


High Performance ParalleX (HPX) and runtime support boosts application performance

- N-body Simulation



Comparisons/Results

1.4x over MPI (16,384 cores)

- Mini-Ghost: Boundary Exchange Mini-app

1.13x over MPI+OpenMP (1024 cores)


- Kernel: Stream benchmark
- Kernel: Matrix transposition

1.4x over OpenMP
2.5x over MPI+OpenMP

- Data analytics miniapp (miniTri)

1.14x over OpenMP

- Distributed GPU work



1.5x over native CUDA on 16 GPUs

- Lulesh: Shock Hydrodynamics



1.2x over MPI on Cori (128 cores)
Same as MPI on Cori (4000 cores)

- DSEL and MTL for HPX


- Same Portable code GPU / CPU



DSL for linear algebra through
DOE NNSA DE-NA0002377 (PSAAP2)

Same as MPI (256 cores)

- Photon: Integrated Communication Library



1.34x for 16 byte puts
1.37x for 16 byte gets
over MPI-3 one-sided