LA-UR-14-24408
Approved for public release; distribution is unlimited.

Title: LANL SDAV Visualization Update

Author(s): Sewell, Christopher Meyer

Intended for: SDAV VIsualization Teleconference, June 17, 2014

Issued: 2014-06-15

SDAV

Scalable Data Management, Analysis, and Visualization

U.S. DEPARTMENT OF ENERGY | Office of Science

LANL SDAV VISUALIZATION UPDATE

Christopher Sewell
Ollie Lo
John Patchett
Jim Ahrens
Los Alamos National Laboratory

SDAV Visualization Conference Call
June 2014

Results: Xeon Phi (MIC) on Stampede

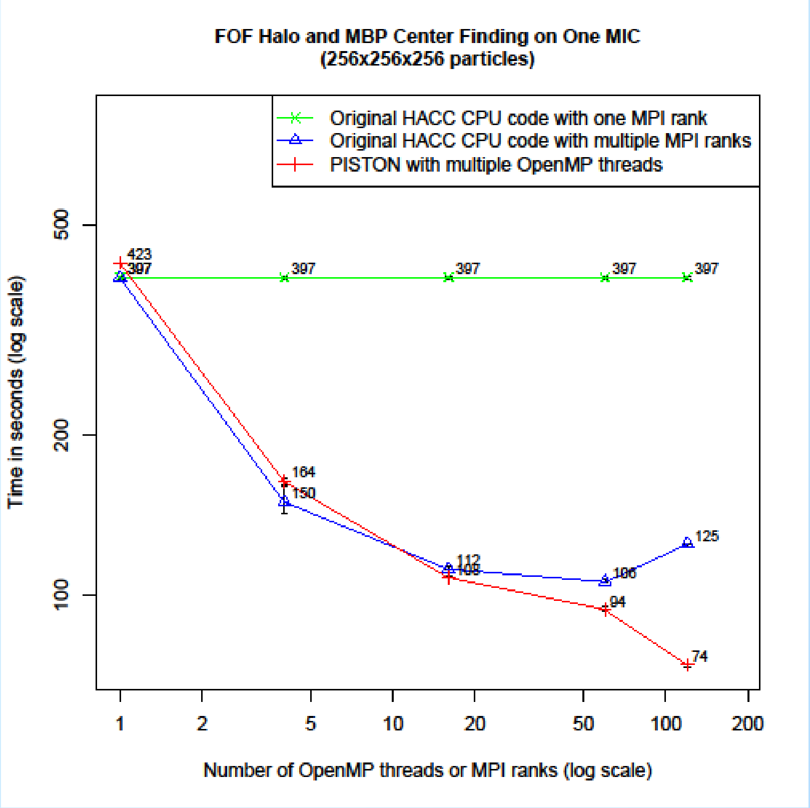

• *To demonstrate the portability of our algorithms, the same code was compiled to the Thrust OpenMP backend (including our own OpenMP implementation of scan) and run on a 256^3 particle data set on an Intel Xeon Phi SE10P (MIC) Coprocessor on a single node of Stampede at TACC*

• *PISTON version scales to more cores than running the existing serial algorithms with multiple MPI processes*

