

CMSC 714

Lecture 24

Applications – Gordon Bell Prize

Alan Sussman

Notes

- Final report due Tuesday, December 11, 6PM
 - Any problems/questions
- Course evaluation web site open through Tuesday
 - <https://www.CourseEvalUM.umd.edu>

Petascale Climate Modeling on Sunway TaihuLight

- CAM (Community Atmospheric Model)
 - Widely used model for weather/climate forecasting
 - Part of CESM (Community Earth System Model)
 - Over 1.5M LOC (lines of code), developed over many years by many people
 - 2 parts to the code
 - Dynamical core – models fluid flow through the simulation grid
 - Physics – clouds, precipitation, long/short-wave radiation, turbulent mixing at each grid point

Sunway TaihuLight

- Fastest supercomputer in the world when this work was done
 - Now #2 on Top500, maxing out at 125Pflops, 93Pflops for HPL benchmark
 - CPU is homegrown, with 260 cores, in 4 core groups, ~3Tflops/processor
 - Each core group has 1 MPE (management PE) and 8x8 cluster of CPEs (computing PEs) connected by a network-on-chip (NOC)
 - While MPE is pretty standard RISC core, a CPE has an L1 instruction cache and a 64KB scratch-pad memory (think of it as a program-controlled cache)
 - Each CG connected to 32GB memory, so 128GB per processor/node (same as Deepthought2!)
 - Supernode has 256 processors, fully connected through a custom network (I think that means a crossbar)
 - Supernodes connected via central network switch
- Programming model is MPI+X
 - X is either OpenACC or Athread, which is similar to pthreads

Implementation

- Goal is to parallelize the whole code, not just a kernel here and there
- Started with OpenACC directives
 - Note that someone had to build the OpenACC compiler for the machine
 - Performance not so good, because of limited memory per core, and compiler limitations
- Redesign to directly use Athreads on performance-limiting kernels
 - To improve data locality – to better use the scratchpad memory
 - To improve vectorization – each core has SIMD units
 - To overlap computation and communication between nodes

Performance results

- Overall goal is to maximize SYPD (simulation years per day)
- Results show for different grid resolutions, with some down to 75m resolution for highest performance – more usual is 3 to 100km per grid element
- Graphs show good speedup for Athreads version, compared to an Intel processor
 - And up to 3.3Pflops for high resolution grid (75m) on whole machine (131K processes/processors, over 10M cores), with very high parallel efficiency

Outcomes

- Able to simulate behavior of hurricane Katrina
 - Both overall structure, and track path over entire duration
 - Requires high resolution grid (25km) to be able to resolve accurately, previously only possible for a regional climate model (which has other problems)
- They don't say how much human effort was involved in building CAM for Sunway, but it was a lot
 - Modified 20% of the >750K LOC in the CAM version they worked on
 - And added >57K LOC