

Efficient 2D Tensor Network Simulation of Quantum Systems

Yuchen Pang*, Tianyi Hao*, Annika Dugad*, Yiqing Zhou* and Edgar Solomonik*

* Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Email: {yuchenp2, tianyi2, dugad2, yiqing2, solomon2}@illinois.edu

Abstract—Simulation of quantum systems is challenging due to the exponential size of the state space. Tensor networks provide a systematically improvable approximation for quantum states. 2D tensor networks such as Projected Entangled Pair States (PEPS) are well-suited for key classes of physical systems and quantum circuits. However, direct contraction of PEPS networks has exponential cost, while approximate algorithms require computations with large tensors. We propose new scalable algorithms and software abstractions for PEPS-based methods, accelerating the bottleneck operation of contraction and refactorization of a tensor subnetwork. We employ randomized SVD with an implicit matrix to reduce cost and memory footprint asymptotically. Further, we develop a distributed-memory PEPS library and study accuracy and efficiency of alternative algorithms for PEPS contraction and evolution on the Stampede2 supercomputer. We also simulate a popular near-term quantum algorithm, the Variational Quantum Eigensolver (VQE), and benchmark Imaginary Time Evolution (ITE), which compute ground states of Hamiltonians.

Index Terms—Numerical simulation, Parallel algorithms, Quantum mechanics, Quantum computing

I. INTRODUCTION

The degrees of freedom required to directly represent a quantum state grow exponentially with respect to the size of the system. However, for slightly entangled quantum states, polynomial-size representations are possible by using tensor networks [1], [2]. 1D tensor networks, also known as *matrix product states* (MPS), have proven successful in modeling 1D strongly correlated quantum systems [3], and various types of tensor networks have been invented to generalize this method to higher dimensions [4]–[7]. Among these variations, the 2D tensor lattice networks, known as *projected entangled pair states* (PEPS), enable effective representation of many physical systems of interest due to their 2D structure and are therefore widely used in condensed matter physics [4], [8], [9]. For the same reason, PEPS are also suitable to model near-term quantum computers with 2D architectures [10], [11]. In fact, this method has been recently applied to a large-scale exact simulation of 2D random quantum circuits [12].

Computations with PEPS involve high-order tensors and are computationally demanding. For example, exact contraction of PEPS to a scalar to calculate an expectation value is a #P-complete problem [13]. Approximations must be taken to avoid exponential cost, but the order of complexity is still high. Given a PEPS network of bond dimension r , i.e., composed of tensors that are of dimension $r \times r \times r \times r$, and maintaining approximate intermediates that have bond

dimensions no larger than some $m \propto r^2$, state-of-the-art PEPS contraction algorithms have a cost that scales as $O(m^3 r^2)$ via an iterative optimization procedure.

In search for an efficient parallel implementation for these PEPS algorithms, we propose an abstraction called `einsumsvd` for operations that contract a set of tensors into one and then refactorize it into two new tensors as described in Section II. This abstraction encapsulates the most costly operations in many tensor network applications, and it can be implemented with a variety of algorithms. In Section IV, we discuss an approach for performing `einsumsvd` that leverages randomized *singular value decomposition* (SVD) with implicit applications of the tensor operator. When applying this idea to the contraction of PEPS, we obtain a new algorithm that minimizes the cost relative to state-of-the-art approaches, while requiring only a single pass over the tensors. In the same section, we also present an intermediate caching strategy that further lowers the cost of calculating expectation values by trading space for time.

We implement these algorithms as part of a new Python library called Koala for PEPS-based simulations of quantum systems (Section V). This library supports various PEPS evolution and contraction algorithms, including our improvements, in the sequential/threaded setting using NumPy [14], the GPU setting using CuPy [15], and the distributed-memory setting using Cyclops [16], [17]. To improve the performance with Cyclops, we propose a parallel approach for tensor orthogonalization that avoids matricization of a high-order tensor, reducing synchronization between processors. This approach forms small Gram matrices in local memory so that matricization can be done locally while large-tensor contractions are kept distributed.

We evaluate the proposed algorithms and software in Section VI via a performance study of the PEPS evolution and contraction methods, using single-node and many-node execution on the Stampede2 supercomputer. Our experimental results demonstrate that the Gram matrix method improves the parallel performance of PEPS evolution up to 3.6X, and the randomized SVD approach improves the performance of PEPS contraction up to 22.4X in the sequential/threaded setting and up to 21.3X in the distributed-memory setting (combined with the Gram matrix method). The performance improvement due to the proposed caching optimization is also quantified in this section. For parallel scaling, we perform strong and weak scaling tests with up to 256 nodes with 64 cores per node

and show good weak scaling performance with respect to the tensor size (i.e. bond dimension).

As driver applications, we consider a common method for ground state calculation using tensor networks, *imaginary time evolution* (ITE), as well as a simulation of a prominent candidate algorithm for near-term quantum computers, the *variational quantum eigensolver* (VQE). These two applications, which are representative of many algorithms in the domain, are both implemented by the PEPS evolution and contraction primitives provided by our library. In Section VI-D, we study the numerical accuracy achieved by PEPS simulations with various bond dimensions for ITE and VQE. Our results demonstrate that higher accuracy is achieved by the use of larger tensors (i.e. larger bond dimension) with PEPS.

Overall, this paper makes the following contributions:

- a new approach for PEPS contraction that minimizes cost and does not require successive rounds of optimization,
- an automatic mechanism to cache intermediate contraction results during PEPS contraction,
- novel distributed-memory implementations of multiple PEPS contraction algorithms,
- a software library, Koala, providing sequential and parallel primitives for PEPS-based approximate simulations of quantum systems as well as encapsulating above algorithms via new software abstractions,
- to the best of our knowledge, the first study of accuracy and scalability of massively-parallel approximate PEPS algorithms with distributed-memory tensors.

A comparison between this work and related work is presented in Section VII.

II. BACKGROUND

In this section, we provide a brief review on tensor networks and their applications with a focus on the concepts used in this paper. Broader surveys on these topics are available [1]. We list notational conventions employed in the paper in Table I.

Notation	Meaning
\mathcal{A} (bold, calligraphic capital letters)	Tensor network
\mathbf{A} (bold capital letters)	Matrix or high-order tensor
\mathbf{a} (bold lowercase letters)	Vector
a_{ijk}	Element of the tensor \mathbf{A}
a^*, \mathbf{A}^*	Complex conjugate of a scalar a or conjugate transpose of an operator \mathbf{A}
$\mathbf{A} \in \mathbb{C}^{d_1 \times \dots \times d_n}$	Complex tensor \mathbf{A} of dimension $d_1 \times \dots \times d_n$
$\mathbf{A} : \mathbb{C}^{p_1 \times \dots \times p_s} \mapsto \mathbb{C}^{q_1 \times \dots \times q_t}$	Complex tensor \mathbf{A} treated as a linear operator
$ \psi\rangle$	Quantum state labeled by ψ
$\langle\psi $	Conjugate transpose of $ \psi\rangle$

Table I: Common notation used in this paper.

A. Quantum States as Tensors

The state $|\psi\rangle$ of a quantum computer with n qubits can be described by a unit vector in \mathbb{C}^{2^n} . By choosing 2^n orthonormal basis vectors/states to be denoted as $|\mathbf{i}\rangle$ with $\mathbf{i} = i_1 \dots i_n \in \{0, 1\}^n$, $|\psi\rangle$ can be written as

$$|\psi\rangle = \sum_{\mathbf{i} \in \{0,1\}^n} t_{\mathbf{i}}^{(\psi)} |\mathbf{i}\rangle.$$

Here, the amplitudes $t_{\mathbf{i}}^{(\psi)}$ are elements of an order n tensor $\mathbf{T}^{(\psi)} \in \mathbb{C}^{2 \times \dots \times 2}$. The *inner product* of two quantum states $|\psi\rangle$ and $|\phi\rangle$ is given by

$$\langle\phi|\psi\rangle = \sum_{\mathbf{i} \in \{0,1\}^n} (t_{\mathbf{i}}^{(\phi)})^* t_{\mathbf{i}}^{(\psi)},$$

and an *operator* \mathbf{H} acting on a quantum state is a linear transformation given by

$$|\phi\rangle = \mathbf{H} |\psi\rangle \Rightarrow t_{\mathbf{i}}^{(\phi)} = \sum_{\mathbf{j} \in \{0,1\}^n} h_{\mathbf{ij}} t_{\mathbf{j}}^{(\psi)}.$$

The *expectation value* of the Hermitian operator \mathbf{H} for a quantum state ψ is defined by $\langle\psi|\mathbf{H}|\psi\rangle$, following the definitions of operators and inner products.

A quantum gate is an operator that can act on a small subset of qubits. For example, a *single qubit gate* $G^{(k)}$ acting on the k th qubit gives

$$|\phi\rangle = G^{(k)} |\psi\rangle \Rightarrow t_{\mathbf{i}}^{(\phi)} = \sum_{j_k=0}^1 g_{i_k j_k}^{(k)} t_{i_1 \dots i_{k-1} j_k i_{k+1} \dots i_n}^{(\psi)}, \quad (1)$$

while a *2-qubit gate* $G^{(k,l)}$ acting on qubits k, l with $k < l$ gives

$$|\phi\rangle = G^{(k,l)} |\psi\rangle \Rightarrow t_{\mathbf{i}}^{(\phi)} = \sum_{j_k=0}^1 \sum_{j_l=0}^1 g_{i_k i_l j_k j_l}^{(k,l)} t_{i_1 \dots j_k \dots j_l \dots i_n}^{(\psi)}. \quad (2)$$

These operator applications are particular examples of *tensor contractions*, which allow for different types of products among tensors.

