SOFTWARE

PRODUCT BRIEF

Mellanox ScalableUPC

Support for Berkeley UPC Parallel Programming Language over InfiniBand

Unified Parallel C (UPC) is an extension of the C programming language designed for high performance computing on large-scale parallel computers. The language provides a uniform programming model for both shared and distributed memory hardware.

UPC Details

There are two types of models for parallel programming. The first is the shared memory model, in which all processes interact through a globally addressable memory space. The other is a distributed memory model, in which each processor has its own memory, and interaction with another processors memory is done though message communication. The PGAS model, or Partitioned Global Address Space, uses a combination of these two methods, in which each process has access to its own private memory, and also to shared variables that make up the global memory space.

The processor memory h a single shared, partitioned address space, where variables may be directly read and written by any processor, but each variable is physically associated with a single processor. UPC uses a Single Program Multiple Data (SPMD) model of computation in which the amount of parallelism is fixed at program startup time, typically with a single thread of execution per processor.

Mellanox ScalableUPC

Mellanox ScalableUPC 2.0 is based on the Berkeley Unified Parallel C project. Berkeley UPC library includes an underlying communication conduit called GASNET, which works over the OpenFabrics RDMA for Linux stack (OFED). Mellanox has optimized this GASNET layer with the inclusion of their Mellanox Messaging libraries (MXM) as well as Mellanox Fabric Collective Accelerations (FCA), providing an unprecedented level of scalability for UPC programs running over InfiniBand.

The use of Mellanox FCA provides for collective optimizations by taking advantage of the high performance features within the InfiniBand fabric, including topology aware coalescing, hardware multicast and separate quality of service levels for collective operations. The figure below shows the type of scalability that can be gained when the collective operations take advantage of FCA.

In addition, the Mellanox MXM acceleration library allows for a high level of performance and scalability for the underlying put/get messages that UPC uses for its node-node communications. MXM is integrated with the MLNX_OFED software stack, and ScalableUPC will automatically take advantage of the performance improvements offered by it when MXM presence is detected. To further understand how these various pieces fit into the software stack refer to Figure 1 below.

HIGHLIGHTS

FEATURES

- Use of symmetric variables and one-sided communication (put/get)
- RDMA for performance optimizations for one-sided communications
- Provides shared memory data transfer operations (put/get), collective operations, and atomic memory operations
- Ability to use as in the standard GASNET method, or with both MXM and FCA for optimized performance

BENEFITS

- Provides a programming library based on Berkeley UPC with Mellanox accelerator extension.
- Seamless integration with MPI libraries and job schedulers
- Improve collective scalability through integrations with Mellanox Fabric Collective Accelerator (FCA).

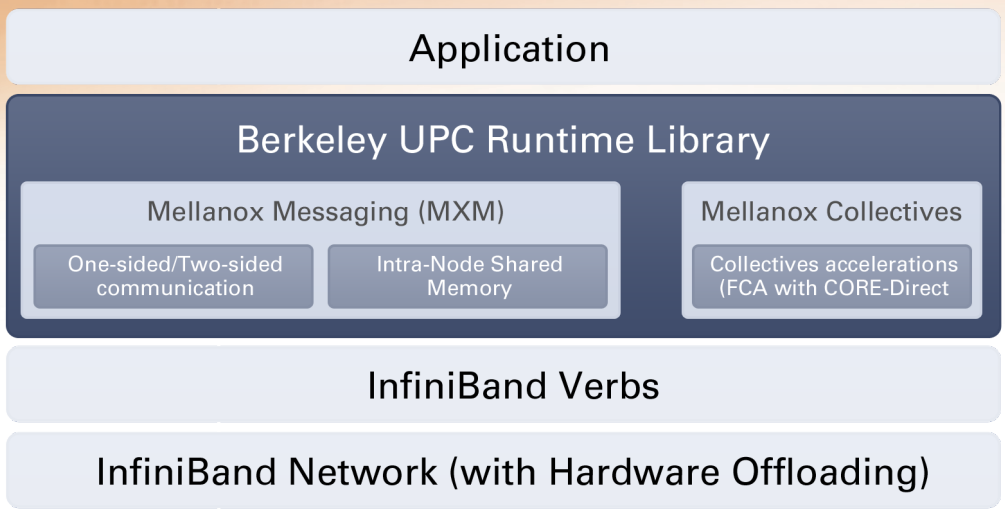

Figure 1.

Mellanox Advantage

Mellanox Technologies is a leading supplier of end-to-end servers and storage connectivity solutions to optimize high performance computing performance and efficiency. Mellanox InfiniBand adapters, switches, and software are powering the largest supercomputers in the world. For the best in server and storage performance and scalability with the lowest TCO, Mellanox interconnect products are the solution.

350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

3888PB Rev 1.0