OOKAMI PROJECT APPLICATION

Date: **04/10/2021**

Project Title: Scalable Hierarchical Algorithms for Preconditioning Large-Scale Sparse and Data-Sparse Systems

Usage:

☒ Testbed

☐ Production

Principal Investigator:

University/Company/Institute: Stony Brook University

Mailing address including country:

Institute for Advanced Computational Science
Stony Brook University
Stony Brook, NY 11794

Phone number: 631-632-2339

Email: Xiangmin.jiao@stonybrook.edu

Names & Email of initial project users:

Nilesh Chaturvedi, nilesh.chaturvedi@stonybrook.edu
Qiao Chen, qiao.chen@stonybrook.edu
Hongji Gao, hongji.gao@stonybrook.edu
Xiangmin Jiao, Xiangmin.jiao@stonybrook.edu
Xuebin Wang, xuebin.wang@stonybrook.edu

Usage Description:

This project aims to develop scalable parallel algorithms and an open-source software package for preconditioning large-scale sparse or data-sparse matrices. We have developed robust and efficient preconditioners for large-scale sparse linear systems that scale nearly linearly in the problem sizes. The serial version of the software will be released in May. We plan to extend the algorithm to hybrid distributed-and-shared memory machines in the coming months using a hybrid MPI+OpenMP programming model. In the later part of the project, we also plan to extend our algorithm to support data-sparse matrices, such as those from quantum chemistry. The Ookmai platform will be an essential testbed for our developed algorithms and software to be future-proof instead of being limited to the Beowulf clusters. We plan to test our developed algorithms on Ookami to solve nonlinear partial differential equations, nonlinear optimization methods, eigenvalue problems, etc.

Computational Resources:

Total node hours per year: 1,000

Size (nodes) and duration (hours) for a typical batch job: 2--5 nodes for 1 hour

Disk space (home, project, scratch):
Home: 40GB for analysis, visualization, and batch scripts
Project: 4 TB for storing test results
Scratch: 4 TB for the code and output data

Personnel Resources:
N/A

Required software:

- C++ compilers with OpenMP support
- LAPACK
- MPI
- HDF5
- Python with SciPy, mpi4py

If your research is supported by US federal agencies:

Agency: N/A

Grant number(s):

Production projects:

Production projects should provide an additional 1-2 pages of documentation about how
(a) the code has been tuned to perform well on A64FX (ideally including benchmark data comparing performance with other architectures such as x86 or GPUs)
(b) it can make effective use of the key A64FX architectural features (notably SVE, the high-bandwidth memory, and NUMA characteristics)
(c) it can accomplish the scientific objectives within the available 32 Gbyte memory per node