# The Past, Present and Future of High Performance <br> Numerical Linear Algebra: <br> Minimizing Communication 

Jim Demmel
EECS \& Math Departments
UC Berkeley

## The Past, the Present ...

- LAPACK - sequential/shared memory parallel dense linear algebra library



## The Past, the Present ...

- LAPACK contributors
- Jack Dongarra, Julien Langou, Julie Langou, Ed Anderson, Zhaojun Bai, David Bailey, David Bindel, Chris Bischof, Susan Blackford, Zvonomir Bujanovic, Karen Braman, Ralph Byers, Inderjit Dhillon, Zlatko Drmac, Peng Du, Jeremy Du Croz, Mark Fahey, Anne Greenbaum, Ming Gu, Fred Gustavson, Deaglan Halligan, Sven Hammarling, Greg Henry, Yozo Hida, Nick Higham, Bo Kagstrom, William Kahan, Daniel Kressner, Ren-Cang Li, Xiaoye Li, Craig Lucas, Osni Marques, Peter Mayes, Alan McKenney, Beresford Parlett, Antoine Petitet, Peter Poromaa, Enrique Quintana-Orti, Gregorio Quintana-Orti, Giuseppe Radicati, Huan Ren, Jason Riedy, Jeff Rutter, Danny Sorensen, Ken Stanley, Xiaobai Sun, Brian Sutton, Francoise Tisseur, Robert van de Geijn, Kresimir Veselic, Christof Voemel, Jerzy Wasniewski + many undergrads


## The Past, the Present ...

- LAPACK - seq/shared mem dense linear algebra
- ScaLAPACK - distributed memory parallel dense linear algebra library
- Jack Dongarra, Susan Blackford, Jaeyoung Choi, Andy Cleary, Ed D'Azevedo, Inderjit Dhillon,
Sven Hammarling, Greg Henry, Antoine Petitet, Ken Stanley, David Walker, Clint Whaley, and many other contributors



## The Past, the Present ...

- LAPACK - seq/shared mem dense linear algebra
- ScaLAPACK - dist mem dense linear algebra
- SuperLU - sparse direct solver for $A x=b$
- Xiaoye Sherry Li



## The Past, the Present

- LAPACK - seq/shared mem dense linear algebra
- ScaLAPACK - dist mem dense linear algebra
- SuperLU - sparse direct solver for Ax=b
- Autotuning
- PhiPAC for matrix multiplication
- Jeff Bilmes, Krste Asanovic, Rich Vuduc, Sriram Iyer, CheeWhye Chin, Dominic Lam

- OSKI for sparse-matrix-vector multiplication (SpMV)
- Rich Vuduc, Kathy Yelick, Rajesh Nishtala, Ben Lee, Shoaib Kamil, Jen Hsu


## The Past, the Present ...

- LAPACK - seq/shared mem dense linear algebra
- ScaLAPACK - dist mem dense linear algebra
- SuperLU - sparse direct solver for $A x=b$
- Autotuning - matmul and SpMV
- Prometheus - parallel unstructured FE solver
- Mark Adams
- Gordon Bell Prize, 2004



## ... The Future

- Why we need to "avoid communication," i.e. avoid moving data
- "Direct" Linear Algebra
- Lower bounds on communication for linear algebra problems like $A x=b$, least squares, $A x=\lambda x$, SVD, etc
- New algorithms that attain these lower bounds
- Not in libraries like Sca/LAPACK (yet!)
- Large speed-ups possible
- "Iterative" Linear Algebra
- Ditto for Krylov Subspace Methods


## Why avoid communication? (1/2)

Algorithms have two costs:

1. Arithmetic (FLOPS)
2. Communication: moving data between

- levels of a memory hierarchy (sequential case)
- processors over a network (parallel case).



