The purpose of computing is insight, not numbers.

Richard Hamming

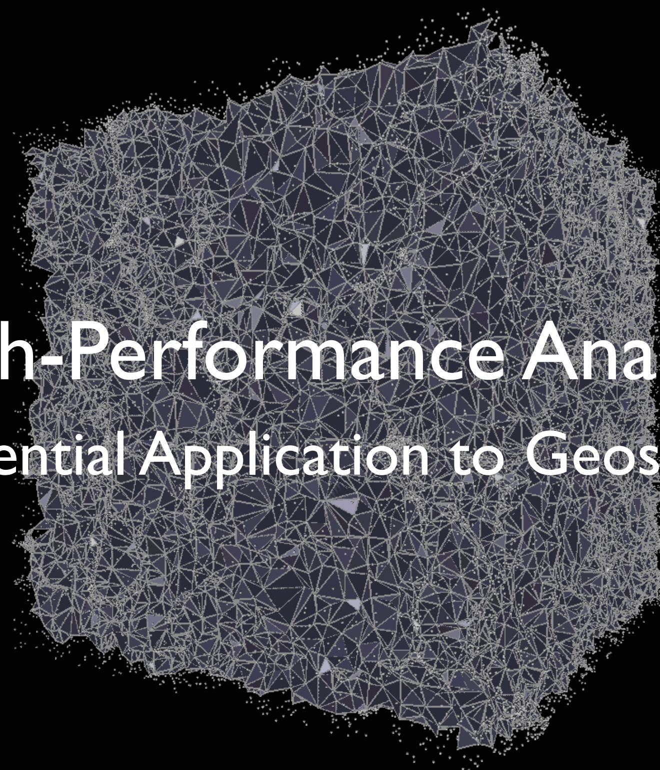

High-Performance Analytics

with Potential Application to Geospatial Data

Tom Peterka

tpeterka@mcs.anl.gov

Mathematics and Computer Science Division

HPC Geospatial Analytics Workshop

April 29, 2014

Preliminaries

Moving from Postprocessing to Run-Time Scientific Data Analysis in HPC

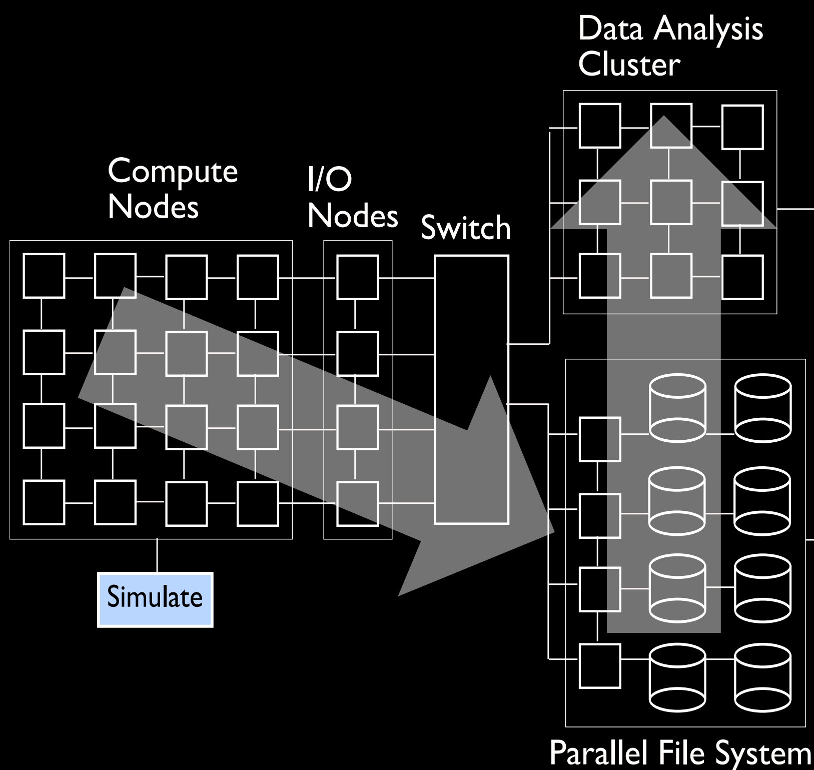

Postprocessing analysis and visualization

Run-time analysis and visualization

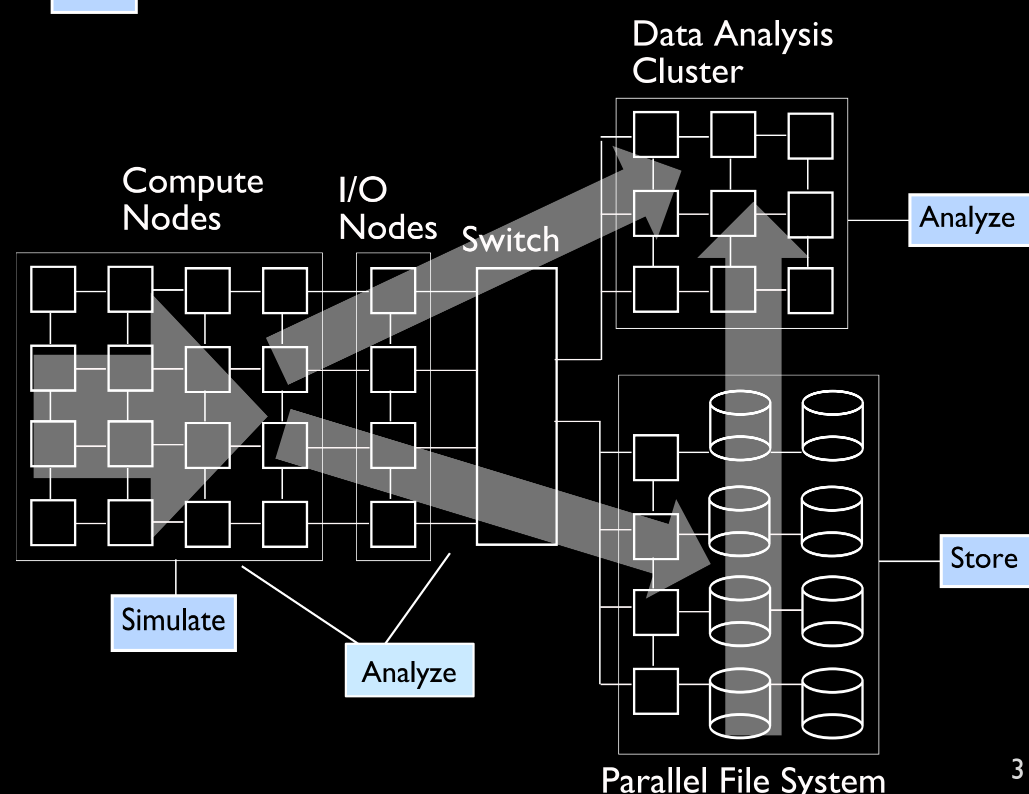

Definition of Data Analysis

- Any data transformation, or a network or transformations.
- Anything done to original data beyond its original generation.
- Can be visual, analytical, statistical, or data management.