B. Tensor Networks for Quantum States

The contraction of a set of tensors defines a *tensor network*, which could serve to provide a tensor decomposition [18], [19] of a quantum state.

The *matrix product state* (MPS) [20], [21] is a 1D tensor network that describes a quantum state ψ with n sites by n tensors $\mathbf{M}^{(1)}, \dots, \mathbf{M}^{(n)}$, such that

$$t_{i_1 \dots i_n}^{(\psi)} = \sum_{k_1 \dots k_{n-1}} m_{i_1 k_1}^{(1)} m_{i_2 k_1 k_2}^{(2)} \dots m_{i_{n-1} k_{n-2} k_{n-1}}^{(n-1)} m_{i_n k_{n-1}}^{(n)}.$$

By convention, i_1, \dots, i_n are called *physical indices* and k_1, \dots, k_{n-1} are called *bond indices*, while their dimensions are called *physical dimensions* and *bond dimensions* respectively.

An operator H may also be described as a 1D tensor network, which is known as the *matrix product operator* (MPO) [1], [22], [23]:

$$h_{i_1 \dots i_n j_1 \dots j_n} = \sum_{k_1 \dots k_{n-1}} m_{i_1 j_1 k_1}^{(1)} m_{i_2 j_2 k_1 k_2}^{(2)} \dots m_{i_n j_n k_{n-1}}^{(n)}.$$

The *projected entangled pair state* (PEPS) [21] is a 2D tensor network that represents a quantum state with n^2 sites on a square lattice by n^2 tensors in a way analogous to MPS. These n^2 tensors are arranged on an $n \times n$ grid, so one can denote the tensor at the p 'th row and q 'th column as $M^{(p,q)}$. Then, each amplitude of a quantum state $|\psi\rangle$ is defined as

$$t_{i_1 \dots i_{n^2}}^{(\psi)} = \sum_{\mathbf{k}} \prod_{pq} m_{i_{pn+q} \mathbf{k}^{(p,q)}}^{(p,q)},$$

where \mathbf{k} denotes all the indices shared by tensors in the PEPS and $\mathbf{k}^{(p,q)}$ represents all the indices of the site tensor $M^{(p,q)}$ that are shared with other site tensors. Similarly to MPS, indices $i_1 \dots i_{n^2}$ are called physical indices and $\mathbf{k}^{(p,q)}$ are called bond indices.

A tensor network can be visualized by a graph (*tensor diagram*), in which each vertex denotes a tensor and each edge denotes an index. Tensor diagrams describing MPS, MPO, and PEPS are displayed in Figure 1.




Figure 1: Tensor diagrams for (a) an MPS, (b) an MPO and (c) a 3×3 PEPS.

C. Basic Tensor Network Computations

We now describe common computational primitives to transform tensor networks, focusing on PEPS.

1) *One-Site and Two-Site Operators*: To apply a one-site operator $G^{(1)}$ on the site $M^{(p,q)}$ of a PEPS, it suffices to compute the updated site $\tilde{M}^{(p,q)}$ by performing the contraction,

$$\tilde{m}_{ik_1 k_2 k_3 k_4}^{(p,q)} = \sum_j g_{ij}^{(1)} m_{jk_1 k_2 k_3 k_4}^{(p,q)}. \quad (3)$$

To apply a two-site operator $G^{(2)}$ on two neighboring sites $M^{(p,q)}$ and $M^{(p,q+1)}$, we consider a variant of the simple update algorithm [24]. Two site tensors are contracted with the operator and then decomposed into two new site tensors $\tilde{M}^{(p,q)}$ and $\tilde{M}^{(p,q+1)}$,

$$\begin{aligned} & \sum_{k'_4} \tilde{m}_{i_1 k_1 k_2 k_3 k'_4}^{(p,q)} \tilde{m}_{i_2 k'_4 k_5 k_6 k_7}^{(p,q+1)} \\ & \approx \sum_{j_1 j_2 k_4} g_{i_1 i_2 j_1 j_2}^{(2)} m_{j_1 k_1 k_2 k_3 k_4}^{(p,q)} m_{j_2 k_4 k_5 k_6 k_7}^{(p,q+1)}. \end{aligned} \quad (4)$$

This contraction and refactorization procedure is described by the first and last tensor diagrams in Figure 4. Applying a two-site operator on non-neighboring sites is possible by applying a chain of two-site operators (i.e. SWAP gates) on neighboring sites.

The contraction and refactorization necessary to apply two-site operators may be executed by contracting the tensors (utilizing `einsum` in Python's NumPy), followed by reshaping the result into a matrix, computing a low-rank matrix factorization such as the truncated SVD, and reshaping the resulting factors into the appropriate tensors. We refer to the combined refactorization operation as `einsumsvd`. This primitive takes as input a tensor network and produces a two-site tensor network with a single virtual leg connecting the two sites. With the notion of `einsumsvd`, we can rewrite Equation (4) as

$$\tilde{M}^{(p,q)}, \tilde{M}^{(p,q+1)} \leftarrow \text{einsumsvd}(G, M^{(p,q)}, M^{(p,q+1)}).$$

Figure 2 depicts other examples of `einsumsvd`.




Figure 2: `einsumsvd` examples: (a) a 5-site tensor network is contracted and then refactorized into two sites. (b) PEPS contraction step in which two tensors that share a neighbor are merged and the combined dimension is truncated.

2) *Amplitudes and Expectation Values*: A given amplitude of a state $|\psi\rangle$, such as $t_i^{(\psi)}$, may be obtained by computing $\langle i | \psi \rangle$. When represented by PEPS, the bond dimension of $|i\rangle$ is 1 and each component of $|i\rangle$ can be applied to the respective site of $|\psi\rangle$, resulting in a contraction of a PEPS without physical dimensions, i.e., a *one-layer contraction*.

Consider a Hermitian operator H composed of local terms $H = \sum_i H_i$, where H_i are operators acting on local sites of the PEPS. One way to calculate its expectation value is to calculate it as a sum of the expectation values of H_i ,

$$\langle \psi | H | \psi \rangle = \sum_{i=1}^N \langle \psi | H_i | \psi \rangle = \sum_{i=1}^N \langle \psi | \phi_i \rangle. \quad (5)$$

$|\phi_i\rangle$ can be obtained by the algorithms of applying local operators on $|\psi\rangle$. The calculation of $\langle \psi | \phi_i \rangle$ involves a contraction of two PEPS, i.e. a *two-layer contraction*, as depicted in Figure 3.

Both one-layer and two-layer contractions require exponential cost in the number of sites if done in an exact way. A variety of approximate schemes with polynomial cost exists, which we survey in Section III.

D. Tensor Network Applications

Most tensor network computations are based on the operator application and contraction primitives introduced in the




Figure 3: Tensor diagram depicting the inner product of two PEPS networks.

previous subsection. We provide a technical introduction to two representative methods below.

1) *Imaginary Time Evolution*: Consider a Hamiltonian $\mathbf{H} = \sum_{j=1}^N \mathbf{H}_j$, where \mathbf{H}_j are local operators. *Imaginary time evolution* (ITE) evolves a quantum state toward its ground state by repetitively applying the ITE operator, $e^{-\tau \mathbf{H}}$, in effect computing power iteration on the matrix $e^{-\tau \mathbf{H}}$, which generally converges to the eigenvector with the smallest eigenvalue of \mathbf{H} . ITE performs this iteration via the *time-evolution block decimation* (TEBD) algorithm, which enables the approximate evolution of a tensor network state by decomposing $e^{-\tau \mathbf{H}}$ into a sequence of local operators using the Trotter-Suzuki decomposition: $e^{-\tau \mathbf{H}} = \prod_{j=1}^N e^{-\tau \mathbf{H}_j} + O(\tau^2)$. Since each \mathbf{H}_j is a one-site or two-site operator, so is $e^{-\tau \mathbf{H}_j}$. Consequently, a Hamiltonian with N local terms requires application of N local operators for each ITE step. The Rayleigh quotient $\frac{\langle \psi | \mathbf{H} | \psi \rangle}{\langle \psi | \psi \rangle}$ of the final state is computed as an approximation to the ground state energy.

2) *Quantum Circuit Simulation*: Following the prescription for gate application described in Section II-A, any quantum circuit can be expressed in terms of one-site and two-site operators acting on an initial state (e.g. a basis state). Simulation of quantum circuits is useful for the development and verification of quantum algorithms and software. Approximate simulation additionally sheds insight on the propagation of error within quantum algorithms, which is useful for understanding the effect of noise on computations with near-term quantum devices [25]. Among the quantum algorithms for near-term quantum devices, the *variational quantum eigensolver* (VQE) [26] is one of the most promising.

VQE is a quantum-classical hybrid algorithm that computes the ground state of a Hamiltonian \mathbf{H} by optimizing a quantum state parameterized by a quantum circuit. The quantum part of the algorithm evaluates the objective $\langle \psi(\theta) | \mathbf{H} | \psi(\theta) \rangle$ for given parameters θ , and the classical part of the algorithm takes the measurement results from the quantum device and runs classical optimization algorithms to tune the circuit parameters. A careful design of the circuit structure, depending on the Hamiltonian of interest, is critical in improving the optimization efficiency and achieving desirable accuracy. A number of efforts have been made to design more efficient circuits for different Hamiltonians [27] and to apply the algorithm to quantum chemistry simulations [28]–[30].

