



## Plan of the lecture:

# Parallel algorithms for Scientific Computations

## On parallel factorization methods, communication-avoiding methods, pipelined iterative methods

Maya Neytcheva

PASC, November 2020

- ▶ Techniques to gain parallelism in matrix factorization methods
  - ▶ odd-even elimination (cyclic reduction)
  - ▶ odd-even simultaneous elimination
  - ▶ divide-and-conquer method
  - ▶ multi-frontal methods
- ▶ Communication avoiding techniques
- ▶ Pipelined versions of Krylov subspace iteration methods
- ▶ Extended Krylov subspace methods



## LU factorization, again

Recall (in general):

```

for k := 1 : n
  c = akk-1
  for j = k + 1 : n
    akj = akj * c
  for i = 1 : n and i ≠ k
    aij = aij - aik * akj

```

Note that the basic update step is the linked triad  $a_{.j} - a_{kj}a_{.k}$ .

To solve  $Ax = b$  use the (block-)matrix factorization  $A = LU$ :

$Ly = b$  (forward substitution)

and

$Ux = y$  (back substitution).

The problem is the recursion (purely sequential process)



How to gain parallelism in matrix factorization methods, in particular, for sparse matrices



## Solutions with triangular matrices, again

Solution of  $Ax = b$ ,

where  $A$  has been already factorized as  $A = LU$  or  $A = LDU$ . The matrices  $L = \{l_{i,j}\}$  and  $U = \{u_{i,j}\}$  are of the form

$$L = \begin{bmatrix} 1 & 0 & \dots & 0 \\ a_{21}^{(1)}/a_{21}^{(1)} & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ a_{n1}^{(1)}/a_{11}^{(1)} & a_{n2}^{(2)}/a_{22}^{(2)} & \dots & 1 \end{bmatrix} \quad \text{and} \quad U = \begin{bmatrix} a_{11}^{(1)} & a_{12}^{(1)} & \dots & a_{1n}^{(1)} \\ 0 & a_{22}^{(2)} & \dots & a_{2n}^{(2)} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_{nn}^{(n)} \end{bmatrix}$$

To compute  $x$ , we must perform two steps:

*forward substitution*:  $Lz = b$ , i.e.,

$$z_1 = b_1$$

$$z_i = b_i - \sum_{k=1}^{i-1} l_{i,k} z_k, \quad i = 2, 3, \dots, n. \quad (1)$$

*backward substitution*:  $Ux = z$ , i.e.,

$$x_n = z_n$$

$$x_i = z_i - \sum_{k=i+1}^n u_{i,k} x_k, \quad i = n-1, n-2, \dots, 1. \quad (2)$$

For illustration purposes we consider (block-) tridiagonal matrices and their solution in parallel.



## Odd-even elimination (cyclic reduction) method

## Cyclic reduction in parallel




Figure: Odd-even elimination (reduction) method

- ▶ Eliminate in row  $i$  (independently,  $n$  fine grain tasks)
- ▶ Communicate to 'adjacent' equations  $\log n$  stages, increasing comm distance
- ▶ In this case it is better to assign adjacent rows to each process, instead of cyclic mapping.
- ▶ Local tridiagonal systems (of size  $n/p$  can be handled by serial cyclic reduction, LU or IC).





## Odd-even elimination (reduction) method

$$\begin{bmatrix} a_{11} & a_{12} & & & \cdots \\ a_{21} & a_{22} & a_{23} & & \cdots \\ & a_{32} & a_{33} & a_{43} & \cdots \\ & & a_{43} & a_{44} & a_{45} & \cdots \\ & & \vdots & \ddots & \ddots & \ddots \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ \vdots \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \\ b_3 \\ b_4 \\ \vdots \end{bmatrix}$$

Eliminate  $x_1$  from equation 1.  
Eliminate  $x_3$  from equations 2 and 4:

$$x_3 = \frac{1}{a_{23}}(b_2 - a_{22}^{(1)})x_2$$

$$(a_{32} - \frac{a_{33}a_{22}}{a_{23}})x_2 + a_{34}x_4 = b_2^{(1)}$$




## Odd-even elimination (reduction) method

Advantages: No fill-in,  $\log_2 n$  stages.  
Drawbacks: the system size is reduced, thus, insufficient load and idleness of processors.


Can be done block-wise.