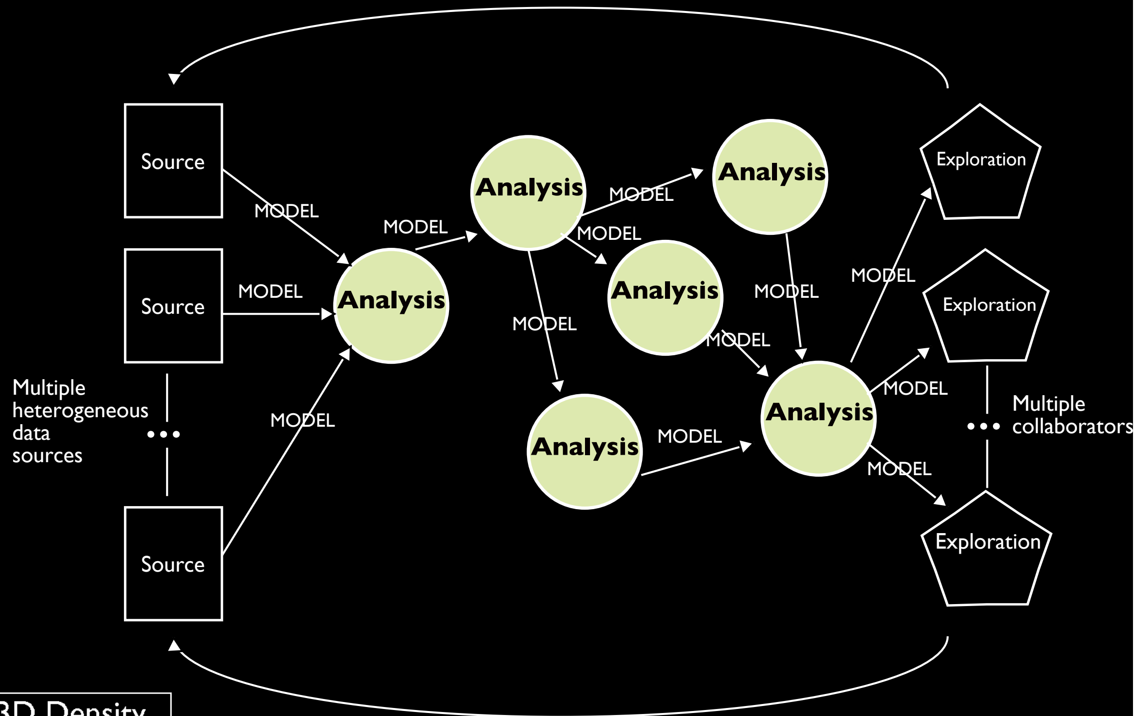

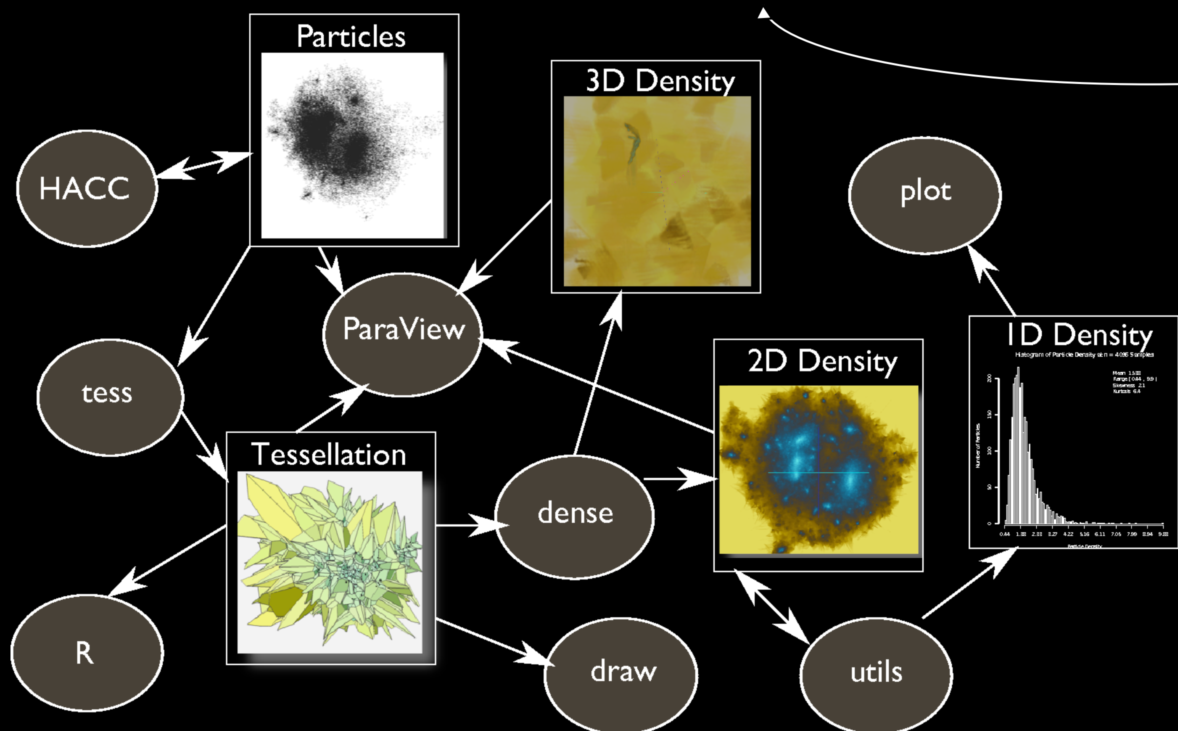

Some Examples

Particle Tracing Streamlines and Pathlines

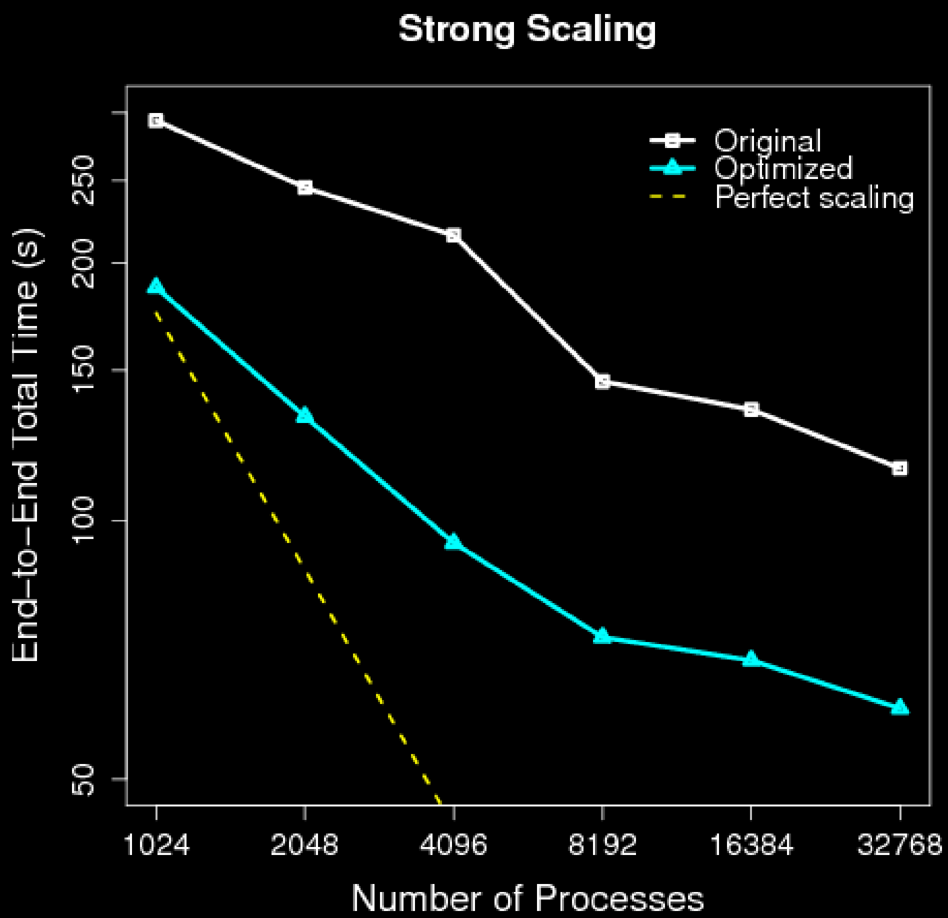

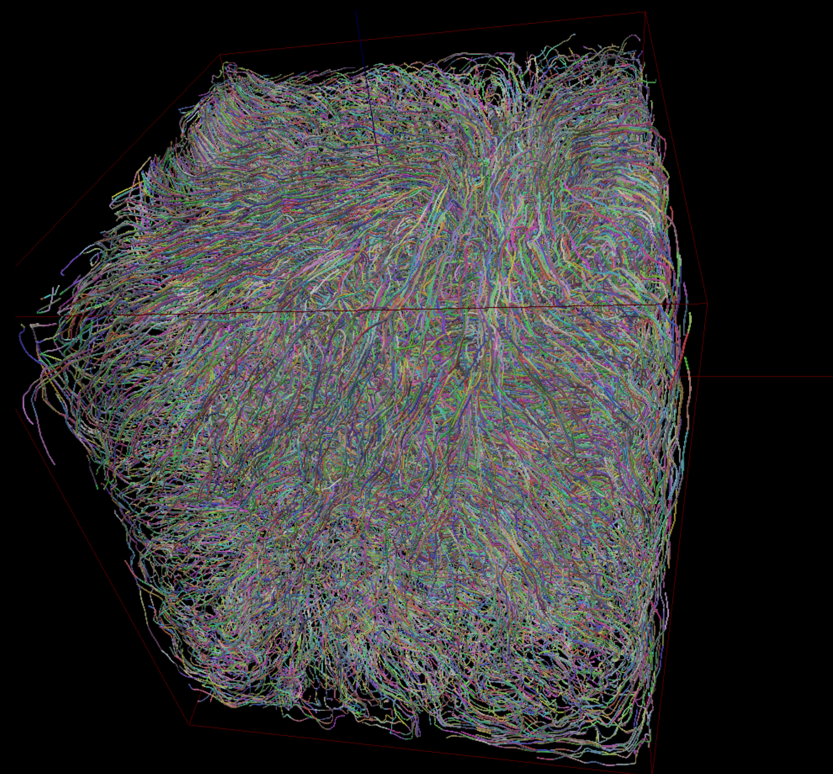

Particle tracing of ¼ million particles in a 2048^3 thermal hydraulics dataset results in strong scaling to 32K processes and an overall improvement of 2X over earlier algorithms

Peterka et al., A Study of Parallel Particle Tracing for Steady-State and Time-Varying Flow Fields. IPDPS '11.

Lagrangian Coherent Structures from FTLE

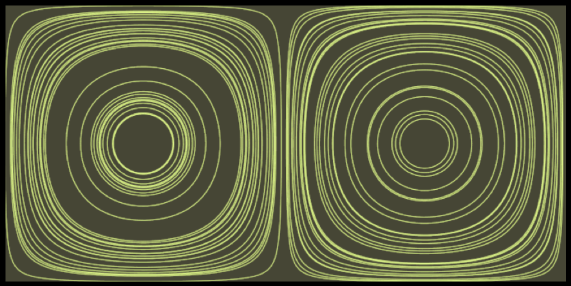

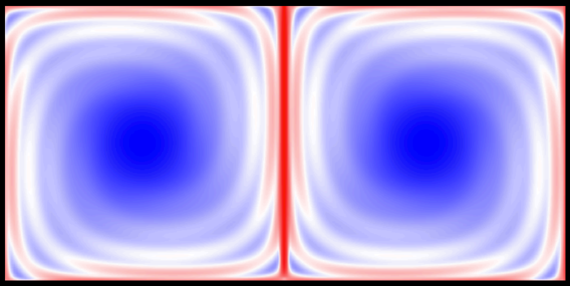

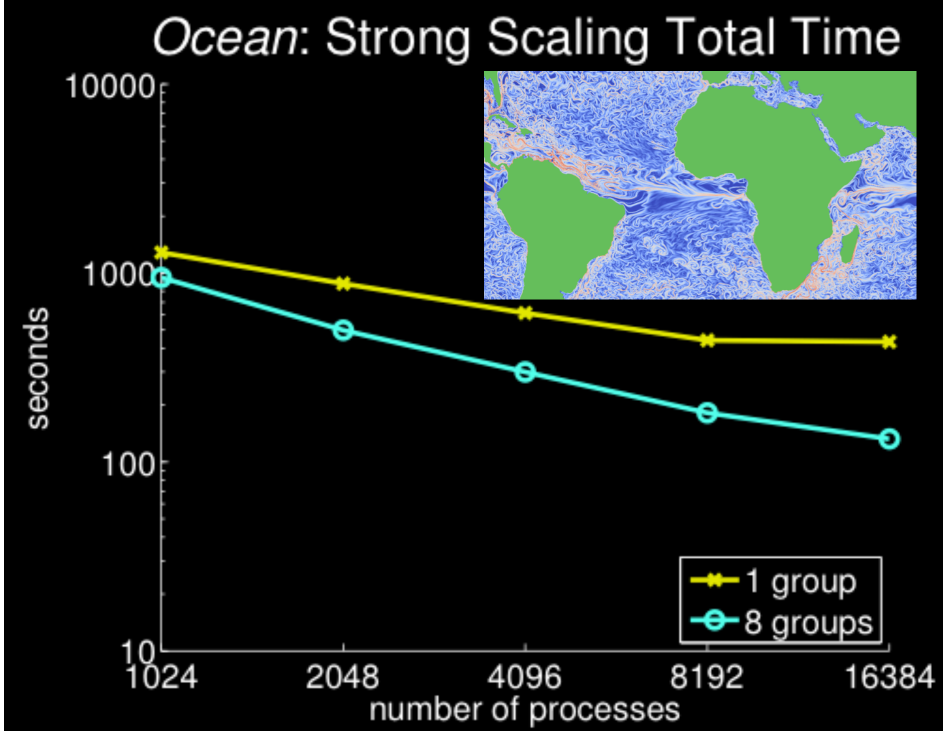

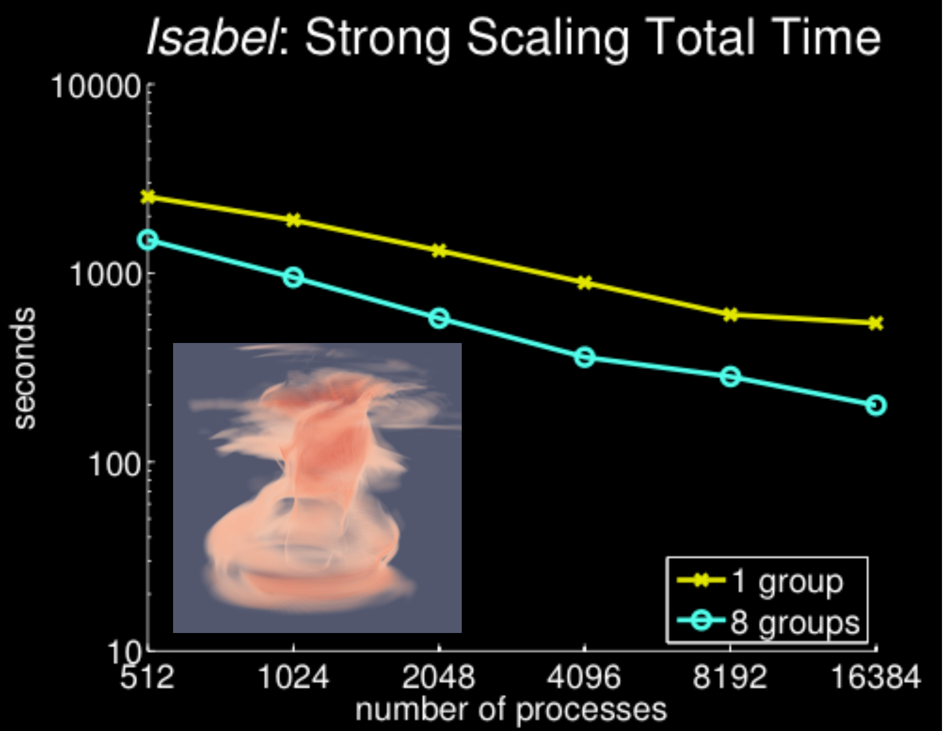

Left: Particle tracing of 288 million particles over 36 time steps in a 3600x2400x40 eddy resolving dataset. Right: 131 million particles over 48 time steps in a 500x500x100 simulation of Hurricane Isabel. Time includes I/O.

Nouanesengsy et al., Parallel Particle Advection and FTLE Computation for Time-Varying Flow Fields, SC12,