III. PREVIOUS WORK ON PEPS ALGORITHMS

Our contributions build directly on existing algorithms for evolution and contraction of PEPS.

A. PEPS Evolution Algorithms

Direct evaluation of Equation 4 reveals a time complexity of $O(d^3 r^9)$ and space complexity of $O(d^2 r^6)$, where d is the size of the physical dimension and r is the size of the bond dimension. Assuming $d \ll r$, the time and space complexity can be reduced to $O(d^2 r^5)$ and $O(dr^4)$ respectively by a QR factorization of the two site tensors prior to the application of the operator [25], as shown in Algorithm 1 and Figure 4.

Algorithm 1 Operator application by QR-SVD

Input: two site tensors $\mathbf{M}^{(p,q)}$ and $\mathbf{M}^{(p,q+1)}$,
a two-site operator \mathbf{G}

Output: updated site tensors $\tilde{\mathbf{M}}^{(p,q)}$ and $\tilde{\mathbf{M}}^{(p,q+1)}$

- 1: $\mathbf{Q}^{(p,q)}, \mathbf{R}^{(p,q)} \leftarrow \text{QR}(\mathbf{M}^{(p,q)})$ {step (1) \rightarrow (2)}
 - 2: $\mathbf{Q}^{(p,q+1)}, \mathbf{R}^{(p,q+1)} \leftarrow \text{QR}(\mathbf{M}^{(p,q+1)})$ {step (1) \rightarrow (2)}
 - 3: $\tilde{\mathbf{R}}^{(p,q)}, \tilde{\mathbf{R}}^{(p,q+1)} \leftarrow \text{einsumsvd}(\mathbf{G}, \mathbf{R}^{(p,q)}, \mathbf{R}^{(p,q+1)})$
{step (2) \rightarrow (4)}
 - 4: $\tilde{\mathbf{M}}^{(p,q)} \leftarrow \mathbf{Q}^{(p,q)} \tilde{\mathbf{R}}^{(p,q)}$ {step (4) \rightarrow (5)}
 - 5: $\tilde{\mathbf{M}}^{(p,q+1)} \leftarrow \mathbf{Q}^{(p,q+1)} \tilde{\mathbf{R}}^{(p,q+1)}$ {step (4) \rightarrow (5)}
 - 6: **return** $\tilde{\mathbf{M}}^{(p,q)}, \tilde{\mathbf{M}}^{(p,q+1)}$
-




Figure 4: Tensor diagrams describing QR-SVD (Algorithm 1).

B. PEPS Contraction Algorithms

Many approximate PEPS contraction algorithms have been proposed in pursuit of reducing memory footprint while maintaining high accuracy. These include the original Boundary MPS (BMPS) [4] algorithm and its variants [31]–[33], Tensor Renormalization Group (TRG) [34] and many TRG-inspired renormalization group methods [35]–[42], Corner Transfer Matrix (CTM) [43], [44], as well as contraction algorithms for new 2D tensor network structures based on PEPS [45], [46]. We focus on a variant of BMPS for one-layer and two-layer contractions as introduced below.

1) *One-Layer BMPS*: Let \mathcal{P} be an n -by- n PEPS of bond dimension r without physical indices and $\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(n)}$ be the rows of \mathcal{P} . Treating $\mathcal{P}^{(1)}$ as an MPS and $\mathcal{P}^{(2)}, \dots, \mathcal{P}^{(n)}$ as MPOs, contracting \mathcal{P} by BMPS can be considered as applying MPOs $\mathcal{P}^{(2)}, \dots, \mathcal{P}^{(n)}$ on the MPS $\mathcal{P}^{(1)}$ approximately.

The approximation is necessary for avoiding exponential cost and it is done by truncating the bond dimension of result MPS to a predefined m . This approach is described in Algorithm 2.

Algorithm 2 Boundary MPS

Input: $n \times n$ PEPS \mathcal{P} consisting of rows $\{\mathcal{P}^{(1)}, \dots, \mathcal{P}^{(n)}\}$,
algorithm ApproxApply, truncation bond dimension m
Output: scalar v

- 1: $\mathcal{S} \leftarrow \mathcal{P}^{(1)}$
- 2: **for** $i = 2, 3, \dots, n$ **do**
- 3: $\mathcal{S} \leftarrow \text{ApproxApply}(\mathcal{S}, \mathcal{P}^{(i)}, m)$
- 4: **end for**
- $\{\mathcal{S} \text{ is an MPS without physical indices at this point}\}$
- 5: $v \leftarrow \text{contracting } \mathcal{S} \text{ to a scalar}$
- 6: **return** v

The approximate application of an MPO on an MPS can be performed in various ways. The original BMPS algorithm [4] involves solving a linear equation $\mathcal{E}T = \mathcal{P}$ for one site T while fixing the other sites (the environment \mathcal{E}) at a time with respect to the target MPS \mathcal{P} . This method, while having a low cost of $O(dm^2r^3) + O(m^3r^2)$ if the MPO has a two-layer structure as discussed below [33], is typically done with multiple rounds of alternating optimization of the sites of the MPS. Our approach is most closely related to the zip-up algorithm proposed in [47], where a series of `einsumsvd`s are conducted to truncate the bond dimension of \mathcal{P} , as shown in Figure 5 and described in Algorithm 3. We also note that this method can incorporate canonicalization (orthogonalization of the tensor network) to minimize amplification of SVD truncation error. Specifically, canonicalization can be achieved by absorbing singular values into $V^{(i-1)}$ or $V^{(i)}$ for different rows of PEPS [47] in an alternating manner.




Figure 5: Apply MPO on MPS approximately (Algorithm 3).

2) *Two-Layer BMPS*: Given two n -by- n PEPS \mathcal{A} and \mathcal{B} of bond dimension r_1 and r_2 respectively, their inner product can be computed in two ways. The naive way is to contract over all the physical indices between \mathcal{A} and \mathcal{B} to create a new n -by- n PEPS \mathcal{P} with bond dimensions $r = r_1r_2$ and no physical dimensions, and then perform one-layer BMPS. However, forming the combined \mathcal{P} requires $O(r_1^4r_2^4)$ memory, while storing \mathcal{A} and \mathcal{B} separately only needs $O(r_1^4 + r_2^4)$. An improvement is to maintain the two-layer structure and only contract the corresponding sites from two layers to one layer when needed, which is numerically equivalent to the naive way but saves memory. This approach is often used in practice

Algorithm 3 Apply MPO on MPS approximately

Input: MPS \mathcal{S} consisting of tensors $\{\mathcal{S}^{(1)}, \dots, \mathcal{S}^{(n)}\}$,
MPO \mathcal{O} consisting of tensors $\{\mathcal{O}^{(1)}, \dots, \mathcal{O}^{(n)}\}$,
algorithm `einsumsvd`, truncation bond dimension m
Output: MPS \mathcal{S}'

- 1: $V^{(1)} \leftarrow \text{Contract } \mathcal{S}^{(1)} \text{ and } \mathcal{O}^{(1)} \{v_{jkl}^{(1)} \leftarrow \sum_i s_{ij}^{(1)} o_{ikl}^{(1)}\}$
- 2: **for** $i = 2, 3, \dots, n$ **do**
- 3: $V^{(i-1)}, V^{(i)} \leftarrow \text{einsumsvd}(V^{(i-1)}, \mathcal{S}^{(i)}, \mathcal{O}^{(i)}, m)$
- 4: **end for**
- 5: $\mathcal{S}' \leftarrow \{V^{(1)}, \dots, V^{(n)}\}$
- 6: **return** \mathcal{S}'

and referred to interchangeably with the naive approach in the literature. In Section IV, we show that when `einsumsvd` is equipped with implicit randomized SVD, the implicit structure of the two-layer BMPS approach yields reduced computational complexity.

IV. ALGORITHMIC IMPROVEMENTS TO PEPS SIMULATION

In this section, we describe our algorithmic improvements to basic PEPS computations, specifically `einsumsvd` and the expectation value calculation introduced in Section II-C.

A. Implicit Tensor Network Refactorization

Randomized SVD is an algorithm that approximates truncated SVD of rank r of an m -by- n matrix with time complexity $O(mnr)$ [48] by the use of orthogonal iteration. This algorithm has been directly applied to replace the full SVD algorithm to apply operators on MPS [49], which reduces the cost by one order of the physical dimension. Here we apply this idea more generally to `einsumsvd`, resulting in asymptotic reductions in cost for one-layer and two-layer BMPS contraction.

