Accelerating HPC and Deep Learning (DL) Applications with MVAPICH2-DPU, X-ScaleHPL-DPU, and X-ScaleAI-DPU Packages

Donglai Dai

d.dai@x-scalesolutions.com

X-ScaleSolutions

http://x-scalesolutions.com

Outline

Overview of X-ScaleSolutions

Overview of Products

- MVAPICH2-DPU: High-Performance MVAPICH2 for Accelerating Applications with NVIDIA's DPU technology
- X-ScaleHPL-DPU: Accelerating High-Performance Linpack Code (HPL) Benchmark with DPU Offload
- X-ScaleAI-DPU: Accelerating DL Training with DPU Offload

Overview of X-ScaleSolutions

- Started in 2018, bring innovative and efficient end-to-end solutions, services, support, and training to our customers
- Commercial support and training for the state-of-the-art communication libraries
  - Platform-specific optimizations and tuning
  - Application-specific optimizations and tuning
  - Obtaining guidelines on best practices
  - Timely support for installation and operational issues encountered with the library
  - Flexible Service Level Agreements
  - Web portal interface to submit issues and tracking their progress
  - Information on major releases and periodic information on major fixes and updates
  - Help with upgrading to the latest release
- Winner of multiple U.S. DOE SBIR grants
- Market these products for HPC and AI applications with commercial support
- A Silver ISV member of the OpenPOWER Consortium

Outline

Overview of X-ScaleSolutions

Overview of Products

- MVAPICH2-DPU: High-Performance MVAPICH2 for Accelerating Applications with NVIDIA's DPU technology
- X-ScaleHPL-DPU: Accelerating High-Performance Linpack Code (HPL) Benchmark with DPU Offload
- X-ScaleAI-DPU: Accelerating DL Training with DPU Offload

MVAPICH2-DPU Library 2023.05 Release

- Based on MVAPICH2 2.3.7
- Supports all features available with the MVAPICH2 2.3.7 release (http://mvapich.cse.ohio-state.edu)
- Novel GVMI-based framework to offload non-blocking collectives to DPU
- Offloads non-blocking Alltoall (MPI_Ialltoall)
- Offloads non-blocking Bcast (MPI_Ibcast)

Overlap of Communication and Computation with osu_Ialltoall (32 nodes)

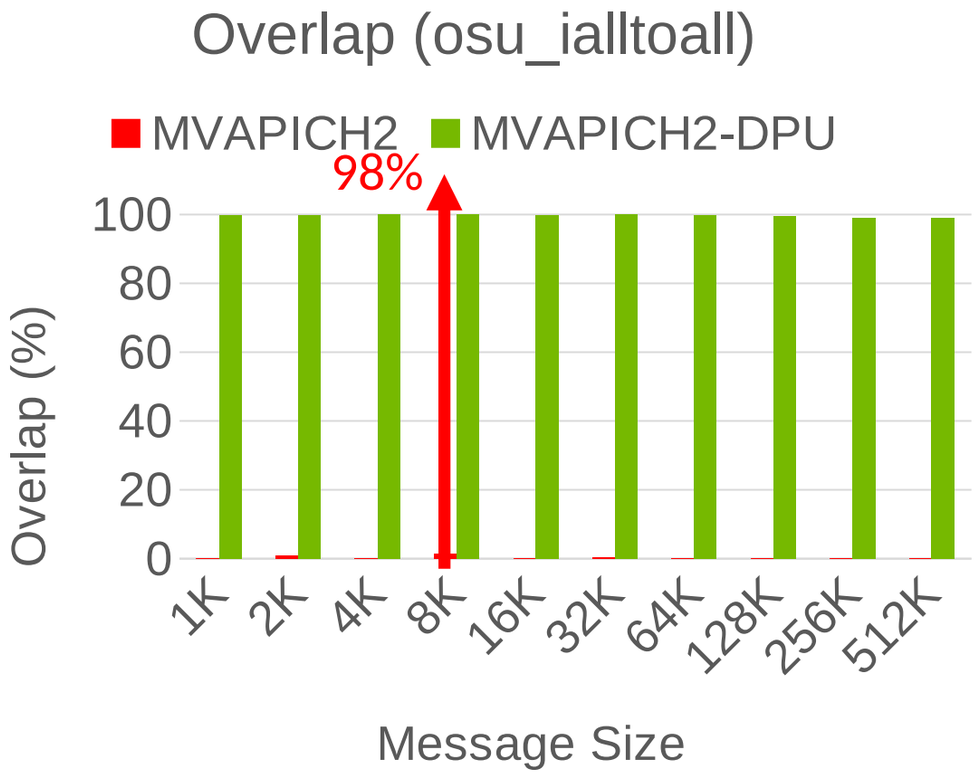

32 Nodes, 16 PPN

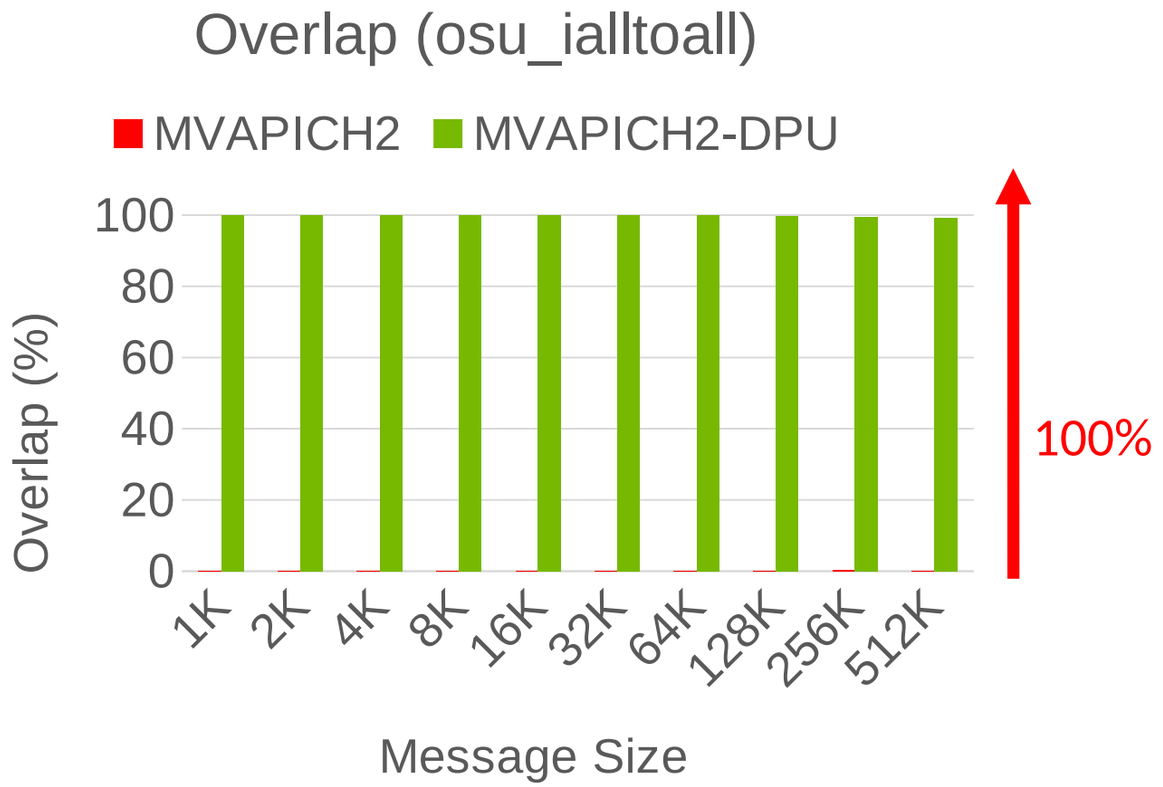

32 Nodes, 32 PPN

Delivers peak overlap

Total Execution Time with osu_Ialltoall (32 nodes), BF-2 100Gbps, Intel Platform, Medium Messages

