



Data Movement: Broadcast



- Broadcast copies data from the memory of one processor to that of other processors
 - One to all operation
- Syntax: `MPI_Bcast (send_buffer, send_count, send_type, rank, comm)`

Data Movement: Gather



- Gather copies data from each process to one process, where it is stored in rank order
 - One to all operation
- Syntax: `MPI_GATHER (send_buffer, send_count, send_type, rcv_buffer, rcv_count, rcv_type, rcv_rank, comm)`

Example: MPI_Gather

```
...
integer,allocatable :: array(:),array_r(:)
! Initialize MPI
call mpi_init(ierr)
call mpi_comm_size(mpi_comm_world,nprocs,ierr)
call mpi_comm_rank(mpi_comm_world,myid,ierr)
! Initialize the array
allocate(array(2),array_r(2*nprocs))
array(0)=2*myid
array(1)=2*myid+1
! Gather data at the root process
call mpi_gather(array,2,mpi_integer, &
               array_r,2,mpi_integer, &
               0,mpi_comm_world)
if (myid.eq.0) then
  write(*,*) "The content of the array_r:"
  write(*,*) array_r
```



Example: MPI_Gather

```

...
integer, allocatable :: array(:), array_r(:)
! Initialize MPI
call mpi_init(ierr)
call mpi_comm_size(mpi_comm_world, nprocs, ierr)
call mpi_comm_rank(mpi_comm_world, myid, ierr)
! Initialize the array

```

```
[lyan1@qb563 ex]$ mpirun -np 4 ./a.out
```

```
The content of the array:
```

```

      0      1      2      3      4      5
      6      7

```

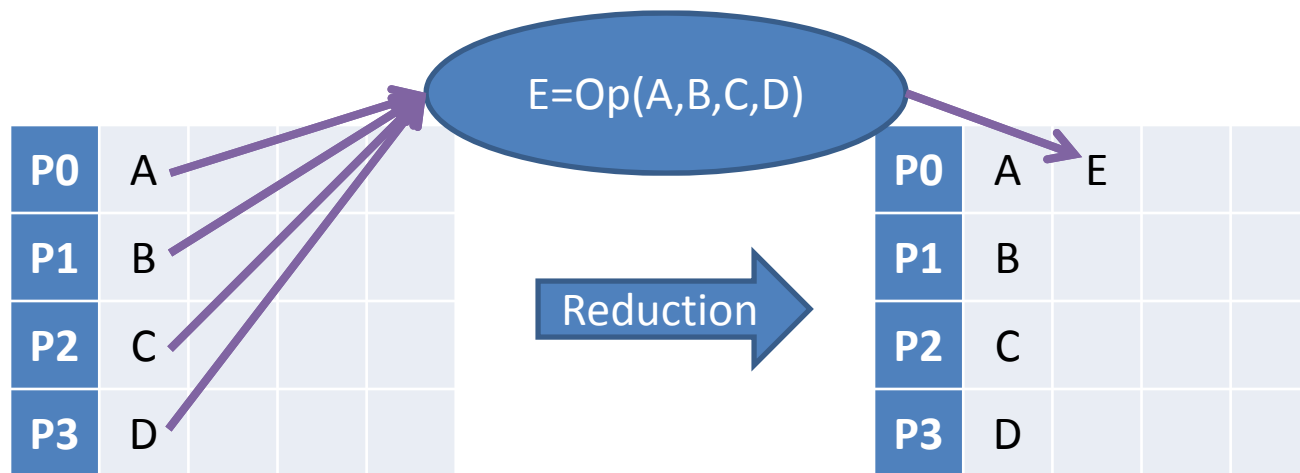
```

...
array_r, 2, mpi_integer, &
0, mpi_comm_world)

if (myid.eq.0) then
write(*,*) "The content of the array_r:"
write(*,*) array_r

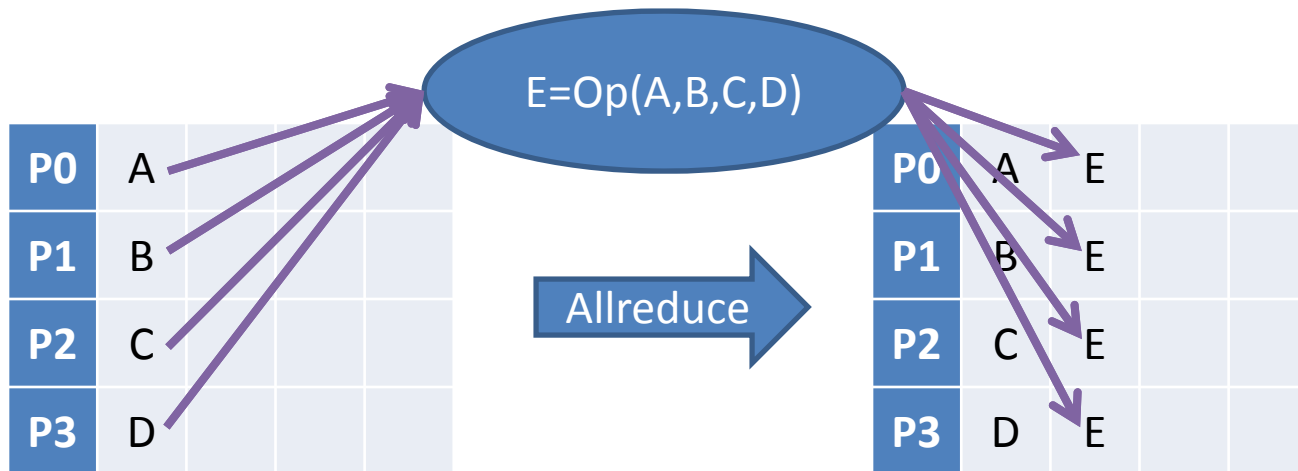
```


