U.S. DEPARTMENT OF ENERGY | Sandia National Laboratories

ECP WBS	STPR04 SNL ATDM PMR
PIs	Christian T[illegible] (SNL)
Members	SNL

SAND2018-10929PE

Kokkos R&D: Remote Memory Spaces

Scope and objectives

- Develop Remote Memory Spaces to integrate PGAS like models with Kokkos
- Explore different backend options such as SHMEM and MPI One-Sided
- Evaluate usability and initial performance in a proxy application.

Technical Approach

- Mostly transparent to applications: change a typedef and allocations and data access are distributed
- Four backends are available: libQUO (on-node only), SHMEM, MPI-OneSided and NVSHMEM (supports NVIDIA GPUs in the same node)
- Implemented as Add-On library to Kokkos.

Impact

- Many applications could potentially benefit from PGAS like communication models – in particular on systems with tightly coupled GPUs.
- EXAALT could benefit from this approach if it helps strong scaling small particle ensembles to multiple GPUs
- QMCPACK is interested in this approach in order to have a distributed spline table, which would require random access and is thus not amenable to typical MPI message passing approaches.

Performance Evaluation

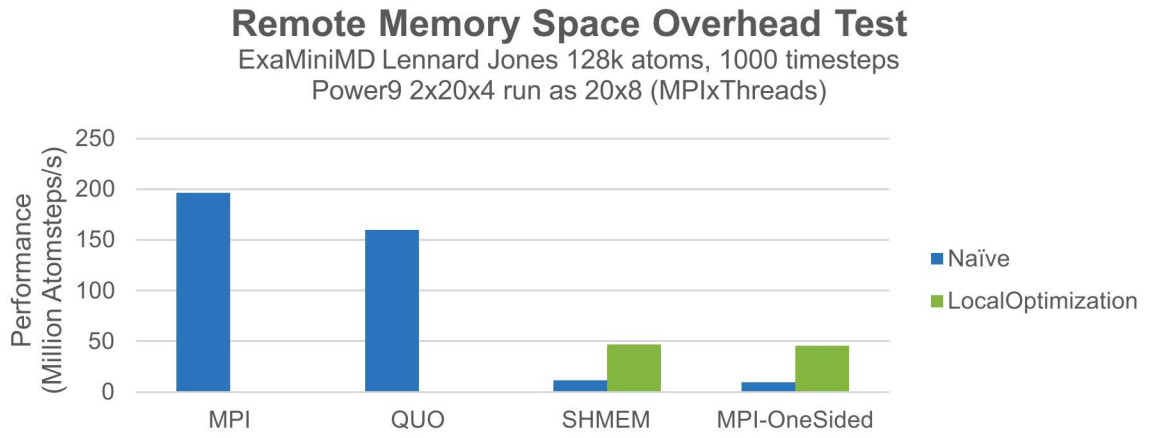

Deliverables STPM12 Milestone 7 Report available at https://confluence.exascaleproject.org
Code Repository Available at https://github.com/kokkos/kokkos-remote-spaces upon request

ECP EXASCALE COMPUTING PROJECT

Acknowledgments

This research was supported by the Exascale Computing Project (ECP), Project Number: 17-SC-20-SC, a collaborative effort of two DOE organizations—the Office of Science and the National Nuclear Security Administration—responsible for the planning and preparation of a capable exascale ecosystem—including software, applications, hardware, advanced system engineering, and early testbed platforms—to support the nation's exascale computing imperative.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.