## MS76: Communication-Avoiding Algorithms - Part I of II

4:00-4:25 Communication-Avoiding Algorithms:
Challenges and New Results
Erin C. Carson
4:30-4:55 Communication-Optimal Loop Nests
Nicholas Knight
5:00-5:25 Communication Lower Bounds for MatricizedTensor Times Khatri-Rao Product
Grey Ballard, Nicholas Knight, Kathryn Rouse
5:30-5:55 Matrix Multiplication, a Little Faster
Oded Schwartz and Elaye E. Karstadt

## MS93: Communication-Avoiding Algorithms - Part

 II of IIFriday, July 14, 4:00 PM - 6:00 PM, Spirit of Pittsburgh B
4:00-4:25 Communication-Avoiding Primal and Dual Methods for Regularized Least-Squares
Aditya Devarakonda, Kimon Fountoulakis, James Demmel, and Michael Mahoney

4:30-4:55 Communication-Avoiding Sparse Inverse Covariance Matrix Estimation
Penporn Koanantakool, Sang-Yun Oh, Dmitriy Morozov, Aydin Buluc, and Leonid Oliker, Katherine Yelick

5:00-5:25 Black-box Communication Optimal Low Rank Approximations
Alan Ayala, Laura Grigori
5:30-5:55 Performance of S-step and Pipelined Krylov Methods Ichitaro Yamazaki, Mark Hoemmen, Jack J. Dongarra, Piotr Luszczek

## GOIIIIUIICEIIOIRAVOIOIII Algorithms: Challenges and New Results

Erin Carson<br>New York University

SIAM Annual Meeting 2017
Pittsburgh, PA, USA

## What is communication?

- Runtime of an algorithm is the sum of
- \#flops x (time/flop) $\longrightarrow$ computation
- \#words moved x (1/bandwidth)
- \#messages x latency
- communicati on

Sequential


Parallel


- Communication is expensive, computation is cheap
- in terms of both time and energy! (time $\rightarrow$ joules)


## Future_exascale_svstems

|  | Petascale Systems (2009) |
| :---: | :---: |
| System Peak | $2 \cdot 10 \uparrow 15$ flops/s |
| Node Memory Bandwidth | 25 GB/s |
| Total Node Interconnect Bandwidth | 3.5 GB/s |
| Memory Latency | 100 ns |
| Interconnect Latency | $1 \mu \mathrm{~s}$ |
| *Sources: from P. Beckman (ANL), J. Shalf (LBL), and D. Unat (LBL) |  |
| - Gaps between communication/computation cost only growing larger in future systems |  |
| - Avoiding communication will be essential for applications at exascale! |  |

## Motivation for CA Algorithms

- This problem will not be solved (in the foreseeable future) in hardware
- Needs to be addressed higher in the computing stack - at the algorithm level
- A paradigm shift in the way the numerical algorithms are designed is required
- We can't only think about computational complexity, we must also think about communication complexity
- Communication-avoiding algorithms:
- Minimize communication volume (total words moved)
- Minimize number of messages
- Minimize over multiple levels of memory/parallelism
- Allow redundant computations


## Work in CA algorithms

- For numerical linear algebra computations...

1. Prove lower bounds on communication cost
2. Design new (stable) algorithms that attain these lower bounds
3. Implement algorithms on various machines, for various applications

## Brief History: Dense Linear Algebra

- Communication long recognized as a bottleneck
- BLAS1 $\rightarrow$ BLAS2 $\rightarrow$ BLAS3
- EISPACK, LINPACK $\rightarrow$ LAPACK, ScaLAPACK
- Lower bounds for (dense) matrix multiply
- Hong and Kung (1981): for sequential matmul, must move $\Omega$ ( $n \uparrow 3 /$ $M \uparrow 1 / 2$ ) words ( $M$ is size of fast memory)
- Irony, Tiskin, Toledo (2004): Generalized to parallel case: when each of $P$ processors stores $M=O(n \uparrow 2 / P)$ words, $\Omega(n \uparrow 3 / P \uparrow 1 / 2)$ words moved.
- Attainable by Cannon's Algorithm (Cannon 1969), SUMMA (van de Geijn and Watts 1997)
- Demmel, Grigori, Hoemmen, Langou (2008): same lower bounds apply to LU, QR factorization


