

Writing Parallel Scientific Applications

Parallel Matrix Computation

Dheeraj Bhardwaj
Department of Computer Science & Engineering
Indian Institute of Technology, Delhi -110 016 India
<http://www.cse.iitd.ac.in/~dheerajb>

Mathematical Models

- ✓ **Many Mathematical models which attempt to interpret real problems can be formulated in terms of rate of change of one or more variables as such naturally lead to partial differential equations.**
- ✓ **Methods for solution (Analytical / Numerical) of PDEs tend to be very problem dependent, so PDEs are usually solved by custom written analytical methods or/and software to take maximum advantage of particular features of given problem.**

Classification of PDEs

Following terms are often used to describe PDEs even when meaning is not so precise.

- ✓ **Hyperbolic PDEs describe time-dependent physical processes such as wave motion, that are not evolving towards steady state.**
- ✓ **Parabolic PDEs describe time-dependent physical processes, such as diffusion of heat, that are evolving toward steady state.**
- ✓ **Elliptic PDEs describe processes that have already reached steady state, and hence are time independent.**

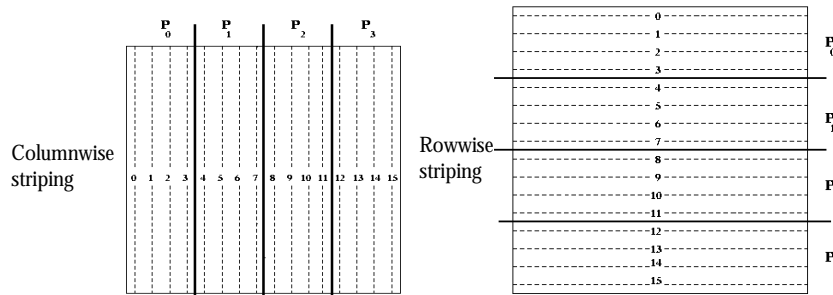
Parallel Matrix Computations

Mapping of Matrices onto Processors

Striped Partitioning

- Division into group of complete rows or columns
- Assign each processor one such group
- If each group contains an equal number of rows or columns -
Uniform partitioning

Block-striped partitioning



Mapping of Matrices onto Processors

Block-striped partitioning

- Each processor is assigned contiguous rows or columns
- Processor P_i contains columns with indices
 $(n/p)i, (n/p)i+1, \dots, (n/p)(i+1) - 1.$

$n \times n =$ size of the matrix; $p =$ No. of processors

Cyclic-striped partitioning

- Distribution of rows or columns among the processors
in wraparound manner
- Processor P_i will have rows with indices
 $i, i+p, i+2p, \dots, i+n - p.$

Mapping of Matrices onto Processors

Block-Cyclic striped partitioning

- Hybrid between block and Cyclic distribution
- Strip matrix into blocks of q rows ($q < n/p$)
- Distribution of these blocks among processors in cyclic manner
- ✓ We can partition an $n \times n$ matrix among a maximum of 'n' processors

Checkerboard Partitioning

- Division of matrix into smaller square or rectangular block or submatrices
- Distribution of such blocks or submatrices among processors
- All submatrices of same size - - Uniform partitioning

Block-checkerboard partitioning

| | | | | | | | |
|-----------------------|--------|-----------------------|--------|-----------------------|--------|-----------------------|--------|
| (0, 0) | (0, 1) | (0, 2) | (0, 3) | (0, 4) | (0, 5) | (0, 6) | (0, 7) |
| P₀ | | P₁ | | P₂ | | P₃ | |
| (1, 0) | (1, 1) | (1, 2) | (1, 3) | (1, 4) | (1, 5) | (1, 6) | (1, 7) |
| (2, 0) | (2, 1) | (2, 2) | (2, 3) | (2, 4) | (2, 5) | (2, 6) | (2, 7) |
| P₄ | | P₅ | | P₆ | | P₇ | |
| (3, 0) | (3, 1) | (3, 2) | (3, 3) | (3, 4) | (3, 5) | (3, 6) | (3, 7) |
| (4, 0) | (4, 1) | (4, 2) | (4, 3) | (4, 4) | (4, 5) | (4, 6) | (4, 7) |
| P₈ | | P₉ | | P₁₀ | | P₁₁ | |
| (5, 0) | (5, 1) | (5, 2) | (5, 3) | (5, 4) | (5, 5) | (5, 6) | (5, 7) |
| (6, 0) | (6, 1) | (6, 2) | (6, 3) | (6, 4) | (6, 5) | (6, 6) | (6, 7) |
| P₁₂ | | P₁₃ | | P₁₄ | | P₁₅ | |
| (7, 0) | (7, 1) | (7, 2) | (7, 3) | (7, 4) | (7, 5) | (7, 6) | (7, 7) |

