



Parallel sparse linear algebra (2)

Prof. Richard Vuduc

Georgia Institute of Technology

CSE/CS 8803 PNA, Spring 2008

[L.10] Thursday, February 7, 2008



Sources for today's material

- Mike Heath (UIUC)
- CS 267 (Yelick & Demmel, UCB)
- John Gilbert (UCSB)
- Xiaoye (Sherry) Li (LBNL)
- *Direct methods for sparse matrices*, by Duff, Erisman, and Reid
- Esmond Ng (LBNL)

Algorithms for 2-D (3-D) Poisson, $N=n^2$ ($=n^3$)

Algorithm	Serial	PRAM	Memory	# procs
Dense LU	N^3	N	N^2	N^2
<i>Band LU</i>	N^2 ($N^{7/3}$)	N	$N^{3/2}$ ($N^{5/3}$)	N ($N^{4/3}$)
Jacobi	N^2 ($N^{5/3}$)	N ($N^{2/3}$)	N	N
<i>Explicit inverse</i>	N^2	$\log N$	N^2	N^2
Conj. grad.	$N^{3/2}$ ($N^{4/3}$)	$N^{1/2(1/3)} \log N$	N	N
RB SOR	$N^{3/2}$ ($N^{4/3}$)	$N^{1/2}$ ($N^{1/3}$)	N	N
Sparse LU	$N^{3/2}$ (N^2)	$N^{1/2}$	$N \log N$ ($N^{4/3}$)	N
<i>FFT</i>	$N \log N$	$\log N$	N	N
Multigrid	N	$\log^2 N$	N	N
Lower bound	N	$\log N$	N	

PRAM = idealized parallel model with zero communication cost.

Source: Demmel (1997)




Anatomy of a sparse direct solver

$$P_r \cdot A \cdot P_c^T = LU$$


- *Order equations & variables to minimize fill in L, U* [Combinatorial]
- Symbolic factorization [Allocate memory for L, U]
- Numerical factorization [Dominates run-time]
- Triangular solves [Small fraction, unless many RHS]




Review: Fill-reducing orderings





A




L






Total nnz = 233



Total nnz = 207




Ordering: Markowitz criterion ('57)

- At elimination stage k :
 - Let $r_i^{(k)} = \text{nnz}(\text{row } i)$, $c_j^{(k)} = \text{nnz}(\text{col } j)$, in uneliminated submatrix
 - Choose pivot entry a_{ij} (*i.e.*, swap row & col) that minimizes:

$$(r_i^{(k)} - 1) \times (c_j^{(k)} - 1)$$

- Don't forget about numerical values, too! (*E.g.*, threshold)
- **Problem: Expensive!**


No fill in natural ordering:




Minimum degree ordering



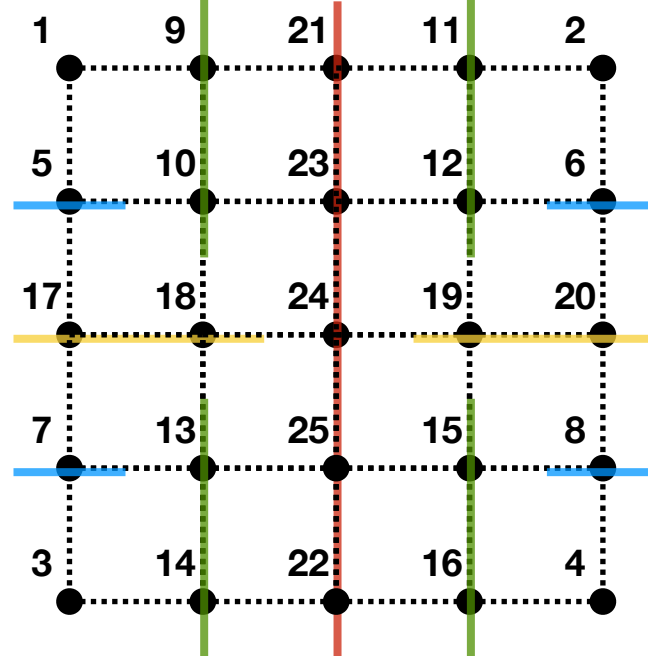


Min-degree node: 5



Choosing 5 leads to fill-in:



Nested dissection










Finding good separators

- Multilevel schemes
 - Chaco [Hendrickson & Leland '94]
 - Metis [Karypis & Kumar '95]
- Spectral bisection [Simon, *et al.* '90+]
- Geometric and spectral bisection [Chan, Gilbert, & Teng '94]




Sources of parallelism: Elimination trees






Cholesky factor



$G^+(A)$






$T(A)$
"Elimination tree"

$$T(A) : \text{parent}(j) = \min\{i > j : (i, j) \in G^+(A)\}$$


Elimination tree defines column dependencies in A.

