ChaNGa

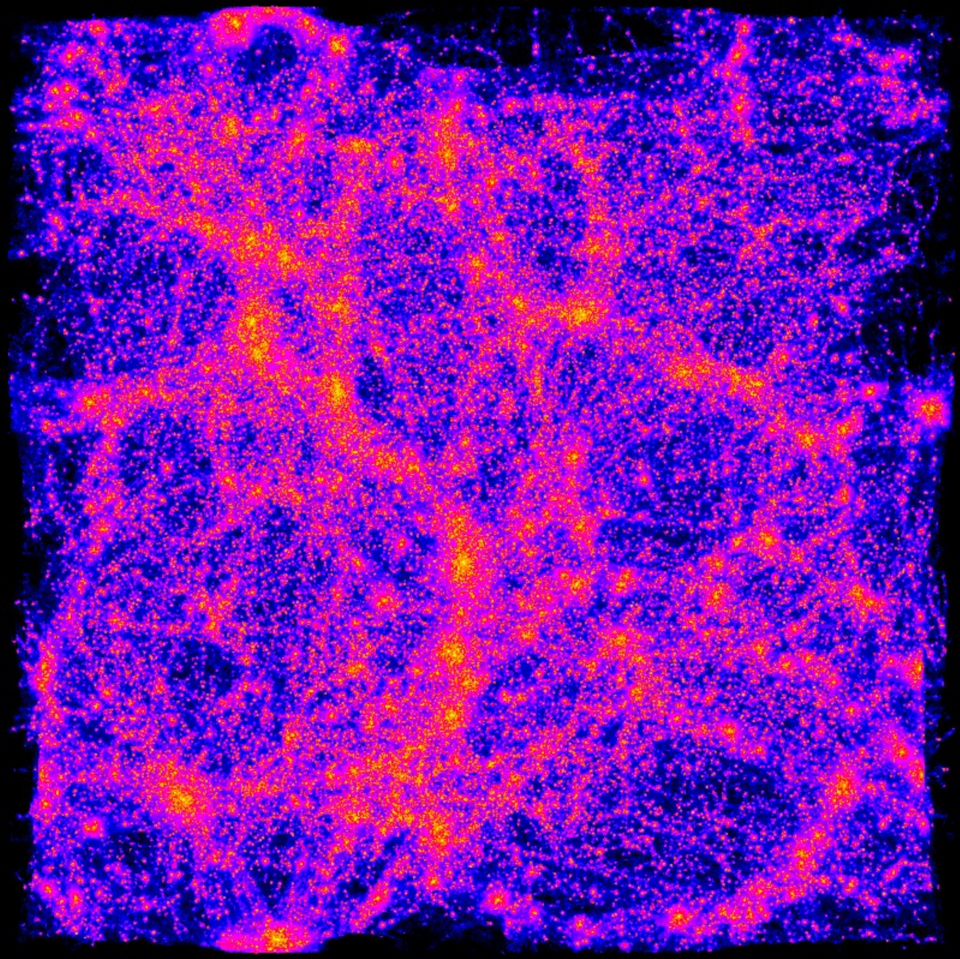

CHArm N-body GrAvity

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
  - Needed Simulations
  - Technology Challenges

Cosmology: How does this ...

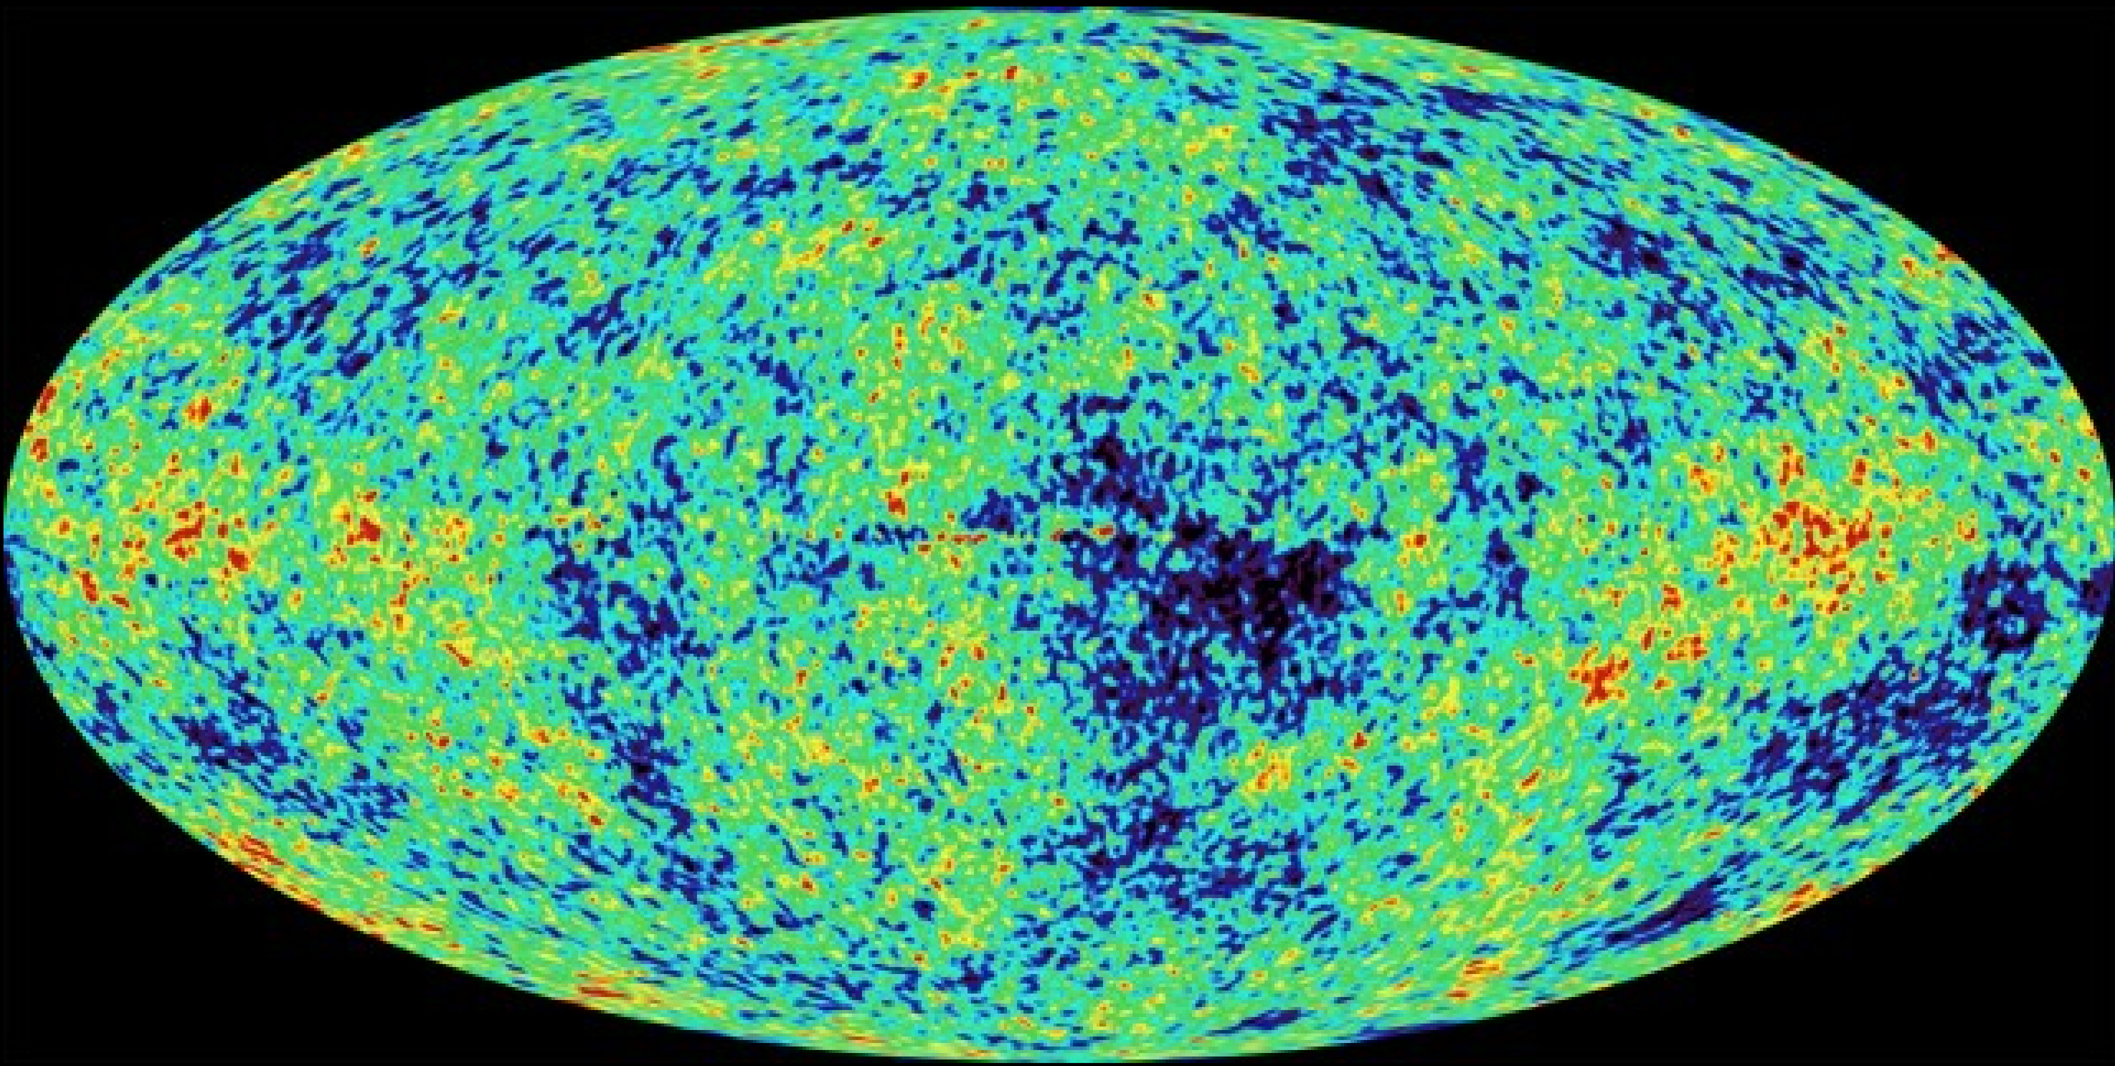

Image courtesy NASA/WMAP

... 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

1. Simulate 100 Mpc volume at 10-100 kpc resolution
2. Pick candidate galaxies for further study
3. Resimulate galaxies with same large scale structure but with higher resolution, and lower resolution in the rest of the computational volume.
4. At higher resolutions, include gas physics and star formation.

