

## Set 9 - Sparse Linear Algebra with MPI

Issued: November 23, 2018

Hand in (optional): December 03, 2018 23:59

### Question 1: CSR format

Sparse matrices have a limited number of nonzero elements. Specialized formats are used for efficient storage and operations on them. The CSR format (compressed sparse row) represents a matrix  $A_{ij}$  by three one-dimensional arrays:

- $A[k]$ , nonzero coefficients in the row-major order,  $0 \leq k < nnz$ ;
- $K[i]$ , extents of rows, row  $i$  consists of  $A[K[i]:K[i+1]]$ ;
- $J[k]$ , column indices.

Write down a representation of the following matrix in the CSR format

$$A = \begin{bmatrix} 2 & -1 & 0 & -1 \\ -1 & 2 & -1 & 0 \\ 0 & -2 & 4 & -2 \\ -1 & 0 & -1 & 2 \end{bmatrix}.$$

From that representation, compute the product  $Au$  with column  $u = [0, 1, 2, 3]^T$ .

### Question 2: Matrix-vector product with MPI

Sparse matrices and corresponding vectors can be distributed among multiple processors. Consider the matrix-vector product

$$b = Au$$

with a square matrix  $A \in \mathbf{R}^{n \times n}$  and vectors (columns)  $u, b \in \mathbf{R}^n$ . Each of  $p$  processors stores  $n/p$  rows of the matrix, the corresponding chunks of vector  $b$  and the same chunks of vector  $u$ . In the worst case, the communication pattern is all-to-all. However, for a banded matrix (e.g. tridiagonal) the communication is required only between neighbouring chunks. For instance, such matrices result from discretization of PDEs.

Here you will implement the sparse matrix-vector product and apply it for the solution of the two-dimensional diffusion equation. The skeleton code provides structures for storing the matrix and converting between local and global indices. Distributed CSR format uses global indices for columns and local indices for extents of rows. The algorithm should involve the following stages:

- traverse rows of the matrix, multiply the elements stored locally and collect global indices of columns that require communication; use `GlbToLoc()` to convert the global indices;
  - use `GlbToRank()` to find the ranks of processors storing the required columns;
  - use `MPI_Allreduce()` to compute the number of messages that every processor needs to receive;
  - send and receive the indices of columns from other processors;
  - send and receive the corresponding elements of vector  $u$ ;
  - add the received data to the product.
- a) Implement the matrix-vector product in function `Mul()`. For testing use the provided discretization of the Laplace operator applied for solving the two-dimensional diffusion equation

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}$$

as well as other operators found in `op.h`. For plotting the results use `make plot` which calls the Python script `plot`.

- b) Implement function `WriteMpi()` in `io.h` with either MPI I/O or gathering the data on one processor.
- c) Report the weak scaling on Euler.
- d) (Optional) Add a new linear operator of your choice to `op.h` and apply for discretization of a different equation (e.g. advection equation). You can combine multiple operators in one time step to discretize non-linear equations.