SDAV
Scalable Data Management, Analysis, and Visualization

U.S. DEPARTMENT OF ENERGY | Office of Science

Results: Titan

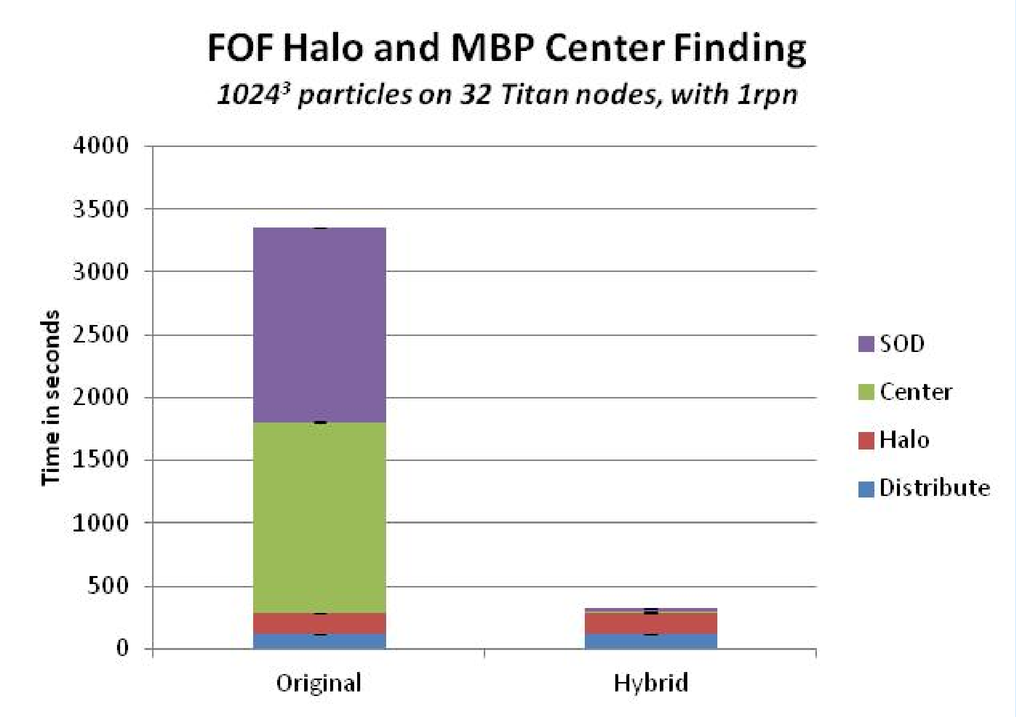

- *This test problem has ~90 million particles per process.*
- *Due to memory constraints on the GPUs, we utilize a hybrid approach, in which the halos are computed on the CPU but the centers on the GPU.*
- *The PISTON MBP center finding algorithm requires much less memory than the halo finding algorithm but provides the large majority of the speed-up, since MBP center finding takes much longer than FOF halo finding with the original CPU code.*

Results: Large Run on Titan

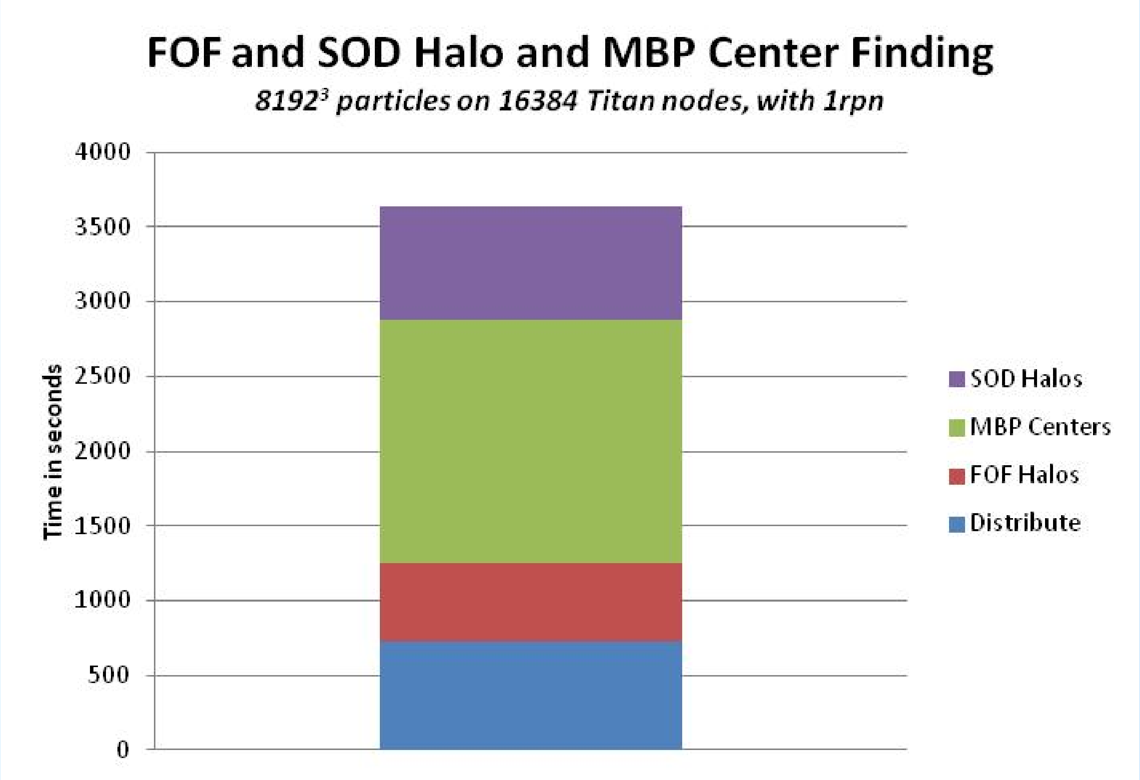

- *Because the memory requirements increase with the number of MPI processes due to overload regions, the HACC simulation with this data can only use one MPI process per node, along with the associated GPU, given the memory available on Titan.*
- *Thus, there is an even greater potential for speed-up by utilizing the GPUs.*
- *The performance improvements using PISTON on GPUs allowed halo analysis to be performed on 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.*

