Tensor Optimization Algorithms and Libraries for Quantum Simulation

Edgar Solomonik

L P. N A @CS@Illinois

Department of Computer Science University of Illinois at Urbana-Champaign

Quantum Information for Mathematics, Economics, and Statistics IMSI

Laboratory for Parallel Numerical Algorithms

Recent/ongoing research topics (*-covered today)

- parallel matrix computations
 - matrix factorizations
 - eigenvalue problems
 - preconditioners
- tensor computations
 - tensor decomposition*
 - sparse tensor kernels*
 - tensor completion*
- simulation of quantum systems
 - tensor networks*
 - quantum chemistry
 - quantum circuits*
- fast bilinear algorithms
 - convolution algorithms
 - tensor symmetry*



http://lpna.cs.illinois.edu

Outline

Introduction

- 2 Tensor Contractions
- 3 Tensor Decompositions

4 Tensor Networks



Tensors

A tensor is a collection of elements

- its dimensions define the size of the collection
- its order is the number of different dimensions
- specifying an index along each tensor mode defines an element of the tensor
- A few examples of tensors are
 - Order 0 tensors are scalars, e.g., $s \in \mathbb{R}$
 - Order 1 tensors are vectors, e.g., $oldsymbol{v} \in \mathbb{R}^n$
 - Order 2 tensors are matrices, e.g., $oldsymbol{A} \in \mathbb{R}^{m imes n}$
 - An order 3 tensor with dimensions $s_1 \times s_2 \times s_3$ is denoted as $\mathcal{T} \in \mathbb{R}^{s_1 \times s_2 \times s_3}$ with elements t_{ijk} for $i \in \{1, \ldots, s_1\}, j \in \{1, \ldots, s_2\}, k \in \{1, \ldots, s_3\}$



A tensor contraction multiplies elements of two tensors and computes partial sums to produce a third, in a fashion expressible by pairing up modes of different tensors, which may be expressed in Einstein summation notation (einsum)

| tensor contraction | einsum |
|-----------------------|-------------------------------------|
| inner product | $w = \sum_{i} u_i v_i$ |
| outer product | $w_{ij} = u_i v_{ij}$ |
| pointwise product | $w_i = u_i v_i$ |
| Hadamard product | $w_{ij} = u_{ij}v_{ij}$ |
| matrix multiplication | $w_{ij} = \sum_k u_{ik} v_{kj}$ |
| batched matmul. | $w_{ijl} = \sum_k u_{ikl} v_{kjl}$ |
| tensor times matrix | $w_{ilk} = \sum_{j} u_{ijk} v_{lj}$ |

General Tensor Contractions

Given tensor \mathcal{U} of order s + v and \mathcal{V} of order v + t, a tensor contraction summing over v modes can be written as

$$w_{i_1\dots i_s j_1\dots j_t} = \sum_{k_1\dots k_v} u_{i_1\dots i_s k_1\dots k_v} v_{k_1\dots k_v j_1\dots j_t}$$

• Other contractions can be mapped to this form after transposition

Unfolding the tensors reduces the tensor contraction to matrix multiplication

- Combine consecutive indices in appropriate groups of size s ,t, or v
- If all tensor modes are of dimension n, obtain matrix-matrix product C = AB where $C \in \mathbb{R}^{n^s \times n^t}$, $A \in \mathbb{R}^{n^s \times n^v}$, and $B \in \mathbb{R}^{n^v \times n^t}$
- Assuming classical matrix multiplication, contraction requires n^{s+t+v} elementwise products and $n^{s+t+v}-n^{s+t}$ additions

Emulation of Quantum Gates and Quantum Circuits

• Consider an *n*-qubit quantum state

• Quantum circuits generally consist of 1-qubit and 2-qubit gates

$$\begin{split} |\phi\rangle &= U^{(s)} |\psi\rangle \Rightarrow t^{\phi}_{i_{1}\cdots i_{n}} = \sum_{j_{s}=0}^{1} u^{(s)}_{i_{s}j_{s}} t^{\psi}_{i_{1}\cdots i_{s-1}j_{s}i_{s+1}\cdots i_{n}} \\ \phi\rangle &= U^{(s,t)} |\psi\rangle \Rightarrow t^{\phi}_{i_{1}\cdots i_{n}} = \sum_{j_{s}=0}^{1} \sum_{j_{t}=0}^{1} u^{(s)}_{i_{s}i_{t}j_{s}j_{t}} t^{\psi}_{i_{1}\cdots i_{s-1}j_{s}i_{s+1}\cdots i_{t-1}j_{t}i_{t+1}\cdots i_{n}} \end{split}$$

