



Mellanox Technologies



HPC User Forum – Fall 2013



Comprehensive End-to-End Software Accelerators and Management

 <p>MPI</p>	 <p>SHMEM</p>	 <p>PGAS</p>	<h3>Management</h3>  <p>UFM Unified Fabric Management</p> 		<h3>Storage and Data</h3> <p>VSA Storage Accelerator</p>  <p>UDA Unstructured Data Accelerator</p> 	
<p>MXM Mellanox Messaging Acceleration</p>	<p>FCA Fabric Collectives Acceleration</p> 					

Comprehensive End-to-End InfiniBand and Ethernet Solutions Portfolio

<h3>ICs</h3> 	<h3>Adapter Cards</h3> 	<h3>Switches/Gateways</h3> 	<h3>Long-Haul Systems</h3>  <p>metroX™</p>	<h3>Cables/Module</h3> 
--	--	--	--	--

- World's first 100Gb/s interconnect adapter
 - PCIe 3.0 x16, dual FDR 56Gb/s InfiniBand ports to provide >100Gb/s
- Highest InfiniBand message rate: 137 million messages per second
 - 4X higher than other InfiniBand solutions
- <0.7 micro-second application latency
- Supports GPUDirect RDMA for direct GPU-to-GPU communication
- Unmatchable Storage Performance
 - 8,000,000 IOPs (1QP), 18,500,000 IOPs (32 QPs)
- New Innovative Transport – Dynamically Connected Transport Service
- Supports Scalable HPC with MPI, SHMEM and PGAS/UPC offloads




Connect **IB**[™]

Enter the World of Boundless Performance

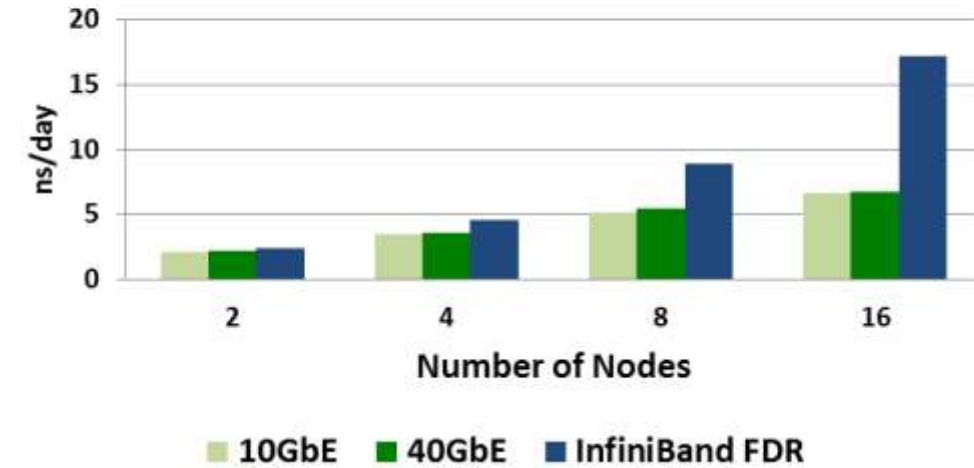
FDR InfiniBand Delivers Highest Application Performance




OpenFOAM Performance (Lid-driven Cavity)




NAMD Benchmark (Platform MPI, ApoA1)



