



Linear Algebra

on a distributed environment

Jing Liu

TDB & LMB, Uppsala University

Programming of Parallel Computers, Feb 2016



Topics in LA

- Solve linear equations
- Matrices operations
 - ✱ Matrices addition/ multiplication / transformation
 - ✱ Eigenvalue/ Eigenvector
 - ✱ Transpose, projection ...
- Vector space
- ...

Scalar Vector Matrix



Loops

- LA operations are often basic building blocks in scientific applications
- Three basic types of loops
 - ✿ Perfectly parallel loops
 - ✿ Reduction loops
 - ✿ Recursive loops
 - ✿ Combination of different loops



Perfectly parallel loops

- Example $Z_m = \lambda X_m + Y_m$
for (i = 0; i < m; i ++) {
 $Z[i] = \lambda * X[i] + Y[i];$
}




- MPI Scatter and MPI Gather



Reduction loops

- Limited parallelism
- Example: Dot production $s = X \cdot Y^T$
for (i = 0; i < m; i ++) {
 s += X[i] + Y[i];
}



- MPI Reduce, MPI Allreduce



Recursive loops

- Each iteration depends on the previous one
- Hardly parallelize, “serial” loop
- Example

```
for ( i = 1; i < m; i ++ ) {  
    X[ i ] = X[ i ] + X[ i - 1 ];  
}
```



Nested loops

- Often the order of loops can be interchanged → for maximal parallelism, choose the perfectly-parallel loops as outmost, and parallelize over it.
- Example: Matrix-Vector multiplication

```
for ( i = 0; i < m; i++){  
    for ( j = 0; j < m ; j++){  
        Y[i] += A[i][j] * X[j];  
    }  
}
```



Nested loops – Alt 1

- Row-wise partition



- All processors have a copy of X, one piece of A and Y.



Nested loop – Alt. 2

- Block algorithm with 1D partition

$$\begin{matrix} Y_1 \\ Y_2 \\ Y_3 \\ Y_4 \end{matrix} = \begin{matrix} A_{00} & A_{01} & A_{02} & A_{03} \\ A_{10} & A_{11} & A_{12} & A_{13} \\ A_{20} & A_{21} & A_{22} & A_{23} \\ A_{30} & A_{31} & A_{32} & A_{33} \end{matrix} * \begin{matrix} X_0 \\ X_1 \\ X_2 \\ X_3 \end{matrix} \quad \begin{matrix} \curvearrowright \\ \downarrow \end{matrix}$$

- Step 1: Compute $Y[i] = A[i][i] * X[i]$ in process i , and then shift $X[i]$ circular one step up.
- Step 2: Compute again, in which $j=(i+1) \bmod p$, shift X circular one step up.
- Repeat, in total $(p-1)$ step



Nested loop – Alt. 2 cont.

- Non-blocking communication to shift X, before computation. MPI_Isend, MPI_Irecv, MPI_wait
- Which one is more efficient?
 - ✱ Alt. 2 is more memory efficient.
 - ✱ CPU efficient is all depends on the problem size, computer systems, implementations of MPI_functions, etc.



Nested loop – Alt. 3

- Block algorithm – 2D partition
- Processor block $\sqrt{p} * \sqrt{p}$,
- Step 1: Divide A_{mn} to $\sqrt{p} * \sqrt{p}$ blocks, X to \sqrt{p} parts
- Step 2: Processor P_{ij} get block A_{ij} and X_j , and hold $Y_i^{(j)} = \mathbf{0}$
- Step 3: P_{ij} computes $Y_i^{(j)} = A_{ij} * X_j$
- Step 4: Accumulate Y_i in each row.




Nested loop – Alt. 3

Y_0	$\leftarrow \sum_j$	$Y_0^0 = A_{00} * X_0$	$Y_0^1 = A_{01} * X_1$	$Y_0^2 = A_{02} * X_2$	$Y_0^3 = A_{03} * X_3$
Y_1	$\leftarrow \sum_j$	$Y_1^0 = A_{01} * X_0$
Y_2	$\leftarrow \sum_j$	$Y_2^0 = A_{02} * X_0$
Y_3	$\leftarrow \sum_j$	$Y_3^0 = A_{03} * X_0$	$Y_3^1 = A_{13} * X_1$	$Y_3^2 = A_{23} * X_2$	$Y_3^3 = A_{33} * X_3$

- Efficient for large matrices.
- Scalability? 2D > 1D. For many processors, 1D partition strips become so thin and communications time increases fast.



Matrix-Matrix Multi.



$$C(i,j) = \sum_k A(i,k) B(k,j)$$





More nested Loops


- Example : Matrix-Matrix Multiplication
- i and j are perfectly parallel loops, k is reduction loop

```
for ( _ = 0; _ < m; _++){  
    for ( _ = 0; _ < m ; _++){  
        for ( _ = 0 ; _ < m; _ ++){  
            C[i][j] += A[i][k] * B[k][j];  
        }  
    }  
}
```



Matrix-Matrix Multi.