Algorithm 4 Randomized SVD

Input: an operator $A : \mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^{m_1 \times \dots \times m_s}$, rank r ,
number of iteration k
Output: an approximated truncated SVD of the input operator
tensor operator $U \in \mathbb{C}^{m_1 \times \dots \times m_s \times r} : \mathbb{C}^r \mapsto \mathbb{C}^{m_1 \times \dots \times m_s}$,
diagonal matrix $\Sigma \in \mathbb{C}^{r \times r}$,
tensor operator $V \in \mathbb{C}^{r \times n_1 \times \dots \times n_t} : \mathbb{C}^r \mapsto \mathbb{C}^{n_1 \times \dots \times n_t}$

- 1: $Q \leftarrow \text{draw a random tensor from } [-1, 1]^{n_1 \times \dots \times n_t \times r}$
 $\{Q : \mathbb{C}^r \mapsto \mathbb{C}^{n_1 \times \dots \times n_t}\}$
- 2: $P \leftarrow \text{orthogonalize}(AQ) \{P : \mathbb{C}^r \mapsto \mathbb{C}^{m_1 \times \dots \times m_s}\}$
- 3: **for** $i = 1, 2, \dots, k$ **do**
- 4: $Q \leftarrow \text{orthogonalize}(A^*P) \{Q : \mathbb{C}^r \mapsto \mathbb{C}^{n_1 \times \dots \times n_t}\}$
- 5: $P \leftarrow \text{orthogonalize}(AQ) \{P : \mathbb{C}^r \mapsto \mathbb{C}^{m_1 \times \dots \times m_s}\}$
- 6: **end for**
- 7: $\tilde{U}, \Sigma, V \leftarrow \text{SVD}(P^*A) \{P^*A : \mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^r\}$
- 8: $U \leftarrow P\tilde{U}$
- 9: **return** U, Σ, V

As shown in Algorithm 4, randomized SVD does not require an explicit form of the operator A but only needs to know

	BMPS	IBMPS	Two-layer IBMPS
Time complexity	$O(n^2 m^3 r^4)$	$O(n^2 m^2 r^4 + n^2 m^3 r^2)$	$O(n^2 d m^2 r^3 + n^2 d m^3 r^2)$
Space complexity	$O(\max\{m^2 r^3, r^4\})$	$O(\max\{m^2 r^2, r^4\})$	$O(\max\{m^2 r^2, r^4\})$

Table II: The asymptotic time and space complexity of computing $\langle \mathcal{P} | \mathcal{P} \rangle$ using BMPS, IBMPS and two-layer IBMPS for an n -by- n PEPS \mathcal{P} of bond dimension \sqrt{r} , physical dimension d , and truncation bond dimension m .

how to apply \mathbf{A} and \mathbf{A}^* . Thus, the operator \mathbf{A} could be given implicitly as an uncontracted tensor network. Depending on the structure of the tensor network \mathbf{A} , it is possible to perform $\mathbf{A}\mathbf{Q}$ and $\mathbf{A}^*\mathbf{P}$ more efficiently and with less memory than by explicitly forming \mathbf{A} and applying it.

A similar implicit method has been specifically applied to the TRG algorithm [50]. Here, we apply this method to the `einsumsvd` step of the BMPS contraction algorithm (Algorithm 3). With the best choice of contraction order, the use of implicit methods in BMPS yields an asymptotic reduction in cost and, in some cases, memory footprint, when compared with a naive implementation of `einsumsvd`. We refer to this approach as *implicit* BMPS (IBMPS) contraction. Two-layer BMPS can achieve greater benefit with the use of implicit randomized SVD, as it naturally maintains tensors in implicit forms. We compare the asymptotic costs for BMPS, IBMPS, and two-layer IBMPS in Table II.

B. Intermediate Caching for PEPS Expectation Values

To compute expectation values of PEPS efficiently, we develop a caching strategy for consecutive executions of BMPS (Algorithm 2) when evaluating Equation (5). We observe that these contractions share common intermediates due to the locality of \mathbf{H}_i , and our approach reuses partial contractions of rows (or columns) of PEPS to compute $\langle \psi | \mathbf{H}_i | \psi \rangle$.

For example, consider a 4×4 PEPS and two local operators \mathbf{H}_1 and \mathbf{H}_2 both acting on the third row of the PEPS, as shown in Figure 6. We cache the two boundary MPS that represent the partial contraction results of the first two rows and last row to accelerate the calculation of $\langle \psi | \mathbf{H}_1 | \psi \rangle$ and $\langle \psi | \mathbf{H}_2 | \psi \rangle$.




Figure 6: Shared intermediates for expectation values of two local operators on a 4×4 PEPS.

When applied to an $n \times n$ PEPS, this caching strategy requires two full two-layer PEPS contractions to generate all the cached intermediates, and then each local-operator expectation value is calculated with a $3 \times n$ PEPS contraction. This strategy can be easily extended to compute a collection of expectation values, e.g., to compute gradients of parameters

of local operators (as used in the Adapt-VQE method [29]). In Section VI, we demonstrate the speed-up achieved by this strategy (see Figure 9).

Besides evaluating Equation (5) with intermediate caching, $\langle \psi | \mathbf{H} | \psi \rangle$ can be calculated in an alternative approach. Using the Trotter-Suzuki decomposition, $e^{\tau \mathbf{H}} = \prod_{j=1}^N e^{\tau \mathbf{H}_j} + O(\tau^2)$, as in ITE, and the Taylor expansion, $e^{\tau \mathbf{H}} = \mathbf{I} + \tau \mathbf{H} + O(\tau^2)$, one can approximate

$$\langle \psi | \mathbf{H} | \psi \rangle = \frac{1}{\tau} \left(\langle \psi | \prod_{j=1}^N e^{\tau \mathbf{H}_j} | \psi \rangle - \langle \psi | \psi \rangle \right) + O(\tau). \quad (6)$$

In comparison with the previous method, Equation (6) requires one two-layer PEPS contraction as opposed to two, but application of an additional ITE step would grow the bond dimension of $|\phi\rangle = \prod_{j=1}^N e^{\tau \mathbf{H}_j} |\psi\rangle$ or require approximation.

V. KOALA: HIGH-PERFORMANCE PEPS SIMULATION

We implement the algorithms that are discussed above within an open-source Python library called “Koala”, which is available at <https://github.com/cyclops-community/koala>.

A. Interface and Features

Koala provides explicit primitives for constructing PEPS networks, applying operators, and computing expectation values. Tensor library backends, as well as different algorithms for applying operators and contracting PEPS, can be specified as arguments to these primitive routines. An example code for constructing a PEPS, applying operators, and computing an expectation value is below.

```
from koala import peps, Observable
from koala.peps import QRUpdate, BMPS
from tensorbackends.interface import \
    ImplicitRandomizedSVD

# Create a 2-by-3 PEPS in distributed memory
qstate = peps.computational_zeros(
    nrow=2, ncol=3, backend='ctf'
)

# Apply one-site and two-site operators with QR-SVD
Y = qstate.backend.astensor([...]).reshape(2, 2)
CX = qstate.backend.astensor([...]).reshape(2, 2, 2, 2)
qstate.apply_operator(Y, [1])
qstate.apply_operator(CX, [1, 4], QRUpdate(rank=2))

# Calculate the expectation value with IBMPS
H = Observable.ZZ(3, 4) + 0.2 * Observable.X(1)
result = qstate.expectation(
    H, use_cache=True,
    contract_option=BMPS(ImplicitRandomizedSVD(rank=4))
)
```

B. Distributed Tensor Parallelization

Koala supports both sequential and parallel execution. This portability is achieved by an abstraction layer of tensor data types and operations that enables different tensor libraries to be used with little change in the source code. Currently, the tensor libraries we support are NumPy, CuPy, and Cyclops. Among them, NumPy provides sequential and threaded routines for tensor computations [14], CuPy implements such routines that run on GPUs [15], and Cyclops allows similar operations in distributed memory [16]. Cyclops distributes each tensor over all processors and provides a Python interface for `einsum` operations, which can specify the contraction of any tensor network into a single tensor. Cyclops also provides a front-end to routines in the ScaLAPACK library [51] for distributed matrix factorization.

C. Avoiding Tensor Reshaping

The standard approach to `einsumsvd` involves performing tensor contractions, matricizing (unfolding) the result, computing a low-rank matrix factorization, and folding the factors into tensors of appropriate shape. Folding and unfolding of this type are logical transformations that need not modify the data and often have negligible cost in sequential execution. However, these two operations can be nontrivial if each tensor is distributed on a processor grid (via any standard distribution, including cyclic and blocked layouts). Since both the matrix and the tensor unfoldings must be mapped to different processor grids for corresponding computations, folding/unfolding requires expensive redistribution of data. Therefore, while calls to NumPy's `reshape` are practically free, Cyclops' `reshape` can become a bottleneck.

We address this bottleneck within implicit `einsumsvd` (Algorithm 4) and QR-SVD (Algorithm 1), by performing orthogonalization via eigendecomposition of a Gram matrix formed by tensor contraction. For the QR factorizations done in QR-SVD, forming the Gram matrix requires a reshape of a small tensor. For orthogonalization in `einsumsvd`, no reshape of the Gram matrix is necessary. In both cases, this matrix is small, so the latency overhead of using distributed memory is alleviated by performing computations on the Gram matrix sequentially. Algorithm 5 describes this general approach. When Algorithm 5 is applied to QR-SVD, we can also perform `einsumsvd` for operator application sequentially by leveraging the fact that the operator G and the R factors are all small enough to fit into local memory.

VI. RESULTS

We benchmark our implementation of the PEPS algorithm variants that are discussed above. The overall accuracy of PEPS-based simulation for two example applications is also studied to demonstrate how accuracy is affected by the bond dimension. Parallel experiments are all done on the Stampede2 supercomputer using 64 cores per node for Cyclops and 32 MKL [52] threads on one node for NumPy. The number of *processes per node* (PPN) for Cyclops is always chosen as 64 unless otherwise specified.