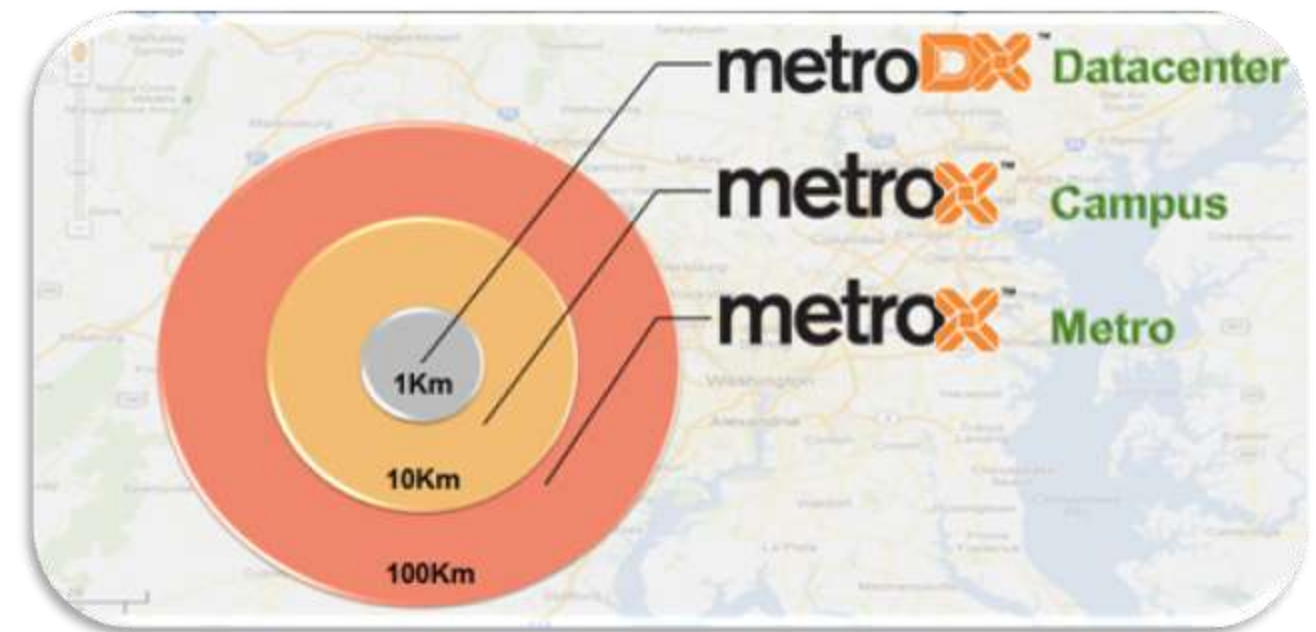
Message Rate



RADIOSS Benchmark (NEON1M11, MPI)



- MetroX™ and MetroDX™ extends InfiniBand and Ethernet RDMA reach
- Enable same subnet on both sides – one system
- Supporting multiple distances – 1KM, 10KM, 40KM
- High port density