Can get $T(A)$ from $G(A)$ in $O(nnz * \alpha(nnz))$, and $G^+(A)$ from $T(A)$ in $O(nnz(L))$.






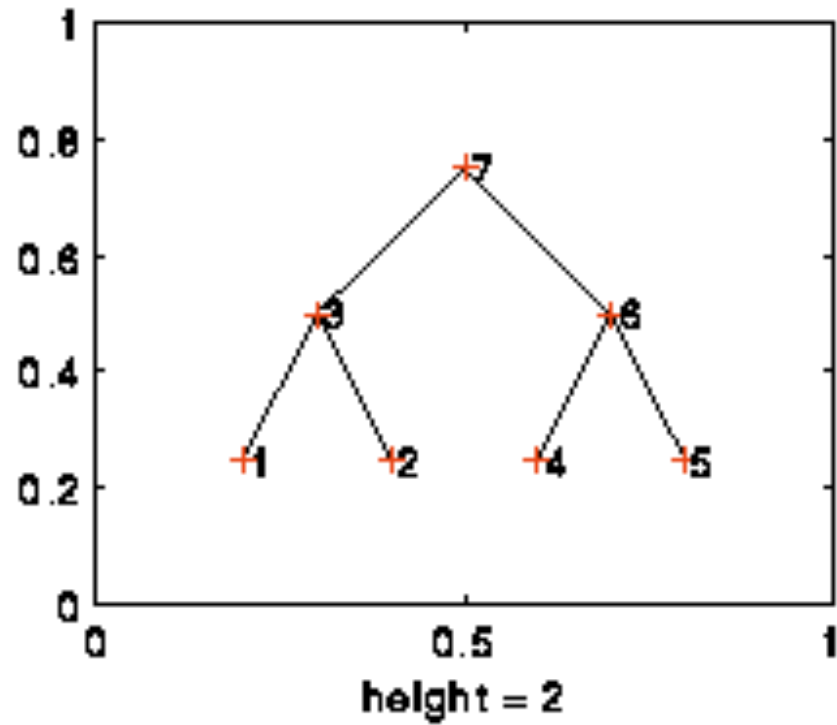
A after minimum degree



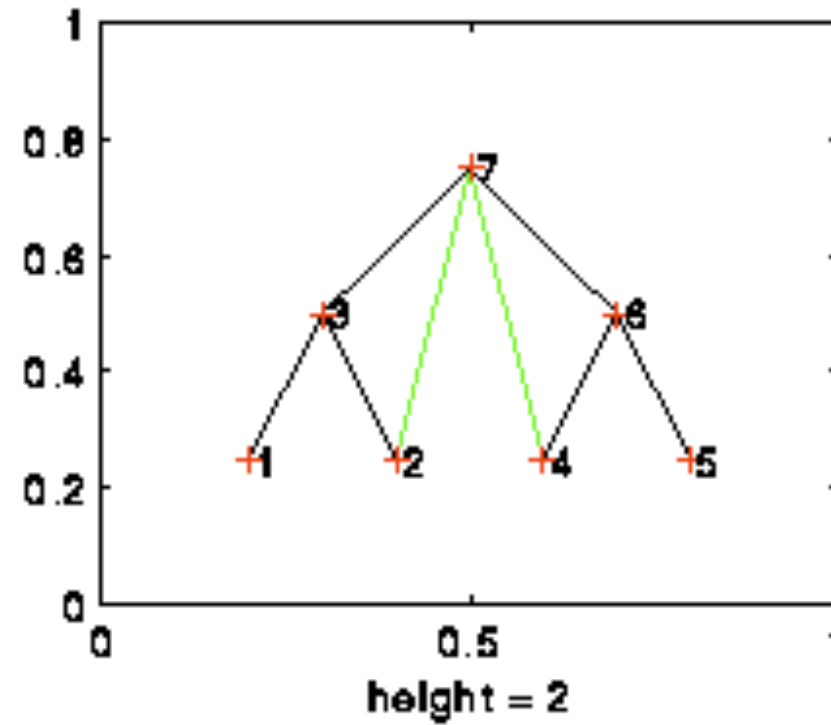
Cholesky factor, flops=35



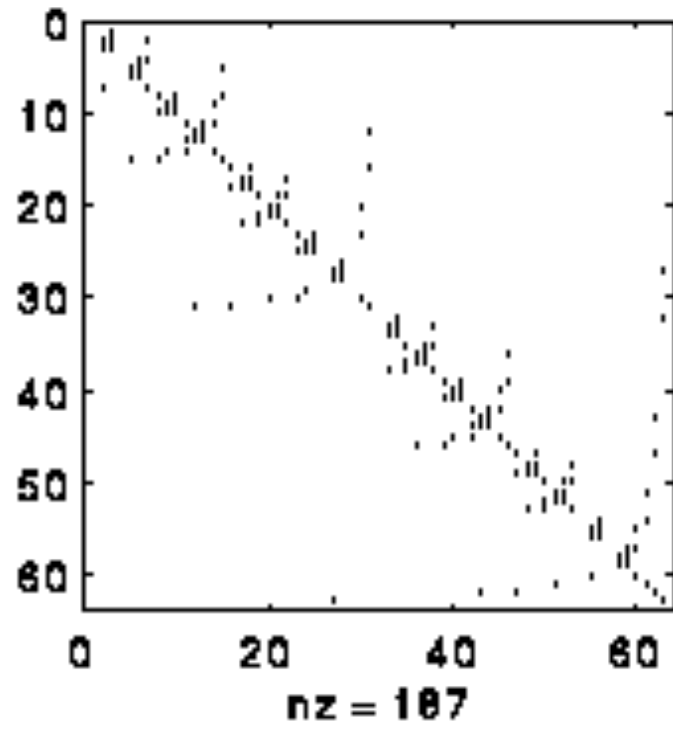
Elimination tree




Graph of Cholesky factor




A after nested dissection




Cholesky factor, flops=509



Elimination tree



Graph of Cholesky factor



Multifrontal methods (e.g., [Liu '92])




$$\begin{aligned} A &= \begin{pmatrix} B & V^T \\ V & C \end{pmatrix} \\ &= \begin{pmatrix} L_B & 0 \\ VL_B^{-T} & I \end{pmatrix} \begin{pmatrix} I & 0 \\ 0 & C - VB^{-1}V^T \end{pmatrix} \begin{pmatrix} L_B^T & L_B^{-1}V^T \\ 0 & I \end{pmatrix} \end{aligned}$$

where:

$$B = L_B L_B^T \text{ (Cholesky factorization)}$$


Observe:

$$-VB^{-1}V^T = -(VL_B^{-T})(L_B^{-1}V^T) = -\sum_{k=1}^{j-1} \begin{pmatrix} l_{j,k} \\ \vdots \\ l_{n,k} \end{pmatrix} (l_{j,k} \cdots l_{n,k})$$




A






For each node of T from leaves to root:



- Sum own row/col of A with children's *Update* matrices into *Frontal* matrix
- Eliminate current variable from *Frontal* matrix, to get *Update* matrix
- Pass *Update* matrix to parent






For each node of T from leaves to root:



- Sum own row/col of A with children's *Update* matrices into *Frontal* matrix
- Eliminate current variable from *Frontal* matrix, to get *Update* matrix
- Pass *Update* matrix to parent