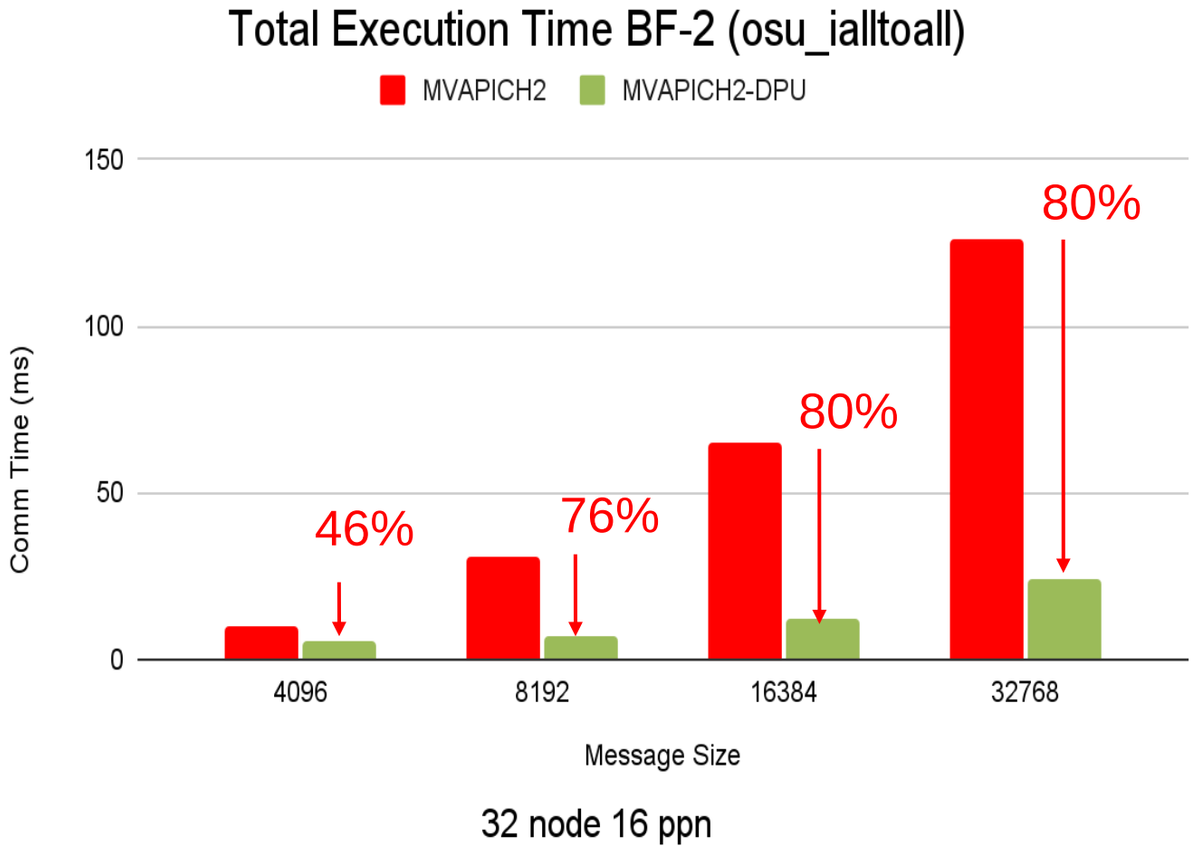

32 node 16 ppn

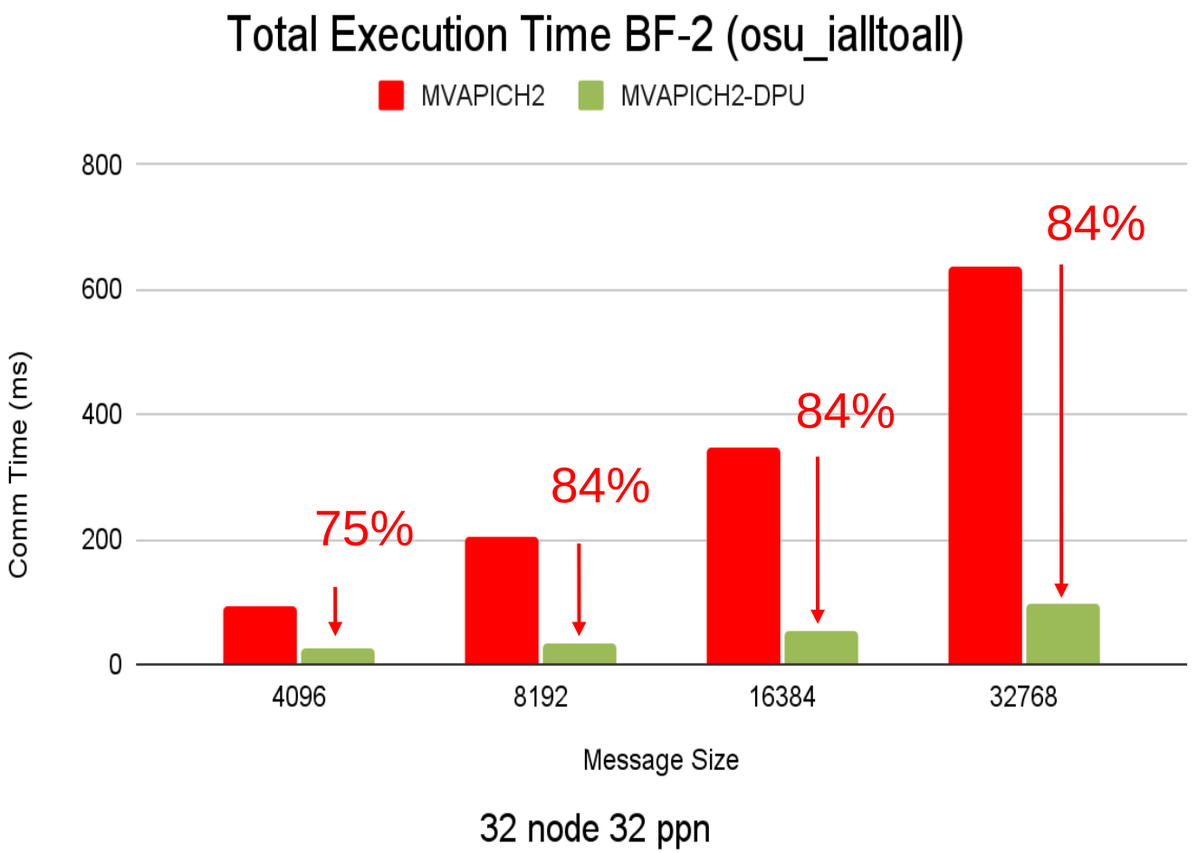

32 node 32 ppn

Benefits in Total execution time (Compute + Communication)

Total Execution Time with osu_Ialltoall (32 nodes), BF-2 100Gbps, Intel Platform, Large Messages