• A quantum gate can be emulated as an $O(2^n)$ -cost tensor contraction

• An *n*-qubit quantum circuit with depth D and O(nD) gates can be simulated classically with $O(nD2^n)$ cost and $O(2^n)$ storage

Symmetric Tensor Contractions

- Recall a symmetric tensor is defined by e.g., $t_{ijk} = t_{ikj} = t_{kij} = t_{jki} = t_{jik} = t_{kji}$
- Tensors can also have skew-symmetry (also known as antisymmetry, permutations have +/- signs), partial symmetry (only some modes are permutable), or group symmetry (blocks are zero if indices satisfy modular equation)
- The simplest example of a symmetric tensor contraction is

$$\boldsymbol{y} = \boldsymbol{A} \boldsymbol{x}$$
 where $\boldsymbol{A} = \boldsymbol{A}^T$

it is not obvious how to leverage symmetry to reduce cost of this contraction

Permutational Symmetry in Tensor Contractions



New contraction algorithms reduce cost via permutational symmetry¹

- Symmetry is hard to use in contraction e.g. $m{y}=Am{x}$ with A symmetric
- For contraction of order s + v and v + t tensors to produce an order s + t tensor, previously known approaches reduce cost by s!t!v!
- New algorithm reduces number of *products* by ω ! where $\omega = s + t + v$, leads to same reduction in *cost* for partially-symmetric contractions

$$\boldsymbol{C} = \boldsymbol{A}\boldsymbol{B} + \boldsymbol{B}\boldsymbol{A} \Rightarrow c_{ij} = \sum_{k} [(a_{ij} + a_{ik} + a_{jk}) \cdot (b_{ij} + b_{ik} + b_{jk})] - \dots$$

LPNA

¹E.S, J. Demmel, CMAM 2020

Group Symmetry

- Tensors arising in physical simulations often have group structure that reflects conservation laws
 - Abelian group symmetries can be mapped to cyclic group, which can be used to define a block-sparse form of the tensors (here represented using extra modes), e.g.,

$$w_{aA,bB,iI,jJ} = \sum_{k,K,l,L} u_{aA,bB,kK,lL} v_{kK,lL,iI,jJ}.$$

where for some group size G, we have symmetries, e.g.,

$$\begin{aligned} & w_{aA,bB,iI,jJ} \neq 0 \text{ if } A + B - I - J \equiv 0 \pmod{G}, \\ & u_{aA,bB,kK,lL} \neq 0 \text{ if } A + B + K + L \equiv 0 \pmod{G}, \\ & v_{kK,lL,iI,jJ} \neq 0 \text{ if } K + L - I - J \equiv 0 \pmod{G}. \end{aligned}$$

• We can write each of these tensors using a reduced form and an irrep map,

$$w_{aA,bB,iI,jJ} = r_{aA,bB,iI,j}^{(W)} m_{ABIJ}$$

where $m_{ABIJ} = 1$ if $A + B - I - J \equiv 0 \pmod{G}$ and $m_{ABIJ} = 0$ otherwise

LPNA

Tensor Algorithms and Software

Block Contraction Approach to Group Symmetry

Algorithm 3.1 Loop nest to perform group symmetric contraction $w_{aA,bB,iI,jJ} = \sum_{k,K,l,L} u_{aA,bB,kK,lL}v_{kK,lL,iI,jJ}$. for A = 1, ..., G do for B = 1, ..., G do $J = A + B - I \mod G$ for K = 1, ..., G do $L = -A - B - K \mod G$ $\forall a, b, i, j, \quad w_{aA,bB,iI,jJ} = \sum_{k,l} w_{aA,bB,iI,jJ} + u_{aA,bB,kK,lL}, v_{kK,lL,iI,jJ}$ end for end for end for

Group Symmetry in Tensor Contractions

New contraction algorithm, *irreducible representation alignment* uses new reduced form to handle group symmetry (momentum conservation, spin, quantum numbers, etc.) without looping over blocks or sparsity¹

$$w_{aA,b,i,jJ,Q} = w_{aA,b,Q-A \mod G,i,-J-Q \mod G,jJ},$$
$$= \sum_{L,k,l} \hat{u}_{aA,b,k,lL,Q} \hat{v}_{k,lL,i,jJ,Q}.$$