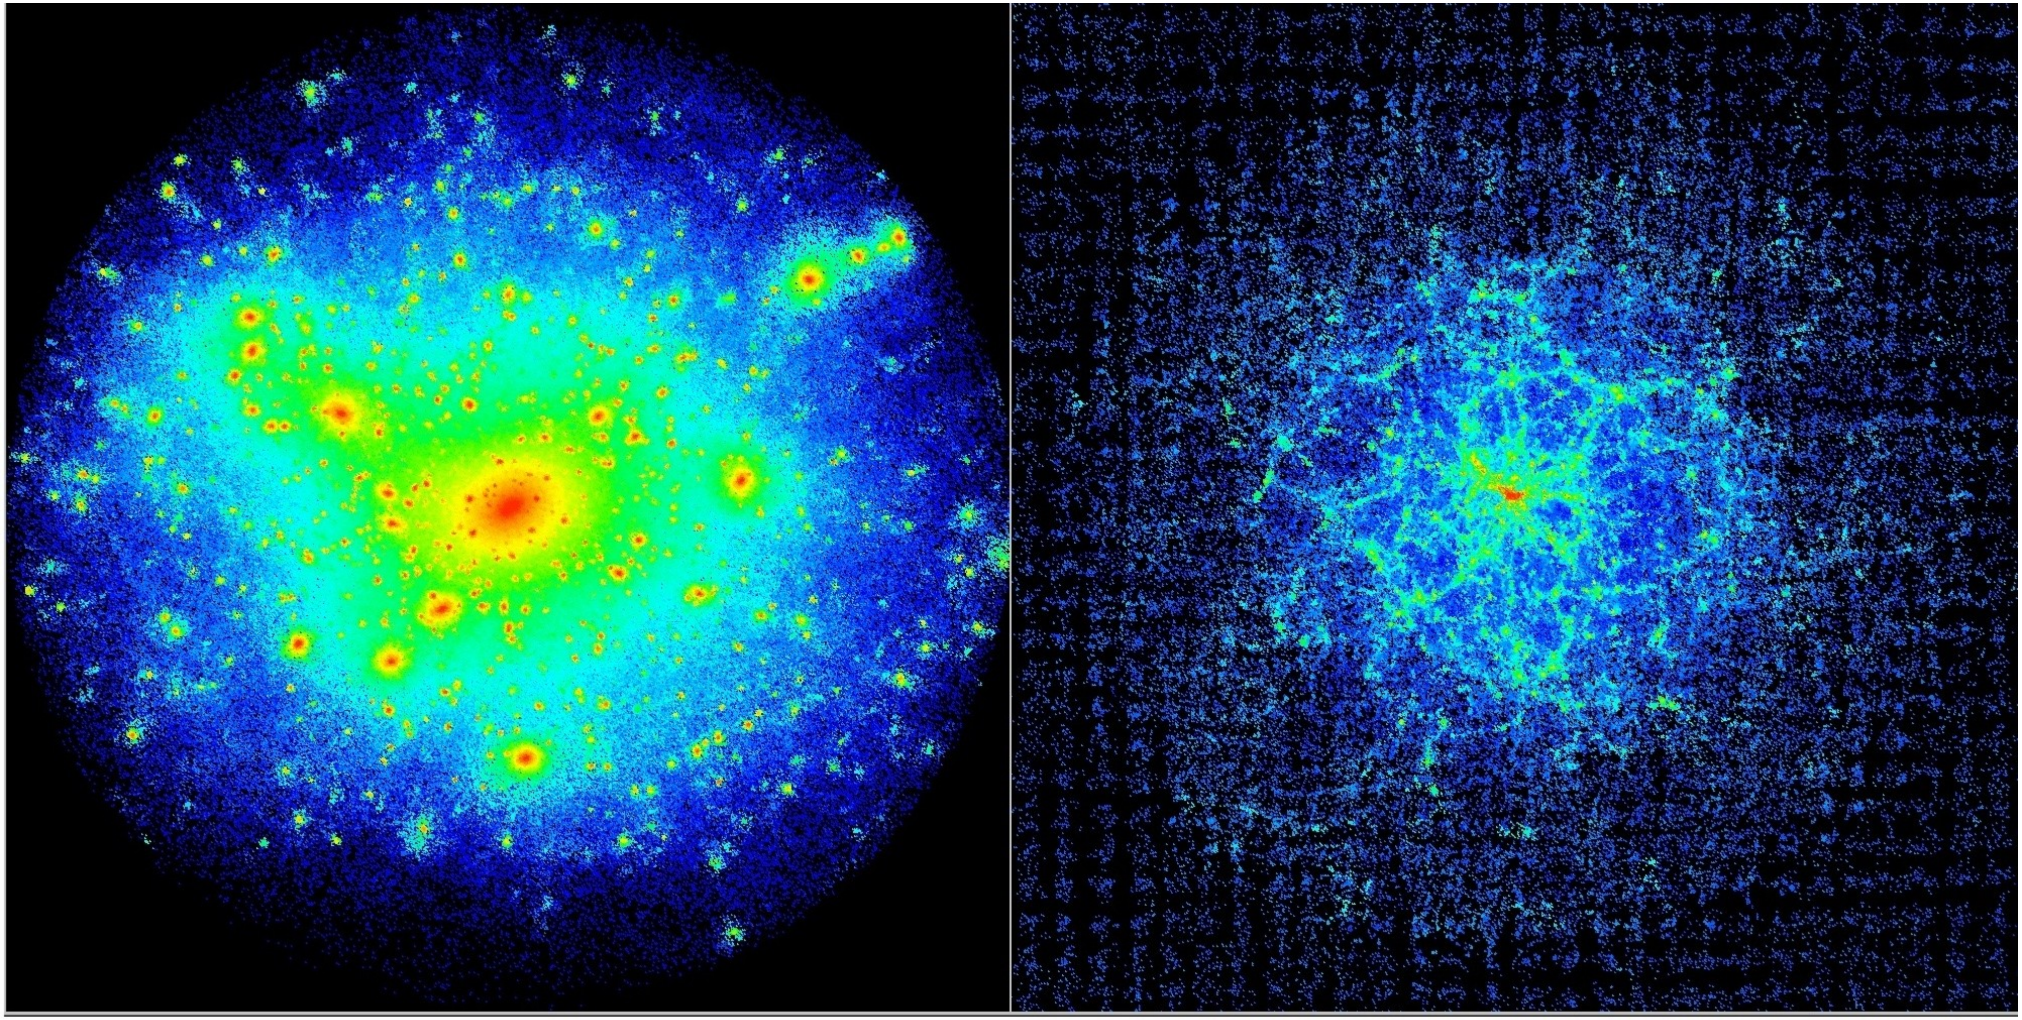

<-----------6 Mpc Sphere-----------> <-----------1000 Mpc Box----------->

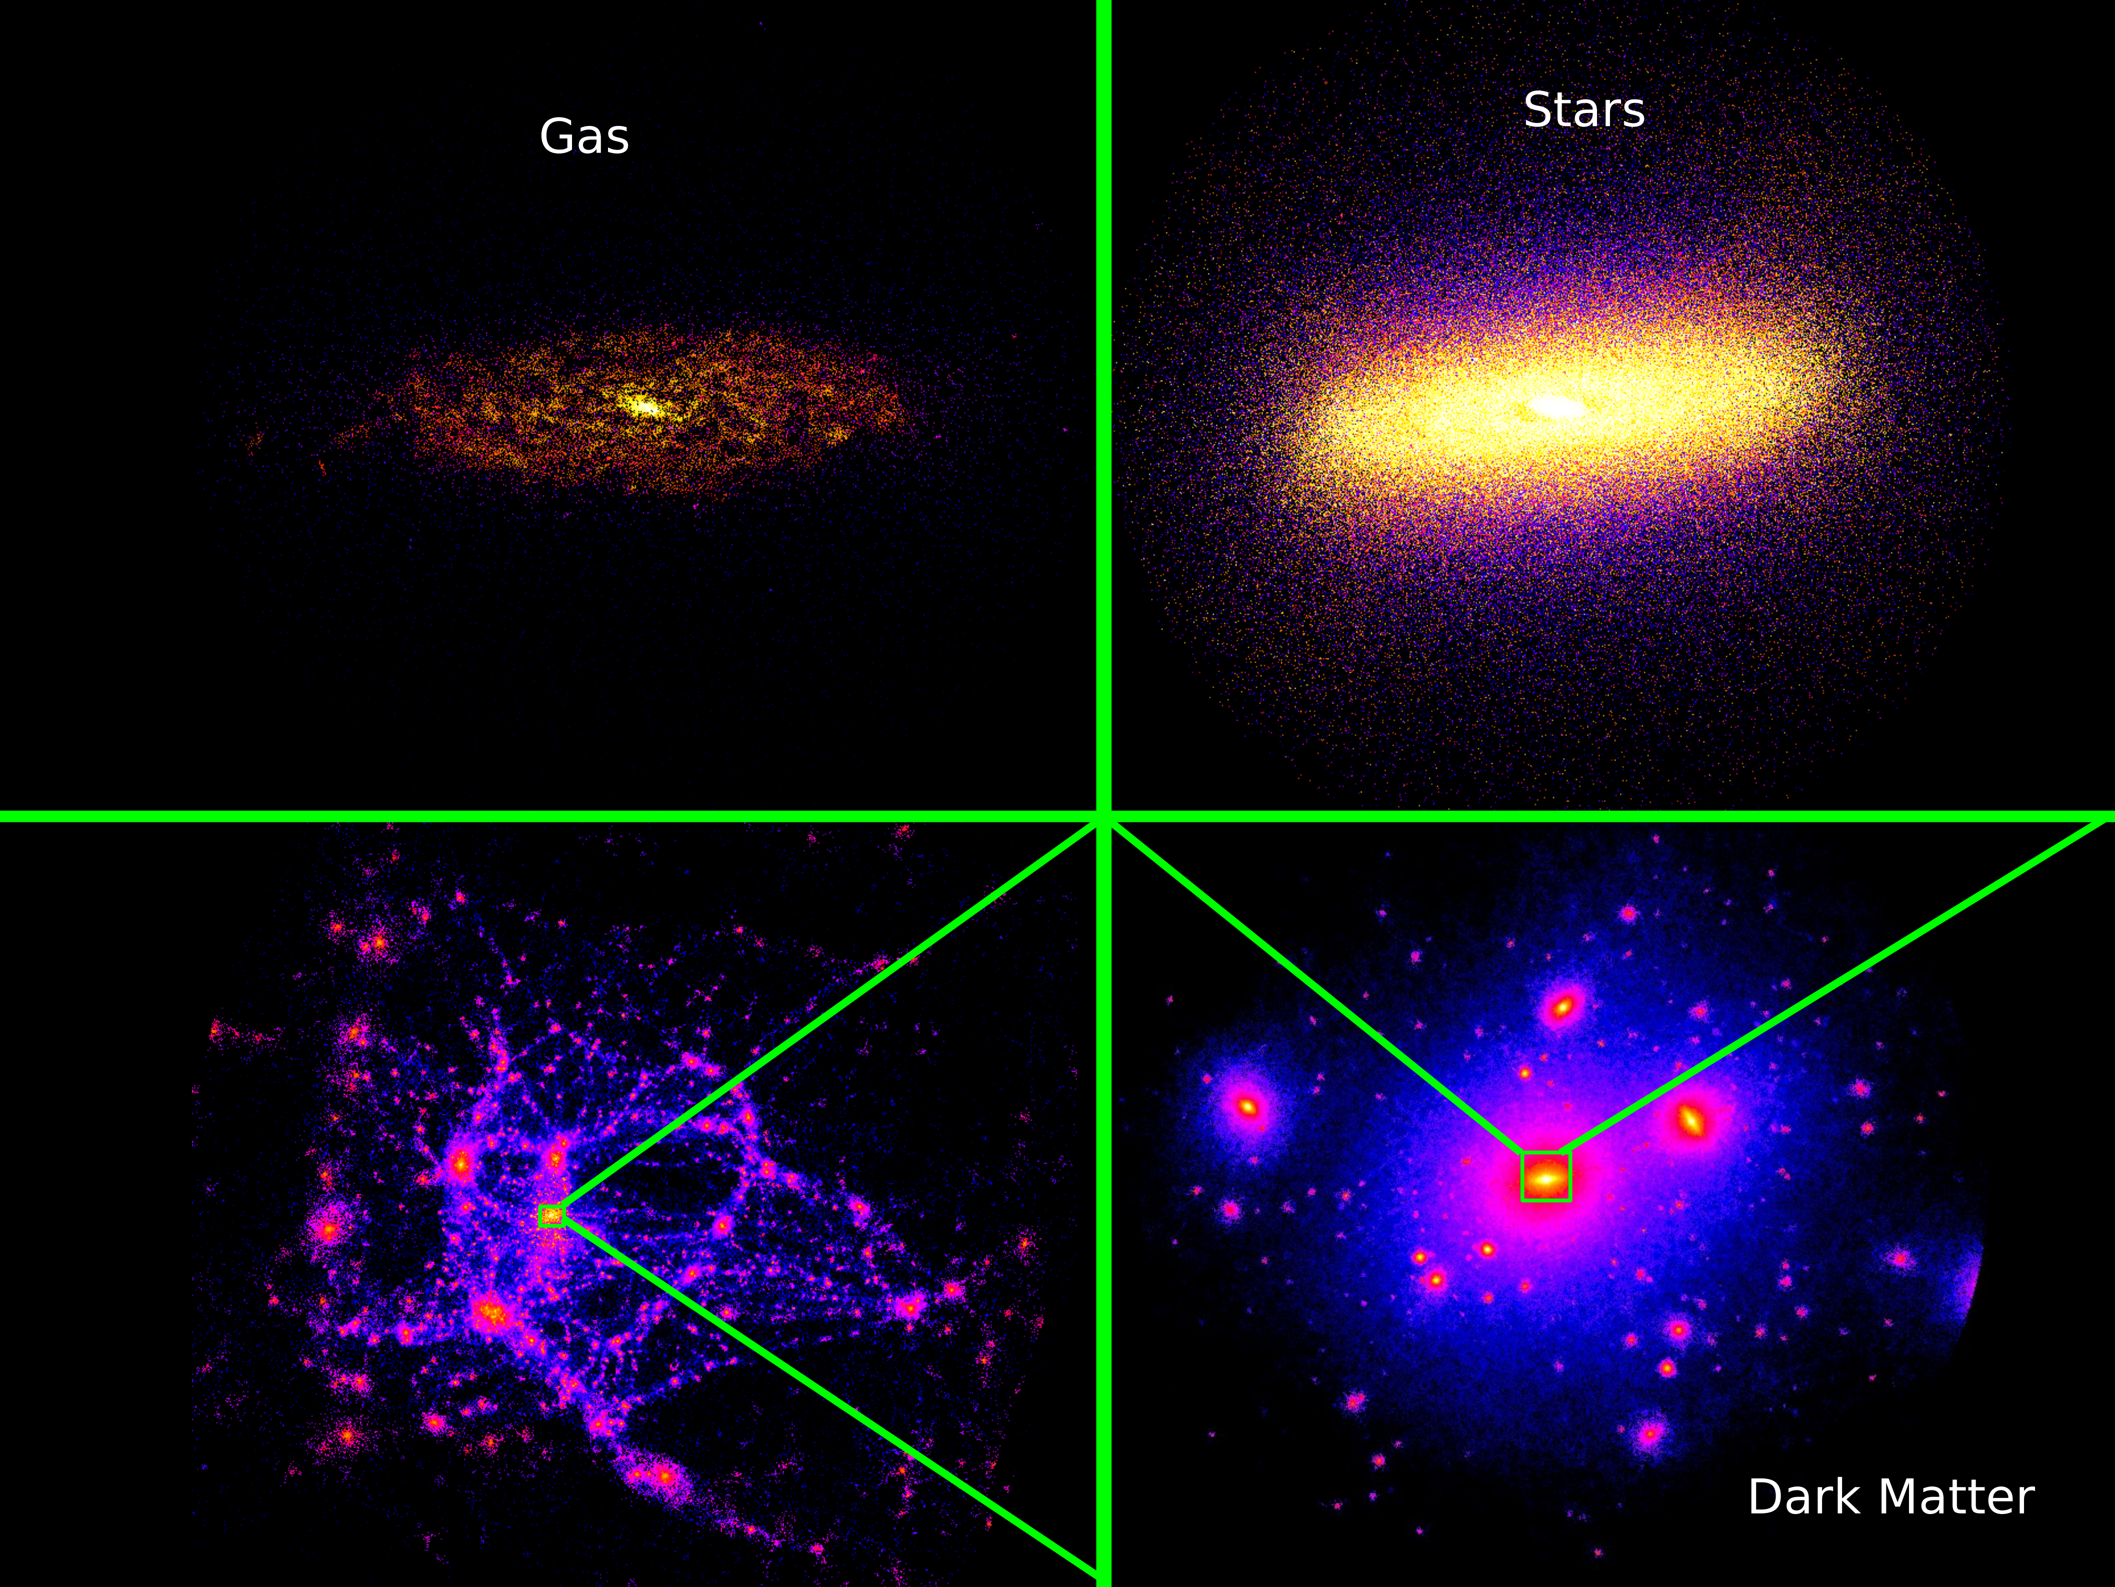
Gas
Stars
Dark Matter

Types of simulations

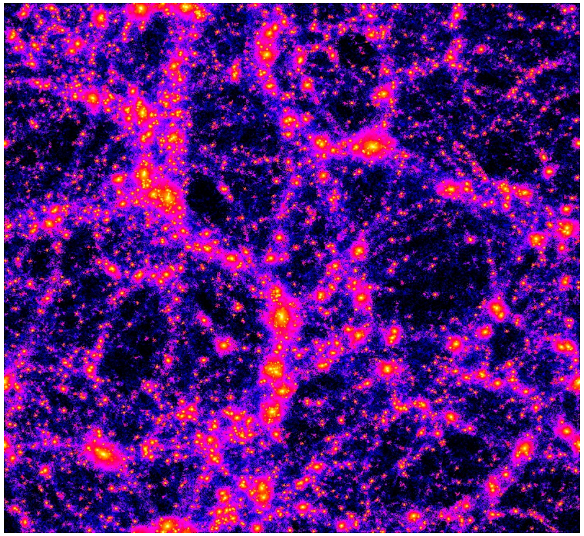

“Uniform” Volume

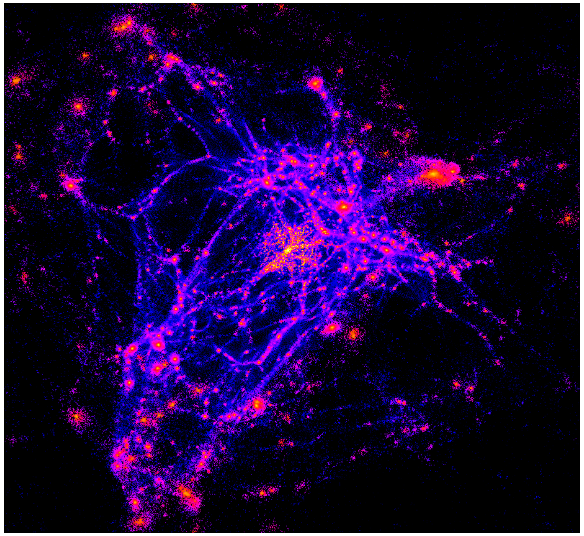

Zoom In

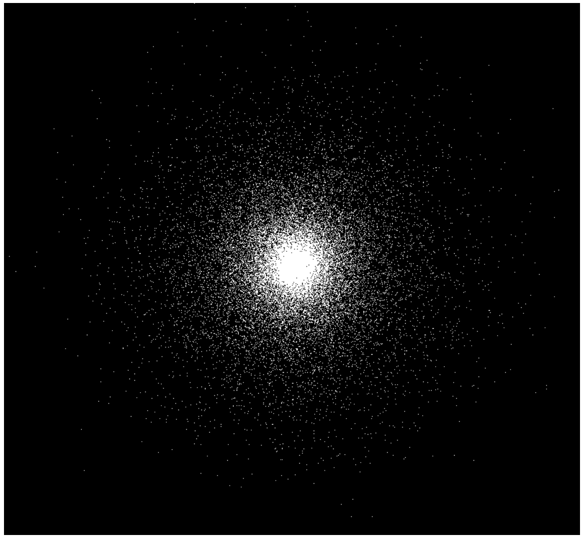

Star Cluster

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
  - Hierarchal information needs to go across processor domains

The existing code:

Gasoline

- Multi-Platform
- Massively Parallel (100s; 1000s on large sims)
- Treecode with periodic boundary conditions
- Multi-stepping (but bad load balancing)
- Hydrodynamics (via SPH) with radiative cooling
- UV background
- Star Formation
- Supernovae feedback into thermal energy

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
- Visualization
- Built from the ground up on Charm++

