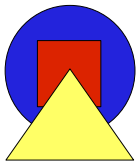

Many-Core CPUs - Parallel Computing in HEP

Alf Wachsmann

Stanford Linear Accelerator Center (SLAC)

alfw@slac.stanford.edu

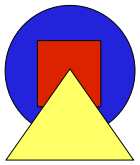

The Problem...

- The clock speed race is over and the core race has begun
- CPUs will move from multi-core to many-core parallel computers
- Today's few-core systems can be equipped with enough memory, disks, network to make them look like several sequential CPUs in one computer
- This will not work for many-core CPUs:
  - Not enough memory slots
  - Not enough bandwidth from single core to memory
  - Not enough network connectivity
  - Not enough local disk space

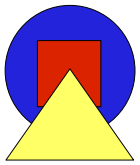

...The Problem...

- **Every** program that needs speedup cannot rely on clock frequency but must efficiently use many-core CPUs
  - Simple threading is not enough
  - These are **cores** that share a lot of infrastructure inside the CPU
  - Parallelization of numerical routines inside programs
- Solution: classical shared-memory parallel programming inside each application!

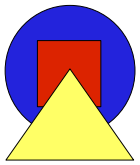

...The Problem

- NVIDIA's GeForce 8800 GTX graphics card has GPU with 128 cores
- Some Cisco CPUs in networking gear have 192 cores
- Sun has many-core SPARC CPUs on their road map
- Dual-socket quad-core systems will already be a problem

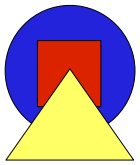

Many-Core CPUs

- The "good" news: Everybody is in the same boat
  - Classical HPC
  - Office applications
  - Web browsers
  - Etc.

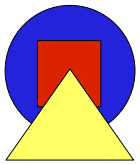

Parallel Computing Today...

- Two different techniques for two different type of computers:
  - Shared memory programming inside a single computer with OpenMP (Open specifications for Multi Processing)
  - Cluster computing with MPI (Message Passing Interface)
- Both APIs are standardized with a wide variety of (partly vendor specific) implementations
- Both will work on many-core CPU systems but...

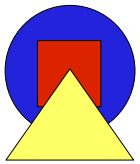

...and Tomorrow

- Some approaches underway to make parallel programming easier
  - DARPA's High Productivity Computing Systems http://www.highproductivity.org/
    - Sun's Fortress Language http://research.sun.com/projects/plrg/
  - Berkeley's "Landscape of Parallel Computing Research" http://view.eecs.berkeley.edu/wiki/Main_Page
- Don't hold your breath
- It will still be complicated to parallelize applications

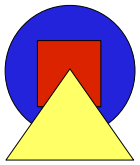

Parallel Computing in HEP

- There is Lattice QCD
- It is still true that most of the HEP computing is "embarrassingly" parallel, i.e., there is no communication among compute jobs
- However, ...

Parallel Computing Use at SLAC…

- Several groups are using parallel computing techniques
  - Astrophysics (KIPAC)
    - Code development at SLAC
    - Smaller runs at SLAC
    - Large runs at NCSA
    - Forming of structures in the universe from "first principles"
  - Accelerator Development Group (ACD)
    - "Mafia" in parallel; partly sponsored by SciDAC
    - Code development at SLAC
    - Smaller runs at SLAC
    - Large runs at NERSC
    - First full length, high time and spatial resolution cavity simulation

