

Comparision of Benchmarks Across CIRC Supercomputing Architectures

Peter Fogg, Baowei Liu
Center for Integrated Research Computing

Abstract


The Center for Integrated Research Computing provides technology and support for researchers at the University in all disciplines, using high-performance computing systems Bluehive and Blue Gene. Here we present the benchmark testing results done on these systems to demonstrate the performance of CIRC supercomputing resources. The software we used include a molecular dynamics software package called Amber, and a standard supercomputing linear algebra benchmark to , Linpack. The results include the performance of both packages on different numbers of cores and the performance of Amber on GPUs.

Systems

Blue Gene



Blue Hive



- 1024 nodes
- 4096 CPU cores
- 2 TB ram
- 152 nodes(Ten of which are equipped with 20 GPU cards)
- 1496 CPU cores
- 3.3 TB ram

High Performance Computing

Modern supercomputers gain their computational power and speed from parallel processing. High performance computing (HPC) is the use of parallel processing for speeding up advanced application programs running. Typical parallel programming tools include MPI, openMP for multicore CPUs and CUDA/OpenCL for GPUs.

Results

Amber Scaling Tests

AMBER is an software package that simulates a family of force fields for molecular dynamics of biomolecules. Figures 1 and 2 show the scaling test results of one typical Amber benchmark Cellulose run with over 400,000 atoms on BlueHive. As more cores are added, the running time decreases. Above eight cores, though, the simulation slows. This is due to the fact that one node contains eight cores, and above that number, multiple nodes must communicate across Ethernet. This is a relatively slow process, and the running time increases accordingly. Blue Gene contains a high-speed optical network for communication, and a library which allows many processors to communicate and share data so the same benchmark scales much better (Figures 3 and 4)




Figure 1: Cellulose on BlueHive




Figure 2: Cellulose speedup on BlueHive




Figure 3: Cellulose on Blue Gene

Linpack Benchmarks

Linpack is a software library for performing numerical linear algebra on computers. Linpack benchmarks measure a computer system's floating computing power by checking how fast the computer solves a dense n by n system of linear equations $Ax = b$, which is a common task in engineering. The computing performance is given in FLOPS, floating-point operations per second. One GFLOPS is equal to 1 billion FLOPS. ON BlueHive, we see the rise in speed slowing down due to internode communication, similar to the Amber benchmarks (Figure 5). On Blue Gene, however, Linpack scales much better, without any significant slowdown as more processors are added (Figure 6). Again, this is due to Blue Gene's high-speed communications between nodes.




Figure 4: Cellulose speedup on Blue Gene




Figure 5: Linpack on BlueHive




Figure 6: Linpack on Blue Gene

Amber Benchmarks on GPU

Modern GPUs (Graphics Processing Unit), which are initially used for very compute-intensive graphics tasks like decoding real-time high-resolution video, have evolved into highly parallel multi-core systems and can be used to do high performance general purpose computations. Nvidia's CUDA platform was the earliest widely adopted programming model for GPU computing. Amber benchmarks can be run on a Nvidia GPU with CUDA. Figure 7 compares the fastest running times of Amber simulations using CPUs(8 or 16 processors per run) and the running time on a single GPU on BlueHive. The GPUs complete the benchmarks much faster than CPUs, due to their highly parallel architecture.




Figure 7: Amber Benchmarks on GPU and CPUs of BlueHive




Figure 8: Amber Benchmarks Speedup on GPU

Conclusion

We run Amber and Linpack benchmarks on the CIRC supercomputing resources. Both the Amber and the Linpack scaling tests show that Blue Gene has a better scaling performance due to its high-speed optical network for communication. Tests on BlueHive GPU nodes also show that Amber Benchmarks run faster on GPUs than on CPUs.