¹ collaboration with Yang Gao, Phillip Helms, and Garnet Chan at Caltech, to appear on arxiv, July 2020

Tensor Algorithms and Software

Automation of Group Symmetric Contractions



Group symmetric tensors represented programmatically by

- a dense reduced tensor (containing unique data)
- an implicit sparse tensor (irrep map) describing the group symmetry
- At contraction time reduced form is transformed so as to align irrep maps (introduce Q index)
- Users can write symmetry-oblivious code

Library for Massively-Parallel Tensor Computations

Cyclops Tensor Framework¹ sparse/dense generalized tensor algebra

- Cyclops is a C++ library that distributes each tensor over MPI
- Used in chemistry (PySCF, QChem)², quantum circuit simulation (IBM/LLNL)³, and graph analysis (betweenness centrality)⁴
- Summations and contractions specified via Einstein notation

E["aixbjy"] += X["aixbjy"] - U["abu"]*V["iju"]*W["xyu"]

- Best distributed contraction algorithm selected at runtime via models
- Support for Python (numpy.ndarray backend), OpenMP, and GPU
- Simple interface to core ScaLAPACK matrix factorization routines

¹https://github.com/cyclops-community/ctf

²E.S., D. Matthews, J. Hammond, J.F. Stanton, J. Demmel, JPDC 2014

E. Pednault, J.A. Gunnels, G. Nannicini, L. Horesh, T. Magerlein, E. S., E. Draeger, E. Holland, and R. Wisnieff, 2017
 E.S., M. Besta, F. Vella, T. Hoefler, SC 2017

CP Decomposition



• For a tensor $T \in \mathbb{R}^{n \times n \times n}$, the CP decomposition¹ is defined by matrices U, V, and W such that

$$t_{ijk} = \sum_{r=1}^{R} u_{ir} v_{jr} w_{kr}$$

• For an order N tensor, the decomposition generalizes as follows,

$$t_{i_1\dots i_d} = \sum_{r=1}^R \prod_{j=1}^d u_{i_j r}^{(j)}$$

• Its rank is generally bounded by $R \le n^{d-1}$

¹Kolda and Bader, SIAM Review 2009

LPNA

Tensor Algorithms and Software

Alternating Least Squares for CP Decomposition

Consider rank R CP decomposition of an $s \times s \times s \times s$ tensor $x_{ijkl} \approx \sum_{r=1}^{R} u_{ir} v_{jr} w_{kr} z_{lr}$ $(\bigvee_{l_j} \bigvee_{l_j} \bigvee_{l_k} \bigvee_{l_j} \bigvee_{l_j} \bigvee_{l_k} \bigvee_{l_j} \bigvee_{l_k} \bigvee_{l_j} \bigvee_{l_k} \bigvee_{l$

Alternating Least Squares for CP Decomposition

Consider rank R CP decomposition of an $s \times s \times s \times s$ tensor

ALS updates factor matrices in an alternating manner

$$\min_{\boldsymbol{A}^{(n)}} f(\boldsymbol{A}^{(1)}, \dots, \boldsymbol{A}^{(N)}) = \frac{1}{2} ||\boldsymbol{\mathcal{X}} - [\![\boldsymbol{A}^{(1)}, \dots, \boldsymbol{A}^{(n)}, \dots, \boldsymbol{A}^{(N)}]\!]||_{F}^{2},$$



Alternating Least Squares for CP Decomposition

Consider rank R CP decomposition of an $s \times s \times s \times s$ tensor

ALS updates factor matrices in an alternating manner

$$\min_{\mathbf{A}^{(n)}} f(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) = \frac{1}{2} || \mathbf{\mathcal{X}} - [\![\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(n)}, \dots, \mathbf{A}^{(N)}]\!] ||_{F}^{2},$$



Each quadratic subproblem is typically solved via normal equations



Tensor Contractions in CP ALS

The normal equations are cheap to compute



Tensor Contractions in CP ALS

The normal equations are cheap to compute



But forming the right-hand sides $(M^{(n)})$ requires expensive MTTKRP (matricized tensor-times Khatri-Rao product)



Pairwise Perturbation Algorithm



New algorithm: pairwise perturbation $(PP)^1$ approximates ALS

- based on perturbative expansion of ALS update to approximate MTTKRP
- approximation is accurate when ALS updates stagnate
- rank $R < s^{N-1}$ CP decomposition:
 - ALS sweep cost $O(s^N R) \Rightarrow O(s^2 R)$, up to 33x speed-up