Collective Computation: Reduction



- MPI reduction collects data from each process, reduces them to a single value, and store it in the memory of one process
 - One to all operation
- Syntax: `MPI_Reduce(send_buffer, rcv_buffer, count, data_type, reduction_operation, rank_of_receiving_process, communicator)`

Collective Computation: Allreduce



- MPI allreduce collects data from each process, reduces them to a single value, and store it in the memory of EVERY process
 - All to all operation
- Syntax: `MPI_Allreduce(send_buffer, recv_buffer, count, data_type, reduction_operation, communicator)`

Why Derived Data Types

Process 0



Process 1



Solution 1

Call `MPI_SEND/MPI_RECEIVE` multiple times (once for each contiguous segment)

Why Derived Data Types

Process 0



Process 1



Solution 2

Process 0: copy data into contiguous buffer, then call MPI_SEND

Process 1: call MPI_RECEIVE, then copy data to destination

Why Derived Data Types

Process 0



Process 1



Solution 3


Define a new data type to describe the data to be transferred,
then call MPI_SEND/MPI_RECV

- Most efficient




Derived Data Type: Vector

- Allows replication of a data type into locations that consist of equally spaced blocks
- Syntax: `MPI_Type_Vector(count, block size, stride, old_type, new_type)`

```
Call mpi_type_vector(3,2,4,mpi_real8,my_vector_type,ierr)
Call mpi_type_commit(my_vector_type,ierr)
```



Data Ordering: C vs. Fortran



C



Fortran

Example: Broadcasting Submatrix

	1	2	3	4	5	6	7	8
1								
2								
3								
4								
5								
6								
7								
8								

```
!Number of blocks
```

```
nblocks=6
```

```
!Block length
```

```
blocklen=6
```

```
!Stride
```

```
stride=8
```

```
!Define the new data type
```

```
CALL MPI_TYPE_VECTOR(NBLCK, BLCKLEN, STRD, &  
MPI_INTEGER, submat_type, IERR)
```


```
CALL MPI_TYPE_COMMIT(submat_type, IERR)
```

```
!Call broadcast
```

```
CALL MPI_BCAST(AMAT(2,2), 1, submat_type, 0, &  
MPI_COMM_WORLD, IERR)
```


Nested Derived Data Type

- New data types can be created out of user-defined data types




```

Call mpi_type_extent(mpi_integer, &
    size_of_int, ierr)
call mpi_type_vector(2,1,2,mpi_integer, &
    column_type,ierr)
call mpi_type_hvector(4,1,6*size_of_int, &
    column_type,new_type,ierr)
    
```

Derived Data Type: Indexed

- Allows replication of a data type into locations that consist of unequally spaced blocks with varying length
- Syntax: `MPI_Type_indexed (count, blocklens[], offsets[], old_type, new_type)`
 - `blocklens` and `offsets` are array of size `count` that specify the length and displacement of each block, respectively



Count=4
 blocklens=[4,1,4,2]
 Offsets=[0,6,12,18]

