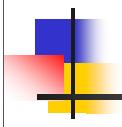

Programming Using the Message-Passing Paradigm (cont.)

Alexandre David
B2-206

You should know what these operations do.

Parallel programs should be written such that they behave correctly even if a global synchronization is performed before and after the collective call.

Barrier

- Communicator: Group of processes that are synchronized.
- The function returns after all processes in the group have called the function.

```
int MPI_Barrier(MPI_Comm comm)
```

Barrier

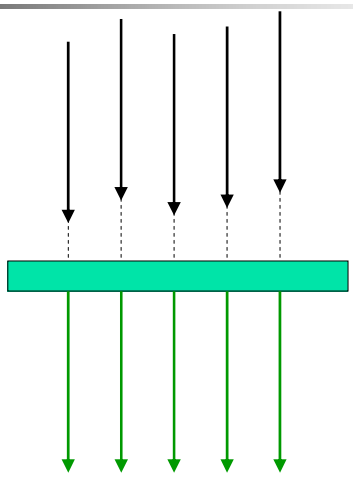

One-to-All Broadcast

- All the processes must call this function, even the receivers.

```
int MPI_Bcast(void *buf,
        int count, MPI_Datatype datatype,
        int source, MPI_Comm comm)
```

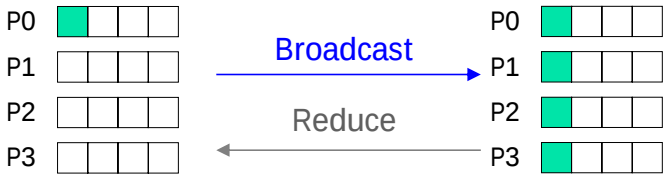

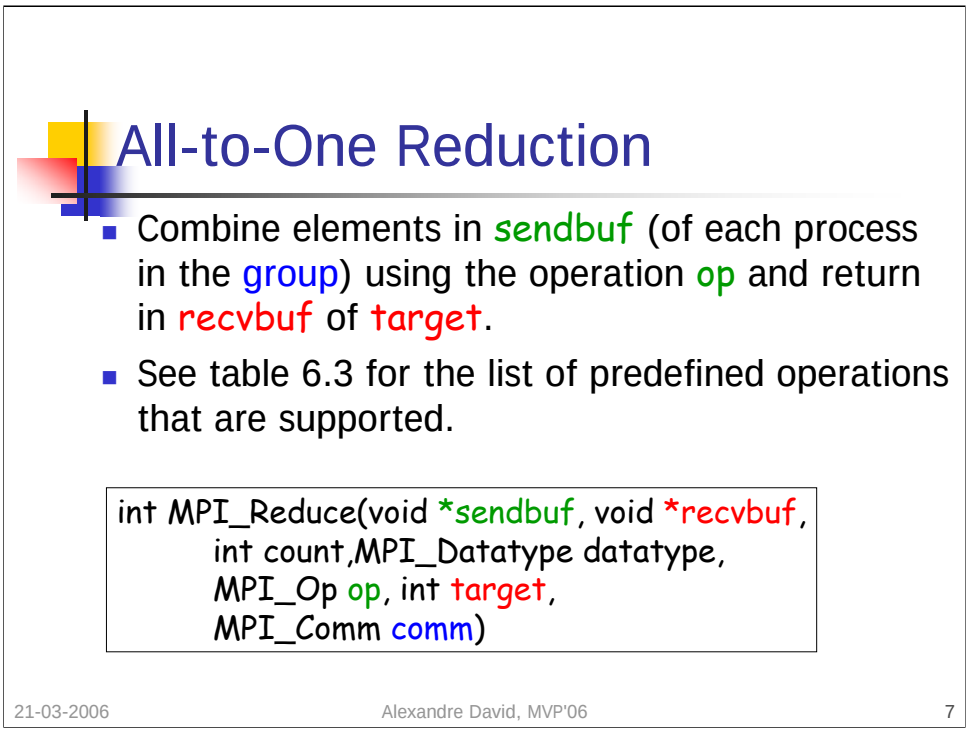

Constraint on the count of items of type datatype. All the processes call this function even those that are not the target and they all provide a recvbuf. When count > 1, the operation is applied element-wise. **Why do they all need a recvbuf?**