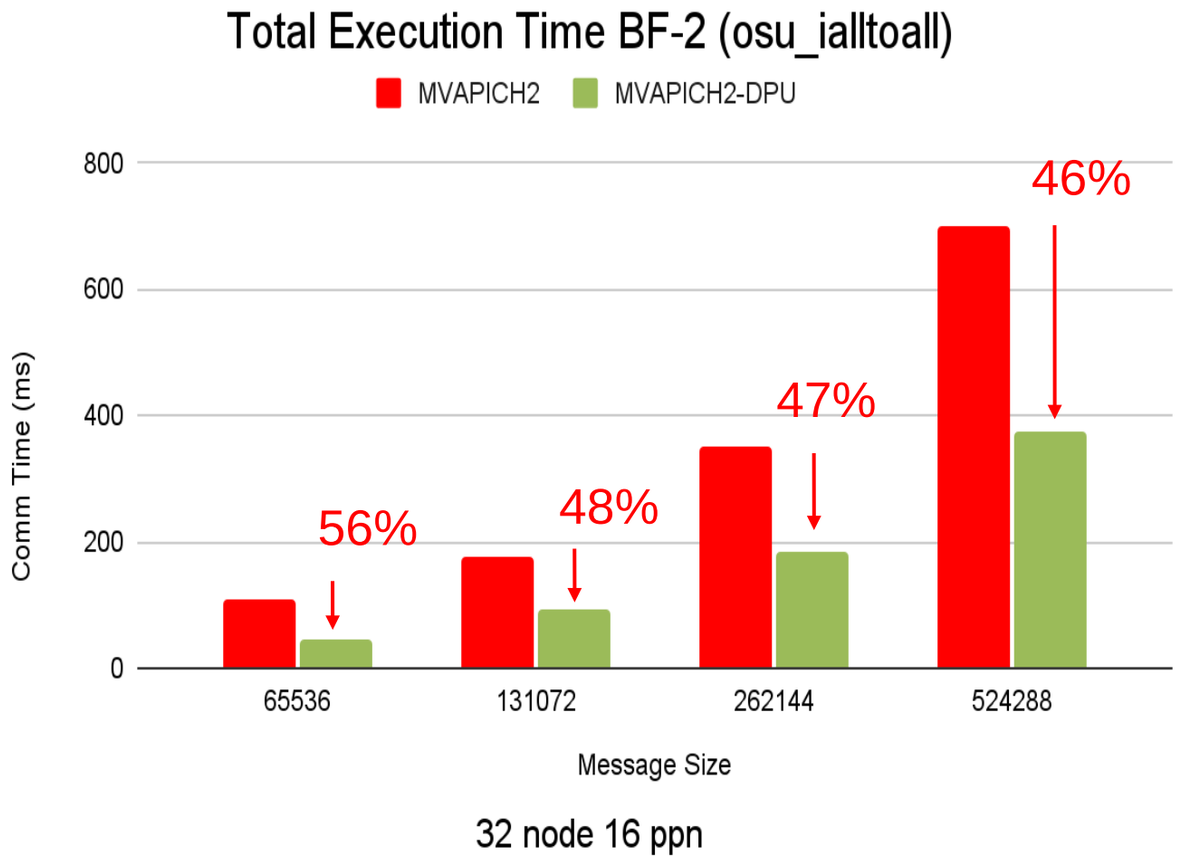

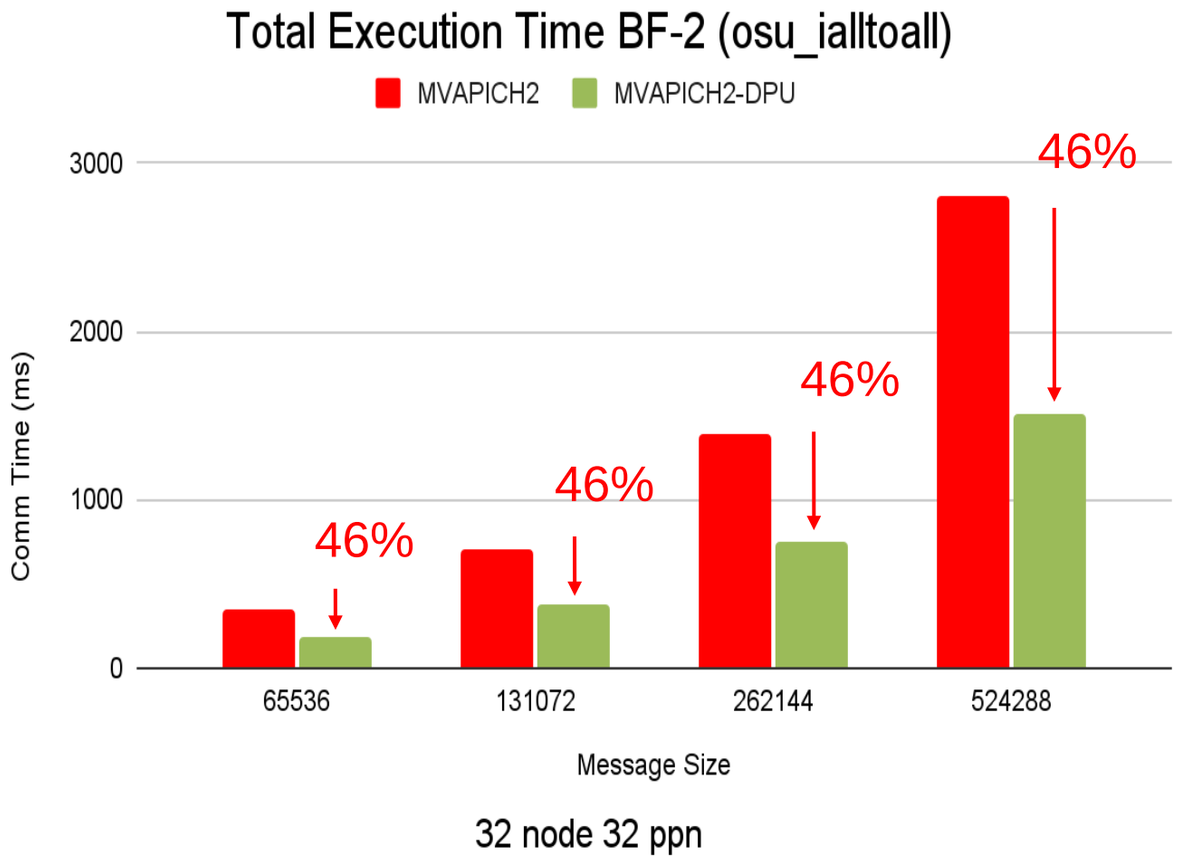

Benefits in Total execution time (Compute + Communication)

Total Execution Time with osu_Ibcast (16 nodes, 16 ppn and 32 ppn)

