COSC 462 Homework 3: Matrix Matrix Multiply with MPI

Piotr Luszczek

September 21, 2016

The objective of HW3 is to implement parallel matrix-matrix multiplication in MPI:

$$C = A \times B \qquad A,B,C \in \mathbb{R}^{N\times N} \tag{1}$$

The MPI ranks should form a square grid of processes $P = S\times S\,(\mathbb{N} \ni S = \sqrt{P})$. In other words, your code only needs to only work for $2\times2=4$ processes, $3\times3=9$ processes, $4\times4=16$ processes, and so on.

The matrices are square, of size N by N. With $N=2$, there will be 4 processes and the three matrices (A, B, and C) will be 2 by 2 each. With $N=3$, there will be 9 processes and the three matrices (A, B, and C) will be 3 by 3 each. And so on.

The distribution of matrix elements is fixed and is exactly the same for all three matrices: each MPI process is mapped to exactly one element of matrix A, B, and C. Rank 0 is mapped to element 0,0; rank 1 – to 1,0 and so forth. The mapping code from a rank to row and column is:

```c
void
rank2rowcol(int N, int rank, int *row, int *col) {
  *row = rank % N;
  *col = rank / N;
}
```

The input matrix data (matrices A and B) are read from a file on rank 0 and the result data (matrix C) is written on rank 0. You should read the data and write data with a code like this (assume that matrices A, B, and C are stored in row-major order):

```c
int rank;

FILE *fd;

double localA, localB, localC;

double *A = (double *)malloc(sizeof(double) * N * N);
double *B = (double *)malloc(sizeof(double) * N * N);
double *C = (double *)malloc(sizeof(double) * N * N);

MPI_Comm_size( MPI_COMM_WORLD, &rank );

if (0 == rank) {
  fd = fopen( "A.dat", "rb" );
  fread( A, sizeof(double), N*N, fd);
  fclose( fd );

  fd = fopen( "B.dat", "rb" );
  fread( B, sizeof(double), N*N, fd);
  fclose( fd );
}

distribute(N, A, &localA); // FROM rank 0 TO all ranks
distribute(N, B, &localB); // FROM rank 0 TO all ranks

matmatmul(N, localA, localB, &localC);

collect(N, localC, C); // FROM all ranks TO rank 0

if (0 == rank) {
  fd = fopen( "C.dat", "wb" );
  fwrite( C, sizeof(double), N*N, fd);
  fclose( fd );
}
```