

Phloem MPI Benchmarks

Summary Version

1.0

Purpose of Benchmark

The Phloem MPI Benchmarks provide a collection of independent MPI benchmarks which can be used to measure various aspects of MPI performance including interconnect messaging rate, point-to-point latency, aggregate bandwidth, and collective latencies. Please see the documentation provided with each benchmark for additional information.

Characteristics of Benchmark

The benchmark suite currently consists of four benchmarks:

- linktest: MPI aggregate link bandwidth benchmark
- mpiBench: MPI collective latency benchmark
- mpiGraph: MPI task pair bandwidth benchmark
- Presta com: MPI latency and aggregate bandwidth benchmark
- SQMR: MPI messaging rate benchmark
- torustest: MPI link bandwidth range benchmark

Limitations of Benchmark

The various benchmarks may require specific task and node counts or specific node allocation.

Mechanics of Building Benchmark

Platform-specific build configurations can be created by editing the top-level Makefile.inc file and running 'make'.

Mechanics of Running Benchmark

Each benchmark will have specific task, node, and argument requirements. The script "run_script" has been provided with the intent of starting minimal example runs of the individual benchmarks by running 'make run' at the command-line. Example benchmark run commands can be generated by running 'make commands'.

Verification of Results

All reports should generate text report output. The run_script script saves benchmark results in a file with verbose run information in the file name.