Parallel Data Storage, Analysis, and Visualization of a Trillion Particles

Suren Byna[1], J. Chou[2], O. Rübel[1], Prabhat[1], H. Karimabadi[3], W. S. Daughton[4], V. Roytershteynz[3], E. W. Bethel[1], M. Howison[5], K.-J. Hsu[2], K.-W. Lin[2], A. Shoshani[1], A. Uselton[1], and K. Wu[1]

(1) Lawrence Berkeley National Laboratory (2) Tsinghua University, Taiwan (3) University of California - San Diego (4) Los Alamos National Laboratory (5) Brown University

Motivation

Modern scientific discoveries are driven by data

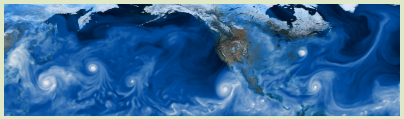

Climate modeling
Expected to be hundreds of Exabytes by 2020

CERN Large Hadron Collider
Produces Petabytes of data per year

Square Kilometer Array
Expected to produce 1 Exabyte of data per day

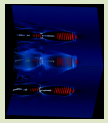

Laser Wakefield Acceleration
Simulates Billions of Particles

Storing, analyzing, and visualizing large data are big challenges

VPIC Trillion Particle Simulation

- State-of-the-art 3D plasma physics simulation
- An Exascale problem that scales well on large systems
- Ran a 2-trillion particle magnetic reconnection simulation on Hopper supercomputer at NERSC
  - Uses 5000 nodes of ~6400 nodes of Hopper
  - Stores ~1 trillion electron particles (File size: 32TB to 39TB)
- Scientists never observed characteristics of all particles

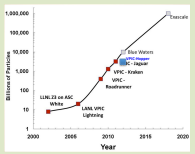

Research Objectives

- Finding scalable data management solutions for the following challenges

- **A scalable I/O strategy for storing massive particle data output**
- **A scalable strategy for conducting analysis on these datasets**
- **A visualization strategy for examining the datasets**

Our Techniques and Tools

- Scalable I/O strategy for storing particle data
  - H5Part: A simple API on top of HDF5 to read/write particle data
- Scalable strategy for conducting analysis on these datasets
  - FastBit: Bitmap index generation and querying software
  - Hybrid Parallel FastQuery
    - API to generate bitmap indexes
    - API to query indexed or data from different data formats (HDF5, NetCDF, and ADIOS-BP)
- Visualization strategy for examining the datasets
  - Query-driven visualization using VisIt

Strategy for Writing > 32TB Data

- Parallel HDF5: Standard data format used by many data analysis and visualization tools
- H5Part: A simple API on top of HDF5 to read/write data to the HDF5 file format
  - Portable high performance parallel data interface
  - Building blocks of H5Part
    - Particles
    - Block structured data (H5Block and H5MultiBlock)
- One file for all data

- Tuning for performance on Lustre parallel file system – Achieved 27 GB/s at a concurrency of 128K cores

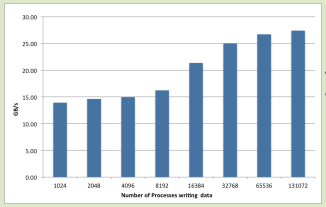

VPICBench: VPIC I/O kernel weak scaling study

- Performance of writing 32TB particle data (Each particle contains 8 variables)
- Reached I/O peak rate in writing each variable
- Amortized I/O rate of 26 GB/s on the Lustre parallel file system with 35 GB/s peak bandwidth

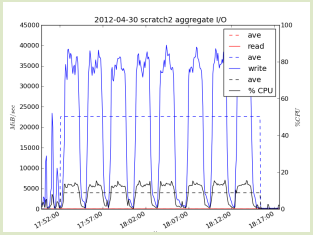

Parallel I/O performance of writing VPIC's particle data;

Indexing and Querying with FastBit & FastQuery

- **FastBit**
  - Provides querying capability for fast scientific data access
  - Indexes and stores each column data separately
  - Generates compressed bitmap indexes
  - Search speeds are 10X-100X better than the best known bitmap indexing methods
- **Hybrid Parallel FastQuery**
  - Parallel querying and index building
  - Works with array data in different data formats (HDF5, NetCDF, ADIOS-BP)
  - MPI + Pthreads implementation to better utilize multicore CPUs

- **10 minutes to index 30TB data and 3 seconds to query high energy particles**

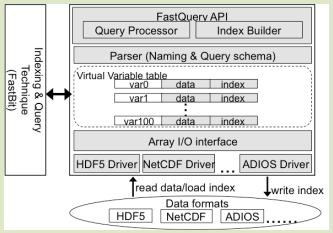

Query-Driven Visualization

- Visualizing too much data on computer displays is a challenge
  - Typical computer displays have O(1M) pixels
  - Brute force rendering a trillion particles implies an overdraw factor of O(1M)
- Reduced the number of particles before rendering by down-selecting the scientifically interesting features
- Highly energetic particles (energy > 1.3)
- **Agyrotropy** visualized for the first time when magnetic reconnection occurs

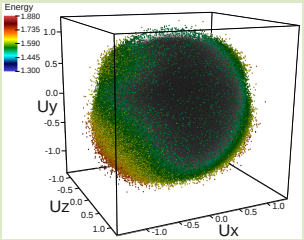

Visualization of all particles with energy > 1.3

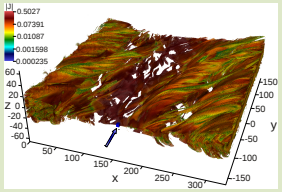

Iso-surface plot of the positron particle density

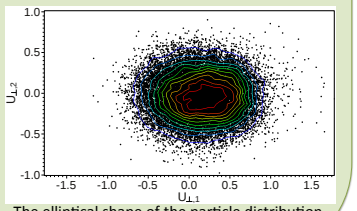

The elliptical shape of the particle distribution is indicative of agyrotropy

This work was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research was supported in part by National Science Foundation under NSF grant OCI 0904734. This research used resources of the National Energy Research Scientific Computing Center (NERSC). Simulations were also supported by an allocation of computing resources provided the NSF at the National Institute for Computational Sciences and by NASA (Pleiades), and the National Center for Computational Sciences at Oak Ridge National Laboratory (Jaguar). The authors would like to thank NERSC and Cray staff for troubleshooting I/O issues on Hopper. We would also like to thank members of the HDF Group and Burlen Loring.