Nehalem-EX scalability testing

Sverre Jarp, Andrzej Nowak

February 16th 2010

non-NDA version

CERN openlab Minor Review 02.2010

Test description

- **Setup: Nehalem-EX BETA – NDA APPLIES**
  - 2.26 GHz, 64 GB memory, Bios #20
  - 8 cores per socket x 4 sockets = 32 cores
  - 32 cores x 2 threads per core = 64 threads
  - Scientific Linux CERN 5.4
- **Multi-threaded Geant4 prototype**
  - Author: Xin Dong (NEU), Gene Cooperman (NEU)
  - A multi-threaded prototype of the popular Geant4 particle simulation framework. Based on the FullCMS example, it resembles very closely a real LHC workload
  - 100 pi- events per thread @ 300 GeV
- **Multi-threaded ROOT minimization**
  - Developer in charge: Alfio Lazzaro (INFN/CERN)
  - A multi-threaded MPI-based implementation of ROOT minimization. A real world example of a High Energy Physics analysis reducing latency
  - mpirun -np XX ./ShapeAll -nll - x 2 (Elapsed time)

Multi-threaded Geant4 prototype

- **Results differ between BIOS versions**
- **TP = throughput**
- **Efficiency (% of max theoretical TP)**
  - > 95% @ 8 cores
  - < 95% @ 16 cores
  - ~90% @ 32 cores
    - nearly 29x speedup in a multi-threaded HEP program
- **SMT benefit @ 64 threads:**
  - 27% more real TP than all cores loaded
  - 15% more real TP then theoretical all cores loaded
- **Excellent and unprecedented scalability**
  - However above average SMT results indicate that there is potential for core scaling to be better

Multi-threaded ROOT minimization

- **Up to 8 cores: scaling on par with Nehalem-EP**
- **Excellent scaling up to 32 cores**
  - > 90% efficiency @ 8 cores
  - ~90% efficiency @ 16 cores
  - > 85% efficiency @ 32 cores
    - Over 27x speedup
    - Analysis time drops from 38 minutes to less than 1.5 minutes
- **SMT benefit @ 64 threads:**
  - ~40% efficiency: 2x performance drop
  - Results suggest a different bottleneck and saturation with 64 threads – we may be at the limit of scalability

Q & A

CERN openlab