

High Speed Interconnection Network Performance Studies

By: Erich Doll UNM, Ruben Salazar UNM, Taylor Owen NMSU

Mentors: Hsing-Bung Chen HPC-5 Andree Jacobson UNM Carolyn Connor HPC-5

Infiniband

Used for connections between processor nodes and I/O nodes such as storage devices.

- Single Data Rate (SDR)
 - Transfers data at 10 Gbps at 4x
- Dual Data Rate (DDR)
 - Transfers data at 20 Gbps at 4x
- Pros
 - Low latency
 - High bandwidth
- Cons
 - Not routable

10GigE

Used in data centers where traffic routing is important.

- Pros
 - Routable
 - Easy to implement in current architectures
- Cons
 - High latency
 - Low bandwidth

Importance

This project is to investigate how to increase application performance on a network cluster with a real-world benchmark. The real world benchmark implemented in this project uses the Fastest Fourier Transform in the West (FFTW) to measure latency and bandwidth. These results help applications run faster and more efficient on a network cluster.

Procedure

- Fastest Fourier Transform in the West (FFTW) is used for taking signal samples and removing the noise using an efficient algorithm. A practical example is using FFTW to filter radio telescope data.
- Message Passing Interface Profiling (MPIP) records data about the amount of communication performed during the run time of the application.
- A tweaked version of SDR and DDR were used to increase the bandwidth and lower the latency.
- A variety of data sizes and repetitions were tested to measure the performance trends of the FFTW benchmarks.
- Used Ohio State University (OSU) benchmarks to test the bandwidth and the latency of the different interfaces.

