Supercomputing 2024

MPI@Intel

Presenter: Maria J. Garzaran

intel

intel® MPI 2021.14 library CPU/GPU features

New **GPU** aware MPI_Allreduce algorithm

IMB-MPI1-GPU **allreduce**. 2 x Intel® Xeon® Platinum 8480+ Processor + 4 x Intel® Data Center GPU Max 1550. Higher is better

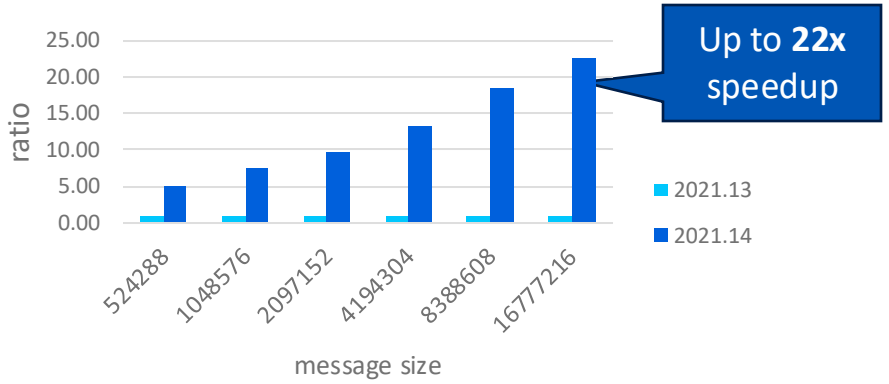

New **NIC** pinning. CPU/GPU/NIC "trinity"

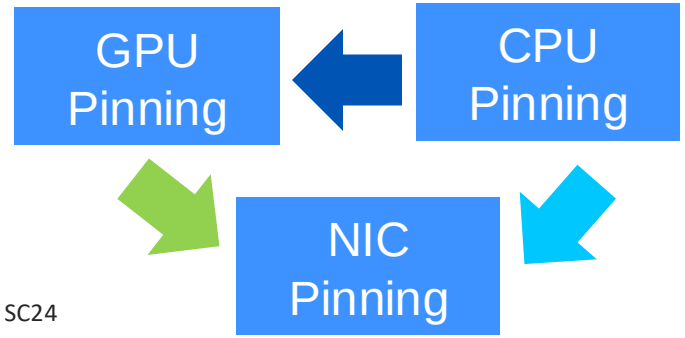

Intel® Xeon® 6 Processor Family enabling

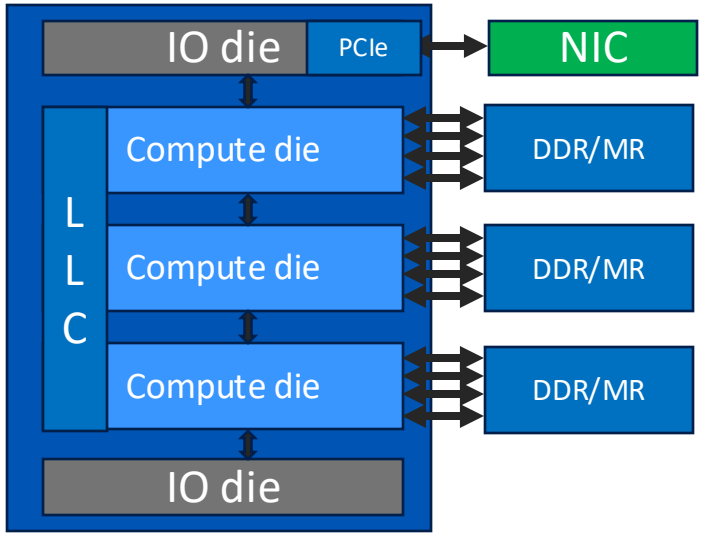

New CPU pinning (IPL2)
- New topology handling infrastructure

New SHM transport for Xeon 6E and Xeon 6P

Native BF16/FP16 support

Intel MPI and oneCCL interoperability optimizations

MPI 5.0+ Standard proposal: **CPU/GPU initiated** MPI_Put_notify

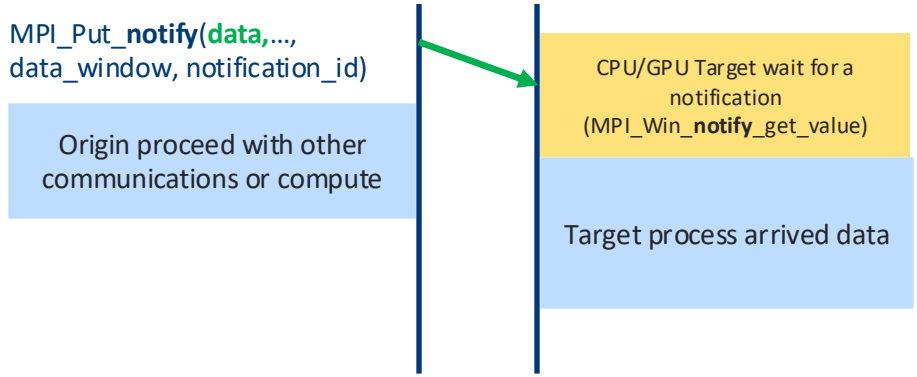

Allows target-side message-level completion semantics to one-sided operations.

Available on **CPU and GPU sides**

Sample code: link
MPI standard proposal (Note: requires MPI Forum github access): link

MPICH for Aurora

- Intel has done significant contributions to MPICH for the Aurora supercomputer (high-radix algorithms for collectives, support for multiple NICs, support for PVC and XeLink usage for point to pint and collective operations, among others)
- During 2024, Intel has worked closely with the MPICH team at Argonne to enable applications to run in the Aurora system
- MPICH was used for runs at scale of HPL in Aurora, which is the second supercomputer in theTop 500 list
  - Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11
- MPICH was used as backend for oneCCL for one of the finalist papers for the Gordon Bell Award in SC24
  - MProt-DPO: Breaking the ExaFLOPS Barrier for Multimodal Protein Design Workflows with Direct Preference Optimization, by Gautham Dharuman et al.

Intel® SHMEM v1.2.0 is Released!

- Device-initiated OpenSHMEM operations on Intel GPUs with SYCL
- Enables multiple host back-end options: Sandia OpenSHMEM, Intel® MPI
  - Experimental support for OSHMPI to allow any MPI runtime as back-end
- Supports most OpenSHMEM 1.5 features: Teams, Non-blocking Atomics, Vector pt-to-pt synchronizations, etc.
- Includes work-group and sub-group extensions for co-operative thread execution
- Extends with on-queue APIs allowing queueing SHMEM operations from host
- Implements fast synchronization algorithm between CPU and GPU threads
- Is tested and validated on Intel® Tiber™ AI Cloud
- https://github.com/oneapi-src/ishmem


intel