## Odd-even simultaneous elimination method



## Odd-even reordering (nested dissection for triangular systems)



$$w_0 = [1, 3, 5, 7, 9, 11, 13, 15, 2, 4, 6, 8, 10, 12, 14]$$

$$w = [1, 5, 9, 13, 3, 7, 11, 15, 2, 6, 10, 14, 4, 8, 12]$$

Figure: Odd-even elimination (reduction) method.





## Divide and conquer approach when solving




Figure: Divide-and-conquer algorithm

As a first step, all off-diagonal entries in all diagonal blocks are eliminated in parallel. The side effect is that a fill-in in the off-diagonal blocks occur. However, the fill-in entries in the off-diagonal blocks reside in only one column.



## Multifrontal elimination

(Domain decomposition, matrix-graph based approach)

The multifrontal methods utilizes the structure to find parts to be treated concurrently.




Figure: Matrix for two-way factorization



## Multifrontal elimination, using nested dissection





Figure: Plus-shaped domain and matrix for four-way factorization

Solve on each diagonal block in parallel using four processors. Final block corresponding to the reduced system for the middle part of the structure we solve the corresponding linear system of equations after Again perform the back substitution in parallel for the four separate parts of the structure.



Ideal structure: the 'arrow' matrix:

$$\begin{bmatrix} A_{11} & & & A_{14} \\ & A_{22} & & A_{24} \\ & & A_{33} & A_{34} \\ A_{41} & A_{42} & A_{43} & A_{44} \end{bmatrix}$$



In practice: much more complicated structures for general sparse matrices.





# Multifrontal

# Multifrontal: Distributed memory setting



Constraints:

- ▶ Enough work per node
- ▶ Dependencies defined by the tree
- ▶ PEs must get roughly the same amount of work
- ▶ Minimize communication and idling of processors



# Multifrontal: What do we have per node?

# Multifrontal:


Frontal matrix:

$$A1 = \begin{bmatrix} \text{F1} & B1 \\ C1 & D1 \end{bmatrix}$$

$$A2 = D2 - C1 * \text{inv}(F1) * B1$$

## Nested Dissection Partitioning Method

- Recursive bisection to partition into >2 partitions
- Use *nested dissection*!





## Communication-avoiding, dense matrices

Idea: *Tall skinny QR, TSQR*

$$A = \begin{bmatrix} A_0 \\ A_1 \\ A_2 \\ A_3 \end{bmatrix} = \begin{bmatrix} Q_0^* R_{00} \\ Q_1^* R_{01} \\ Q_2^* R_{02} \\ Q_3^* R_{03} \end{bmatrix} = \begin{bmatrix} Q_0 & & & \\ & Q_1 & & \\ & & Q_2 & \\ & & & Q_3 \end{bmatrix} \begin{bmatrix} R_{00} \\ R_{01} \\ R_{02} \\ R_{03} \end{bmatrix}$$

$$\begin{bmatrix} R_{00} \\ R_{01} \\ R_{02} \\ R_{03} \end{bmatrix} = \begin{bmatrix} Q_{10}^* R_{10} \\ \dots \\ Q_{11}^* R_{11} \end{bmatrix} \dots$$

J. Demmel, LG, M. Hoemmen, J. Langou, 08 References: Golub, Plemmons, Sameh 88, Pothen, Raghavan, 89,

Da Cunha, Becker, Patterson, 02

## Communication avoiding techniques



## Communication-avoiding, dense matrices, TSQR

## Communication-avoiding

Handwritten derivation of TSQR decomposition:

$$A(n, m) = \begin{bmatrix} A_1(m_1, m) \\ A_2(m_2, m) \end{bmatrix} = \begin{bmatrix} Q_1 & R_1 \\ 0 & R_2 \end{bmatrix}$$

where  $Q_1$  is  $n \times m$ ,  $R_1$  is  $m \times m$ ,  $Q_2$  is  $m_2 \times m$ , and  $R_2$  is  $m \times m$ .

The diagram shows the block structure of the matrices and their dimensions. The total height of the matrix is  $n$ , and the total width is  $m$ . The first block  $A_1$  has height  $m_1$  and width  $m$ . The second block  $A_2$  has height  $m_2$  and width  $m$ . The decomposition shows  $Q_1$  (height  $m_1$ , width  $m$ ) and  $R_1$  (width  $m$ ) in the first row, and  $0$  (width  $m$ ) and  $R_2$  (width  $m$ ) in the second row. The total height of the matrix is  $n$ , and the total width is  $m$ .