Need for high multipole order

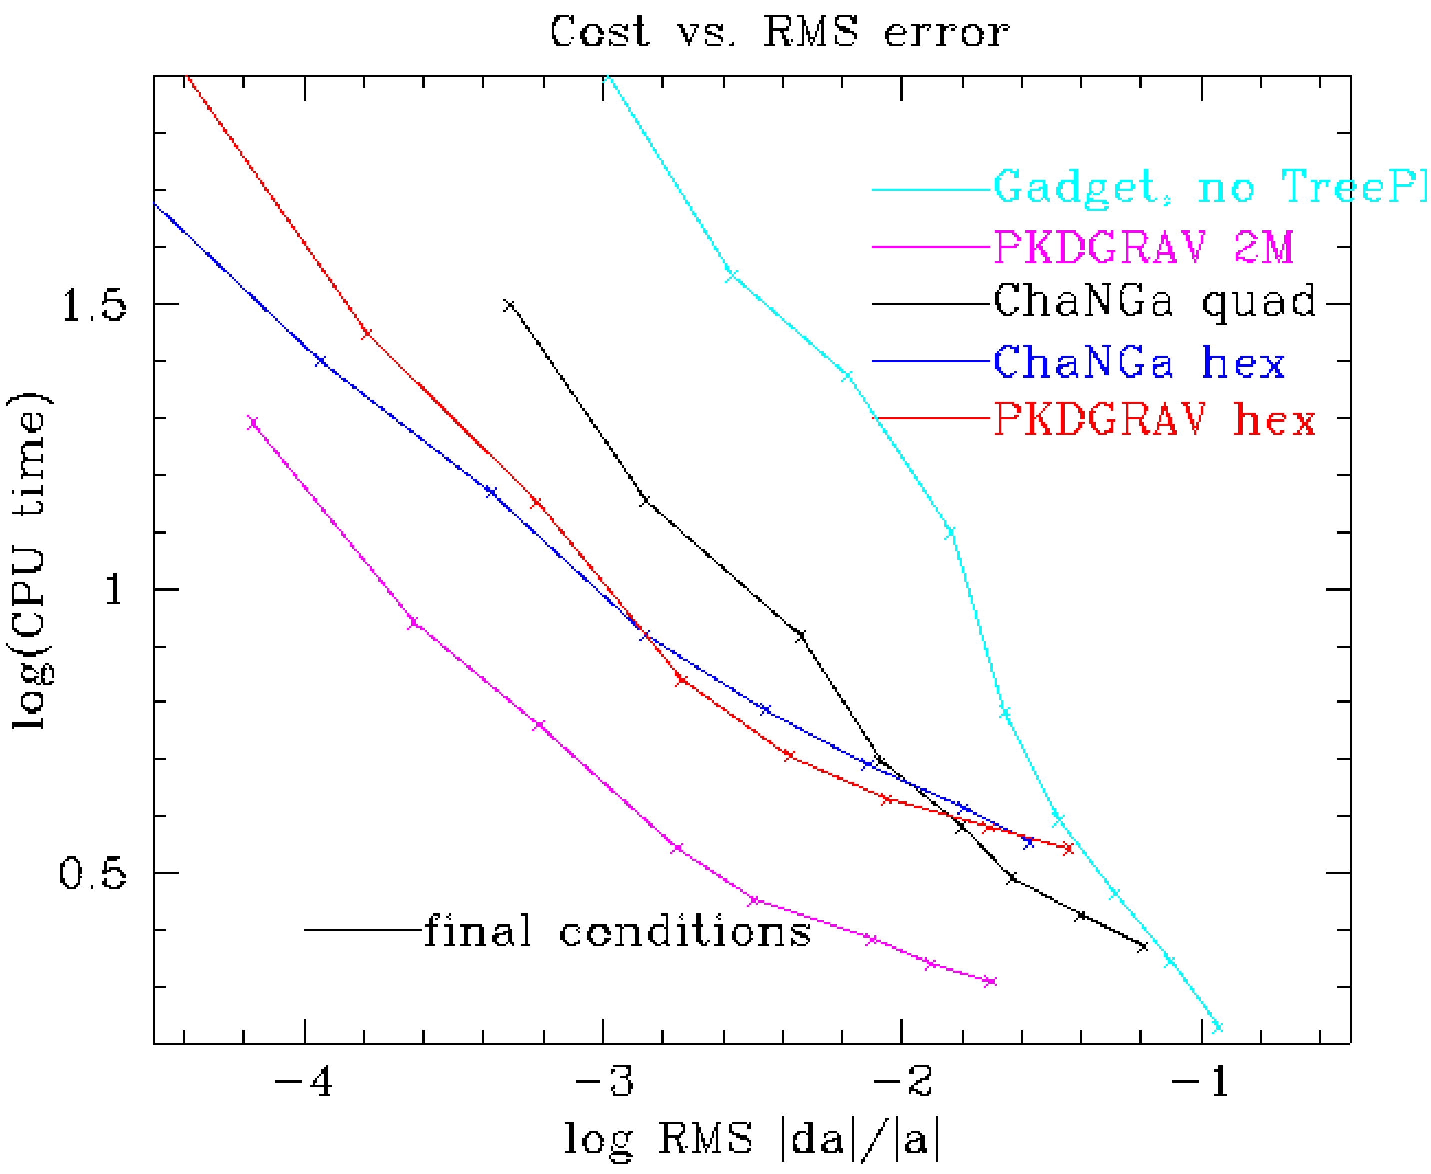

Space decomposition

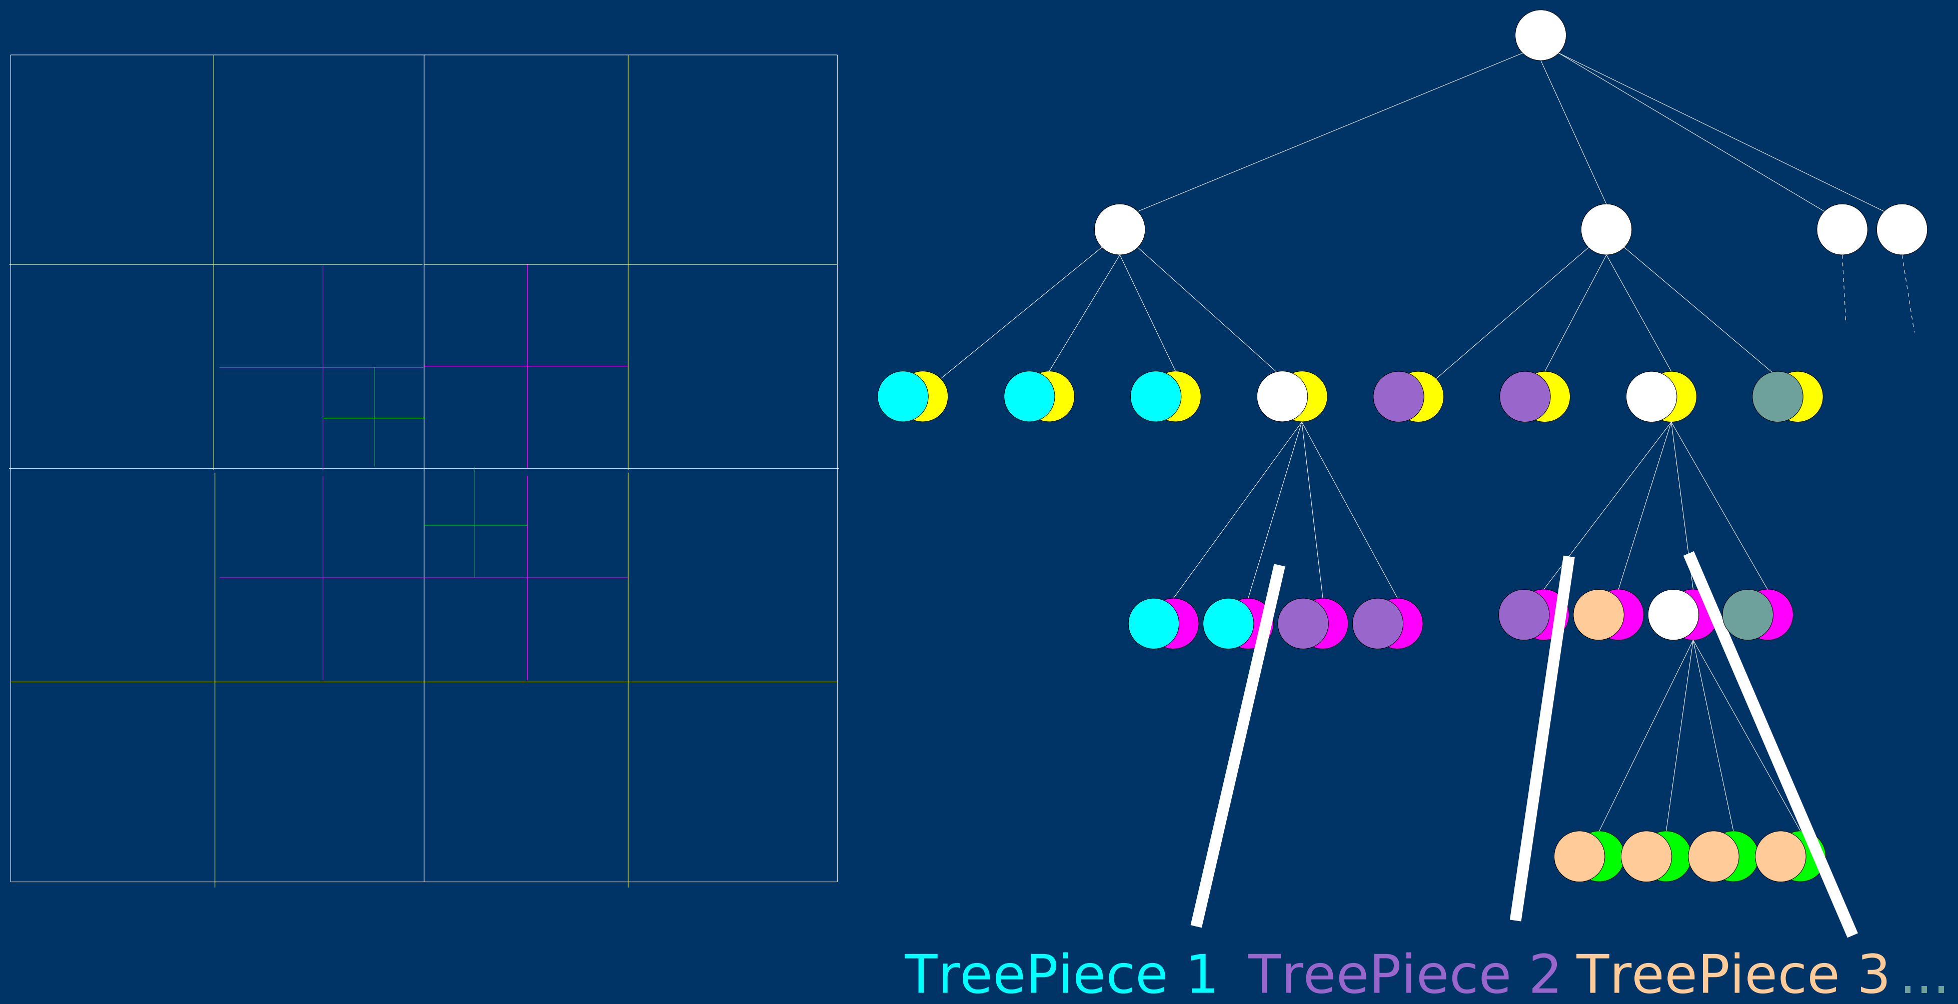

Basic algorithm ...

