Outline

- Motivation, overview for Dense Linear Algebra
- Review Gaussian Elimination (GE) for solving Ax=b
- Optimizing GE for caches on sequential machines
 - using matrix-matrix multiplication (BLAS)
- LAPACK library overview and performance
- Data layouts on parallel machines
- Parallel Gaussian Elimination
- ScaLAPACK library overview
- Eigenvalue problems
- Current Research

Motivation (1)

- **3 Basic Linear Algebra Problems**
 - 1. Linear Equations: Solve Ax=b for x
 - 2. Least Squares: Find x that minimizes $||\mathbf{r}||_2 = \sqrt{\Sigma r_i^2}$ where r=Ax-b
 - Statistics: Fitting data with simple functions
 - **3a. Eigenvalues: Find** λ and x where Ax = λ x
 - Vibration analysis, e.g., earthquakes, circuits
 - **3b. Singular Value Decomposition:** $A^{T}Ax = \sigma^{2}x$
 - Data fitting, Information retrieval

Lots of variations depending on structure of A

• A symmetric, positive definite, banded, ...

Motivation (2)

- Why dense A, as opposed to sparse A?
 - Many large matrices are sparse, but ...
 - Dense algorithms easier to understand
 - Some applications yields large dense matrices
 - LINPACK Benchmark (www.top500.org)
 - "How fast is your computer?" =
 "How fast can you solve dense Ax=b?"
 - Large sparse matrix algorithms often yield smaller (but still large) dense problems

Gaussian Elimination (GE) for solving Ax=b

- Add multiples of each row to later rows to make A upper triangular
- Solve resulting triangular system Ux = c by substitution

```
... for each column i

... zero it out below the diagonal by adding multiples of row i to later rows

for i = 1 to n-1

... for each row j below row i

for j = i+1 to n

... add a multiple of row i to row j

tmp = A(j,i);

for k = i to n

A(j,k) = A(j,k) - (tmp/A(i,i)) * A(i,k)
```



Refine GE Algorithm (1)

Initial Version

```
... for each column i
... zero it out below the diagonal by adding multiples of row i to later rows
for i = 1 to n-1
... for each row j below row i
for j = i+1 to n
... add a multiple of row i to row j
tmp = A(j,i);
for k = i to n
A(j,k) = A(j,k) - (tmp/A(i,i)) * A(i,k)
```

 Remove computation of constant tmp/A(i,i) from inner loop.





Refine GE Algorithm (2)

Last version

for i = 1 to n-1for j = i+1 to n m = A(j,i)/A(i,i)for k = i to n A(j,k) = A(j,k) - m * A(i,k)

 Don't compute what we already know: zeros below diagonal in column i



Refine GE Algorithm (3)

Last version

for i = 1 to n-1 for j = i+1 to n m = A(j,i)/A(i,i) for k = i+1 to n A(j,k) = A(j,k) - m * A(i,k)

 Store multipliers m below diagonal in zeroed entries for later use



Refine GE Algorithm (4)

Last version

for i = 1 to n-1 for j = i+1 to n A(j,i) = A(j,i)/A(i,i)for k = i+1 to n A(j,k) = A(j,k) - A(j,i) * A(i,k)

Split Loop

for i = 1 to n-1 for j = i+1 to n A(j,i) = A(j,i)/A(i,i)for j = i+1 to n for k = i+1 to n A(j,k) = A(j,k) - A(j,i) * A(i,k)



Store all m's here before updating rest of matrix

02/23/2007

CS267 DLA2

Refine GE Algorithm (5)

Last version

- for i = 1 to n-1for j = i+1 to n A(j,i) = A(j,i)/A(i,i)for j = i+1 to n for k = i+1 to n A(j,k) = A(j,k) - A(j,i) * A(i,k)
- Express using matrix operations (BLAS) Work at step i of Gaussian Elimination



for i = 1 to n-1A(i+1:n,i) = A(i+1:n,i) * (1 / A(i,i))A(i+1:n,i+1:n) = A(i+1:n, i+1:n) - A(i+1:n , i) * A(i , i+1:n)

CS267 DLA2

What GE really computes

for i = 1 to n-1 A(i+1:n,i) = A(i+1:n,i) / A(i,i) A(i+1:n,i+1:n) = A(i+1:n , i+1:n) - A(i+1:n , i) * A(i , i+1:n)