Stream Surfaces

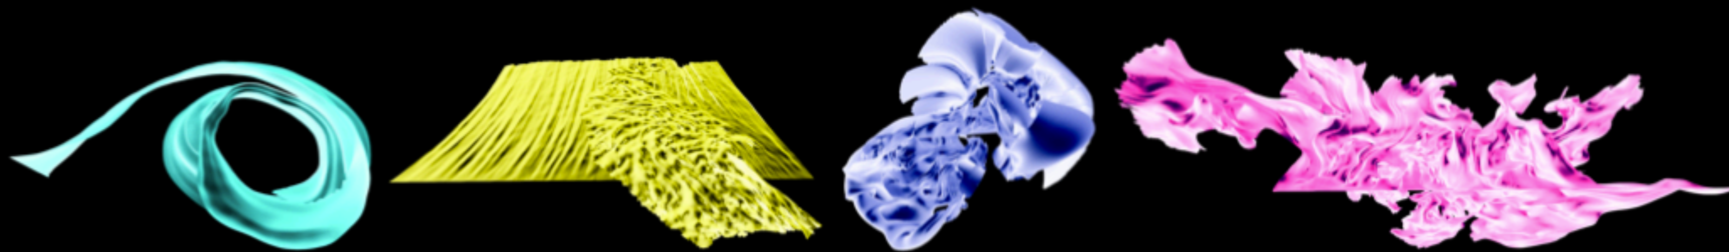

(a) Isabel (b) MJO (c) Plume (d) Nek

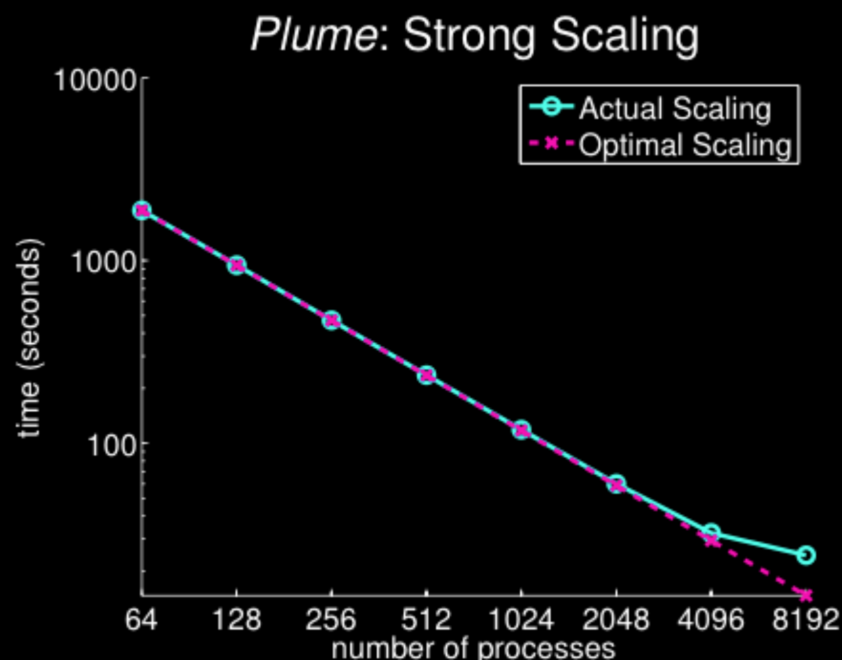

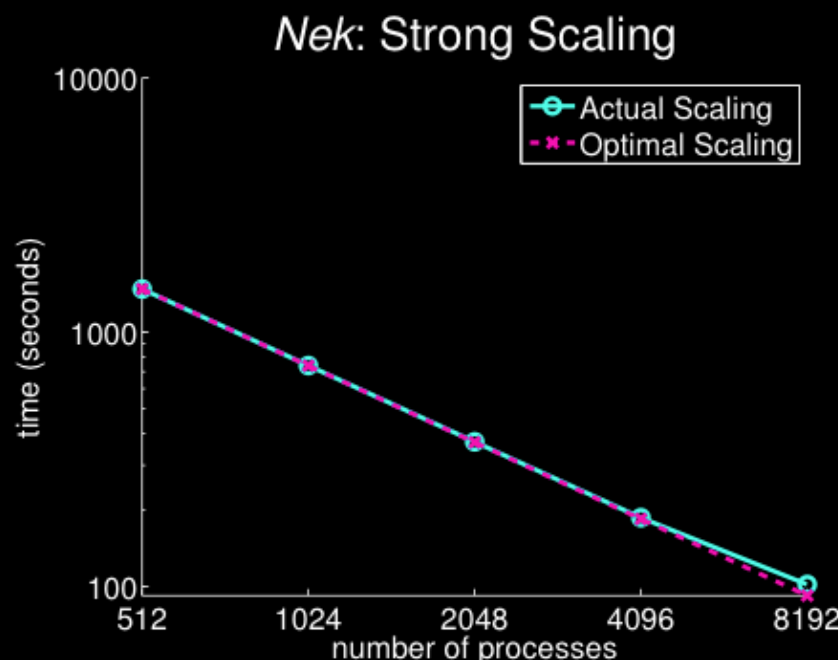

Left: 64 surfaces each seeded with 512 particles are advected in a 504x504x2048 simulation of a solar flare. Right: 64 surfaces each with 2K seeds in a 2K x 2K x 2K Nek5000 thermal hydraulics simulation. Time excludes I/O.

Lu et al., Scalable Computation of Stream Surfaces on Large Scale Vector Fields, submitted to SC14.

Information Entropy

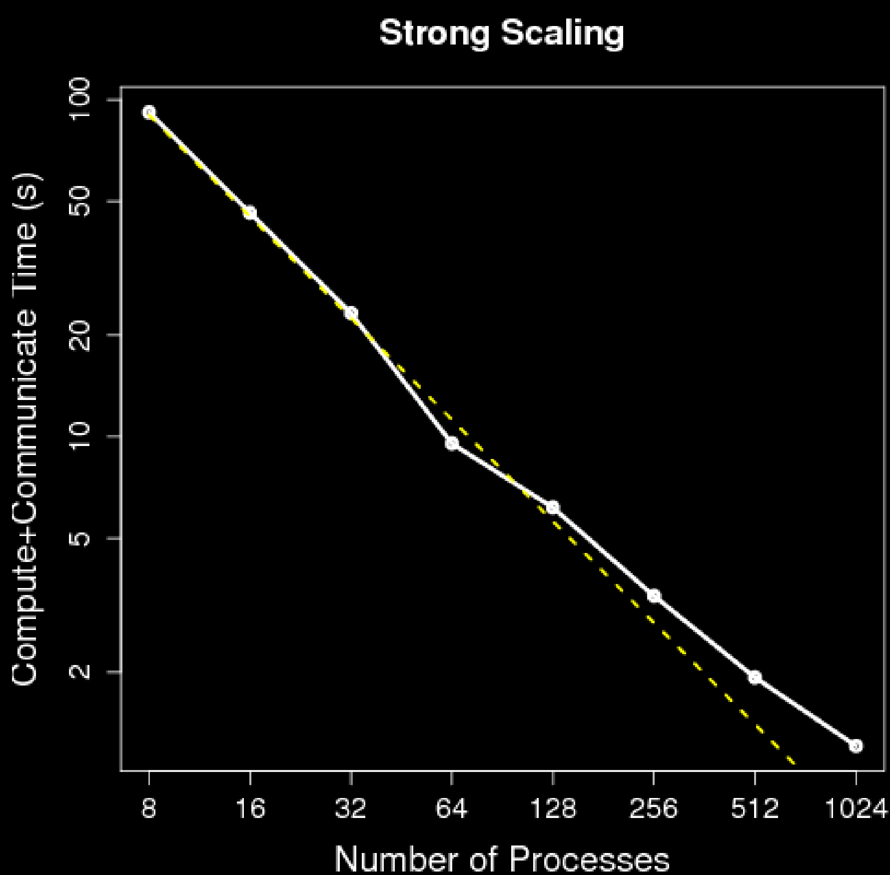

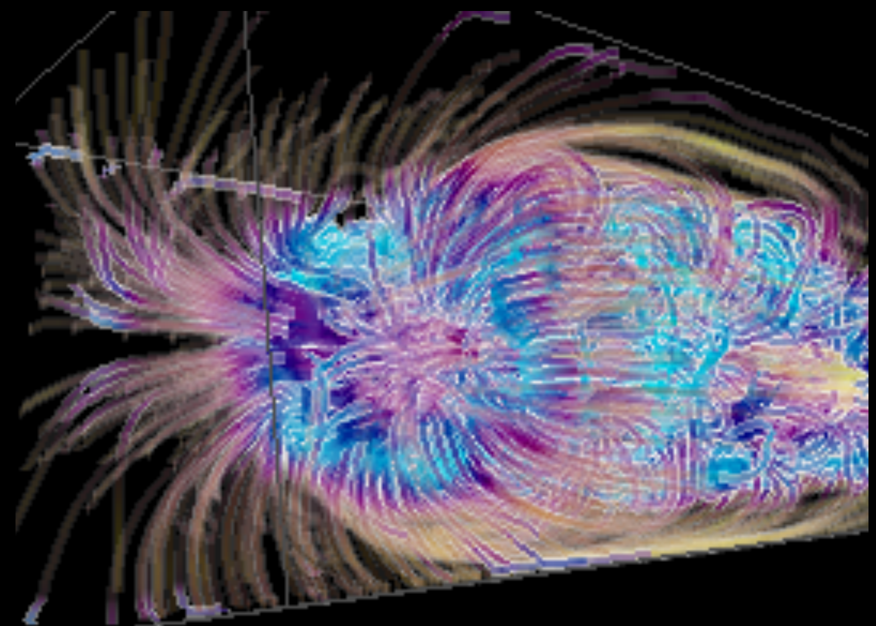

Computation of information entropy in 126x126x512 solar plume dataset shows 59% strong scaling efficiency. Time excludes I/O.

Chaudhuri et al., Scalable Computation of Distributions from Large Scale Data Sets, LDAV '12.

Topological Analysis

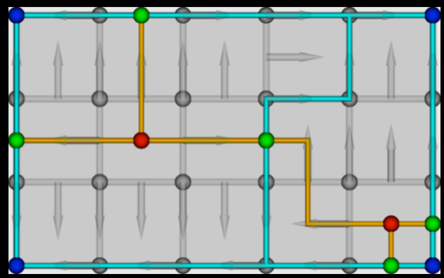

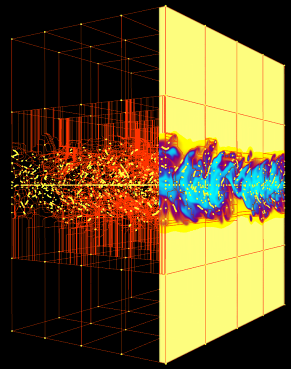

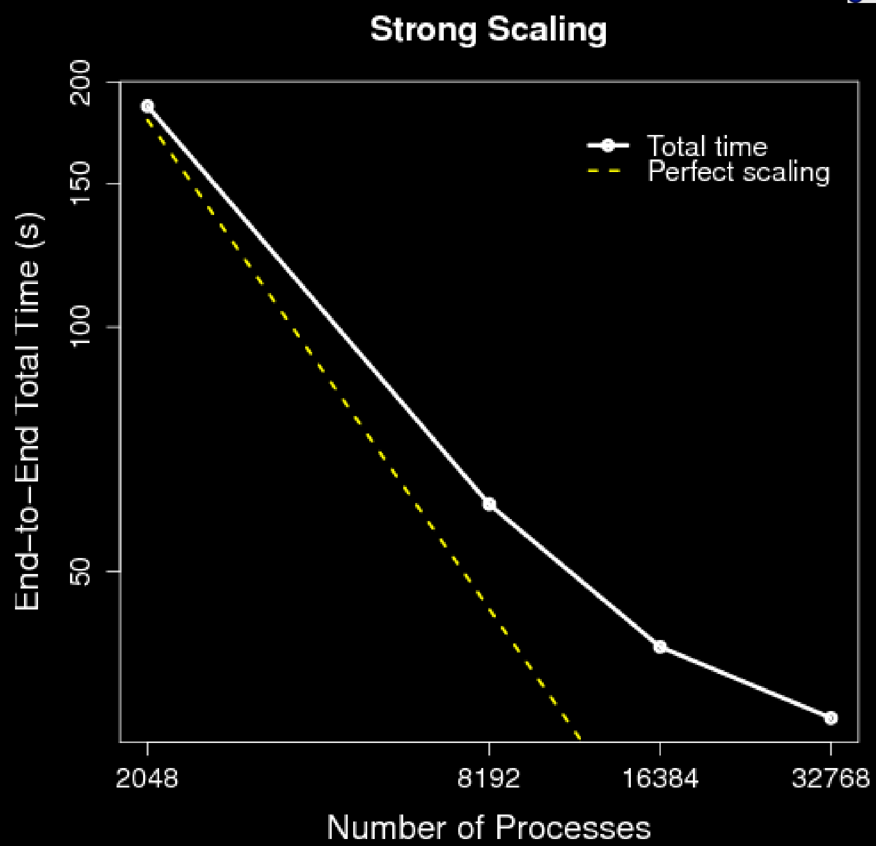

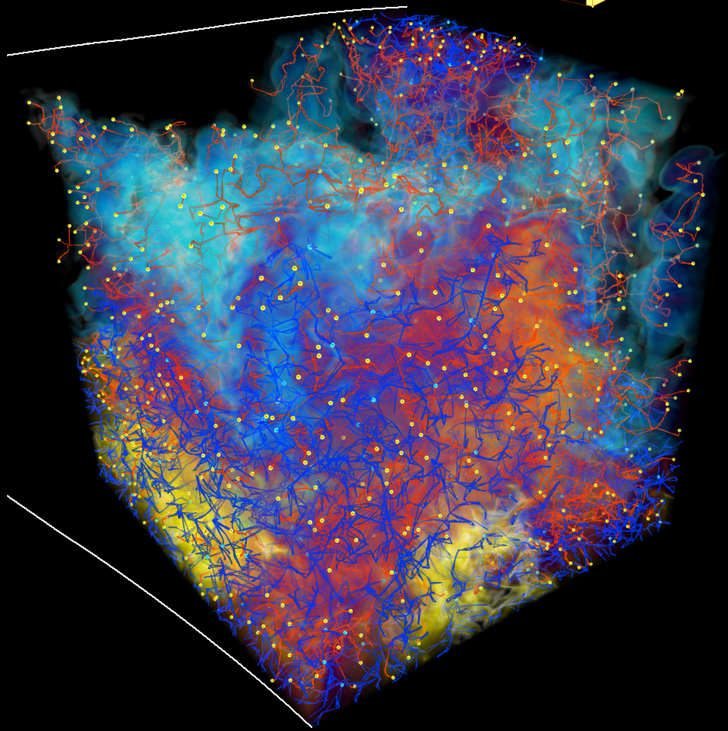

Computation of Morse-Smale complex in 1152^3 Rayleigh-Taylor instability data set results in 35% end-to-end strong scaling efficiency, including I/O.

Gyulassy et al., The Parallel Computation of Morse-Smale Complexes, IPDPS ’12.

