ChaNGa

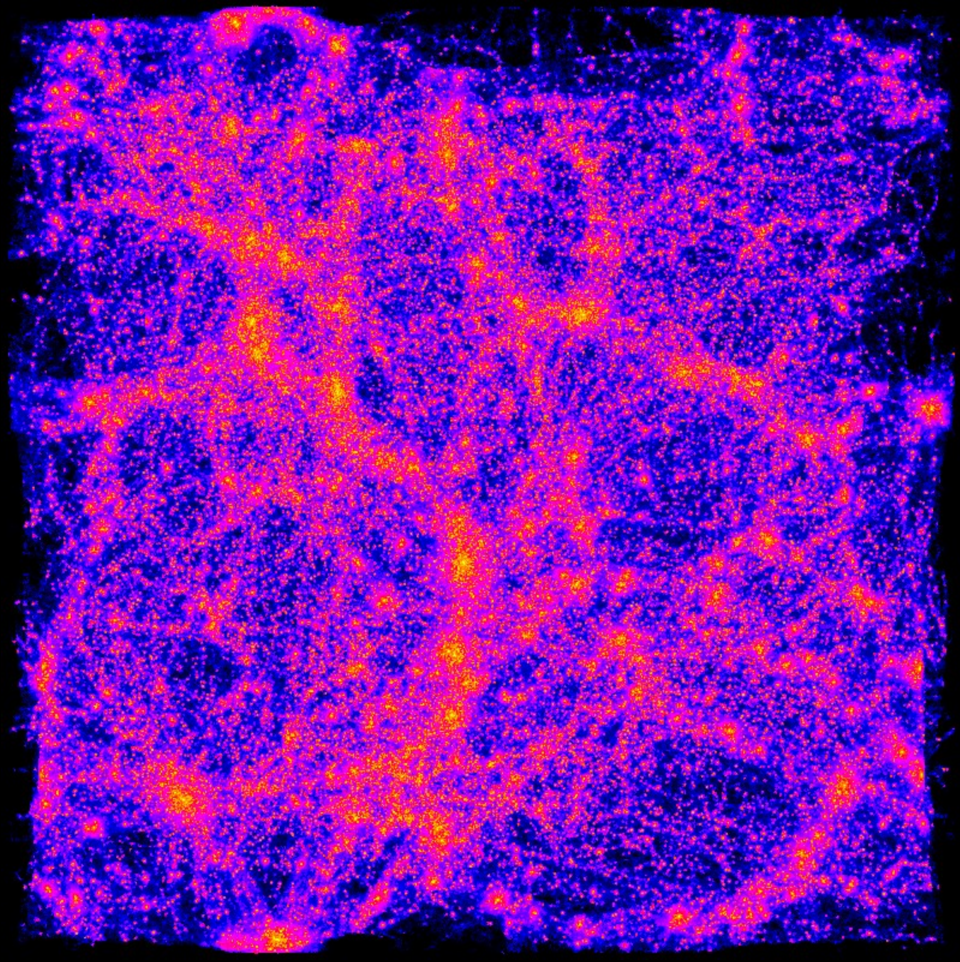

CHArm N-body GrAvity

PPL PARALLEL PROGRAMMING LAB
DEPT OF COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS

Thomas Quinn

Graeme Lufkin

Joachim Stadel

Laxmikant Kale

Filippo Gioachin

Pritish Jetley

Celso Mendes

Amit Sharma

Outline

- Scientific background
  - How to build a Galaxy
  - Types of Simulations
  - Simulation Challenges
- ChaNGa and those Challenges
  - Features
  - Tree gravity
  - Load balancing
  - Multistepping
- Future Challenges

Cosmology: How does this ...

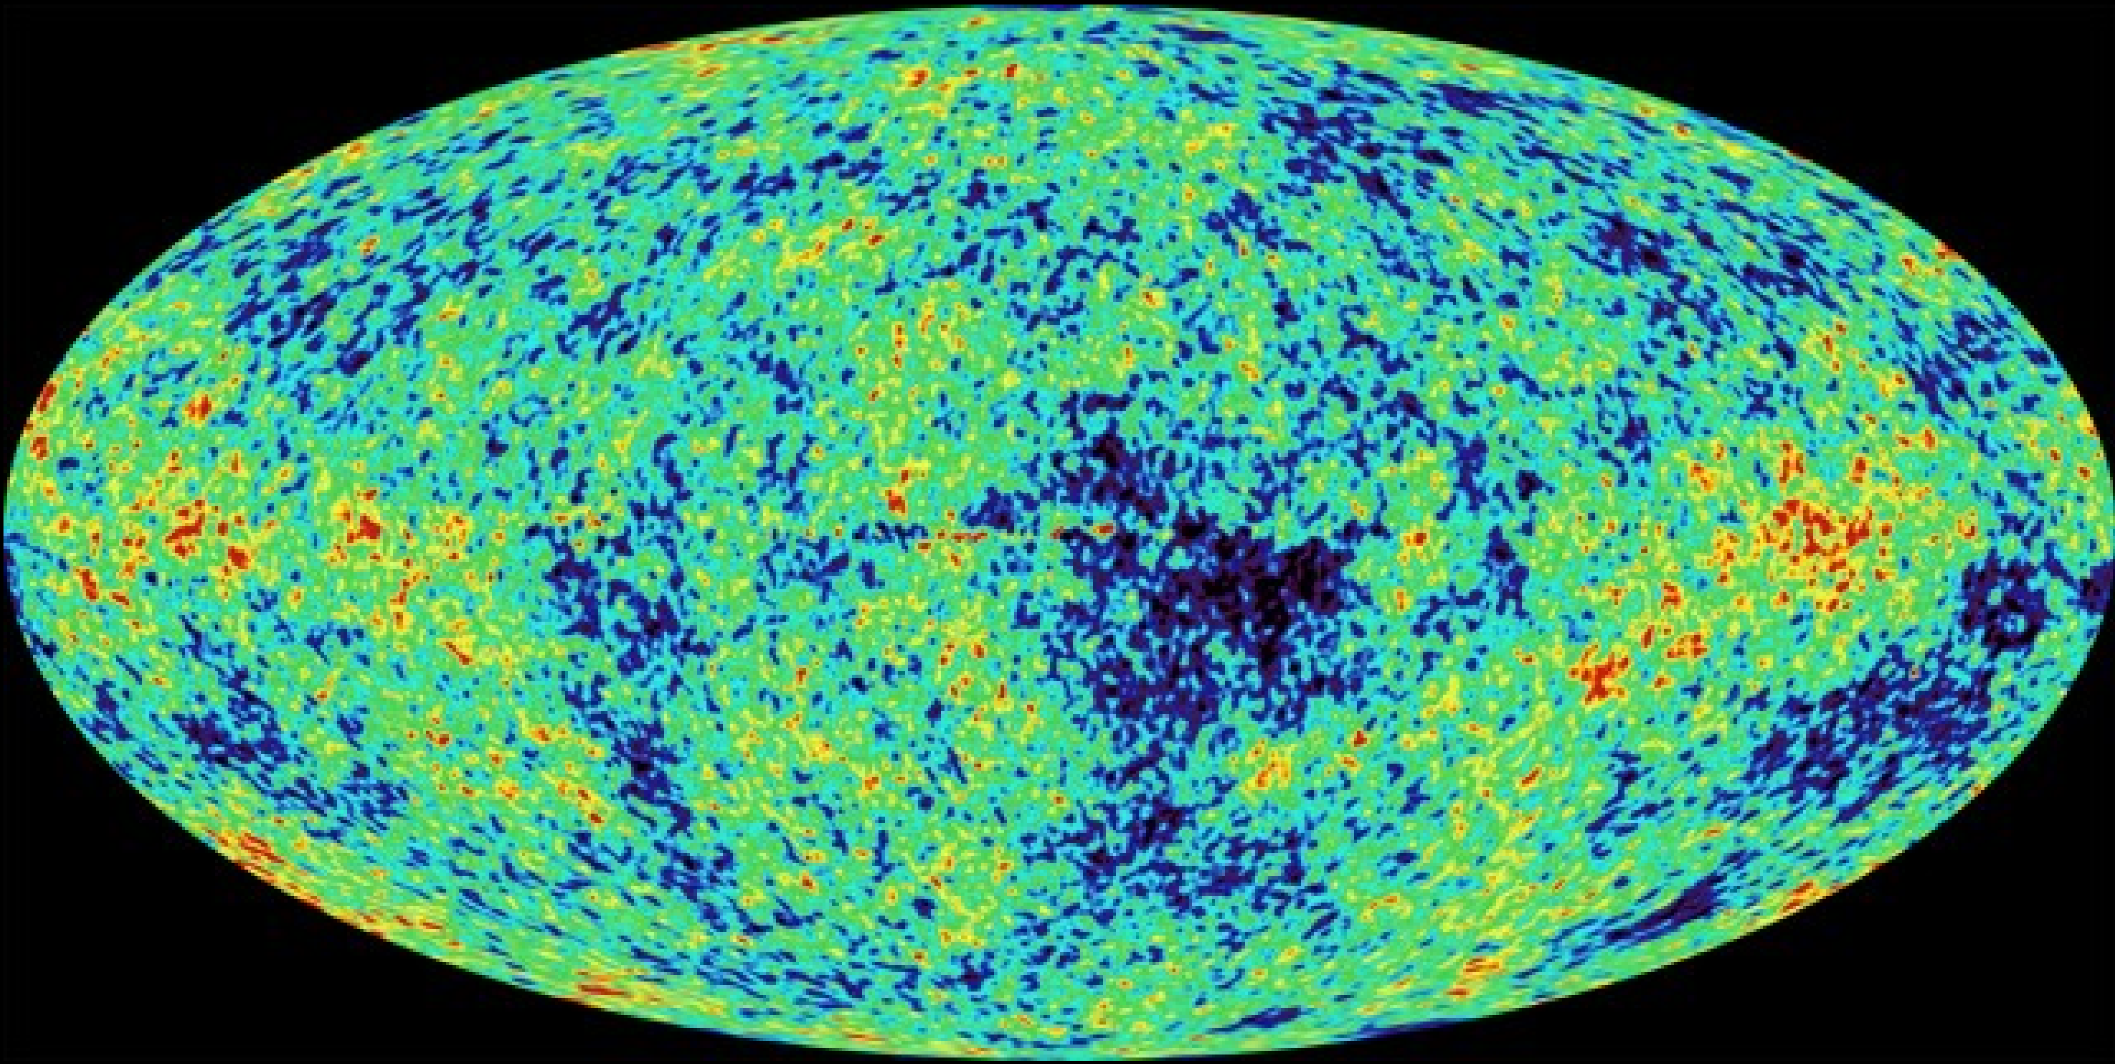

Image courtesy NASA/WMAP

... turn into this?

