

# MPI Communicators and Groups

---

- 
-

Copyright (c) 2012 Young W. Lim.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Please send corrections (or suggestions) to [youngwlim@hotmail.com](mailto:youngwlim@hotmail.com).

This document was produced by using OpenOffice and Octave.

# Communicator & Group Overview

---

To limit communication to a subset of the processors,  
a group are created and are associated with a communicator

A group is an ordered set of processes.

Each process in a group is associated with a unique integer rank. (0 ~ N-1)

The group routines are used to specify processes involved in a communicator

A communicator is for a self-contained communication "universe".

A message sent with a given communicator can only be received by a process specifying the same communicator.

A communicator incorporates an instance of a group, and also includes contexts.

Both groups and communicators are

MPI objects that are accessed by handles (opaque object)

# Context

A message sent in one context cannot be received in another context. Collective operations may be independent of pending p2p operations.

Distinct communicators in the same process have distinct contexts. A context is essentially a system-managed tag (or tags) to ensure that collective and p2p communication within one communicator do not interfere, and that communications over distinct communicators don't interfere.

A possible implementation:  
a supplemental tag attached to messages on send and matched on receive.

Each intra-communicator stores the value of its two tags (one for p2p and one for collective communication).

In inter-communication (which is strictly p2p communication), two context tags are stored per communicator, one used by group A to send and group B to receive, and a second used by group B to send and for group A to receive.

Since contexts are not explicit objects, other implementations are also possible.

# Communicator

---

**Communicators** and **Groups** defines **collection of processes** that may communicate with each other.

Need to specify a communicator as an argument.

## **MPI\_COMM\_WORLD**

predefined communicator that includes **all MPI processes**.

## **Rank (Task ID)**

within a communicator, every process has its own unique integer identifier

- used to specify the source and destination.
- can be used in conditional statements.

# Group

---

Group: an ordered set of processes

Rank: 0 to N-1

One process can belong to many groups

Opaque group object, and cannot directly transferred to other process

Used within a communicator

To describe the participants in a communication universe and to rank them

Include the same local process

The source and destination of a message is identified by process rank within that group

Dynamic – can be created and destroyed during execution

Need to specify a communicator as an argument.

# Message Aggregation

---

## References

- [1] <http://en.wikipedia.org/>
- [2] [http://static.msi.umn.edu/tutorial/scicomp/general/MPI/mpi\\_coll\\_new.html](http://static.msi.umn.edu/tutorial/scicomp/general/MPI/mpi_coll_new.html)
- [3] <https://computing.llnl.gov/tutorials/mpi/>
- [4] <https://computing.llnl.gov/tutorials/mpi/>
- [5] Hager & Wellein, Introduction to High Performance Computing for Scientists and Engineers
- [6] <http://www.mpi-forum.org/docs/mpi-11-html>
- [7] <http://www.ib.cnea.gov.ar/~ipc/ipd2007/htmlcourse/>