

Strassen-like algorithms for symmetric tensor contractions

Edgar Solomonik

Theory Seminar
University of Illinois at Urbana-Champaign

September 18, 2017

- 1 Introduction
- 2 Applications of tensor symmetry
- 3 Exploiting symmetry in matrix products
- 4 Exploiting symmetry in tensor contractions
- 5 Numerical error analysis
- 6 Communication cost analysis
- 7 Summary and conclusion

Terminology

A tensor $\mathbf{T} \in \mathbb{R}^{n_1 \times \dots \times n_d}$ has

- order d (i.e. d modes / indices)
- dimensions n_1 -by- \dots -by- n_d (in this talk, usually each $n_i = n$)
- elements $\mathbf{T}_{i_1 \dots i_d} = T_{\mathbf{i}}$ where $\mathbf{i} \in \{1, \dots, n\}^d$

We say a tensor is **symmetric** if for any $j, k \in \{1, \dots, n\}$

$$\mathbf{T}_{i_1 \dots i_j \dots i_k \dots i_d} = \mathbf{T}_{i_1 \dots i_k \dots i_j \dots i_d}$$

A tensor is **antisymmetric** (skew-symmetric) if for any $j, k \in \{1, \dots, n\}$

$$\mathbf{T}_{i_1 \dots i_j \dots i_k \dots i_d} = (-1) \mathbf{T}_{i_1 \dots i_k \dots i_j \dots i_d}$$

A tensor is **partially-symmetric** if such index interchanges are restricted to be within subsets of $\{1, \dots, n\}$, e.g.

$$\mathbf{T}_{kl}^{ij} = \mathbf{T}_{kl}^{ji} = \mathbf{T}_{lk}^{ji} = \mathbf{T}_{lk}^{ij}$$

Tensor contractions

We work with contractions of tensors

- \mathbf{A} of order $s + v$, and
- \mathbf{B} of order $v + t$ into
- \mathbf{C} of order $s + t$, defined as

$$C_{ij} = \sum_{k \in \{1, \dots, n\}^v} A_{ik} B_{kj}$$

- requires $O(n^{s+t+v})$ multiplications and additions, $\omega := s + t + v$
- assumes an index ordering, but does not lose generality
- works with any symmetries of \mathbf{A} and \mathbf{B}
- is extensible to symmetries of \mathbf{C} via [symmetrization](#) (sum all permutations of modes in \mathbf{C} , denoted $[\mathbf{C}]_{ij}$)
- generalizes simple matrix operations, e.g.

$$\underbrace{(s, t, v) = (1, 0, 1)}_{\text{matrix-vector product}} \quad \underbrace{(s, t, v) = (1, 1, 0)}_{\text{vector outer product}} \quad \underbrace{(s, t, v) = (1, 1, 1)}_{\text{matrix-matrix product}}$$

Applications of symmetric tensor contractions

Symmetric and Hermitian matrix operations are part of the BLAS

- matrix-vector products: `symv` (`symm`), `hemv`, (`hemm`)
- symmetrized outer product: `syr2` (`syr2k`), `her2`, (`her2k`)
- these operations dominate symmetric/Hermitian diagonalization

Hankel matrices are order $2 \log_2(n)$ partially-symmetric tensors

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_{11} & \mathbf{H}_{21}^T \\ \mathbf{H}_{21} & \mathbf{H}_{22} \end{bmatrix}$$

where \mathbf{H}_{11} , \mathbf{H}_{21} , \mathbf{H}_{22} are also Hankel.

In general, **partially-symmetric** tensors are **nested symmetric** tensors

- a nonsymmetric matrix is a vector of vectors
- $\mathbf{T}_{kl}^{ij} = \mathbf{T}_{kl}^{ji} = \mathbf{T}_{lk}^{ji} = \mathbf{T}_{lk}^{ij}$ is a symmetric matrix of symmetric matrices