## Why avoid communication? (2/2)

- Running time of an algorithm is sum of 3 terms:
- \# flops * time_per_flop
- \# words moved / bandwidth
- \# messages * latency

```
communication
```

- Time_per_flop << 1/ bandwidth << latency
- Gaps growing exponentially with time (FOSC, 2004)

| Annual improvements |  |  |  |
| :---: | :---: | :---: | :---: |
| Time_per_flop |  | Bandwidth | Latency |
| $59 \%$ | Network | $26 \%$ | $15 \%$ |
|  | DRAM | $23 \%$ | $5 \%$ |

- Goal : reorganize linear algebra to avoid communication
- Between all memory hierarchy levels
- $\mathrm{L} 1 \longleftrightarrow \mathrm{~L} 2 \longleftrightarrow$ DRAM $\longleftrightarrow$ network, etc
- Not just hiding communication (speedup $\leq 2 x$ )
- Arbitrary speedups possible


## Direct linear algebra: Prior Work on Matmul

- Assume $\mathrm{n}^{3}$ algorithm for $\mathrm{C}=\mathrm{A}$ *B (i.e. not Strassen-like)
- Sequential case, with fast memory of size M
- Lower bound on \#words moved to/from slow memory $=\Omega\left(\mathrm{n}^{3} / \mathrm{M}^{1 / 2}\right) \quad$ [Hong, Kung, 81]
- Attained using "blocked" or cache-oblivious algorithms
- Parallel case on P processors:
- Assume load balanced, one copy each of A, B, C
- Lower bound on \#words communicated $=\Omega\left(n^{2} / P^{1 / 2}\right) \quad[$ Irony, Tiskin, Toledo, 04]
- Attained by Cannon's Algorithm


## Lower bound for all "direct" linear algebra

- Let M = "fast" memory size (per processor)
\#words_moved (per processor) = $\Omega$ (\#flops (per processor) / $\mathbf{M}^{1 / 2}$ )
\#messages_sent (per processor) $=\Omega$ (\#flops (per processor) / $\mathbf{M}^{3 / 2}$ )
- Parallel case: assume either load or memory balanced
- Holds for
- BLAS, LU, QR, eig, SVD, tensor contractions, ...
- Some whole programs (sequences of these operations, no matter how individual ops are interleaved, eg $\mathrm{A}^{k}$ )
- Dense and sparse matrices (where \#flops << $n^{3}$ )
- Sequential and parallel algorithms
- Some graph-theoretic algorithms (eg Floyd-Warshall)


## Can we attain these lower bounds?

- Do conventional dense algorithms as implemented in LAPACK and ScaLAPACK attain these bounds?
- Mostly not
- If not, are there other algorithms that do?
- Yes, for dense linear algebra
- Only a few sparse algorithms so far (eg Cholesky)


## One-sided Factorizations (LU, QR), so far

- Classical Approach for $\mathrm{i}=1$ to n
update column i
update trailing matrix
- \#words_moved = O( $\mathrm{n}^{3}$ )
- Recursive Approach
func factor(A)
if $A$ has 1 column, update it else
factor(left half of $A$ )
update right half of $A$
factor(right half of $A$ )
- $\#$ words moved $=0\left(n^{3} / M^{1 / 2}\right)$
- Blocked Approach (LAPACK) for $\mathrm{i}=1$ to $\mathrm{n} / \mathrm{b}$
update block i of b columns
update trailing matrix
- \#words moved $=0\left(n^{3} / M^{1 / 3}\right)$
- None of these approaches minimizes \#messages
- Need another idea


## TSQR: QR of a Tall, Skinny matrix

$$
\mathrm{W}=\left(\begin{array}{l}
\mathrm{W}_{0} \\
\hline \mathrm{~W}_{1} \\
\hline \frac{\mathrm{~W}_{2}}{\mathrm{~W}_{3}}
\end{array}\right)
$$

$$
\left(\begin{array}{l}
R_{00} \\
\frac{R_{10}}{R_{20}} \\
R_{30}
\end{array}\right)=\left(\frac{Q_{01} R_{01}}{Q_{11} R_{11}}\right)
$$

$$
\left(\frac{R_{01}}{R_{11}}\right)=\left(Q_{02} R_{02}\right)
$$

## TSQR: QR of a Tall, Skinny matrix

$$
W=\left(\begin{array}{l}
W_{0} \\
\hline W_{1} \\
\hline \frac{W_{2}}{W_{3}}
\end{array}\right)=\left(\frac{Q_{00} R_{00}}{\frac{Q_{10} R_{10}}{Q_{20} R_{20}}} \begin{array}{|}
\mathrm{Q}_{30} R_{30}
\end{array}\right)=\binom{\frac{\mathrm{Q}_{00}}{\mathrm{Q}_{10}}}{\frac{Q_{20}}{Q_{30}}} \cdot\binom{\frac{R_{00}}{R_{10}}}{\frac{R_{20}}{R_{30}}}
$$