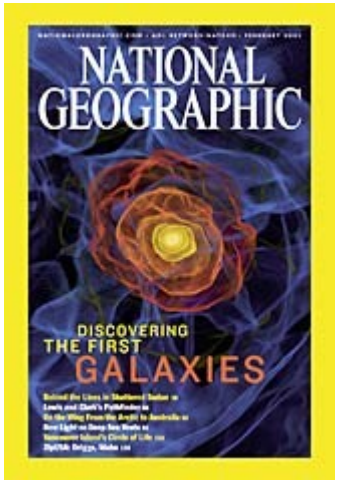

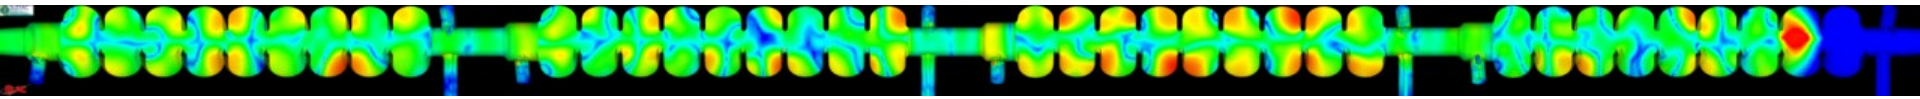

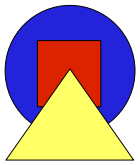

...Parallel Computing Use at SLAC

- Synchrotron people:
  - Structural molecular dynamics using Gaussian™
  - Runs as sequential application
  - Also runs in parallel with the "Linda" shared memory simulation system underneath

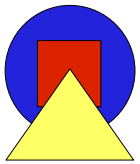

Headline News

- *SLAC Today* reported on March 23, 2007 (http://today.slac.stanford.edu/feature/2007/physics-in-alps.asp):

> Among the pyramidal Alps in the sleepy town of La Thuile, Italy, members of the BaBar collaboration are currently presenting spectacular results at the premier winter conference, Rencontres de Moriond.
> [...]
> to find this rare "mixing" in D mesons. The analysis also depended on an innovative use of a parallel computing system [sic!] at SLAC that sped up crucial calculations.

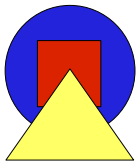

What happened

- Lecture serious at SLAC about parallel computing
- Young BaBar physicist attended
- Brian Aagaard Petersen: "I am using OpenMPI to minimize a likelihood function. It is an almost trivially parallelizable problem, since the likelihood is just a large sum of integrals."
- Gregory P. Dubois-Felsmann (BaBar's Computing Coordinator): "He's been able to reduce the time it takes him to perform certain analysis tasks from four hours to ~20 minutes by using MPI"

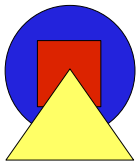

Today's Tools for Parallel Comp.

- Different hardware
  - Low-latency high-bandwidth interconnect can be anything:
    - Infiniband
    - Myrinet
    - Quadrics
- and wide variety of parallel software:
  - MPICH
  - LAM-MPI
  - MPICH-GM
  - OpenMPI
  - OpenMP
  - …
- Large number of libraries with parallel functions

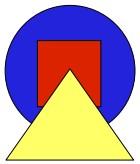

OpenMPI

- Message Passing Interface (MPI) is standardized
- API for sending messages between machines
- MPI's execution environment is not standardized
  - Subtle differences make porting annoying
- Many vendors, incl. Red Hat, Sun, Cisco, are now supporting **OpenMPI** (http://www.open-mpi.org/)
- Relatively new implementation
- Runtime switch to choose communication fabric
  - No need to use different MPI implementations for different fabrics
  - No need to recompile just because you want to run your application over Ethernet today even though it ran over Infiniband yesterday

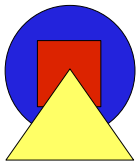

OpenMP

- Shared memory programming inside a single machine
- Several implementations are readily available:
  - GNU's Compiler Collection version ≥ 4.2
  - Intel compilers
  - Sun compilers (SPARC and x86)
  - See http://www.openmp.org/
- Intel's Cluster OpenMP
  http://www.intel.com/cd/software/products/asmona/eng/329023.htm

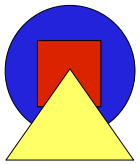

References

- SLAC's resources about parallel computing:
  - My lecture series about parallel computing
    http://www.slac.stanford.edu/~alfw/Parallel.html
  - High Performance Computing at SLAC
    http://www.slac.stanford.edu/comp/unix/unix-hpc.html
  - Parallel Computing at SLAC
    http://www.slac.stanford.edu/comp/unix/farm/parallel_at_SLAC.html
  - SLAC specific MPI tutorial
    http://www.slac.stanford.edu/comp/unix/farm/mpi.html
  - SLAC specific OpenMP tutorial
    http://www.slac.stanford.edu/comp/unix/farm/openmp.html
  - Mixing MPI and OpenMP
    http://www.slac.stanford.edu/comp/unix/farm/mpi_and_openmp.html