- 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

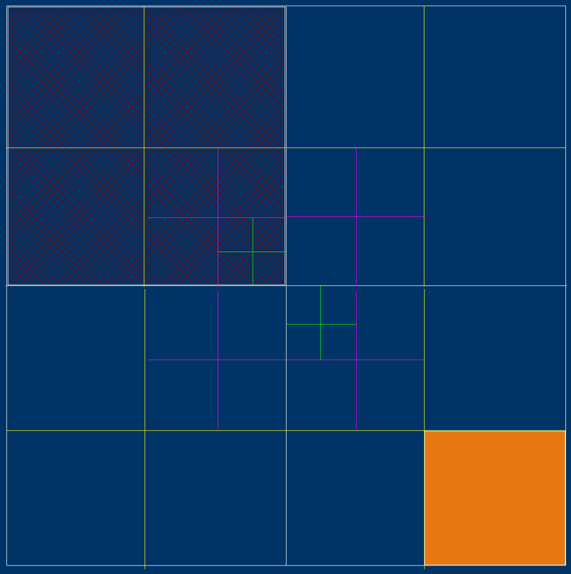

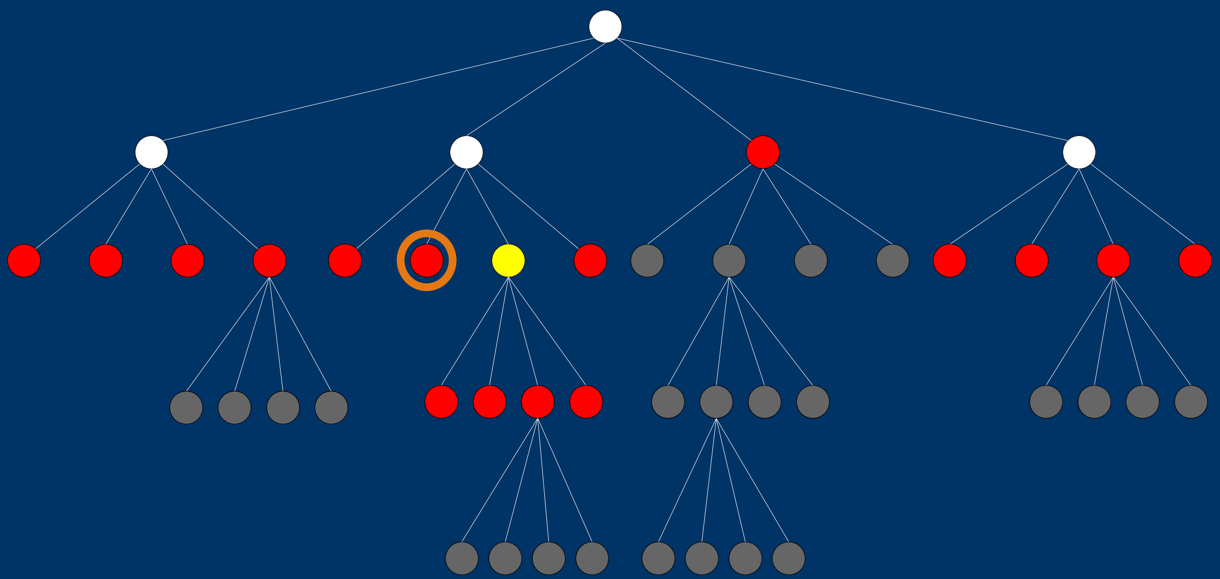

... in parallel

- Remote data
  - need to fetch from other processors
- Data reusage
  - same data needed by more than one particle