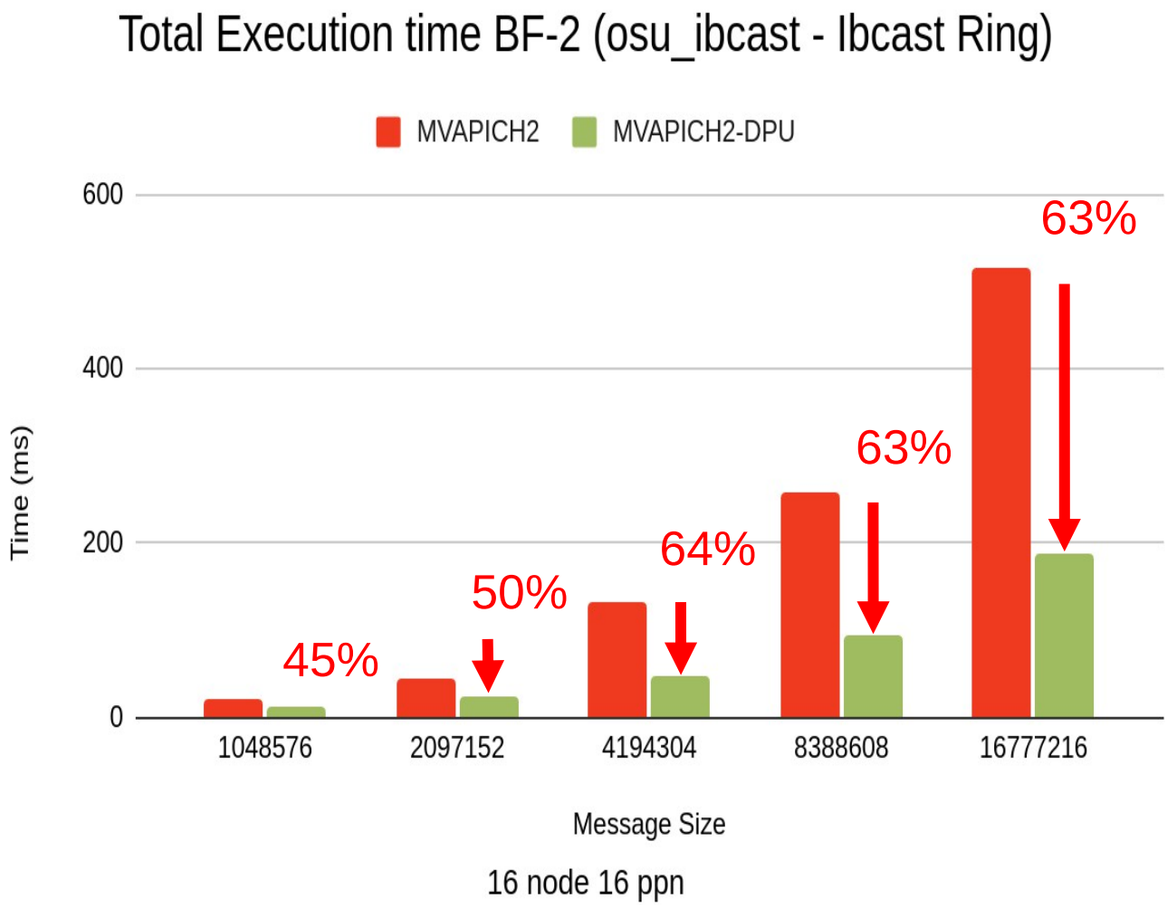

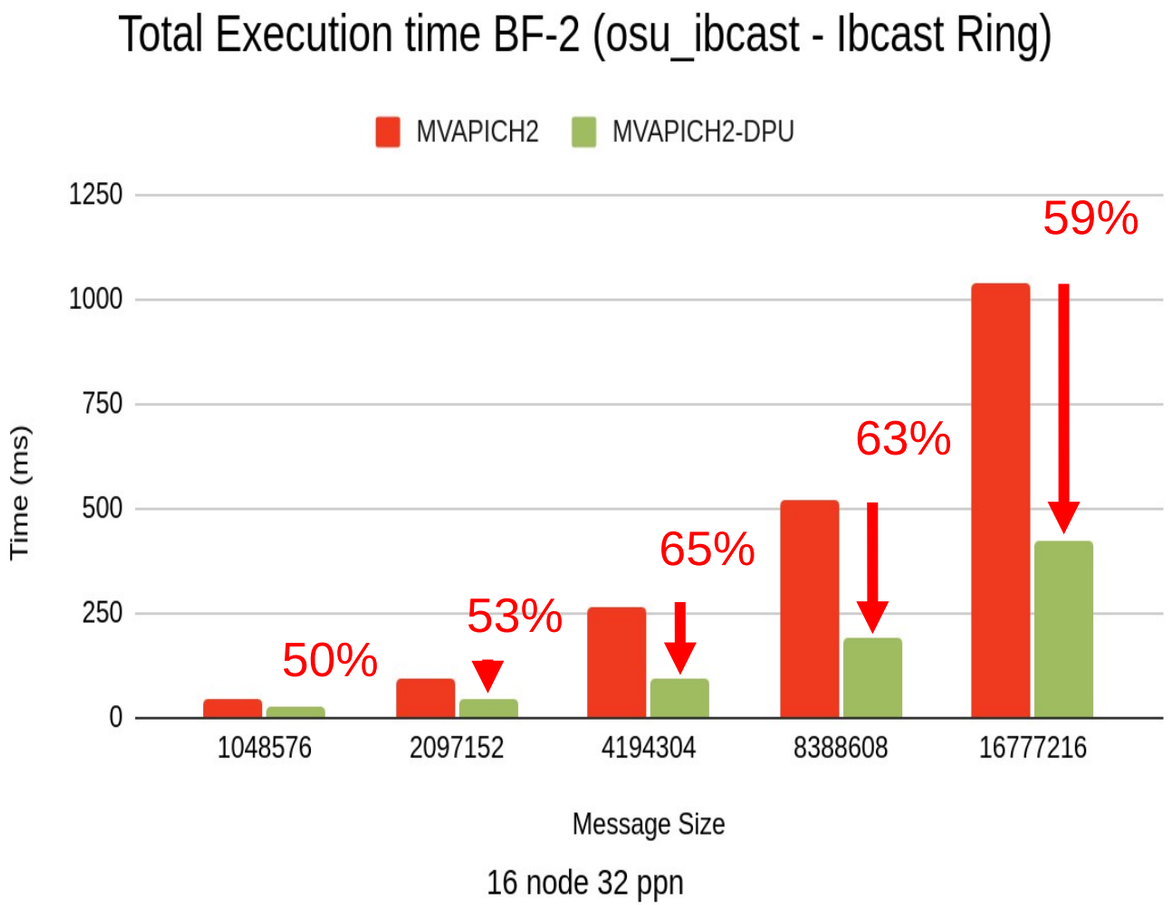

Benefits in Total execution time (Compute + Communication)

P3DFFT Application Execution Time (32 nodes), BF-2 100Gbps, Intel Platform

32 nodes with 32 ppn
(1,024 processes)

32x32 process grid

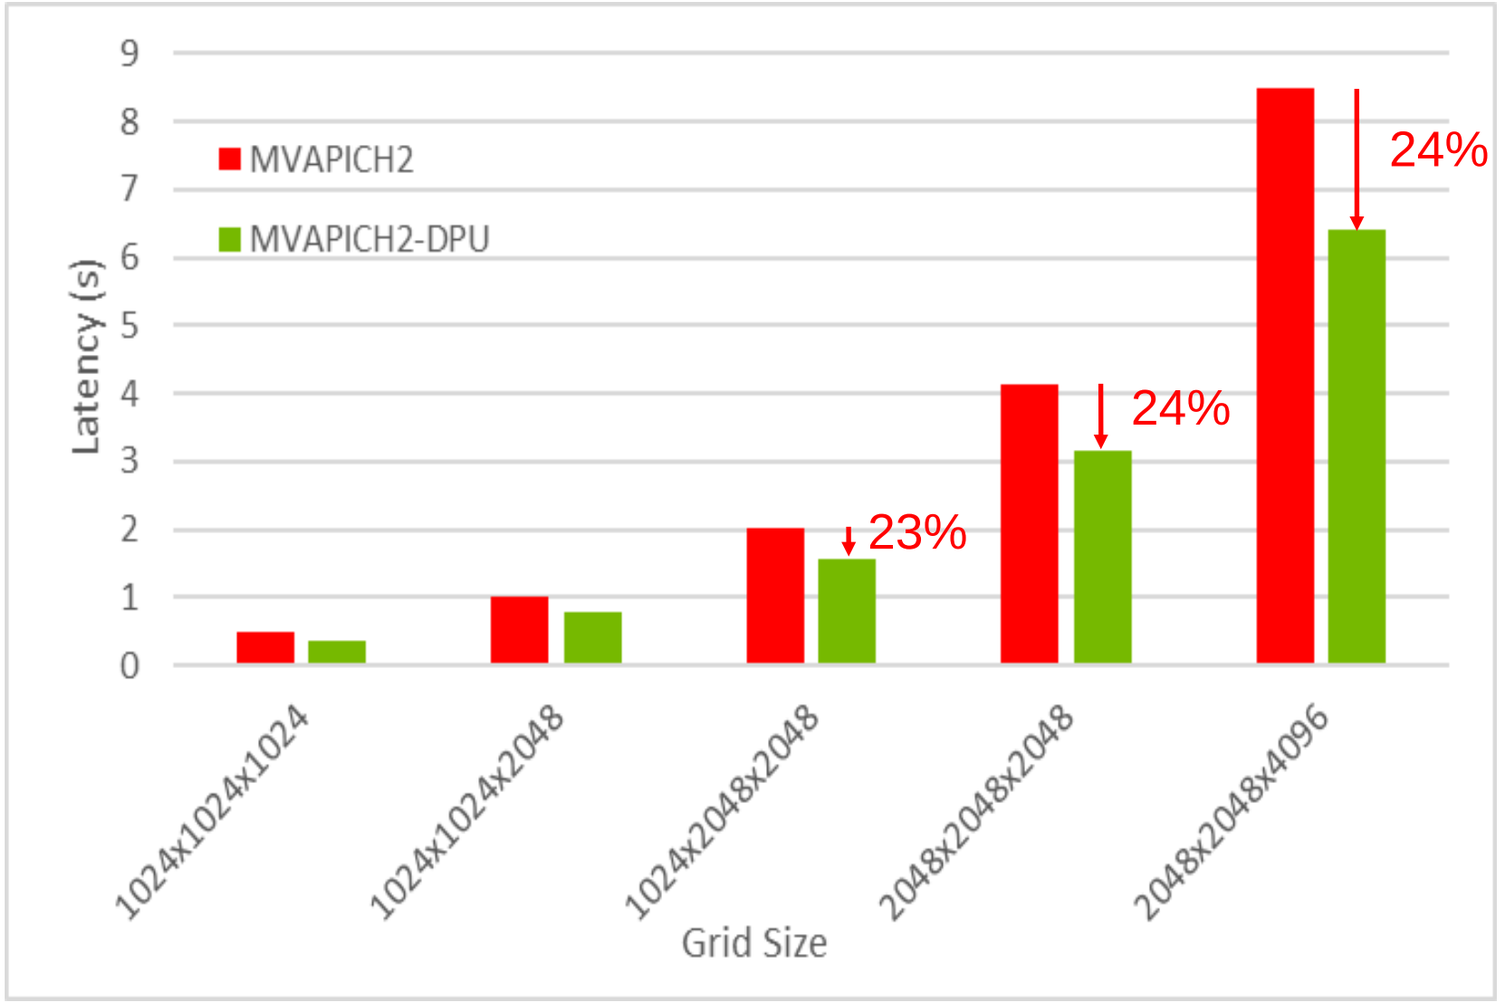

Benefits in application-level execution time

Outline

Overview of X-ScaleSolutions

Overview of Products

- MVAPICH2-DPU: High-Performance MVAPICH2 for Accelerating Applications with NVIDIA's DPU technology
- X-ScaleHPL-DPU: Accelerating High-Performance Linpack Code (HPL) Benchmark with DPU Offload
- X-ScaleAI-DPU: Accelerating DL Training with DPU Offload

X-ScaleHPL-DPU Package 2023.05 Release

- DPU Optimized version of the High-Performance Linpack Code (HPL) v2.3
- Co-designed with MVAPICH2-DPU library 2023.05 release
- Can be run in two modes: DPU mode and Host mode
  - In DPU mode, communication offloading to DPU is enabled
  - In Host mode, no such offloading occurs

Available from X-ScaleSolutions, please send a note to contactus@x-scalesolutions.com to get a trial license.