Applications of partially-symmetric tensor contractions

High-accuracy methods in computational quantum chemistry

- solve the multi-electron Schrödinger equation $\mathbf{H}|\Psi\rangle = E|\Psi\rangle$, where \mathbf{H} is a linear operator, but Ψ is a function of *all* electrons
- use wavefunction ansatz like $\Psi \approx \Psi^{(k)} = e^{\mathbf{T}^{(k)}} |\Psi^{(k-1)}\rangle$ where $\Psi^{(0)}$ is a mean-field (averaged) function and $\mathbf{T}^{(k)}$ is an order $2k$ tensor, acting as a multilinear excitation operator on the electrons
- **coupled-cluster** methods use the above ansatz for $k \in \{2, 3, 4\}$ (CCSD, CCSDT, CCSDTQ)
- solve iteratively for $\mathbf{T}^{(k)}$, where each iteration has cost $O(n^{2k+2})$, dominated by contractions of partially antisymmetric tensors
- for example, a dominant contraction in CCSD ($k = 2$) is

$$\mathbf{z}_{i\bar{c}}^{a\bar{k}} = \sum_{b=1}^n \sum_{j=1}^n \mathbf{T}_{ij}^{ab} \cdot \mathbf{v}_{b\bar{c}}^{j\bar{k}}$$

where $\mathbf{T}_{ij}^{ab} = -\mathbf{T}_{ij}^{ba} = \mathbf{T}_{ji}^{ba} = -\mathbf{T}_{ji}^{ab}$. We'll show an algorithm that requires n^6 rather than $2n^6$ operations.

Fast algorithms

Strassen's algorithm

$$\begin{bmatrix} \mathbf{C}_{11} & \mathbf{C}_{12} \\ \mathbf{C}_{21} & \mathbf{C}_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \cdot \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{bmatrix}$$

$$\mathbf{M}_1 = (\mathbf{A}_{11} + \mathbf{A}_{22}) \cdot (\mathbf{B}_{11} + \mathbf{B}_{22})$$

$$\mathbf{C}_{11} = \mathbf{M}_1 + \mathbf{M}_4 - \mathbf{M}_5 + \mathbf{M}_7$$

$$\mathbf{M}_2 = (\mathbf{A}_{21} + \mathbf{A}_{22}) \cdot \mathbf{B}_{11}$$

$$\mathbf{C}_{21} = \mathbf{M}_2 + \mathbf{M}_4$$

$$\mathbf{M}_3 = \mathbf{A}_{11} \cdot (\mathbf{B}_{12} - \mathbf{B}_{22})$$

$$\mathbf{C}_{12} = \mathbf{M}_3 + \mathbf{M}_5$$

$$\mathbf{M}_4 = \mathbf{A}_{22} \cdot (\mathbf{B}_{21} - \mathbf{B}_{11})$$

$$\mathbf{C}_{22} = \mathbf{M}_1 - \mathbf{M}_2 + \mathbf{M}_3 + \mathbf{M}_6$$

$$\mathbf{M}_5 = (\mathbf{A}_{11} + \mathbf{A}_{12}) \cdot \mathbf{B}_{22}$$

$$\mathbf{M}_6 = (\mathbf{A}_{21} - \mathbf{A}_{11}) \cdot (\mathbf{B}_{11} + \mathbf{B}_{12})$$

$$\mathbf{M}_7 = (\mathbf{A}_{12} - \mathbf{A}_{22}) \cdot (\mathbf{B}_{21} + \mathbf{B}_{22})$$

By minimizing number of products, minimize number of recursive calls

$$T(n) = 7T(n/2) + O(n^2) = O(7^{\log_2 n}) = O(n^{\log_2 7})$$

For convolution, DFT matrix reduces from naive $O(n^2)$ products to $O(n)$, both of these are [bilinear algorithms](#)

How can we formally define an algorithm?

Formally defining a space of algorithms enables systematic exploration.

Definition (Bilinear algorithms (V. Pan, 1984))

A bilinear algorithm $\Lambda = (\mathbf{F}^{(A)}, \mathbf{F}^{(B)}, \mathbf{F}^{(C)})$ computes

$$\mathbf{c} = \mathbf{F}^{(C)}[(\mathbf{F}^{(A)\top} \mathbf{a}) \circ (\mathbf{F}^{(B)\top} \mathbf{b})],$$

where \mathbf{a} and \mathbf{b} are inputs and \circ is the Hadamard (pointwise) product.

