

SC22 MPICH BoF

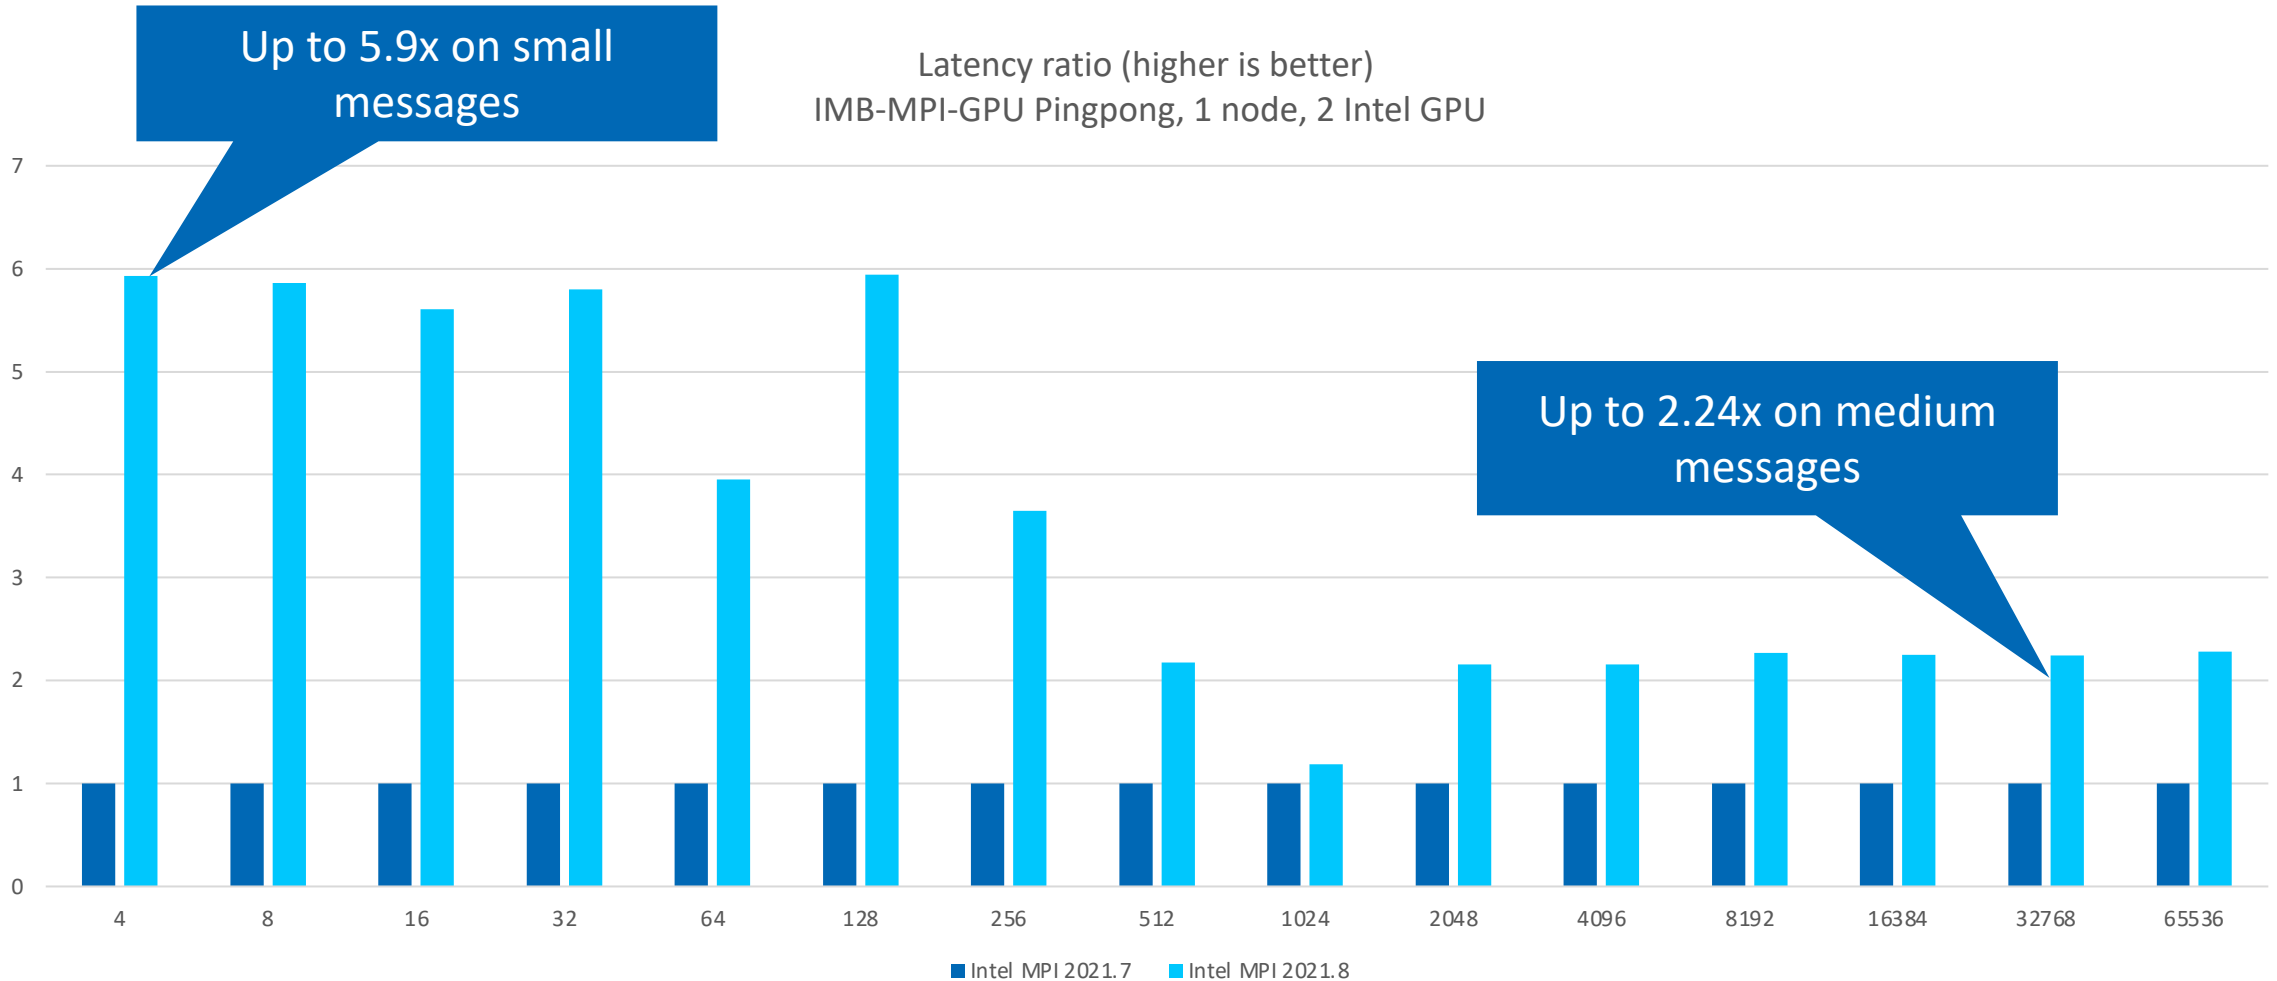
MPI@Intel

Maria Garzaran


Intel® MPI 2021.8 Update

- What's new:
 - Intel® Xeon® Max Series support and optimizations
 - Intra/Inter node transport support for Intel® Data Center GPU Max Series:
 - Fast memcpy (I_MPI_OFFLOAD_FAST_MEMCPY) for Intra and Inter node
 - Xe IPC (I_MPI_OFFLOAD_IPC) Intra node
- Enhancements:
 - Additional optimizations for Intel® Data Center GPU Max Series for Intra and Inter node paths


Intel® MPI 2021.8. Intra node GPU to GPU latency



Intel® MPI 2021.8. Intra node GPU to GPU BW




Intel® MPI 2021.8. Inter node GPU to GPU latency



Intel® MPI 2021.8. Inter node GPU to GPU BW

Unidirectional Bandwidth ratio (higher is better)
IMB-MPI-GPU uniband, 2 nodes, Intel GPU with HDR



MPICH for Aurora Update

- Contributions to Collectives
 - High radix algorithms for collectives (see our EuroMPI 2021 paper)
 - Support for multi-leader algorithms (Alltoall, Allreduce, Allgather)
- Support the CXI provider
 - Validation and bug fixes
- Coming soon: Improvements for Intel® Data Center GPU Max Series
 - Optimizations for intra-node communication using oneAPI Level Zero Inter Process Communication (IPC) primitives
 - Collective algorithms that use Xe-Links between GPUs inside the node



Thank You

Maximize Possibilities



intel.
XEON®
MAX SERIES

Maximize
Bandwidth



intel.
DATA CENTER
GPU
MAX SERIES

Maximize
Impact