$$
\left(\begin{array}{l}
R_{00} \\
\frac{R_{10}}{R_{20}} \\
R_{30}
\end{array}\right)=\left(\frac{Q_{01} R_{01}}{Q_{11} R_{11}}\right)=\left(\frac{Q_{01}}{Q_{11}}\right) \cdot\left(\frac{R_{01}}{R_{11}}\right)
$$

$$
\left(\frac{R_{01}}{R_{11}}\right)=\left(Q_{02} R_{02}\right)
$$

## Minimizing Communication in TSQR

Parallel: $W=\left[\begin{array}{l}W_{0} \\ W_{1} \\ W_{2} \\ W_{3}\end{array}\right] \rightarrow R_{00} \rightarrow R_{10} \rightarrow R_{01} \longrightarrow R_{02} \longrightarrow R_{11} \longrightarrow R_{02}$

Sequential: $W=\left[\begin{array}{l}W_{0} \\ W_{1} \\ W_{2} \\ W_{3}\end{array}\right] \longrightarrow R_{00} \longrightarrow R_{01} \longrightarrow R_{02} \longrightarrow R_{03}$

Dual Core: $\quad W=\left[\begin{array}{c}W_{0} \\ W_{1} \\ W_{2} \\ W_{3}\end{array}\right] \longrightarrow R_{00} \longrightarrow R_{01} \longrightarrow R_{02} \longrightarrow R_{11} \longrightarrow R_{03}$
Multicore / Multisocket / Multirack / Multisite / Out-of-core: ?
Can choose reduction tree dynamically

## TSQR Performance Results

- Parallel
- Intel Clovertown
- Up to $8 x$ speedup (8 core, dual socket, 10M x 10)
- Pentium III cluster, Dolphin Interconnect, MPICH
- Up to $6.7 x$ speedup ( 16 procs, $100 \mathrm{~K} \times 200$ )
- BlueGene/L
- Up to 4 x speedup ( 32 procs, $1 \mathrm{M} \times 50$ )
- Grid - $4 x$ on 4 cities (Dongarra et al)
- Cloud - early result - up and running using Mesos
- Sequential
- "Infinite speedup" for out-of-Core on PowerPC laptop
- As little as $2 x$ slowdown vs (predicted) infinite DRAM
- LAPACK with virtual memory never finished
- What about QR for a general matrix?

Data from Grey Ballard, Mark Hoemmen, Laura Grigori, Julien Langou, Jack Dongarra, Michael Anderson

# CAQR on a GPU (Fermi C2050) (1/2) \# rows = 8192 



## CAQR on a GPU (Fermi C2050) (2/2) \#rows = 8192



Biggest speedup over MAGMA QR is $13 x$ for $1 \mathrm{M} \times 192$

## Partial History of dense sequential QR algorithms attaining communication lower bounds

- Algorithms shown minimizing \# Messages assume (recursive) block layout
- Many references (see reports), only some shown, plus ours
- Oldest reference may or may not include analysis
- Cache-oblivious are underlined, Green are ours, ? is unknown/future work
- $\mathrm{b}=$ block size, $\mathrm{M}=$ fast memory size, $\mathrm{n}=$ dimension

| Algorithm | 2 Levels of Memory |  | Multiple Levels of Memory |  |
| :---: | :---: | :---: | :---: | :---: |
|  | \#Words Moved | and \# Messages | \#Words Moved | and \# Messages |
| QR | [Elmroth, <br> Gustavson, 98] <br> Recursive, may do more flops <br> LAPACK (only for $n \geq M$ and $b=M^{1 / 2}$, or $\mathrm{n}^{2} \leq \mathrm{M}$ ) <br> [DGHLO8] | [Frens,Wise,03] explicit Q only, $3 x$ more flops <br> [DGHLO8] implicit Q , but represented differently, ~ same flop count | [Elmroth, Gustavson,98] <br> [DGHLO8]? | [Frens,Wise,03] <br> [DGHLO8]? <br> Can we get the same flop count? |