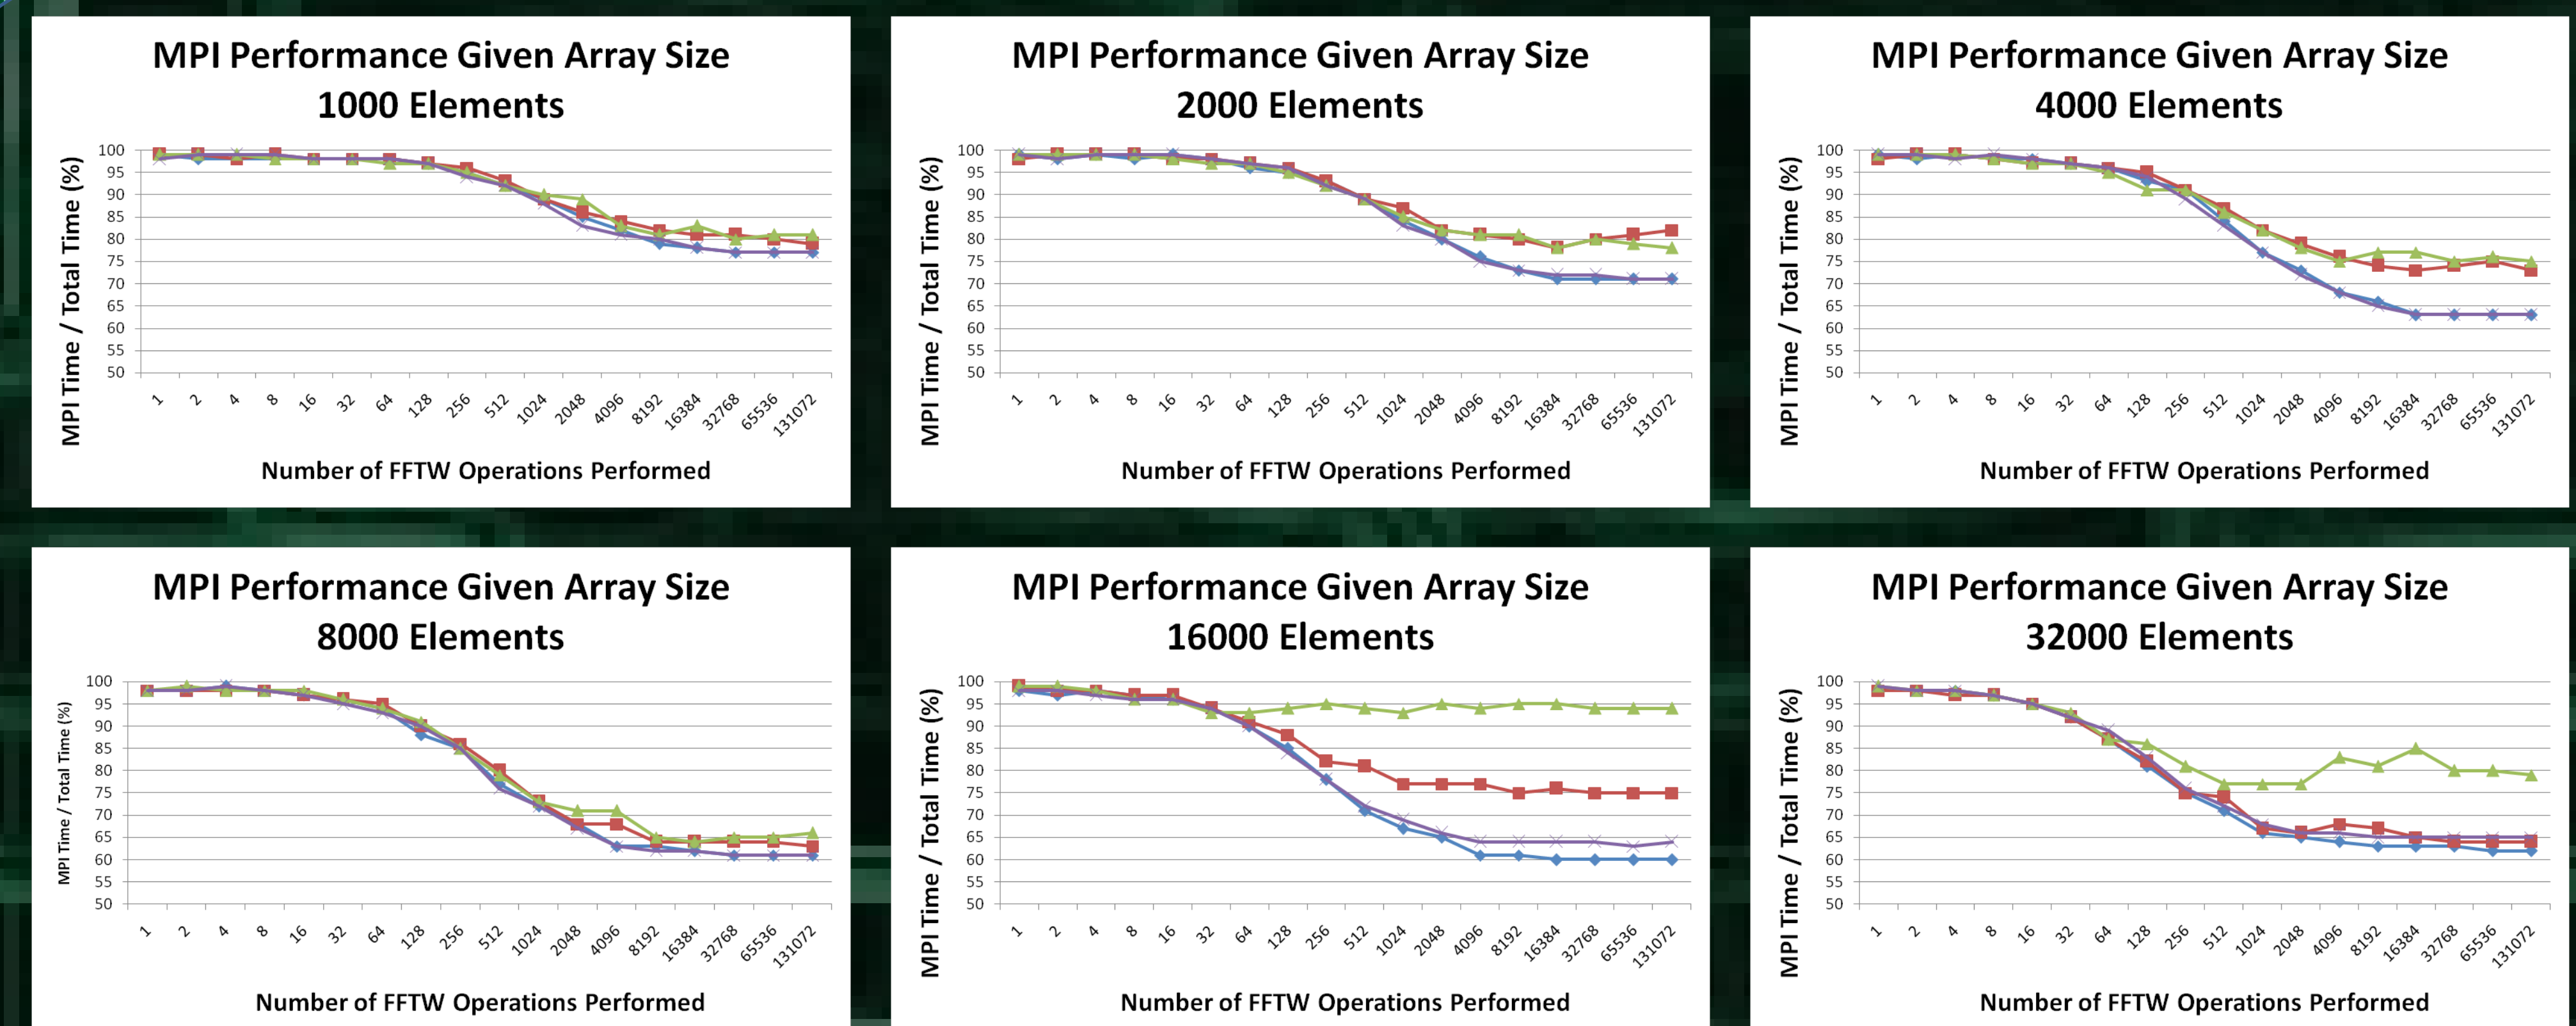
Conclusion

•With FFTW the graphs show the settling effect of the runs due to inconsistent connections. This is why the graphs start flat and then slopes to the stable state. This is where the MPI percentage is constant for any number of repetitions. Thus the optimal point for operation, given the element size, is where the number of repetitions are small and where the stable MPI percentage is reached. The 10 Gigabit Ethernet is not present in these tests due to driver issues.

•The OSU benchmarks test the bandwidth and latency in extreme conditions. Since DDR has high bandwidth and low latency it performed the best. After tweaking the settings of DDR it out performed SDR and 10 Gigabit Ethernet by a significant amount.


Benchmark Results

FFTW with MPIP



Legend: SDR (blue), DDR (red), DDRT (green), SDRT (purple). Lower MPI percentage is better.

Bandwidth and Latency (OSU Benchmark)



Legend: 10GigE (blue), DDR (red), DDRT (green), SDR (purple), SDRT (black). Higher Bandwidth is better. Lower latency is better.