Algorithm 5 Reshape-avoiding orthogonalization

Input: a distributed tensor $A \in \mathbb{C}^{m_1 \times \dots \times m_s \times n_1 \times \dots \times n_t}$ such that $\prod_{i=1}^s m_i \gg \prod_{i=1}^t n_i$
 (A represents an operator $\mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^{m_1 \times \dots \times m_s}$)

Output:

- distributed isometric tensor operator $Q \in \mathbb{C}^{m_1 \times \dots \times m_s \times n_1 \times \dots \times n_t}$;
 $\mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^{m_1 \times \dots \times m_s}$
 distributed tensor operator $R \in \mathbb{C}^{n_1 \times \dots \times n_t \times n_1 \times \dots \times n_t}$;
 $\mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^{n_1 \times \dots \times n_t}$
- 1: In parallel, Compute $G \leftarrow A^* A$
 $\{G : \mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^{n_1 \times \dots \times n_t}\}$
 - 2: Gather G data among processors
 - 3: Reshape G into a matrix $G \in \mathbb{C}^{n_1 \dots n_t \times n_1 \dots n_t}$
 - 4: Compute eigendecomposition X, Λ of G such that
 $G = X \Lambda X^*$
 - 5: Compute $R \leftarrow \sqrt{\Lambda} X^* \in \mathbb{C}^{n_1 \dots n_t \times n_1 \dots n_t}$
 - 6: Compute $P \leftarrow R^{-1} \in \mathbb{C}^{n_1 \dots n_t \times n_1 \dots n_t}$
 - 7: Reshape R into a tensor $R \in \mathbb{C}^{n_1 \times \dots \times n_t \times n_1 \times \dots \times n_t}$
 $\{R : \mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^{n_1 \times \dots \times n_t}\}$
 - 8: Reshape P into a tensor $P \in \mathbb{C}^{n_1 \times \dots \times n_t \times n_1 \times \dots \times n_t}$
 $\{P : \mathbb{C}^{n_1 \times \dots \times n_t} \mapsto \mathbb{C}^{n_1 \times \dots \times n_t}\}$
 - 9: Distribute R and P among different processors
 - 10: In parallel, compute $Q \leftarrow AP$
 - 11: **return** Q, R
-


A. PEPS Evolution Benchmark

Figure 7 studies the performance of applying one layer of TEBD operators on all neighboring sites on PEPS of two different sizes: 8×8 and 15×15 . For the smaller system size, which fits into memory on 1 node, we compare the performance of NumPy and Cyclops, as shown in Figure 7a. We observe that while NumPy outperforms Cyclops for small bond dimensions, their performance becomes similar as bond dimension grows. We also compare three different algorithms for a 15×15 PEPS on 16 nodes using Cyclops in Figure 7b, which shows that reshape-avoiding orthogonalization (Algorithm 5) can accelerate the parallel execution of QR-SVD (Algorithm 1) by factors of up to 3.6X.

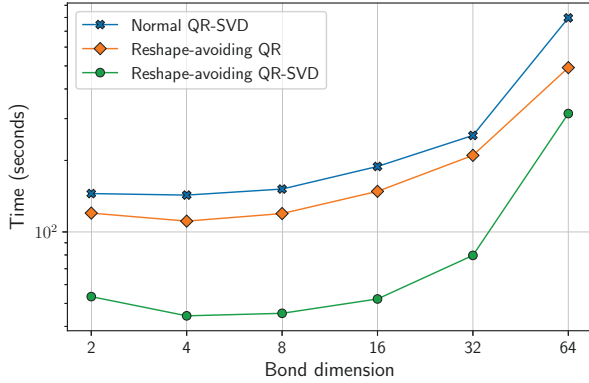
B. PEPS Contraction Benchmark

Figure 8a shows the performance of contracting an 8×8 PEPS on a single KNL node of Stampede2 using BMPS, IBMPS, two-layer IBMPS, and an exact algorithm [12] with NumPy and Cyclops as backends. Since contracting the inner product between two identical PEPS limits the selection of bond dimensions, we directly generate a PEPS without physical indices to obtain more data points. However, two-layer IBMPS is only applicable for the inner product, so we include fewer data points.

The contraction bond dimension is set equal to the initial bond dimension, which varies from 2 to 64. The trade-off between backends is similar to the PEPS evolution benchmark in that Cyclops is more scalable despite NumPy being faster



(a) 8×8 PEPS using 1 node.



(b) 15×15 PEPS using 16 nodes.


Figure 7: Running time versus bond dimension for applying one-layer TEBD on PEPS using variants of QR-SVD.

“Reshape-avoiding QR” denotes the variant where orthogonalization is done locally; “reshape-avoiding QR-SVD” denotes the variant where both orthogonalization and `einsumsvd` are done locally.

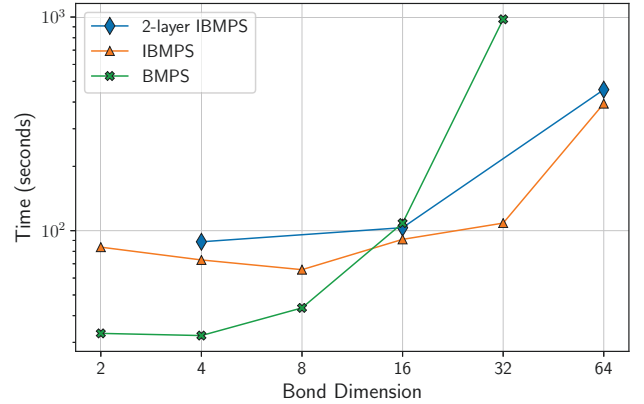
for small bond dimensions. Further, we show that IBMPS and two-layer IBMPS not only have an advantage on asymptotic computational complexity, but are also more memory-efficient compared to BMPS. In practice, we observe that only IBMPS and two-layer IBMPS are able to contract an 8×8 PEPS with initial and contraction bond dimension both equal to 64 on a single node. Moreover, we test the highest achievable bond dimension for contracting a 6×6 PEPS using various algorithms on a single node. The exact algorithm and BMPS can only contract such a PEPS with bond dimension less than 30 and 40 respectively, while IBMPS can achieve a bond dimension of 95 and two-layer IBMPS can perform contraction with a bond dimension of more than 100.

We show the performance of contracting a 15×15 PEPS on 16 nodes using Cyclops in Figure 8b. As in the 8×8 case, we observe that IBMPS has an asymptotic advantage over BMPS, agreeing with our analysis in Table II.

Figure 9 studies the performance of intermediate caching for the expectation value calculation introduced in Section IV-B. The expectation operator is composed of one-site operators



(a) 8×8 PEPS using 1 node.



(b) 15×15 PEPS using 16 nodes.

Figure 8: Running time for fully contracting a PEPS as the bond dimension grows.

acting on all sites and two-site operators acting on all pairs of neighboring sites. With more PEPS sites, the speed-up that caching brings becomes greater. For 12×12 PEPS with bond dimension 4, the expectation value calculation is 4.7X faster.

We benchmark the accuracy of our contraction algorithms with *random quantum circuits* (RQC) [53], [54], which are designed to be difficult to simulate classically and have been proven to satisfy both average-case hardness and anti-concentration property [25], [55]–[57]. It is especially challenging to simulate RQCs using tensor networks, since the states created by RQCs are strongly entangled and very sensitive to approximations. By applying our algorithms to this problem, we show that approximate contraction is effective even for one of the hardest simulations.

Following the construction procedure proposed in [54], we apply iSWAP gates for all pairs of neighboring sites every four layers, increasing the bond dimension by a factor of 4. For the RQC benchmarking, we use 4×4 , 5×5 , 6×6 , and 7×7 PEPS generated by 8 layers of RQC with exact evolution, which has an initial bond dimension of 16. Then, BMPS and IBMPS with varying contraction bond dimension are used to compute one amplitude of the circuit. The result is

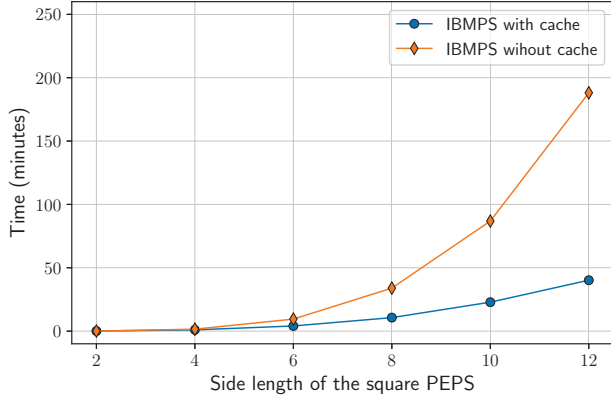


Figure 9: Running time of the expectation value calculation with and without caching for a square PEPS with bond dimension 4 on 1 node of Stampede2.




Figure 10: Relative error of contracting an RQC-generated $n \times n$ PEPS of bond dimension 16 using BMPS and IBMPS with varying contraction bond dimension.

compared to the exact contraction algorithm to determine the relative error, as shown in Figure 10. We show that the use of implicit randomized SVD in IBMPS does not incur additional error as compared to the naive SVD. We also observe that after increasing the contraction bond dimension above a certain threshold, the relative error quickly drops to near machine epsilon. This threshold is positively correlated to the PEPS size and the initial bond dimension. For 6×6 and 7×7 PEPS, we achieved one digit of accuracy with a contraction bond dimension of 256 on only 64 nodes with 7808 gigabytes of memory in total.