40Gb/s over Campus and Metro

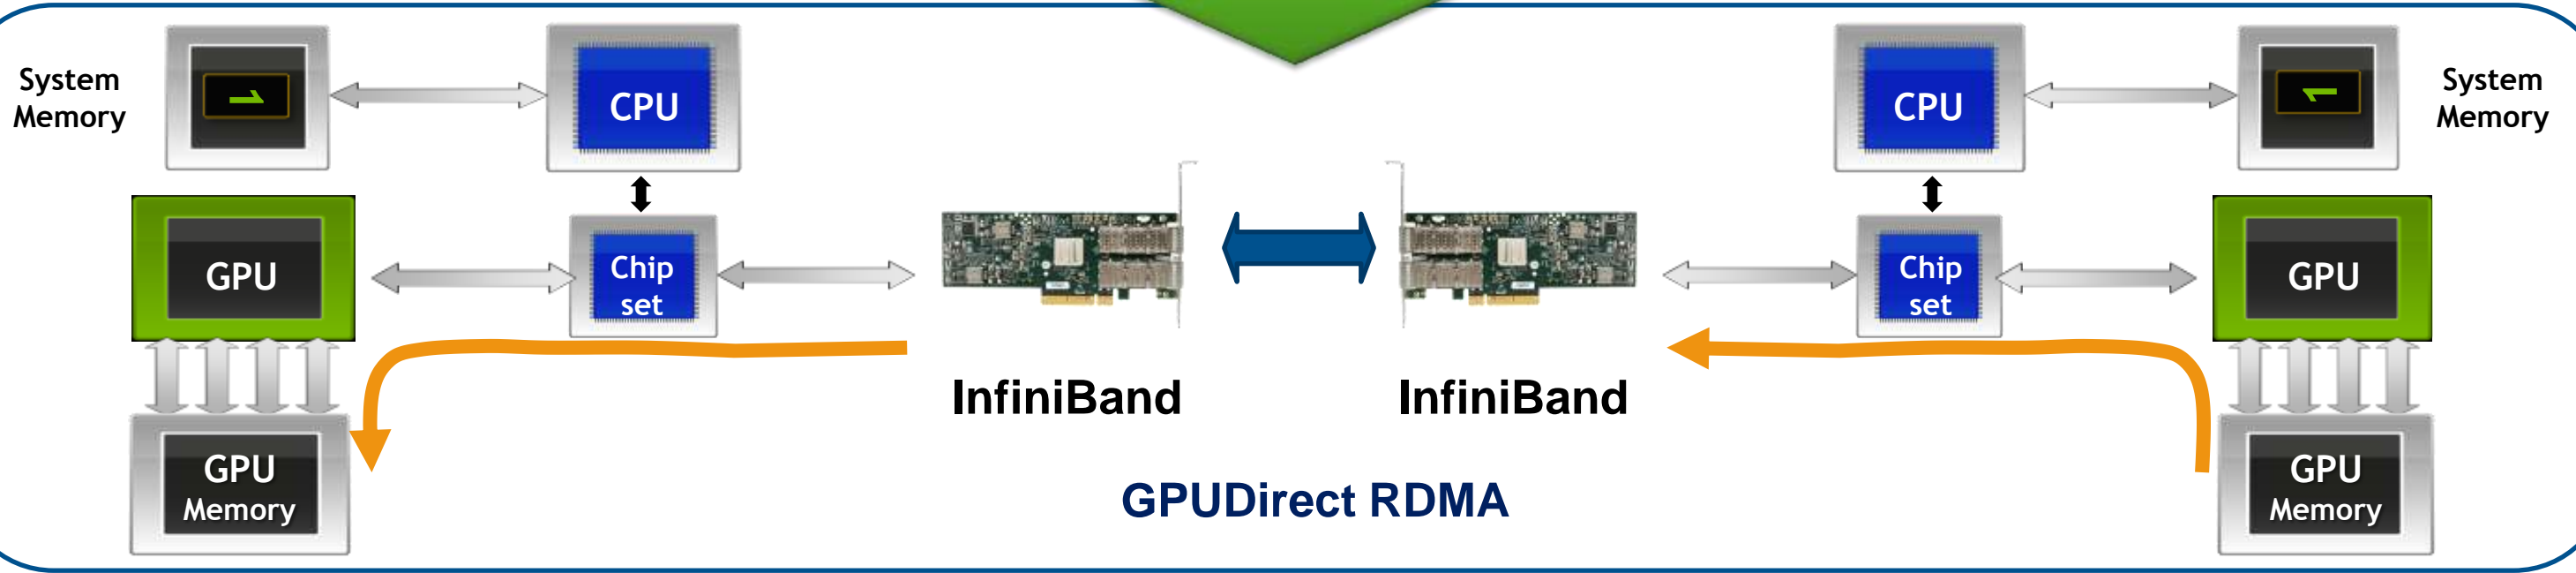