HPL Benchmark Performance (8 EPYC nodes, 128 ppn)

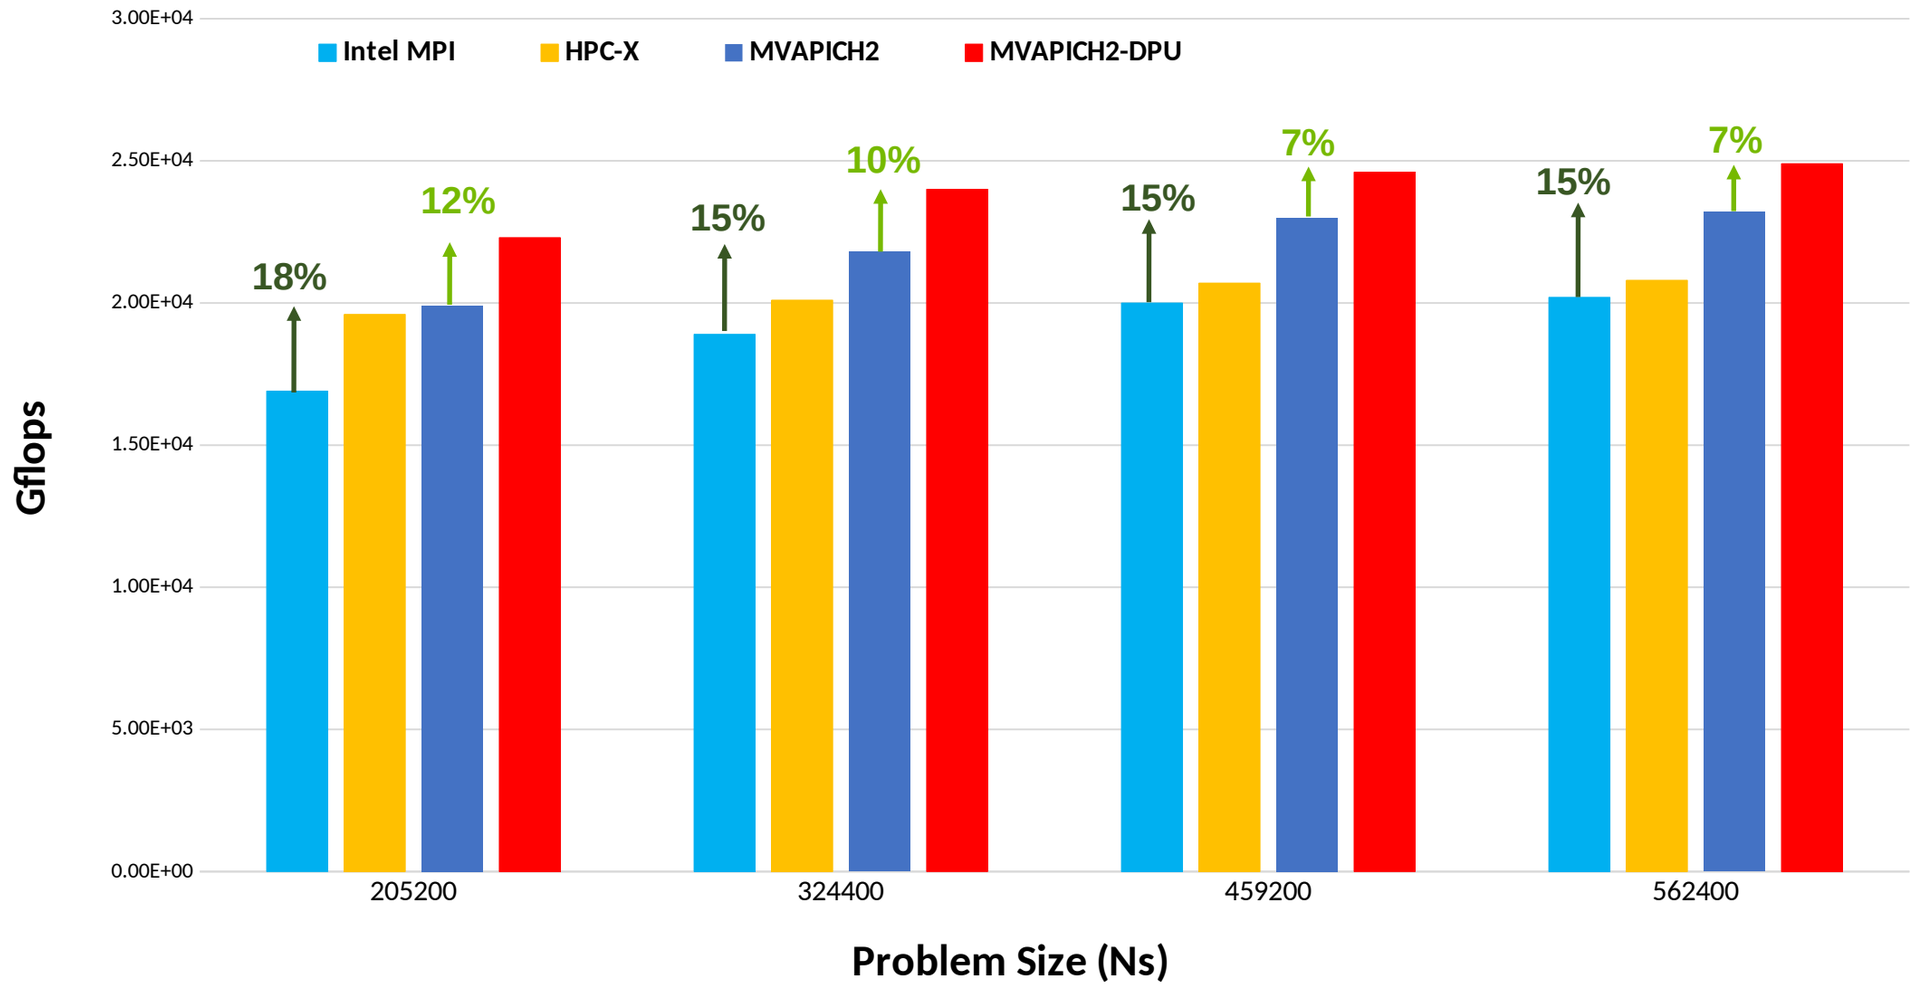

Performance benefits at application-level

HPL Benchmark Performance (16 nodes and 31 nodes, Intel Platform)