C. Parallel Scaling

We study the parallel scaling of applying one layer of TEBD operators (PEPS evolution) and IBMPS (PEPS contraction) in Figure 11 and Figure 12. In the strong scaling analysis shown in Figure 11, we test a smaller problem that occupies most of the memory on a single node and a larger problem that occupies most of the memory on 16 nodes for both




Figure 11: Strong scaling for PEPS evolution (applying one layer of TEBD operators) and PEPS contraction (IBMPS for PEPS with no physical indices); the dashed line separates single-node and multiple-node execution.

PEPS evolution and contraction. For better performance with Cyclops, we choose to use PPN=16 on 256 nodes (2^{14} cores). When executing on a single node with multiple cores (less than 2^6 cores), the running time roughly halves when number of cores doubles, and good strong scaling efficiency is achieved up until all cores of the node are utilized.

When the benchmarks are executed on the smaller PEPS with multiple nodes, we obtain a speed-up of 4.5X on 32 nodes (2^{10} cores) for PEPS evolution and a speed-up of 13.5X on 64 nodes for PEPS contraction with respect to a single node (2^6 cores), after which the performance deteriorates. We observe that the fraction time in local matrix multiplication (GEMM), in the smaller PEPS contraction experiments, is 62% on 1 node, and 16% on 16 nodes. For the larger problem starting from 16 nodes, we observe that initially the performance scales well, but for PEPS evolution, the performance deteriorates after 128 nodes, and for PEPS contraction, the performance roughly decreases by one-half on 256 nodes. PEPS evolution has a cost that is only slightly superlinear in the size of the tensor sites, so this kernel tends to be communication-bound.

In the weak scaling analysis shown in Figure 12, we focus on increasing the bond dimension of PEPS while keeping the memory usage per node constant. This case better reflects the practical usage of large PEPS simulations, where the major bottleneck is memory and not execution time. As shown in the figure, we observe sustained weak scaling for the full PEPS evolution and contraction benchmark up to 64 nodes (2^{12} cores). Similarly, the time spent in GEMM is always around 60-70% for PEPS contraction. For the experiments with 128 nodes (2^{13} cores) and 256 nodes (2^{14} cores), we instead run microbenchmarks, where only the main computational primitives are benchmarked. Overall, our approach shows good weak scalability for dealing with large bond dimensions that are prohibitive in cost or infeasible in memory footprint for typical single-node architectures.




Figure 12: Weak scaling for PEPS evolution (applying one layer of TEBD operators) and PEPS contraction (IBMPS for PEPS with no physical indices) with evolution bond dimension r and contraction bond dimension m :
 $r = 70, 83, 98, 117, 140, 166, 197, 235, 280$,
 $m = 80, 95, 113, 134, 160, 190, 226, 269, 320$.

D. Applications

We apply the PEPS algorithms to imaginary time evolution and variational quantum eigensolver simulation to solve for the ground states of quantum systems.

1) *Imaginary Time Evolution*: We simulate the spin- $\frac{1}{2}$ J1-J2 Heisenberg Model [58] with the Hamiltonian defined as the operator,

$$\begin{aligned}
 H = & \sum_{\langle ij \rangle} (J_1^x X_i X_j + J_1^y Y_i Y_j + J_1^z Z_i Z_j) \\
 & + \sum_{\langle\langle ij \rangle\rangle} (J_2^x X_i X_j + J_2^y Y_i Y_j + J_2^z Z_i Z_j) \\
 & + \sum_i (h^x X_i + h^y Y_i + h^z Z_i), \quad (7)
 \end{aligned}$$

where X , Y , and Z are Pauli operators and the indices/subscripts appearing in each Pauli operator indicate the site the operators act on. The notation $\sum_{\langle ij \rangle}$ denotes a sum over pairs of sites (i and j) that are adjacent on a 2D lattice, while $\sum_{\langle\langle ij \rangle\rangle}$ is the same for pairs of sites that are diagonally adjacent (e.g., sites at lattice points $i = (k, l)$ and $j = (k + 1, l - 1)$). Further, J_n defines the coupling constants between neighbors, and h represents the strength of a transverse magnetic field along a particular axis.

We perform ITE steps as described in Section II-D1 on PEPS with evolution bond dimension r and contract the resulting PEPS using IBMPS with contraction bond dimension m . Figure 13 shows the iterations of PEPS ITE for small bond dimensions and the result of PEPS ITE after 150 steps as the bond dimension grows for a 4×4 J1-J2 model. The energy of PEPS calculated by IBMPS converges to the ground state calculated by the exact simulation (using the full state vector) after 1000 ITE steps when we increase the bond

dimension. Also, we compare the choice of contraction bond dimension $m = r^2$ and $m = r$ and observe that for this model, their accuracy is similar while the latter requires much less computation.

2) *Variational Quantum Eigensolver Simulation*: We benchmark VQE with the spin- $\frac{1}{2}$ J1-J2 Heisenberg Model (see Equation (7)). The parameters are set as $J_1^x = J_1^y = J_1^z = 1.0$, $J_2^x = J_2^y = J_2^z = 0.5$ and $h_x = h_y = h_z = 0.2$.

We utilize the *constrained optimization by linear approximation* (COBYLA) algorithm [59], provided via the `scipy.optimize.minimize` function in Python. Our variational ansatz, which is passed to the optimizer, applies a parameterized quantum circuit composed of repeated layers to construct the state prior to calculating the expectation value of the energy. Each circuit layer consists of single rotation gates $R_y(\theta) = e^{-i\theta Y/2}$ applied to each qubit. The rotation gates are followed by a set of CNOT gates, which are applied to each nearest neighbor pair.

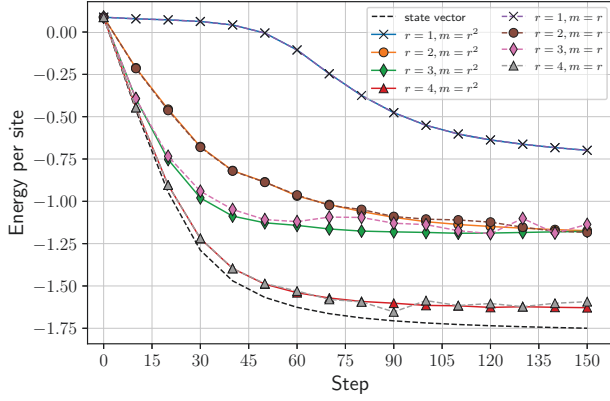
Each VQE iteration involves the evolution and contraction of PEPS in order to calculate the expectation value, which creates a substantial bottleneck in terms of running time. VQE itself involves running hundreds or even thousands of objective function evaluations prior to convergence. Consequently, we focus our accuracy experiments on a small number of qubits with bounded iterations.

As can be seen in Figure 14, for a 4×4 qubit system increasing the evolution and contraction bond dimensions used in the PEPS simulation improves the lowest energy per site that each system reaches. The energy obtained is -0.55990 , -0.84822 , and -1.46659 for a maximum evolution bond dimension of 1, 2, and 4, and a maximum contraction bond dimension of 4, 16, and 64, respectively. By comparison, the state vector system reaches a value of -1.59796 , while the approximate ground state energy per site (computed from a state vector ITE run) is -1.77688 .


VII. RELATED WORK

While algorithms and parallel computations with PEPS are a relatively new area, there exist several works that have taken similar approaches to ours. The PEPS++ method has been proposed by He et al. as a computational method that enables parallel computation for quantum systems simulation [60]. To overcome the high order of complexity, i.e. $O(r^{10})$, they combine quantum Monte Carlo methods with PEPS algorithms to avoid the two-layer PEPS contraction. Instead, they perform one-layer PEPS contraction along with a sampling procedure, which reduces the asymptotic cost to $O(r^6)$ for each sampling step. In the implementation, they use processes-level parallelism for sampling steps and thread-level parallelism for tensor computations, which differs from the distributed-memory tensor computation considered in this work.

Like our work, Guo et al. have carried out large-scale parallel quantum circuit simulations utilizing PEPS [12]. They simulate random quantum circuits exactly by applying gate operators on PEPS without truncation, and they calculate amplitudes by projecting the resulting PEPS onto the correspond-



(a) Energies calculated by IBMPS at each step for small bond dimensions.



(b) Energies calculated by IBMPS after 150 steps as the bond dimension grows.

Figure 13: PEPS ITE results of the 4×4 J1-J2 model with $J_1^x = J_1^y = J_1^z = 1.0$, $J_2^x = J_2^y = J_2^z = 0.5$ and $h^x = h^y = h^z = 0.2$.