## CALU: Using similar idea for TSLU as TSQR:

 Use reduction tree, to do "Tournament Pivoting"$$
\begin{aligned}
& W^{n \times b}=\left(\begin{array}{l}
W_{1} \\
\hline W_{2} \\
\hline W_{3} \\
\hline W_{4}
\end{array}\right)=\left(\begin{array}{l}
\left.P_{1} \cdot L_{1} \cdot U_{1}\right) \\
\hline P_{2} \cdot L_{2} \cdot U_{2} \\
\hline P_{3} \cdot L_{3} \cdot U_{3} \\
\hline P_{4} \cdot L_{4} \cdot U_{4}
\end{array}\right) \begin{array}{l}
\text { Choose } b \text { pivot rows of } W_{1} \text {, call them } W_{1}^{\prime}, \\
\text { Choose } b \text { pivot rows of } W_{2}, \text { call them } W_{2}^{\prime} \\
\text { Choose } b \text { pivot rows of } W_{3} \text {, call them } W_{3}^{\prime} \\
\text { Choot rows of } W_{4}, \text { call them } W_{4}^{\prime}
\end{array} \\
&\left(\begin{array}{l}
W_{1}^{\prime} \\
W_{2}^{\prime} \\
\hline W_{3}^{\prime} \\
W_{4}^{\prime}
\end{array}\right)=\binom{P_{12} \cdot L_{12} \cdot U_{12}}{P_{34} \cdot L_{34} \cdot U_{34}} \quad \begin{array}{l}
\text { Choose b pivot rows, call them } W_{12}^{\prime}
\end{array} \\
&\binom{W_{12}^{\prime}}{W_{34}^{\prime}}=P_{1234} \cdot L_{1234} \cdot U_{1234} \quad \text { Choose b pivot rows, call them } W_{34}^{\prime}
\end{aligned}
$$

- Go back to W and use these b pivot rows (move them to top, do LU without pivoting)
- Repeat on each set of $b$ columns to do CALU


## Stability of CALU using Tournament Pivoting

- Thm: Tournament Pivoting (TP) as stable as Partial Pivoting (PP) in following sense: TP gets same results as PP applied to different input matrix whose entries are blocks taken from input A
- Worst case pivot growth factor on $\mathrm{n} \times \mathrm{n}$ matrix - Proven: $2^{\mathrm{nH}-1}$ where $\mathrm{H}=1+$ height of reduction tree - Attained: $2^{\mathrm{n}-1}$, same as PP
- There are examples where PP is exponentially less stable than TP, and vice-versa, so neither is always better than the other
- Extensive numerical testing confirms stability


## Exascale Machine Parameters

- 2^30 $\approx 1,000,000$ nodes
- 1024 cores/node (a billion cores!)
- $100 \mathrm{~GB} /$ sec interconnect bandwidth
- 400 GB/sec DRAM bandwidth
- 1 microsec interconnect latency
- 50 nanosec memory latency
- 32 Petabytes of memory
- 1/2 GB total L1 on a node
- 20 Megawatts !?


## Exascale predicted speedups for CA-LU vs ScaLAPACK-LU



## History of Spectral Divide-and-Conquer for Eigenvalue problems and the SVD

- Ideas go back to Bulgakov, Godunov, Malyshev [BG88,Mal89]
- Bai, D., Gu [bDG97]
- Reduced to matmul, QR, generalized QR with pivoting (bug)
- D., Dumitriu, Holtz [DDH07]
- Uses RURV (randomized URV), not pivoted QR (same bug)
- D., Grigori, Hoemmen, Langou [DGHL08]
- Communication-avoiding QR (CAQR)
- New Communication-avoiding algorithm [BDD10]
- Use generalized RURV for better rank-detection (fixes bug)
- Uses communication-avoiding matmul and QR
- Uses randomization in divide-and-conquer
- Ballard, D., Dumitriu


## Overview of Spectral Divide-and-Conquer

- Implicitly compute $\left(1+A^{-2^{k}}\right)^{-1}$
- Maps spectrum to 0 and 1 (as $k \rightarrow \infty$ )
- Converges to spectral projector of A for $|\lambda|>1$
- Rank Revealing QR yields invariant subspace

$$
A_{\text {new }}=Q^{*} A Q=\left(\begin{array}{rr}
A_{11} & A_{12} \\
\varepsilon & A_{22}
\end{array}\right)
$$