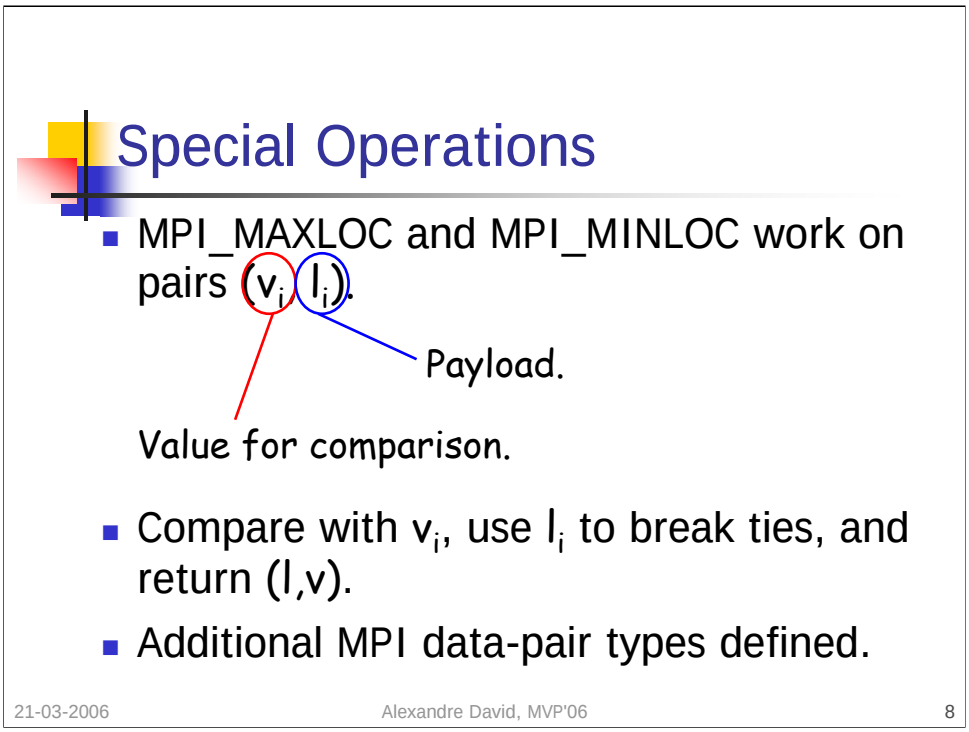

See table 6.4 for the different pair data types.

Example

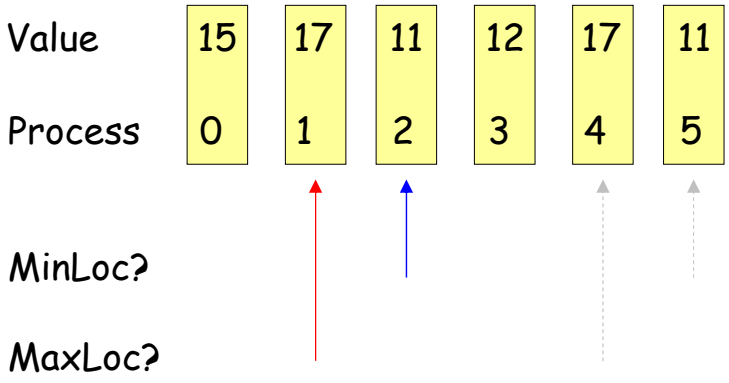

All-Reduce

- No target argument since all processes receive the result.

```
int MPI_Allreduce(void *sendbuf, void *recvbuf,
int count, MPI_Datatype datatype,
MPI_Op op, MPI_Comm comm)
```

Prefix-Operations

- Not only sums.
- Process j has prefix s_j as expected.

```
int MPI_Scan(void *sendbuf, void *recvbuf,
        int count, MPI_Datatype datatype, MPI_Op op,
        MPI_Comm comm)
```

