{\alpha}^4}{2S} \sum_{c=1}^{2C-1} \|\mathbf{M}_0^c\|^2 + \frac{8\bar{\alpha}^2\bar{\beta}^4 + 8\bar{\alpha}^4\bar{\beta}^2 + (P+Q)\bar{\alpha}^4\bar{\beta}^4}{2S} \sum_{c=1}^{2C-1} \max_m \|\mathbf{M}_m^c\|^2.$$

□

A.5 PQC architectures

The tested PQC architectures are reported in Table 2.