- Call the strictly lower triangular matrix of multipliers M, and let L = I+M
- Call the upper triangle of the final matrix U
- Lemma (LU Factorization): If the above algorithm terminates (does not divide by zero) then A = L*U
- Solving A*x=b using GE
 - Factorize A = L*U using GE (cost = 2/3 n³ flops)
 - Solve $L^*y = b$ for y, using substitution (cost = n^2 flops)
 - Solve U*x = y for x, using substitution (cost = n² flops)
- Thus A*x = (L*U)*x = L*(U*x) = L*y = b as desired

Problems with basic GE algorithm

- What if some A(i,i) is zero? Or very small?
 - Result may not exist, or be "unstable", so need to pivot
- Current computation all BLAS 1 or BLAS 2, but we know that BLAS 3 (matrix multiply) is fastest (earlier lectures...)

for i = 1 to n-1A(i+1:n,i) = A(i+1:n,i) / A(i,i) ... BLAS 1 (scale a vector) A(i+1:n,i+1:n) = A(i+1:n, i+1:n) ... BLAS 2 (rank-1 update) - A(i+1:n , i) * A(i , i+1:n) RS2: Level 1, 2 and 3 BLAS 300 Peak 250 **BLAS 3** s200 In Megatlops 150 i paad s 100 BLAS 2 BLAS 1 50 100 300 500 200 400 600 Order of vectors/matrices CS267 DLA2

02/23/2007

Pivoting in Gaussian Elimination

- A = [0 1] fails completely because can't divide by A(1,1)=0
 [1 0]
- But solving Ax=b should be easy!
- When diagonal A(i,i) is tiny (not just zero), algorithm may terminate but get completely wrong answer
 - Numerical instability
 - Roundoff error is cause
- Cure: Pivot (swap rows of A) so A(i,i) large

Gaussian Elimination with Partial Pivoting (GEPP)

• Partial Pivoting: swap rows so that A(i,i) is largest in column

```
for i = 1 to n-1

find and record k where |A(k,i)| = \max\{i \le j \le n\} |A(j,i)|

... i.e. largest entry in rest of column i

if |A(k,i)| = 0

exit with a warning that A is singular, or nearly so

elseif k != i

swap rows i and k of A

end if

A(i+1:n,i) = A(i+1:n,i) / A(i,i) ... each quotient lies in [-1,1]

A(i+1:n,i+1:n) = A(i+1:n, i+1:n) - A(i+1:n, i) * A(i, i+1:n)
```

- Lemma: This algorithm computes A = P*L*U, where P is a permutation matrix.
- This algorithm is numerically stable in practice
- For details see LAPACK code at

http://www.netlib.org/lapack/single/sgetf2.f

Parallelizing Gaussian Elimination

- Parallelization steps
 - Decomposition: identify enough parallel work, but not too much
 - Assignment: load balance work among threads
 - Orchestrate: communication and synchronization
 - Mapping: which processors execute which threads
- Decomposition
 - In BLAS 2 algorithm nearly each flop in inner loop can be done in parallel, so with n² processors, need 3n parallel steps

for i = 1 to n-1 A(i+1:n,i) = A(i+1:n,i) / A(i,i) ... BLAS 1 (scale a vector) A(i+1:n,i+1:n) = A(i+1:n , i+1:n) ... BLAS 2 (rank-1 update) - A(i+1:n , i) * A(i , i+1:n)

- This is too fine-grained, prefer calls to local matmuls instead
- Need to use parallel matrix multiplication
- Assignment

- Which processors are responsible for which submatrices? 02/23/2007 CS267 DLA2



Review: BLAS 3 (Blocked) GEPP







Row and Column Block Cyclic Layout



- processors and matrix blocks are distributed in a 2d array
 - prow-by-pcol array of processors
 - brow-by-bcol matrix blocks
- pcol-fold parallelism in any column, and calls to the BLAS2 and BLAS3 on matrices of size brow-by-bcol
- serial bottleneck is eased
- prow \neq pcol possible, even desirable
- brow ≠ bcol more complicated