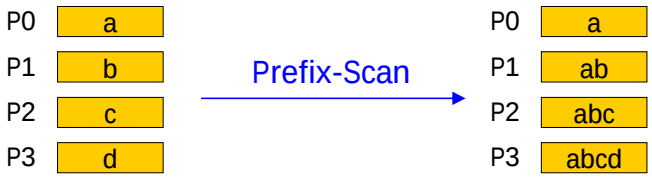

Scatter and Gather

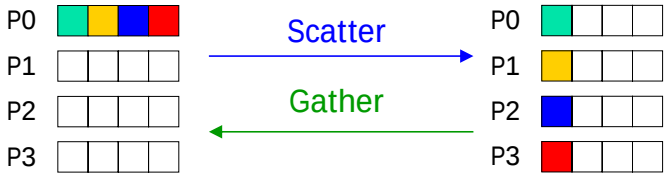

All-Gather

- Variant of gather.

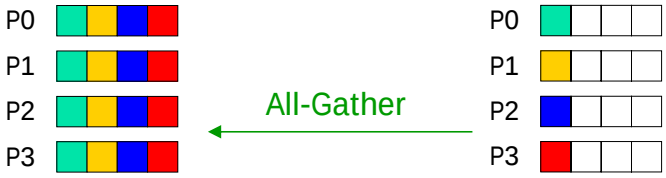

All-to-All Personalized

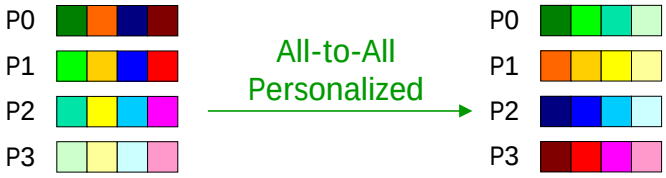

Example Matrix*Vector (Program 6.4)

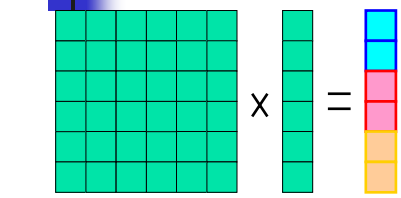

Partition on rows.

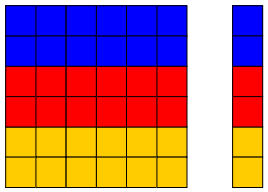

Allgather (All-to-all broadcast)

