SAND2019-13462R

SAND2019- 9827 R

ROCM+Intel-PathForward+RemoteSpaces Development

WBS STPR 04 Milestone 25

Authors

Christian Trott (1

Author Affiliations

(1) Sandia National Laboratories

September 29th 2019

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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-NA0003525.

EXECUTIVE SUMMARY

This report documents the completion of milestone STPR04-25 Harden and optimize the ROCm based AMD GPU backend, develop a prototype backend for the Intel ECP Path Forward architecture, and improve the existing prototype Remote Memory Space capabilities. The ROCM code was hardened up to the point of passing all Kokkos unit tests - then AMD deprecated the programming model, forcing us to start over in FY20 with HIP. The Intel ECP Path Forward architecture prototype was developed with some initial capabilities on simulators - but plans changed, so that work will not continue. Instead SYCL will be developed as a backend for Aurora. Remote Spaces was improved. Development is ongoing as part of a collaboration with NVIDIA.

1. INTRODUCTION

The Milestone is tracking development for new architectures relevant for DOE exascale machines and beyond.

2. MILESTONE OVERVIEW

2.1 DESCRIPTION

We will harden the ROCm backend to a point where all Kokkos Core tests are getting passed, and help downstream customers such as KokkosKernels and Trilinos to start exploration of AMD GPUs via the ROCm backend.

We will develop a prototype backend for the Intel ECP platform, with a goal of passing at least 50% of all Kokkos tests.

The remote memory space capabilities will be further developed, with demonstrations in miniApps using at least a SHMEM based backend for CPU architectures and NVSHMEM for NVIDIA GPus.

2.2 EXECUTION PLAN

2.3 COMPLETION CRITERIA

Code being made available to relevant partners via github.

3. TECHNICAL WORK SCOPE, APPROACH, RESULTS

3.1 ROCM BACKEND

The ROCM backend was intended as the primary backend for AMD GPUs. Development on the backend had started in FY17 and the goal was to harden it in FY19 to the point where it could be used for customer applications in preparation for potential Exascale machines using such technology. By Q2 FY19 the work had progressed enough that about 95% of all unit tests passed, and we were looking forward to starting work with downstream packages. Considering that ORNL also announced that Frontier would be based on AMD hardware this work gained additional urgency. Unfortunately, in March 2019 AMD announced the the HCC ROCm backend will be dropped from their plans, and not supported going forward. This means that the successful work on the backend up to that point would be largely useless going forward. AMD said that they will focus all their resources on the HIP model, which previously was only considered as a porting help for CUDA codes. The expanded ECP Kokkos team has started that work, with developers at ORNL leading the effort. The ROCM backend is getting removed from the Kokkos code base with the 3.0 release. We are now also committed via the ECP Kokkos team to deliver the OpenMPTarget backend, which will enable AMD GPU utilization with the Cray compiler. The OpenMPTarget backend has significant pieces working, already but we did not yet have access to compilers able to target AMD GPUs. As an aside: issues with AMD drivers on RedHat based systems, delayed the availability of AMD development clusters (as opposed to some isolated workstations with a single GPU) until the very end of FY20. But it appears those problems has been resolved, helping with accelerated development of the backends for AMD in FY20.

3.2 INTEL ECP PATHFORWARD BACKEND

Information on this work is still considered covered by a RSNDA and can not be described in detail here. Progress had been made on simulators with initial Kokkos capabilities working. Intel plans changed though, and work had to get refocused on the OpenMP Target backend and the later announced Intel Xe Compute architecture with the related SYCL backend. Initial development efforts have started and already resulted in discovery of severe compiler issues which were reported to Intel. We are now committed to enable both OpenMP Target and SYCL on the ANL Aurora platform, with the expectation that future Intel architectures will be able to also execute SYCL and/or OpenMP-Target code as part of Intels OneAPI effort.

3.3 REMOTE SPACES

Remote space development has continued. We developed capabilities to interact with filesystems and checkpointing, and improved the PGAS capabilities. We are now working with NVIDIA directly to improve the NVSHMEM based implementation as part of the workstreams created due to the Summit on Summit meetings. These implementations were made available to various collaborators including NVIDIA to help with further development.

4. RESOURCE REQUIREMENTS

The work performed here required 0.5 FTE, distributed over 8 developers: Steven Bova, Nathan Ellingwood, David Hollman, Dan Ibanez, Duane Labreche, Jeff Miles, Dan Sunderland and Christian Trott.

5. CONCLUSIONS AND FUTURE WORK

One of the primary conclusions from this work is that vendor plans can still change rapidly, even on timeframes which result in severe risks for timely deployment on contracted machines. Radical changes in architecture design, tool chains and programming models happened during FY19 which were not expected. This means the Kokkos team needs to be agile, and improve its ability to react to such changes fast. To help with that we started development of a more systematic approach for backend development. In particular we are now developing a set of unit tests which are significantly more isolated in terms of capabilities which are required to run them. While, the existing unit-tests are very comprehensive they are not designed to help with incremental backend development. For example the tests for parallel dispatch usually use Views for data management, while the unit tests for Views use parallel dispatch as part of the tests. I.e. one has to implement both before either test can be executed. The goal is now to develop a series of tests which can allow the incremental testing of backends while they are developed. In the end we expect to have something like a 70 step plan, which can be followed to implement a full backend by simply implementing the capabilities needed for each test in order.

6. 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.

This work was performed under US Government contract DE-AC52-06NA25396 for Los Alamos National Laboratory, which is operated by Los Alamos National Security, LLC for the U.S. Department of Energy.