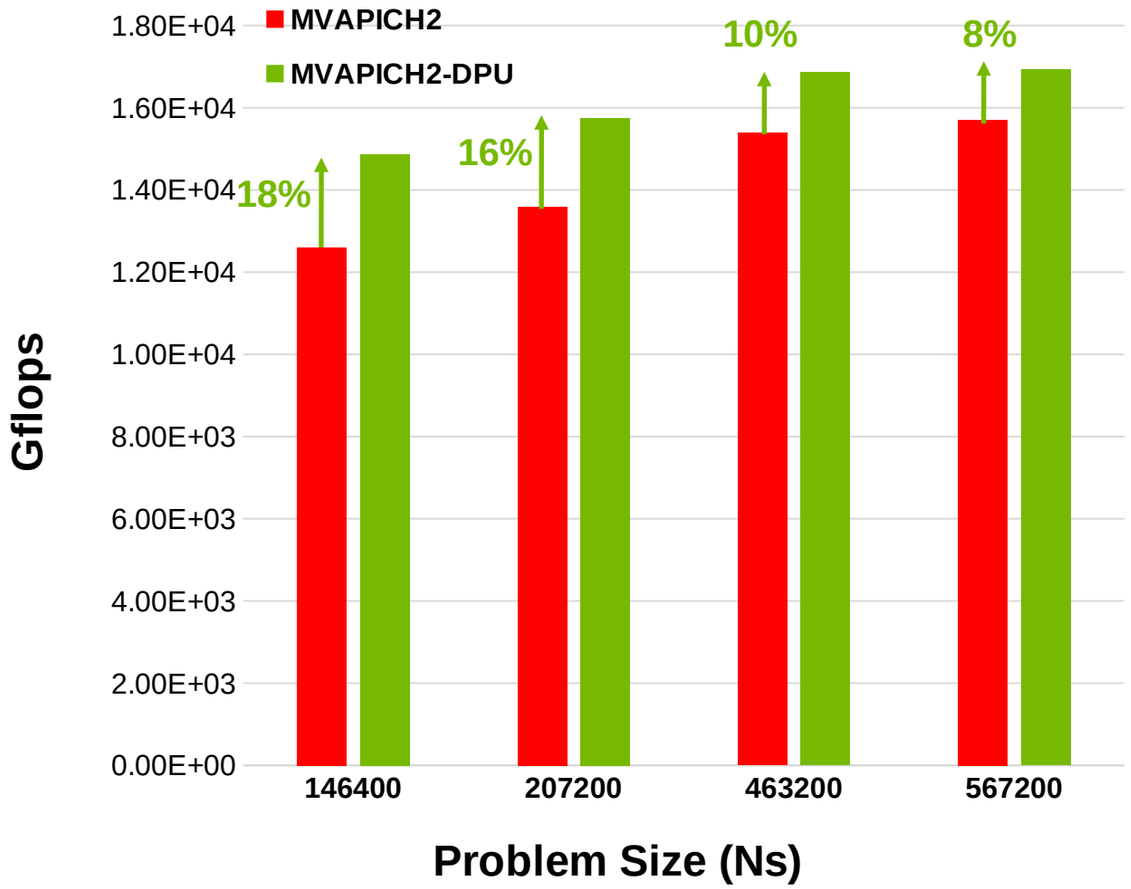

16x32 process grid

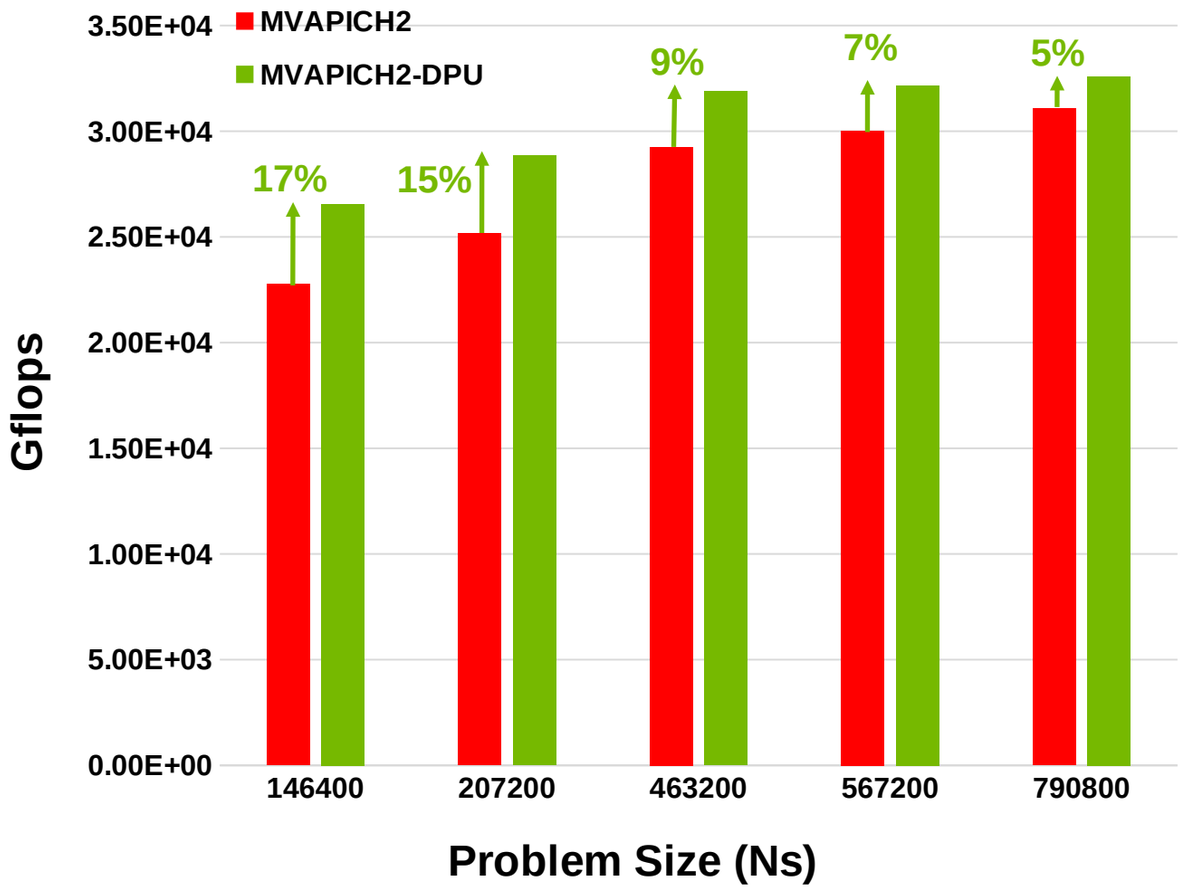

31x32 process grid

Performance benefits at application-level

Outline

Overview of X-ScaleSolutions

Overview of Products

- MVAPICH2-DPU: High-Performance MVAPICH2 for Accelerating Applications with NVIDIA's DPU technology
- X-ScaleHPL-DPU: Accelerating High-Performance Linpack Code (HPL) Benchmark with DPU Offload
- X-ScaleAI-DPU: Accelerating DL Training with DPU Offload

X-ScaleAI-DPU Package 2023.05 Release

- Accelerating CPU-based DNN training with DPU support
  - Easy to use (deployment and execution)
  - Scalable High-Performance
- Based on MVAPICH2-DPU 2023.05 with PyTorch 1.12.0 and Horovod 0.25.0
- Supports all features available with the MVAPICH2 2.3.7 release
- Supports PyTorch framework for Deep Learning
- Main Innovations:
  - Offloading some computation steps of DL training to DPU
  - Offloading checkpointing during long running DL training to DPU

Available from X-ScaleSolutions, please send a note to contactus@x-scalesolutions.com to get a trial license.

Training of ResNet-20v1 model on CIFAR10 dataset

System Configuration

- Two Intel(R) Xeon(R) 16-core CPUs (32 total) E5-2697A V4 @ 2.60 GHz
- NVIDIA BlueField-2 SoC, HDR100 100Gb/s InfiniBand/VPI adapters
- Memory: 256GB DDR4 2400MHz RDIMMs per node
- 1TB 7.2K RPM SSD 2.5" hard drive per node
- NVIDIA ConnectX-6 HDR/HDR100 200/100Gb/s InfiniBand/VPI adapters with Socket Direct