Checkerboard Partitioning

Cyclic-checkerboard partitioning

Mapping the rows onto processors in a cyclic manner followed by the columns or vice-versa

| | | | | | | | |
|-----------------------|--------|-----------------------|--------|-----------------------|--------|-----------------------|--------|
| (0, 0) | (0, 4) | (0, 1) | (0, 5) | (0, 2) | (0, 6) | (0, 3) | (0, 7) |
| P₀ | | P₁ | | P₂ | | P₃ | |
| (4, 0) | (4, 4) | (4, 1) | (4, 5) | (4, 2) | (4, 6) | (4, 3) | (4, 7) |
| (1, 0) | (1, 4) | (1, 1) | (1, 5) | (1, 2) | (1, 6) | (1, 3) | (1, 7) |
| P₄ | | P₅ | | P₆ | | P₇ | |
| (5, 0) | (5, 4) | (5, 1) | (5, 5) | (5, 2) | (5, 6) | (5, 3) | (5, 7) |
| (2, 0) | (2, 4) | (2, 1) | (2, 5) | (2, 2) | (2, 6) | (2, 3) | (2, 7) |
| P₈ | | P₉ | | P₁₀ | | P₁₁ | |
| (6, 0) | (6, 4) | (6, 1) | (6, 5) | (6, 2) | (6, 6) | (6, 3) | (6, 7) |
| (3, 0) | (3, 4) | (3, 1) | (3, 5) | (3, 2) | (3, 6) | (3, 3) | (3, 7) |
| P₁₂ | | P₁₃ | | P₁₄ | | P₁₅ | |
| (7, 0) | (7, 4) | (7, 1) | (7, 5) | (7, 2) | (7, 6) | (7, 3) | (7, 7) |

Hybrid-checkerboard partitioning

- **Block-cyclic checkerboard partitioning**
- **Divide matrix into mxn blocks and map these blocks of size pxq in a cyclic manner**

Checkerboard Partitioning

- **Partitioned square matrix maps naturally onto a 2D square mesh of processors**
- **It is often convenient to visualize the ensemble of processors as a logical 2D mesh**
- **Unlike striping, the lowest level of granulate in checker- -boarding is one matrix per element**
- **Can exploit more concurrency then striping**

Mapping of Matrices onto Processors

- **Block partitioning leads to high granularity, which is recommended for high latency networks (EtherNet)**

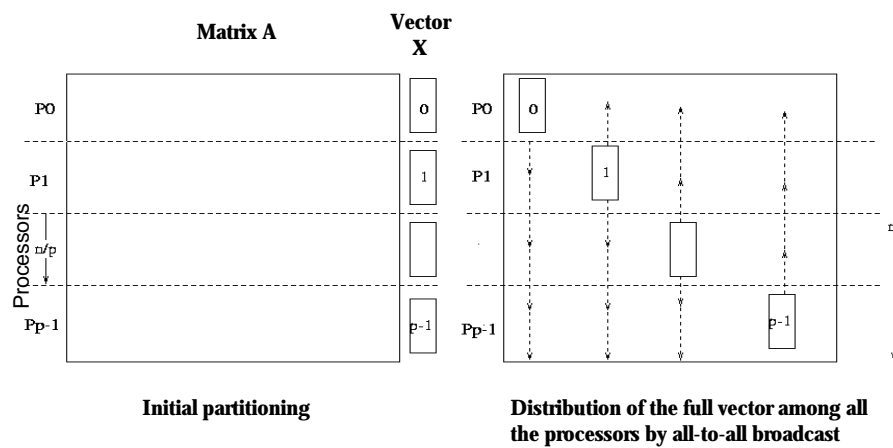
- **Cyclic distribution provides an opportunity for load balancing on network of workstations. If load balancing is not performed then overall efficiency of the code will be decided by the slowest processor OR it might lead to idling of the processors.**

- **Most of the commercial Linear Algebra libraries make use of block-cyclic distribution (e.g. ScaLAPACK & PetSc)**