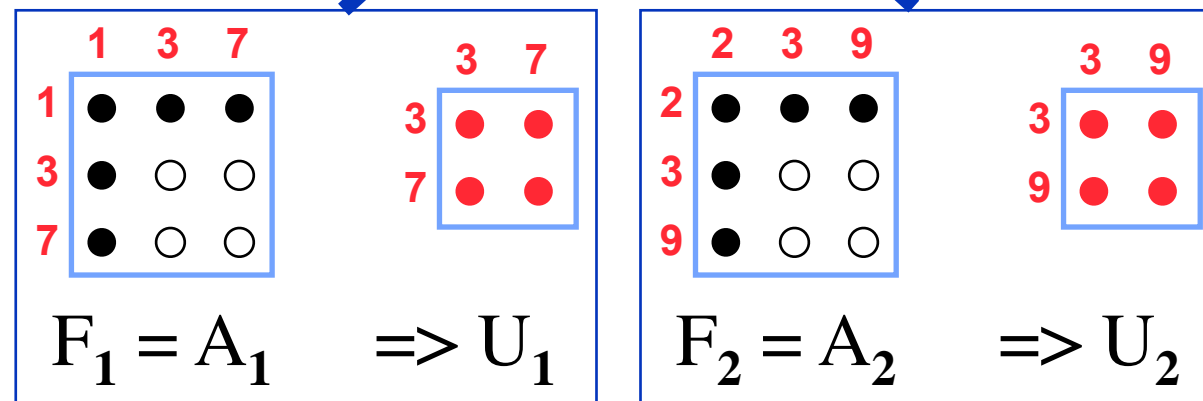










For each node of T from leaves to root:

- Sum own row/col of A with children's *Update* matrices into *Frontal* matrix
- Eliminate current variable from *Frontal* matrix, to get *Update* matrix
- Pass *Update* matrix to parent