- Algorithm:
$\mathrm{A}_{0}=\mathrm{A}, \mathrm{B}_{0}=\mathrm{I}$
Repeat

$$
\begin{aligned}
& \left(\begin{array}{ll}
Q_{11} & Q_{12} \\
Q_{12} & Q_{22}
\end{array}\right] \cdot\left[\begin{array}{l}
R_{k} \\
0
\end{array}\right]=\operatorname{qr}\left[\begin{array}{l}
B_{k} \\
A_{k}
\end{array}\right] \\
& A_{k+1}=Q_{12}{ }^{*} \cdot A_{k} \\
& B_{k+1}=Q_{22}{ }^{*} \cdot B_{k}
\end{aligned}
$$

- Fact: $A_{k}^{-1} B_{k}=A^{-2^{k}}$
- So need RRQR of

$$
\left(I+A^{-2^{k}}\right)^{-1}=\left(A_{k}+B_{k}\right)^{-1} A_{k}
$$

without forming $\left(A_{k}+B_{k}\right)^{-1} A_{k}$
until $R_{k}$ converges

## Rank Revealing CA-Decompositions

- Randomized URV of A
- $\left[\mathrm{V}, \mathrm{R}^{\prime}\right]=\operatorname{caqr}(\operatorname{randn}(\mathrm{n}, \mathrm{n})) \ldots \mathrm{V}$ random orthogonal
$[U, R]=\operatorname{caqr}\left(A \cdot V^{\top}\right)$
- Thm: $A=U R V$ reveals the rank of $A$ with high probability
- Randomized URV of $\left(A_{k}+B_{k}\right)^{-1} A_{k}$
- $\left[\mathrm{V}, \mathrm{R}^{\prime}\right]=\operatorname{caqr}(\operatorname{randn}(\mathrm{n}, \mathrm{n}))$... V random orthogonal
$\left[Q_{1}, R_{1}\right]=\operatorname{caqr}\left(A_{k} \cdot V^{\top}\right)$
$\left[\mathrm{R}_{2}, \mathrm{U}\right]=\operatorname{carq}\left(\mathrm{Q}_{1}{ }^{\top} \cdot\left(\mathrm{A}_{\mathrm{k}}+\mathrm{B}_{\mathrm{k}}\right)\right)$
- Thm: $\left(A_{k}+B_{k}\right)^{-1} A_{k}=U^{\top}\left(R_{2}{ }^{-1} R_{1}\right) V$ reveals rank with high probability
- Only need $U$, not $R_{2}^{-1} R_{1}$
- Applies to $\mathrm{M}_{1}{ }^{ \pm 1} \cdot \mathrm{M}_{2}{ }^{ \pm 1} \cdot \mathrm{M}_{3}{ }^{ \pm 1} \cdot \cdots=$ URV
- QR with Tournament Pivoting
- Emulate column pivoting by using a "tournament" to pick next group of $b$ columns
- Only applies to single matrix, not $\left(A_{k}+B_{k}\right)^{-1} A_{k}$


## Splitting the spectrum in different places

- Splitting the spectrum along unit circle not enough
- Nonsymmetric case
- Moebius transformation $z \rightarrow(\alpha z+\beta) /(\gamma z+\delta)$ maps unit circle to any line or circle in complex plane
- Apply splitting algorithm to $\left(\mathrm{A}_{0}, \mathrm{~B}_{0}\right)=(\alpha \mathrm{A}+\beta \mathrm{I}, \gamma \mathrm{A}+\delta \mathrm{I})$ instead of ( $\mathrm{A}, \mathrm{I}$ ), to split along any desired line or circle
- Symmetric case - easier
- Just shift A to A- $\sigma$ I
- SVD analogous


## Randomized Spectral Divide-and-Conquer

- Pick random place to divide spectrum
- Two progress metrics
- Spectrum splits, with $\mathrm{f} \cdot \mathrm{n}$ and (1-f) $\cdot \mathrm{n}$ evals per part, with $0<\mathrm{f}<1$
- Region splits, with fractions $f$ and 1-f of area in each
- Progress guaranteed with high probability; always stable
- SymEig and SVD: random point near center of spectrum
- NonsymEig: random line/circle near "center" of spectrum



## Successive Band Reduction

- Works just for eig(A=A ${ }^{\top}$ ) or SVD
- Bischof, Lang Sun
- Step 1: Reduce Dense $\rightarrow$ Narrow Band
- If bandwidth $\approx \mathrm{M}^{1 / 2}$, can minimize communication
- Step 2: Reduce Narrow Band $\rightarrow$ Tridiagonal
- Trickier...