Matrix - Vector Multiplication

Parallel Formulation

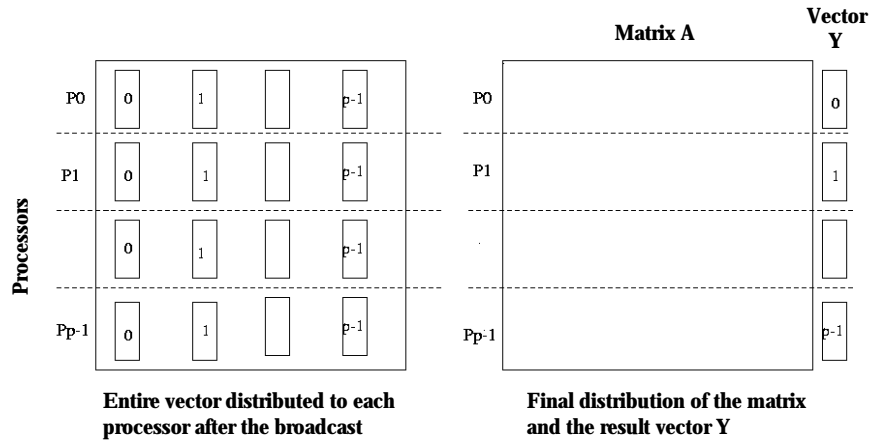
Rowwise Striping



Matrix - Vector Multiplication

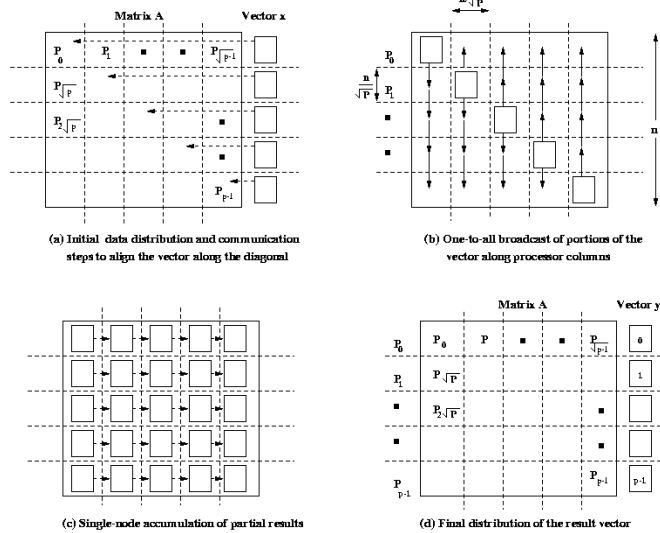
Parallel Formulation

Rowwise Striping




Matrix - Vector Multiplication

Checkerboard



Matrix - Matrix multiplication



Checkerboard partitioning

Matrix - Matrix multiplication

- **A simple Parallel Algorithm**
- **Cannon's Algorithm**
- **Fox's Algorithm**
- **The DNS algorithm**

How to Decide the Partitioning ?

- Efficient memory access pattern
- Most of the RISC based processors have hierarchical system of memory. Data access from high hierarchy is order of magnitude faster than lower hierarchy.

Checkerboard partitioning leads to efficient memory access pattern thus achieve better performance of serial part on parallel machines.

- Minimum inter-processor communication
- On conflict between memory access and inter-processor communication, we give preference to memory access.

Reason : 1. Memory is no more a problem

2. Even on shared memory computers inter-processor communication takes hundreds of cycles which is prohibitively expensive.

$$[A]\{X\} = [B]$$

Matrix Computations

Dense Matrix
Full Matrix with few
zero entries

Sparse Matrix
Majority of elements
zero

Solvers

1. Direct Solvers
2. Iterative Solvers

Why do we need direct solvers

- **Direct solvers give accuracy to the machine precision**
- **Iterative Solvers in general are very inefficient unless used with good preconditioners. Selection of preconditioner is one of the most critical phase of designing. They may lead to large overheads on matrix multiplication.**

Example : In general most of the iterative solvers act like smoothers, which reduces high frequency errors very fast but their convergence is very slow for low frequency error.

For unsymmetrical matrices arises from convective terms in fluid flows (They have directional biasness) —→ Direct methods are more reliable.

Why we don't use direct solvers ?