Computational Geometry

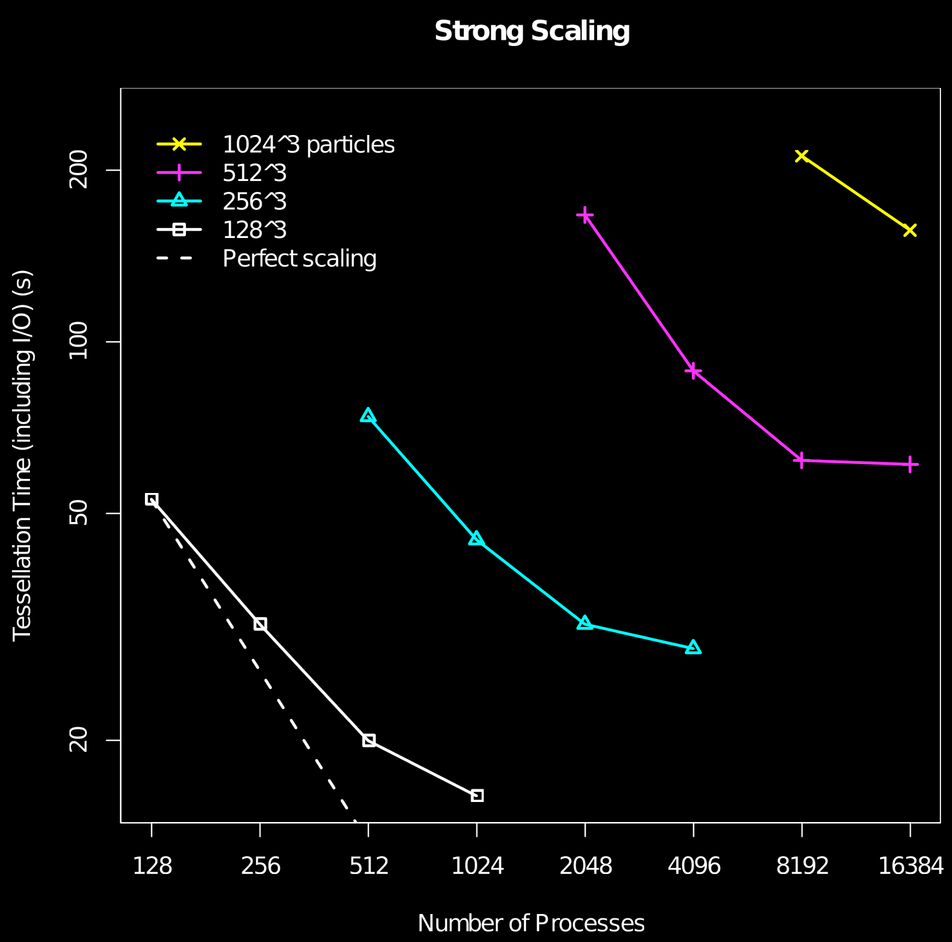

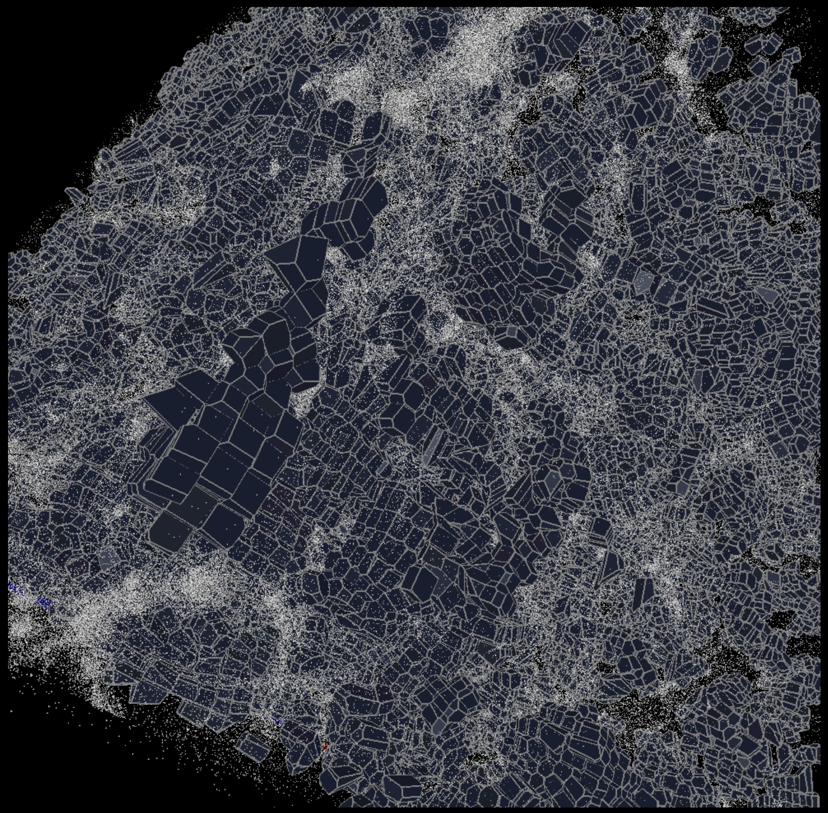

For 128^3 particles, 41 % strong scaling for total tessellation time, including I/O; comparable to simulation strong scaling.

Peterka et al., High-Performance Computation of Distributed-Memory Parallel 3D Voronoi and Delaunay Tessellation, submitted to SC14,

Common Denominators

- Big science => big data, big machines
- Most analysis algorithms are not up to speed
  - Either serial, or
  - Overheads kill scalability
- Solutions
  - Process data closer to the source
  - Write scalable analysis algorithms
  - Parallelize in various forms
  - Build software stacks of useful and reusable layers
- Usability and workflow
  - Develop libraries rather than tools
  - Users write small main programs and call into libraries

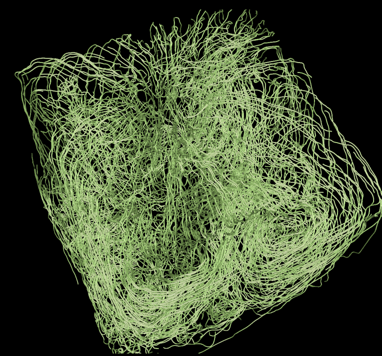

Streamlines and pathlines

Stream surfaces

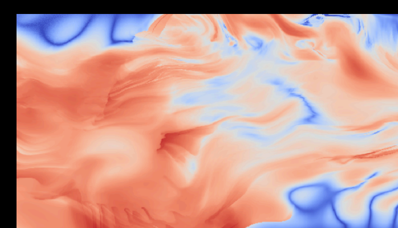

FTLE

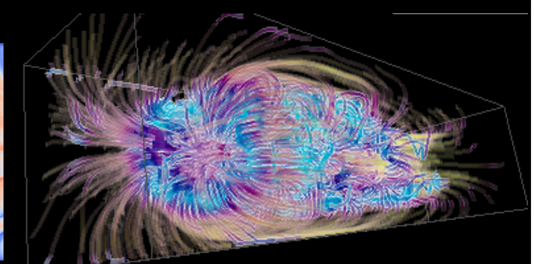

Information entropy

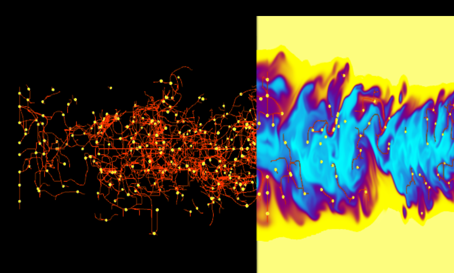

Morse-Smale complex

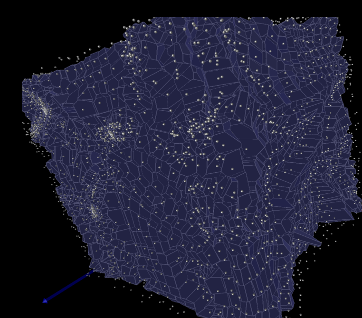

Voronoi Tessellation

Core Infrastructure

A library with a small ℓ

Features

Parallel I/O to/from storage

Domain decomposition

Network communication

Utilities

Library

Written in C++ with C bindings

Autoconf build system (configure, make, make install)

Lightweight: libdiy.a 800KB

Maintainable: ~15K lines of code, including examples

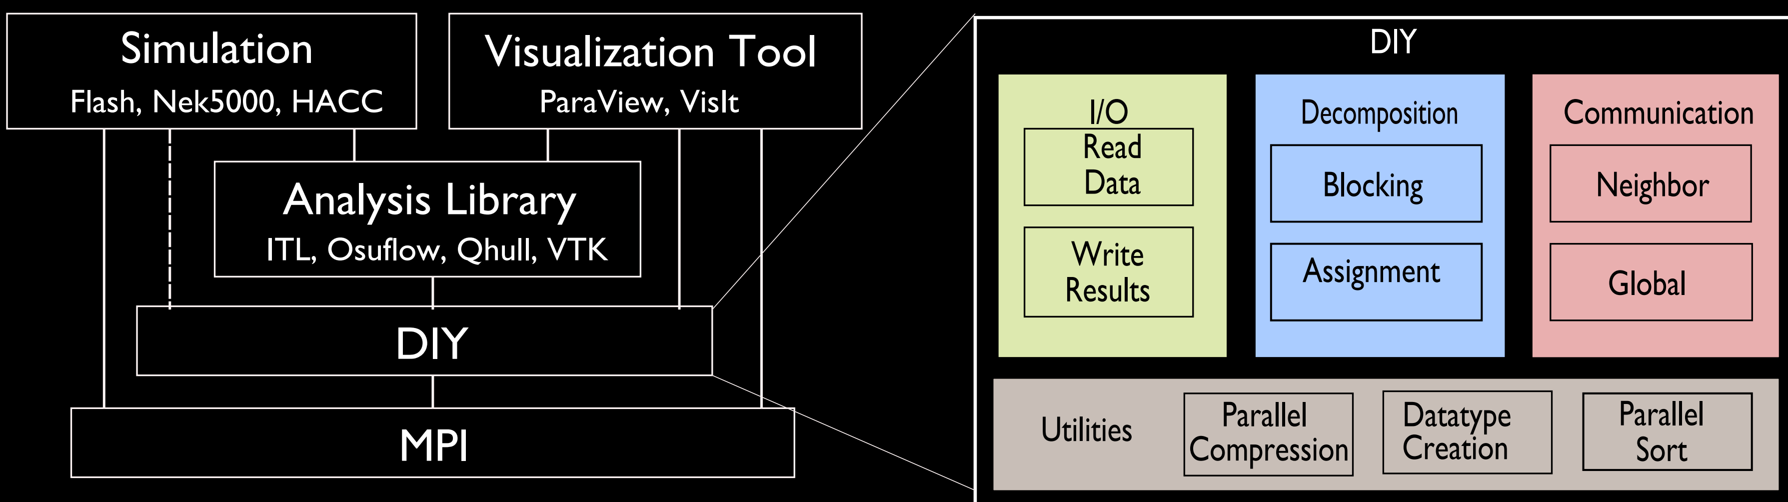

DIY usage and library organization

Write Your Own Main

Documentation

- README for installation
- User's manual with description, examples of custom datatypes, complete API reference

Tutorial Examples

- Block I/O: Reading data, writing analysis results
- Static: Merge-based, Swap-based reduction, Neighborhood exchange
- Time-varying: Neighborhood exchange
- Spare thread: Simulation and analysis overlap
- MOAB: Unstructured mesh data model
- VTK: Integrating DIY communication with VTK filters
- R: Integrating DIY communication with R stats algorithms
- Multimodel: multiple domains and communicating between them

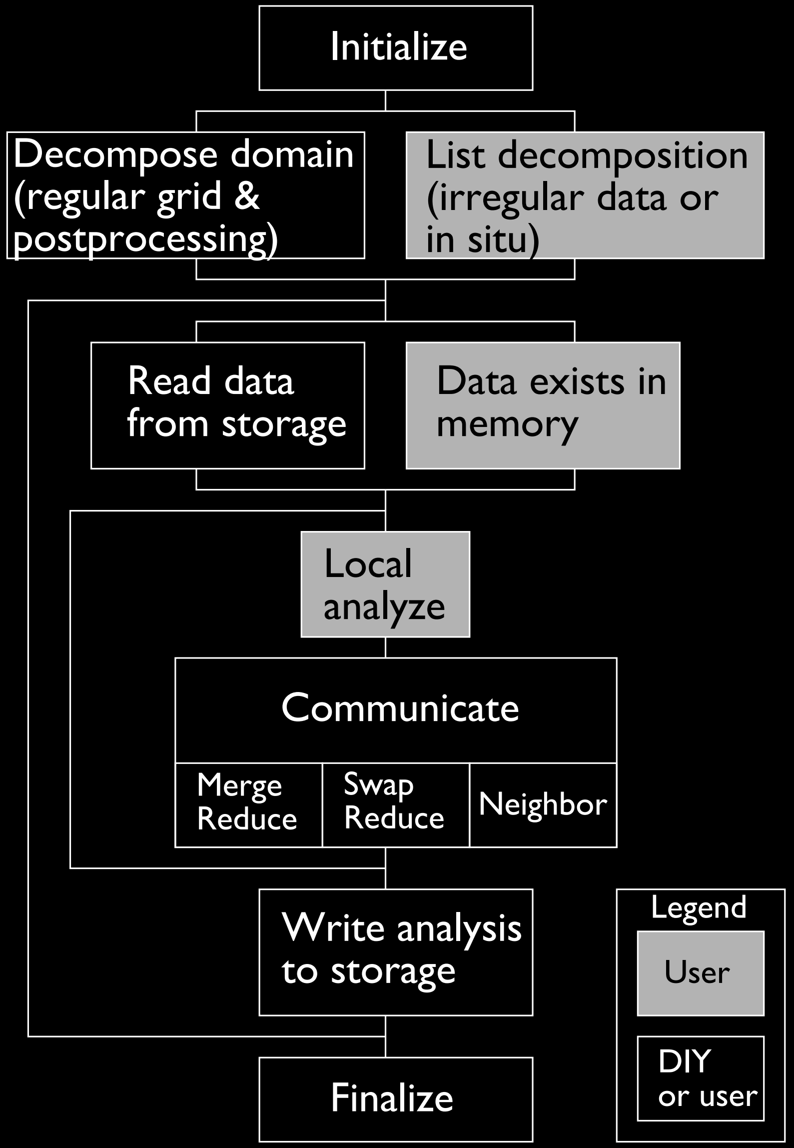

One Example in Greater Detail

Parallel Tessellation

We developed a prototype library for computing in situ Voronoi and Delaunay tessellations from particle data and applied it to cosmology, molecular dynamics, and plasma fusion.

