

OOKAMI PROJECT APPLICATION

Date: September 10, 2021

Project Title: Measuring Performance and Scaling of Kokkos-accelerated Athena++ on Ookami

Usage:

- Testbed
- Production

Principal Investigator: Brian O'Shea

University/Company/Institute: Michigan State University Department of Computational Mathematics, Science and Engineering and Department of Physics and Astronomy

Mailing address including country:

Professor Brian W. O'Shea
Department of Computational Mathematics, Science and Engineering
Michigan State University
Engineering Building
428 South Shaw Lane, 1508F Engineering Building
East Lansing, MI 48824

Phone number: 517-432-0331

Email: oshea@msu.edu

Names & Email of initial project users:

Brian O'Shea, oshea@msu.edu
Forrest Glines, glinesfo@msu.edu
Philipp Grete, gretephi@msu.edu

Usage Description: Over the past three years, we have developed a version of the widely-used Athena++ astrophysical fluid code that achieves performance portability using Kokkos (<https://kokkos.org/>; part of the Exascale Computing Project). This project began by focusing on calculations with large, monolithic grids (K-Athena; Grete, Glines & O'Shea 2021, IEEE Transactions on Parallel and Distributed Systems, 32, 1, 85-97, <https://ieeexplore.ieee.org/document/9143480>; source code at <https://gitlab.com/pgrete/kathena>), and demonstrated true performance portability on a wide variety of architectures and individual systems. Since then, this project has grown into Parthenon, a standalone performance portable adaptive mesh refinement (AMR) framework developed in collaboration with researchers at Los Alamos National Laboratory (<https://github.com/lanl/parthenon>) and the Kokkos developers, and Athena-PK, an application using Parthenon to duplicate the functionality of Athena++ (<https://gitlab.com/theias/hpc/jmstone/athena-parthenon/athenapk>). This code replicates K-Athena's behavior for simulations with large grid patches on both CPUs and GPUs, but has been extensively optimized for high performance for small grid patches (i.e., for calculations

requiring AMR). Tests of Athena-PK on CPUs and NVIDIA and AMD GPUs indicate that this code achieves similar relative performance for large vs. small grid blocks as Athena++ has demonstrated on CPUs alone.

We are applying for time on Oookami because we have not yet had an opportunity to profile and optimize Parthenon/Athena-PK on the A64FX architecture, and would like to do so to ensure that our code works well there (and as a test of Kokkos as a performance portability platform). Ultimately, if we can achieve acceptabler single-node performance and scaling on Oookami, we hope to use the machine for production calculations for science in the future.

Computational Resources:

Total node hours per year: 5,000

Size (nodes) and duration (hours) for a typical batch job: Initially, most runs will be single-node performance tests that will be one node and run for an hour. Scaling and multi-node performance tests will go up to roughly the size of the full machine (~160 nodes) for less than an hour at a time.

Disk space (home, project, scratch): Home: 100 GB (source code, Kokkos, analysis tools, etc.). We do not need project space. Scratch: 10-20 TB, intermittently (for data written out during scaling/performance testing; none of it needs to be saved).

Personnel Resources (assistance in porting/tuning, or training for your users): Minimal. We anticipate that we may occasionally consult with the Oookami support team for help with performance measurement tools for the A64FX architecture, but that's about it.

Required software: Standard libraries: MPI, HDF5. We will also need Athena-PK and Kokkos, but generally install those ourselves on new systems.

If your research is supported by US federal agencies:

Agency: Los Alamos National Laboratory subcontract to Michigan State University (Philipp Grete); National Science Foundation

Grant number(s): LANL PD# 46669; NSF OAC grant #1835426 and AAG #2106575