$$\begin{bmatrix} \mathbf{c} \end{bmatrix} = \begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix} \left[\left(\begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix}^\top \begin{bmatrix} \mathbf{a} \end{bmatrix} \right) \circ \left(\begin{bmatrix} \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times \end{bmatrix}^\top \begin{bmatrix} \mathbf{b} \end{bmatrix} \right) \right]$$

Bilinear algorithms as tensor factorizations

A bilinear algorithm corresponds to a CP tensor decomposition

$$\begin{aligned}c_i &= \sum_{r=1}^R \mathbf{F}_{ir}^{(C)} \left(\sum_j \mathbf{F}_{jr}^{(A)} \mathbf{a}_j \right) \left(\sum_k \mathbf{F}_{kr}^{(B)} \mathbf{b}_k \right) \\ &= \sum_j \sum_k \left(\sum_{r=1}^R \mathbf{F}_{ir}^{(C)} \mathbf{F}_{jr}^{(A)} \mathbf{F}_{kr}^{(B)} \right) \mathbf{a}_j \mathbf{b}_k \\ &= \sum_j \sum_k \mathbf{T}_{ijk} \mathbf{a}_j \mathbf{b}_k \quad \text{where} \quad \mathbf{T}_{ijk} = \sum_{r=1}^R \mathbf{F}_{ir}^{(C)} \mathbf{F}_{jr}^{(A)} \mathbf{F}_{kr}^{(B)}\end{aligned}$$

For multiplication of $n \times n$ matrices,

- \mathbf{T} is $n^2 \times n^2 \times n^2$
- classical algorithm has rank $R = n^3$
- Strassen's algorithm has rank $R \approx n^{\log_2(7)}$

Symmetric matrix times vector

Lets consider the simplest tensor contraction with symmetry

- let \mathbf{A} be an n -by- n symmetric matrix ($\mathbf{A}_{ij} = \mathbf{A}_{ji}$)
- the symmetry is not preserved in matrix-vector multiplication

$$\mathbf{c} = \mathbf{A} \cdot \mathbf{b}$$

$$c_i = \sum_{j=1}^n \underbrace{\mathbf{A}_{ij} \cdot \mathbf{b}_j}_{\text{nonsymmetric}}$$

- generally n^2 additions and n^2 multiplications are performed
- we can perform only $\binom{n+1}{2}$ multiplications using

$$c_i = \sum_{j=1, j \neq i}^n \underbrace{\mathbf{A}_{ij} \cdot (\mathbf{b}_i + \mathbf{b}_j)}_{\text{symmetric}} + \underbrace{\left(\mathbf{A}_{ii} - \sum_{j=1, j \neq i}^n \mathbf{A}_{ij} \right)}_{\text{low-order}} \cdot \mathbf{b}_i$$

Symmetrized outer product

Consider a rank-2 outer product of vectors \mathbf{a} and \mathbf{b} of length n into symmetric matrix \mathbf{C}

$$\mathbf{C} = \mathbf{a} \cdot \mathbf{b}^T + \mathbf{b} \cdot \mathbf{a}^T$$
$$\mathbf{C}_{ij} = \underbrace{\mathbf{a}_i \cdot \mathbf{b}_j}_{\text{nonsymmetric}} + \underbrace{\mathbf{a}_j \cdot \mathbf{b}_i}_{\text{permutation}}$$

usually computed via the n^2 multiplications and n^2 additions
new algorithm requires $\binom{n+1}{2}$ multiplications

$$\mathbf{C}_{ij} = \underbrace{(\mathbf{a}_i + \mathbf{a}_j) \cdot (\mathbf{b}_i + \mathbf{b}_j)}_{\mathbf{z}_{ij}} - \underbrace{\mathbf{a}_i \cdot \mathbf{b}_i}_{\mathbf{w}_i} - \underbrace{\mathbf{a}_j \cdot \mathbf{b}_j}_{\mathbf{w}_j}$$

