Data-Parallel Cosmology Analysis on Titan

Christopher Sewell (LANL), Katrin Heitmann (ANL), Li-ta Lo (LANL), Salman Habib (ANL), James Ahrens (LANL)

Objectives

Milestones

- Implement application-specific visualization and/or analysis operators needed for in-situ use by LCF science codes
- Use VTK-m to take advantage of multi-core and many-core hardware technologies

Target Application

- The Hardware/Hybrid Accelerated Cosmology Code (HACC) simulates the distribution of dark matter in the universe over time
- An important and time-consuming analysis function within this code is finding halos (high density regions) and centers of those halos

Impact

Portable, Data-Parallel Halo and Center Finding Algorithms

- Data-parallel algorithms for halo and center finding implemented using VTK-m allow the code to take advantage of parallelism on accelerators such as GPUs, and for the code to be portable across architectures
- Can be used for post-processing or in-situ (including with CosmoTools library)

Streaming Workflow for Large Halos

- At late time steps, particles become more concentrated in some nodes, resulting in great load imbalance for center finding
- During initial analysis task, particles in halos larger than specified size are output to disk using HACC's GenericIO library, which can then be streamed into memory from file to find centers one at a time as single-node jobs

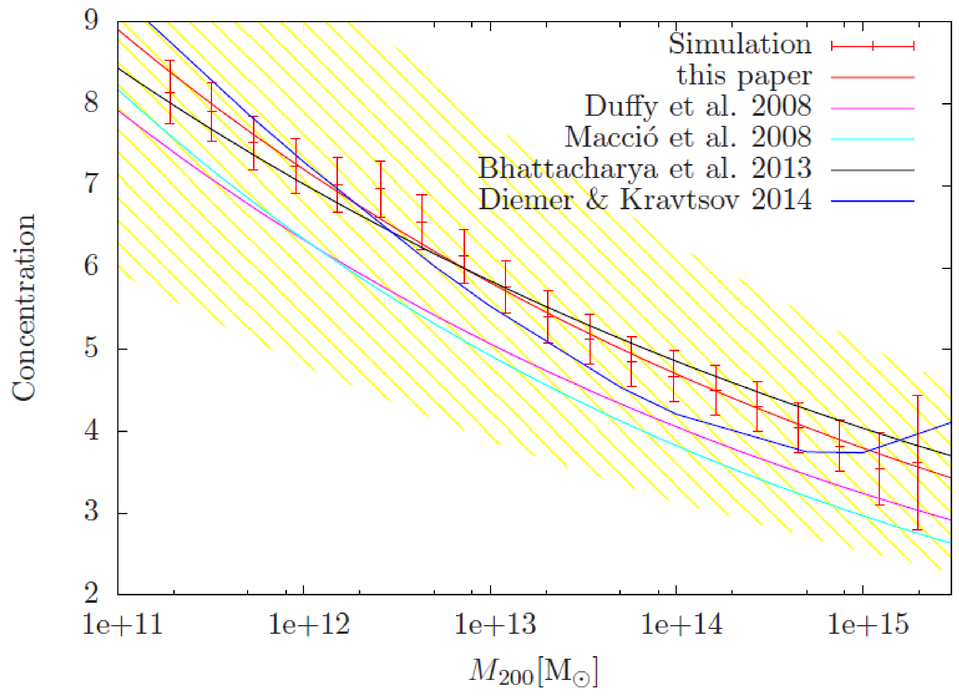

Concentration-mass relation over the full mass range covered by the Q Continuum simulation at redshift z = 0 (points with error bars) and the predictions from various groups. The yellow shaded region shows the intrinsic scatter. All predictions and the simulation results are well within that scatter.

Accomplishments

Results

- On Titan, this enabled MBP centers to be found on the GPU ~50x faster than using the pre-existing algorithms on the CPU (with one rank per node)
- Due to memory constraints on the GPUs, we utilize a hybrid approach for the large runs, in which the halos are computed on the CPU but the centers on the GPU
- Since no nodes wait idly while others compute centers of large halos, and since the jobs can be run even on another machine, resource allocations are not wasted
- The portability of VTK-m allowed us to run the same code on an Intel Xeon Phi

Science Impact

- This work allowed halo analysis to be completed on all time steps of a very large 8192^3 particle data set across 16,384 nodes on Titan for which analysis using the existing CPU algorithms was not feasible
- This is the first time that the c-M relation has been measured from a single simulation volume over such an extended mass range (see graph at left)

Publications

- Submitted to Astrophysical Journal Supplement Series: "The Q Continuum Simulation: Harnessing the Power of GPU Accelerated Supercomputers"

U.S. DEPARTMENT OF ENERGY | Office of Science

LA-UR-15-21392

SDAV Scalable Data Management, Analysis, and Visualization