$L+U$








Improving memory behavior: Supernodes









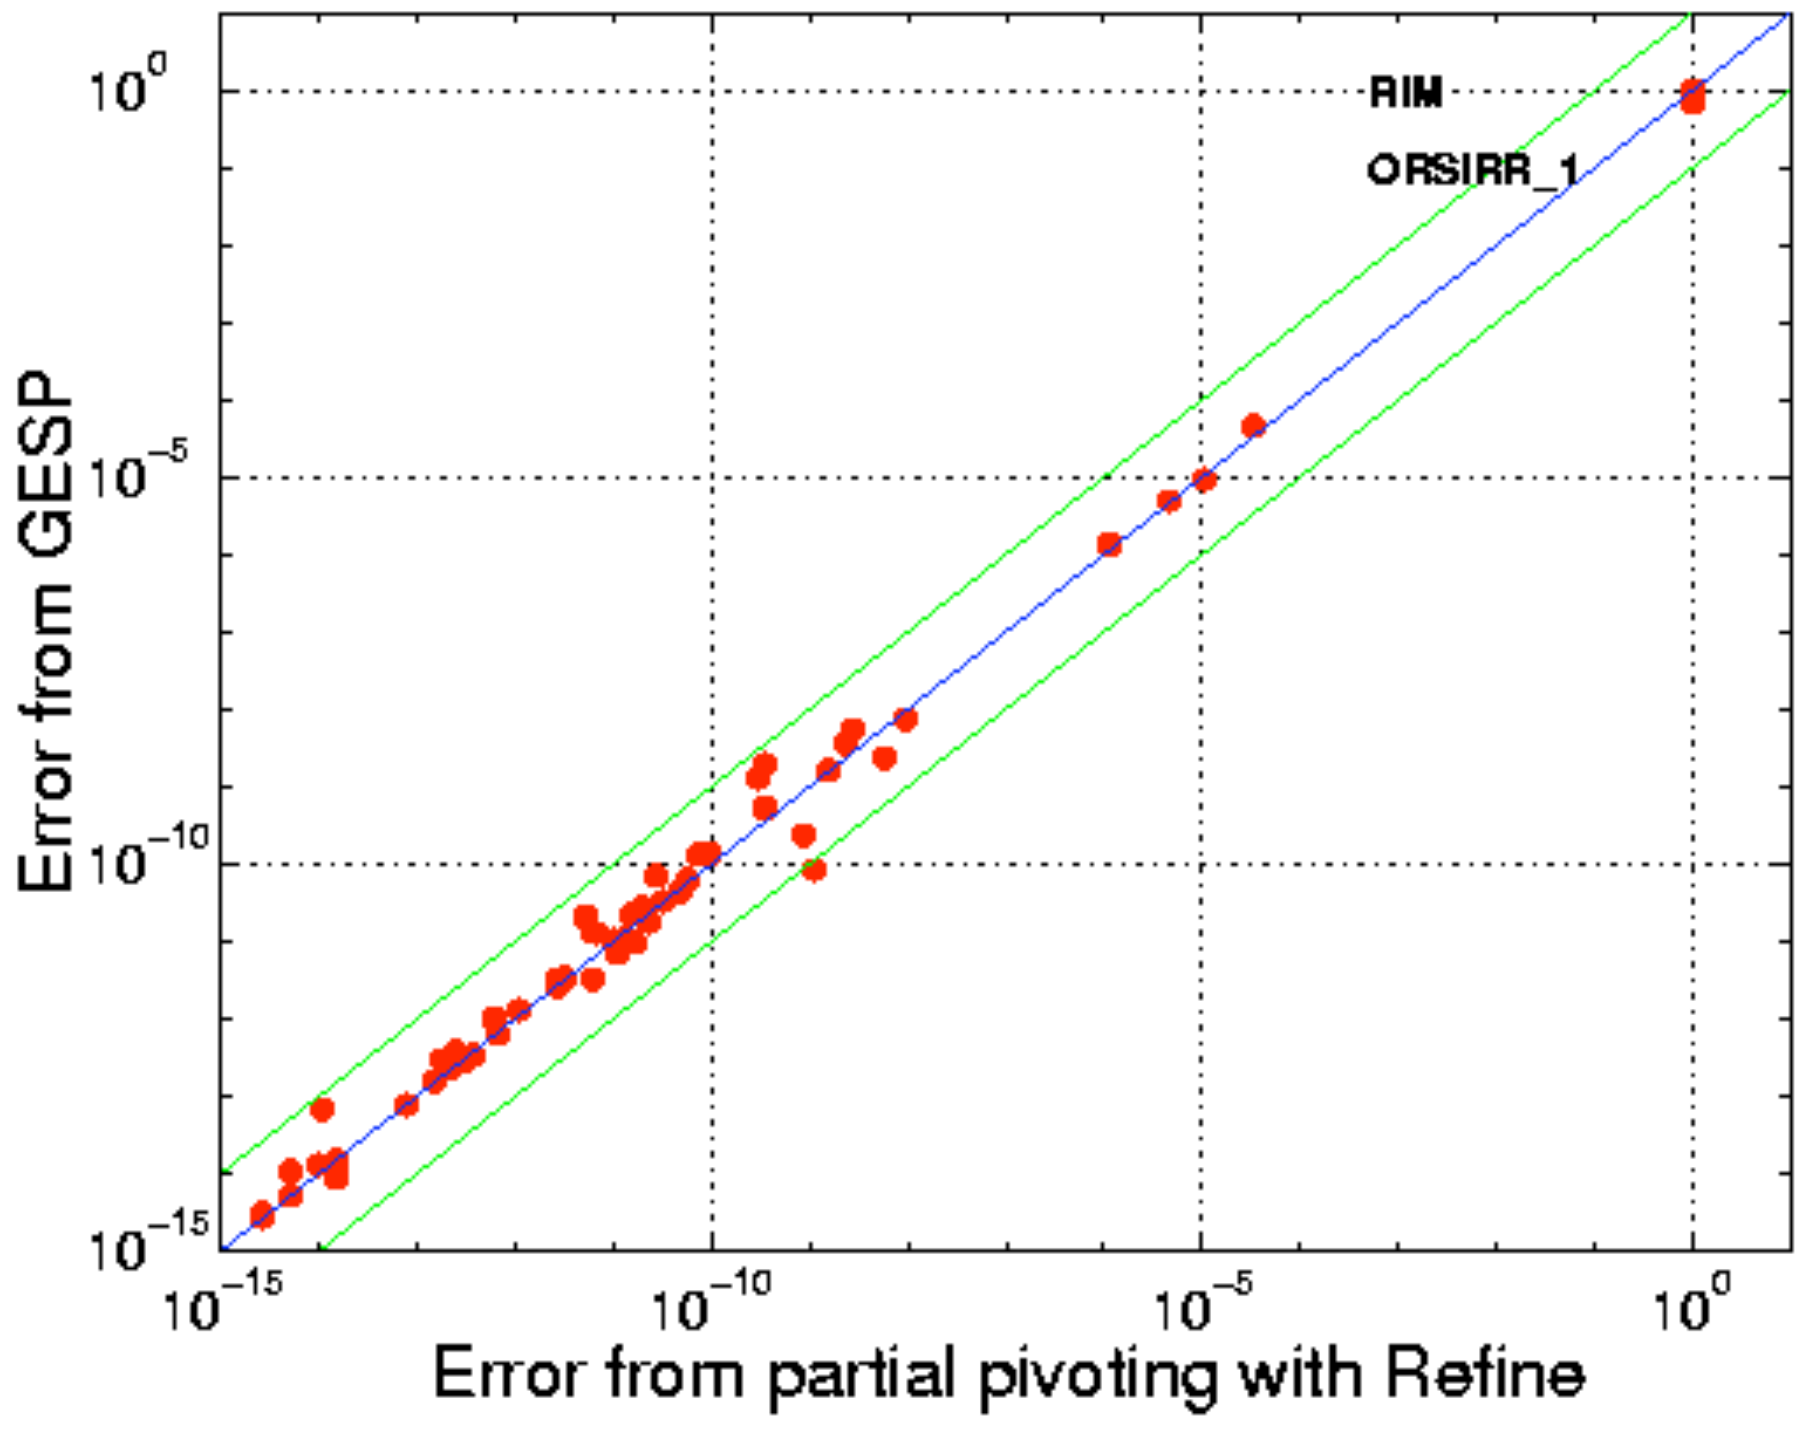


Complications in the unsymmetric (sparse LU) case



Need for numerical pivoting in LU

- Partial pivoting: Hard to implement scalably for sparse factorization
 - Not needed if A is SPD
 - Low cost in dense case, but not in sparse
- Alternative: **Static pivoting** (e.g., SuperLU_DIST)
 - Before factor, scale and permute A to maximize diagonal: $P_r^* D_r^* A^* D_c = A'$
 - Find P_r by **weighted bipartite matching** on $G(A)$
 - During factorization, replace tiny pivots with $\sqrt{\epsilon} \|A\|$
 - Iterative refinement if necessary (future lecture)



Source: X. Li (2006)





Ordering for sparse LU

- Earlier examples implicitly show or assume symmetric structure

- Options in the unsymmetric case


- Symmetric ordering for $A^T A$, based on theorem [George & Ng '87]:

$$R^T R = A^T A \quad \text{and} \quad PA = LU$$


$$\implies \text{struct}(L + U) \subseteq \text{struct}(R + R^T) \text{ for any } P$$

- Static pivoting (shortly): Similar theorem holds for $R^T R = A^T + A$ (without P)


- “Symmetrization-free” [Amestoy, Li, & Ng '03]




A



$A^T A$






Problem: Elimination tree (DAG) not known until pivoting is done


- Symbolic and numerical factorization stages become interleaved
- Option 1: Use elimination tree of $A^T A$
 - Bound sparsity, parallelism
 - Update tree on-the-fly [Demmel, Gilbert, & Li '97]
- Option 2: Pivot statically