symmetric low-order

Symmetrized matrix multiplication

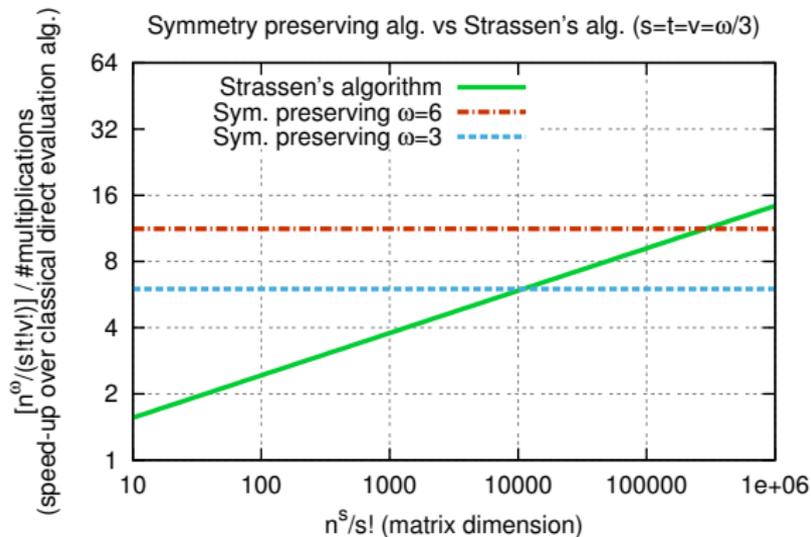
For symmetric matrices \mathbf{A} and \mathbf{B} , compute

$$\mathbf{C}_{ij} = \sum_{k=1}^n \left(\underbrace{\mathbf{A}_{ik} \cdot \mathbf{B}_{kj}}_{\text{nonsymmetric}} + \underbrace{\mathbf{A}_{jk} \cdot \mathbf{B}_{ki}}_{\text{permutation}} \right)$$

New algorithm requires $\binom{n}{3} + O(n^2)$ multiplications rather than n^3

$$\begin{aligned} \mathbf{C}_{ij} &= \sum_k \underbrace{(\mathbf{A}_{ij} + \mathbf{A}_{ik} + \mathbf{A}_{jk}) \cdot (\mathbf{B}_{ij} + \mathbf{B}_{ik} + \mathbf{B}_{jk})}_{\mathbf{Z}_{ijk} - \text{symmetric}} \\ &\quad - \underbrace{\mathbf{A}_{ij} \cdot \left(\sum_k \mathbf{B}_{ij} + \mathbf{B}_{ik} + \mathbf{B}_{jk} \right)}_{\mathbf{U}_{ij} - \text{low-order}} - \underbrace{\mathbf{B}_{ij} \cdot \left(\sum_k \mathbf{A}_{ij} + \mathbf{A}_{ik} + \mathbf{A}_{jk} \right)}_{\mathbf{V}_{ij} - \text{low-order}} \\ &\quad - \underbrace{\sum_k \mathbf{A}_{ik} \cdot \mathbf{B}_{ik}}_{\mathbf{w}_i - \text{low-order}} - \underbrace{\sum_k \mathbf{A}_{jk} \cdot \mathbf{B}_{jk}}_{\mathbf{w}_j - \text{low-order}} \end{aligned}$$

Comparison to Strassen's algorithm



Fully symmetric tensor contractions

For general symmetric tensor contraction algorithms,

- \mathbf{A} of order $s + v$, and
- \mathbf{B} of order $v + t$ into
- \mathbf{C} of order $s + t$, defined as

we define the (nonsymmetrized) contraction as $\mathbf{C} = \mathbf{A} \odot_v \mathbf{B}$ where

$$C_{ij} = \sum_{k \in \{1, \dots, n\}^v} A_{ik} B_{kj}$$

then define the symmetrized tensor contraction as

$$\mathbf{C}_i = [\mathbf{A} \odot_v \mathbf{B}]_i$$

The usual method first computes $\mathbf{A} \odot_v \mathbf{B}$ with

$$\binom{n}{s} \binom{n}{t} \binom{n}{v} \approx \frac{n^{s+t+v}}{s!t!v!}$$

multiplications and additions

Fast symmetrized product of symmetric matrices

Using symmetrization notation for $s = t = v = 1$, we have fast algorithm:

