Efficient Breadth First Search on Multi-GPU Systems Using GPU-Centric OpenSHMEM

Achievement: Developed Breadth First Search (BFS) implementation for multi-GPU systems. Additionally, improved and evaluated NVSHMEM, an implementation of the OpenSHMEM programming model for GPU-based systems, to efficiently support graph algorithms.

Significance and Impact: This work enables applications using NVIDIA based GPU-systems to use OpenSHMEM programming model for Graph algorithms.

Research Details:

- Improved NVSHMEM, an implementation of the OpenSHMEM programming model for GPU-based systems, for graph algorithms.
- Developed BFS for GPUs using NVSHMEM, and evaluated with the GPU and NVLINK based systems.

Sponser/Facility: Work was performed with support from ORNL and DOD

PI and affiliation: Manjunath Gorentla Venkata from CSMD – Oak Ridge National Laboratory

Team: Manjunath Gorentla Venkata (ORNL), Neena Imam (ORNL), Sreeram Potluri (NVIDIA), Anshuman Goswami (NVIDIA)

Publication: Sreeram Potluri, Anshuman Goswami, Manjunath Gorentla Venkata, Neena Imam. Efficient Breadth First Search on Multi-GPU Systems using GPU-centric OpenSHMEM . In Proceedings of the Fourth OpenSHMEM Workshop, Annapolis, MD, USA, August 2017.

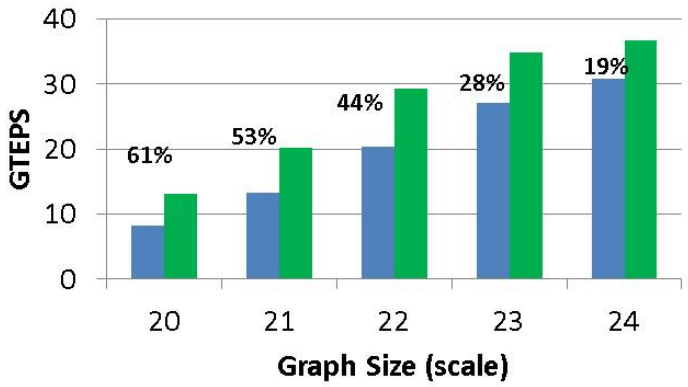

Comparison between CUDA-aware MPI version and NVSHMEM of BFS implementation, on 8 GPUs with NVLINK

Overview:

NVSHMEM is an implementation of OpenSHMEM for NVIDIA GPUs which enables communication from the CUDA kernels. This work is focused on the implementation of BFS for multi-GPU using NVSHMEM interfaces and semantics. Our experimental result shows that NVSHMEM-based BFS implementation outperforms MPI-based implementation by 60%. Also,

our results presented in the paper shows that we achieve 19% peak GTEPS improvement on a 8-GPU system with NVLINK.