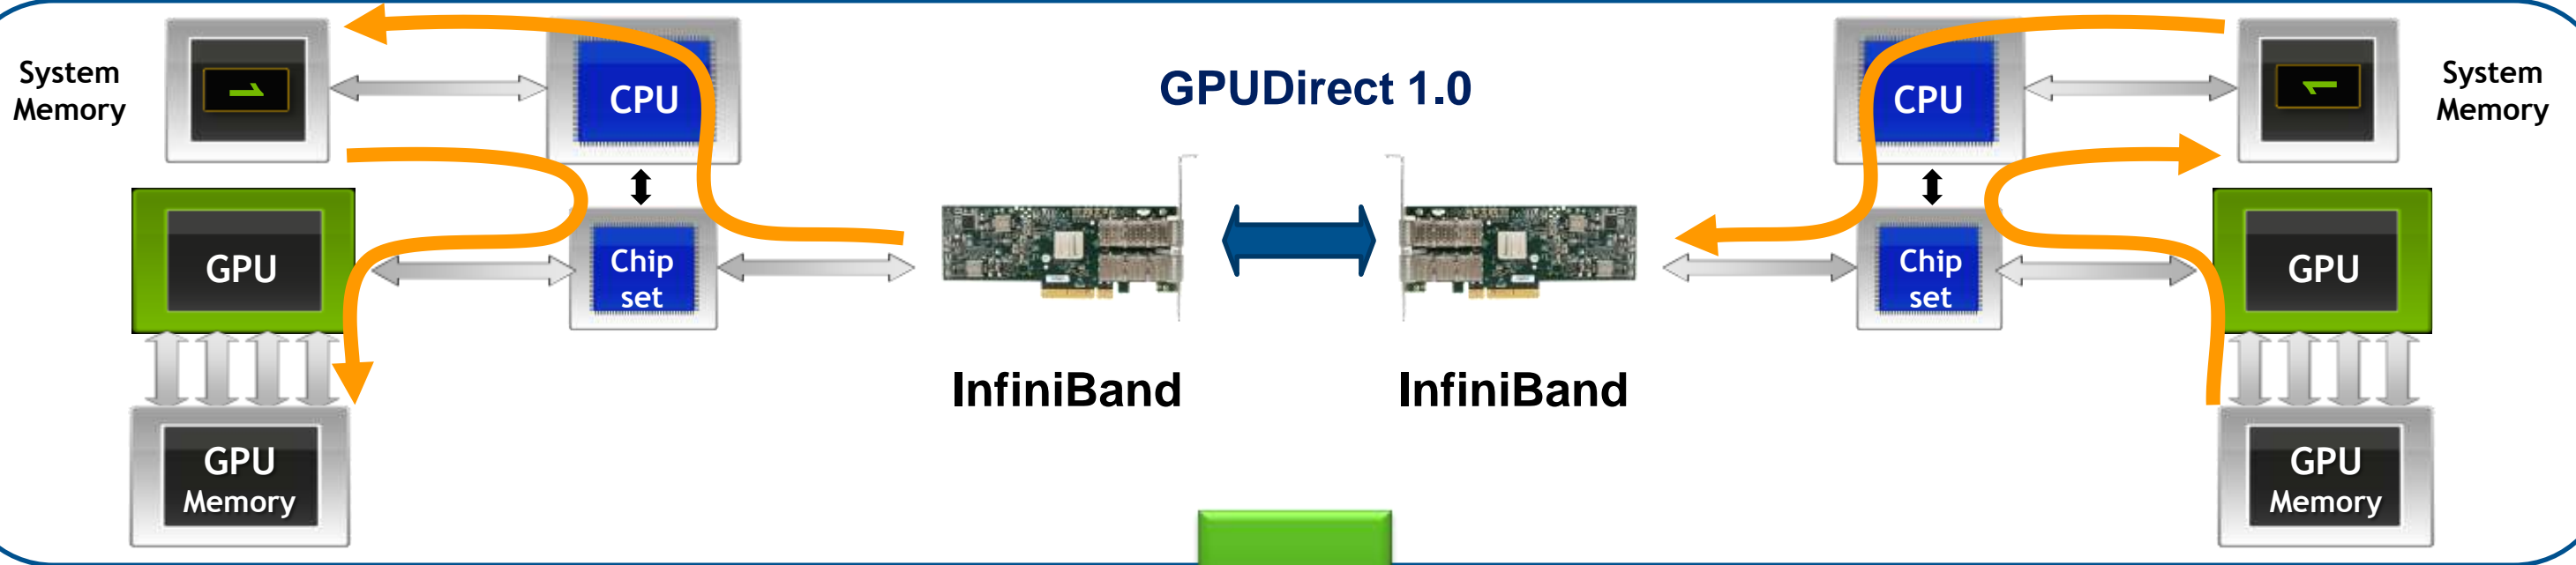
Receive