$$\begin{aligned} C_{ij} &= [AB]_{ij} = \underbrace{\sum_k (A_{ij} + A_{ik} + A_{jk}) \cdot (B_{ij} + B_{ik} + B_{jk})}_{\sum_k [A]_{ijk} \cdot [B]_{ijk}} \\ &\quad - \underbrace{A_{ij} \cdot \left(\sum_k B_{ij} + B_{ik} + B_{jk} \right)}_{A_{ij} \cdot \sum_k [B]_{ijk}} - \underbrace{B_{ij} \cdot \left(\sum_k A_{ij} + A_{ik} + A_{jk} \right)}_{B_{ij} \cdot \sum_k [A]_{ijk}} \\ &\quad - \underbrace{\sum_k A_{ik} \cdot B_{ik} - \sum_k A_{jk} \cdot B_{jk}}_{[A \circ_1 B]_{ij}} \\ &= \sum_k [A]_{ijk} \cdot [B]_{ijk} - A_{ij} \sum_k [B]_{ijk} - B_{ij} \sum_k [A]_{ijk} - [A \circ_1 B]_{ij} \end{aligned}$$

where $A \circ_1 B = \sum_k A_{ik} \cdot B_{jk}$

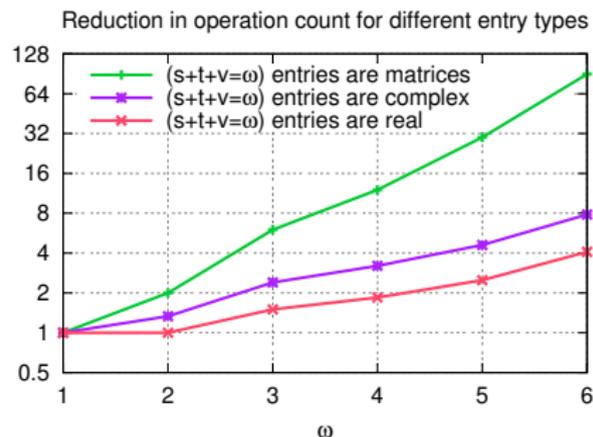
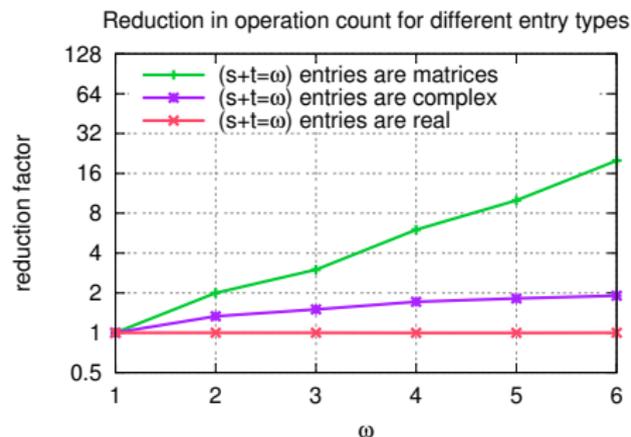
Fast fully-symmetric contraction algorithm

The fast algorithm is defined as follows (using $\omega = s + t + v$)

$$\begin{aligned} \mathbf{C}_i = & \underbrace{\sum_{k \in \{1, \dots, n\}^v} [\mathbf{A}]_{ik} \cdot [\mathbf{B}]_{ik}}_{\text{symmetric, requires } \binom{n+\omega-1}{\omega} \text{ multiplications}} \\ & - \underbrace{\sum_{p+q=v} \sum_{k \in \{1, \dots, n\}^{v-p-q}} \left(\sum_{p \in \{1, \dots, n\}^p} [\mathbf{A}]_{ikp} \right) \cdot \left(\sum_{q \in \{1, \dots, n\}^q} [\mathbf{B}]_{ikq} \right)}_{\text{requires } O(n^{\omega-1}) \text{ multiplications}} \\ & - \underbrace{\sum_{r=1}^{\min(s,t)} [\mathbf{A} \odot_{v+r} \mathbf{B}]_i}_{\text{requires } O(n^{\omega-1}) \text{ multiplications}} \end{aligned}$$

Overall need $\binom{n}{\omega} + O(n^{\omega-1}) = \frac{n^{s+t+v}}{(s+t+v)!} + O(n^{s+t+v-1})$ multiplications, i.e. $(s+t+v)!/(s!t!v!)$ factor better

