

STRIDE

Summary Version

1.0

Purpose of Benchmark

STRIDE is designed to severely stress the memory subsystem on a node using a series of sequential kernels.

Characteristics of Benchmark

The STRIDE benchmark consists of eight separate benchmark kernels designed to severely test and stress the memory subsystem of a single node of a computational platform. The tests are STRID3, VECOP, CACHE, STRIDOT, and CACHEDOT. The first three benchmarks include C and Fortran language versions. All of the benchmarks utilize combinations of loops of scalar and vector operations and measure the MFLOP rate delivered as a function of the memory access patterns and length of vector utilized. The observed rates of computation of the various access patterns within an individual test can then be compared to provide an understanding of the performance implications.

Mechanics of Building the Benchmark

A script named “buildit” is provided to build the benchmark kernels. Users are free to change the compiler used as well as any compiler options when building the application

Mechanics of Running the Benchmark

In the root directory for the benchmark, a script name “runit” is provided to run the benchmark. In the “src” subdirectory there is a file named “script.cache” that needs to be edited to correctly reflect the number of CPU cores on a node.

The eight tests are run by executing the “runit” script, generating an output file for each core in a multicore node. On a 2.2-GHz AMD, quad socket, dual core AMD Opteron system, the tests take approximately 45hours to complete.