


OPEN MP

- More threads than processes possible
- Each thread **unique id**
- **Sync. constructs** to prevent **DATARACE**
- **Shared & private** vars poss
- **SPMD**
- **Fork-join** thread model
- **thread** fork & create **threads**
- When done: **threads join** **thread**



pThreads

- **Fork-join** parallelism
- → can spawn new threads
- **symmetric peers**
- → threads can wait for others to complete
- **SPMD**
- **scheduled** by OS
- **No implicit sync.**
- → threads are independent
- **Share same memory**
- **Construct for sync. &**
- **coordination**
- → **MUTEX**
- → **2 locks**
- → **condition vars**


MSA PASSING MODEL

- **Final** processes with **LM**
- **Expl. Comm**
- **No implicit sync.** → only **comm.**
- **Nothing shared**

MPI

- **Final** processes in **comm. domain** (more than 1 pro)
- Each process **identified by rank** in domain
- **Ranks** successive $(0, \dots, (P-1))$
- **LOCAL DATA**
- **EXPLICIT COMM** → **Reliable & ordered**
- Structure of **com. data** ⊥ to **node/mod**
- **Comm. d.** **reflected physical topology**
- **No comm. cost model**
- 3 basic **comm models**
 - → **Point-to-Point** (diff. modes, local & ^{compt} remote)
 - → **1-sided** (diff. sync, local **compt. semantics**)
- **collective ops** (**local** **compt. semantics**)

Point-to-Point




1-sided

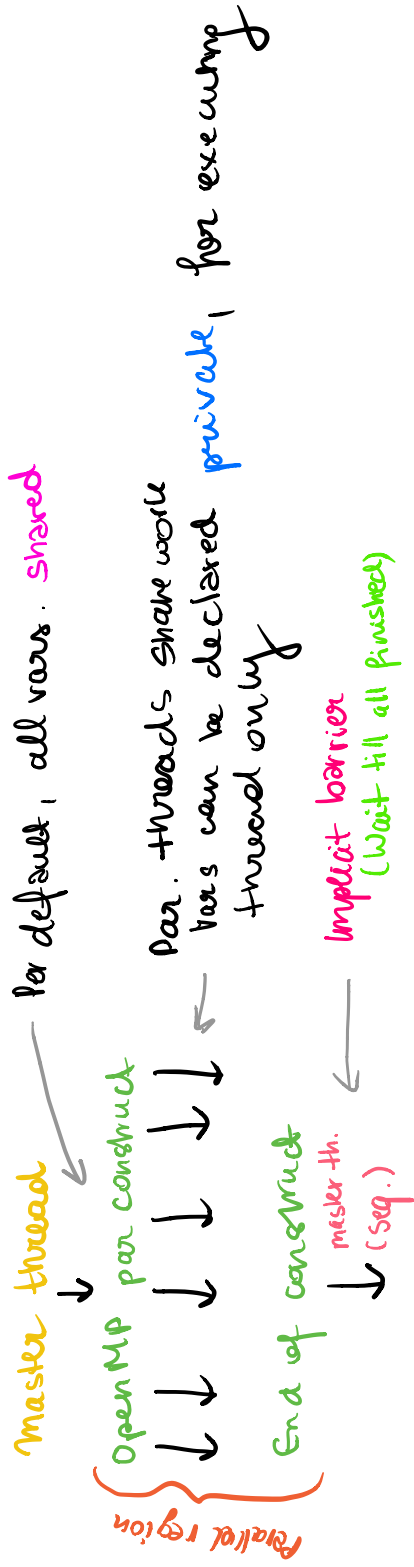


Collective

- **MPI-Barrier**
- **Bcast (root)**
- **Scatter(V)**
- **(All) Gather(V)**
- **(All) Reduce**
- **MPI-AlltoAll** (Ex) Scan



OPENMP




DeLo transfer between s. & p. copies var. is transparent, impl. $\mathcal{K}$ can be controlled

Parallel region

#pregn and parallel

5

۱۰



↓ m_{max} ← "memory model"

290K

→ Memory is consistent by end of card.

↳ updates to shared var visible

Threads asynchron

MPI Equivalents

MPI_Send	} = MPI_Sendrecv
MPI_Recv	
MPI_request	
MPI_Irecv	
MPI_Send	
MPI_Wait	

MPI_Isend } = MPI_Send
MPI_Wait

MPI_Recv } = MPI_Gather
MPI_Sendrecv MPI_Send

MPI_gather	} = MPI_Allgather (= P <u>Bcast</u>)
MPI_Bcast	
MPI_Bcast	
memcpy	

blocking
notsync

All collective func. blocking & notsync

Can do:

i...root

Bcast i
Bcast j

Bcast } i
gather }
Bcast } k
gather } j

Bcast } i
Recv } j
Bcast } i
Bcast } j

POINT-TO-POINT

Reliable

Deadlock-free

Non-overtaking (ordered)

MPI-SEND → MPI-RECV } Blocking
MPI-RECV ← MPI-SEND

Successful:

0 ≤ rank < size

① Sender specifies valid rank within comm.

& valid (allocated) buffer

② Receive with matching source rank & tag is eventually posted on same comm.

③ Sent data < | = to received amount

④ Type sig of sent & received data matches

ERROR: Caught by MPI-SEND: error

↳ IF NOT: DEADLOCK

↳ OTHERWISE: MPI_ERR_TRUNCATE // memory corruption

MPI-SEND

Short msg buffered at Sender receiver

Medium msg may be buffered locally

Large msg participation of receiving process needed

Sender receiver has been posted

Processed later

MPI-SEND can return fast

IMMEDIATELY

Large msg

participation of receiving process needed

MPI-SEND cannot return

before MPI-RECV becomes active

before MPI-RECV becomes active

before MPI-RECV becomes active

before MPI-RECV becomes active

before MPI-RECV becomes active

before MPI-RECV becomes active

before MPI-RECV becomes active

COST $\propto \alpha + \beta m$
Latency: Time to get data from source to dest process
bandwidth: How much data/amount (s) can be transferred from source to dest process

1-sided

Comm. between 2 processes

ONLY 1 Expiry involved with comm

MPI-GET → MPI-PUT ← Non-Blocking

only origin initiator comm & provides all args
target not involved
MPI-ACCUMULATE ()

local completion semantics

No guarantee that data have arrived before Sync. operation

Comm. window

origin puts/gets data from stand. MPI-buffer (buf, count, type)

Collective Operations

All processes involved Different for $\uparrow 0 \dots 8 \rightarrow$

Same root → if diff. roots DEADLOCK

#sent = #received

Matching type signatures (pairwise same size)

Difficult to impl: perf. correct. safety

COLLECTIVE OPERATIONS

MPI_Bcast: Data from root $\xrightarrow{a}$ all
 not evenly $\rightarrow$ MPI_Scatter: Indiv. data root
 MPI_Gather: Indiv. data all $\rightarrow$ root

MPI_Alltoall: Indiv. all $\rightarrow$ all

MPI_Reduce: Apply op. function(t, r, ...) to data from each process
 res at root
 MPI_Allreduce: res to all
 MPI_Reduce_Scatter: res scattered / dist.
 P i (non-root)
 r (sub, rbuf, ... root, comm);
 P j (root)
 r (sub, rbuf, root, comm);
 only sig. for root

MPI_scan: Prefix sum (reduction)

MPI_Exscan: value of process not used in reduction

MPI_Barrier: Sync. (waits until have reached routine)