- 1D partitioning – choose j as the outmost loop \rightarrow partition data column wise



- $\rightarrow C_0 = A_0 * B_{00} + A_1 * B_{10} + A_2 * B_{20} + A_3 * B_{30}$

- ...



$C = A * B$, 1D partition

- A is needed in every processor.
- Alt. 1 : Every processor has completed A,
→ Not scalable (memory?!)
- Alt. 2: Shift A around.
 - Similar idea to matrix-vector alt. 2.
 - For many processors, the stripes (block-columns) become thin and comm. overhead becomes large.



C = A*B, 2D partition

- Choose both i and j outmost.
- $\sqrt{p} * \sqrt{p}$ blocks, each processor gets one block of each matrix.
- In processor P_{ij} , compute $C_{ij} = \sum_{k=0}^{\sqrt{p}-1} A_{ik} * B_{kj}$
 - ➔ P_{ij} need all blocks A_{ik} in block row i, and B_{kj} in block column j
 - ➔ Communications needed.



$C = A * B$, 2D partition, Alt. 1

- Simple and naïve method.
- Simply distribute A in each block row, and distribute of B in each block column, using MPI_functions
 - limited scalability due to memory.
 - Bad performance if data don't fit in cache



$C = A * B$, 2D partition, Alt. 2 Cannon's Algorithm (1969)

- Shift and compute. $M * M$ mesh ($\sqrt{p} * \sqrt{p}$ blocks processors, data).
- Phase 1: shift
 - ✱ Shift the i th block row of A i steps cyclically to the left.
 - ✱ Shift the j th block column of B j steps cyclically upwards

A_{00}	A_{01}	A_{02}	A_{03}
A_{11}	A_{12}	A_{13}	A_{10}
A_{22}	A_{23}	A_{20}	A_{21}
A_{33}	A_{30}	A_{31}	A_{32}

B_{00}	B_{11}	B_{22}	B_{33}
B_{10}	B_{21}	B_{32}	B_{03}
B_{20}	B_{31}	B_{02}	B_{13}
B_{30}	B_{01}	B_{12}	B_{23}



$C = A * B$, 2D partition, Alt. 2 Cannon's Algorithm Cont.

- Phase 2: Compute and shift
- For each iteration do:
 - ✱ Compute $C_{ij} = A_{ik} * B_{kj}$ in each processor P_{ij} , where $k = (i+j+l) \bmod M$, where l is the number of iterations (start from 0).
 - ✱ Shift A one step left, B one step upwards
- In total, $M-1$ steps. We can do shift with non-blocking communication, and compute while sending.
- Read more on-line [Cannon's algorithm](#).



$C = A * B$, 2D partition, Alt. 3 Fox's Algorithm

- In total $M-1$ step.
- For each step k ($k = 0, 1, \dots, M-1$)
 - ✱ Broadcast block n of A within each block row i ($n = (i+k) \bmod M$)
 - ✱ Multiply the broadcasted block with B -block in each processor ($C_{ij} += A_{in} * B_{nj}$)
 - ✱ Shift blocks of B , one step upwards.