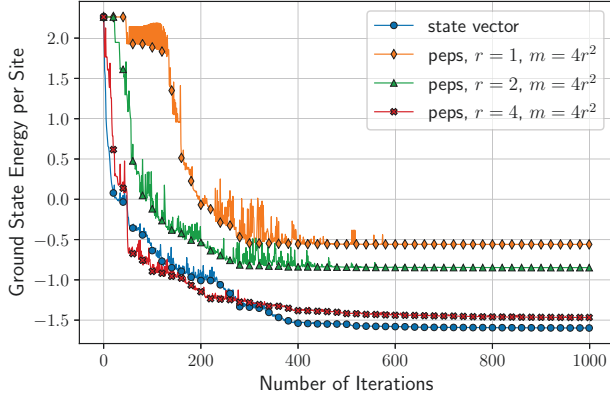


Figure 14: VQE simulation for the 4×4 J1-J2 model with product state initialization. Energies are calculated by IBMPS. Every 20 iterations, markers show the minimum energies found by the simulation up to that point.

ing basis states. The cost of this method is associated with the entanglement generated by the quantum circuit instead of the circuit size, but could still be exponential because it requires exact PEPS contraction. They also use Cyclops for distributed-memory tensor computations, and their results include an exact contraction of a 7×7 PEPS with bond dimension 32 using 4096 nodes of the Tianhe-2 supercomputer. Our work differs from theirs in that we focus on approximate PEPS algorithms with polynomial time complexity, which allows for larger bond dimensions with a controlled error and consequently serves different application needs.

A related algorithm, the *higher-order tensor renormalization group* (HOTRG) has also been studied in a distributed parallel setting [61], where an optimal reordering procedure for tensor contractions used in HOTRG is proposed. In comparison, our work provides a more general approach via the

use of Cyclops, which automates the performance optimization within tensor contractions. We also consider a broad set of PEPS primitives as opposed to a single component of a PEPS contraction algorithm.

There are also a number of libraries available for quantum system simulation using tensor networks, such as ITensor [62], NCON [63], quimb [64], TeNPy [65], Uni10 [66], and qTorch [67]. These libraries implement useful algorithms in the field with good sequential, threaded, or GPU performance. Our library, Koala, contributes to this community with the use of distributed-memory parallelism.

VIII. CONCLUSION

As demonstrated by our application studies, tensor network methods can systematically improve the accuracy for approximate simulation of quantum systems and near-term quantum devices by using larger bond dimensions (i.e. larger tensors). The new algorithms and Koala library we introduced provide an efficient approach for PEPS evolution and contraction and enable the use of distributed-memory tensors. Performance studies on the Stampede2 supercomputer indicate that our approach provides scalable PEPS primitives and systematically improvable accuracy.

ACKNOWLEDGMENTS

We would like to thank Phillip Helms, Ryan Levy, Garnet Chan, and Bryan Clark for helpful conversations and insights on tensor network methods. Yuchen Pang, Annika Dugad, and Edgar Solomonik were supported via the US NSF OAC RAISE/TAQS program, award number 1839204. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. We used XSEDE to employ Stampede2 at the Texas Advanced Computing Center (TACC) through allocation TG-CCR180006.

REFERENCES

- [1] R. Orús, “A practical introduction to tensor networks: Matrix product states and projected entangled pair states,” *Annals of Physics*, vol. 349, pp. 117 – 158, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0003491614001596>
- [2] G. Vidal, “Efficient classical simulation of slightly entangled quantum computations,” *Phys. Rev. Lett.*, vol. 91, p. 147902, Oct 2003. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.91.147902>
- [3] U. Schollwöck, “The density-matrix renormalization group in the age of matrix product states,” *Annals of Physics*, vol. 326, no. 1, p. 96–192, Jan 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.aop.2010.09.012>
- [4] F. Verstraete and J. I. Cirac, “Renormalization algorithms for quantum-many body systems in two and higher dimensions,” *arXiv preprint cond-mat/0407066*, 2004.
- [5] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, “Matrix product density operators: Simulation of finite-temperature and dissipative systems,” *Physical Review Letters*, vol. 93, no. 20, Nov 2004. [Online]. Available: <http://dx.doi.org/10.1103/PhysRevLett.93.207204>
- [6] G. Vidal, “Class of quantum many-body states that can be efficiently simulated,” *Phys. Rev. Lett.*, vol. 101, p. 110501, Sep 2008. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.101.110501>
- [7] Y.-Y. Shi, L.-M. Duan, and G. Vidal, “Classical simulation of quantum many-body systems with a tree tensor network,” *Physical review a*, vol. 74, no. 2, p. 022320, 2006.
- [8] V. Murg, F. Verstraete, and J. I. Cirac, “Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states,” *Phys. Rev. A*, vol. 75, p. 033605, Mar 2007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.75.033605>
- [9] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, “Classical simulation of infinite-size quantum lattice systems in two spatial dimensions,” *Phys. Rev. Lett.*, vol. 101, p. 250602, Dec 2008. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.101.250602>
- [10] J. Preskill, “Quantum computing in the NISQ era and beyond,” *Quantum*, vol. 2, p. 79, Aug 2018. [Online]. Available: <http://dx.doi.org/10.22331/q-2018-08-06-79>
- [11] A. Shafaei, M. Saeedi, and M. Pedram, “Qubit placement to minimize communication overhead in 2D quantum architectures,” in *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, 2014, pp. 495–500.
- [12] C. Guo, Y. Liu, M. Xiong, S. Xue, X. Fu, A. Huang, X. Qiang, P. Xu, J. Liu, S. Zheng, H.-L. Huang, M. Deng, D. Poletti, W.-S. Bao, and J. Wu, “General-purpose quantum circuit simulator with projected entangled-pair states and the quantum supremacy frontier,” *Phys. Rev. Lett.*, vol. 123, p. 190501, Nov 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.123.190501>
- [13] N. Schuch, M. M. Wolf, F. Verstraete, and J. I. Cirac, “Computational complexity of projected entangled pair states,” *Phys. Rev. Lett.*, vol. 98, p. 140506, 4 2007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.98.140506>
- [14] T. E. Oliphant, *A guide to NumPy*. Trelgol Publishing USA, 2006, vol. 1.
- [15] R. Nishino and S. H. C. Loomis, “CuPy: A NumPy-compatible library for NVIDIA GPU calculations,” *31st conference on neural information processing systems*, p. 151, 2017.
- [16] E. Solomonik, D. Matthews, J. R. Hammond, J. F. Stanton, and J. Demmel, “A massively parallel tensor contraction framework for coupled-cluster computations,” *Journal of Parallel and Distributed Computing*, vol. 74, no. 12, pp. 3176–3190, 2014.
- [17] Z. Zhang, X. Wu, N. Zhang, S. Zhang, and E. Solomonik, “Enabling distributed-memory tensor completion in python using new sparse tensor kernels,” *ArXiv e-prints arXiv:1910.02371*, 2019.
- [18] L. Grasedyck, D. Kressner, and C. Tobler, “A literature survey of low-rank tensor approximation techniques,” *GAMM-Mitteilungen*, vol. 36, no. 1, pp. 53–78, 2013.
- [19] T. Kolda and B. Bader, “Tensor decompositions and applications,” *SIAM Review*, vol. 51, no. 3, pp. 455–500, 2009. [Online]. Available: <http://epubs.siam.org/doi/abs/10.1137/07070111X>
- [20] I. V. Oseledets, “Tensor-train decomposition,” *SIAM Journal on Scientific Computing*, vol. 33, no. 5, pp. 2295–2317, 2011.
- [21] F. Verstraete, V. Murg, and J. I. Cirac, “Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems,” *Advances in Physics*, vol. 57, no. 2, pp. 143–224, 2008.
- [22] B. N. Khoromskij and I. V. Oseledets, “QTT approximation of elliptic solution operators in higher dimensions,” *Russian Journal of Numerical Analysis and Mathematical Modelling*, vol. 26, no. 3, pp. 303–322, 2011.
- [23] C. Hubig, I. McCulloch, and U. Schollwöck, “Generic construction of efficient matrix product operators,” *Physical Review B*, vol. 95, no. 3, p. 035129, 2017.
- [24] H. C. Jiang, Z. Y. Weng, and T. Xiang, “Accurate determination of tensor network state of quantum lattice models in two dimensions,” *Phys. Rev. Lett.*, vol. 101, p. 090603, Aug 2008. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.101.090603>
- [25] Y. Zhou, E. M. Stoudenmire, and X. Waintal, “What limits the simulation of quantum computers?” *ArXiv e-prints arXiv:2002.07730*, Feb 2020.
- [26] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, “A variational eigensolver on a photonic quantum processor,” *Nature communications*, vol. 5, p. 4213, 2014.
- [27] J.-G. Liu, Y.-H. Zhang, Y. Wan, and L. Wang, “Variational quantum eigensolver with fewer qubits,” *Physical Review Research*, vol. 1, no. 2, Sep 2019. [Online]. Available: <http://dx.doi.org/10.1103/PhysRevResearch.1.023025>
- [28] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, “Scalable quantum simulation of molecular energies,” *Phys. Rev. X*, vol. 6, p. 031007, Jul 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.6.031007>
- [29] H. R. Grimsley, S. E. Economou, E. Barnes, and N. J. Mayhall, “An adaptive variational algorithm for exact molecular simulations on a quantum computer,” *Nature communications*, vol. 10, no. 1, pp. 1–9, 2019.
- [30] R. M. Parrish, E. G. Hohenstein, P. L. McMahon, and T. J. Martínez, “Quantum computation of electronic transitions using a variational quantum eigensolver,” *Phys. Rev. Lett.*, vol. 122, p. 230401, Jun 2019. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.122.230401>
- [31] I. Pižorn, L. Wang, and F. Verstraete, “Time evolution of projected entangled pair states in the single-layer picture,” *Physical Review A*, vol. 83, no. 5, p. 052321, 2011.
- [32] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, “Unifying projected entangled pair state contractions,” *New Journal of Physics*, vol. 16, no. 3, p. 033014, 2014.
- [33] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, “Algorithms for finite projected entangled pair states,” *Physical Review B*, vol. 90, no. 6, Aug 2014. [Online]. Available: <http://dx.doi.org/10.1103/PhysRevB.90.064425>
- [34] M. Levin and C. P. Nave, “Tensor renormalization group approach to two-dimensional classical lattice models,” *Phys. Rev. Lett.*, vol. 99, p. 120601, 9 2007. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.99.120601>
- [35] Z.-C. Gu, M. Levin, and X.-G. Wen, “Tensor-entanglement renormalization group approach as a unified method for symmetry breaking and topological phase transitions,” *Phys. Rev. B*, vol. 78, p. 205116, 11 2008. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.78.205116>
- [36] Z.-C. Gu and X.-G. Wen, “Tensor-entanglement-filtering renormalization approach and symmetry-protected topological order,” *Phys. Rev. B*, vol. 80, p. 155131, Oct 2009. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.80.155131>
- [37] Z. Y. Xie, H. C. Jiang, Q. N. Chen, Z. Y. Weng, and T. Xiang, “Second renormalization of tensor-network states,” *Phys. Rev. Lett.*, vol. 103, p. 160601, Oct 2009. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.103.160601>
- [38] H. H. Zhao, Z. Y. Xie, Q. N. Chen, Z. C. Wei, J. W. Cai, and T. Xiang, “Renormalization of tensor-network states,” *Phys. Rev. B*, vol. 81, p. 174411, May 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.81.174411>
- [39] Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang, and T. Xiang, “Coarse-graining renormalization by higher-order singular value decomposition,” *Phys. Rev. B*, vol. 86, p. 045139, Jul 2012. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.86.045139>