Computational Cosmology

CMB gives fluctuations of 1e-5

Galaxies are overdense by 1e7

It happens through **Gravitational Collapse**

Making testable predictions from a cosmological hypothesis requires

Non-linear, dynamic calculation

e.g. **Computer simulation**

Simulation process

Start with fluctuations based on Dark Matter properties

Follow model analytically (good enough to get CMB)

Create a realization of these fluctuations in particles.

Follow the motions of these particles as they interact via gravity.

Compare final distribution of particles with observed properties of galaxies.

Simulating galaxies: Procedure

Simulate 100 Mpc volume at 10-100 kpc resolution

Pick candidate galaxies for further study

Resimulate galaxies with same large scale structure but with higher resolution, and lower resolution in the rest of the computational volume.

At higher resolutions, include gas physics and star formation.

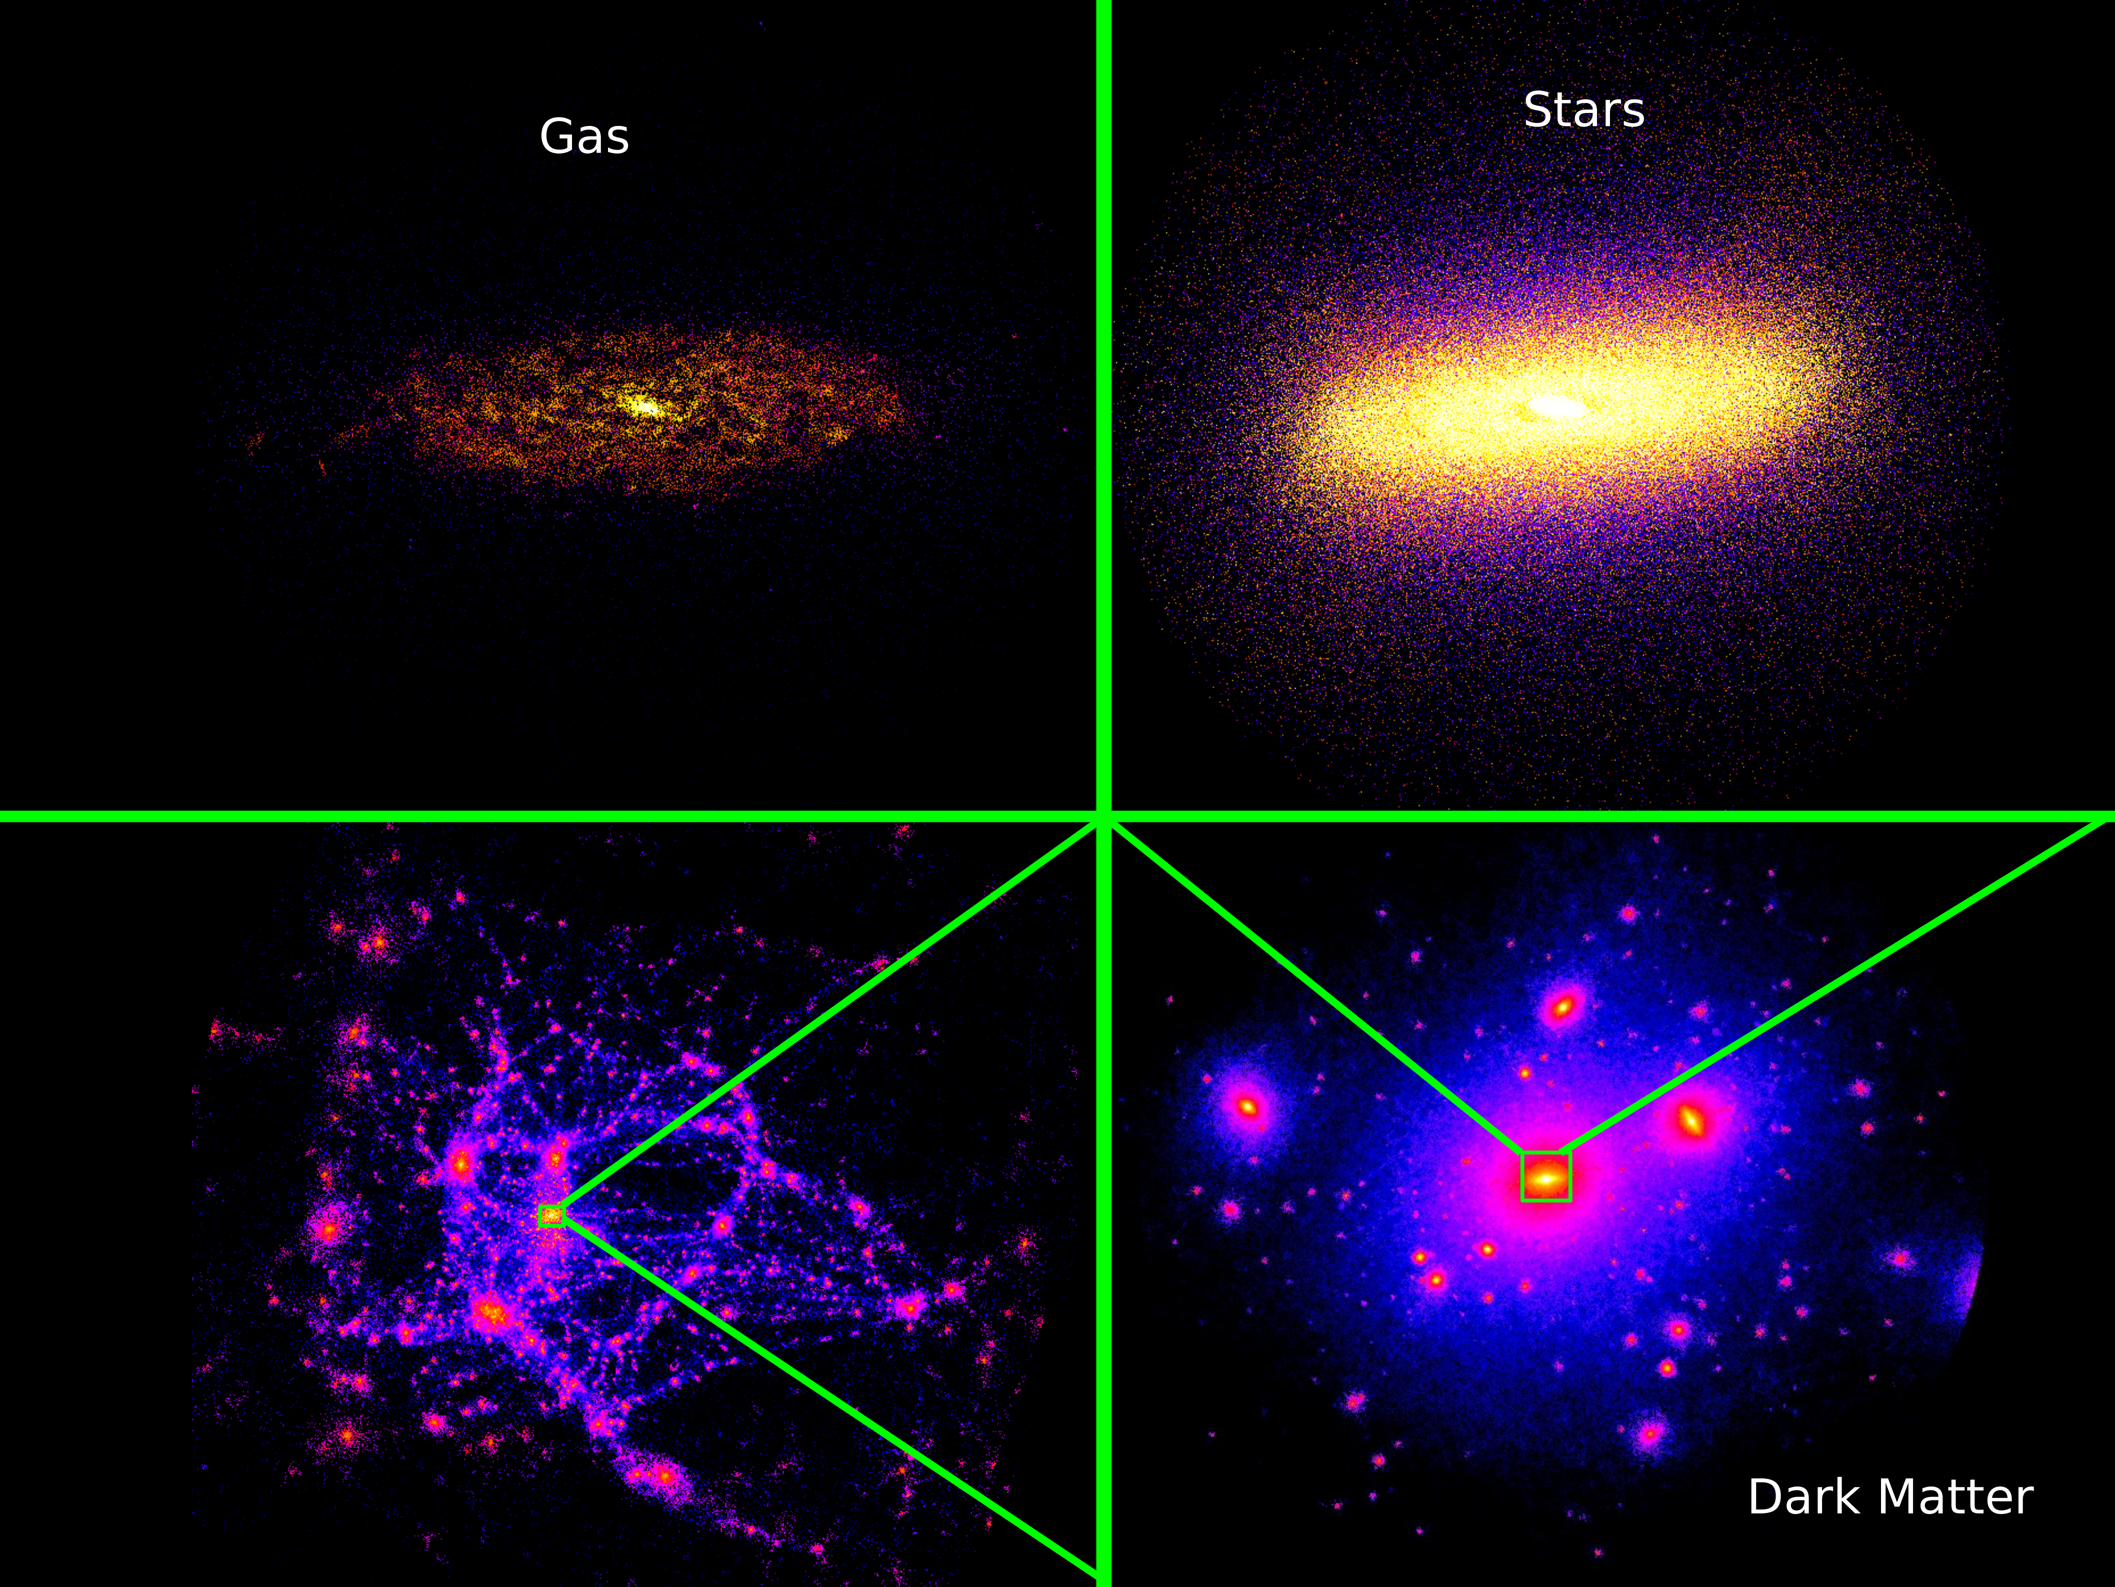
Gas
Stars
Dark Matter

Types of simulations

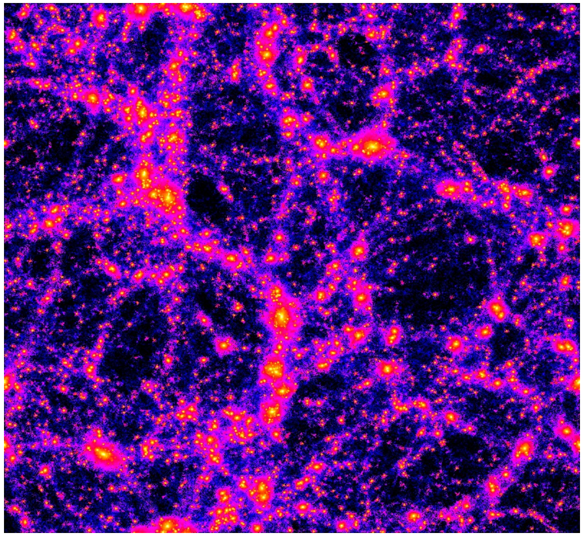

"Uniform" Volume

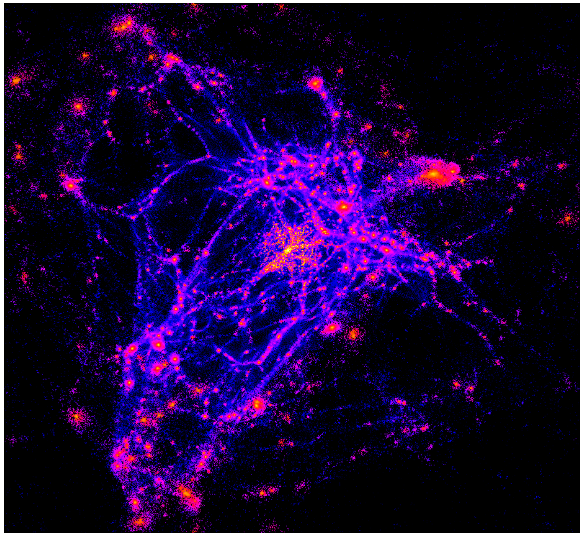

Zoom In

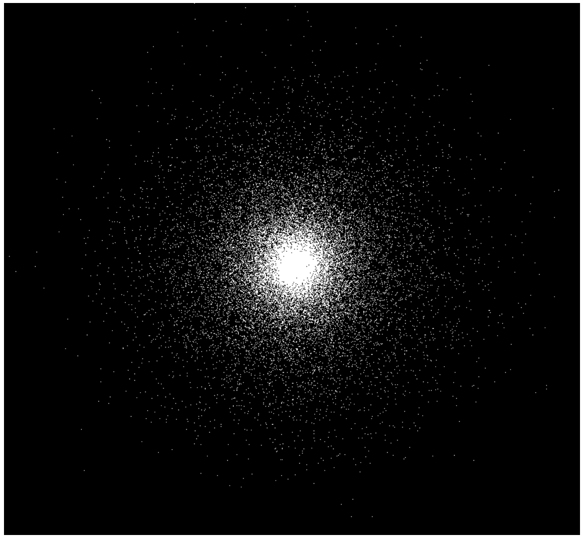

Isolated Cluster/ Galaxy

Star Formation History: What's needed

1 Million particles/galaxy for proper morphology/heavy element production