## Successive Band Reduction



## Summary of dense parallel algorithms attaining communication lower bounds

- Assume nxn matrices on $P$ processors, memory per processor $=O\left(n^{2} / P\right)$
- ScaLAPACK assumes best block size b chosen
- Many references (see reports), Green are ours
- Recall lower bounds:
\#words_moved $=\Omega\left(n^{2} / P^{1 / 2}\right) \quad$ and $\quad$ \#messages $=\Omega\left(P^{1 / 2}\right)$

| Algorithm | Reference | Factor exceeding lower bound for \#words moved | Factar exceeding bound for \#nessages |
| :---: | :---: | :---: | :---: |
| Matrix multiply | [Cannon, 69] | 1 | 1 |
| Cholesky | ScaLAPACK | +og | $\log P$ |
| LU | [GDX08] ScaLADANㅏㄴ | $\begin{aligned} & \log P \\ & \log P \end{aligned}$ | $\begin{gathered} \log P \\ \left(N / P^{1 / 2}\right) \cdot \log P \end{gathered}$ |
| QR | PCHLO8] <br> ScaLAPACK | $\begin{aligned} & \log P \\ & \log P \end{aligned}$ | $\begin{gathered} \log ^{3} P \\ \left(N / P^{1 / 2}\right) \cdot \log P \end{gathered}$ |
| Sym Eig, SVD | [BDD10] <br> ScaLAPACK | $\begin{aligned} & \log P \\ & \log P \end{aligned}$ | $\begin{aligned} & \log ^{3} P \\ & N / P^{1 / 2} \end{aligned}$ |
| Nonsym Eig | [BDD10] <br> ScaLAPACK | $\begin{gathered} \hline \log P \\ \mathrm{P}^{1 / 2} \cdot \log \mathrm{P} \end{gathered}$ | $\begin{gathered} \log ^{3} P \\ N \cdot \log P \end{gathered}$ |

## Summary of dense parallel algorithms attaining communication lower bounds

- Assume nxn matrices on $P$ processors, memory per processor $=O\left(n^{2} / P\right)$ ? Why?
- ScaLAPACK assumes best block size b chosen
- Many references (see reports), Green are ours
- Recall lower bounds:
\#words_moved $=\Omega\left(n^{2} / P^{1 / 2}\right) \quad$ and $\quad$ \#messages $=\Omega\left(P^{1 / 2}\right)$

| Algorithm | Reference | Factor exceeding lower bound for \#words_moved | Factor exceeding bound for nessaages |
| :---: | :---: | :---: | :---: |
| Matrix multiply | [Cannon, 69] | 1 | 1 |
| Cholesky | ScaLAPACK | Hog P | $\log P$ |
| LU | $\begin{aligned} & \text { [GDX08] } \\ & \text { ScaLADA } \end{aligned}$ | $\begin{aligned} & \log P \\ & \log P \end{aligned}$ | $\begin{gathered} \log P \\ \left(N / P^{1 / 2}\right) \cdot \log P \end{gathered}$ |
| QR | paHLO8] <br> ScaLAPACK | $\begin{aligned} & \log P \\ & \log P \end{aligned}$ | $\begin{gathered} \log ^{3} P \\ \left(N / P^{1 / 2}\right) \cdot \log P \end{gathered}$ |
| Sym Eig, SVD | [BDD10] <br> ScaLAPACK | $\begin{aligned} & \log P \\ & \log P \end{aligned}$ | $\begin{aligned} & \log ^{3} P \\ & N / P^{1 / 2} \end{aligned}$ |
| Nonsym Eig | [BDD10] <br> ScaLAPACK | $\begin{gathered} \log \mathrm{P} \\ \mathrm{P}^{1 / 2} \cdot \log \mathrm{P} \end{gathered}$ | $\begin{gathered} \log ^{3} P \\ N \cdot \log P \end{gathered}$ |