- [40] G. Evenbly and G. Vidal, "Tensor network renormalization," *Phys. Rev. Lett.*, vol. 115, p. 180405, Oct 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.115.180405>
- [41] H.-H. Zhao, Z.-Y. Xie, T. Xiang, and M. Imada, "Tensor network algorithm by coarse-graining tensor renormalization on finite periodic lattices," *Phys. Rev. B*, vol. 93, p. 125115, Mar 2016. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.93.125115>
- [42] S. Yang, Z.-C. Gu, and X.-G. Wen, "Loop optimization for tensor network renormalization," *Phys. Rev. Lett.*, vol. 118, p. 110504, Mar 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.118.110504>
- [43] T. Nishino and K. Okunishi, "Corner transfer matrix renormalization group method," *Journal of the Physical Society of Japan*, vol. 65, no. 4, pp. 891–894, 1996.
- [44] R. Orús and G. Vidal, "Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction," *Phys. Rev. B*, vol. 80, p. 094403, Sep 2009. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.80.094403>
- [45] Z. Y. Xie, J. Chen, J. F. Yu, X. Kong, B. Normand, and T. Xiang, "Tensor renormalization of quantum many-body systems using projected entangled simplex states," *Phys. Rev. X*, vol. 4, p. 011025, Feb 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevX.4.011025>
- [46] Z. Y. Xie, H. J. Liao, R. Z. Huang, H. D. Xie, J. Chen, Z. Y. Liu, and T. Xiang, "Optimized contraction scheme for tensor-network states," *Phys. Rev. B*, vol. 96, p. 045128, Jul 2017. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevB.96.045128>
- [47] E. Stoudenmire and S. R. White, "Minimally entangled typical thermal state algorithms," *New Journal of Physics*, vol. 12, no. 5, p. 055026, 2010.
- [48] N. Halko, P. G. Martinsson, and J. A. Tropp, "Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions," *SIAM Review*, vol. 53, no. 2, pp. 217–288, 2011. [Online]. Available: <https://doi.org/10.1137/090771806>
- [49] D. Tamascelli, R. Rosenbach, and M. B. Plenio, "Improved scaling of time-evolving block-decimation algorithm through reduced-rank randomized singular value decomposition," *Phys. Rev. E*, vol. 91, p. 063306, Jun 2015. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevE.91.063306>
- [50] S. Morita, R. Igarashi, H.-H. Zhao, and N. Kawashima, "Tensor renormalization group with randomized singular value decomposition," *Physical Review E*, vol. 97, no. 3, p. 033310, 2018.
- [51] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet *et al.*, *ScaLAPACK users' guide*. Siam, 1997, vol. 4.
- [52] E. Wang, Q. Zhang, B. Shen, G. Zhang, X. Lu, Q. Wu, and Y. Wang, "Intel math kernel library," in *High-Performance Computing on the Intel® Xeon Phi™*. Springer, 2014, pp. 167–188.
- [53] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, "Characterizing quantum supremacy in near-term devices," *Nature Physics*, vol. 14, no. 6, pp. 595–600, 2018.
- [54] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [55] S. Aaronson and L. Chen, "Complexity-theoretic foundations of quantum supremacy experiments," in *32nd Computational Complexity Conference (CCC 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.
- [56] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya *et al.*, "A blueprint for demonstrating quantum supremacy with superconducting qubits," *Science*, vol. 360, no. 6385, pp. 195–199, 2018.
- [57] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, "Quantum supremacy and the complexity of random circuit sampling," *arXiv preprint arXiv:1803.04402*, 2018.
- [58] J.-F. Yu and Y.-J. Kao, "Spin-1/2 J_1 - J_2 Heisenberg antiferromagnet on a square lattice: A plaquette renormalized tensor network study," *Physical Review B*, vol. 85, no. 9, Mar 2012. [Online]. Available: <http://dx.doi.org/10.1103/PhysRevB.85.094407>
- [59] M. J. D. Powell, "A direct search optimization method that models the objective and constraint functions by linear interpolation," *Advances in Optimization and Numerical Analysis*, 1994.
- [60] L. He, H. An, C. Yang, F. Wang, J. Chen, C. Wang, W. Liang, S. Dong, Q. Sun, W. Han, W. Liu, Y. Han, and W. Yao, "PEPS++: Towards extreme-scale simulations of strongly correlated quantum many-particle models on sunway tainhulight," 2018.
- [61] H. Yamada, A. Imakura, T. Imamura, and T. Sakurai, "Optimization of reordering procedures in hotrg for distributed parallel computing," in *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 2018, pp. 957–966.
- [62] ITensor Library <http://itensor.org>.
- [63] R. N. C. Pfeifer, G. Evenbly, S. Singh, and G. Vidal, "NCON: A tensor network contractor for matlab," 2014.
- [64] J. Gray, "quimb: A python package for quantum information and many-body calculations," *Journal of Open Source Software*, vol. 3, no. 29, p. 819, 2018.
- [65] J. Hauschild and F. Pollmann, "Efficient numerical simulations with tensor networks: Tensor Network Python (TeNPy)," *SciPost Physics Lecture Notes*, 2018.
- [66] Y.-J. Kao, Y.-D. Hsieh, and P. Chen, "Uni10: An open-source library for tensor network algorithms," in *Journal of Physics: Conference Series*, vol. 640, no. 1. IOP Publishing, 2015, p. 012040.
- [67] E. S. Fried, N. P. Sawaya, Y. Cao, I. D. Kivlichan, J. Romero, and A. Aspuru-Guzik, "qTorch: The quantum tensor contraction handler," *PLoS one*, vol. 13, no. 12, 2018.

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

We implemented the algorithms discussed in the paper as an open-source Python library. We benchmarked our implementation and compared the performance of different variants for PEPS evolution and PEPS contraction. We also ran experiments to study the accuracy of approximated PEPS contraction and the overall accuracy of imaginary time evolution (ITE) and variational quantum eigensolver (VQE) simulation using PEPS. Parallel experiments are all done on the Stampede2 supercomputer KNL nodes with 64 processes per node for Cyclops Tensor Framework (CTF) and 32 MKL threads on one node for NumPy.

ARTIFACT AVAILABILITY

Software Artifact Availability: All author-created software artifacts are maintained in a public repository under an OSI-approved license.

Hardware Artifact Availability: There are no author-created hardware artifacts.

Data Artifact Availability: There are no author-created data artifacts.

Proprietary Artifacts: None of the associated artifacts, author-created or otherwise, are proprietary.

Author-Created or Modified Artifacts:

Persistent ID: <https://doi.org/10.5281/zenodo.4008217>
Artifact name: Koala

Persistent ID: <https://doi.org/10.5281/zenodo.4008214>
Artifact name: TensorBackends

with the results from analytical calculations. Moreover, we utilized models that had been previously simulated in existing literature, and we referred to several published works to verify our results.

Accuracy and precision of timings: We ran experiments multiple times to get an average running time. When we recorded the running times that are not consistent, we would look into them and study the cause.

BASELINE EXPERIMENTAL SETUP, AND MODIFICATIONS MADE FOR THE PAPER

Relevant hardware details: Stampede2 KNL nodes

Operating systems and versions: Linux

Compilers and versions: Python 3.7; gcc v18.0.2

Libraries and versions: Intel MPI v18.0.2

Key algorithms: Randomized SVD, Tensor contraction, Time-evolving block decimation for PEPS evolution and Boundary MPS for PEPS contraction

ARTIFACT EVALUATION

Verification and validation studies: We developed unit tests for each functionality of our library to ensure the correctness of each step. Also, we compared our results for easy cases that can be solved directly with widely used libraries like NumPy. When it's possible, we also validated the accuracy of our results by comparing them