

The Cell Messaging Layer

The world's fastest MPI implementation for the Cell Broadband Engine

Simplify programming of Cell clusters

One MPI rank per SPE across an entire cluster; SPEs communicate directly with other SPEs regardless of distance

Coordinate heterogeneous processors

RPC facility lets SPEs invoke functions on PPEs and Opterons and receive the results

Keep local store largely free

5–15 KB footprint; only functions actually called occupy space in SPE local store

Invoke familiar MPI functions

MPI_Abort(), MPI_Allreduce(), MPI_Barrier(), MPI_Bcast(), MPI_Comm_get_attr(), MPI_Comm_rank(), MPI_Comm_size(), MPI_Finalize(), MPI_Init(), MPI_Recv(), MPI_Reduce(), MPI_Send(), MPI_Wtime()

Transfer messages at breakneck speeds

Messaging latencies and bandwidths at nearly the performance of raw RDMA

