

Tuesday 04/24/2:

HW plan:

Before Wed. download files and put in proper place

tar zxvf v0r0.0 suppliedfiles.tar.gz  
↑                    ↖ "expand"  
zip/unzip



Have a look at the files before lecture and ask questions  
Test "make".

Fri: Have working code and show to Stefan!  
→ Sit together as team!

Mo.: Have all results and assembled as table

Wed. morning: <sup>submit HW</sup> time in seconds:

	1 node	2 nodes	4 nodes
1 ppm	16	8	6
2 ppm	9	4	2
4 ppm	5	3	1
8 ppm	2	1	0.5



Use as many processes per node as possible, as the red numbers above  
or as in Figure 4.5 of tech. rep.


Def.:  $S_p = \frac{T_1}{T_p}$ ,  $E_p = \frac{S_p}{p}$

What does the best possible speedup plot look like?

basic answer: "optimal  $S_p$ " is linear

more subtle answer:

- We want several cases of problem size to get a feel for performance
- Specifically, we want one bad case, one excellent case, and one or two in between



Use Matlab script for the plot?

Idea is to reuse this script in future HW by just changing the numbers?

Currently, the public partition has a mixture of 4 core nodes (2 dual-core CPUs) and 12 core nodes (2 six-core CPUs).

Any mixture in one run would behave strangely.

Also, any mixture from one run to the next does not allow to compare performance, such as speedup.

⇒ Use either --constraint = 4 cores or  
= 12 cores for all studies

## Ch. 5 Collective Communications

Last time MPI\_Bcast, MPI\_Reduce, MPI\_Allred.

### Sec. 5.7 Gather and Scatter:

Example with vectors last time

$x \in \mathbb{R}^n$ ,  $l_x \in \mathbb{R}^{l_n}$  on each process

$l_n = n/np$

Want to assemble  $x$  on Process 0 (for test output or some other purpose)

MPI\_Gather (

void	* send_data
int	send_count
MPI_Datatype	send_type
void	* recv_data
int	recv_count
MPI_Datatype	recv_type
int	root
MPI_Comm	comm

MPI\_Gather (  $l_x$ ,  $l_n$ , MPI\_DOUBLE,  $x$ ,  $l_n$ , MPI\_DOUBLE, 0, MPI\_COMM\_WORLD )

Notice: recv\_count is the number of elements from each process  $\Rightarrow l_n$ , not  $n$

Since MPI\_Gather is collective, it needs to be executed on all processes  $\Rightarrow$   $x$  needs to be defined on all processes. But it needs to have memory allocated only on Process 0

```
double *x  
if (id == 0)  
    x = malloc(n * sizeof(double))
```

other options:

```
x = calloc(n, sizeof(double))
```

yet another option:

```
double x[n]
```

would work if  $n$  is constant or under some conditions. Clearly, one problem would be that  $x$  would be defined and allocated on all processes

MPI\_Scatter is the inverse operation of MPI\_Gather

```
MPI_Scatter(x, l_n, MPI_DOUBLE,  
           l_x, l_n, MPI_DOUBLE,  
           0, MPI_COMM_WORLD)
```

---

MPI\_Allgather does exactly the same as MPI\_Gather, but it makes the result available on all processes

Same syntax, except no 'root' argument.

→ See man pages! Look for how it defines the send\_count, rec\_count.

Do and contrast to Pacheco examples:

MPI-Gather (  $l_A, n \times l_n, \text{MPI\_DOUBLE},$   
 $A, n \times l_n, \text{MPI\_DOUBLE},$   
 $0, \text{MPI\_COMM\_WORLD} )$

matrix  $A \in \mathbb{R}^{n \times n}$  (square), assume  $n_p$   
divides  $n \Rightarrow l_n = n/n_p$  is integer

$$A = \begin{bmatrix} \hline l_A \\ \hline l_A \\ \hline l_A \\ \hline l_A \end{bmatrix} \text{ with } l_A \in \mathbb{R}^{l_n \times n}$$

$$\text{or } A = \left[ \begin{array}{c|c|c|c} l_A & l_A & l_A & l_A \end{array} \right] \text{ with } l_A \in \mathbb{R}^{n \times l_n}$$

Pacheco shows rowwise division  $A = \left[ \begin{array}{c} \hline \hline \hline \hline \end{array} \right]$

since C does it like that

double  $A[n][n], l_A[l_n][n]$

dynamic memory allocation for matrix

double \*\*l\_A

l\_A = malloc( l\_n \* sizeof(double\*))

for (i = 0; i < l\_n; i++) {

    l\_A[i] = malloc( n \* sizeof(double))  
}

⇒ l\_A<sub>ij</sub> mathematically is l\_A[i][j]

In this sense, this is row-oriented

Only each row is consecutive in memory, not l\_A overall ⇒ This means

that the above example for MPI Gather is wrong; it would only be applicable

for double l\_A[l\_n][n]

But we want to use dynamic memory allocation and MPI Gather

⇒ Need to use 1-D arrays? <sup>storage</sup>

Then, we may as well do column-oriented