## Beating \#words_moved $=\Omega\left(\mathrm{n}^{2} / \mathrm{P}^{1 / 2}\right)$

- \#words_moved $=\Omega\left(\left(\mathrm{n}^{3} / \mathrm{P}\right) /\right.$ local_mem $\left.{ }^{1 / 2}\right)$
- If one copy of data, local_mem = $\mathrm{n}^{2} / \mathrm{P}$
- Can we use more memory to communicate less?
- Johnson's Matmul Algorithm on $\mathrm{P}^{1 / 3} \times \mathrm{P}^{1 / 3} \times \mathrm{P}^{1 / 3}$ processor grid
- Broadcast A in j direction ( $\mathrm{P}^{1 / 3}$ redundant copies)
- Broadcast B in i direction (ditto) "C face"
- Local multiplies
- Reduce (sum) in k direction
- Communication volume
$=O\left(\left(n / P^{1 / 3}\right)^{2} \log (P)\right)$ - optimal
- Number of messages
$=\mathrm{O}(\log (\mathrm{P}))$ - optimal
- Extends to $1 \leq \mathrm{C} \leq \mathrm{P}^{1 / 3}$ copies $_{\mathrm{i}} \leftarrow$
- What about LU, etc?

"A face"


## Predicted Exascale

2.5D Matmul - Sveedun over Cannon
2.5D/Cannon time speed up


## Exascale 2.5D Matmul - optimal \# copies

2.5D Best choice of c for effieciency


## Strassen-like Algorithms

- M = fast memory size
- Conventional matmul
- \#flops = $2 n^{3}$
- \#words_moved $=\Omega\left(n^{3} / M^{1 / 2}\right)=\Omega\left(M\left(n / M^{1 / 2}\right)^{3}\right)$, attainable
- Strassen matmul
- \#flops $=\Theta\left(\mathrm{n}^{\omega}\right)$ where $\omega=\log _{2}(7) \approx 2.81$
- \#words_moved $=\Omega\left(\mathrm{M}\left(\mathrm{n} / \mathrm{M}^{1 / 2}\right)^{\omega}\right)$, attainable too
- Applies to other sufficiently "Strassen-like" algorithms
- How broadly does it apply?
- We know rest of linear algebra can be done in $O\left(n^{\omega}\right)$ flops
- Also with \#words_moved $=\Omega\left(M\left(n / M^{1 / 2}\right)^{\omega}\right) \quad$ (sequential)
- Lower bounds?


## Summary of Direct Linear Algebra

- New lower bounds, optimal algorithms, big speedups in theory and practice
- Lots of other progress, open problems
- New ways to "pivot"
- Heterogeneous architectures
- Some sparse algorithms
- Autotuning. . .


## Avoiding Communication in Iterative Linear Algebra

- $k$-steps of iterative solver for sparse $A x=b$ or $A x=\lambda x$
- Does k SpMVs with A and starting vector
- Many such "Krylov Subspace Methods"
- Conjugate Gradients (CG), GMRES, Lanczos, Arnoldi, ...
- Goal: minimize communication
- Assume matrix "well-partitioned"
- Serial implementation
- Conventional: $\mathrm{O}(\mathrm{k})$ moves of data from slow to fast memory
- New: O(1) moves of data - optimal
- Parallel implementation on p processors
- Conventional: O(k $\log p$ ) messages ( $k$ SpMV calls, dot prods)
- New: $\mathbf{O}(\log p)$ messages - optimal
- Lots of speed up possible (modeled and measured)
- Price: some redundant computation


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

$$
\begin{align*}
& A \cdot x \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \\
& X \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \\
& 1234 \ldots
\end{align*}
$$

- Example: A tridiagonal, n=32, k=3
- Works for any "well-partitioned" A


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

$$
\begin{aligned}
& A^{3} \times \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet
\end{aligned}
$$

$$
\begin{align*}
& x \text { ••! } \\
& 1234 \text {... }
\end{align*}
$$

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

$$
\begin{aligned}
& A^{3} \times \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet \bullet
\end{aligned}
$$

$$
\begin{align*}
& \text { x •••••• } \\
& 1234 \text {... }
\end{align*}
$$

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Sequential Algorithm



## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Sequential Algorithm



## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Sequential Algorithm



## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Sequential Algorithm

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$


## Communication Avoiding Kernels: The Matrix Powers Kernel : [ $\left.\mathrm{Ax}, \mathrm{A}^{2} \mathrm{x}, \ldots, \mathrm{A}^{k} \mathrm{x}\right]$

