

Hybrid Parallel Programs

Henry Neeman and Charlie Peck
Blue Waters/UPEP Institute @ NCSA
Shodor Education Foundation
May, 2010

How Did We Get Here?

Almost all of the clusters provisioned now and for the foreseeable future are constellations, that is they are composed of nodes, each with 1 or more sockets holding CPUs with 2 or more cores, connected by a high-speed network fabric.

With their additional levels of memory hierarchy these constellation clusters are more difficult to use efficiently than their predecessors, yet they offer the promise of significantly greater cycles available for science. The Blue Waters computational resource is a good example of this trend.

Often a hybrid approach, utilizing 2 or more of MPI, OpenMP, pthreads and/or GPGPU techniques, can utilize these computational resource more efficiently and effectively than any one of them on their own.

Common Parallel Programming Paradigms, Strengths and Weaknesses

- OpenMP
 1. Strengths - relatively easy to use, portable, relatively easy to adapt to existing serial software, standard standard, low latency/high bandwidth communication, implicit communication model, dynamic load balancing
 2. Weaknesses - shared memory model only, *i.e.* one “node”, support for C/C++ and FORTRAN only, no explicit control over thread creation and rundown

- Message Passing Interface (MPI)
 1. Strengths - support for C/C++, FORTRAN, Perl, Python and other languages, scales beyond one shared memory image “node”, widely used in the scientific software community, enables you to harness more memory which translates to bigger science
 2. Weaknesses - non-standard standard to some extent, can support shared memory model but not all bindings do it efficiently, can be difficult to program, high latency/low bandwidth communication, explicit communication model, load balancing can be difficult

- GPGPU with CUDA

1. Strengths - ability to harness significant computational cycles, becoming widely used in the scientific software community
2. Weaknesses - can be difficult to program, non-portable proprietary language, requires re-thinking many problems to take advantage of the high degree of parallel cardinality

- pthreads

1. Strengths - better granularity of control over thread model than OpenMP, fairly portable, explicit control over thread creation and rundown
2. Weaknesses - can be difficult to program, only fairly portable

- Field Programmable Gate Array (FPGA)
 1. Strengths - encode any algorithm in hardware
 2. Weaknesses - can be difficult to program, not portable, expensive hardware

Hybrid Parallel Algorithms

- OpenMP + CUDA
- OpenMP + FPGA
- MPI + OpenMP
- MPI + CUDA
- MPI + FPGA
- MPI + pthreads
- MPI + OpenMP + CUDA
- MPI + pthreads + CUDA

Designing, Building and Debugging Hybrid Parallel Programs

- The basic approach is to find the most efficient way to do the work, whether it be OpenMP (CPU), CUDA (GPGPU) or pthreads or a blend of them, and then the most efficient way to distribute the data and harvest the results via MPI.
- Study the communication pattern(s) to insure that the algorithm maps to the architecture.
 1. Embarrassingly parallel
 2. Loosely coupled through tightly coupled
- Look for opportunities to overlap computation and communication, this is a key attribute of efficient hybrid parallel programs.

- Don't design/implement an algorithm that requires more than one type of parallelism to be enabled for it to run. This will make testing and debugging *much* harder than it needs to be.
- Better to {OpenMP, CUDA, Pthread} your MPI code than the other way around. MPI is harder, do that first, then work on the "on-node" parallelism.
- The simplest and least error prone architecture is to use MPI calls only outside of any parallel regions {OpenMP, pthreads} and only allow the master thread to communicate between MPI processes. It's also possible to use MPI calls within parallel regions if you are using a thread-safe MPI binding (not all of them are).
- Debug by running on one node and testing the {OpenMP, CUDA, Pthreads}, then on 2-n nodes with just the MPI enabled to verify the data transfer.

- Take care to have no more than about one process/thread doing network communication contemporaneously, contention for the network port will quickly become a bottleneck unless there is a minimal amount of communication.
- Take care to have no more than about two processes/threads performing CUDA calculations contemporaneously, contention for the I/O bandwidth to/from the GPGPU card and the GPGPU's cores can become a bottleneck.