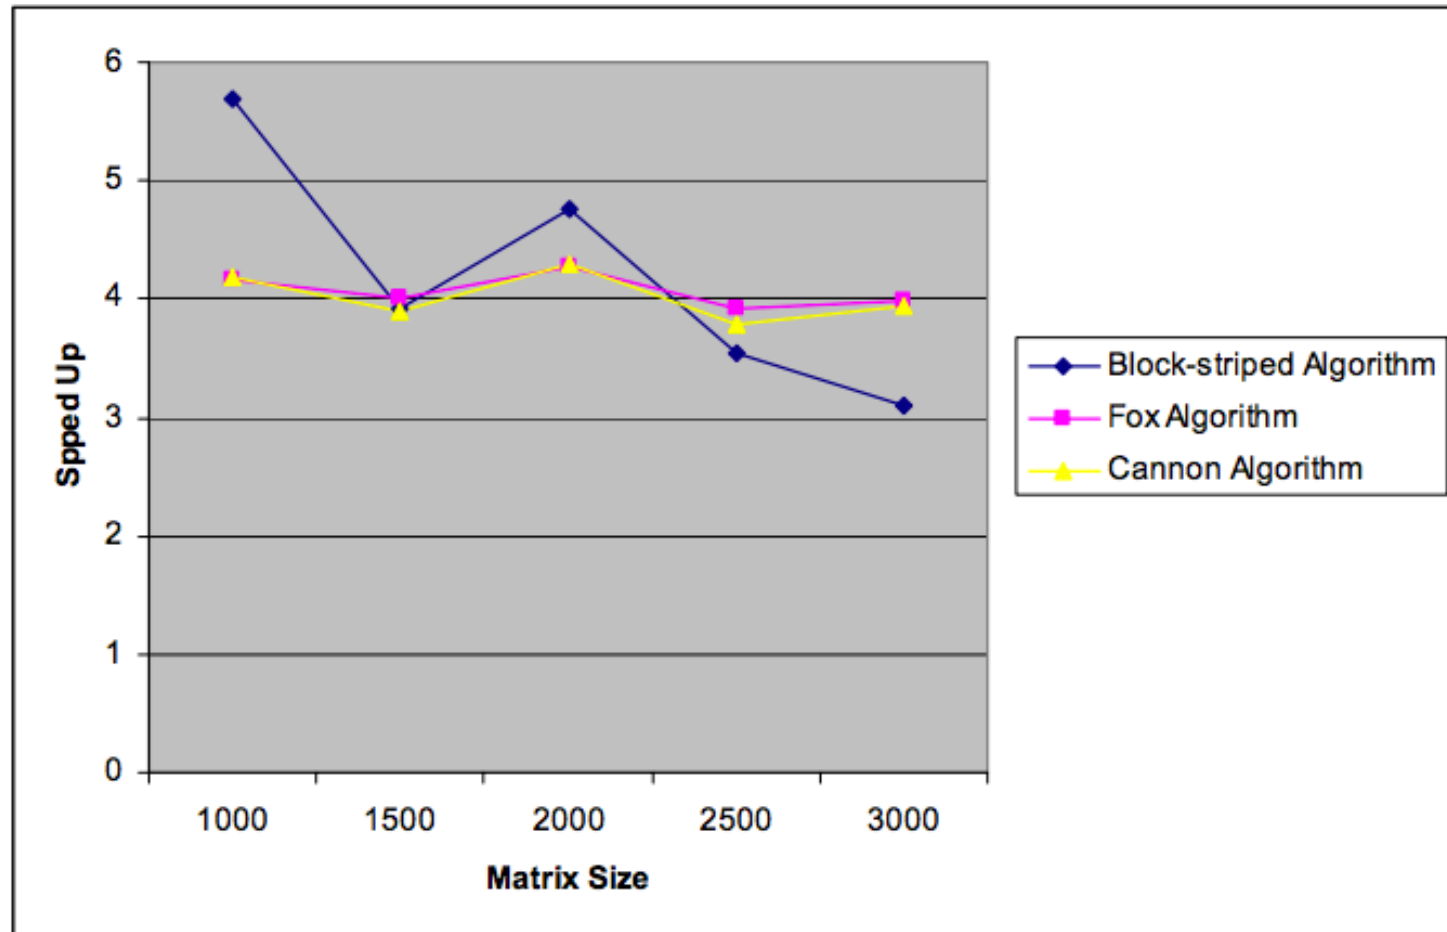
$C = A * B$, 2D partition

- Both Cannon's and fox's algorithm is scalable.
- Which is more efficient?
 - ✱ Depends on problem size, computer system, efficiency of MPI, etc



UPPSALA
UNIVERSITET

$C = A * B$, 2D partition





Assignment 1

- Dense matrix-matrix multiplication.
- Fortran/C/C++ and MPI.
- Two parameters:
 - ✱ The number of process
 - ✱ The size of matrices
- Randomly generate A and B
- Distribute data
- Implement Fox's algorithm
- Collect data and output.



Assignment 1, cont.

- Data generation (at rank 0): `srand()`, `rand()` / `CALL RANDOM_SEED()`, `CALL RANDOM_NUMBER()`
- Data distribution: use `MPI_Type_vector`, `MPI_Cart_rank`, `MPI_Isend`, `MPI_Recv`
- Data Collection: `MPI_Probe`, `MPI_Cart_coords`, `MPI_Recv`, `MPI_wait`



Assignment 1, cont.


- C structure / C++ class is helpful to make a nicer code.
 - ✱ Name space works for large project.
- Good [coding style](#) makes your code more understandable and maintainable.
- Write comments in your code to help yourself and others.
- Demo code at https://github.com/JinLi971/MPI_DEMO



Advanced Topic: Recursive loop

■ Example:

```
for ( i=1; i<n; i++){  
    X[i] += X[i-1];  
}
```





Solve linear system with Gaussian elimination

A linear equation with n unknown variables

$$a_0x_0 + a_1x_1 + \dots + a_{n-1}x_{n-1} = b$$

A finite set of n linear equations is called a *system of linear equations* or a *linear system*

$$a_{0,0}x_0 + a_{0,1}x_1 + \dots + a_{0,n-1}x_{n-1} = b_0$$

$$a_{1,0}x_0 + a_{1,1}x_1 + \dots + a_{1,n-1}x_{n-1} = b_1$$

...