Key Ideas

- Mesh tessellations convert sparse point data into continuous dense field data.
- Meshing output of simulations is data-intensive and requires supercomputing resources
- No large-scale data-parallel tessellation tools exist.
- We developed such a library, tess.
- We achieved good parallel performance and scalability.
- Widespread GIS applicability in addition to the datasets we tested.

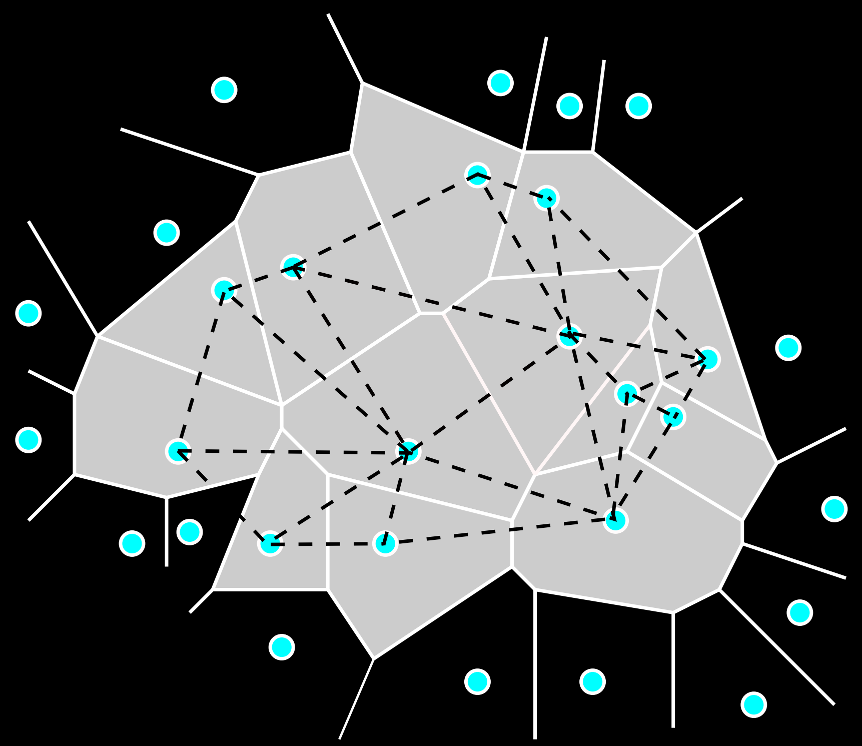

Tess Library

Tess is our parallel library for large-scale distributed-memory Voronoi and Delaunay tessellation.

Dense, our density estimator, currently reads the tessellation from disk and estimates density onto a regular grid. Eventually dense will be converted to a library that can be coupled in memory to tess output, saving the tessellation storage.

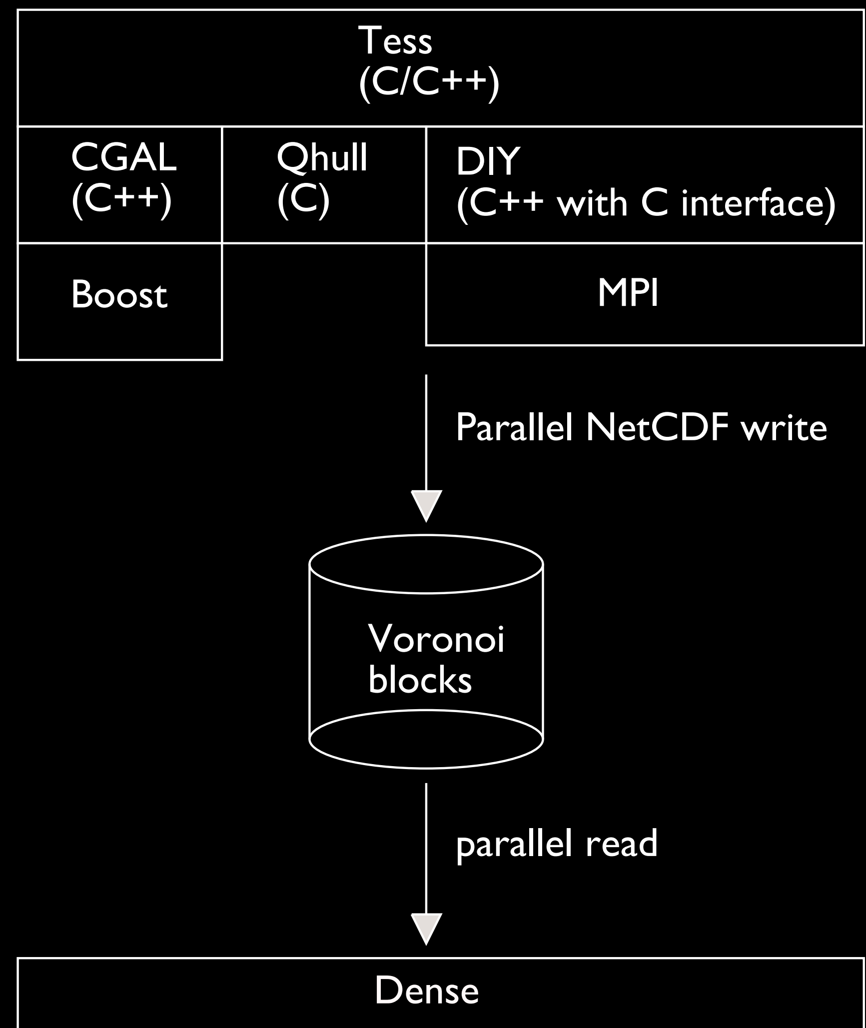

Scalability

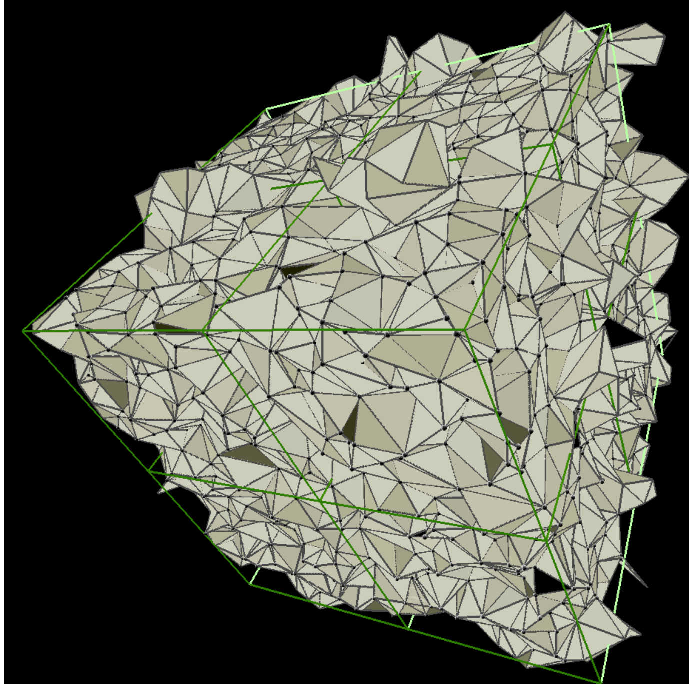

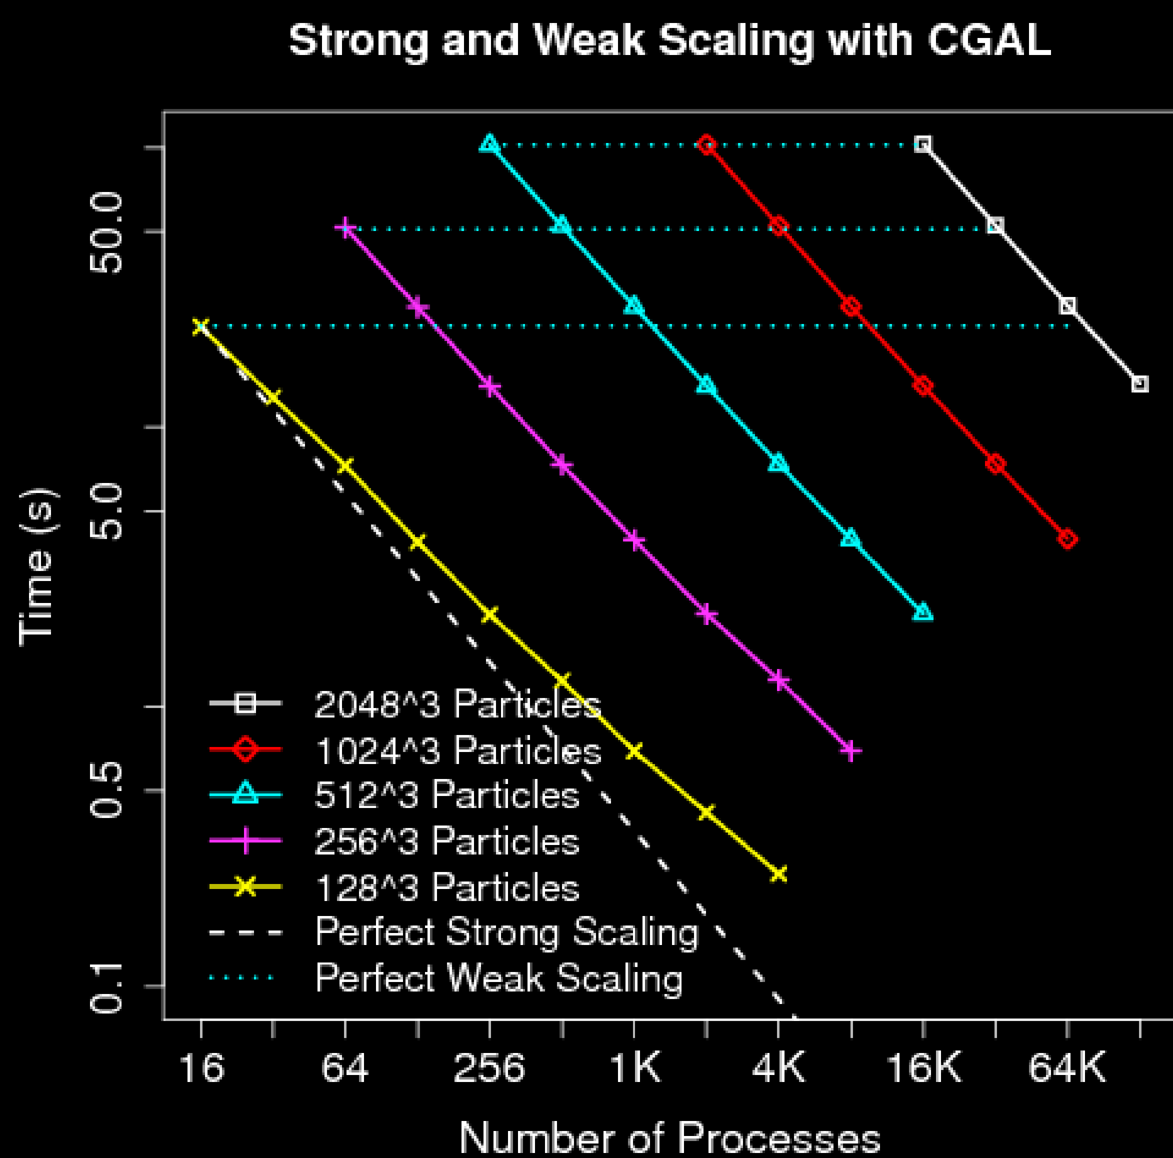

Strong and weak scaling for up to 2048^3 synthetic particles and up to 128K processes (excluding I/O) shows up to 90% strong scaling and up to 98% weak scaling.