In-situ Integration in HACC

- Successfully ran 500 time-step, 512^3 particle simulation on Moonlight using our GPU halo and center finders integrated with HACC in-situ
- Completed prototype integration with CosmoTools

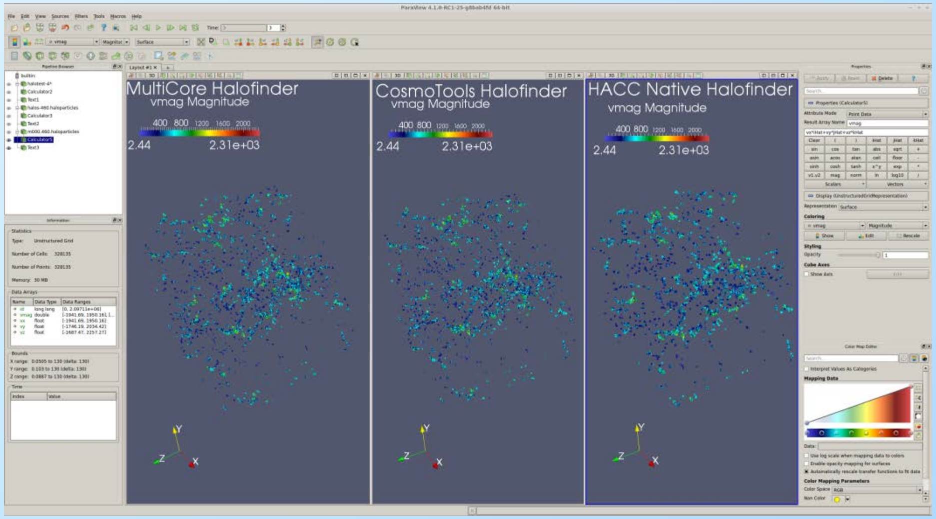

Extension: Potential Field Algorithm

- Optimization for MBP center finder: for each halo…
  - Superimpose a grid over the (extended) extents of the halo
  - Estimate particle density on the grid using binning
  - Solve Poisson equation for the potential on the grid using FFT (actually DST for zero boundary conditions approximation)
  - Find the grid point with minimum potential
  - Search around the neighborhood of the minimum potential grid for the particle with minimum potential
  - Return the position of such particle as the halo center