¹Linjian Ma, E.S. arXiv:1811.10573



Linjian Ma

Parallel Pairwise Perturbation Algorithm



Effective parallelization by decomposing MTTKRP into local MTTKRPs ¹

$$U = \mathsf{MTTKRP}(\mathcal{T}, \mathcal{V}, \mathcal{W}) \Rightarrow U_i = \sum_{j,k} \mathsf{MTTKRP}(\mathcal{T}_{ijk}, \mathcal{V}_j, \mathcal{W}_k)$$

• processor (i,j,k) owns $oldsymbol{\mathcal{T}}_{ijk}$, $oldsymbol{V}_{j}$, and $oldsymbol{W}_{k}$

- pairwise perturbation can be used to approximate local MTTKRPs
- multi-sweep dimension-tree (MSDT) amortizes terms across sweeps

¹Linjian Ma, E.S. IPDPS 2021

LPNA

Tucker Decomposition



- The Tucker decomposition¹ expresses an order *d* tensor via a smaller order *d* core tensor and *d* factor matrices
 - For a tensor $\mathcal{T} \in \mathbb{R}^{n \times n \times n}$, the Tucker decomposition is defined by core tensor $\mathcal{Z} \in \mathbb{R}^{R_1 \times R_2 \times R_3}$ and factor matrices U, V, and W with orthonormal columns, such that

$$t_{ijk} = \sum_{p=1}^{R_1} \sum_{q=1}^{R_2} \sum_{r=1}^{R_3} z_{pqr} u_{ip} v_{jq} w_{kr}$$

- If an exact Tucker decomposition exists, it can be computed via SVD (HoSVD)
- HOOI method optimizes in an alternating manner among (U, \mathcal{Z}) , (V, \mathcal{Z}) , (W, \mathcal{Z})

¹Kolda and Bader, SIAM Review 2009

Randomized Methods for Sparse Tensor Decomposition

- When seeking a low-rank R = O(1) decomposition for a sparse tensor, sketching schemes have been shown to be efficient
- In this regime, Tucker can be used to construct a CP decomposition
- Leverage score sampling on the rank-constrained least squares problem $\min_{\mathbf{X}, \operatorname{rank}(\mathbf{X}) \leq R} \| \mathbf{A}\mathbf{X} - \mathbf{B} \|_F$ leads to a state-of-the-art cost-accuracy trade-off¹ to effectively approximate HOOI (ALS)

| Algorithm for Tucker | LS solve cost | Sample size (m) |
|-----------------------------|--|---|
| ALS | $O(\mathrm{nnz}(\boldsymbol{\mathcal{T}})R^{N-1})$ | / |
| $ALS + TensorSketch^2$ | $\tilde{O}(mR^N + msR)$ | $O(R^{2(N-1)} \cdot 3^{N-1}/(\epsilon^2 \delta))$ |
| $ALS + TTMTS^2$ | $\tilde{O}(msR^{N-1})$ | $O(R^{2(N-1)} \cdot 3^{N-1} / (\epsilon^2 \delta))$ |
| $ALS + TensorSketch^1$ | $\tilde{O}(mR^{2N-2} + sR^{N-1})$ | $O((R^{(N-1)} + 1/\epsilon^2) \cdot (3R)^{(N-1)}/\delta)$ |
| $ALS + Ieverage \ scores^1$ | $\tilde{O}(mR^{2N-2} + sR^{N-1})$ | $O(R^{(N-1)}/(\epsilon^2\delta))$ |

¹Linjian Ma and E.S., arXiv:2104.01101

LPNA

²O. Malik and S. Becker, 2018 (assuming unconstrained LSQ)

Tensor Completion

- The tensor completion problem seeks to build a model (e.g., CP decomposition) for a partially-observed tensor
- For an order three tensor $\mathcal{T} \in \mathbb{R}^{n \times n \times n}$, given a set of observed entries t_{ijk} for $(i, j, k) \in \Omega$, we seek to minimize

$$\sum_{(i,j,k)\in\Omega} \underbrace{(t_{ijk} - \sum_{r} u_{ir}v_{jr}w_{kr})^2}_{\text{loss function}} + \lambda^2 (\|\boldsymbol{U}\|_2^2 + \|\boldsymbol{V}\|_2^2 + \|\boldsymbol{W}\|_2^2)$$

- Completion objective differs from decomposition of a sparse tensor, as it excludes unobserved entries
- Other loss functions than quadratic loss are often interest for different tensor data

