Title: SDAV Visualization Area: Highlights at Los Alamos with HACC

Author(s): Sewell, Christopher Meyer
Ahrens, James P.
Patchett, John M.

Intended for: Web

Issued: 2014-03-27

Los Alamos
NATIONAL LABORATORY
—— EST.1943 ——

SDAV Visualization Area: Highlights at Los Alamos with HACC

Jim Ahrens, Chris Sewell, and John Patchett (Los Alamos National Laboratory)

Scalable Data Management, Analysis, and Visualization SciDAC Institute

**LANL Project Lead: Jim Ahrens**

Objectives

- **Milestone**
  - Implement application-specific visualization and/or analysis operators needed for in-situ use by LCF science codes
  - Use PISTON to take advantage of multi-core and many-core 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 the centers of those halos

Impact

- **VTK-m framework**
  - The PISTON component of VTK-m focuses on developing data-parallel algorithms that are portable across multi-core and many-core architectures for use by LCF codes of interest
  - PISTON consists of a library of visualization and analysis algorithms implemented using Thrust, and our extensions to Thrust

- **Halo and Center Finders**
  - Data-parallel algorithms for halo and center finding implemented using VTK-m (PISTON) allow the code to take advantage of parallelism on accelerators such as GPUs
  - Can be used for post-processing or in-situ, with in-situ integration directly into HACC or via the CosmoTools library
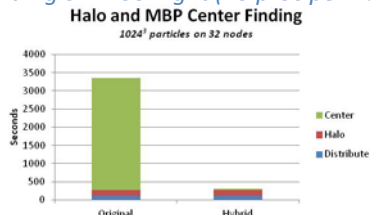
Accomplishments

- **Performance Improvements**

  - On Moonlight with 1024^3 particles on 128 nodes with 16 processes per node, PISTON 4.9x faster for halo + most bound particle center finding

  - On Titan with 1024^3 particles on 32 nodes with 1 process per node, PISTON 11x faster for halo + most bound particle center finding

  - Prototyped grid-based most bound particle center finder that performs $O(n\log n)$ total work instead of $O(n^2)$

- **Science Impact**

  - These performance improvements allow exact center finding to be used (instead of approximations) frequently in very large runs

- **Publications**

  - In preparation: "Utilizing Many-Core Accelerators for Halo and Center Finding within a Cosmology Simulation"

  - "The SDAV Software Frameworks …" Ultrascale Vis '12.

Timing results for halo and center finding on Moonlight (16 proc per node)

Timing results for halo and center finding on Titan (1 proc per node)
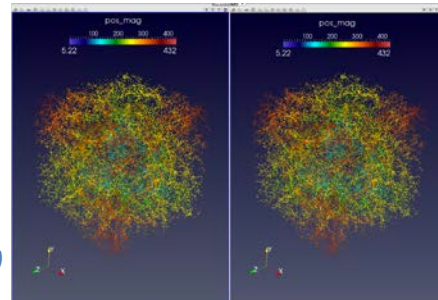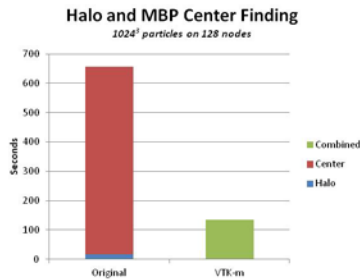
Visual comparison of halos computed by the original HACC algorithms (left) and the PISTON algorithms (right)