

OOKAMI PROJECT APPLICATION

Date: March 20, 2021

Project Title: OpenMP in LLVM (llvm-openmp)

Usage:

- Testbed

Principal Investigator: Barbara Chapman

- University/Company/Institute: Stony Brook University
- Mailing address including country:

Institute for Advanced Computational Science,
Stony Brook University,
Stony Brook,
NY 11794-5250,
USA

- Phone number: (631)-632-2351
- Email: barbara.chapman@stonybrook.edu

Names & Email of initial project users:

- Barbara Chapman <barbara.chapman@stonybrook.edu>
- Tony Curtis <anthony.curtis@stonybrook.edu>
- Shilei Tian <shilei.tian@stonybrook.edu>

Usage Description:

OpenMP is considered the most popular intra-node parallel programming interface in high-performance computing (HPC). Numerous applications, computational libraries, and runtime systems have been successfully parallelized with OpenMP. The LLVM Project is a collection of modular and reusable compiler and toolchain technologies, and may be best known as the basis of the Clang compiler family. Stony Brook University is a member of the OpenMP Architecture Review Board (ARB), and has already contributed many optimizations to

LLVM OpenMP that significantly improve application performance. The current LLVM OpenMP runtime (`libomp`) is well optimized for the Intel platform. Considering the difference between x86-64 and A64FX, optimizations for the Intel platform will likely not work well on A64FX. We plan to optimize `libomp` on A64FX, and to use Ookami to verify our optimization.

Reference:

- <https://www.openmp.org>
- <https://github.com/llvm/llvm-project/tree/main/openmp>

Computational Resources:

- Total node hours per year: estimate 1000
- Size (nodes) and duration (hours) for a typical batch job: There are generally no typical batch jobs, usually single node for development testing. Runs often range from a few minutes to an hour (e.g. building the LLVM project and running some benchmarks).
- Disk space (home, project, scratch): 40GB, 4TB, 4TB

Personnel Resources:

None anticipated.

Required software:

None extra.

If your research is supported by US federal agencies:

- Agency: Department of Energy
- Grant number(s): 17-SC-20-SC

Production projects:

Production projects should provide an additional 1-2 pages of documentation about how

1. the code has been tuned to perform well on A64FX (ideally including benchmark data comparing performance with other architectures such as x86 or GPUs)
2. it can make effective use of the key A64FX architectural features (notably SVE, the high-bandwidth memory, and NUMA characteristics)

3. it can accomplish the scientific objectives within the available 32 Gbyte memory per node