- **High memory requirements**
Even in the case of sparse matrices Gauss elimination or LU factorization may lead to storage requirements well beyond most of the present machines.
- **Hard to Parallelize**
The parallelization of best serial algorithm is often leads to poor efficiency on distributed parallel machines.
- **Good serial algorithm may not be good parallel algorithm**

We need a non- conventional thinking for developing efficient parallel algorithm for direct solvers.

LU Factorization

System of Linear Equations

$$Ax = b,$$

A is $n \times n$ matrix, b is given n -vector, and x is unknown solution n -vector to be determined.

To solve a linear system, we transform it into one whose solution is same but easier to compute.

One such form is LU factorization, $A = LU$, where L is unit lower triangular and U is upper triangular.

LU factorization of general nonsingular matrix A can be computed by Gaussian elimination.

LU Factorization

System of Linear Equations

If $A = LU$, then system $Ax = b$ becomes

$$Ax = LUx = b,$$

Which can be solved by forward-substitution in lower triangular system

$$Ly = b,$$

followed by back-substitution in upper triangular system

$$Ux = y.$$

In general, row interchanges (Pivoting) may be necessary for existence and numerical stability of LU factorization

In the case of irreducible, Symmetric Positive Definite (SPD) matrix, the system can be solved by Cholesky factorization.

Difficulties in efficient Parallel Implementations

Substantial parallelism inherent in Sparse direct methods, but limited success has been achieved in developing efficient general purpose parallel formulation because :

- **The amount of computation relative to the size of the system to be solved is very small.**
- **Even the modest communication might lead to poor efficiency due to poor computation to communication ratio. This becomes even severe for parallel computation on NOW / COW.**
- **Poor design criterion: Most of the serial formulation give stress on minimizing memory use and operation count. These criteria may not lead to scalable parallel formulation.**

Gaussian elimination




Fig. A typical computation in Gaussian elimination

Gaussian elimination

| | | | | | |
|-------|---|-------|---|-------|---|
| P_0 | 1 (0,1) (0,2) (0,3) (0,4) (0,5) (0,6) (0,7) | P_0 | 1 (0,1) (0,2) (0,3) (0,4) (0,5) (0,6) (0,7) | P_0 | 1 (0,1) (0,2) (0,3) (0,4) (0,5) (0,6) (0,7) |
| P_1 | 0 1 (1,2) (1,3) (1,4) (1,5) (1,6) (1,7) | P_1 | 0 1 (1,2) (1,3) (1,4) (1,5) (1,6) (1,7) | P_1 | 0 1 (1,2) (1,3) (1,4) (1,5) (1,6) (1,7) |
| P_2 | 0 0 1 (2,3) (2,4) (2,5) (2,6) (2,7) | P_2 | 0 0 1 (2,3) (2,4) (2,5) (2,6) (2,7) | P_2 | 0 0 1 (2,3) (2,4) (2,5) (2,6) (2,7) |
| P_3 | 0 0 0 (3,3) (3,4) (3,5) (3,6) (3,7) | P_3 | 0 0 0 1 (3,3) (3,4) (3,5) (3,6) (3,7) | P_3 | 0 0 0 1 (3,4) (3,5) (3,6) (3,7) |
| P_4 | 0 0 0 (4,3) (4,4) (4,5) (4,6) (4,7) | P_4 | 0 0 0 (4,3) (4,4) (4,5) (4,6) (4,7) | P_4 | 0 0 0 (4,3) (4,4) (4,5) (4,6) (4,7) |
| P_5 | 0 0 0 (5,3) (5,4) (5,5) (5,6) (5,7) | P_5 | 0 0 0 (5,3) (5,4) (5,5) (5,6) (5,7) | P_5 | 0 0 0 (5,3) (5,4) (5,5) (5,6) (5,7) |
| P_6 | 0 0 0 (6,3) (6,4) (6,5) (6,6) (6,7) | P_6 | 0 0 0 (6,3) (6,4) (6,5) (6,6) (6,7) | P_6 | 0 0 0 (6,3) (6,4) (6,5) (6,6) (6,7) |
| P_7 | 0 0 0 (7,3) (7,4) (7,5) (7,6) (7,7) | P_7 | 0 0 0 (7,3) (7,4) (7,5) (7,6) (7,7) | P_7 | 0 0 0 (7,3) (7,4) (7,5) (7,6) (7,7) |

(a) Computation:

- (i) $A[k,j] := A[k,j]/A[k,k]$ for $k < j < n$
- (ii) $A[k,k] := 1$

(b) Communication:

One-to-all broadcast of row $A[k,*]$

(c) Computation:

- (i) $A[i,j] := A[i,j] - A[i,k] * A[k,j]$ for $k < j < n$ and $k < i < n$
- (ii) $A[i,k] := 0$ for $k < i < n$

Fig. Computation load on different processors in block and cyclic-striped partitioning of an 8 X 8 matrix onto 16 processors during Gaussian elimination iteration corresponding to $k = 3$.