Reduction in operation count of fast algorithm with respect to standard



(s, t, v) values for left and right graph tabulated below

ω	1	2	3	4	4	6
Left graph	(1, 0, 0)	(1, 1, 0)	(2, 1, 0)	(2, 2, 0)	(3, 2, 0)	(3, 3, 0)
Right graph	(1, 0, 0)	(1, 1, 0)	(1, 1, 1)	(2, 1, 1)	(2, 2, 1)	(2, 2, 2)

Nesting the fast algorithm

For partially-(anti)symmetric contractions we can

- nest the new algorithm over each group of symmetric modes
- reduction in mults can translate to reduction in the number of operations
- for Hankel matrices, yields $O(n^{1.585})$ algorithm, which is better than naive ($O(n^2)$) but worse than DFT ($O(n)$)
- for coupled-cluster contractions, significant reductions in cost (number of operations) can be achieved
 - CCSD 1.3X on a typical system
 - CCSDT 2.1X on a typical system
 - CCSDTQ 5.7X on a typical system

Theoretical error bounds

We express error bounds in terms of $\gamma_n = \frac{n\epsilon}{1-n\epsilon}$, where ϵ is the machine precision.

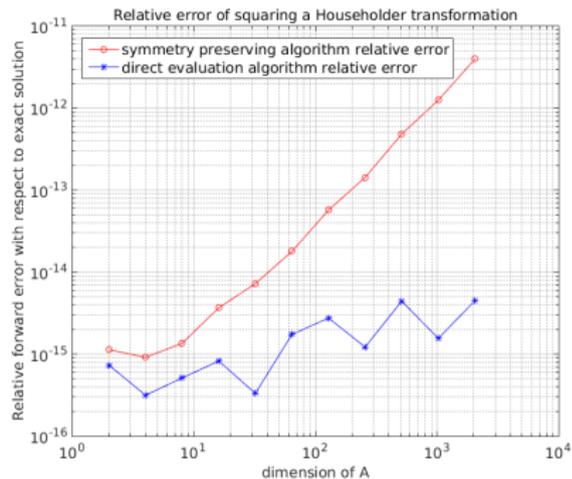
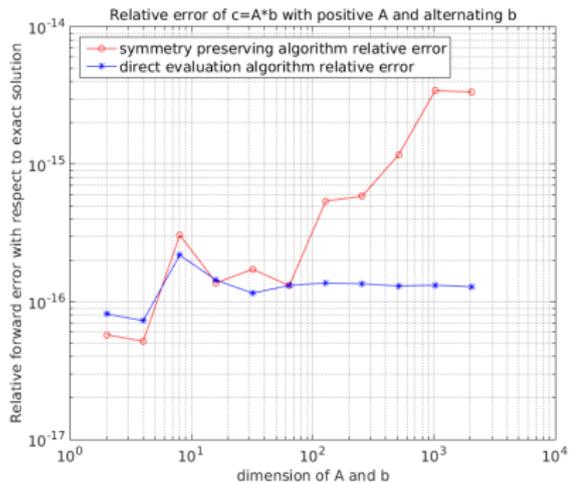
Let Ψ be the standard algorithm and Φ be the fast algorithm. The error bound for the standard algorithm arises from matrix multiplication

$$\|fl(\Psi(\mathbf{A}, \mathbf{B})) - \mathbf{C}\|_\infty \leq \gamma_m \cdot \|\mathbf{A}\|_\infty \cdot \|\mathbf{B}\|_\infty \quad \text{where } m = \binom{n}{v} \binom{\omega}{v}.$$

The following error bound holds for the fast algorithm

$$\|fl(\Phi(\mathbf{A}, \mathbf{B})) - \mathbf{C}\|_\infty \leq \gamma_m \cdot \|\mathbf{A}\|_\infty \cdot \|\mathbf{B}\|_\infty \quad \text{where } m = 3 \binom{n}{v} \binom{\omega}{t} \binom{\omega}{s}.$$

Stability of symmetry preserving algorithms



Communication complexity

Parallelization is crucial for BLAS-like operations and scientific-applications

- can assess **parallel scalability** by considering **communication complexity**
- various notions of communication *complexity* or *cost* exist
- for the problems studied (which have high degree of concurrency), a simple measure is most important

W – maximum number of words sent or received by any processor

- can derive lower and upper bounds for W for a given bilinear algorithm
- generally assume that tensor data is evenly distributed initially

Expansion in bilinear algorithms

The **communication complexity** of a bilinear algorithm depends on the amount of data needed to compute subsets of the bilinear products.