- Replace $k$ iterations of $y=A \cdot x$ with $\left[A x, A^{2} x, \ldots, A^{k} x\right]$
- Parallel Algorithm

- Example: A tridiagonal, $\mathrm{n}=32, \mathrm{k}=3$
- Each processor works on (overlapping) trapezoid


## Communication Avoiding Kernels: The Matrix Powers Kernel : $\left[A x, A^{2} x, \ldots, A^{k} x\right]$

Remotely Dependent Entries for $\left[x, A x, A^{2} x, A^{3} x\right]$, A irregular, multiple processors


## Minimizing Communication of GMRES to solve $A x=b$

- GMRES: find $x$ in span $\left\{b, A b, \ldots, A^{k} b\right\}$ minimizing || $A x-b \mid \|_{2}$

Standard GMRES
for $\mathrm{i}=1$ to k $\mathrm{w}=\mathrm{A} \cdot \mathrm{v}(\mathrm{i}-1) \quad . . \mathrm{SpM} \mathrm{V}$ MGS(w, v(0),...,v(i-1)) update $\mathrm{v}(\mathrm{i}), \mathrm{H}$
endfor
solve LSQ problem with H

Communication-avoiding GMRES
$W=\left[v, A v, A^{2} v, \ldots, A^{k} v\right]$
$[\mathrm{Q}, \mathrm{R}]=\operatorname{TSQR}(\mathrm{W})$
... "Tall Skinny QR"
build H from R solve LSQ problem with $H$

Sequential case: \#words moved decreases by a factor of $k$ Parallel case: \#messages decreases by a factor of $k$
-Oops - W from power method, precision lost!

Matrix diag-cond-1.000000e-11: rel. 2-nrm resid.


# Speed ups of GMRES on 8-core Intel Clovertown Requires Co-tuning Kernels 

## [MHDY09]

Runtime per kernel, relative to CA-GMRES(k,t), for all test matrices, using 8 threads and restart length 60


Exascale predicted speedups for Matrix Powers Kernel over SpMV for 2D Poisson (5 point stencil)

Akx/regular speed up


## Summary of Iterative Linear Algebra

- New Lower bounds, optimal algorithms, big speedups in theory and practice
- Lots of other progress, open problems
- GMRES, CG, BiCGStab, Arnoldi, Lanczos reorganized
- Other Krylov methods?
- Recognizing stable variants more easily?
- Avoiding communication with preconditioning harder
- "Hierarchically semi-separable" preconditioners work
- Autotuning


## For further information

- www.cs.berkeley.edu/~demmel
- Papers
- bebop.cs.berkeley.edu
- www.netlib.org/lapack/lawns
- 1-week-short course - slides and video
- www.ba.cnr.it/ISSNLA2010


## Collaborators and Supporters

- Collaborators
- Michael Anderson (UCB), Grey Ballard (UCB), Erin Carson (UCB), Jack Dongarra (UTK), Ioana Dumitriu (U. Wash), Laura Grigori (INRIA), Ming Gu (UCB), Mark Hoemmen (Sandia NL), Olga Holtz (UCB \& TU Berlin), Nick Knight (UCB), Julien Langou, (U Colo. Denver), Marghoob Mohiyuddin (UCB), Oded Schwartz (TU Berlin), Edgar Solomonik (UCB), Michelle Strout (Colo. SU), Vasily Volkov (UCB), Sam Williams (LBNL), Hua Xiang (INRIA), Kathy Yelick (UCB \& LBNL)
- Other members of the ParLab, BEBOP, CACHE, EASI, MAGMA, PLASMA, TOPS projects
- Supporters
- NSF, DOE, UC Discovery
- Intel, Microsoft, Mathworks, National Instruments, NEC, Nokia, NVIDIA, Samsung, Sun


## We're hiring!

- Seeking software engineers to help develop the next versions of LAPACK and ScaLAPACK
- Locations: UC Berkeley and UC Denver
- Please send applications to julie@cs.utk.edu


## Summary

## Time to redesign all linear algebra algorithms and software

## Don't Communic...