Tensor Completion



- Via the Cyclops Python interface, we have implemented parallel (over MPI) completion with SGD, CCD, ALS (with iterative and direct solves), and Gauss-Newton, with support for generalized loss¹
- Tensor times tensor product (TTTP) routine enables CP tensor completion

$$r_{ijk} = \sum t_{ijk} u_{ir} v_{jr} w_{kr}$$

• For ALS, explicit parallel direct ^rsolves² are fastest

 ¹Navjot Singh, Zecheng Zhang, Xiaoxiao Wu, Naijing Zhang, Siyuan Zhang, and Edgar Solomonik arXiv:1910.02371
 ²Shaden Smith, Jongsoo Park, and George Karypis, 2016

Hamiltonians as Tensor Network Operators

- Tensor network methods use tensor decompositions to represent quantum systems
- These methods are most natural for lattice spin systems such as the Heisenberg model and the simpler transverse field Ising model

$$H = \sum_{\langle i j \rangle} J^z Z_i Z_j + \sum_i h_x X_i$$

where $\langle i j \rangle$ denote neighboring sites on a 2D lattice

• In the 1D case, 2-qubit operators such as Z_iZ_{i+1} can be written as

 $H = Z \otimes Z \otimes I \otimes \cdots \otimes I + I \otimes Z \otimes Z \otimes I \otimes \cdots \otimes I + \cdots$

• In the 1D case, *H* can be represented as a matrix-product operator (MPO) with constant *bond dimension* (rank)

$$(a) \begin{array}{c} & & & \\ & & \\ & & \\ (b) \end{array} \begin{pmatrix} - & & \\ & - \end{array} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ & - \end{array} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ & - \end{array} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ & - \end{array} \begin{pmatrix} - & & \\ & - \end{array} \end{pmatrix} \begin{pmatrix} - & & \\ \end{array} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ \end{array} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ \end{array} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{array} \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix} \end{pmatrix} \end{pmatrix} \begin{pmatrix} - & & \\ \end{pmatrix} \end{pmatrix}$$

Density Matrix Renormalization Group (DMRG)



Parallel DMRG with Group Symmetry



We have recently developed a parallel DMRG code using Cyclops¹

- compare two approaches to group symmetry
 - iterate over block-wise contractions
 - use CTF's sparse tensor representation
- match ITensor efficiency at scale for spin-system, but significantly lower efficiency for fermionic system with large number of blocks

¹R. Levy, B. Clark, E.S. arXiv:2007.05540

Automatic Generation of Tensor Network Methods

- Automatic differentiation (AD) in principle enables automatic generation of methods such as ALS and DMRG
- Both apply Newton's method on a sequence of subsets of variables
- However, existing AD tools such as Jax (used by TensorFlow) are designed for deep learning and are ineffective for more complex tensor computations
 - focus purely on first order optimization via Jacobian-vector products
 - unable to propagate tensor algebra identities such as $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$ to generate efficient code

AutoHOOT: Automatic High-Order Optimization for Tensors

- AutoHOOT¹ provides a tensor-algebra centric AD engine
- Designed for einsum expressions and alternating minimization common in tensor decomposition and tensor network methods
- Python-level AD is coupled with optimization of contraction order and caching of intermediates
- Generates code for CPU/GPU/supercomputers using high-level back-end interface to tensor contractions



¹Linjian Ma, Jiayu Ye, and E.S. arXiv:2005.04540, 2020

LPNA

Tensor Algorithms and Software

Tensor Network State Evolution

• We can evolve a tensor network state by Trotterization of a Hamiltonian with m local terms

$$e^{-iH\tau} = \prod_{j=1}^{m} e^{-iH_j\tau} + O(\tau^2)$$

- \bullet Dynamics may be simulated by time-evolution $|\psi^{t+\tau}\rangle=e^{-iH\tau}|\psi^t\rangle$
- Ground state calculation can be done via imaginary time evolution, $|\psi^{i(t+\tau)}\rangle = e^{-H\tau}|\psi^{it}\rangle$, maximizing as follows

$$e^{-E\tau} = \max_{\|\psi\|_2} \langle \psi | e^{-H\tau} | \psi \rangle$$

which is equivalent to minimizing E and leads to the same maximizer/minimizer ψ

• If H_j is a local (e.g., one/two-site) operator, so is $e^{-iH_j \tau}$

Quantum Circuit Simulation with Tensor Networks

- Evolution of tensor network states can also simulate quantum circuits
- In fact, a quantum circuit is a direct description of a tensor network¹



• Why use HPC to (approximately) simulate quantum circuits?