Distributed GE with a 2D Block Cyclic Layout

- block sizes b = bcol = brow in the algorithm and in the layout are all equal
- shaded regions indicate processors busy with computation or communication.
- unnecessary to have a barrier between each step of the algorithm, e.g. steps 9, 10, and 11 can be pipelined

Extra Slides

Recursive Algorithms

- Still uses delayed updates, but organized differently
 - (formulas on board)
- Can exploit recursive data layouts
 - 3x speedups on least squares for tall, thin matrices







- Theoretically optimal memory hierarchy performance
- See references at
 - "Recursive Block Algorithms and Hybrid Data Structures," Elmroth, Gustavson, Jonsson, Kagstrom, SIAM Review, 2004
 - <u>http://www.cs.umu.se/research/parallel/recursion/</u>

Gaussian Elimination via a Recursive Algorithm

F. Gustavson and S. Toledo

- LU Algorithm:
 - 1: Split matrix into two rectangles (m x n/2) if only 1 column, scale by reciprocal of pivot & return
 - 2: Apply LU Algorithm to the left part
 - 3: Apply transformations to right part (triangular solve $A_{12} = L^{-1}A_{12}$ and matrix multiplication $A_{22} = A_{22} - A_{21} * A_{12}$)
 - 4: Apply LU Algorithm to right part



Most of the work in the matrix multiply Matrices of size n/2, n/4, n/8, ... CS267 DLA2 38

02/23/2007

Recursive Factorizations

- Just as accurate as conventional method
- Same number of operations
- Automatic variable-size blocking
 - Level 1 and 3 BLAS only !
- Simplicity of expression
- Potential for efficiency while being "cache oblivious"
 - But shouldn't recur down to single columns!
- The recursive formulation is just a rearrangement of the pointwise LINPACK algorithm
- The standard error analysis applies (assuming the matrix operations are computed the "conventional" way).



Source: Jack Dongarra

Recursive Algorithms – Limits

- Two kinds of dense matrix compositions
- One Sided
 - Sequence of simple operations applied on left of matrix
 - Gaussian Elimination: A = L*U or A = P*L*U
 - Symmetric Gaussian Elimination: A = L*D*L^T
 - Cholesky: A = L*L^T
 - QR Decomposition for Least Squares: A = Q*R
 - Can be nearly 100% BLAS 3
 - Susceptible to recursive algorithms
- Two Sided
 - Sequence of simple operations applied on both sides, alternating
 - Eigenvalue algorithms, SVD
 - At least ~25% BLAS 2
 - Seem impervious to recursive approach?

- Some recent progress on SVD (25% vs 50% BLAS2) 02/23/2007 CS267 DLA2

Out-of-Core Performance Results for Least Squares

- Prototype code for Out-of-Core extension
- Linear solvers based on "Left-looking" variants of LU, QR, and Cholesky factorization

=

• Portable I/O interface for reading/writing ScaLA-PACK matrices

Out-of-core means matrix lives on disk; too big for main memory

Much harder to hide latency of disk

QR much easier than LU because no pivoting needed for QR

QR Factorization on 64 processors Intel Paragon



Some contributors (incomplete list)

Participants

Krste Asanovic (UC Berkeley) Zhaojun Bai (U Kentucky) Richard Barrett (U. Tenn) Michael Berry (U Tenn) Chris Bischof (ANL) Jeff Bilmes (UC Berkeley) Susan Blackford (ORNL) Soumen Chakrabarti (UC Berkeley) Tony Chan (UCLA) Chee-Whye Chin (UC Berkeley) Jaeyoung Choi (LBNL) Andy Cleary (LLNL) Ed D'Azeveda (ORNL) Jim Demmel (UC Berkeley) Inderjit Dhillon (UC Berkeley) June Donato (ORNL) Jack Dongarra (U Tenn, ORNL) Zlatko Drmač (U Hagen) Victor Eijkhout (UCLA) Jeremy Du Croz (NAG) Stan Eisenstat (Yale) Vince Fernando (NAG) John Gilbert (Xerox PARC) Ming Gu (UC Berkeley, LBL) Sven Hammarling (NAG) Mike Heath (U Illinois) Greg Henry (Intel) Dominic Lam (UC Berkeley) Steve Huss-Lederman (SRC) Bo Kågström (U Umeå) W. Kahan (UC Berkeley) Youngbae Kim (U Tenn) Rencang Li (UC Berkeley) Xiaoye Li (UC Berkeley) Joseph Liu (York) Beresford Parlett (UC Berkeley) Antoine Petitet (U Tenn) Peter Poromaa (U Umea) Roldan Pozo (U Tenn) Padma Raghavan (U Illinois) Howard Robinson (UC Berkeley) Huan Ren (UC Berkeley) Charles Romine (ORNL) Jeff Rutter (UC Berkeley) Ivan Slapničar (U Split) Dan Sorensen (Rice U) Ken Stanley (UC Berkeley) Xiaobai Sun (ANL) Bernard Tourancheau (U Tenn) Anna Tsao (SRC) Robert van de Geijn (U Texas) Henk van der Vorst (Utrecht U) Paul Van Dooren (U Illinois) Krešimir Veselić (U Hagen) David Walker (ORNL) Clint Whaley (U Tenn) Kathy Yelick (UC Berkeley)