Gaussian elimination

| | | | |
|-------|---|-------|---|
| P_0 | 1 (0,1) (0,2) (0,3) (0,4) (0,5) (0,6) (0,7) | P_0 | 1 (0,1) (0,2) (0,3) (0,4) (0,5) (0,6) (0,7) |
| P_1 | 0 1 (1,2) (1,3) (1,4) (1,5) (1,6) (1,7) | P_1 | 0 0 0 (4,3) (4,4) (4,5) (4,6) (4,7) |
| P_2 | 0 0 1 (2,3) (2,4) (2,5) (2,6) (2,7) | P_2 | 0 0 1 (2,3) (2,4) (2,5) (2,6) (2,7) |
| P_3 | 0 0 0 (3,3) (3,4) (3,5) (3,6) (3,7) | P_3 | 0 0 0 (6,3) (6,4) (6,5) (6,6) (6,7) |
| P_4 | 0 0 0 (4,3) (4,4) (4,5) (4,6) (4,7) | P_4 | 0 0 0 (3,3) (3,4) (3,5) (3,6) (3,7) |
| P_5 | 0 0 0 (5,3) (5,4) (5,5) (5,6) (5,7) | P_5 | 0 0 0 (7,3) (7,4) (7,5) (7,6) (7,7) |
| P_6 | 0 0 0 (6,3) (6,4) (6,5) (6,6) (6,7) | P_6 | 0 0 0 (3,3) (3,4) (3,5) (3,6) (3,7) |
| P_7 | 0 0 0 (7,3) (7,4) (7,5) (7,6) (7,7) | P_7 | 0 0 0 (7,3) (7,4) (7,5) (7,6) (7,7) |

(a) Block-striped mapping

(b) Cyclic-striped mapping

Fig. Computation load on different processors in block and cyclic-striped partitioning of an 8 X 8 matrix onto 16 processors during Gaussian elimination iteration corresponding to $k = 3$.

Gaussian elimination

| | | | | | | | |
|---|-------|-------|-------|-------|-------|-------|-------|
| 1 | (0,1) | (0,2) | (0,3) | (0,4) | (0,5) | (0,6) | (0,7) |
| 0 | 1 | (1,2) | (1,3) | (1,4) | (1,5) | (1,6) | (1,7) |
| 0 | 0 | 1 | (2,3) | (2,4) | (2,6) | (2,7) | |
| 0 | 0 | 0 | (3,3) | (3,4) | (3,5) | (3,6) | (3,7) |
| 0 | 0 | 0 | (4,3) | (4,4) | (4,5) | (4,6) | (4,7) |
| 0 | 0 | 0 | (5,3) | (5,4) | (5,5) | (5,6) | (5,7) |
| 0 | 0 | 0 | (6,3) | (6,4) | (6,5) | (6,6) | (6,7) |
| 0 | 0 | 0 | (7,3) | (7,4) | (7,5) | (7,6) | (7,7) |

(a) Rowwise broadcast of $A[k,k]$ for $(k-1) < n$

| | | | | | | | |
|---|-------|-------|-------|-------|-------|-------|-------|
| 1 | (0,1) | (0,2) | (0,3) | (0,4) | (0,5) | (0,6) | (0,7) |
| 0 | 1 | (1,2) | (1,3) | (1,4) | (1,5) | (1,6) | (1,7) |
| 0 | 0 | 1 | (2,3) | (2,4) | (2,5) | (2,6) | (2,7) |
| 0 | 0 | 0 | (3,3) | (3,4) | (3,5) | (3,6) | (3,7) |
| 0 | 0 | 0 | (4,3) | (4,4) | (4,5) | (4,6) | (4,7) |
| 0 | 0 | 0 | (5,3) | (5,4) | (5,5) | (5,6) | (5,7) |
| 0 | 0 | 0 | (6,3) | (6,4) | (6,5) | (6,6) | (6,7) |
| 0 | 0 | 0 | (7,3) | (7,4) | (7,5) | (7,6) | (7,7) |