## Generalization: 3 nested loops

- Generalized by Ballard, Demmel, Holtz, Schwartz (2011) to any "3-nested loops" algorithm, using Loomis Whitney inequality
- matmul, Cholesky, LU, LDL , QR, FloydWarshall, etc.; for dense and sparse matrices; sequential, parallel, hybrid
- words moved $=\Omega$ (\#flops $/ M \uparrow 1 / 2$ ),


$$
\text { messages }=\Omega(\# \text { flops } / M \uparrow 3 / 2)
$$

- (assuming $M=O(n \uparrow 2 / P)$ )

| Algorithm | Minimizes \# <br> words moved | Minimizes \# words moved and \# messages |
| :---: | :---: | :---: |
| Cholesky | ScaLAPACK | ScaLAPACK |
| LU | ScaLAPACK | CALU |
| QR | ScaLAPACK | (Grigori, Demmel, Xiang, 2008), (Khabou, Demmel, Grigori, Gu, 2012) |
| (Demmel, Grigori, Hoemmen, Langou, 2008), (Ballard et al., 2014) |  |  |

- 3D matmul: $M=O(n \uparrow 2 / P \uparrow 2 / 3), 2.5 D$ matmul: $M=c n \uparrow 2 / P$


## Tall-Skinny QR

- TSQR: QR factorization of a tall skinny matrix using Householder transformations
- QR decomposition of $m \times b$ matrix $W$, m >> b
- P processors, block row layout
- Classic Parallel Algorithm
- Compute Householder vector for each column
- Number of messages $\propto b \log P$
- Communication Avoiding Algorithm
- Reduction operation, with QR as operator
- Number of messages $\propto \log P$

| Parallel |  |
| :---: | :---: |
| $W=[\square W \downarrow 0 \rightarrow \square$ |  |
|  |  |
| $\begin{aligned} & W \downarrow 2 @ \\ & W \downarrow 3] \end{aligned}$ | -@ $\rightarrow 01 \sim R \downarrow$ |
|  | $\begin{array}{cc} \\ >R \downarrow 10 & \\ 01 & R \downarrow \\ & R \downarrow \\ 02\end{array}$ |
|  | @ 111 |
|  | sedutential |
| $W=[\square W \downarrow 0 \rightarrow 8 \downarrow 30$ |  |
|  |  |
| $W \downarrow 2$ @ $00 \quad R \downarrow \stackrel{1}{3}$ |  |
| $W \downarrow 3]$ | $01 R \downarrow \geqslant$ |
|  | $02 R \downarrow$ |
|  | Dual Core 03 |
|  |  |
| @WL1 @ |  |
|  |  |
| $W \downarrow 3$ ] | $R \downarrow 01$ Rı $>_{02}^{R \downarrow}$ |
|  | 11 R】 03 |

## CAQR

- Demmel, Grigori, Hoemmen, Langou (2008)
- TSQR on panels
- Performing trailing matrix update is complicated - $Q$ factor represented implicitly
- Can reconstruct Householder vectors from Tall-Skinny QR (TSQR-HR) (Ballard, Demmel, Grigori, Jacquelin, Knight, Nyugen, 2015)
- Close to optimal
- Assume $M=O(n \uparrow 2 / P)$
- Choose $b$ near $n / P \uparrow 1 / 2$
- Words moved: $O(n \uparrow 2 \log P / P \uparrow 1 / 2)$
- Lower bound: $\Omega(n \uparrow 2 / P \uparrow 1 / 2)$
- Messages: $O(n \log P / b)$
- Lower bound: $\Omega(P \uparrow 1 / 2)$
- As stable as Householder QR


Recent work (Ayala, Claeys, Grigori)

- Talk tomorrow in Part II (MS93)
- Comm.-optimal low-rank approximation of matrices arising from boundary element problems
- Uses CAQR factorization

TSQR/CAQR implemented in

- Intel MKL library
- GNU Scientific Library
- ScaLAPACK
- Spark for data mining


## Tall-Skinny LU

- Similar idea as in TSQR; uses tournament pivoting
- Step 1: preprocessing: select a set of $b$ good pivot rows
- Reduction operation with GEPP being the operator used to select pivot rows at each node of the reduction tree
- Like in TSQR, shape of reduction tree depends on underlying architecture
- Step 2: permute the b pivot rows into the first b positions in the panel; perform LU factorization with no pivoting
- Stability caveat: upper bound on the growth factor is worse than for GEPP
- For GEPP, growth factor bounded by $2 \uparrow n-1$; for CALU, $2 \uparrow n \ell$, where $\ell$ is the depth of the reduction tree


## CALU

- (Grigori, Demmel, Xiang, 2008), (Khabou, Demmel, Grigori, Gu, 2012)
- Assume 2D grid of P processors, using a 2D block cyclic layout with square blocks of size b