> With the cooperation of Cray, IBM, Convex, DEC, Fujitsu, NEC, NAG, IMSL

Supported by ARPA, NSF, DOE

Upcoming related talks

- SIAM Conference on Parallel Processing in Scientific Computing
 - San Francisco, Feb 22-24
 - <u>http://www.siam.org/meetings/pp06/index.htm</u>
 - Applications, Algorithms, Software, Hardware
 - 3 Minisymposia on Dense Linear Algebra on Friday 2/24
 - MS41, MS47(*), MS56
- Scientific Computing Seminar,
 - "An O(n log n) tridiagonal eigensolver", Jonathan Moussa
 - Wednesday, Feb 15, 11-12, 380 Soda
- Special Seminar
 - Towards Combinatorial Preconditioners for Finite-Elements Problems", Prof. Sivan Toledo, Technion
 - Tuesday, Feb 21, 1-2pm, 373 Soda

Performance of ScaLAPACK QR (Least squares)

QR (Least Squares)

Efficiency = MFlops(PDGELS)/MFlops(PDGEMM) Procs Block Machine Ν Size 2000 4000 10000 Cray T3E .54 32 .61 $\mathbf{4}$ 16.46 .55 .60 .26 .47.5464IBM SP2 .51 50 4 .29 16.51 .19 .36 64.54Intel XP/S GP 4 32.61.43 16.63 Paragon .22 64.48.62 Berkeley NOW 32 .51 .77 4 32.49 .66 .71.37 .60 .72 64

Time(PDGELS)/Time(PDGEMM)							
Machine	Proce	Block	N				
		Size	2000	4000	10000		
Cray T3E	4	32	1.2	1.1			
	16		1.5	1.2	1.1		
	64		2.6	1.4	1.2		
IBM SP2	4	50	1.3				
	16		2.3	1.3			
	64		3.6	1.8	1.2		
Intel XP/S GP	4	32	1.1				
Paragon	16		1.6	1.1			
	64		3.0	1.4	1.1		
Berkeley NOW	4	32	1.3	.9			
	32		1.4	1.0	.9		
	64		1.8	1.1	.9		

Scales well, nearly full machine speed

02/23/2007

 \mathbf{C}

Performance of Symmetric Eigensolvers

Current algorithm: Faster than initial algorithm Occasional numerical instability New, faster and more stable algorithm planned

$\operatorname{Time}(\operatorname{PDSYEVX})/\operatorname{Time}(\operatorname{PDGEMM})$						
(bisection + inverse iteration $)$						
Machine	Proces	Block	N			
		Size	2000	4000		
Cray T3E	4	32	10			
	16		13	10		
	64		29	14		
IBB SP2	16	50	24			
	64		40	29		
Intel XP/S GP	16	32	22			
Paragon	64		34	20		
Berkeley NOW	16	32	20			
	32		24	52		

Time(PDSYEV)/Time(PDGEMM) (QR iteration)

Block

Size

32

50

32

32

Ν

4000

35

41

47

55

2000

35

37

57

38

58

99

193

31 35

	Marth	D
	Macume	1.1008
	Cray T3E	4
		16
		64
le	IBM SP2	16
d		64
ŭ	Intel XP/S GP	16
on	Paragon	64
	Berkeley NOW	16
		32

Initial algorithm: Numerically stable Easily parallelized Slow: will abandon

02/23/2007

Scalable Symmetric Eigensolver and SVD

The "Holy Grail" (Parlett, Dhillon, Marques) Perfect Output complexity (O(n * #vectors)), Embarrassingly parallel, Accurate



To be propagated throughout LAPACK and ScaLAPACK CS267 DLA2

02/23/2007

Performance of SVD (Singular Value Decomposition)

Have good ideas to speedup Project available!