(b) $A[k,j] \leftarrow A[k,j]/A[k,k]$ for $k < j < n$

| | | | | | | | |
|---|-------|-------|-------|-------|-------|-------|-------|
| 1 | (0,1) | (0,2) | (0,3) | (0,4) | (0,5) | (0,6) | (0,7) |
| 0 | 1 | (1,2) | (1,3) | (1,4) | (1,5) | (1,6) | (1,7) |
| 0 | 0 | 1 | (2,3) | (2,4) | (2,6) | (2,7) | |
| 0 | 0 | 0 | 1 | (3,4) | (3,5) | (3,7) | |
| 0 | 0 | 0 | (4,3) | (4,4) | (4,5) | (4,7) | |
| 0 | 0 | 0 | (5,3) | (5,4) | (5,5) | (5,7) | |
| 0 | 0 | 0 | (6,3) | (6,4) | (6,5) | (6,7) | |
| 0 | 0 | 0 | (7,3) | (7,4) | (7,5) | (7,7) | |

(c) Columnwise broadcast of $A[k,j]$ for $k < j < n$

| | | | | | | | |
|---|-------|-------|-------|-------|-------|-------|-------|
| 1 | (0,1) | (0,2) | (0,3) | (0,4) | (0,5) | (0,6) | (0,7) |
| 0 | 1 | (1,2) | (1,3) | (1,4) | (1,5) | (1,6) | (1,7) |
| 0 | 0 | 1 | (2,3) | (2,4) | (2,5) | (2,6) | (2,7) |
| 0 | 0 | 0 | 1 | (3,4) | (3,5) | (3,6) | (3,7) |
| 0 | 0 | 0 | (4,3) | (4,4) | (4,5) | (4,6) | (4,7) |
| 0 | 0 | 0 | (5,3) | (5,4) | (5,5) | (5,6) | (5,7) |
| 0 | 0 | 0 | (6,3) | (6,4) | (6,5) | (6,6) | (6,7) |
| 0 | 0 | 0 | (7,3) | (7,4) | (7,5) | (7,6) | (7,7) |

(d) $A[i,j] \leftarrow A[i,j] - A[i,k]A[k,j]$ for $k < i < n$ and $k < j < n$

Fig. Various steps in the Gaussian elimination iteration corresponding to $k = 3$ for an 8×8 matrix on 64 processors of a two dimensional mesh

Cholesky Factorization

A : $n \times n$ Symmetric and Positive Definite (SPD) matrix.

Cholesky factorization : $A = L L^T$

Where L is lower triangular matrix with positive diagonal entries.

Given Cholesky factorization, linear system $Ax = b$, can be solved by successive forward and backward substitutions

$$Ly = b \quad \text{and} \quad L^T x = y$$

Cholesky Factorization


Algorithm Features

Features of Cholesky make it attractive for SPD matrices :

- All n square roots involved are of positive numbers, so algorithms is well defined.
- No pivoting is required for numerical stability.
- Only lower triangular portion of A is accessed, and hence upper triangular portion need not to be stored.
- Factor L is computed in place, overwriting lower triangle of A
- Only $n^3/6$ multiplications and similar number of additions are required.

Cholesky Factorization

Parallel Algorithm



Cholesky Factorization

Parallel Algorithm

```
for k = 1 to min(i,j) - 1
  rcv  ak,j
  rcv  ai,k
end   ai,j = ai,j - ai,k * aj,k
if i = j then
  ai,i =  $\sqrt{a_{i,i}}$ 
  broadcast ai,j to task (k,i) and (i,k), k = i+1, ..., n
else if i < j then
  rcv  ai,j
  ai,j = ai,j / ai,i
  broadcast ai,j to task (k,j) k = i+1, ..., j
  else
  rcv  aj,j
  ai,j = ai,j / aj,j
  broadcast ai,j to task (i,k) k = j+1, ..., i
end
```

Cholesky Implementations

Three choices of index for outer loop yield different algorithms with memory access patterns

- **Row-Cholesky:** With i in outer loop, inner loops solve triangular system for each new row in term of previous computed rows.
- **Column-Cholesky:** with j in outer loop, inner loops compute matrix-vector product that gives effect of previously computed columns on column currently being computed.
- **Checkerboard-Cholesky:** With k in outer loop, inner loops apply current columns as *rank-1* update to remaining unreduced submatrix.