Applications in Cosmology

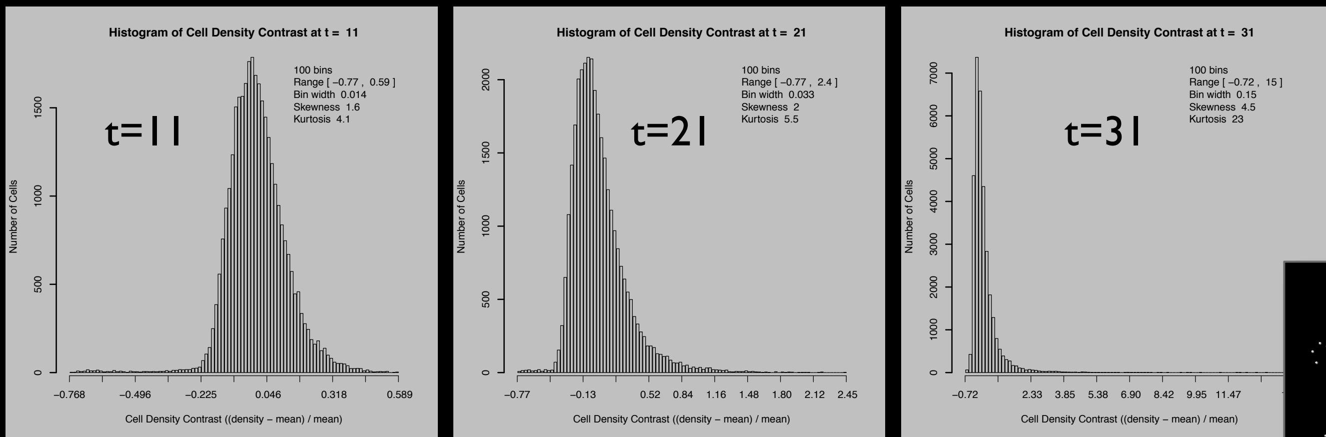

Feature statistics: Total volume, surface area, curvature, topology of connected components of Voronoi cells classify and quantify features.

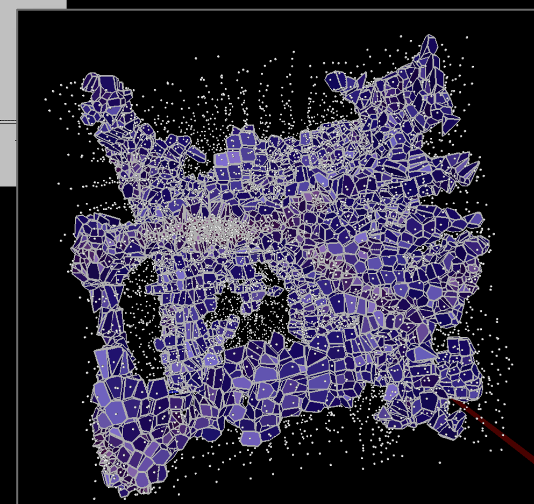

Temporal structure dynamics: As time progresses, the range of cell volume and density expands, kurtosis and skewness increases. These statistics are consistent with the governing physics of the formation of high- and low-density structures over time and can perhaps be used to summarize evolution at given time steps.

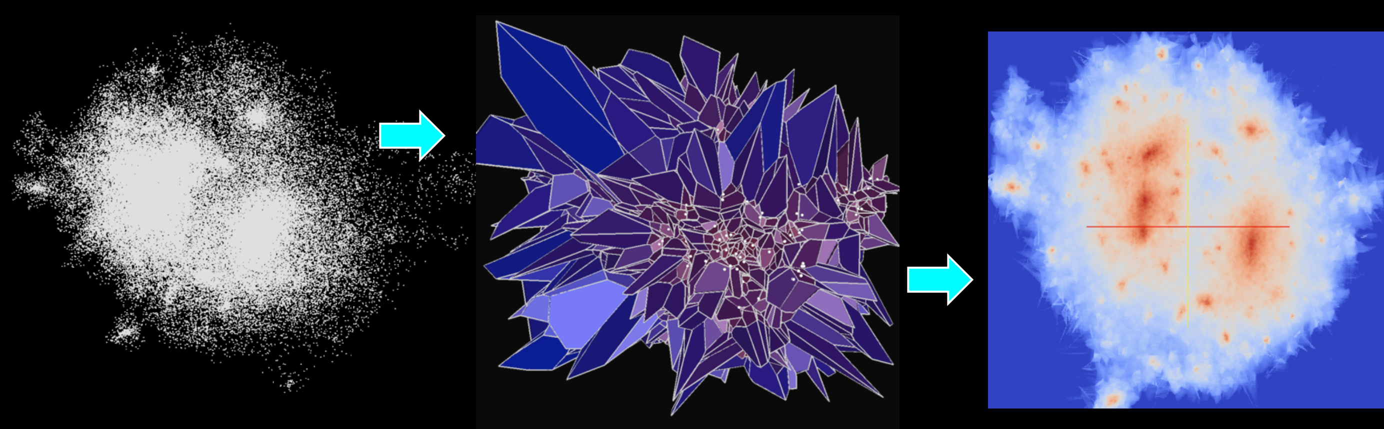

Density estimation: Tessellations as intermediate representations enable accurate regular grid density estimators.

Using Tessellation as a Density Estimator for Regular Grids

sampling a regular density field from a distribution of particle positions using a Voronoi tessellation as an intermediate data model.

Key Ideas

- Convert discrete particle data into continuous function that can be interpolated, differentiated, interpolated, represented as a regular grid (field)
- Automatically adaptive window size and shape
- Comparison with CIC using synthetic and actual data
- Voronoi tessellation and density estimation computed in parallel on distributed-memory HPC machines
- Many visualization and analysis applications are written for regular grids
- Applications in astrophysics, environmental science, social science

Cloud in Cell (CIC)

The mass of point P is distributed among nearest grid points G_0 – G_7.

The volume of of the grid cube with corners G_0 – G_7, $v(G_0, G_7)$ is normalized to 1.0

The mass assigned to grid point G_i is $m(G_i) = 1.0 - v(G_i, P)$

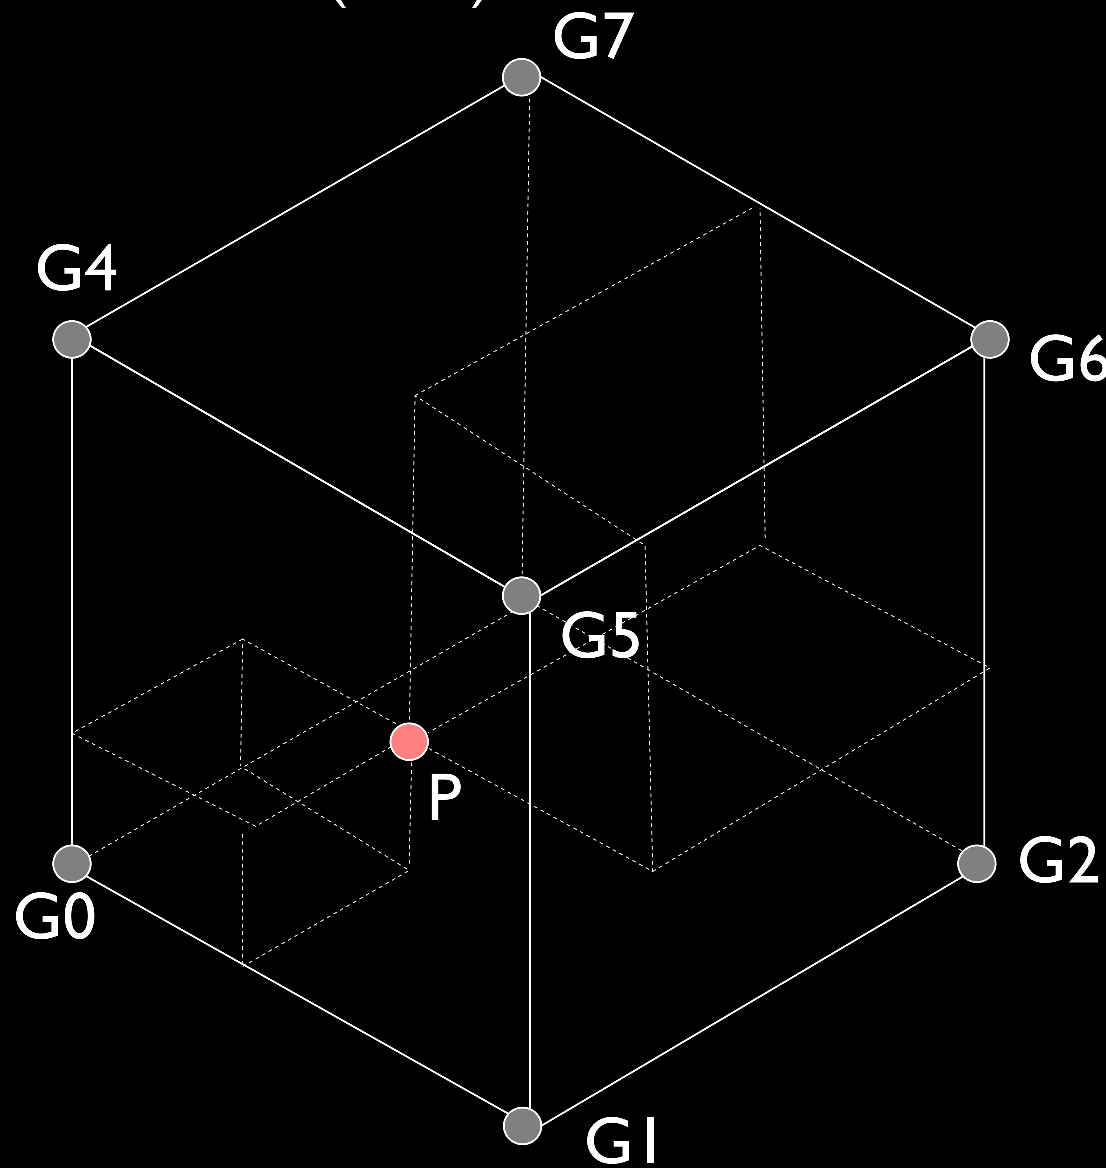

Tessellation (TESS)

Parameter free: no fixed window size determined by grid or number of particles

Kernel free: no smoothing kernel

Shape free: asymmetrical, no window or kernel shape

Automatically adaptive

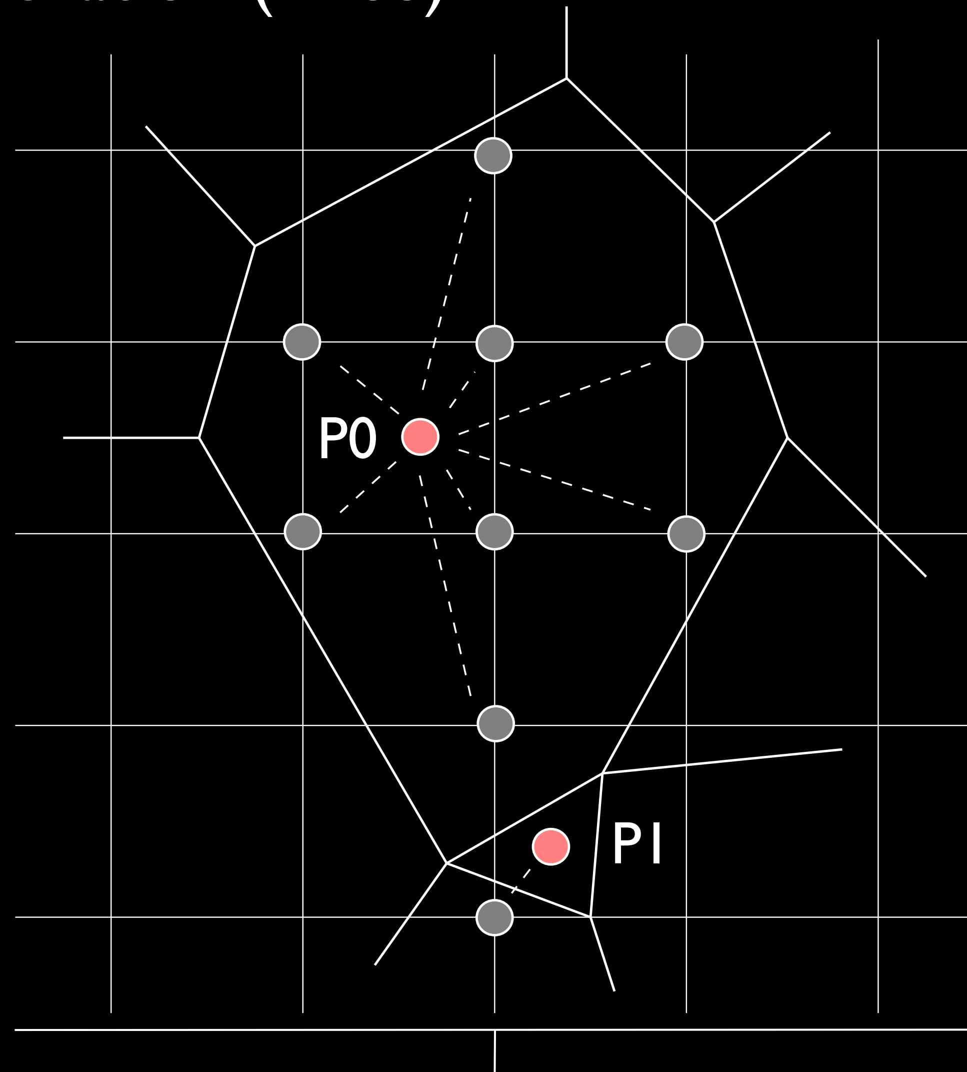

P0 is a particle whose Voronoi cell covers several grid points. Its mass is uniformly distributed (zero-order estimation) to those grid points. P1 is a small cell that covers no grid points. Its mass is assigned to the nearest grid point.