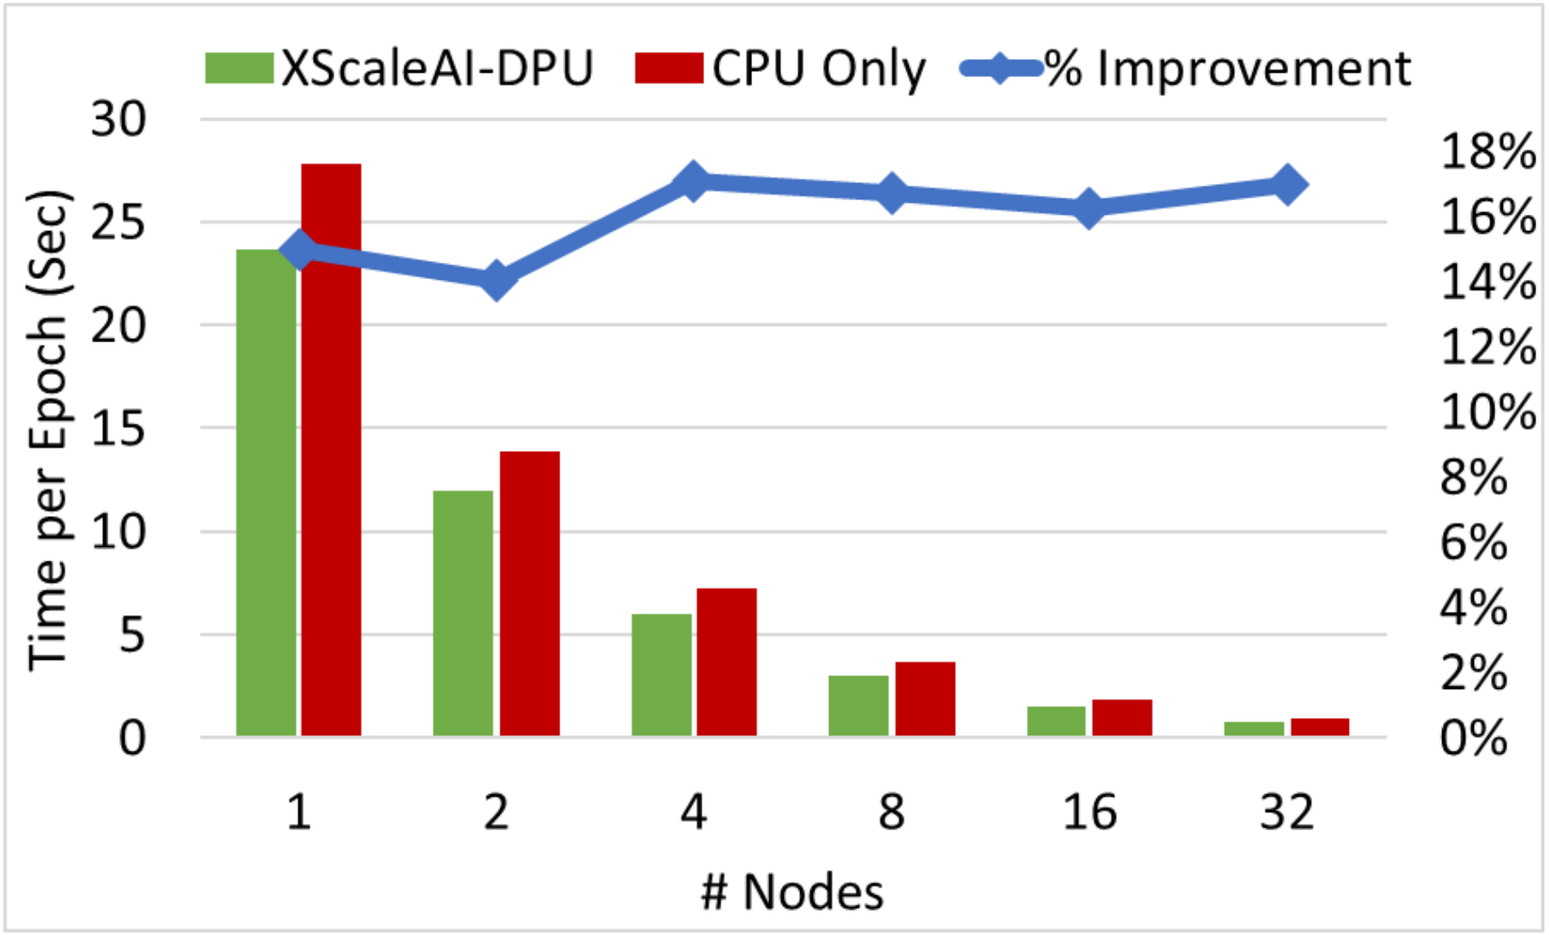

Performance improvement using X-ScaleAI-DPU over CPU-only training on the ResNet-20v1 model on the CIFAR10 dataset

Training of the ShuffleNet model on TinyImageNet Dataset

System Configuration

- Same as the last slide

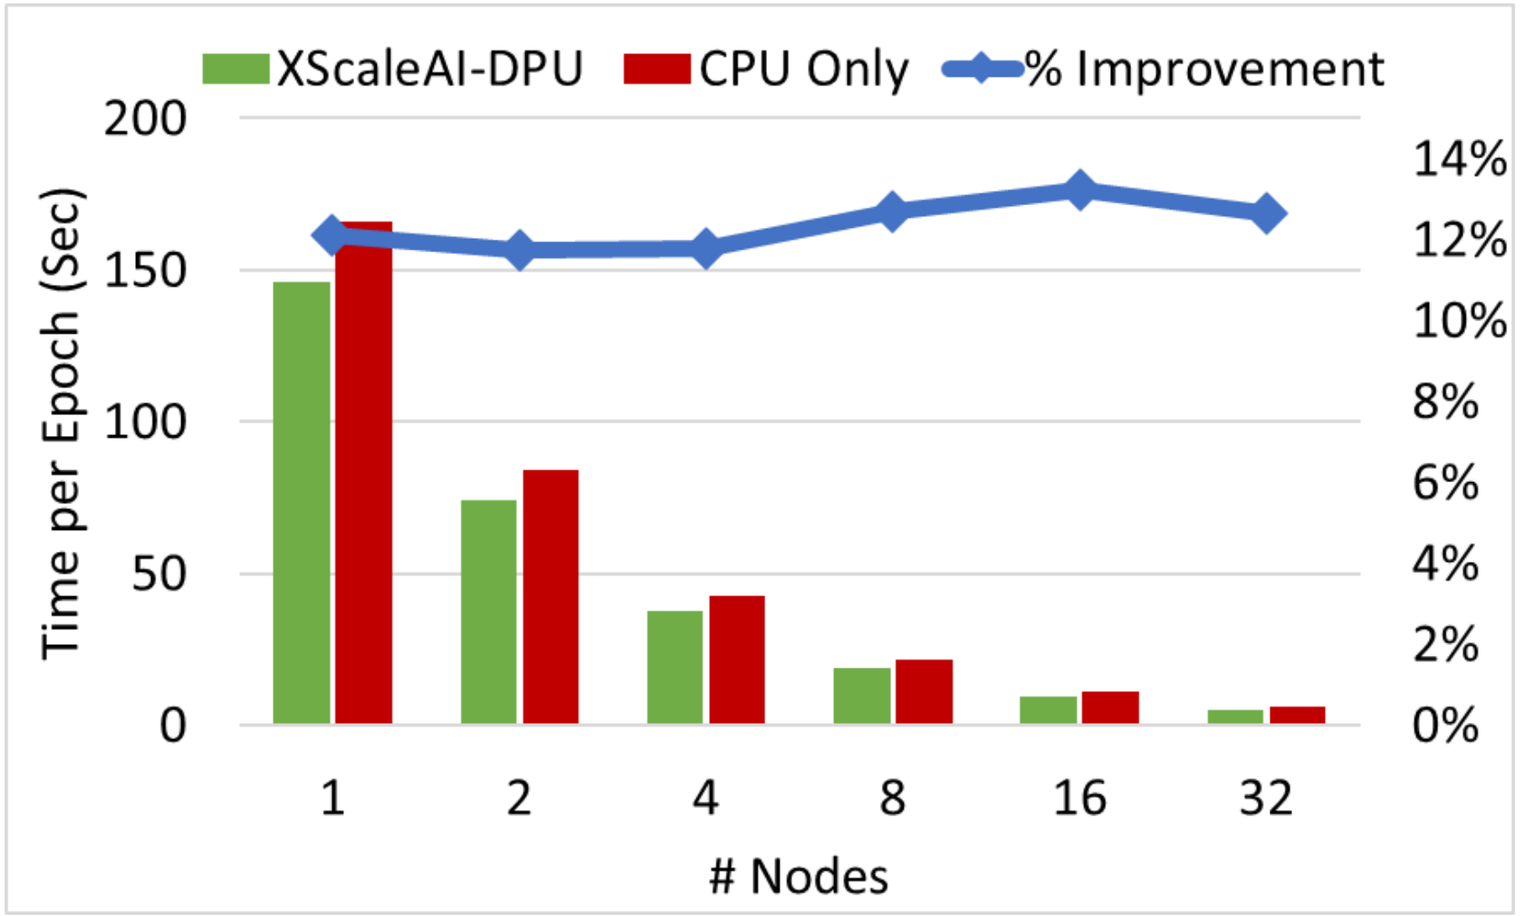

Performance improvement using X-ScaleAI-DPU over CPU-only training on the ShuffleNet model on the TinyImageNet dataset

X-ScaleAI-DPU Checkpointing

- All DNN training runs must save snapshots of in-progress snapshots of the model parameters called a *checkpoint*
  - Unstable HPC/Cloud clusters require frequent checkpointing
  - The checkpoint cost scales with the number of model parameters
- Typically, the root rank saves the checkpoint while all other ranks stall (Called a *root ckpt*)
- By offloading checkpointing to the DPU, we can overlap checkpoint I/O with compute