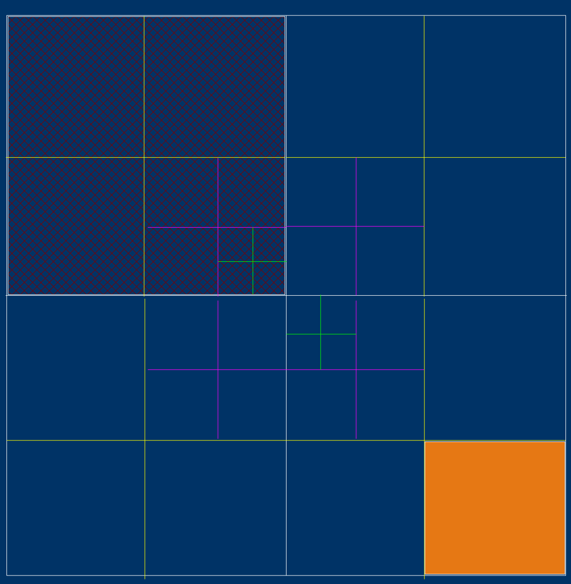

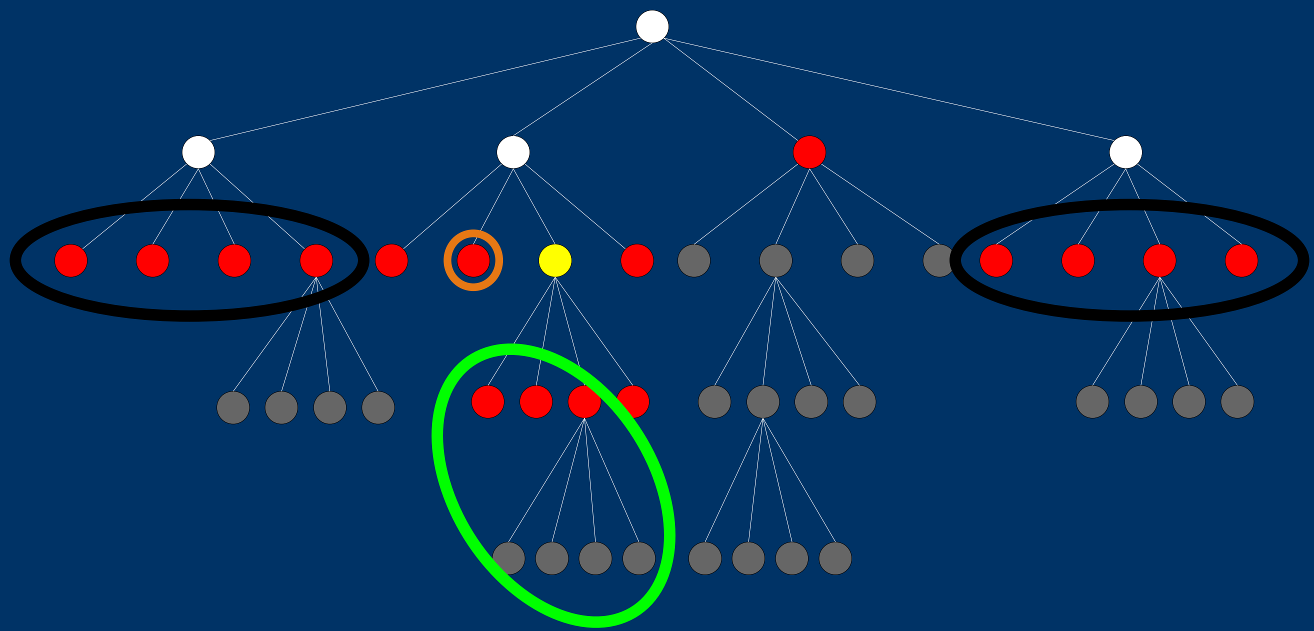

Overall algorithm

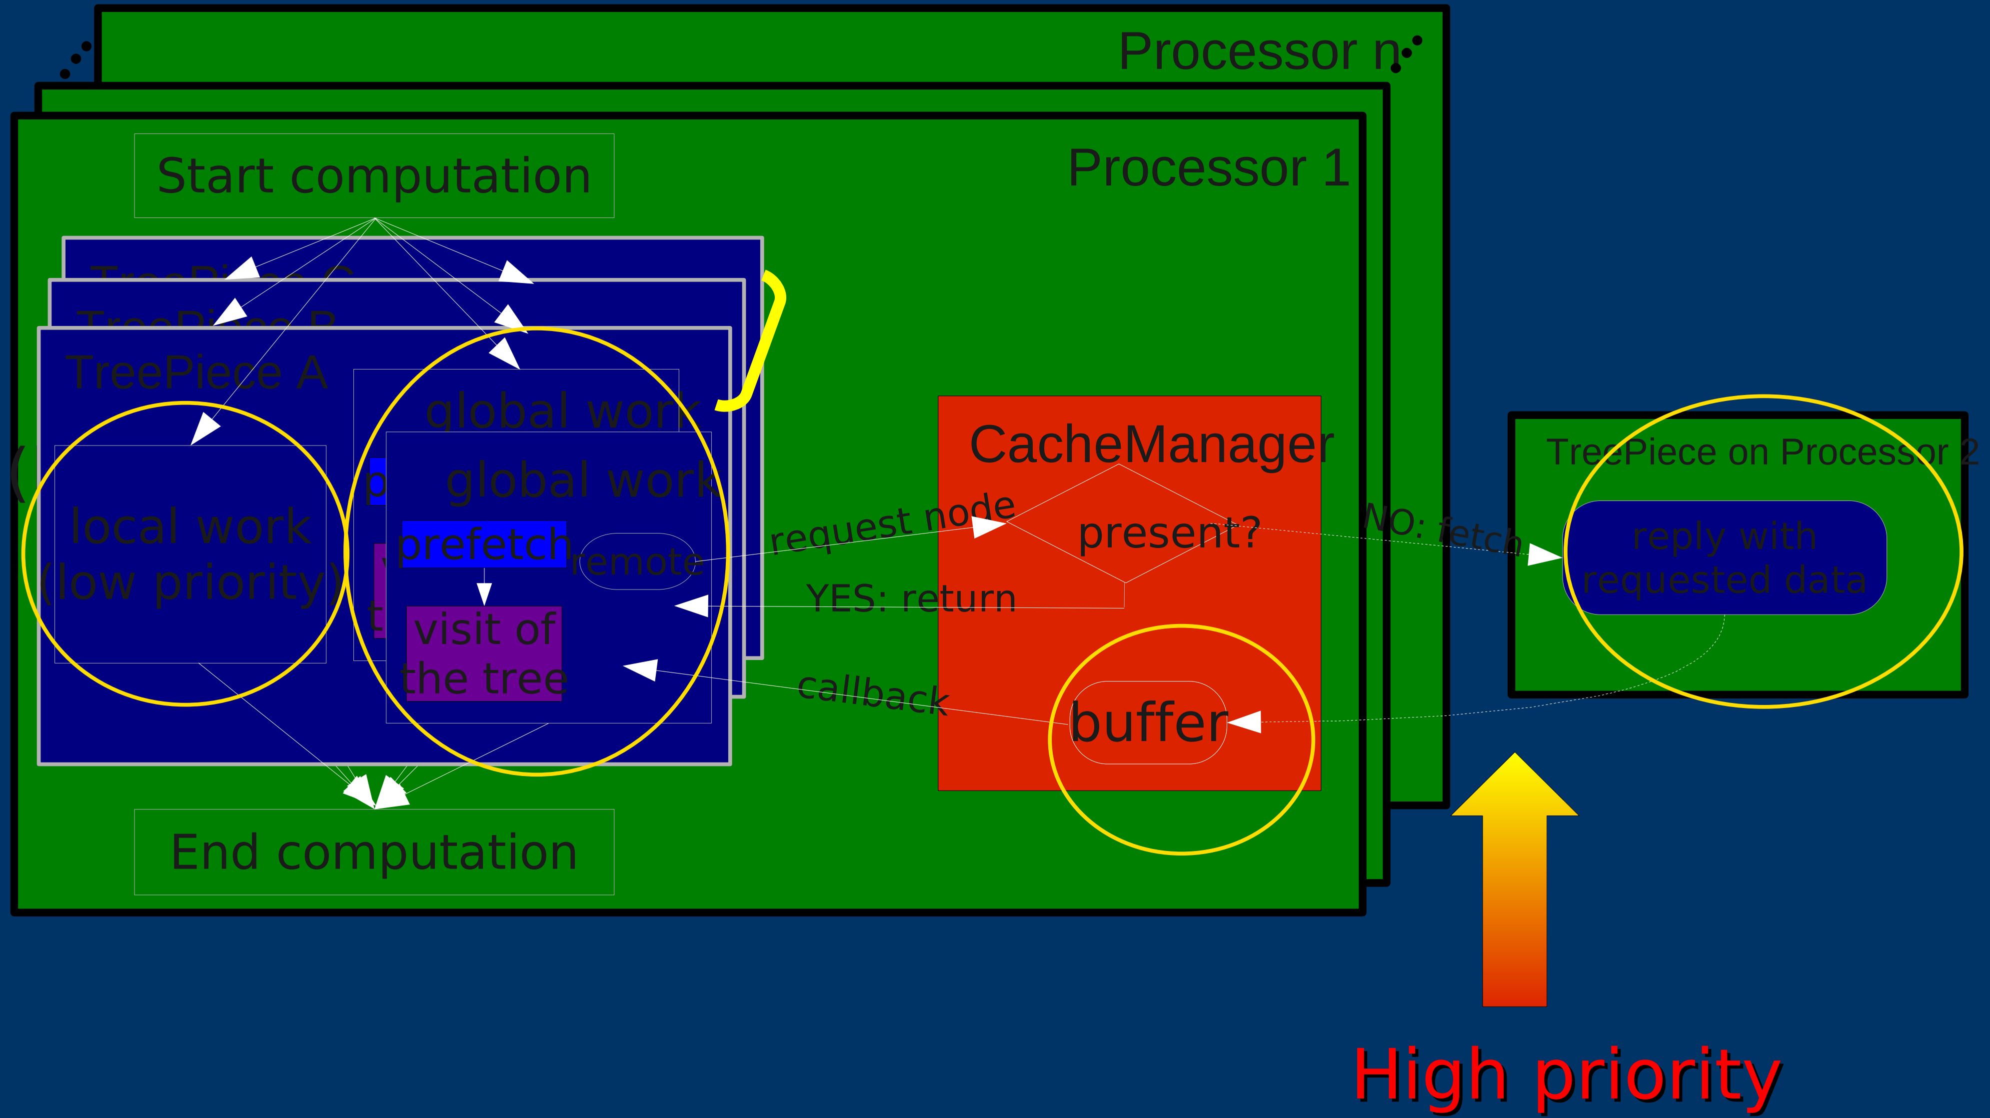

Scaling: comparison

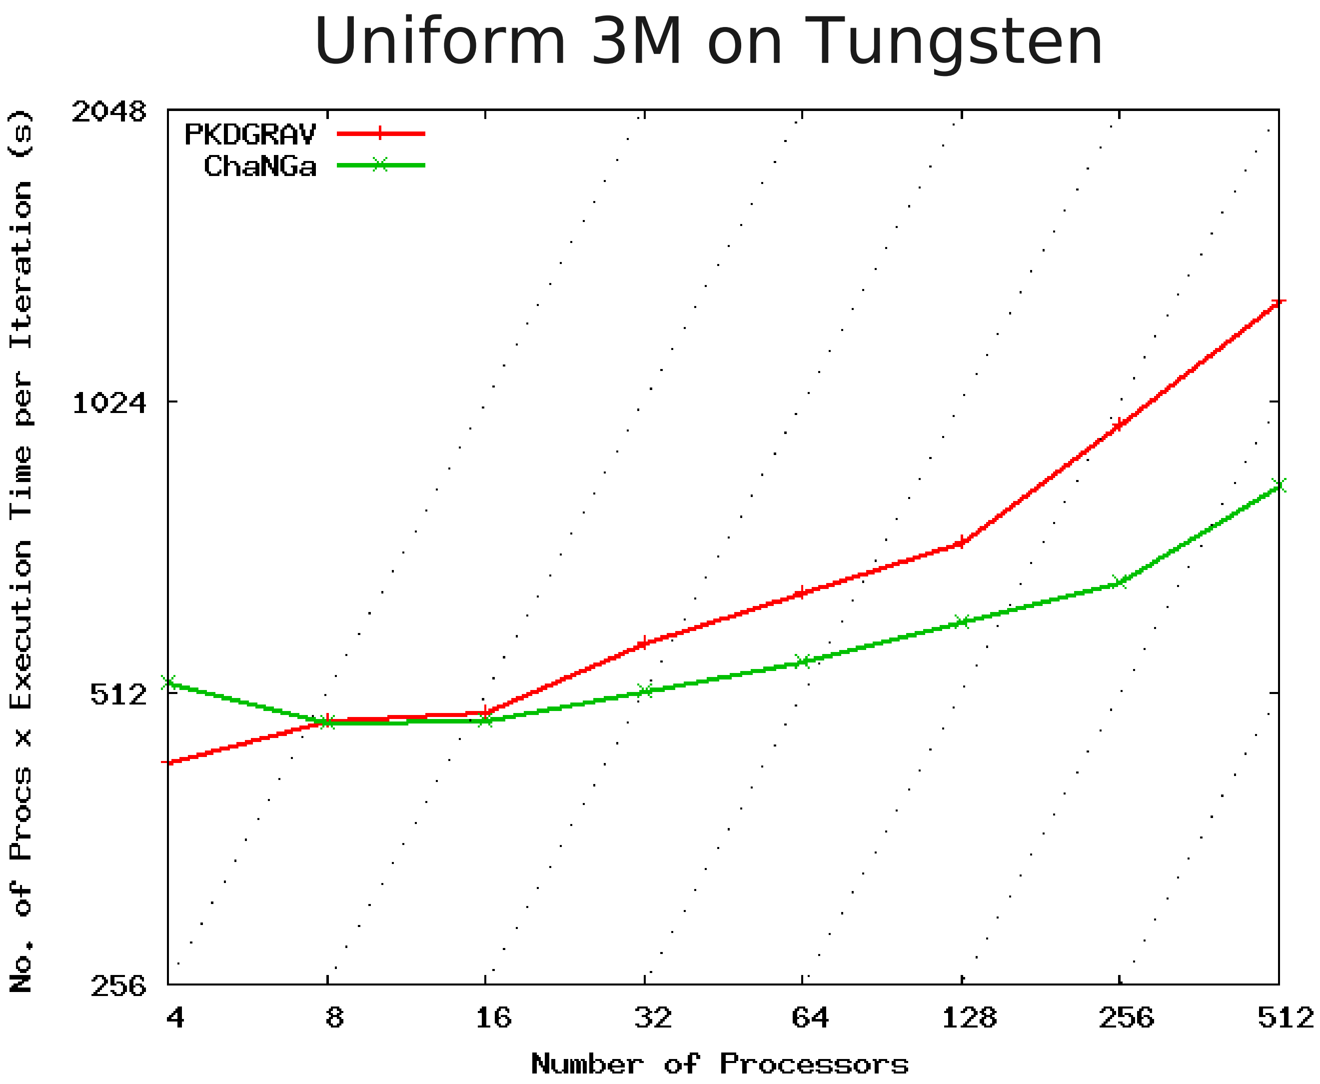

Load balancing with GreedyLB

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

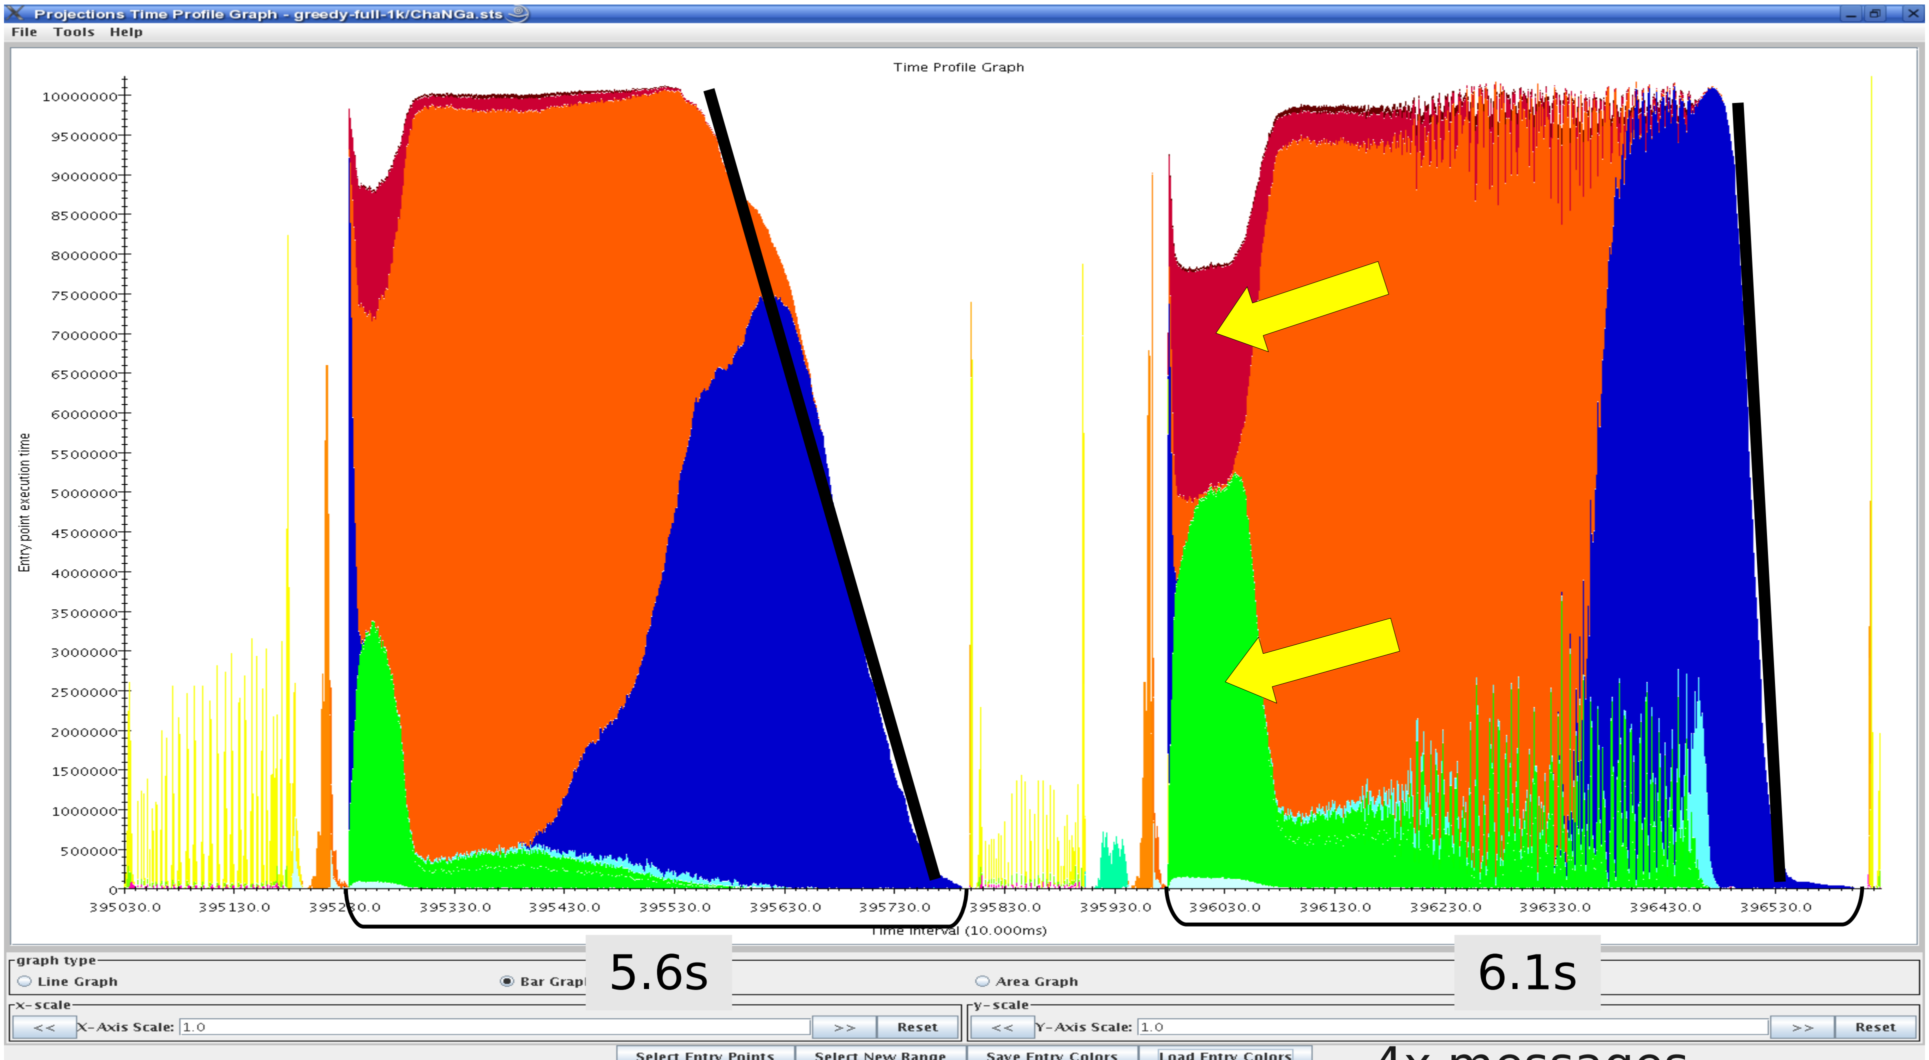

4x messages

Load balancing with OrbRefineLB

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