1 Petaflop-week of computation

Necessary for:

Comparing with Hubble Space Telescope surveys of the local Universe

Interpreting HST images of high redshift galaxies

Large Scale Structure: What’s needed

6.5 Gigaparsec volume

10 Trillion particles (1 Petabyte of RAM)

1 Petaflop week of computation

Necessary for:

Interpreting future surveys (LSST)

Relating Cosmic Microwave Background to galaxy surveys

Halo Simulations: What's needed

Billions of particles in a single halo

Dark Matter detection experiments

Influence on disk

Theories of gravitational collapse (Insight!)

Computational Challenges

Large spacial dynamic range: > 100 Mpc to < 1 kpc

Hierarchical, adaptive gravity solver is needed

Large temporal dynamic range: 10 Gyr to 1 Myr

Multiple timestep algorithm is needed

Gravity is a long range force

Hierarchical information needs to go across processor domains

ChaNGa Features

Tree-based gravity solver

High order multipole expansion

Periodic boundaries (if needed)

Individual multiple timesteps

Dynamic load balancing with choice of strategies

Checkpointing (via migration to disk)

Visualization

Cosmological Comparisons: Mass Function

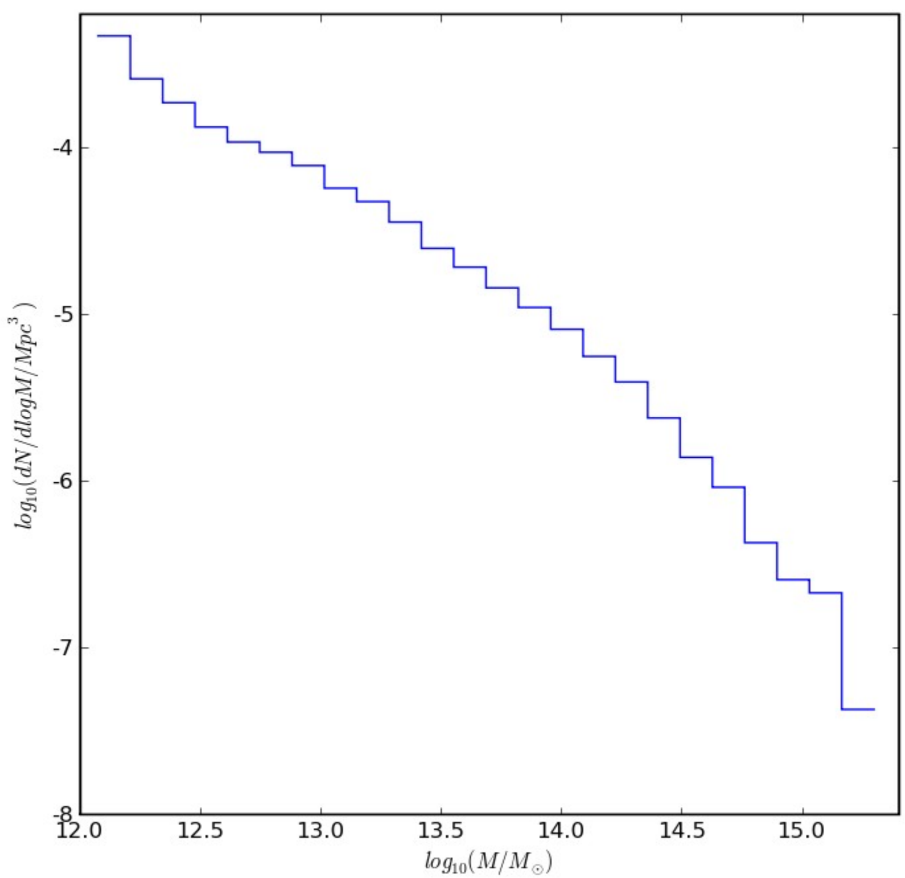

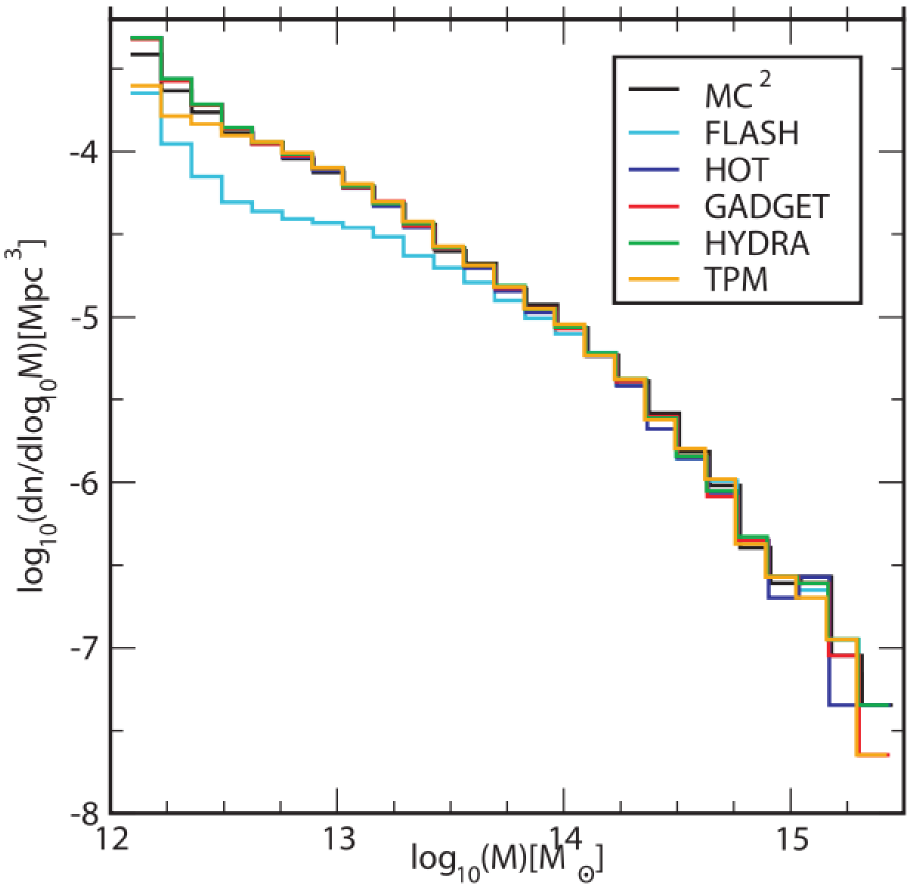

Heitmann, et al. 2005

Basic algorithm ...