Transmit

GPUDirect 1.0

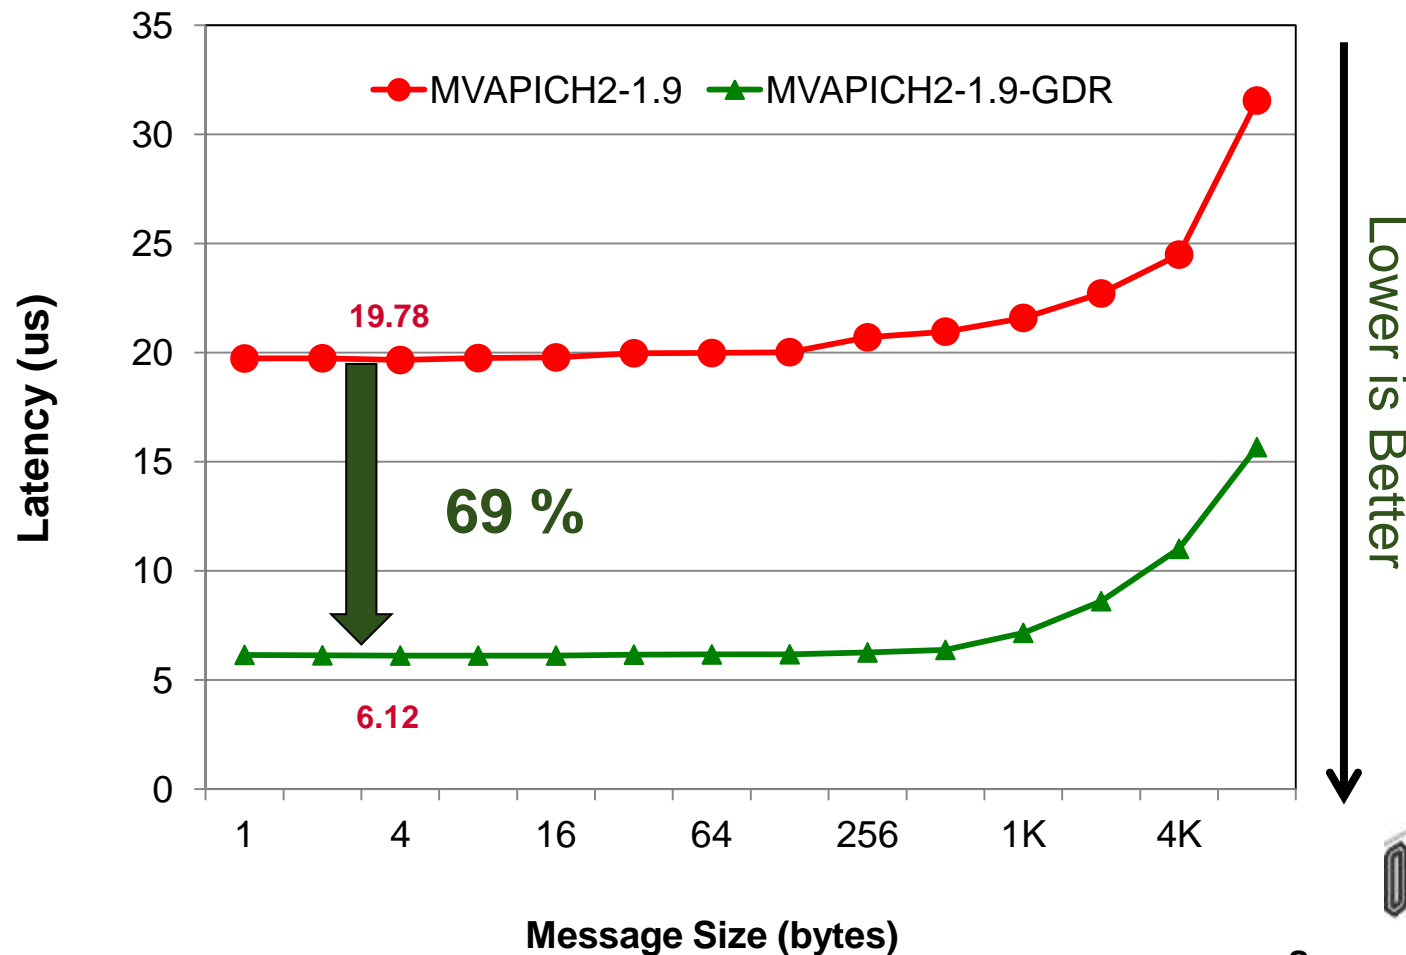


GPUDirect RDMA



GPU-GPU Internode MPI Latency