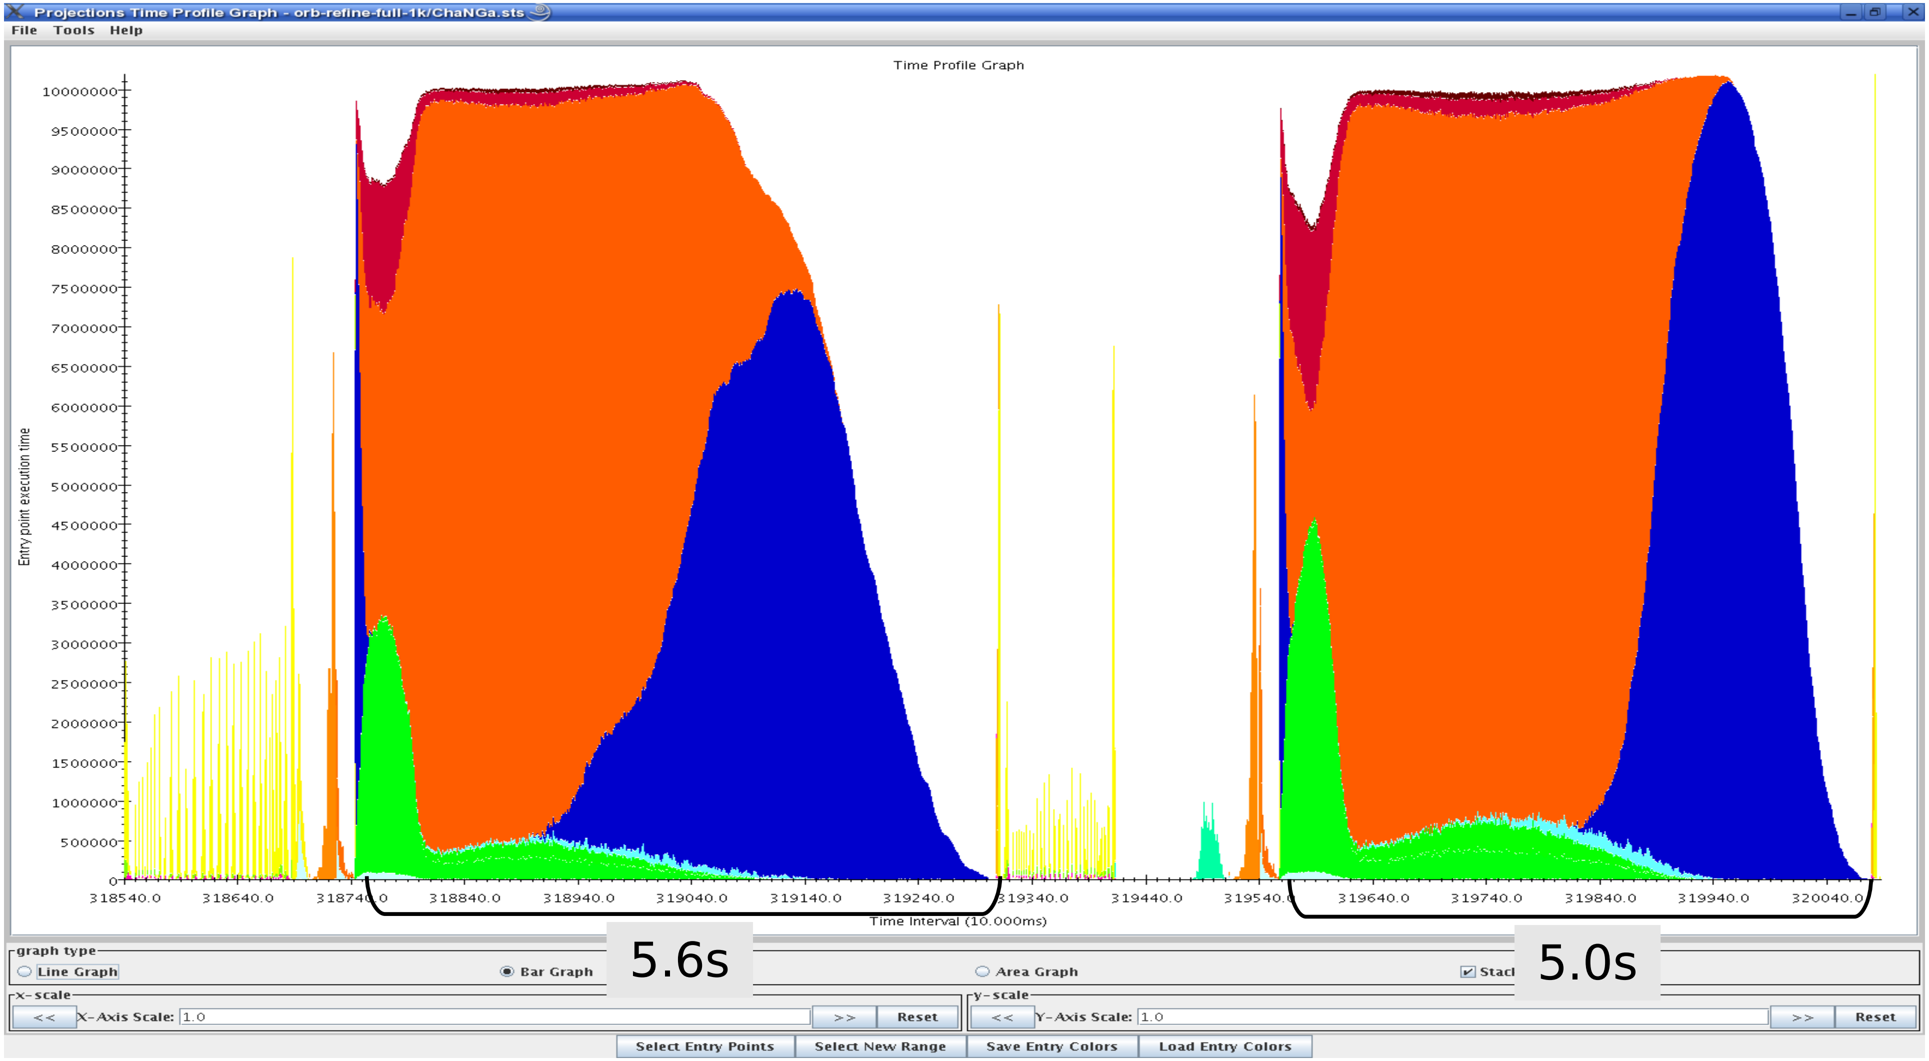

Scaling with load balancing

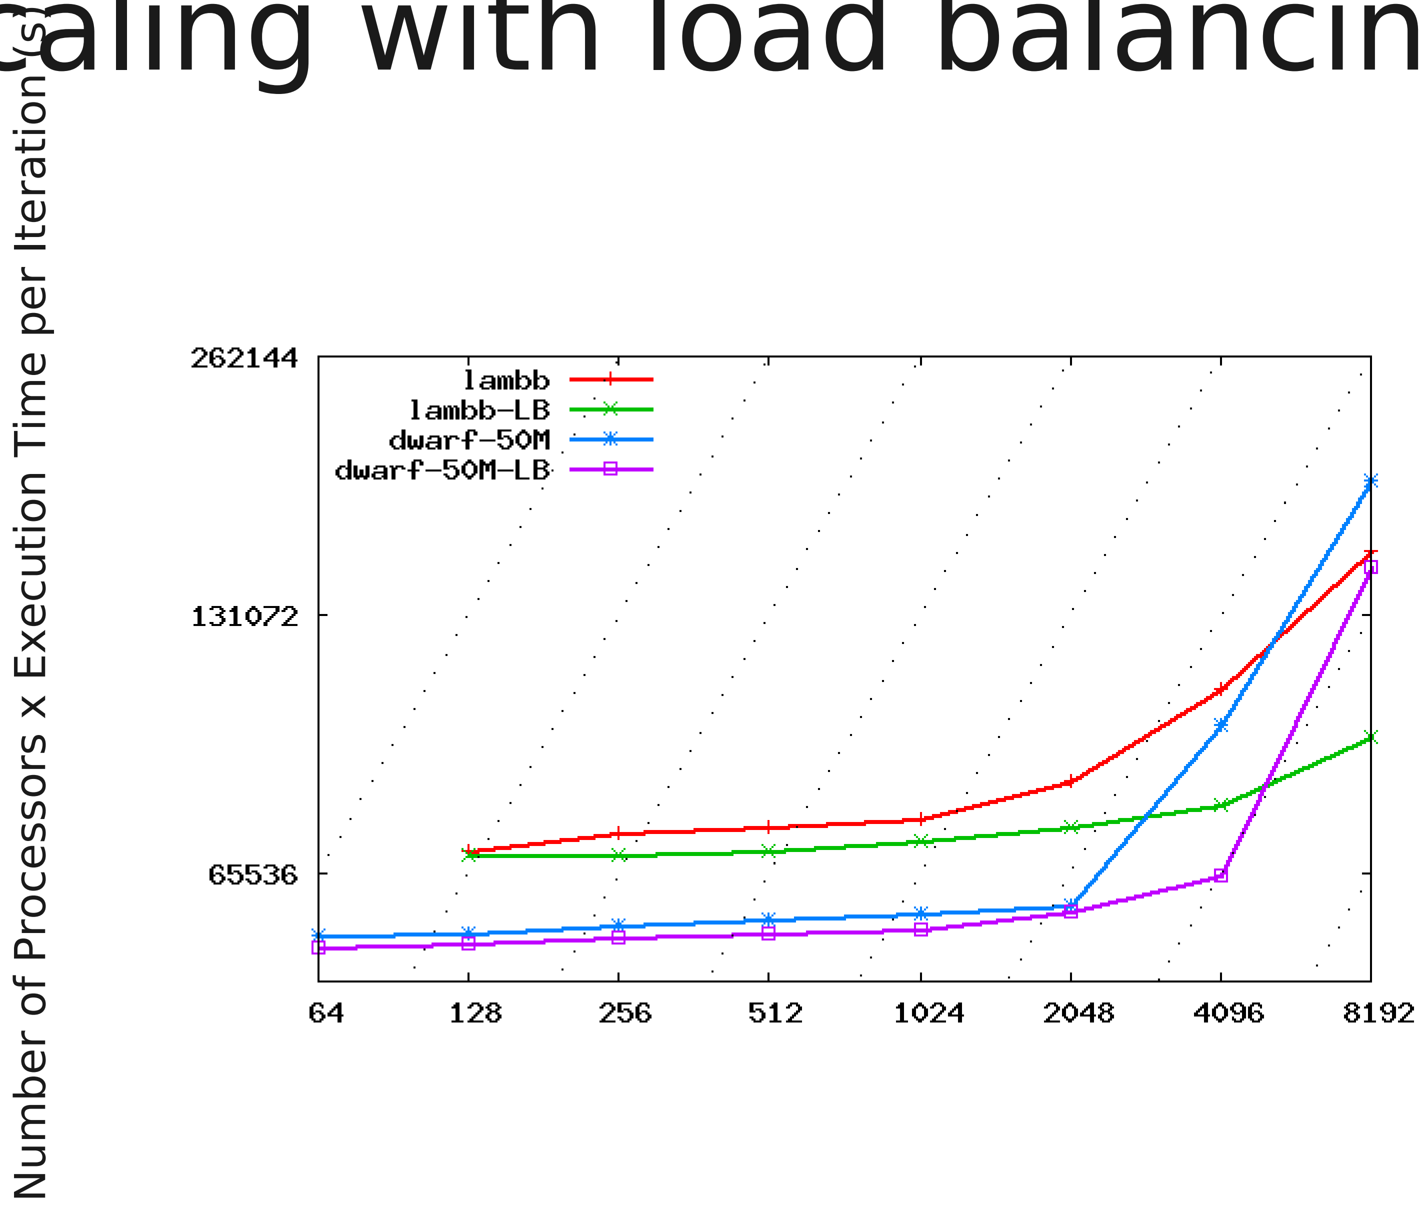

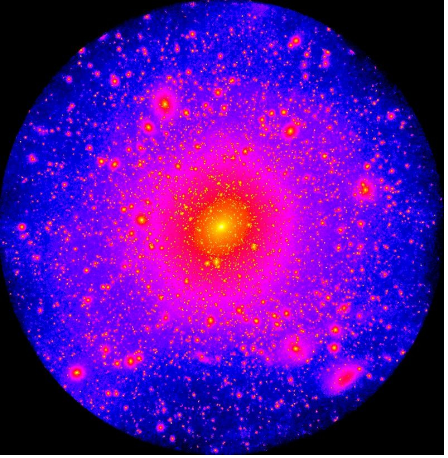

Timestepping Challenges

- 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

Timestepping and Load Balancing

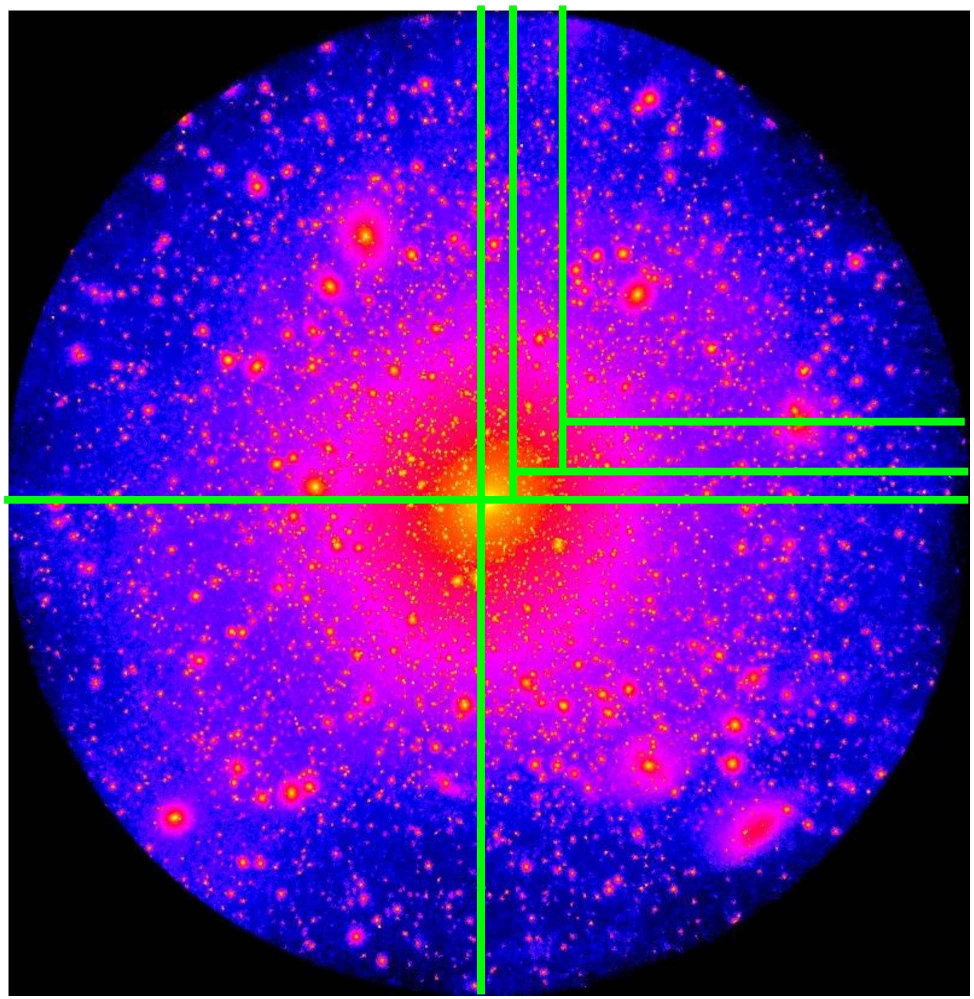

Cosmo Loadbalancer

- 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

Results on 3 rung example