For $\mathrm{ib}=1$ to $\mathrm{n}-1$ step b

1. Find permutation for current panel using TSLU
2. Apply all row permutations (pdlaswp)

- broadcast pivot information along the rows of the grid

3. Compute panel factorization (dtrsm)
4. Compute block row of $U$ (pdtrsm)

- broadcast right diagonal part of $L$ of current panel

5. Update trailing matrix (pdgemm)

- broadcast right block column of $L$
- broadcast down block row of U
$M=O(n \uparrow 2 / P)$, choose $b$ near $n / P \uparrow 1 / 2$
- Words: $O(n \uparrow 2 \log P / P \uparrow 1 / 2$ ) ; Lower bound: $\Omega(n \uparrow 2 /$ P\$1/2 )
- Messages: $O(n \log P / b)$; Lower bound: $\Omega(P \uparrow 1 / 2)$

CALU implemented in

- Cray’s libsci
- To be implemented in LAPACK/ScaLAPACK


## Sparse Matrix Computations

- Sparse Matrix $\times$ Vector (SpMV) $(y=A x)$
- Very communication-bound; no reuse
- Lower bound depends on sparsity structure, algorithm used (1D rowwise/ colwise, 2D, etc.)
- Communication cost depends on partition
- Hypergraph models capture communication dependencies (Catalyurek, Aykanat, 1999)

- minimize hypergraph cut = minimize words moved
- Repeated SpMVs ( $Y=[A x, A \uparrow \mathcal{X}, \ldots, A \uparrow k x])$
- Naive approach: k repeated SpMVs
- Communication-avoiding approach: "matrix powers kernel"
- see, e.g., (Demmel, Hoemmen, Mohiyuddin, Yelick, 2008)
- Avoids communication:
- In serial, by exploiting temporal locality in reading $A$, reading vectors
- In parallel, by doing only 1 'expand’ phase (instead of $k$ ).


## Parallel Matrix Powers Kernel

Example: tridiagonal matrix, $s=3, n=40, p=4$



## Sparse Matrix Computations

- Sparse matrix x sparse matrix
- (Ballard, Druinsky, Knight, Schwartz, 2016)
- Fine-grained and coarse-grained hypergraph models for sparse matrix-matrix multiplication
- Identifying a communication-optimal algorithm for particular input matrices is equivalent to solving a hypergraph partitioning problem
- Sparse matrix x dense matrix
- Building block of an increasing number of applications in many areas such as machine learning and graph algorithms
- (Koanantakool, Azad, Buluc, Morozov, Oh, Oliker, Yelick, 2016)
- Communication lower bounds
- New communication-avoiding algorithms based on 1.5D decomposition
- Outperform 2D and 3D variants in both theory and practice


## Krylov subspace methods

- Krylov Subspace Method is a projection process onto the Krylov subspace $\mathcal{K} \downarrow i(A, r \downarrow 0)=\operatorname{span}\{r \downarrow 0, A r \downarrow 0, A \uparrow 2 r \downarrow 0, \ldots, A \uparrow i-1 r \downarrow 0\}$ where $A$ is an $N \times N$ matrix and $r \downarrow 0=b-A x \downarrow 0$ is a length- $N$ vector
- Linear systems $A x=b$, eigenvalue problems, singular value problems, least squares, etc.
- Ex: Lanczos/Conjugate Gradient (CG), Arnoldi/Generalized Minimum Residual (GMRES), Biconjugate Gradient (BICG), BICGSTAB, GKL, LSQR, etc.
- In each iteration,
- Add a dimension to the Krylov subspace
- Forms nested sequence of Krylov subspaces

$$
\begin{aligned}
& \mathcal{K} \downarrow 1(A, r \downarrow 0) \subset \mathcal{K} \downarrow 2(A, r \downarrow 0) \subset \cdots \\
& \subset \mathcal{K} \downarrow i(A, r \downarrow 0)
\end{aligned}
$$

- Orthogonalize (with respect to some $\mathcal{C} \downarrow i$ )

- Select approximate solution $\not \subset \checkmark \vec{i} \in \times$ SpMV $\mathcal{K} \downarrow i(A, r \downarrow 0)$


## Communication-avoiding (s-step) KSMs

- Idea: Compute blocks of $s$ iterations at once
- Compute updates in a different basis (using matrix powers kernel)
- Communicate every $s$ iterations instead of every iteration
- Reduces number of synchronizations per iteration by a factor of $s$
- An idea rediscovered many times...
- First related work: s-dimensional steepest descent, least squares
- Khabaza ('63), Forsythe ('68), Marchuk and Kuznecov ('68)
- Flurry of work on s-step Krylov methods in '80s/early '90s: see, e.g.,