Communication-avoiding versions exist for many algorithms:

- ▶ QR, SVD
- ▶ LU, Cholesky
- ▶ Sym, nonsym eigenproblems

Attaining the lower bounds of number-of-words-moved and number-of-messages





# Communication-avoiding ideas for iterative solvers

Two ideas:

- ▶ Decrease the number of iterations to decrease the number of global communications
- ▶ Increase arithmetic intensity - compute sparse matrix-set of vectors product


Enlarged Krylov subspace methods  
2014-2016, 2019  
Based on replacing the Krylov subspace with a larger subspace, specifically the enlarged Krylov subspace, leading to better convergence.



## Enlarged Krylov methods, convergence

### Enlarged Krylov subspaces

- Partition the matrix into  $t$  domains
- Split the initial residual into  $t$  vectors corresponding to the  $t$  domains



- Generate  $t$  new basis vectors, obtain an enlarged Krylov subspace  $K_{t,k}(A, r_0) = \text{span}\{T(r_0), AT(r_0), \dots, A^{k-1}T(r_0)\}$
- Search for the solution of the system  $Ax = b$  in  $K_{t,k}(A, r_0)$

Laura Grigori, Sophie Moufawad, Frederic Nataf  
Enlarged Krylov Subspace Conjugate Gradient Methods for Reducing Communication  
SIAM Journal on Matrix Analysis and Applications 37(2):744-773, 2016  
\* Convergence at least as fast as the standard CG.





## Pipelined unpreconditioned CG

Given  $A$ ,  $\mathbf{b}$  and an initial guess  $\mathbf{x}^{(0)}$ .

### Pipelined versions of Krylov subspace iteration methods

```

1    $\mathbf{r}^{(0)} = \mathbf{b} - A\mathbf{x}^{(0)}, \mathbf{w}^{(0)} = A\mathbf{x}^{(0)}$ 
2   for  $k = 0, 1, \dots$  until convergence
3        $\gamma_k = (\mathbf{r}^{(k)}, \mathbf{r}^{(k)})$ 
4        $\delta = (\mathbf{w}^{(k)}, \mathbf{r}^{(k)})$ 
5        $\mathbf{q} = A\mathbf{w}^{(k)}$ 
6       if  $k > 0$ ,
7            $\beta_k = \gamma_k / \gamma_{k-1}; \alpha_k = \gamma_k / (\delta - \beta_k \gamma_k / \alpha_{k-1})$ 
8       else
9            $\beta_k = 0; \alpha_k = \gamma_k / \delta$ 
10      endif
11       $\mathbf{z}^{(k)} = \mathbf{q} + \beta_k \mathbf{z}^{(k-1)}$ 
12       $\mathbf{s}^{(k)} = \mathbf{w}^{(k)} + \beta_k \mathbf{s}^{(k-1)}$ 
13       $\mathbf{p}^{(k)} = \mathbf{r}^{(k)} + \beta_k \mathbf{p}^{(k-1)}$ 
14       $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha_k \mathbf{p}^{(k)}$ 
15       $\mathbf{r}^{(k+1)} = \mathbf{r}^{(k)} - \alpha_k \mathbf{s}^{(k)}$ 
16       $\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} - \alpha_k \mathbf{z}^{(k)}$ 
17  endfor