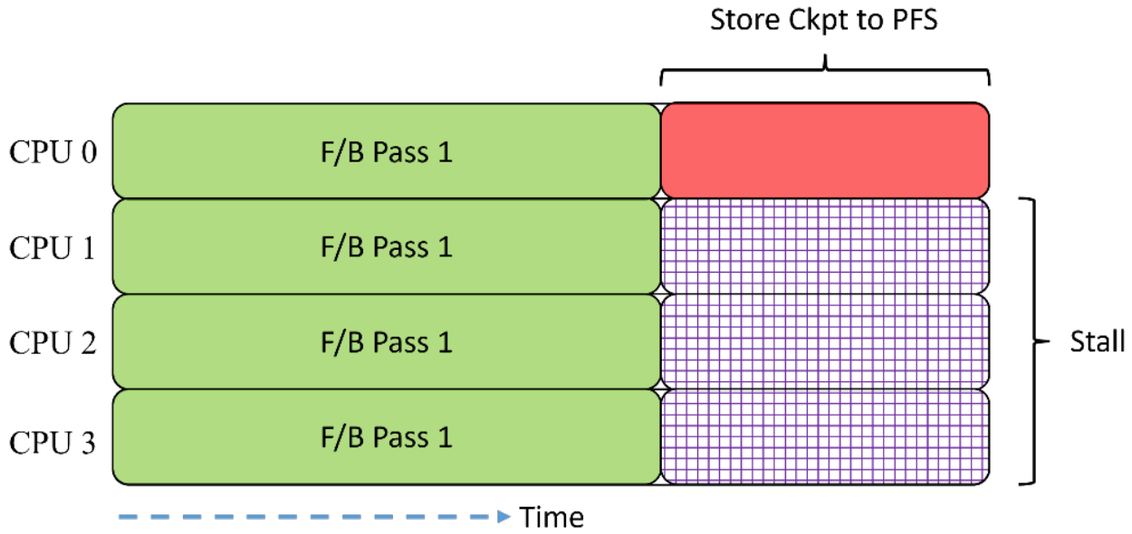

Default Root Checkpointing

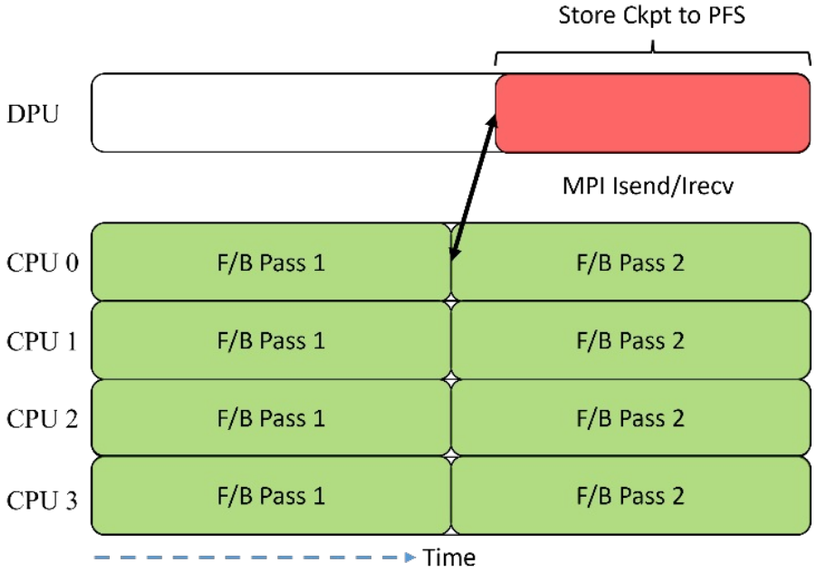

Checkpointing with DPU offload

X-ScaleAI-DPU Checkpointing

- We measure the time per epoch with ResNet152 on the TinyImageNet dataset
  - *Root* and *Offloaded* checkpoints refer to the default and DPU checkpointing schemes, respectively
  - TinyImageNet is a subset of ImageNet containing 100k images downsized to 64x64 pixels
  - We use *4x* and *10x* to refer to checkpointing 4 and 10 times within an epoch, respectively
  - We expect users to checkpoint less frequently (4x) on stable HPC/Cloud systems, and more frequently (10x) on unstable HPC/Cloud systems
- DPU-Offloaded checkpointing outperforms root checkpointing at all node scales
- Up to **52%** reduction in epoch time for ResNet152 on an unstable system where frequent checkpointing is required

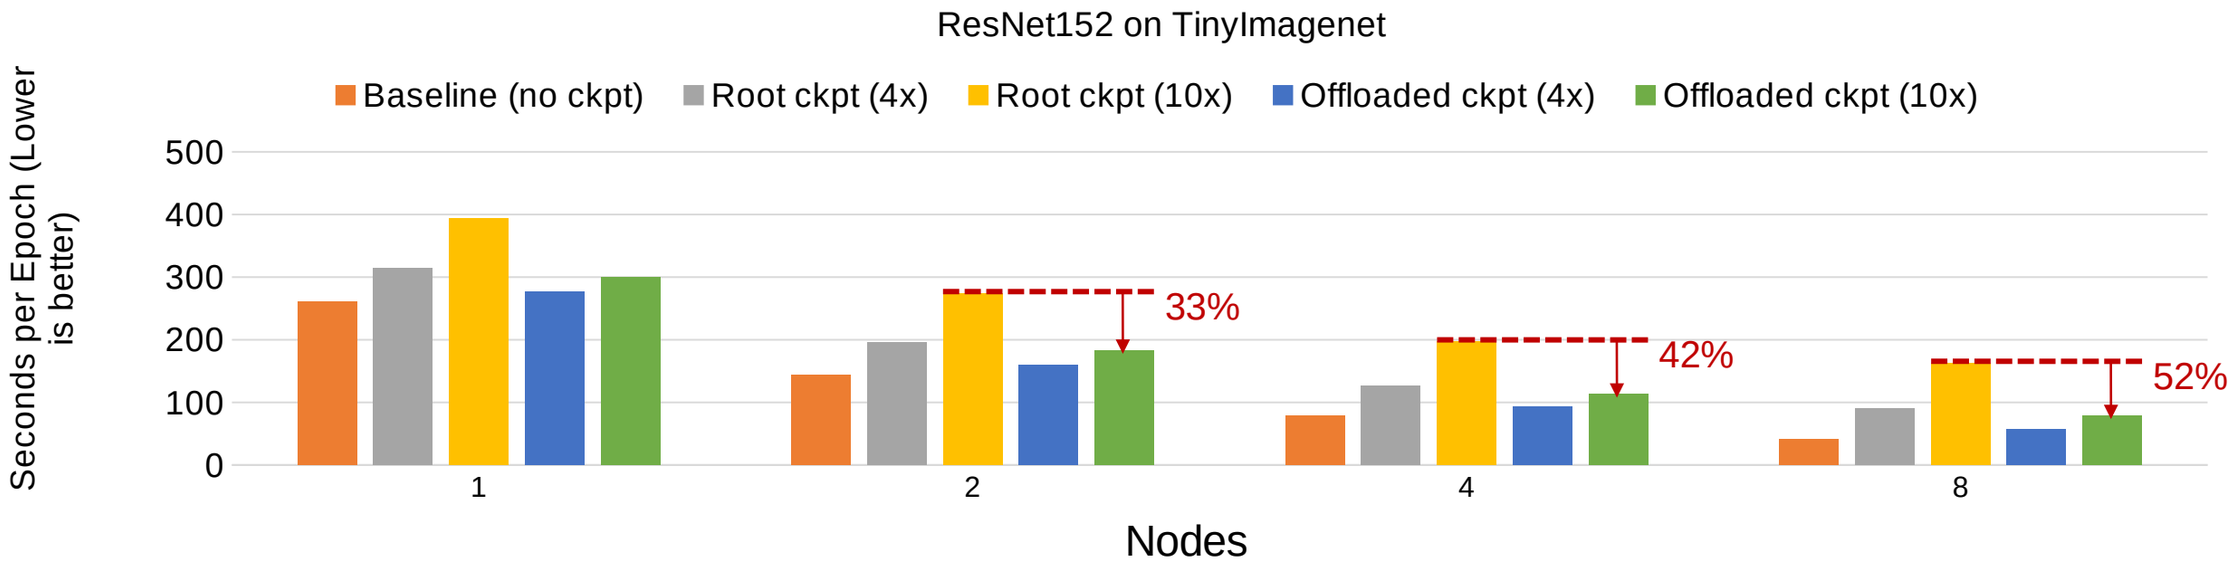

Conclusions

- Provided an overview of the products and services
- Innovative value-added products provide high-performance and scalable solutions for HPC and AI applications while exploiting modern CPU, GPU, and DPU technologies
- Happy to work with interested end customers and/or third-party integrators

Thank You!

Donglai Dai

d.dai@x-scalesolutions.com

contactus@x-scalesolutions.com

X-ScaleSolutions

http://x-scalesolutions.com/