Van Rosendale (1983); Chronopoulos and Gear (1989)

- Another approach: "pipelined" Krylov subspace methods
- Reduce the cost of synchronization by enabling overlapping of AllReduces and other computations

Recent work (Yamazaki, Hoemmen, Dongarra, Luszczek)

- Talk tomorrow in Part II (MS93)
- Performance comparison of s-step and pipelined Krylov subspace methods
- New approach that combines these two techniques


## CA-CG (s-step CG)

$$
r \downarrow 0=b-A x \downarrow 0, p \downarrow 0=r \downarrow 0
$$

for $k=0: \mathrm{nmax}^{\operatorname{m}} S$
Compute $\mathcal{Y} \downarrow k$ and $\mathcal{B} \downarrow k$ such that $A \mathcal{Y}$ $\downarrow k=\mathcal{Y} \downarrow k \mathcal{B} \downarrow k$ and

$$
\operatorname{span}(\mathcal{Y} \downarrow k)=\mathcal{K} \downarrow S+1(A,
$$

$$
p \downarrow s k)+\mathcal{K} \downarrow s(A, r \downarrow s k)
$$

$$
\mathcal{G} \downarrow k=\mathcal{Y} \downarrow k \uparrow T \mathcal{Y} \downarrow k
$$

$$
x \downarrow 0 \uparrow^{\prime}=0, r \downarrow 0 \tau^{\prime}=e \downarrow s+2
$$

$$
p \downarrow 0 \kappa^{\prime}=e \downarrow 1
$$

$$
\text { for } j=1: S
$$

$$
a \downarrow s k+j-1=r \downarrow j-1 \lambda^{\prime} T
$$

$$
\mathcal{G} \downarrow k r \downarrow j-1 \tau^{\prime} / p \downarrow j-1 \tau^{\prime} T \mathcal{G} \downarrow k
$$ $\mathcal{B} \downarrow k p \downarrow j-1 \uparrow^{\prime}$



## CA-CG (s-step CG)

$s$-step CA-CG with monomial basis $(\mathcal{Y}=[p \downarrow i, A p \downarrow i, \ldots, A \uparrow s p \downarrow i, r \downarrow i, A r \downarrow i, \ldots A \uparrow s-1$ $r \downarrow i])$


Changes to how the recurrences are computed can exacerbate finite precision effects of convergence delay and loss of accuracy!
$A$ : bcsstk03 from UFSMC, $b$ : equal components in the eigenbasis of $A$ and $\|b\|=1$

$$
N=112, \kappa(A) \approx 7 \mathrm{e} 6
$$

## Optimizing CA iterative solvers

- CA variants are designed to reduce the time/iteration
- But what we really want to minimize is the runtime, subject to some constraint on accuracy,


## runtime $=($ time/iteration) $\times$ (\# iterations)

- Known that attainable accuracy and convergence rate in CA Krylov subspace methods depends on conditioning of the generated s-step bases
- (C. and Demmel, 2014), (C. and Demmel, 2015), (Philippe and Reichel, 2012)
- Using finite precision analysis, can develop ways to improve numerical behavior while still avoiding communication in each iteration
- e.g., residual replacement (C. and Demmel, 2014), adaptive basis size (C. 2016)
- To reduce communication, can also focus on reducing the number of iterations - Preconditioning
- Challenging to implement in a CA way, but sometimes possible, e.g., ILU (Grigori, Moufawad, 2015), SPAI (Dehnavi, Demmel, Fernández, 2014), DD (Yamazaki, Rajamanickam, Boman, Hoemmen, Heroux, Tomov 2014)


## Extensions: Strassen's Matmul

- Ballard, Demmel, Holtz, Schwartz (2011 SPAA, 2012 JACM, 2014 CACM)

|  | Classical $O(n \nmid 3)$ matmul | Strassen's $O(n \uparrow l g$ 7 ) matmul | Strassen-like $O($ $n \uparrow \omega$ ) matmul |
| :---: | :---: | :---: | :---: |
| words moved: | $\begin{aligned} & \Omega(M / P \cdot(n / \\ & \sqrt{M}) \uparrow() \end{aligned}$ | $\begin{aligned} & \Omega(M / P \cdot(n / \\ & \sqrt{M}) \uparrow \lg 7) \end{aligned}$ | $\begin{aligned} & \Omega(M / P \cdot(n / \\ & \sqrt{M}) \uparrow \omega) \end{aligned}$ |

- New communication-optimal parallelization of Strassen's algorithm
- Outperforms all previous matrix multiplication algorithms