Time(PDGESVD)/Time(PDGEMM)						
Machine	Proces	Block	N			
		Size	2000	4000		
Cray T3E	4	32	67			
	16		66	64		
	64		93	70		
IBM SP2	4	50	97			
	16		60			
	64		81			
Berkeley NOW	4	32	72			
	16		38	16		
	32		59	26		

Performance of Nonsymmetric Eigensolver (QR iteration)

Hardest of all to parallelize

Time(PDLAHQR)/Time(PDGEMM)						
Machine	Proces	Block	N			
		Size	1000	1500		
Intel XP/S MP	16	50	123	97		
Paragon						
CS267 DLA2						

02/23/2007

.

Scalable Nonsymmetric Eigensolver

- $Ax_i = \lambda_i x_i$, Schur form $A = QTQ^T$
- Parallel HQR
 - Henry, Watkins, Dongarra, Van de Geijn
 - Now in ScaLAPACK
 - Not as scalable as LU: N times as many messages
 - Block-Hankel data layout better in theory, but not in ScaLAPACK
- Sign Function
 - Beavers, Denman, Lin, Zmijewski, Bai, Demmel, Gu, Godunov, Bulgakov, Malyshev
 - $A_{i+1} = (A_i + A_i^{-1})/2 \rightarrow \text{shifted projector onto Re } \lambda > 0$
 - Repeat on transformed A to divide-and-conquer spectrum
 - Only uses inversion, so scalable
 - Inverse free version exists (uses QRD)
 - Very high flop count compared to HQR, less stable

Assignment of parallel work in GE

- Think of assigning submatrices to threads, where each thread responsible for updating submatrix it owns
 - "owner computes" rule natural because of locality
- What should submatrices look like to achieve load balance?

Computational Electromagnetics (MOM)

The main steps in the solution process are

- Fill: computing the matrix elements of A
- Factor: factoring the dense matrix A
- Solve: solving for one or more excitations b
- Field Calc: computing the fields scattered from the object

Analysis of MOM for Parallel Implementation

	Task	Work	Parallelism	Parallel Speed	
	Fill	O(n**2)	embarrassing	low	
\rightarrow	Factor	O(n**3)	moderately diff.	very high	
	Solve	O(n**2)	moderately diff.	high	
	Field Calc.	O(n)	embarrassing	high	

BLAS2 version of GE with Partial Pivoting (GEPP)

```
for i = 1 to n-1
find and record k where |A(k,i)| = max{i <= j <= n} |A(j,i)|
... i.e. largest entry in rest of column i
if |A(k,i)| = 0
exit with a warning that A is singular, or nearly so
elseif k != i
swap rows i and k of A
end if
A(i+1:n,i) = A(i+1:n,i) / A(i,i)
... each quotient lies in [-1,1]
... BLAS 1
A(i+1:n,i+1:n) = A(i+1:n, i+1:n) - A(i+1:n, i) * A(i, i+1:n)
... BLAS 2, most work in this line</pre>
```

Computational Electromagnetics – Solve Ax=b

- •Developed during 1980s, driven by defense applications
- •Determine the RCS (radar cross section) of airplane
- Reduce signature of plane (stealth technology)
- •Other applications are antenna design, medical equipment
- •Two fundamental numerical approaches:
 - •MOM methods of moments (frequency domain)
 - •Large dense matrices
 - •Finite differences (time domain)
 - •Even larger sparse matrices

Computational Electromagnetics

- Discretize surface into triangular facets using standard modeling tools

- Amplitude of currents on surface are unknowns



- Integral equation is discretized into a set of linear equations

image: NW Univ. Comp. Electromagnetics Laboratory http://nueml.ece.nwu.edu/

02/23/2007

CS267 DLA2

After discretization the integral equation has the form

$$\mathbf{A} \mathbf{x} = \mathbf{b}$$

where

A is the (dense) impedance matrix,

x is the unknown vector of amplitudes, and

b is the excitation vector.