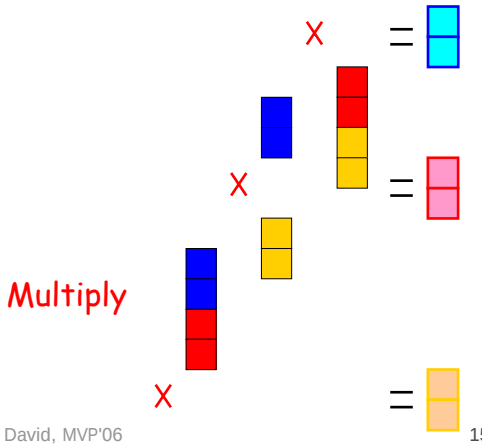

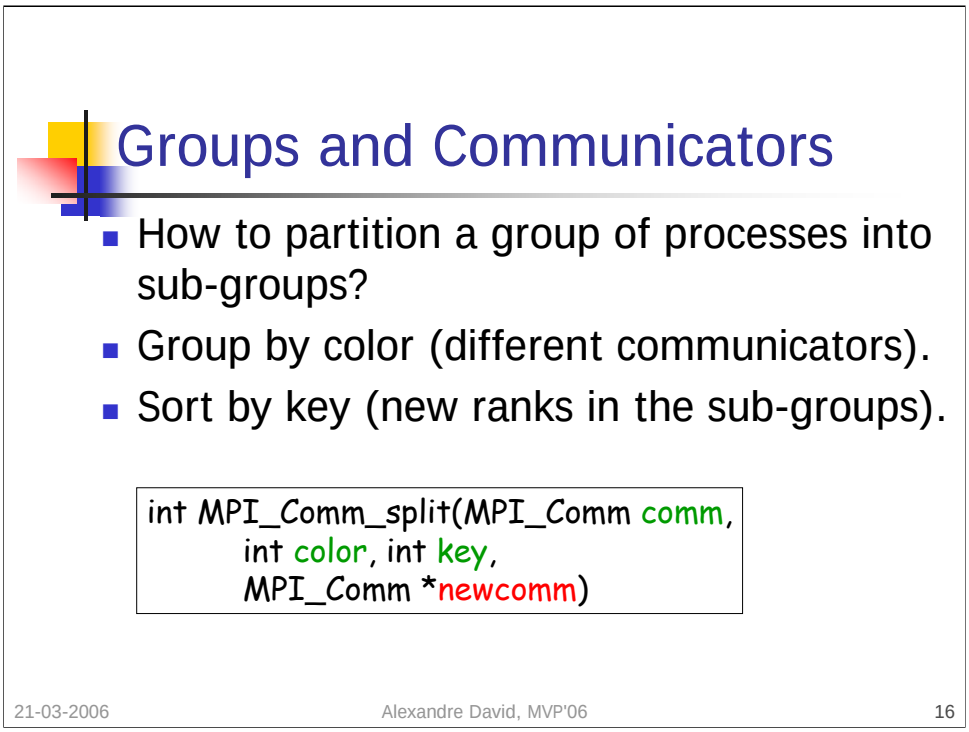

Sometimes, parallel algorithms need a restricted communication to certain subsets of processes.