Recent work: (Schwartz and Karstadt)

- 4th talk in this MS
- New faster version of Strassen's alg.
- Communication costs, parallelization, extension to other Strassen-like algorithms

(Strong scaling on a Cray XT4)


## Extensions: Tensor Computations

- For tensor contractions, generalized lower bounds apply
- Cyclops Tensor Framework (Solomonik, Matthews, Hammond, Stanton, Demmel, 2014)
- A massively parallel tensor contraction framework for coupledcluster computations
- Avoids communication by exploiting replication, dynamically selecting a parallel decomposition with the least communication cost


## Recent work (Ballard, Knight, Rouse)

- 3rd talk in this MS!
- Matricized-tensor times Khatri-Rao product (MTTKRP)
- Bottleneck in CANDECOMP/PARAFAC decomposition of a tensor
- Communication lower bounds for sequential and parallel memory models, communication costs of existing algorithms


## Extension: Machine Learning Alaorithms

- CA-SVMs (You, Demmel, Czechowski, Song, Vuduc, 2015)

Recent work (Devarakonda, Fountoulakis, Demmel, Mahoney)

- Talk tomorrow in Part II (MS93)
- Communication-avoiding primal and dual block coordinate descent methods for regularized least-squares
- Extends results on CA-KSMs

Recent work (Koanantakool, Oh, Morozov, Buluc, Oliker, Yelick)

- Talk tomorrow in Part II (MS93)
- Communication-Avoiding Sparse Inverse Covariance Matrix Estimation
- Main bottleneck: iterative sparse-dense matmul
- For the first time, enables analyzing high-dimensional data sets with millions of variables and arbitrary underlying graph structures
- Experimental results using 3D brain fMRI data sets


## Generalization: Arbitrary Nested Loops

- Beyond Linear Algebra: Affine Array References
- e.g., A(i + 2j; 3k + 4)
- (Christ, Demmel, Knight, Scanlon, Yelick, 2013); (Knight, 2015)
- Extend previous lower bounds to larger class of programs
- Lower bounds are computable - proof via Hölder-Brascamp-Lieb (HBL) theory
- Matching upper bounds (i.e., optimal algorithms) in special cases: linear algebra, tensor contraction, direct N -body, database join, etc. (when array references pick a subset of the loop indices)
- Ongoing work addresses attainability in the general case


## Recent work (Knight)

- 2nd talk in this MS!
- Application of HBL theory to Convolutional Neural Networks (CNNs)



## Generalization: Write-Avoiding Algorithms

- (C., Demmel, Grigori, Knight, Koanantakool, Schwartz, Simhadri, 2016)
- Writes can be much more expensive than reads in some current and emerging storage devices such as nonvolatile memories.

(a) Symmetric model for (b) Asymmetric model for (c) Distributed memory model with NVM disks on each node sequential algorithms accounting reads/writes in sequential algorithms for parallel algorithms. Interprocessor communication occurs between second lowest level of the memories of each node.
"communication-avoiding" (CA): minimize the sum of reads and writes "write-avoiding" (WA): CA and also do asymptotically fewer writes than reads
- Some CA algorithms are also WA
- Matrix multiply, Cholesky, TRSM, direct N-body
- For some algorithms, WA algorithm cannot exist (\#writes must be within a constant factor of the total \# of reads and writes; "bounded reuse")
- Strassen matmul, Cooley-Tukey FFT, cache oblivious (CO) algorithms for classical linear algebra (no WACO algs!)


## Lots of speedups...

- Up to 12x faster for 2.5D matmul on 64K core IBM BG/P
- Up to $3 x$ faster for tensor contractions on 2K core Cray XE/6
- Up to 6.2x faster for All-Pairs-Shortest-Path on 24K core Cray CE6
- Up to 2.1x faster for 2.5D LU on 64K core IBM BG/P
- Up to 11.8x faster for direct N -body on 32K core IBM BG/P
- Up to 13x faster for Tall Skinny QR on Tesla C2050 Fermi NVIDIA GPU
- Up to 6.7x faster for symeig(band A) on 10 core Intel Westmere
- Up to $2 x$ faster for 2.5D Strassen on 38K core Cray XT4
- Up to $4.2 x$ faster for MiniGMG benchmark bottom solver, using CABICGSTAB (2.5x for overall solve), 2.5x / 1.5x for combustion simulation code
- Up to 42x for Parallel Direct 3-Body

These and many more recent papers available at bebop.cs.berkeley.edu

# Thank you! 

erinc@cims.nyu.edu math.nyu.edu/~erinc