Definition (Bilinear subalgorithm)

Given $\Lambda = (\mathbf{F}^{(A)}, \mathbf{F}^{(B)}, \mathbf{F}^{(C)})$, $\Lambda_{\text{sub}} \subseteq \Lambda$ if \exists projection matrix \mathbf{P} , so

$$\Lambda_{\text{sub}} = (\mathbf{F}^{(A)}\mathbf{P}, \mathbf{F}^{(B)}\mathbf{P}, \mathbf{F}^{(C)}\mathbf{P}).$$

The projection matrix extracts $\#\text{cols}(\mathbf{P})$ columns of each matrix.

Definition (Bilinear algorithm expansion)

A bilinear algorithm Λ has expansion bound $\mathcal{E}_{\Lambda} : \mathbb{N}^3 \rightarrow \mathbb{N}$, if for all

$$\Lambda_{\text{sub}} := (\mathbf{F}_{\text{sub}}^{(A)}, \mathbf{F}_{\text{sub}}^{(B)}, \mathbf{F}_{\text{sub}}^{(C)}) \subseteq \Lambda$$

we have $\text{rank}(\Lambda_{\text{sub}}) \leq \mathcal{E}_{\Lambda}(\text{rank}(\mathbf{F}_{\text{sub}}^{(A)}), \text{rank}(\mathbf{F}_{\text{sub}}^{(B)}), \text{rank}(\mathbf{F}_{\text{sub}}^{(C)}))$

For matrix mult., Loomis-Whitney inequality $\rightarrow \mathcal{E}_{\text{MM}}(x, y, z) = \sqrt{xyz}$

Communication cost of the standard algorithm

We consider communication bandwidth cost on a sequential machine with cache size M .

The intermediate formed by the standard algorithm may be computed via matrix multiplication with communication cost,

$$W(n, s, t, v, M) = \Theta \left(\frac{\binom{n}{s} \binom{n}{t} \binom{n}{v}}{\sqrt{M}} + \binom{n}{s+v} + \binom{n}{t+v} + \binom{n}{s+t} \right).$$

The cost of symmetrizing the resulting intermediate is low-order or the same.

Communication cost of the fast algorithm

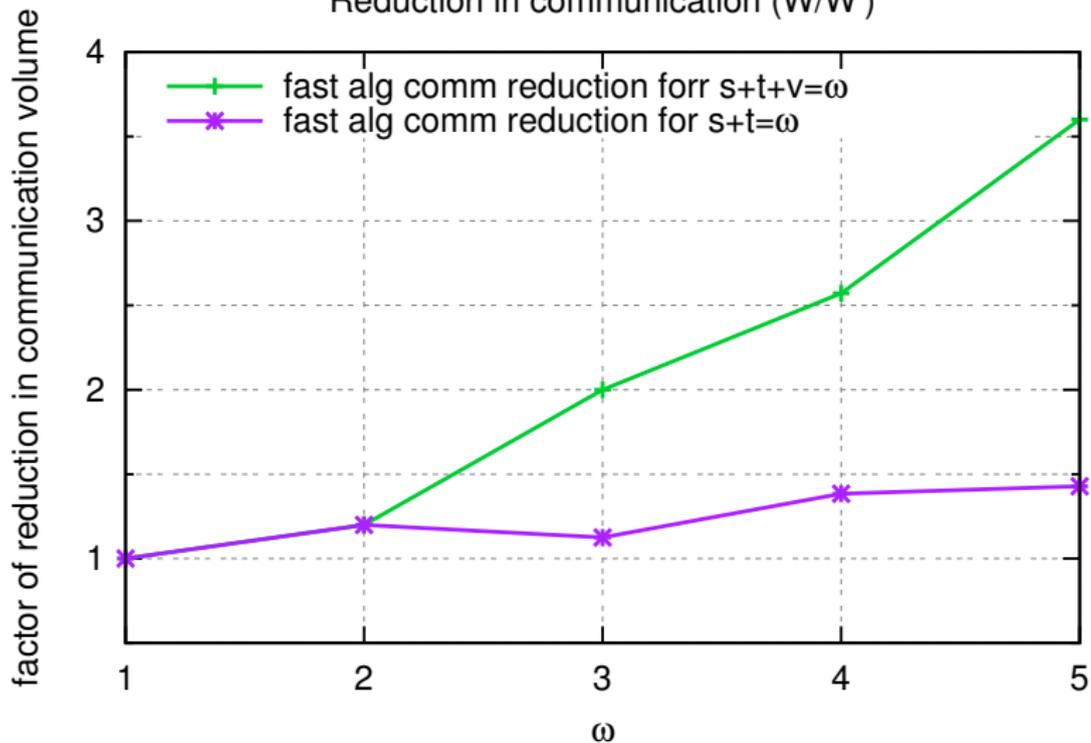
We can lower bound the cost of the fast algorithm using the Hölder-Brascamp-Lieb inequality.