$$a_{n-1,0}x_0 + a_{n-1,1}x_1 + \dots + a_{n-1,n-1}x_{n-1} = b_{n-1}$$

or in the matrix form:

$$Ax = b$$



Solve linear system with Gaussian elimination

- On the first stage of the algorithm, which is called *the Gaussian elimination*, the initial system of linear equations is transformed into an upper triangular system by the sequential elimination of unknowns:


$$Ux = c, \quad U = \begin{pmatrix} u_{0,0} & u_{0,1} & \dots & u_{0,n-1} \\ 0 & u_{1,1} & \dots & u_{1,n-1} \\ & & \dots & \\ 0 & 0 & \dots & u_{n-1,n-1} \end{pmatrix}$$

- On the second stage of the algorithm, which is called *the back substitution*, the values of the variables are calculated



Solve linear system with Gaussian elimination

- Scheme of data at the i -th iteration of the Gaussian elimination





Solve linear system with Gaussian elimination

□ Gaussian elimination:

- At step i , $0 \leq i < n-1$, of the algorithm the nonzero elements below the diagonal in column i are eliminated by replacing each row k , where $i < k \leq n-1$, with the sum of the row k and the row i multiplied by *the value* $(-a_{ki}/a_{ii})$,
- All the necessary calculations are determined by the equations:

$$\begin{aligned} a'_{kj} &= a_{kj} - (a_{ki} / a_{ii}) \cdot a_{ij}, & i \leq j \leq n-1, i < k \leq n-1, 0 \leq i < n-1 \\ b'_k &= b_k - (a_{ki} / a_{ii}) \cdot b_i, \end{aligned}$$



Solve linear system with Gaussian elimination

□ Back substitution

After the matrix of the linear system was transformed to the upper rectangular type, it becomes possible to calculate the unknown variables:

- We can solve the last equation directly, since it has only a single unknown x_{n-1} ,
- After we have determined the x_{n-1} , we can simplify the other equation by substituting the value of x_{n-1} ,
- Then the equation $n-2$ has only the single unknown x_{n-2} and can be solved and so on.

The calculations of the back substitution can be represented as follows:

$$x_{n-1} = b_{n-1} / a_{n-1,n-1},$$

$$x_i = (b_i - \sum_{j=i+1}^{n-1} a_{ij}x_j) / a_{ii}, \quad i = n-2, n-3, \dots, 0$$



Solve linear system with Gaussian elimination

- Parallel Gaussian elimination

$$\begin{aligned} a'_{kj} &= a_{kj} - (a_{ki} / a_{ii}) \cdot a_{ij}, \\ b'_k &= b_k - (a_{ki} / a_{ii}) \cdot b_i, \end{aligned}$$

- Parallel back substitution:


$$x_{n-1} = b_{n-1} / a_{n-1,n-1},$$

$$x_i = (b_i - \sum_{j=i+1}^{n-1} a_{ij}x_j) / a_{ii}, \quad i = n-2, n-3, \dots, 0$$



Solve linear system with Gaussian elimination

- Scaling issue: # processors $<$ matrix size
 - ✳ rowwise cyclic striped Decomposition





Solve linear system with Gaussian elimination

- MPI calls
 - ✱ MPI_Scatterv
 - ✱ MPI_Barrier
 - ✱ MPI_Bcast
 - ✱ MPI_Rsend
 - ✱ MPI_Recv





More Advanced Topic: BLAS

■ CPUs:

- ★ [Armadillo](#) : Matlab style, C++ coding .
- ★ [CBLAS](#) : GNU supported.
- ★ Support: AMD -> ACML; IBM -> ESSL;
Apple -> Accelerate framework; HP -> MLIB;
SUN -> Sun Performance Library,
[Intel Math Kernel Library](#)


■ GPUs: NVIDIA -> [CuBLAS](#)

OPENCL : third part support.



More Advanced Topic: Sparse Matrix

- A sparse matrix is a matrix populated primarily with zeros.
- Save sparse matrix:
 - ✿ Dictionary of keys
 - ✿ List of lists
 - ✿ Coordinate list
 - ✿ Yale format
 - ✿ Etc.





More Advanced Topic: Application using LA

- PageRank: imaging incredible large matrix
- Modern Digital imaging.
 - ✱ Video tracking: Xbox Kinect
- Genetics
- Cryptography
- Economic
- More ...



More Advanced Topic: Application using LA

- Schedule & auto tuning
 - ✱ Test cases and pre-determined
 - ✱ Dynamically schedule
- Kernel and Convolution
 - ✱ Performs in parallel computers, edges of each blocks need to fix, according to the size of the kernel