Supernode structure in L, not U





Original matrix A



Factors L+U

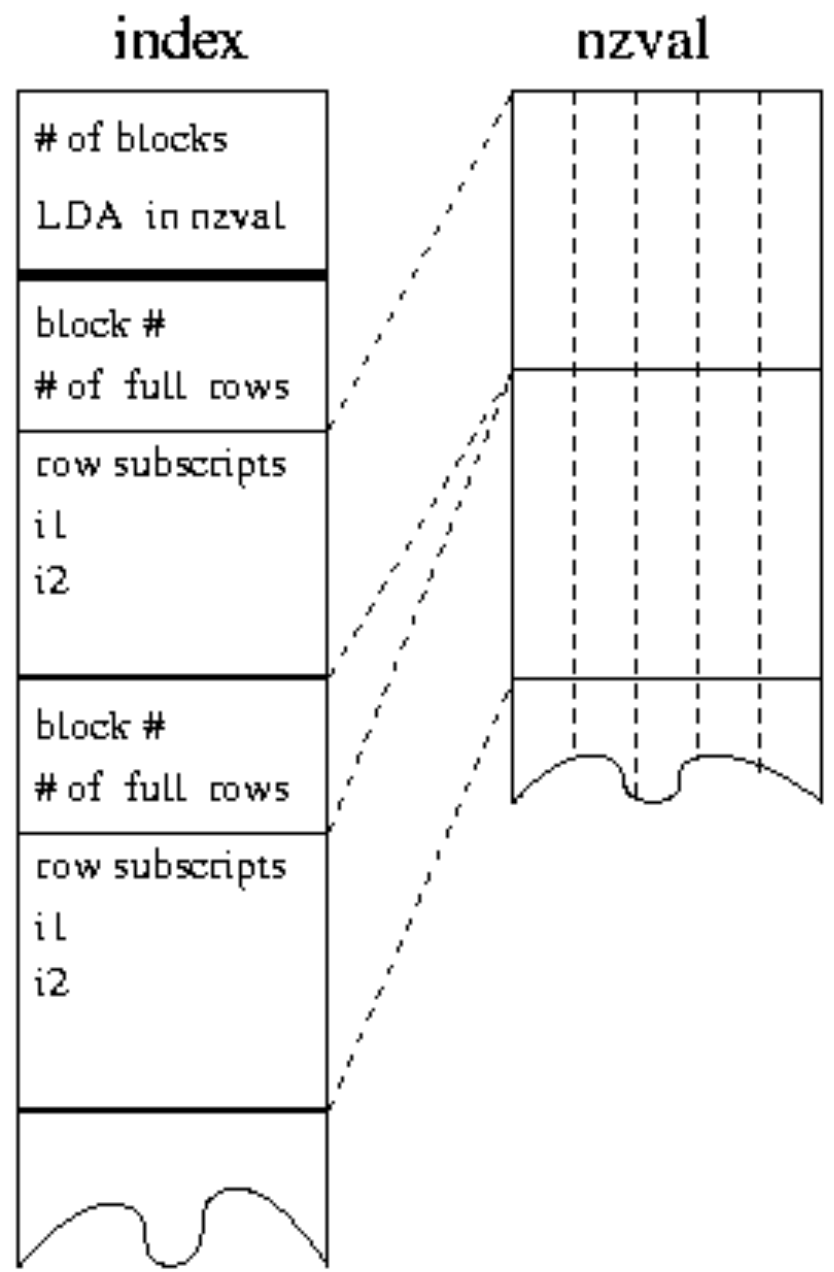


Data layout/distribution



Process Mesh

Storage of block column of L




“In conclusion...”



Sparse LU software

- SuperLU [Xiaoye “Sherry” Li @ LBNL]
- MUMPS [Amestoy, *et al.*]
- UMFPACK [Davis @ UFL]
- WSMP [Gupta @ IBM]