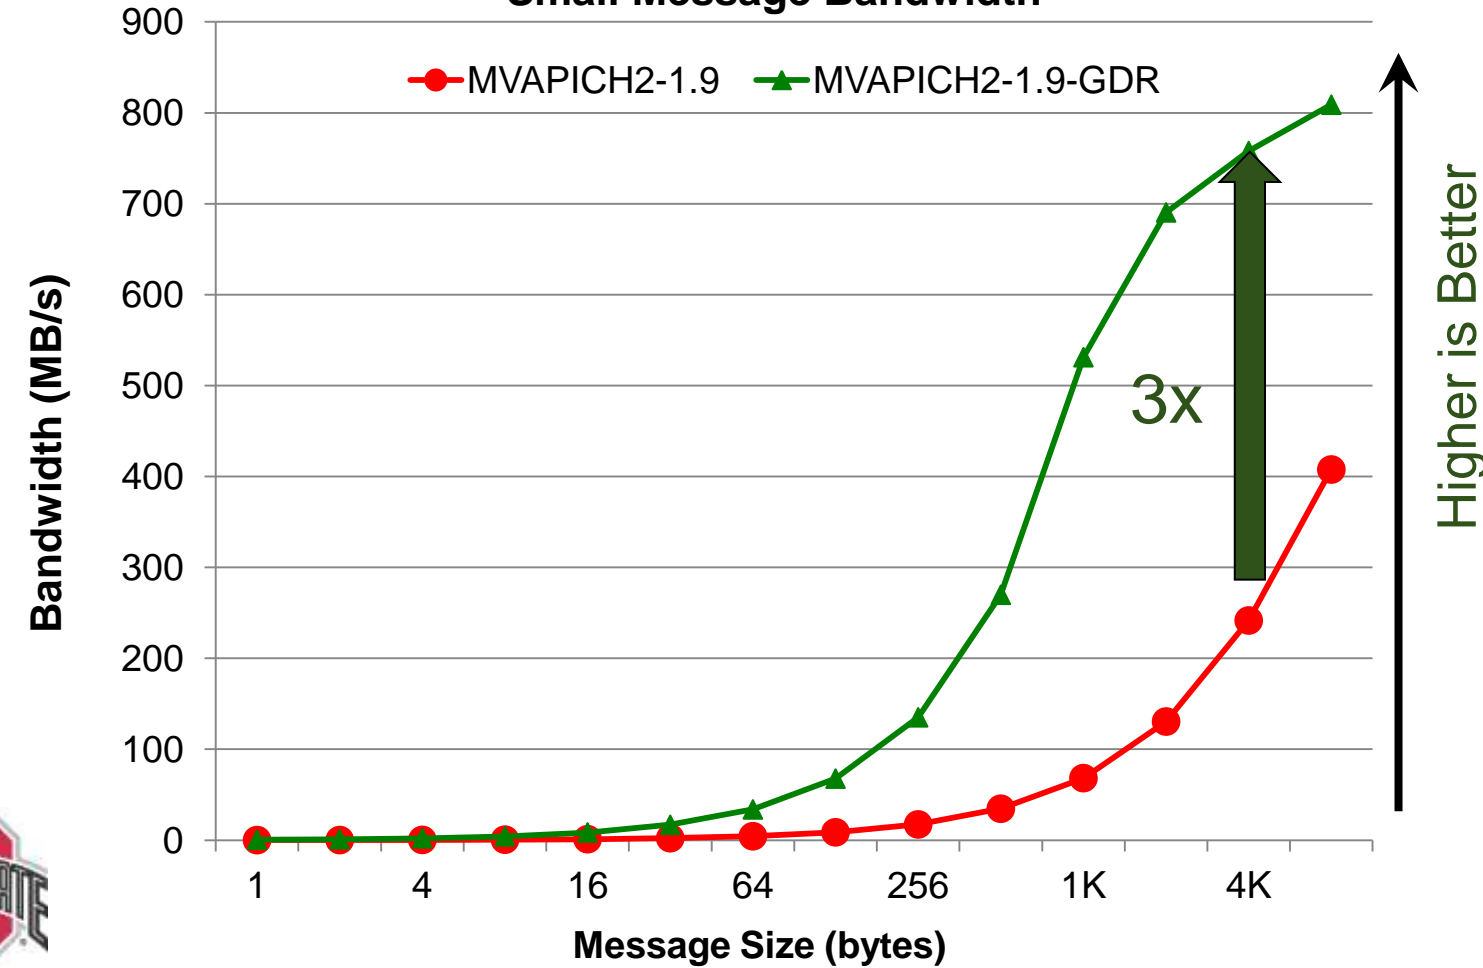
Small Message Latency



Source: Prof. DK Panda

GPU-GPU Internode MPI Bandwidth

Small Message Bandwidth