An algorithm that blocks \mathbf{Z} symmetrically nearly attains the cost

$$W'(n, s, t, v, M) = O\left(\frac{\binom{n}{\omega}}{M^{\omega/(\omega - \min(s, t, v))}} \cdot \left[\binom{\omega}{t} + \binom{\omega}{s} + \binom{\omega}{v} \right] + \binom{n}{s+v} + \binom{n}{t+v} + \binom{n}{s+t}\right).$$

which is not far from the lower bound and attains it when $s = t = v$.

Reduction in communication (W/W')



Communication cost of symmetry preserving algorithms

For contraction of order $s + v$ tensor with order $v + t$ tensor

- symmetry preserving algorithm requires $\frac{(s+v+t)!}{s!v!t!}$ fewer multiplies
- matrix-vector-like algorithms ($\min(s, v, t) = 0$)
 - vertical communication dominated by largest tensor
 - horizontal communication asymptotically greater if only unique elements are stored and $s \neq v \neq t$
- matrix-matrix-like algorithms ($\min(s, v, t) > 0$)
 - vertical and horizontal communication costs asymptotically greater for symmetry preserving algorithm when $s \neq v \neq t$

Summary of results

The following table lists the leading order number of multiplications F required by the standard algorithm and F' by the fast algorithm for various cases of symmetric tensor contractions

ω	s	t	v	F	F'	applications
2	1	1	0	n^2	$n^2/2$	syr2, syr2k, her2, her2k
2	1	0	1	n^2	$n^2/2$	symv, symm, hemv, hemm
3	1	1	1	n^3	$n^3/6$	symmetrized matmul
$s+t+v$	s	t	v	$\binom{n}{s} \binom{n}{t} \binom{n}{v}$	$\binom{n}{\omega}$	any symmetric tensor contraction

High-level conclusions:

- Algebraic complexity result for leveraging symmetry in contractions
- Applications for basic complex arithmetic and partially symmetric contractions
- Caveats: more communication per flop, slightly higher numerical error

Collaborators on various parts:

- James Demmel
- Torsten Hoefler
- Devin Matthews

S., Demmel; Technical Report, ETH Zurich, December 2015.

S., Demmel, Hoefler; Technical Report, ETH Zurich, January 2015.

The fast algorithm for computing \mathbf{C} forms the following intermediates with $\binom{n}{\omega}$ multiplications (where $\omega = s + t + v$),

$$Z_i = \left(\sum_{j \in \mathcal{X}(i)} A_j \right) \cdot \left(\sum_{l \in \mathcal{X}(i)} B_l \right)$$

$$V_i = \left(\sum_{j \in \mathcal{X}(i)} A_j \right) \cdot \left(\sum_{k_1} \sum_{l \in \mathcal{X}(i \cup k)} B_l \right) \\ + \left(\sum_{k_1} \sum_{j \in \mathcal{X}(i \cup k)} A_j \right) \cdot \left(\sum_{l \in \mathcal{X}(i)} B_l \right)$$

$$W_i = \left(\sum_{j \in \mathcal{X}(i)} A_j \right) \cdot \left(\sum_{l \in \mathcal{X}(i)} B_l \right)$$

$$C_i = \sum_k Z_{i \cup k} - \sum_k V_{i \cup k} \\ - \sum_{j \in \mathcal{X}(i)} \left(\sum_k W_{j \cup k} \right)$$