Newtonian gravity interaction

Each particle is influenced by all others: O(n^2) algorithm

Barnes-Hut approximation: O($n\log n$)

Influence from distant particles combined into center of mass

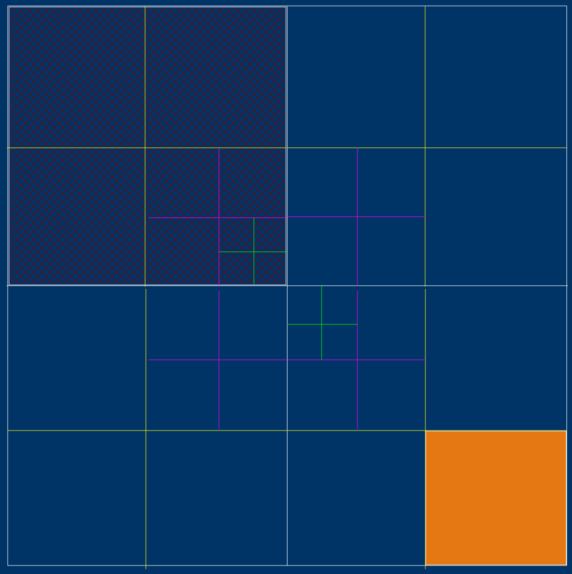

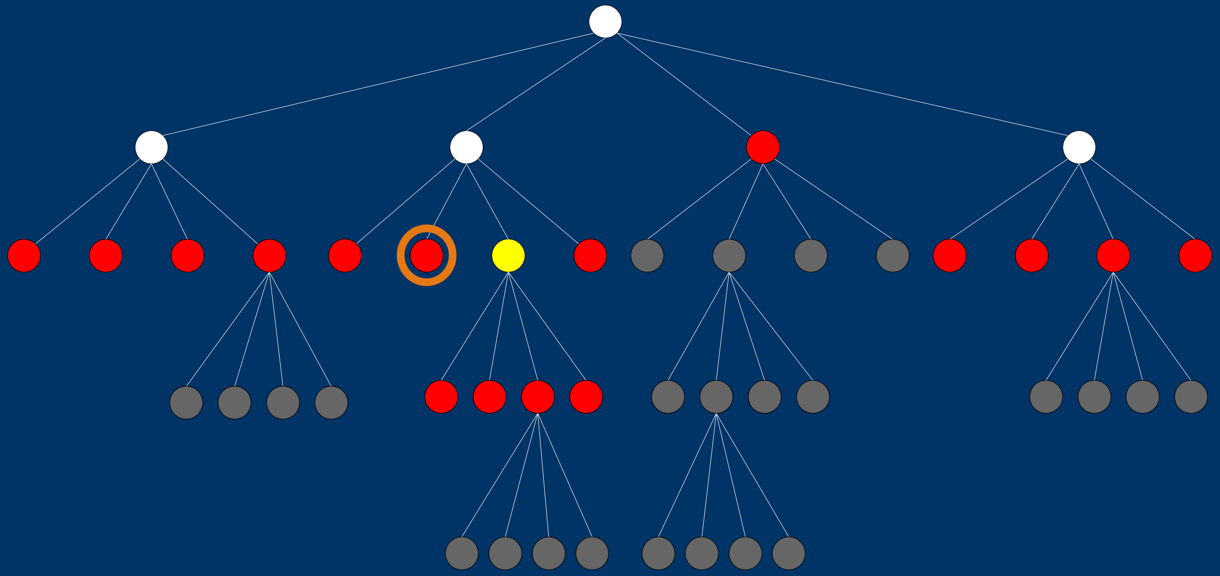

Overall algorithm

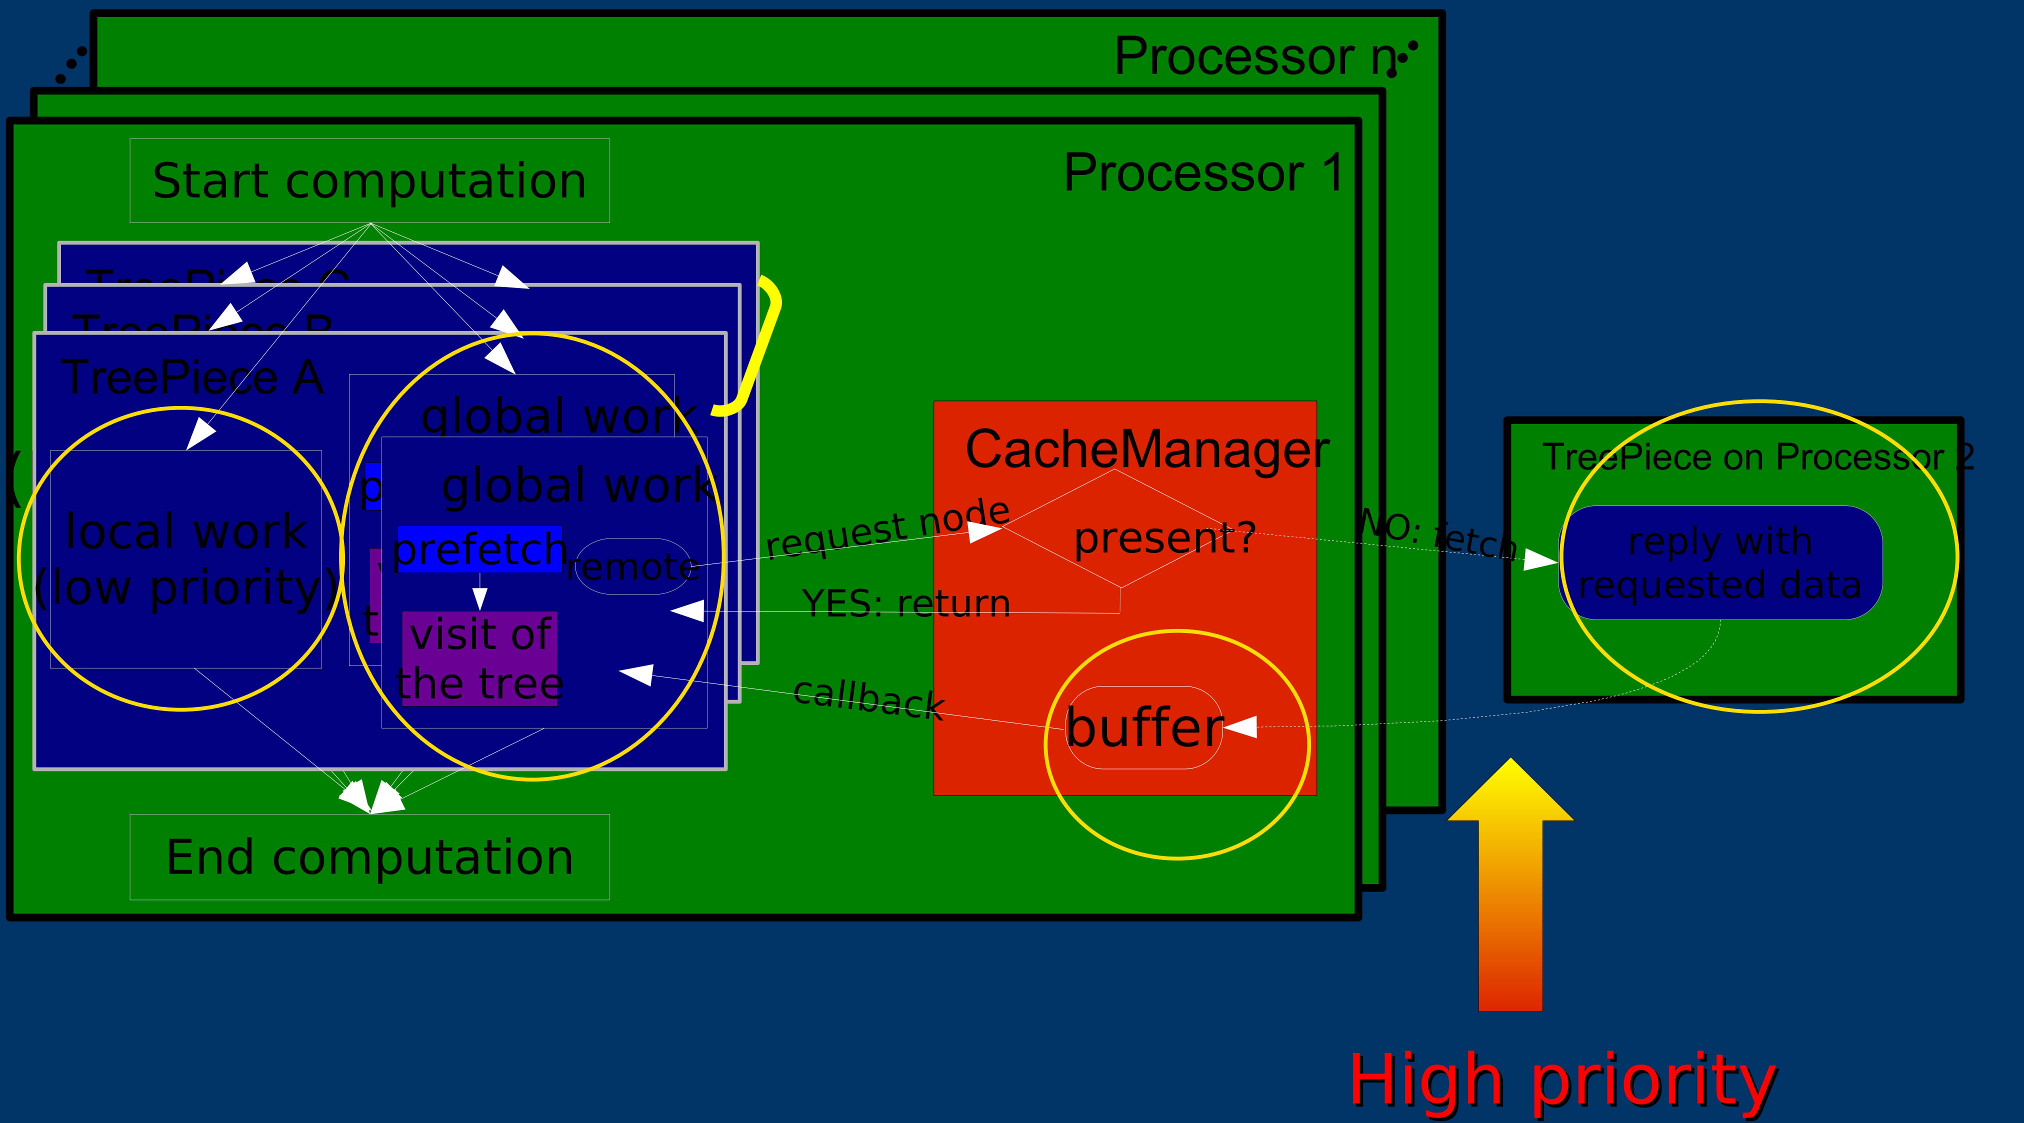