(see Cwik, Patterson, and Scott, Electromagnetic Scattering on the Intel Touchstone Delta, IEEE Supercomputing '92, pp 538 - 542)

02/23/2007

Results for Parallel Implementation on Intel Delta

Time (hours) Task Fill (compute n^2 matrix entries) 9.20 (embarrassingly parallel but slow) Factor (Gaussian Elimination, $O(n^3)$) 8.25 (good parallelism with right algorithm) Solve $(O(n^2))$ 2.17 (reasonable parallelism with right algorithm) 0.12 Field Calc. (O(n)) (embarrassingly parallel and fast) The problem solved was for a matrix of size 48,672.

2.6 Gflops for Factor - The world record in 1991, 02/23/2007

Computational Chemistry – Ax = λ **x**

- Seek energy levels of a molecule, crystal, etc.
 - Solve Schroedinger's Equation for energy levels = eigenvalues
 - Discretize to get $Ax = \lambda Bx$, solve for eigenvalues λ and eigenvectors x
 - A and B large Hermitian matrices (B positive definite)
- MP-Quest (Sandia NL)
 - Si and sapphire crystals of up to 3072 atoms
 - A and B up to n=40000, complex Hermitian
 - Need all eigenvalues and eigenvectors
 - Need to iterate up to 20 times (for self-consistency)
- Implemented on Intel ASCI Red
 - 9200 Pentium Pro 200 processors (4600 Duals, a CLUMP)
 - Overall application ran at 605 Gflops (out of 1800 Gflops peak),
 - Eigensolver ran at 684 Gflops
 - www.cs.berkeley.edu/~stanley/gbell/index.html
 - Runner-up for Gordon Bell Prize at Supercomputing 98

LAPACK and ScaLAPACK

		LAPACK	ScaLAPACK
	Machines	Workstations,	Distributed
		Vector, SMP	Memory, DSM
	Based on	BLAS	BLAS, BLACS
	Functionality	Linear Systems	Linear Systems
		Least Squares	Least Squares
		Eigenproblems	${f Eigenproblems}$
			(less than LAPACK)
	Matrix types	Dense, band	Dense, band,
			out-of-core
	Error Bounds	Complete	A few
	Languages	F77 or C	$\mathbf{F77} \mathbf{and} \mathbf{C}$
	Interfaces to	C++, F90	HPF
	Manual?	Yes	Yes
	Where?	www.netlib.org/	www.netlib.org/
		lapack	scalapack
02/23/2	2007	CS267 DLA2	

Parallelism in ScaLAPACK

- Level 3 BLAS block operations
 - All the reduction routines
- Pipelining
 - QR Iteration, Triangular Solvers, classic factorizations
- Redundant computations
 - Condition estimators
- Static work assignment
 - Bisection

- Task parallelism
 - Sign function eigenvalue computations
- Divide and Conquer
 - Tridiagonal and band solvers, symmetric eigenvalue problem and Sign function
- Cyclic reduction
 - Reduced system in the band solver

Winner of TOPS 500 (LINPACK Benchmark)

Year	Machine	Tflops	Factor faster	Peak Tflops	Num Procs	Ν
2004	Blue Gene / L, IBM	70.7	2.0	91.8	32768	.93M
2002 2003	Earth System Computer, NEC	35.6	4.9	40.8	5104	1.04M
2001	ASCI White, IBM SP Power 3	7.2	1.5	11.1	7424	.52M
2000	ASCI White, IBM SP Power 3	4.9	2.1	11.1	7424	.43M
1999	ASCI Red, Intel PII Xeon	2.4	1.1	3.2	9632	.36M
1998	ASCI Blue, IBM SP 604E	2.1	1.6	3.9	5808	.43M
1997	ASCI Red, Intel Ppro, 200 MHz	1.3	3.6	1.8	9152	.24M
1996	Hitachi CP-PACS	.37	1.3	.6	2048	.10M
1995	Intel Paragon XP/S MP	.28	1	.3	6768	.13M

02/23/2007

CS267 DLA2

Source: Jack Dongarra (UTK)



Performance of LAPACK (n=100)

