

---

## HPC for numerical methods and data analysis

*Fall Semester 2024*

*Prof. Laura Grigori*

*Assistant: Mariana Martínez Aguilar*

**Session 6 – October 15, 2024**

---

## TSQR Factorization and matrix multiplication

### Exercise 1 TSQR

Remember that communication refers to messages between processors. In the recent years we've seen trends causing floating point to become faster than communication. This is why it's important to minimize communication when dealing with parallelism. The TSQR, "Tall Skinny QR" algorithm is a communication avoiding algorithm for matrices with many more rows than columns. In this exercise we are going to assume we're using  $P = 4$  processors and the matrix we want to factorize is  $A \in \mathbb{R}^{m \times n}$  with  $m \geq n$ . The computation can be expressed as a product of intermediate orthonormal factors in a binary tree like structure. We scatter row wise our matrix  $A$  along 4 processors  $A_1, A_2, A_3, A_4 \in \mathbb{R}^{m/4 \times n}$ . At the leaves of the binary tree, we perform in parallel 4 local QR factorizations to get  $A_1 = \hat{Q}_1^{(2)} \hat{R}_1^{(2)}$ ,  $A_2 = \hat{Q}_2^{(2)} \hat{R}_2^{(2)}$ ,  $A_3 = \hat{Q}_3^{(2)} \hat{R}_3^{(2)}$ ,  $A_4 = \hat{Q}_4^{(2)} \hat{R}_4^{(2)}$ . Here  $\hat{Q}_l^{(2)} \in \mathbb{R}^{m/4 \times m/4}$  and  $\hat{R}_l^{(2)} \in \mathbb{R}^{m/4 \times n}$ . In block structure we get

$$\begin{bmatrix} A_1 \\ A_2 \\ A_3 \\ A_4 \end{bmatrix} = \begin{bmatrix} \hat{Q}_1^{(2)} & & & \\ & \hat{Q}_2^{(2)} & & \\ & & \hat{Q}_3^{(2)} & \\ & & & \hat{Q}_4^{(2)} \end{bmatrix} \begin{bmatrix} \hat{R}_1^{(2)} \\ \hat{R}_2^{(2)} \\ \hat{R}_3^{(2)} \\ \hat{R}_4^{(2)} \end{bmatrix}$$

Recall that  $\hat{R}_l^{(2)} \in \mathbb{R}^{m/4 \times n}$  are tall and skinny upper triangular matrices, hence they can be written as

$$\hat{R}_l^{(2)} = \begin{bmatrix} R_l^{(2)} \\ 0 \end{bmatrix},$$

where  $R_l^{(2)} \in \mathbb{R}^{n \times n}$ . In the second level of the binary tree we combine the upper triangular matrices  $R_1^{(2)}$  with  $R_2^{(2)}$  and  $R_3^{(2)}$  with  $R_4^{(2)}$  to get block structured matrices. We perform the QR factorization in parallel to get the following structured matrices

$$\begin{bmatrix} R_1^{(2)} \\ R_{(2)}^{(2)} \end{bmatrix} = \begin{bmatrix} \hat{Q}_{11}^{(1)} & \hat{Q}_{12}^{(1)} \\ \hat{Q}_{21}^{(1)} & \hat{Q}_{22}^{(2)} \end{bmatrix} \begin{bmatrix} R_1^{(1)} \\ 0 \end{bmatrix} \quad \begin{bmatrix} R_3^{(2)} \\ R_{(4)}^{(2)} \end{bmatrix} = \begin{bmatrix} \hat{Q}_{33}^{(1)} & \hat{Q}_{34}^{(1)} \\ \hat{Q}_{43}^{(1)} & \hat{Q}_{44}^{(2)} \end{bmatrix} \begin{bmatrix} R_3^{(1)} \\ 0 \end{bmatrix}.$$

This can be written as

$$\begin{bmatrix} \hat{R}_1^{(2)} \\ \hat{R}_2^{(2)} \\ \hat{R}_3^{(2)} \\ \hat{R}_4^{(2)} \end{bmatrix} = \begin{bmatrix} \hat{Q}_{11}^{(1)} & & \hat{Q}_{12}^{(1)} & & & \\ & I & & & & \\ \hat{Q}_{21}^{(1)} & & \hat{Q}_{22}^{(1)} & & & \\ & & & I & & \\ & & & \hat{Q}_{33}^{(1)} & & \hat{Q}_{34}^{(1)} \\ & & & & I & \\ \hat{Q}_{43}^{(1)} & & & & \hat{Q}_{44}^{(1)} & \\ & & & & & I \end{bmatrix} \begin{bmatrix} R_1^{(1)} \\ 0 \\ R_3^{(1)} \\ 0 \end{bmatrix}.$$

Finally at the root of the tree we compute the last QR factorization

$$\begin{bmatrix} R_1^{(1)} \\ R_3^{(1)} \end{bmatrix} = \begin{bmatrix} \hat{Q}_{11}^{(0)} & \hat{Q}_{13}^{(0)} \\ \hat{Q}_{31}^{(0)} & \hat{Q}_{33}^{(0)} \end{bmatrix} \begin{bmatrix} R_1^0 \\ 0 \end{bmatrix}$$

Which can be written as

$$\begin{bmatrix} R_1^{(1)} \\ 0 \\ R_{(3)}^{(1)} \\ 0 \end{bmatrix} = \begin{bmatrix} \hat{Q}_{11}^{(0)} & & \hat{Q}_{13}^{(0)} & \\ & I & & \\ \hat{Q}_{31}^{(0)} & & \hat{Q}_{33}^{(0)} & \\ & & & I \end{bmatrix} \begin{bmatrix} R_1^{(0)} \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

$\hat{Q}$  is represented implicitly as the product of the intermediate orthogonal matrices.

- Suppose that you are given the implicit representation of  $Q$  as a product of orthogonal matrices. This representation is a tree of sets of Householder vectors,  $\{\hat{Q}_{r,k}^d\}$ . Here,  $k$  indicates the processor number (both where it is computed and where it is stored), and  $d$  indicates the level in the tree. How would you get  $\hat{Q}$  explicitly? We only need the “thin”  $\hat{Q}$ , meaning only its first  $n$  columns. Note that we can do this by applying the intermediate factors  $\{\hat{Q}_{r,k}^d\}$  to the first  $n$  columns of the  $m \times m$  identity matrix. Write a Python script using MPI that does this (assume you are using 4 processors). (*Hint: looking at the graphical representation of parallel TSQR might help.*)
- Optional:* consider more processors. How would you change your code? Do you have a restriction on the number of processors you can use?
- Suppose you want to use the QR factorization to solve the following least squares problem

$$\min_x \|Ax - b\|_2, \quad (1)$$

where  $x \in \mathbb{R}^n, A \in \mathbb{R}^{m \times n}, b \in \mathbb{R}^m$ . Write down the algebra for solving this problem using the QR factorization.

d) Assume you are in the situation where you have an implicit representation of  $Q$  as a product of orthogonal matrices and  $R$  from the TSQR algorithm. This means that each processor ends up with different information. Write down the algebra from solving the least squares problem 1 using the implicit representation of  $Q$ .

### **Exercise 2 Matrix multiplication**

Recall the SUMMA algorithm for parallelizing matrix matrix multiplication,  $A, B \in \mathbb{R}^{n \times n}$  to get  $C = AB$ . Implement this algorithm in  $P = 4, 16$  processors. Compare this to numpy's standard matrix matrix multiplication.