Uniform Volume Scaling

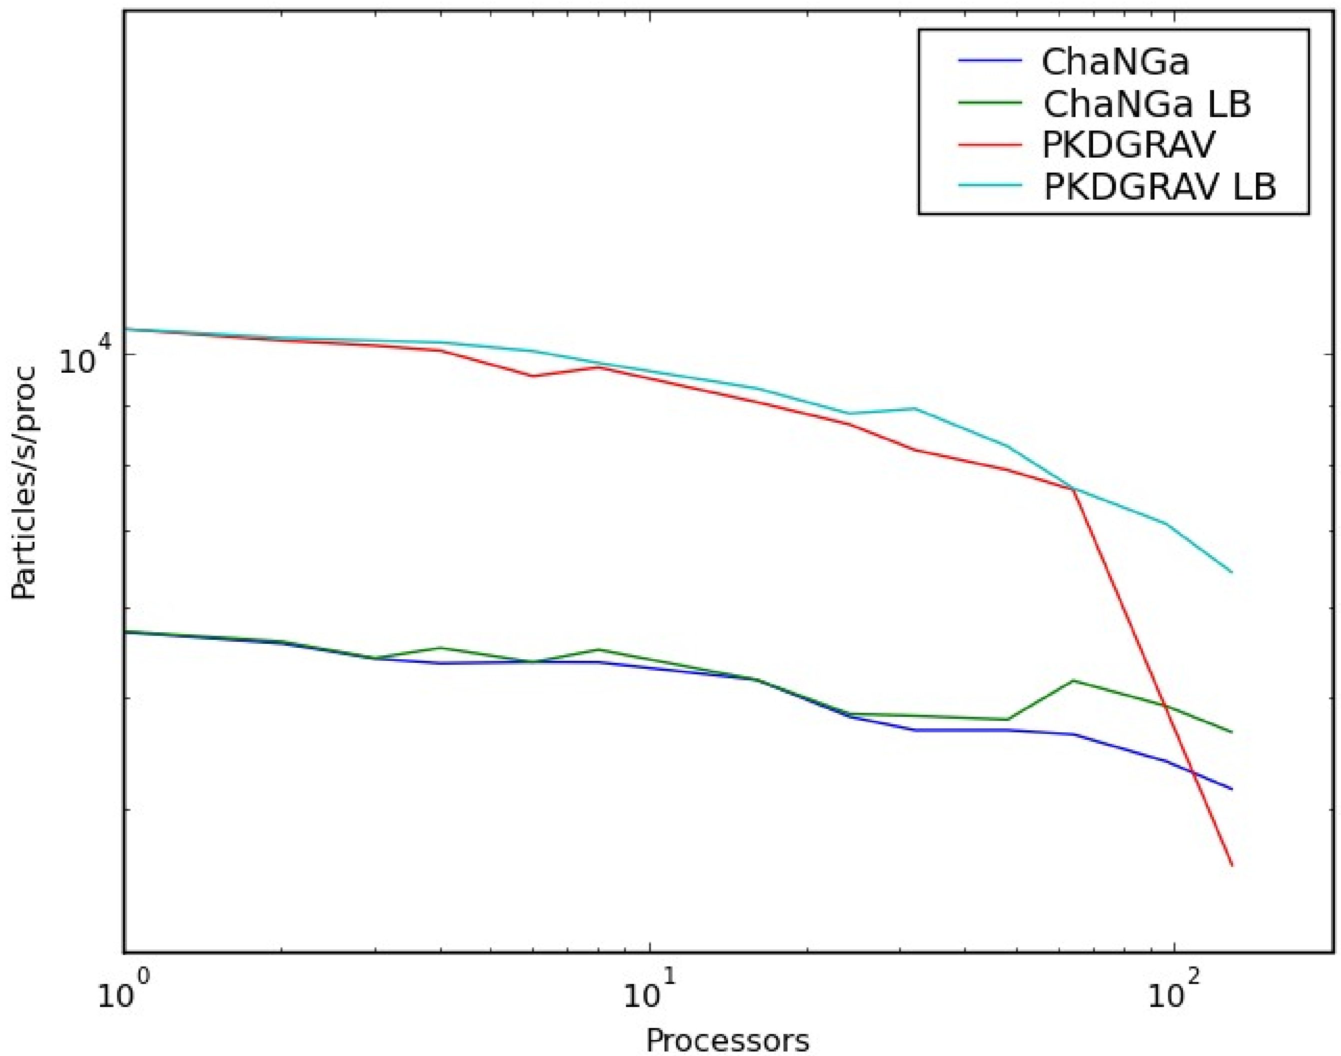

Zoom-in Scaling

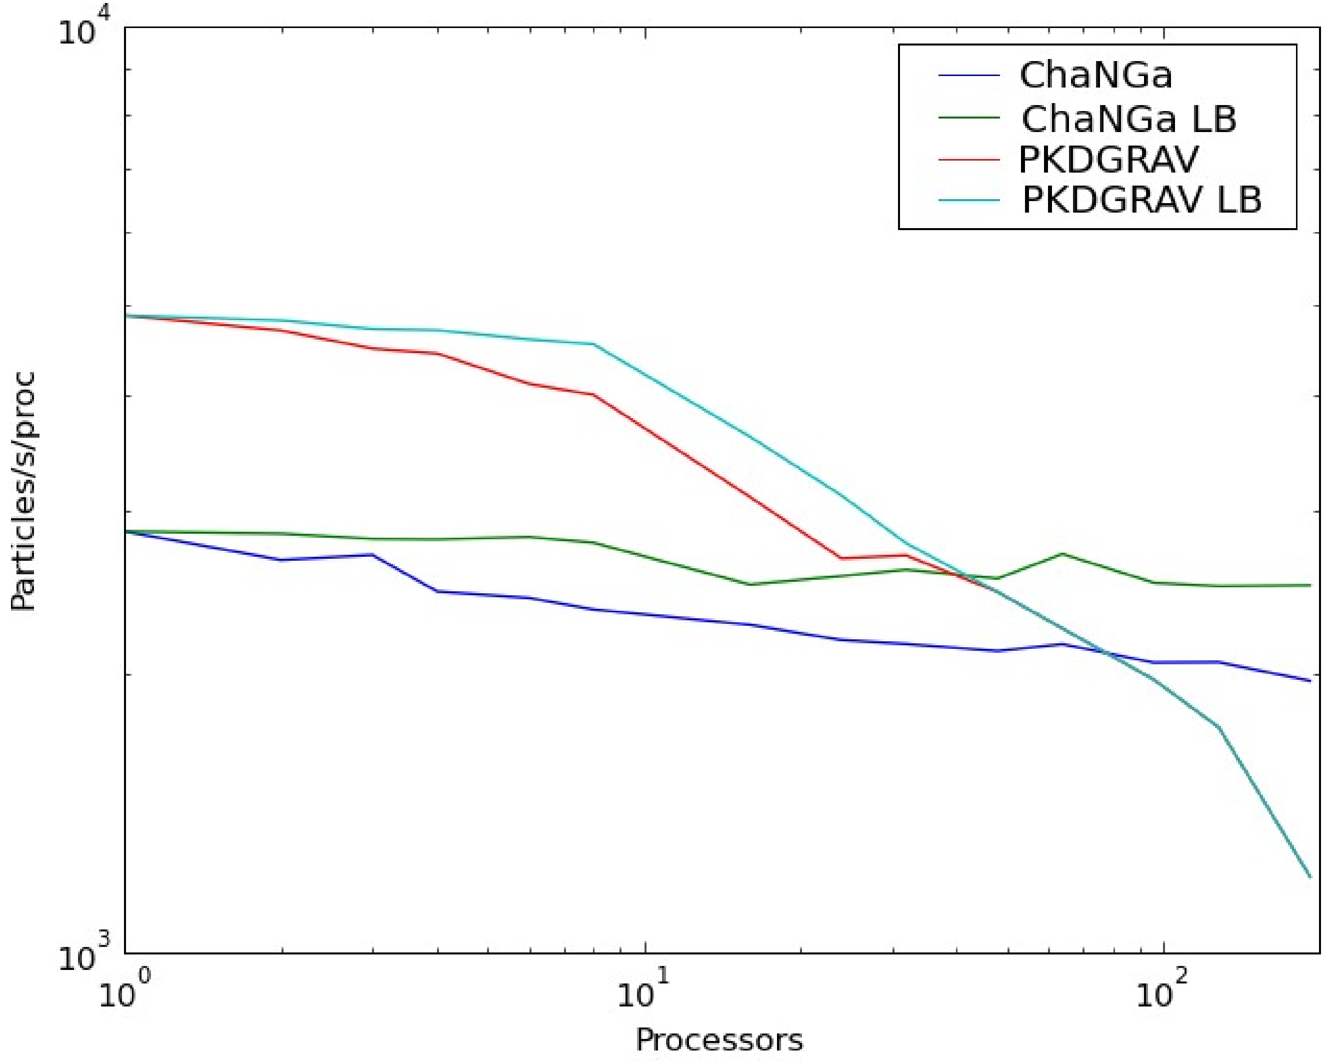

Clustering and Load Balancing

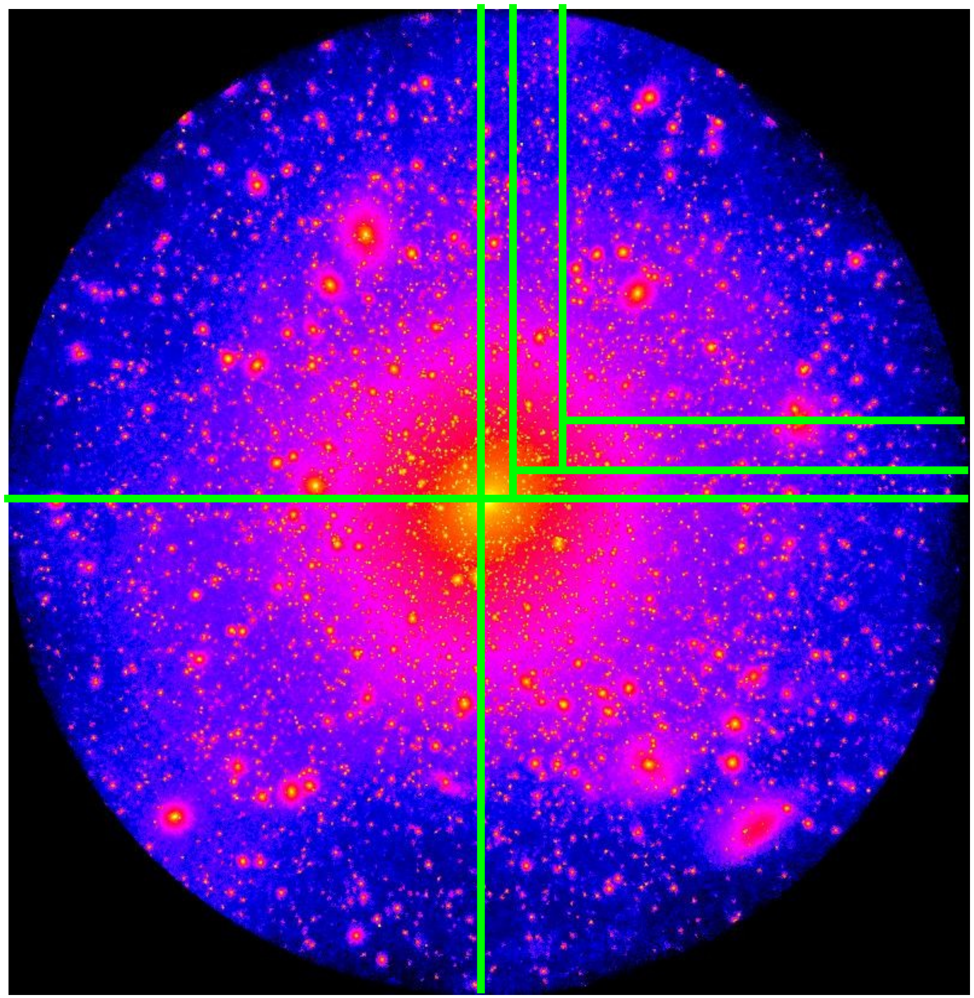

Remote/local latency hiding

Clustered data on 1,024 BlueGene/L processors

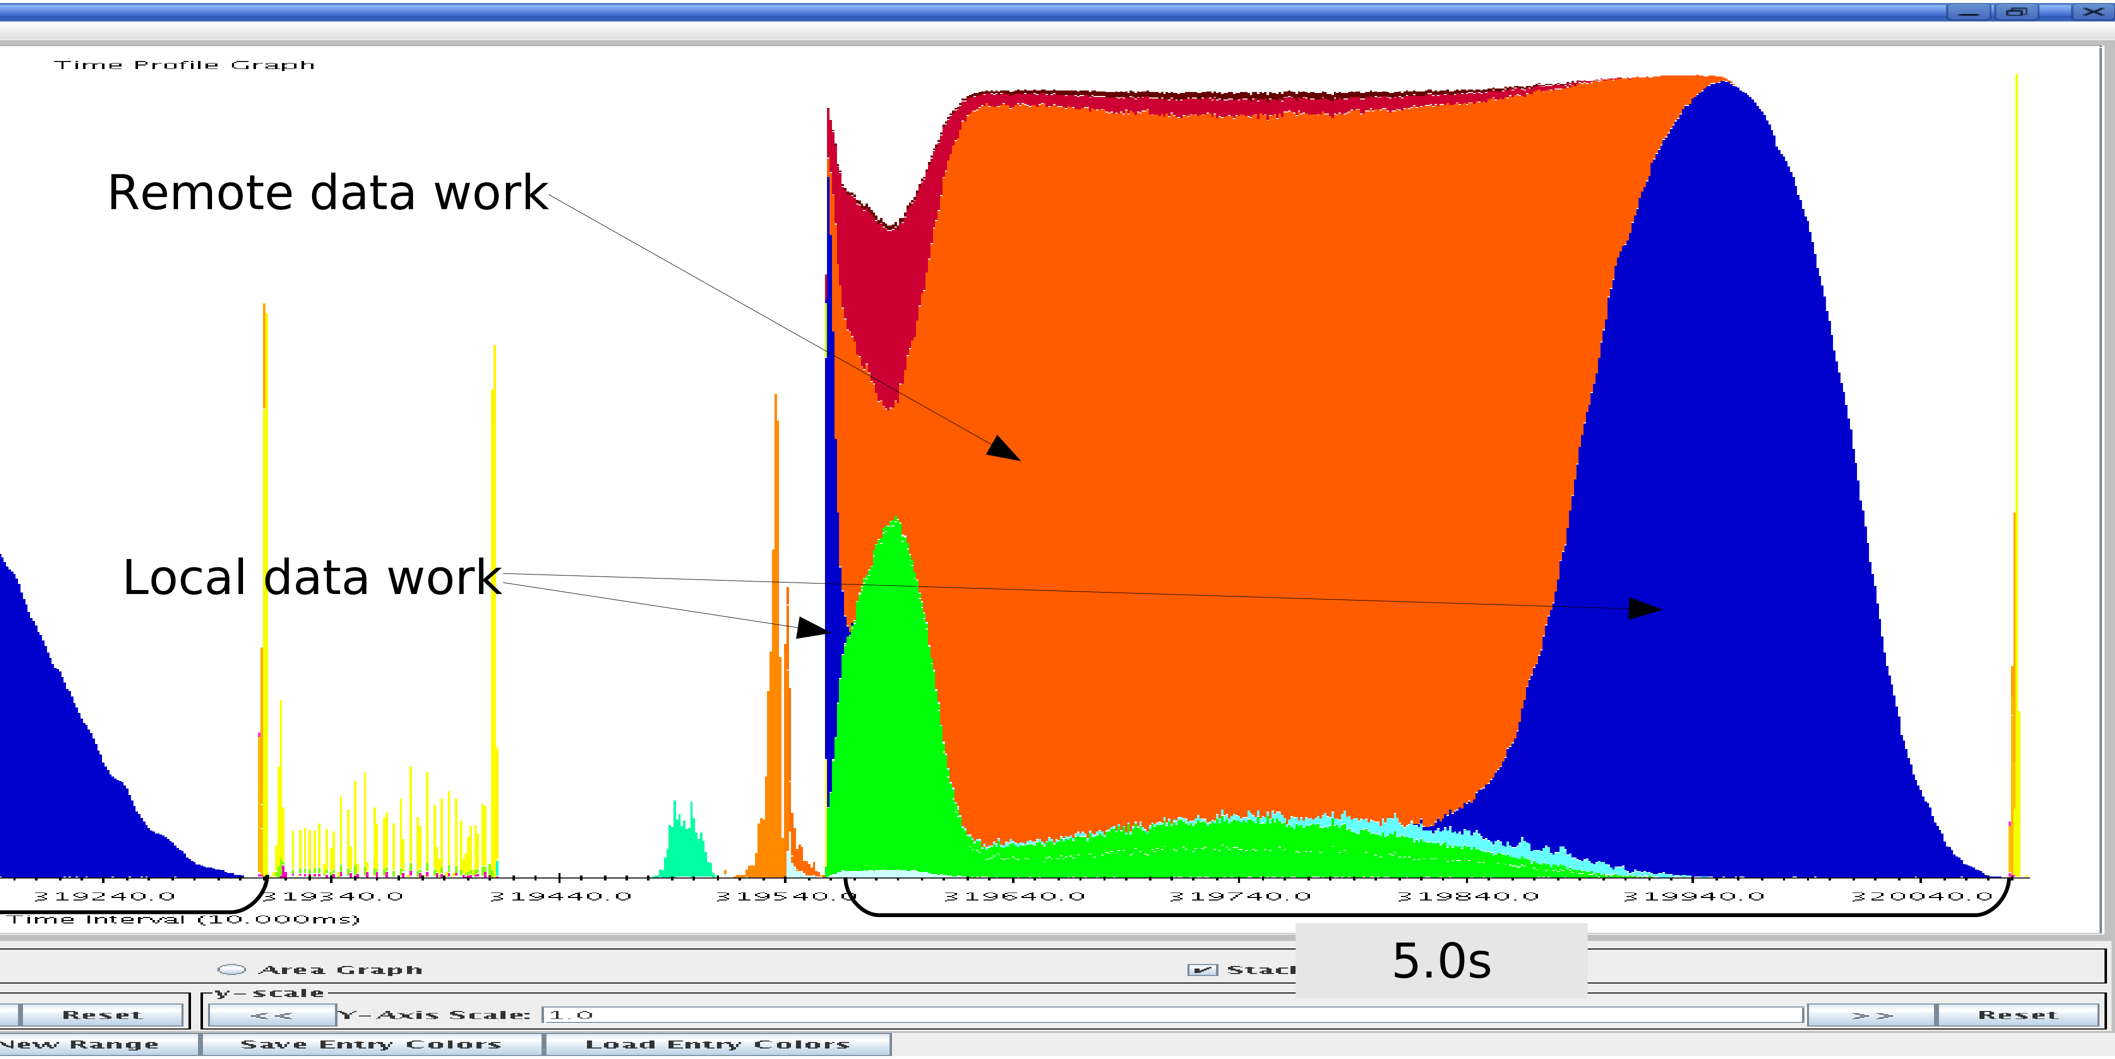

Load balancing with GreedyLB