```



## Pipelined preconditioned CG

```

1    $\mathbf{r}^{(0)} = \mathbf{b} - A\mathbf{x}^{(0)}, \mathbf{u}^{(0)} = M^{-1}\mathbf{r}^{(0)}, \mathbf{w}^{(0)} = A\mathbf{u}^{(0)}$ 
2   for  $k = 0, 1, \dots$  until convergence
3        $\gamma_k = (\mathbf{r}^{(k)}, \mathbf{u}^{(k)})$ 
4        $\delta = (\mathbf{w}^{(k)}, \mathbf{u}^{(k)})$ 
5        $\mathbf{m} = M^{-1}\mathbf{w}^{(k)}$ 
6        $\mathbf{n} = A\mathbf{m}$ 
7       if  $k > 0$ ,
8            $\beta_k = \gamma_k / \gamma_{k-1}; \alpha_k = \gamma_k / (\delta - \beta_k \gamma_k / \alpha_{k-1})$ 
9       else
10           $\beta_k = 0; \alpha_k = \gamma_k / \delta$ 
11      endif
12       $\mathbf{z}^{(k)} = \mathbf{n} + \beta_k \mathbf{z}^{(k-1)}$ 
13       $\mathbf{q}^{(k)} = \mathbf{m} + \beta_k \mathbf{q}^{(k-1)}$ 
14       $\mathbf{s}^{(k)} = \mathbf{w}^{(k)} + \beta_k \mathbf{s}^{(k-1)}$ 
15       $\mathbf{p}^{(k)} = \mathbf{u}^{(k)} + \beta_k \mathbf{p}^{(k-1)}$ 
16       $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha_k \mathbf{p}^{(k)}$ 
17       $\mathbf{r}^{(k+1)} = \mathbf{r}^{(k)} - \alpha_k \mathbf{s}^{(k)}$ 
18       $\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} - \alpha_k \mathbf{q}^{(k)}$ 
19       $\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} - \alpha_k \mathbf{z}^{(k)}$ 
20  endfor

```

Convergence rate, computational cost per iteration

## Pipelined preconditioned Conjugate Residuals

```

1    $\mathbf{r}^{(0)} = \mathbf{b} - A\mathbf{x}^{(0)}, \mathbf{u}^{(0)} = M^{-1}\mathbf{r}^{(0)}, \mathbf{w}^{(0)} = A\mathbf{u}^{(0)}$ 
2   for  $k = 0, 1, \dots$  until convergence
3        $\mathbf{m} = M^{-1}\mathbf{w}^{(k)}$ 
4        $\gamma_k = (\mathbf{w}^{(k)}, \mathbf{u}^{(k)})$ 
5        $\delta = (\mathbf{m}^{(k)}, \mathbf{u}^{(k)})$ 
6        $\mathbf{n} = A\mathbf{m}$ 
7       if  $k > 0$ ,
8            $\beta_k = \gamma_k / \gamma_{k-1}; \alpha_k = \gamma_k / (\delta - \beta_k \gamma_k / \alpha_{k-1})$ 
9       else
10           $\beta_k = 0; \alpha_k = \gamma_k / \delta$ 
11      endif
12       $\mathbf{z}^{(k)} = \mathbf{n} + \beta_k \mathbf{z}^{(k-1)}$ 
13       $\mathbf{q}^{(k)} = \mathbf{m} + \beta_k \mathbf{q}^{(k-1)}$ 
14       $\mathbf{p}^{(k)} = \mathbf{u}^{(k)} + \beta_k \mathbf{p}^{(k-1)}$ 
15       $\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \alpha_k \mathbf{p}^{(k)}$ 
16       $\mathbf{u}^{(k+1)} = \mathbf{u}^{(k)} - \alpha_k \mathbf{q}^{(k)}$ 
17       $\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} - \alpha_k \mathbf{z}^{(k)}$ 
18  endfor

```

Convergence rate, computational cost per iteration







**s-step Krylov subspace iteration methods**

Erin Carson  
The adaptive s-step Conjugate Gradient Method  
SIAM J. MATRIX ANAL. APPL., Vol. 39 2018, pp. 1318–1338  
See the paper.

Erin C. Carson, Miroslav Rozložnik, Zdenek Strakos, Petr Tichy, Miroslav Tuma

*The numerical stability analysis of pipelined conjugate gradient methods: historical context and methodology*  
SIAM J. Sci. Comput., 40(5), A3549–A3580, 2018

**Conclusions:** "... We give an expression for the residual gap which is applicable to any such variant, including the pipelined CG method of Ghysels and Vanroose [18]. We show that, similar to 3-term recurrence CG, any variant of CG which uses an auxiliary recurrence for  $A_{pi}$  in updating the residual vector can suffer from amplifications of local rounding errors due to large residual oscillations.

... modifications borne out of efforts to minimize time per iteration. While an important piece of the puzzle, one must be cautious about using time per iteration as a metric of efficiency. Alone, this fails to capture the whole picture of the method's effectiveness within the context of a scientific application. If the modifications



Bottom line: be precautious when trading numerical efficiency for parallel efficiency!  
There should be a theoretical support for the algorithmic changes to ensure the quality of the numerically computed solutions.