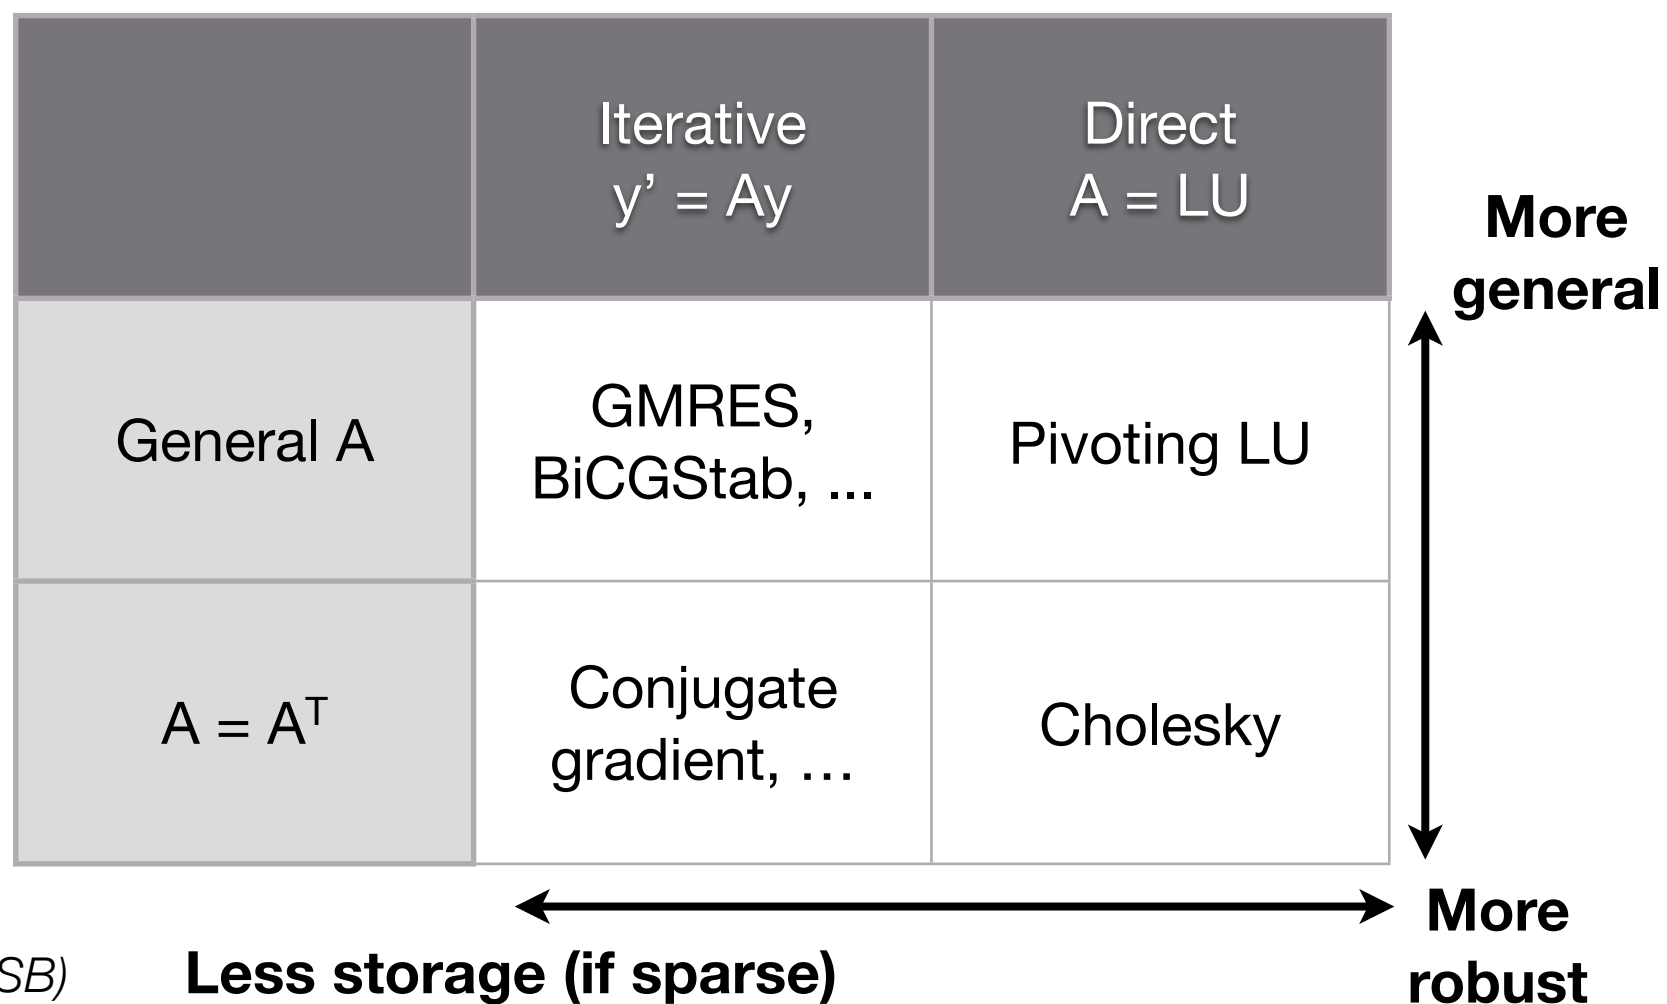


Open problems (suggested by Sherry Li '07)

- Optimizing parallel performance: blocking, tuning, reducing/hiding latency
- Graph partitioning ordering for unsymmetric case
- Scalability of sparse triangular solve
- Efficient incomplete-LU (ILU)
- Optimal complexity sparse factorization
 - Fast multipole for matrix inversion — J. Xia's dissertation (UCB, now @ UCLA)
- Latency-avoidant LU & QR



Landscape of $Ax=b$ solvers



Source: Gilbert (UCSB)



Administrivia



Administrative stuff

- **New room** (dumpier, but cozier?): College of Computing Building **(CCB) 101**
- **Accounts**: Apparently, you already have them
- Front-end login node: **ccil.cc.gatech.edu** (CoC Unix account)
 - We “own” **warp43—warp56**
 - Some docs (**MPI**): <http://www-static.cc.gatech.edu/projects/ihpcl/mpi.html>
 - **Sign-up** for mailing list: <https://mailman.cc.gatech.edu/mailman/listinfo/ihpc-lab>



Homework 1:

Parallel conjugate gradients

- Implement a parallel solver for $Ax = b$ (serial C version provided)
 - Evaluate on three matrices: 27-pt stencil, and two application matrices
 - “Simplified:” No preconditioning
 - **Bonus:** Reorder, precondition
- Performance models to understand scalability of your implementation
 - Make measurements
 - Build predictive models
- Collaboration encouraged: Compare programming models or platforms