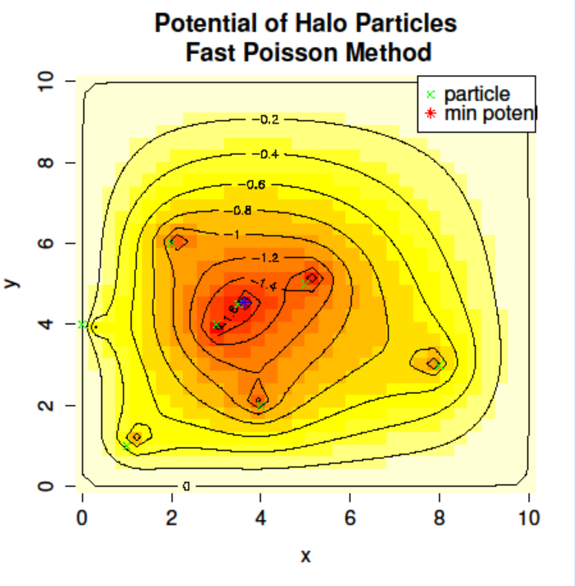

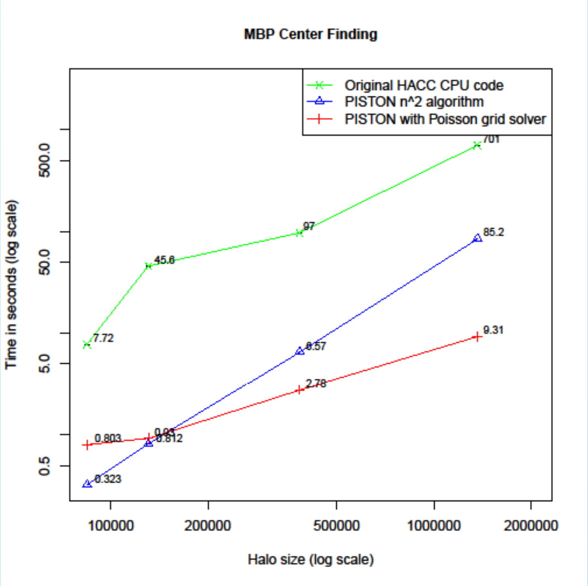

Extension: Dendrogram Halo Finder

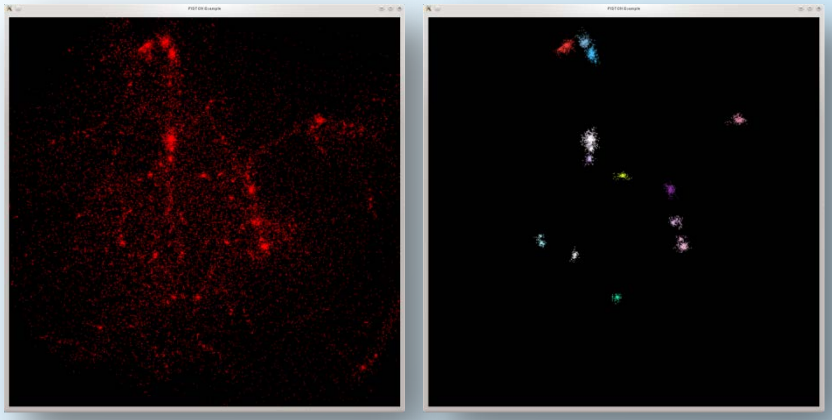

Collaboration with University of Utah (Wathsala Widanagamaachchi, Valerio Pascucci, and Timo Bremer)

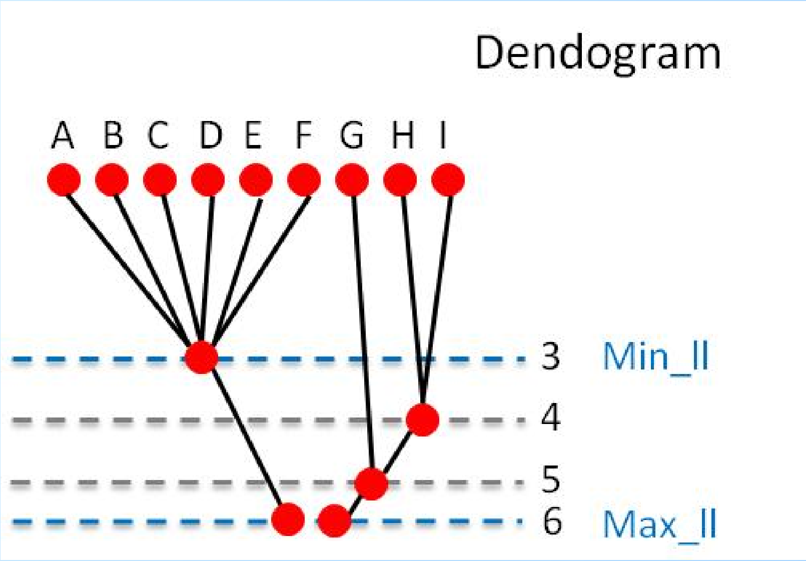

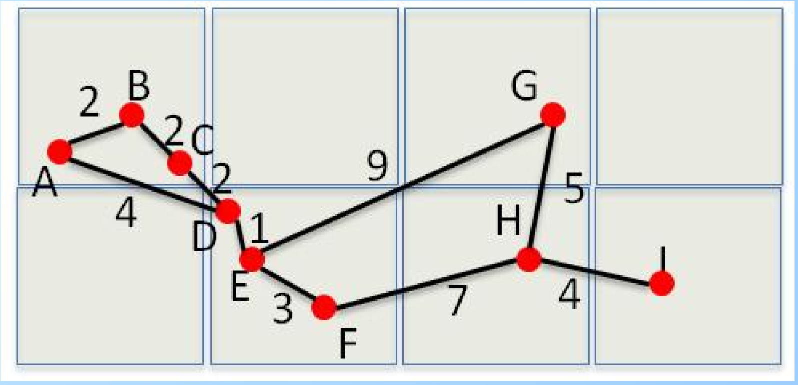

- Implements friend-of-friends halo finding algorithm
- Stores the clustering hierarchy of particles in a dendogram
- Dendogram allows queries for a range of linking lengths
- Uses PISTON to make use of on-node parallelism