Dense Strong and Weak Scaling

- 128^3 synthetic particles
- End-to-end time (including reading tessellation and writing image)
- 3D->2D projection
- 51% strong scaling (End-to-end) for 4096^3 grid

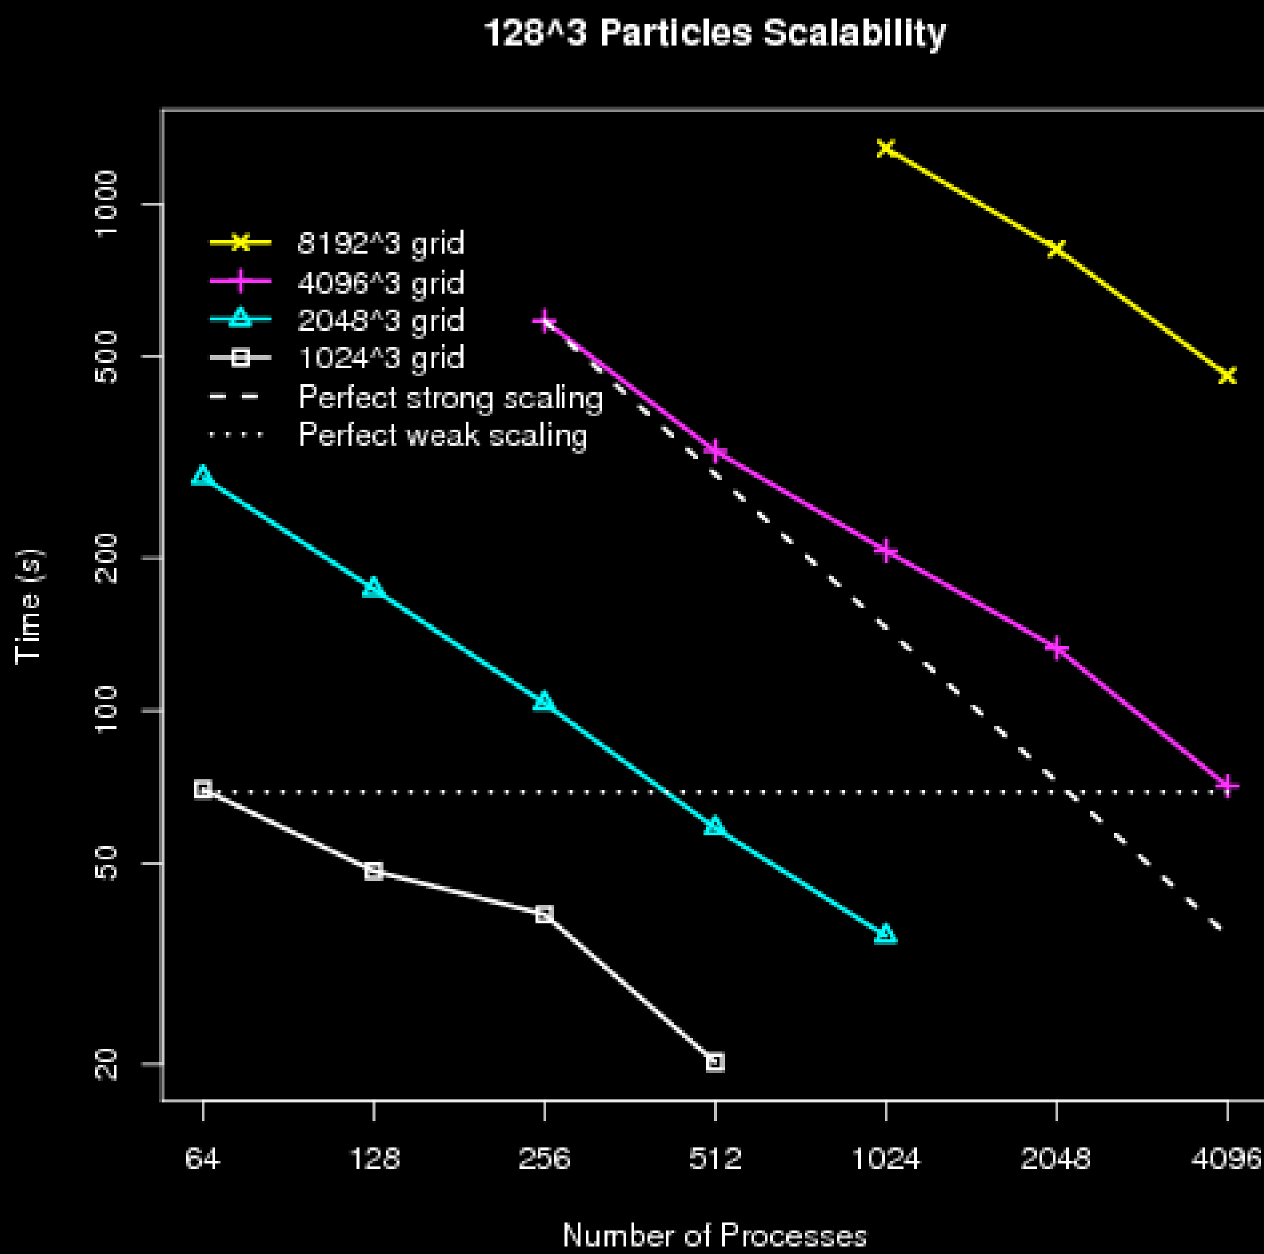

Navarro-Frenk-White (NFW)

Our first synthetic dataset is derived from an analytical density function commonly used in cosmology.

k is a constant, 1 for us

ρ (r) is Monte Carlo sampled to get test set of particles

Ground truth is 2D plot of ρ (r)

We limit r to [-1.5, 1.5] and NFW(r) to 10^6

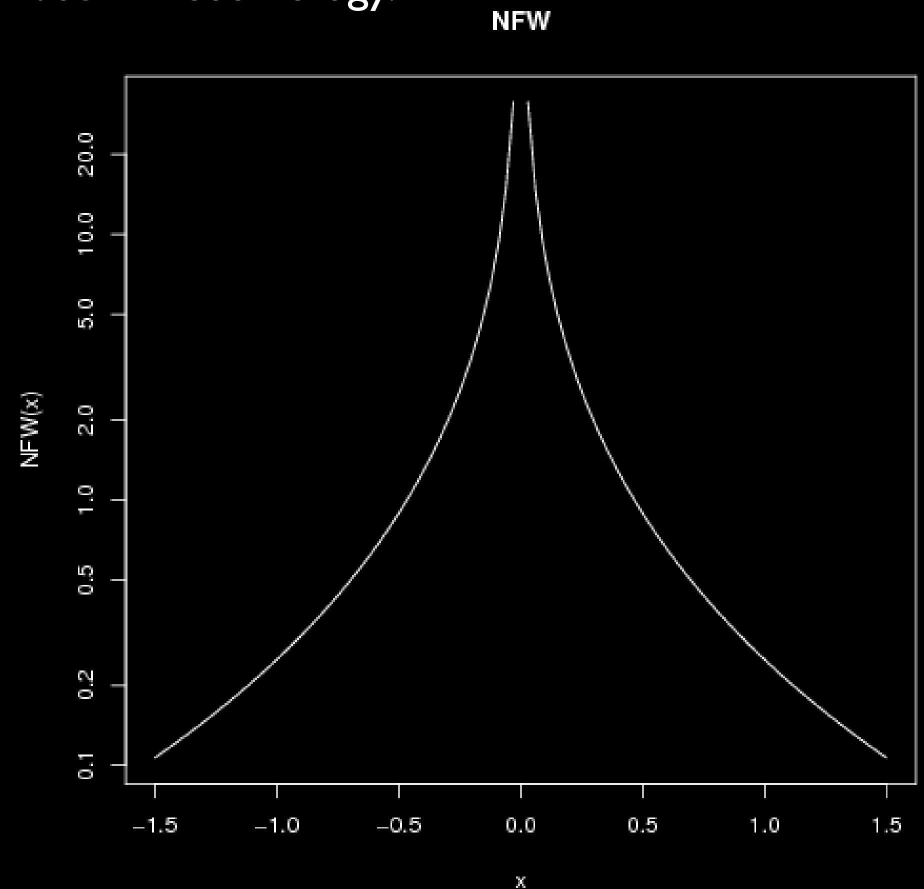

$$\rho(r) = \frac{k}{(r(r+1)^2)}$$

Visual Comparison

Analytical TESS CIC

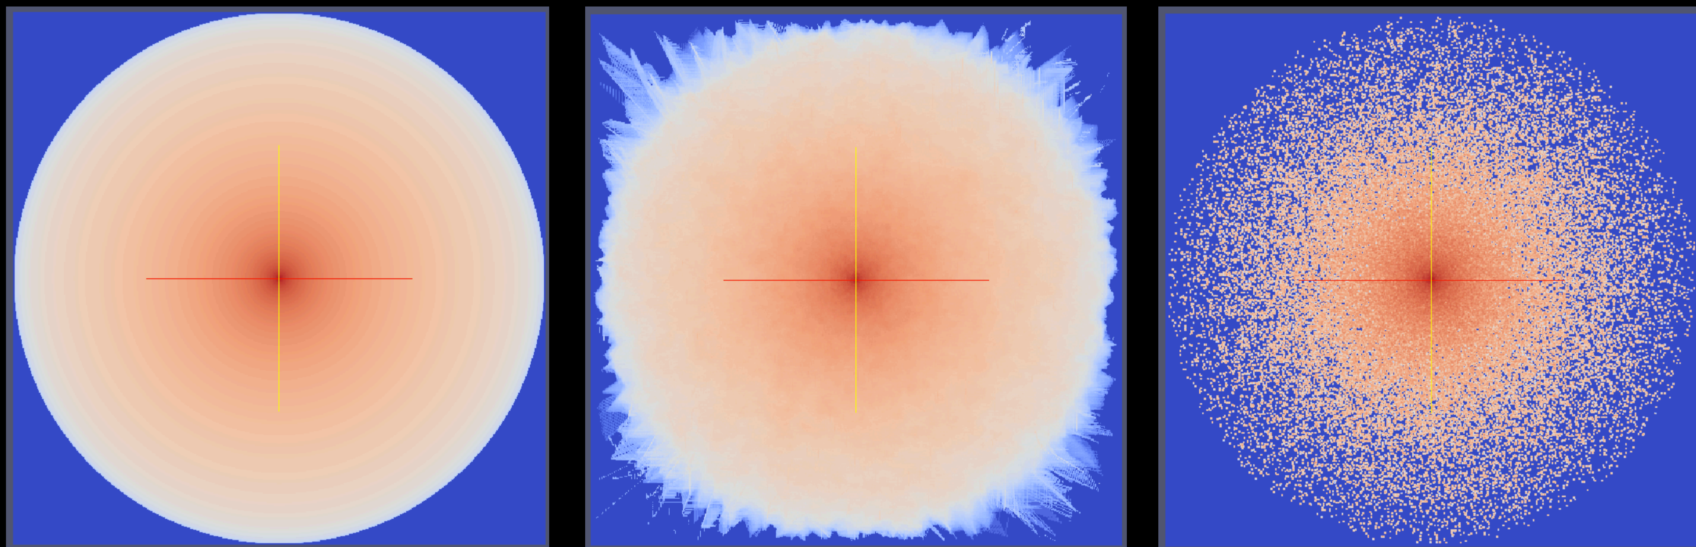

Top row:
1024^3 3D density projected to 1024^2 2D density field and rendered in ParaView

Bottom row:
Ratio of analytical divided by estimated density

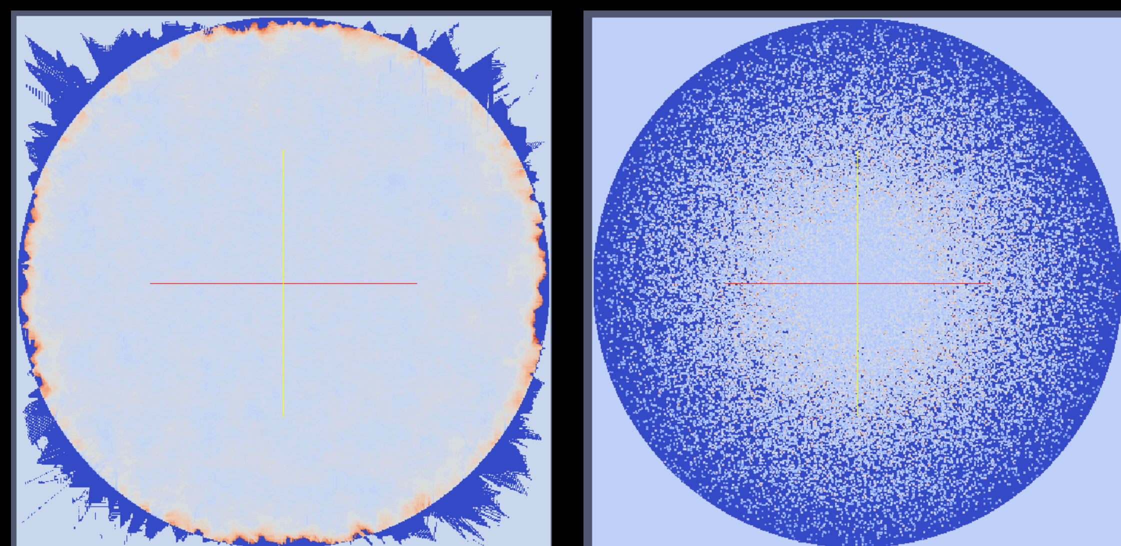

Analytical Comparison

Comparison between analytical 2D density and estimated density at y = 0 cross section for TESS