Split Example

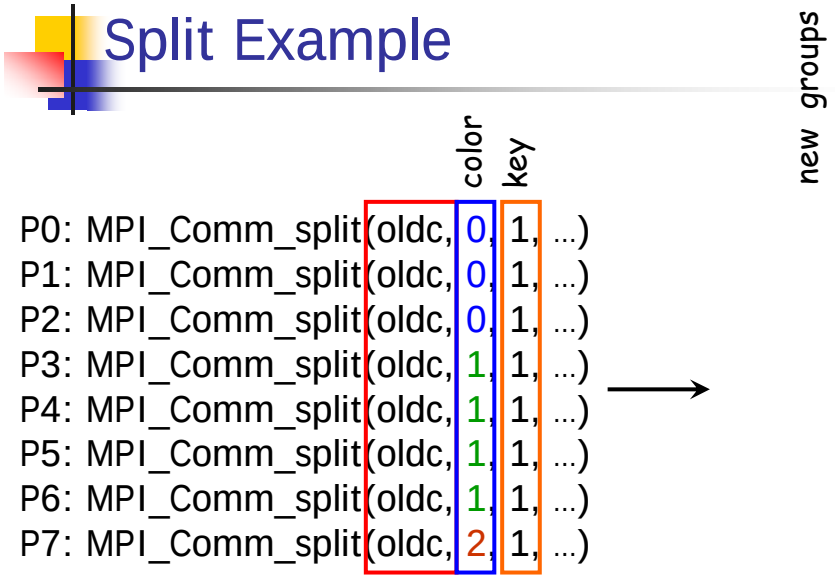

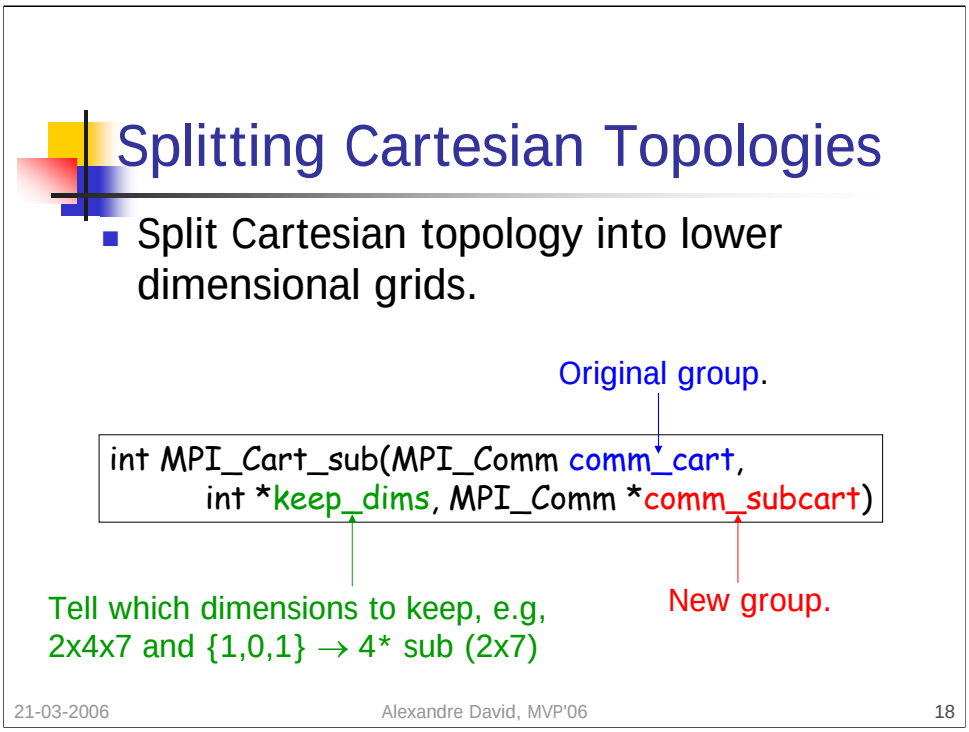

The keep_dims (boolean) array tells which dimensions to keep for the new sub-group partitioning. The coordinate will match, e.g., (1,2,3) in the original will give (1,3) and will be in the 2nd sub-group.

Example

original 2x4x7

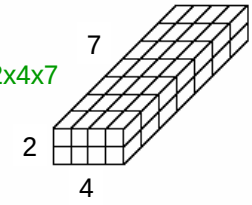

(1,0,1) -> 4* (2x7)

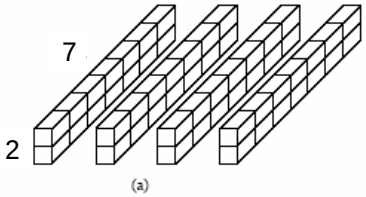

(a)

(0,0,1) -> 2*4* (7)

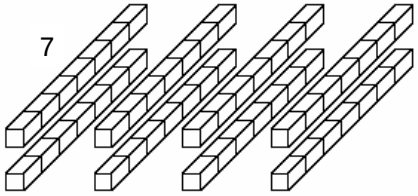

(b)

Example Matrix*Vector (Program 6.8)

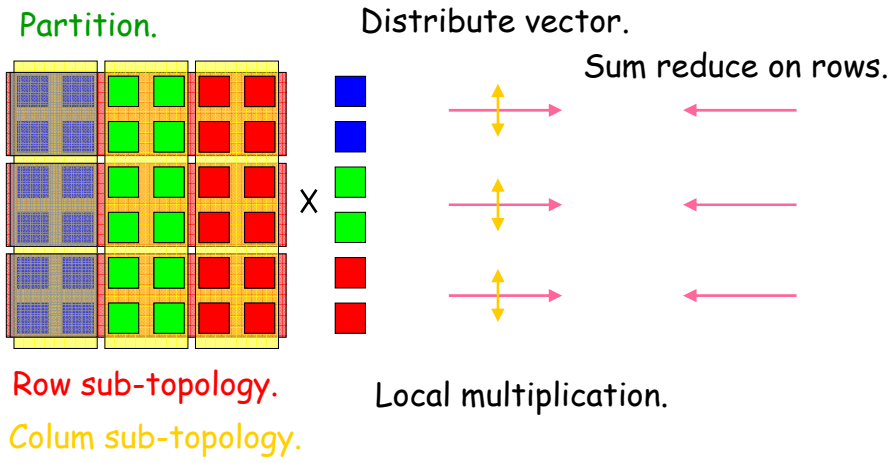

Performance Evaluation

- Elapsed time.

```
double t1, t2;
t1=MPI_Wtime();
...
t2=MPI_Wtime();
printf("Elapsed time is %f sec\n", t2-t1);
```