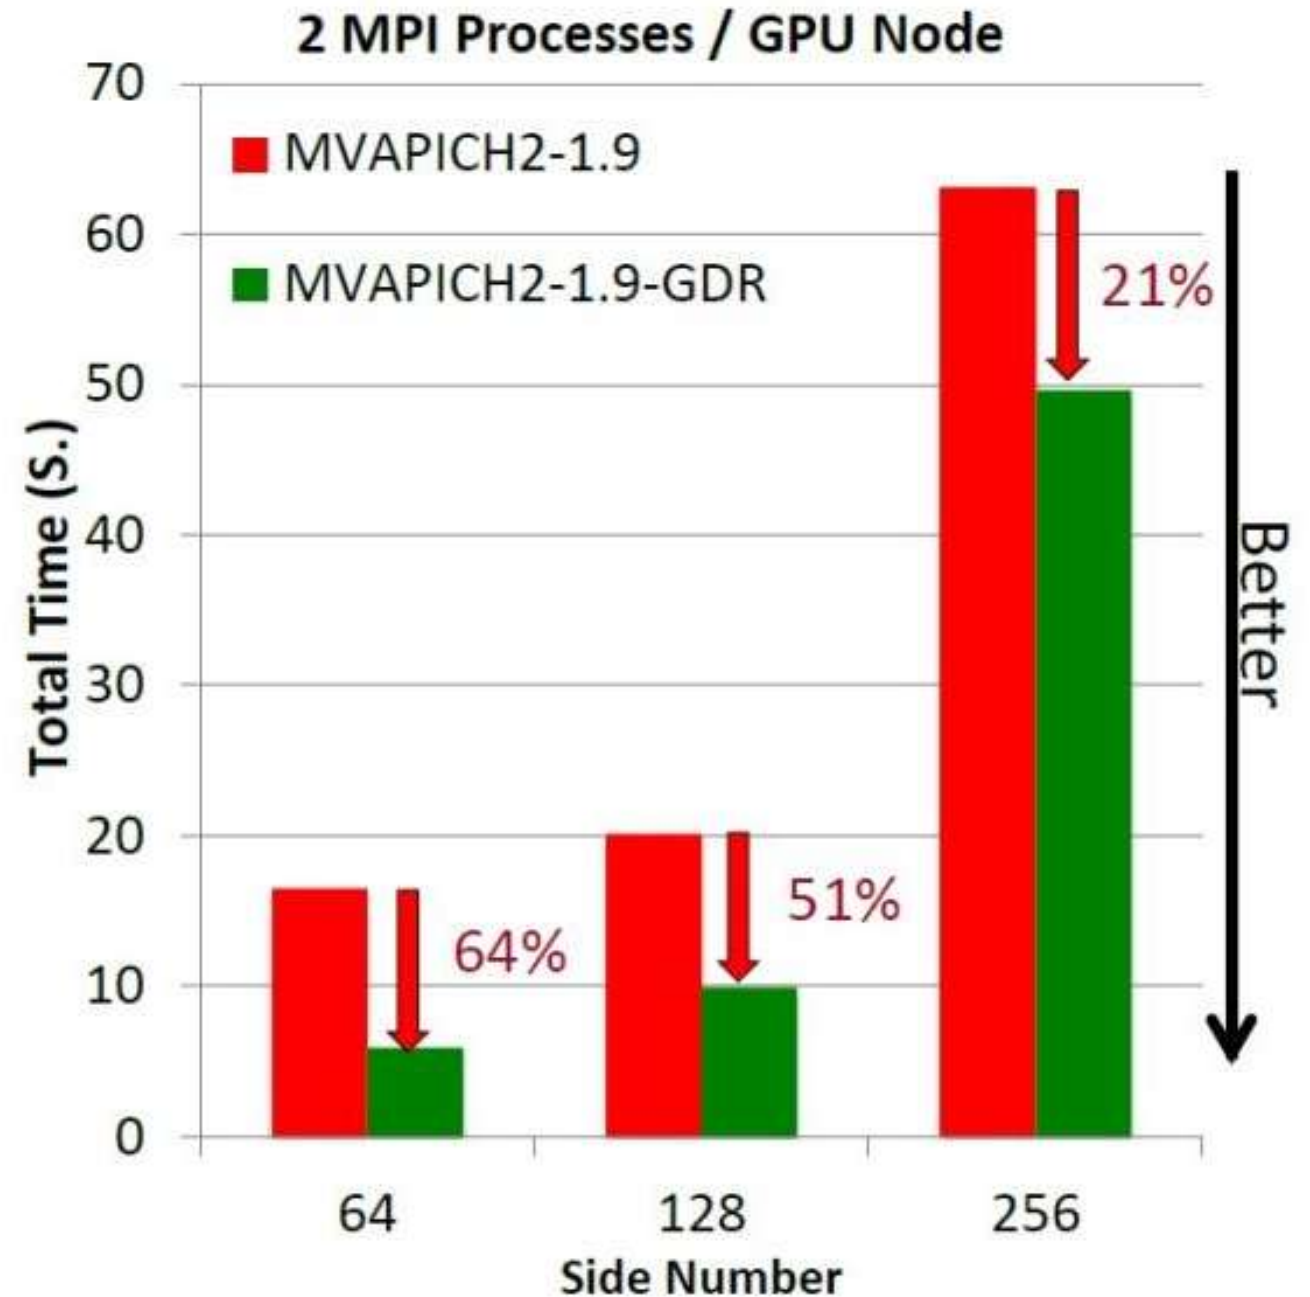
69% Lower Latency

3X Increase in Throughput


Execution Time of HSG (Heisenberg Spin Glass) Application with 2 GPU Nodes




Source: Prof. DK Panda



Bandwidth



Same Software Interface



Latency

Thank You