- enable development/testing/tuning of larger quantum circuits
- understand approximability of different quantum algorithms
- quantify sensitivity of algorithms to noise/error
- potentially enable new hybrid quantum-classical algorithms
- Tensor network states (e.g., MPS/PEPS) provide robust tools for approximate simulation

¹Markov and Shi SIAM JC 2007

LPNA

Quantum Circuit Simulation using PEPS¹

- Near-term quantum architectures mostly connect qubits in a 2D fashion
- Non-local gates can be applied via the use of swap gates (with corresponding overhead)
- 2D tensor networks (projected entangles pair states (PEPS)) provide a natural way to simulate 2D quantum circuits
- Same software/algorithms infrastructure is also effective for (imaginary) time evolution with many Hamiltonians of interest
- Gate application and contraction of PEPS can both have exponential cost in the size of the circuit, so desire effective approximation

¹Yuchen Pang, Tianyi Hao, Annika Dugad, Yiqing Zhou, and E.S., arXiv:2006.15234.

Approximate Application of Two-Site Operators

- Consider application of a two-site operator on neighboring PEPS sites
- Simple update (QR-SVD) algorithm:



- We provide an efficient distributed implementation of QR-SVD
- This operation is an instance of what we'll refer to as einsumsvd and QR-SVD is one algorithm/implementation

Implicit Randomized einsumsvd

• The einsumsvd primitive will also enable effective algorithms for PEPS contraction



- An efficient general implementation is to leverage randomized SVD / orthogonal iteration, which iteratively computes a low-rank SVD by a matrix-matrix product that can be done implicitly via tensor contractions
- This approach improves complexity for PEPS boundary contraction method

Koala

• We introduce a new library, Koala¹, for high-performance simulation of quantum circuits and time evolution with PEPS

```
from koala import peps, Observable
     from tensorbackends.interface import ImplicitRandomizedSVD
 З
     # Create a 2-by-3 PEPS in distributed memory using CTF
 4
     gstate = peps.computational zeros(nrow=2, ncol=3, backend='ctf')
 6
 7
     # Construct operators and apply them to the quantum state
    Y = gstate.backend.astensor([0,-1i,1i,0]).reshape(2,2)
 8
 9
    CX = qstate.backend.astensor([1,0,0,0,0,1,0,0,0,0,0,0,1,0,0,1,0,])
    CX = CX, reshape(2, 2, 2, 2)
10
11
12
     gstate.apply operator(Y, [1])
     qstate.apply_operator(CX, [1,4], update_option=peps.QRUpdate(rank=2))
14
     # Construct an observable and calculate the expectation with IBMPS
15
     observable = 0bservable.ZZ(3, 4) + 0.2 \times 0bservable.X(1)
16
     result = gstate.expectation(
         observable. use cache=True.
18
19
         contract_option=peps.BMPS(ImplicitRandomizedSVD(rank=4)),
20
    )
```



¹https://github.com/cyclops-community/koala

PEPS Benchmark Performance



- Koala achieves good parallel scalability for approximate gate application (evolution) and contraction
- Approximation can be effective even for adversarially-designed circuits such as Google's random quantum circuit model (figure on right)

PEPS Accuracy for Quantum Simulation



- ITE code achieves improvable accuracy with increased PEPS bond dimension, but approximation in PEPS contraction is not variational
- Variational quantum eigensolver (VQE), which represents a wavefunction using a parameterized circuit $U(\theta)$ and minimizes

$\left< U(\theta) \right| H \left| U(\theta) \right>,$

also achieves improvable accuracy with higher PEPS bond dimension

- We have presented innovations to numerical algorithms and software libraries for tensor contractions, decompositions, and tensor networks
- Our research group is developing an ecosystem of algorithms and software for simulation of quantum systems
- This work is relevant to both classical methods for quantum chemistry and physics, as well as quantum computation
- See our website¹ and github pages for access to high performance software for tensor methods

http://lpna.cs.illinois.edu

Acknowledgements

- Laboratory for Parallel Numerical Algorithms (LPNA) at University of Illinois, lpna.cs.illinois.edu and collaborators
- Funding from NSF awards: #1839204 (RAISE-TAQS), #1931258 (CSSI), #1942995 (CAREER)
- Stampede2 resources at TACC via XSEDE



http://lpna.cs.illinois.edu