Zoom In 5M on 1,024 BlueGene/L processors

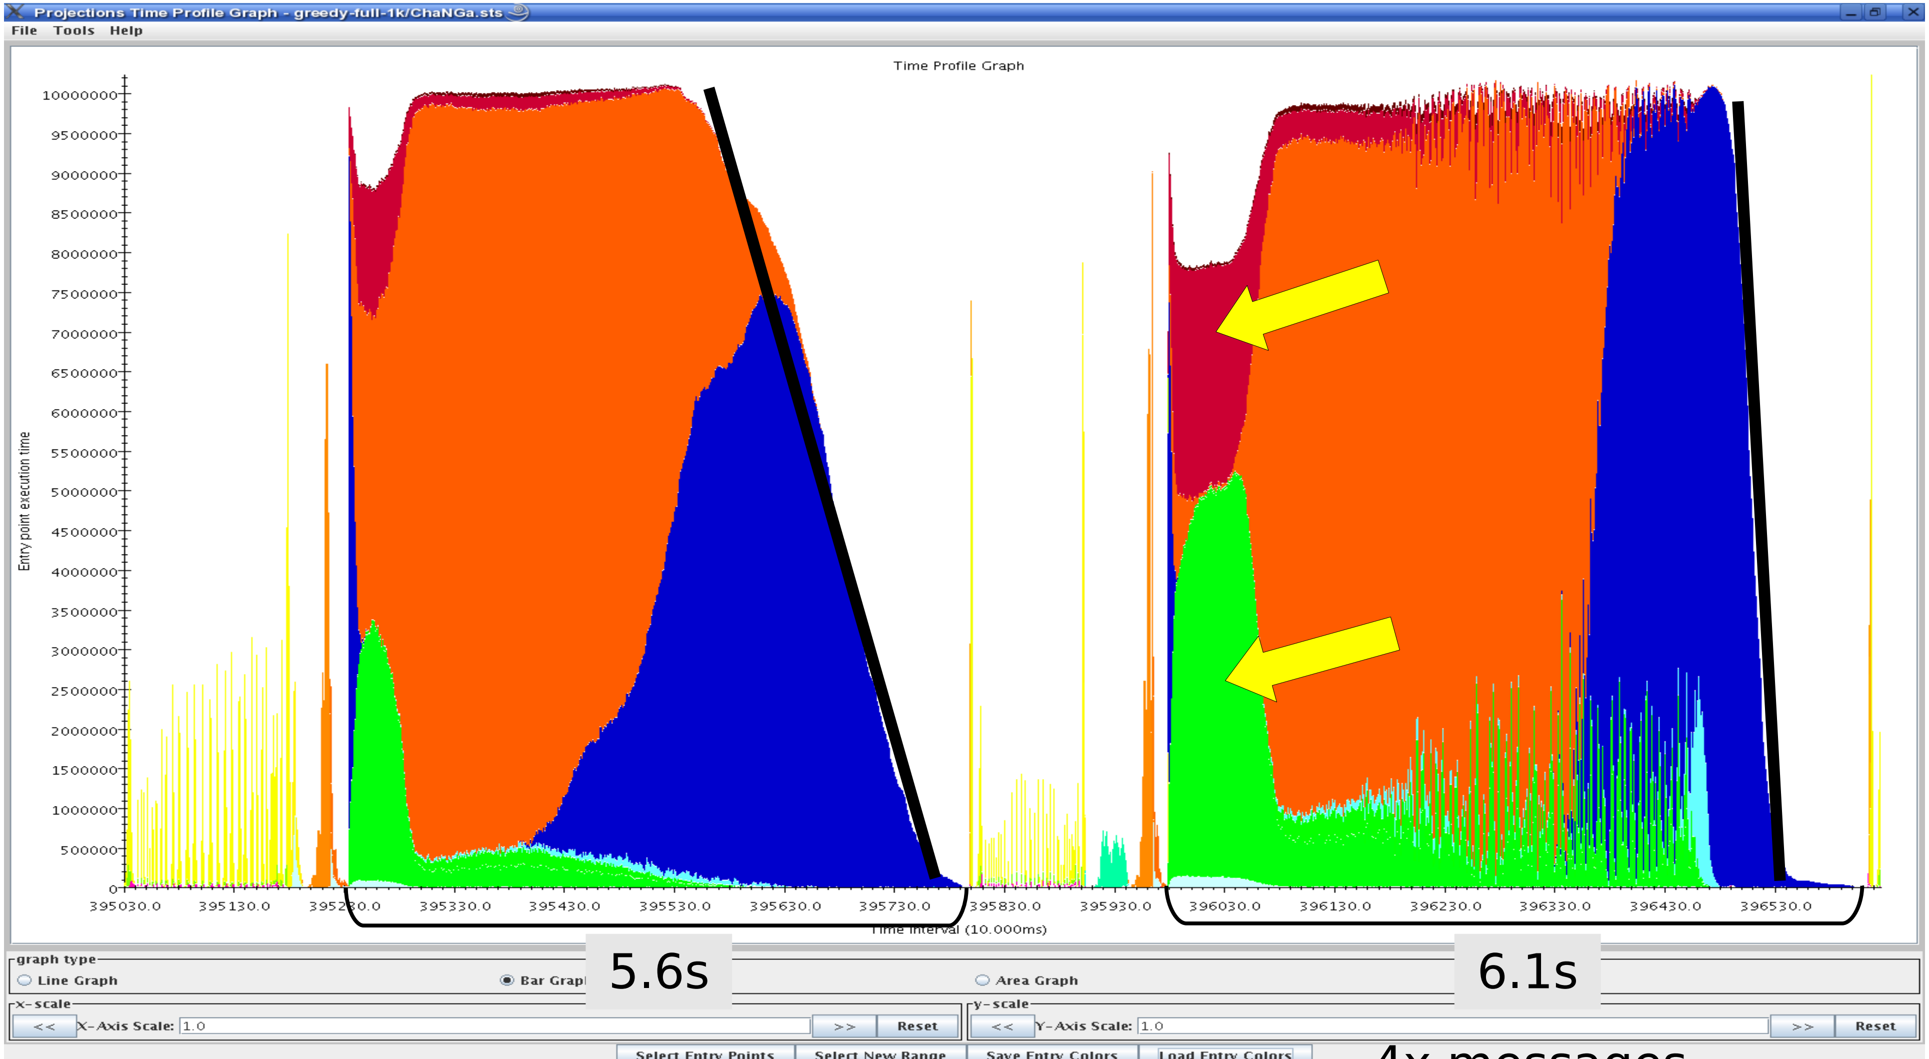

4x messages

Load balancing with OrbRefineLB

Zoom in 5M on 1,024 BlueGene/L processors

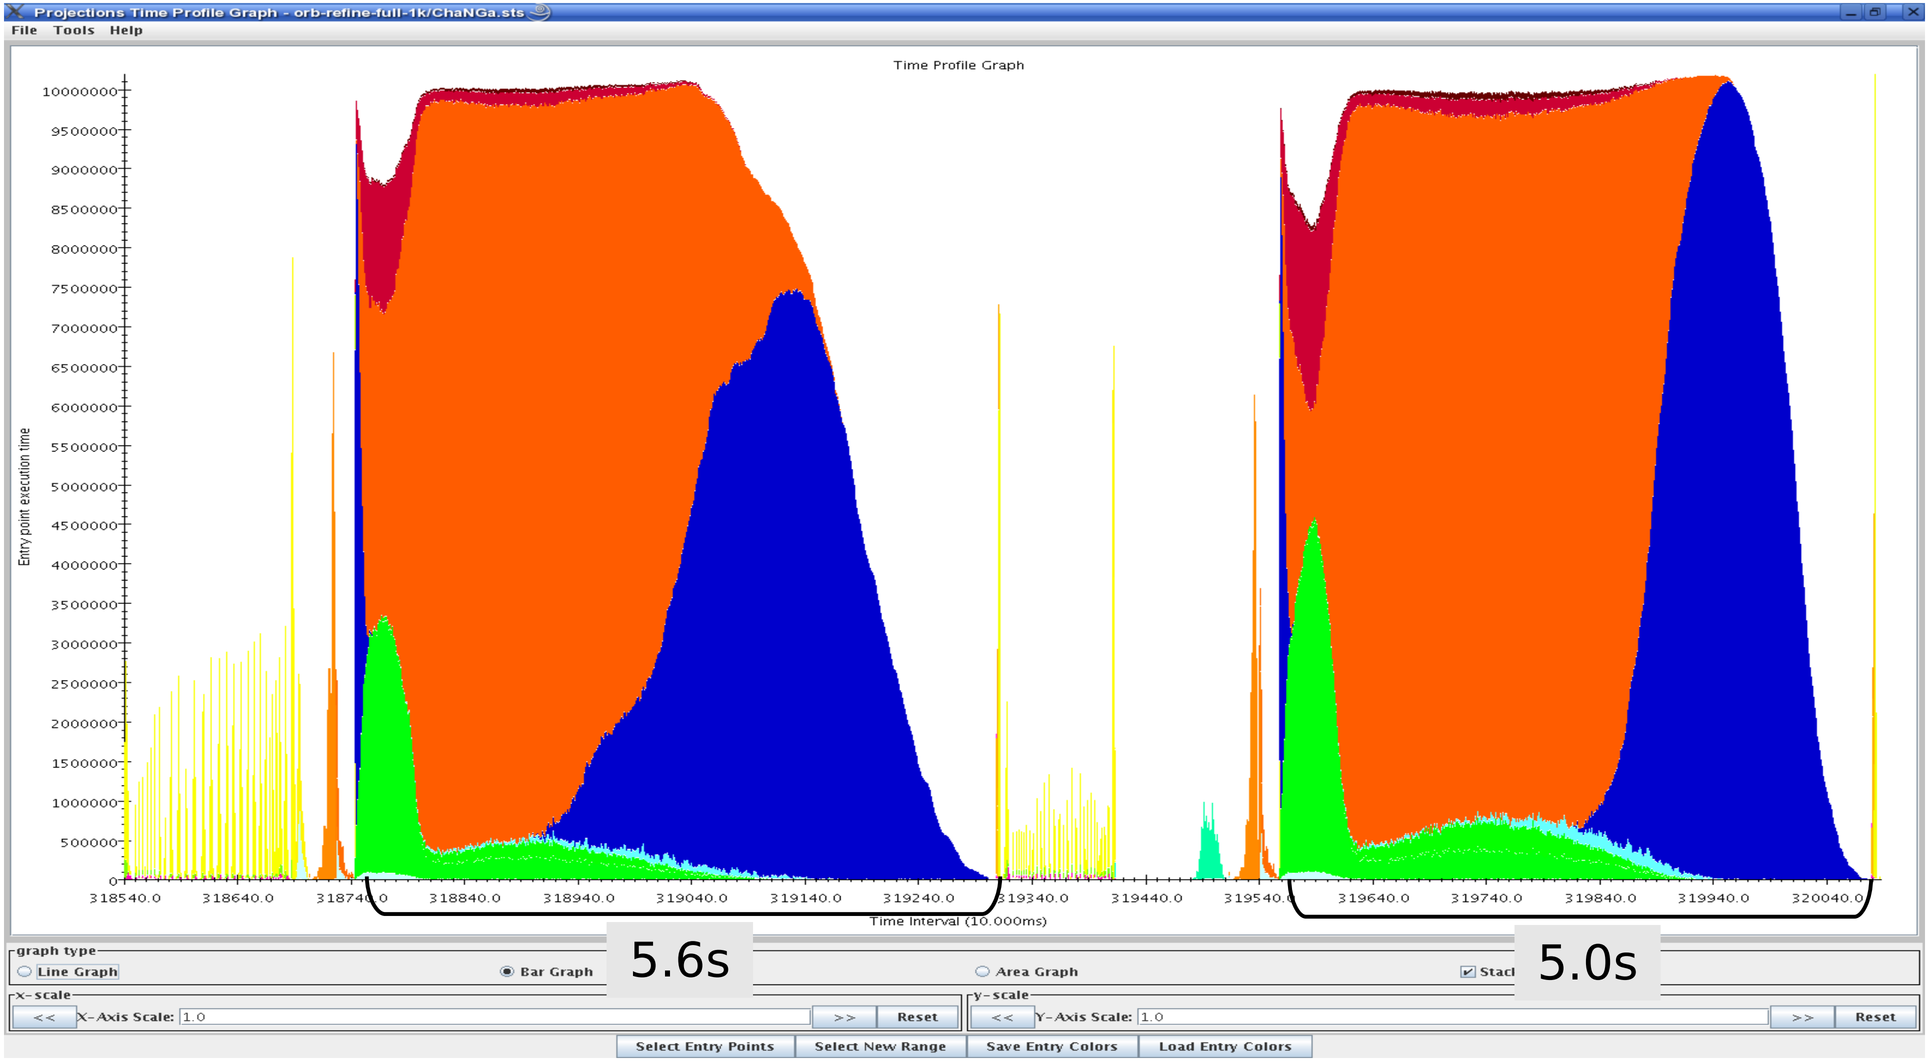

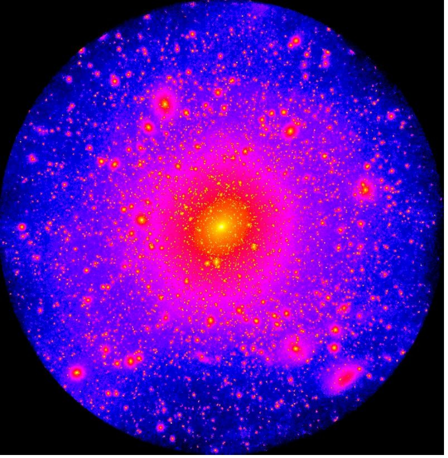

Timestepping Challenges

1/*m* particles need *m* times more force evaluations

Naively, simulation cost scales as N^(4/3)ln(N)

This is a problem when N ~ 1e9 or greater

If each particle an individual timestep scaling reduces to N (ln(N))^2

A difficult dynamic load balancing problem

Multistep Loadbalancer

Use Charm++ measurement based load balancer

Modification: provide LB database with information about timestepping.