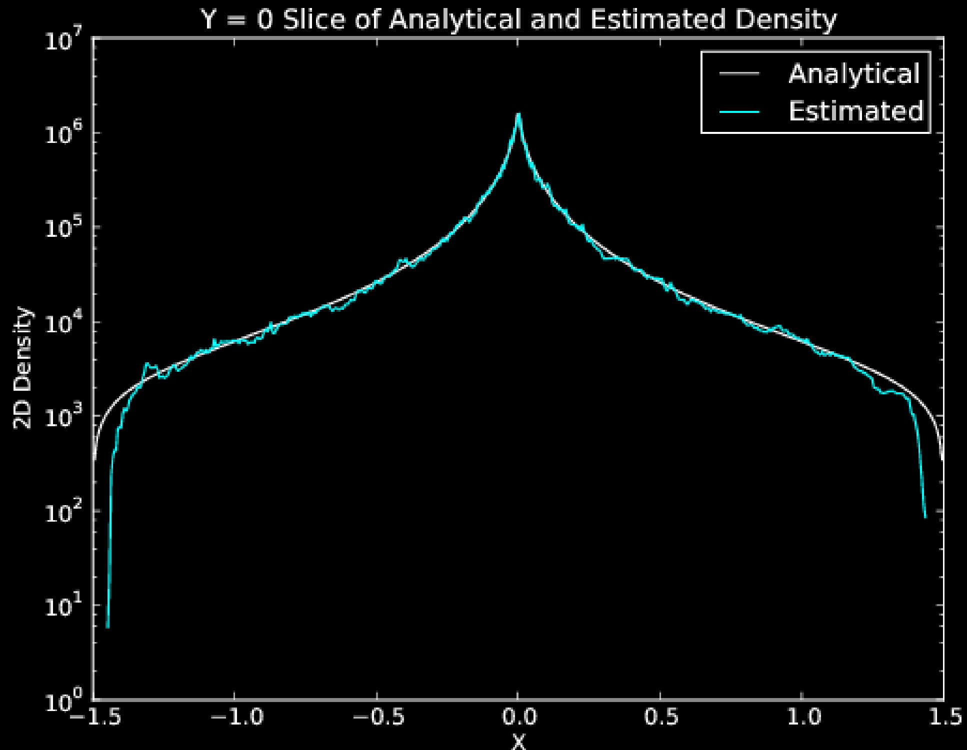

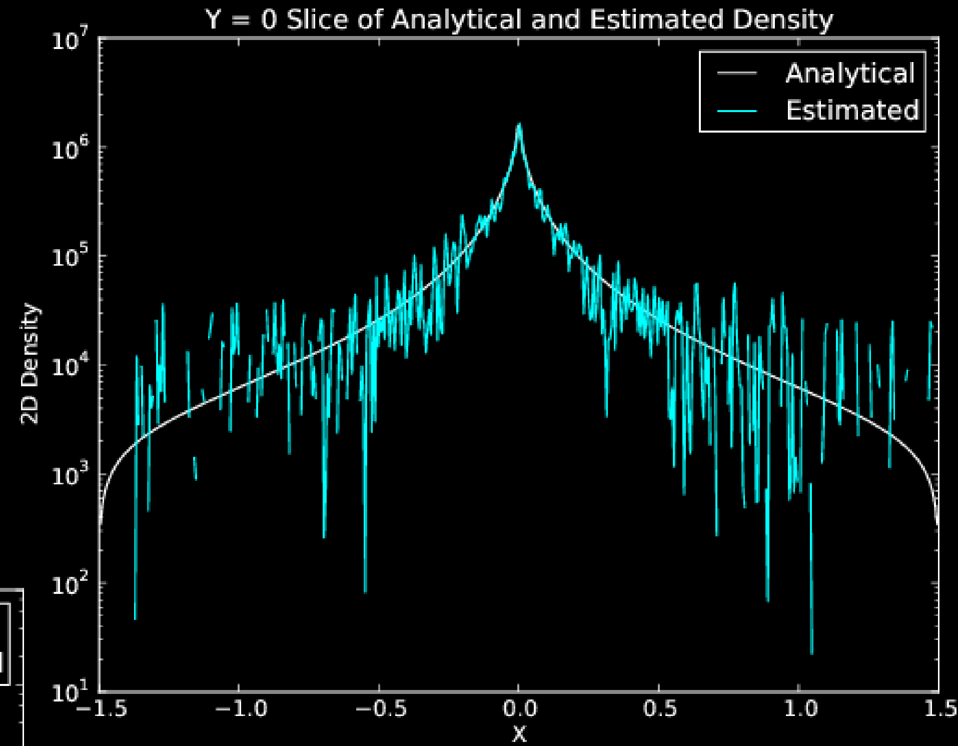

Same for CIC

Application: 2D Density of Halo

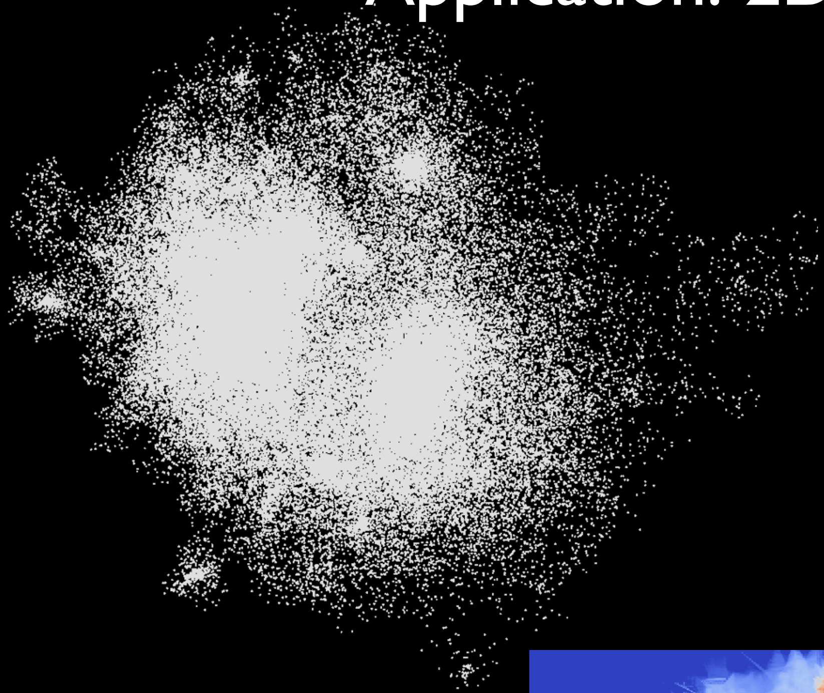

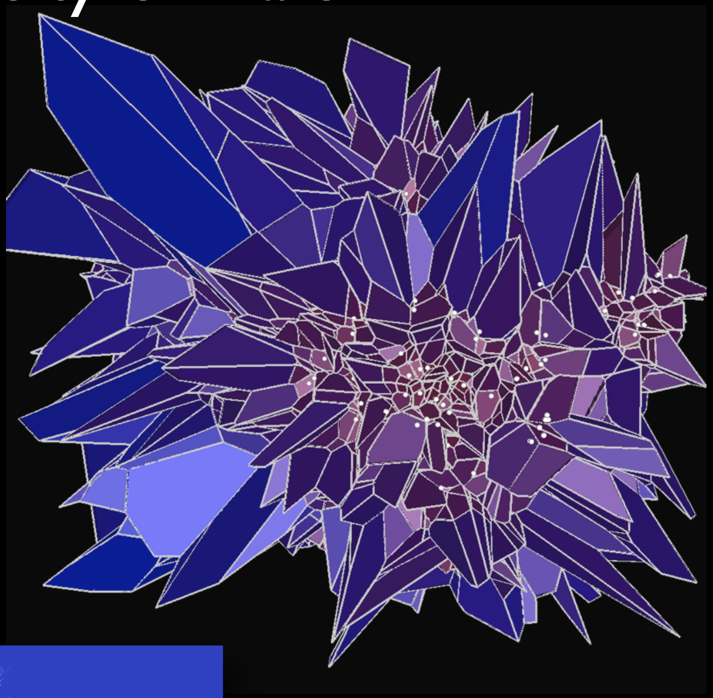

Particle data from HACC N-body cosmology code from halo ID 7445077095

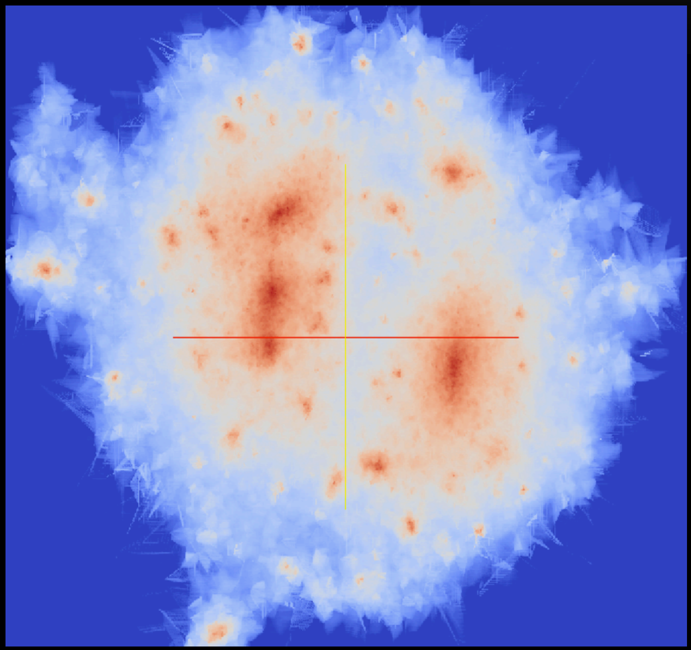

Voronoi tesellation of halo particles colored by cell volume

Final output 2D density field for lensing

Recap and Looking Ahead

Done: Benefits

- Productivity
  - Express complex algorithms flexibly
  - Simplify existing tasks
- Performance
  - Published scalability
  - Less data movement, earlier results
- Applications
  - Spatiotemporal data
  - Regular, unstructured, particle
  - Simulation and experiment

To Do: Research Directions

- Irregular communication patterns
  - Graphs
  - Load balancing, dynamic partitioning
- Stochastic, approximate algorithms
  - Limited resources
- Streaming and out of core algorithms
  - Temporal analysis
- Load balancing
  - Block overloading, dynamic reassignment
- Hybrid programming models
  - MPI + X
- Coupling analysis in workflows
  - Flexible coupling, decoupling

Further Reading

DIY

- Peterka, T., Ross, R., Kendall, W., Gyulassy, A., Pascucci, V., Shen, H.-W., Lee, T.-Y., Chaudhuri, A.: Scalable Parallel Building Blocks for Custom Data Analysis. Proceedings of Large Data Analysis and Visualization Symposium (LDAV'11), IEEE Visualization Conference, Providence RI, 2011.
- Peterka, T., Ross, R.: Versatile Communication Algorithms for Data Analysis. 2012 EuroMPI Special Session on Improving MPI User and Developer Interaction IMUDI'12, Vienna, AT.

DIY applications

- Peterka, T., Ross, R., Nouanesengsey, B., Lee, T.-Y., Shen, H.-W., Kendall, W., Huang, J.: A Study of Parallel Particle Tracing for Steady-State and Time-Varying Flow Fields. Proceedings IPDPS'11, Anchorage AK, May 2011.
- Gyulassy, A., Peterka, T., Pascucci, V., Ross, R.: The Parallel Computation of Morse-Smale Complexes. Proceedings of IPDPS'12, Shanghai, China, 2012.
- Nouanesengsy, B., Lee, T.-Y., Lu, K., Shen, H.-W., Peterka, T.: Parallel Particle Advection and FTLE Computation for Time-Varying Flow Fields. Proeedings of SC12, Salt Lake, UT.
- Peterka, T., Kwan, J., Pope, A., Finkel, H., Heitmann, K., Habib, S., Wang, J., Zagaris, G.: Meshing the Universe: Integrating Analysis in Cosmological Simulations. Proceedings of the SC12 Ultrascale Visualization Workshop, Salt Lake City, UT.
- Chaudhuri, A., Lee-T.-Y., Zhou, B., Wang, C., Xu, T., Shen, H.-W., Peterka, T., Chiang, Y.-J.: Scalable Computation of Distributions from Large Scale Data Sets. Proceedings of 2012 Symposium on Large Data Analysis and Visualization, LDAV'12, Seattle, WA.

I have had my results for a long time: but I do not yet know how I am to arrive at them.
Carl Friedrich Gauss.

Acknowledgments:

Facilities
Argonne Leadership Computing Facility (ALCF)
Oak Ridge National Center for Computational Sciences (NCCS)

Funding
DOE SDMAV Exascale Initiative
DOE Exascale Codesign Center
DOE SciDAC SDAV Institute

https://svn.mcs.anl.gov/repos/diy/trunk

Tom Peterka
tpeterka@mcs.anl.gov
Mathematics and Computer Science Division

HPC Geospatial Analytics Workshop
April 29, 2014