

Parallel I/O Performance Characterization of Columbia and NEC SX-8 Superclusters

Subhash Saini¹, Dale Talcott¹, Rajeev Thakur², Panagiotis Adamidis³, Rolf Rabenseifner⁴ and Robert Ciotti¹

¹*NASA Advanced Supercomputing Division
NASA Ames Research Center
Moffett Field, California, USA
subhash.saini@nasa.gov, dtalcott@mail.arc.nasa.gov,
ciotti@nas.nasa.gov*

²*Mathematics and Computer Science Division
Argonne National Laboratory, Argonne
Argonne, IL, USA
thakur@mcs.anl.gov*

³*German Climate Computing Center
Hamburg, Germany
adamidis@dkrz.de*

⁴*High Performance Computing Center
University of Stuttgart
Stuttgart, Germany
rabenseifner@hlsr.de*

Many scientific applications running on today's supercomputers deal with increasingly large data sets and are correspondingly bottlenecked by the time it takes to read or write the data from/to the file system. We therefore undertook a study to characterize the parallel I/O performance of two of today's leading parallel supercomputers: the Columbia system at NASA Ames Research Center and the NEC SX-8 supercluster at the University of Stuttgart, Germany. On both systems, we ran a total of seven parallel I/O benchmarks, comprising five low-level benchmarks: (i) IO_Bench, (ii) MPI Tile IO, (iii) IOR (POSIX and MPI-IO), (iv) b_eff_io (five different patterns), and (v) SPIOBENCH, and two scalable synthetic compact application (SSCA) benchmarks: (a) HPCS (High Productivity Computing Systems) SSCA #3 and (b) FLASH IO (parallel HDF5). We present the results of these experiments characterizing the parallel I/O performance of these two systems.