

Monday, 01/30/12:

Today: The data structure from homework, i.e., with matrix A stored as

one-dimensional array in C , with storage organized in column-oriented storage \Leftrightarrow first index is consecutive in memory

$$A[(i+n)*j] = A_{ij} \text{ is next to } A[(i+1)+n*j] = A_{i+1,j}$$

$$A = \begin{bmatrix} \downarrow & \downarrow & \downarrow & \dots & \downarrow \\ \downarrow & \downarrow & \downarrow & & \downarrow \\ \downarrow & \downarrow & \downarrow & & \downarrow \end{bmatrix}$$


HW7 #2(b): $B \in \mathbb{R}^{k \times n}$

Question: What structure and size is $\underline{L} \in \mathbb{R}^{? \times ?}$?

$$C = AB = \sum_{g=0}^{k-1} a_g b_g^T, \quad b_g^T \text{ is } g\text{-th row of } B \in \mathbb{R}^{k \times n}$$

Idea of algorithm is to split this sum onto p parallel processes, so we need $\underline{L}_k = k/p$.

\Rightarrow So, on each process, we need several, \underline{L}_k many, rows of B



Have B on Process 0, so we want to program MPI_Send that sends these rows to another process.

```
MPI_Type_vector (
    count
    blocklength
    stride
    old_type
    & new_ptr )
```

```
MPI_Type_vector( n, l_k, k,
    MPI_DOUBLE, &block_row_t )
```

Solution 1: use MPI_Send /
MPI_Recv like Q. 3/4

Question: $l_B \in \mathbb{R}^{l_k \times n}$

```

if (id == 0) {
    for (j = 0; j < n; j++)
        for (l_g = 0; l_g < l_k; l_g++)
            l_B[l_g + l_k * j] = B[l_g + k * j]
    for (dest = 1; dest < npr; dest++)
        MPI_Send(&(B[ ]), 1, block_row_t,
                dest, 0, MPI_COMM_WORLD)
                dest * l_k
} else {
    MPI_Recv(l_B, l_k * n, MPI_DOUBLE,
            0, 0, MPI_COMM_WORLD)
}

```

For the case, recall that we used the type signature to match, namely $l_k * n$ doubles. Those are exactly the data for l_B on each process.

Solution 2: Make this simpler.

Really want to program Ch. 5 style collective call using MPI_Scatter

```
MPI_Scatter ( B, 1, block_row, t,  
            l_B, (l_k * n), MPI_DOUBLE  
            0, MPI_COMM_WORLD )
```

But this will not work $\frac{1}{2}$

Reason: MPI_Scatter sends $l_k * n$ doubles to each process here. The MPI_Scatter automatically uses the next-following element in B as the first element sent to the next process, instead of using $\&(B[dest * l_k])$ which it should be.

In MPI terms, the sending to Process 1 uses as starting index the extent of block_row_t, which is $l_k * n$. But we need l_k as starting index for sending to Process 1. And need $dest * l_k$ in general for all processes.

In other words, MPI_Scatter is designed for consecutive data, but block_row_t is not consecutive.

Solution in Pacheco Sec. 8.4.5

= there is a special datatype that overwrites the extent of a datatype manually, we will put l_k there, and then MPI_Scatter will work.

So, starting point is `block_row_t` as above.

Now define a doubly derived datatype based on that with a modified extent. Use `MPI_Type_struct` with first variable of type `block_row_t` and second one `MPI_UB` = upper bound = modifies the extent

```
int count = 2
```

```
int bl[2]
```

```
MPI_Aint dl[2]
```

```
MPI_Datatypes tl[2], ub_block_row_t
```

```
bl[0] = bl[1] = 1
```

```
dl[0] = 0
```

```
dl[1] =  $l_k * \text{sizeof}(\text{double})$ 
```

```
tl[0] = block_row_t
```

```
tl[1] = MPI_UB
```

```
MPI_Type_struct(count, bl, dl, tl, &ub_block_row_t)
```

```
MPI_Type_commit(&ub_block_row_t)
```

then can use

```
MPI_Scatter(B, 1, ub_block_row_t,  
            l_B, (l_k * n), MPI_DOUBLE, 0, MPI_COMM_WORLD)
```

Load balancing on p parallel processes = HW2

trapezoidal rule, if p does not divide n

Idea: If p does not divide n , that is, n/p is not an integer, then there is a remainder of the integer division, namely $rem = n \% np$.

So, add 1 trapezoid to rem many processes

$$rem = n \% np$$

if ($id < rem$) {

$$l_n = n / np + 1$$

} else {

$$l_n = n / np \quad /* \text{integer division!} */$$

}

for trap. rule, also need l_a and l_b

\Rightarrow Think and debug integers!

Idea: Think of indices l_{ia} and l_{ib} only

$$l_a = a + h * l_{ia} ; l_b = a + h * l_{ib} \quad \left. \vphantom{l_a} \right\} \text{same code}$$

$$l_{ib} = l_{ia} + l_n$$

on all processes!

$$l_{ia} = \begin{cases} l_n * id = (n/np + 1) * id & \text{if } (id < rem) \\ (l_n + 1) * rem + l_n * (id - rem) & \text{if } (id \geq rem) \\ = (n/np + 1) * rem + (n/np) * (id - rem) & \end{cases}$$