"Large timestep": balance based on previous large step

"Small step" balance based on previous small step

Maintains principle of persistence

Results on 3 rung example

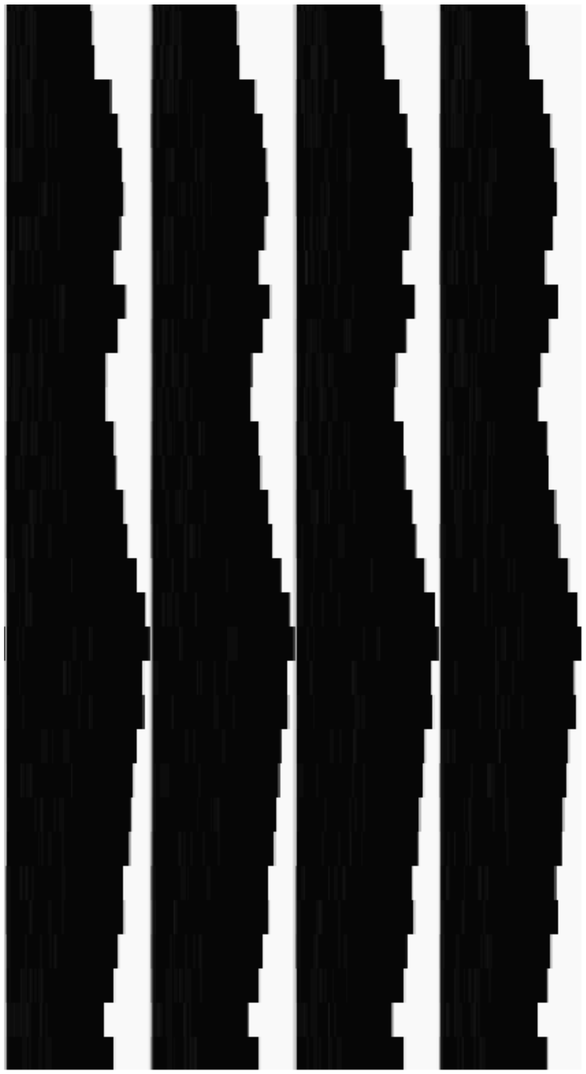

613s

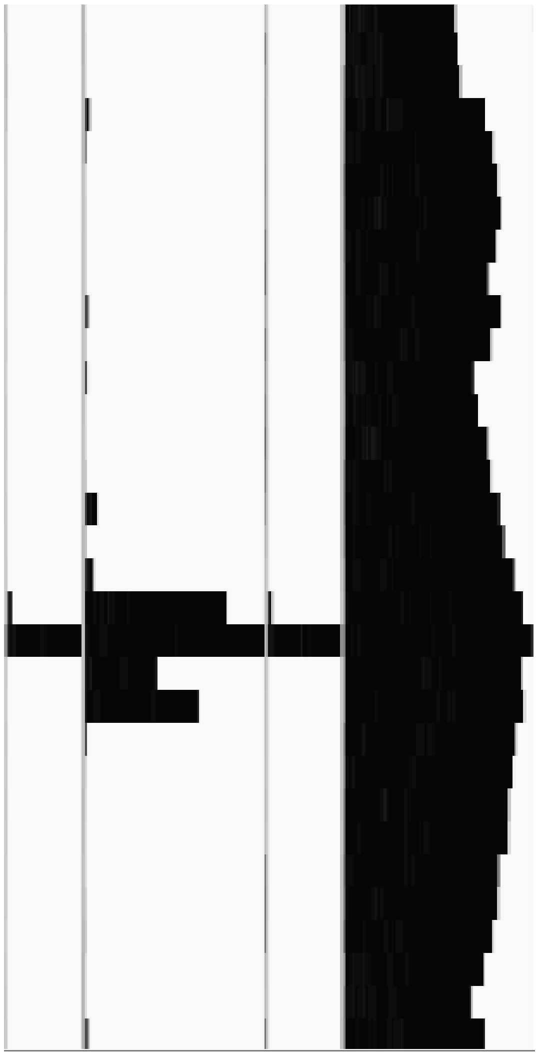

429s

228s

Summary

- Cosmological simulations provide a challenges to parallel implementations
  - Non-local data dependencies
  - Hierarchical in space and time
- ChaNGa has been successful in addressing this challenges using Charm++ features
  - Message priorities
  - New load balancers

Future

New Physics

Smooth particle hydrodynamics

Better gravity algorithms

Fast multipole method

New domain decomposition/load balancing strategies

Heterogeneous Computing

- Multicore
- GPUs

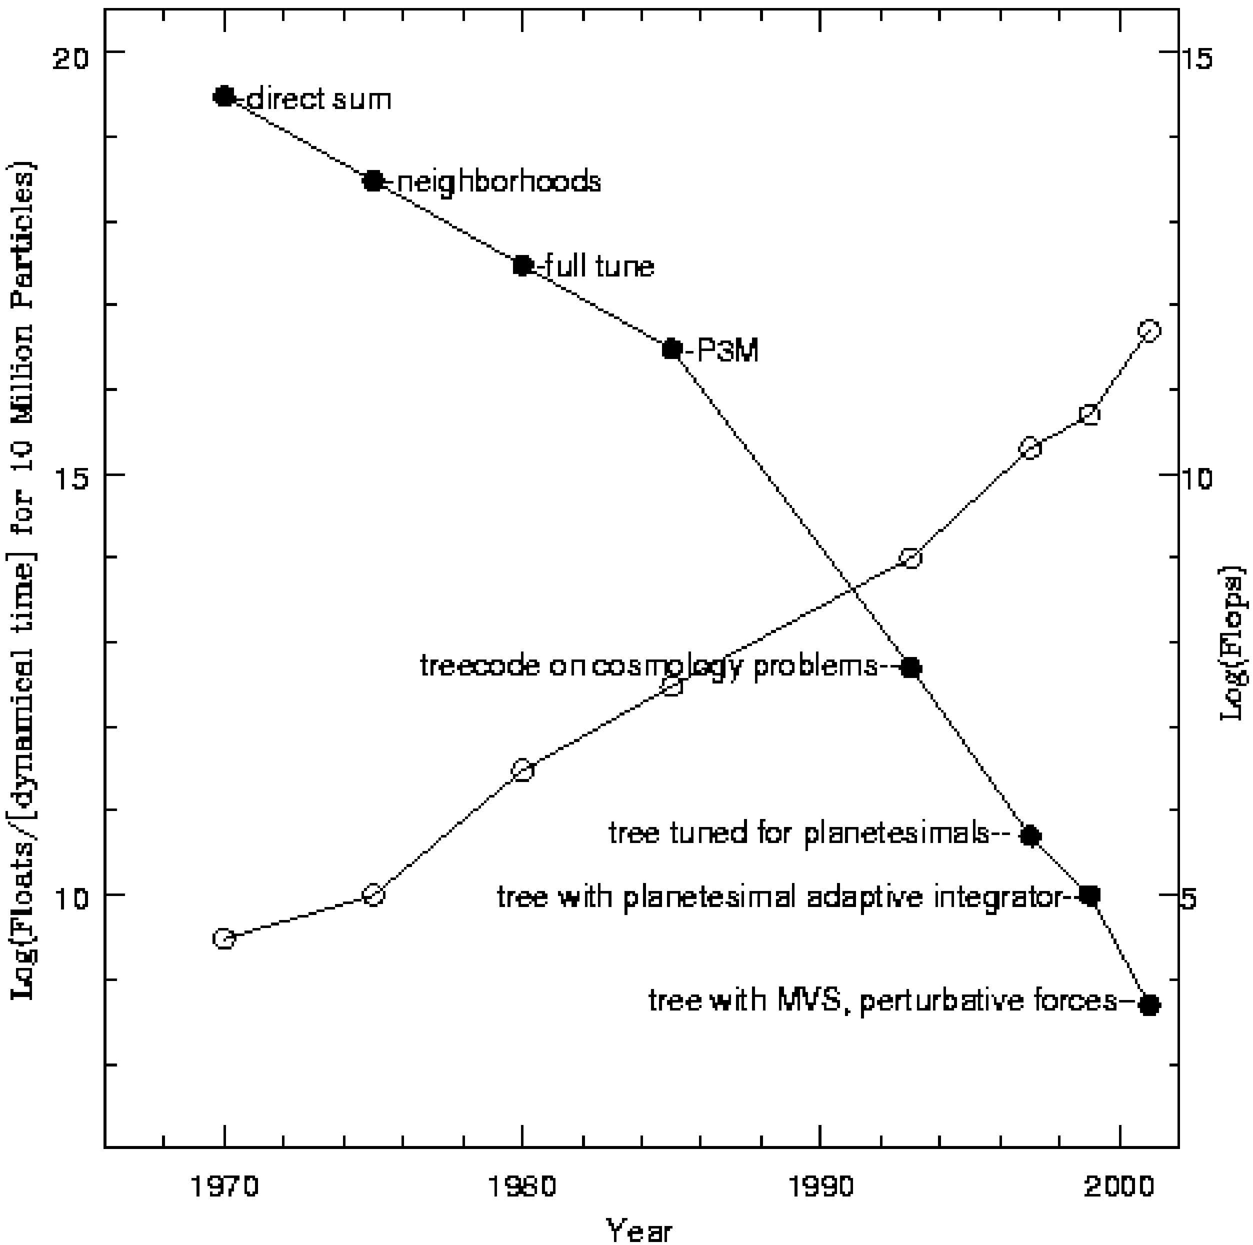

-direct sum
-neighborhoods
-full tune
-P3M
treecode on cosmology problems--
tree tuned for planetesimals--
tree with planetesimal adaptive integrator--
tree with MVS, perturbative forces--
Log(Floats/[dynamical time] for 10 Million Particles)
Log(Flops)
Year
20
15
10
1970
1980
1990
2000

Computing Challenge Summary

The Universe is big => we will always be pushing for more resources

New algorithm efforts will be made to make efficient use of the resources we have

Efforts made to abstract away from machine details

Parallelization efforts need to depend on more automated processes.