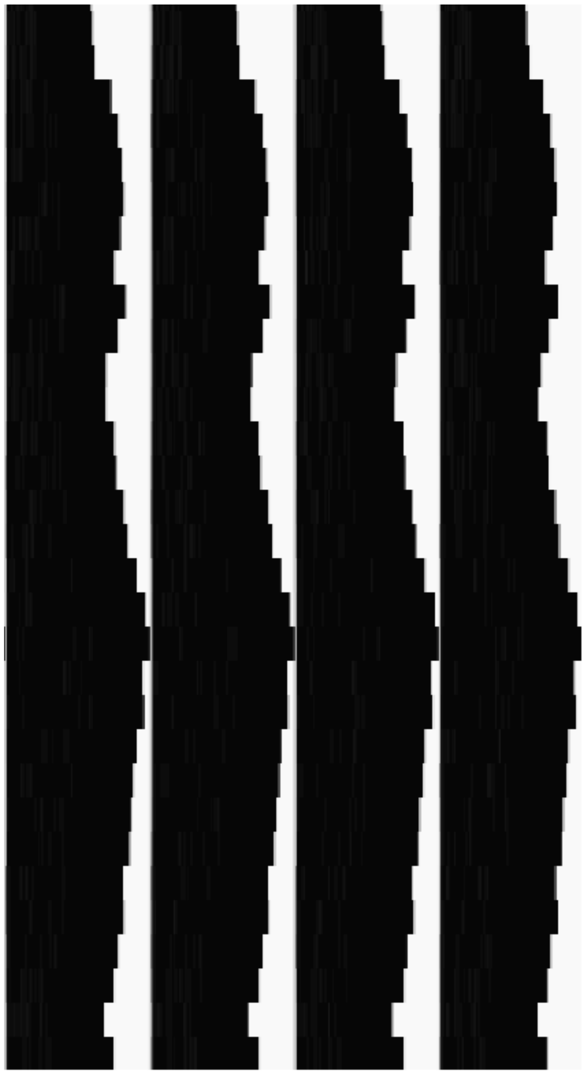

613s

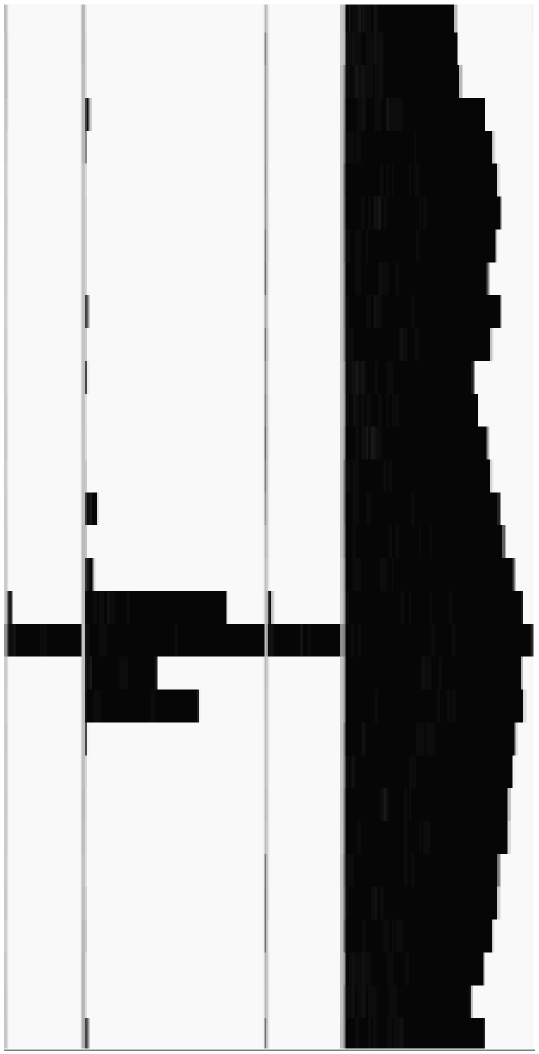

429s

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

- Changa currently in use in high time dynamic range simulations: galactic nuclei
- New Physics
  - Smooth particle hydrodynamics
- Better gravity algorithms
  - Fast multipole method
  - New domain decomposition/load balancing strategies
- Generic tree walk to enable new algorithms

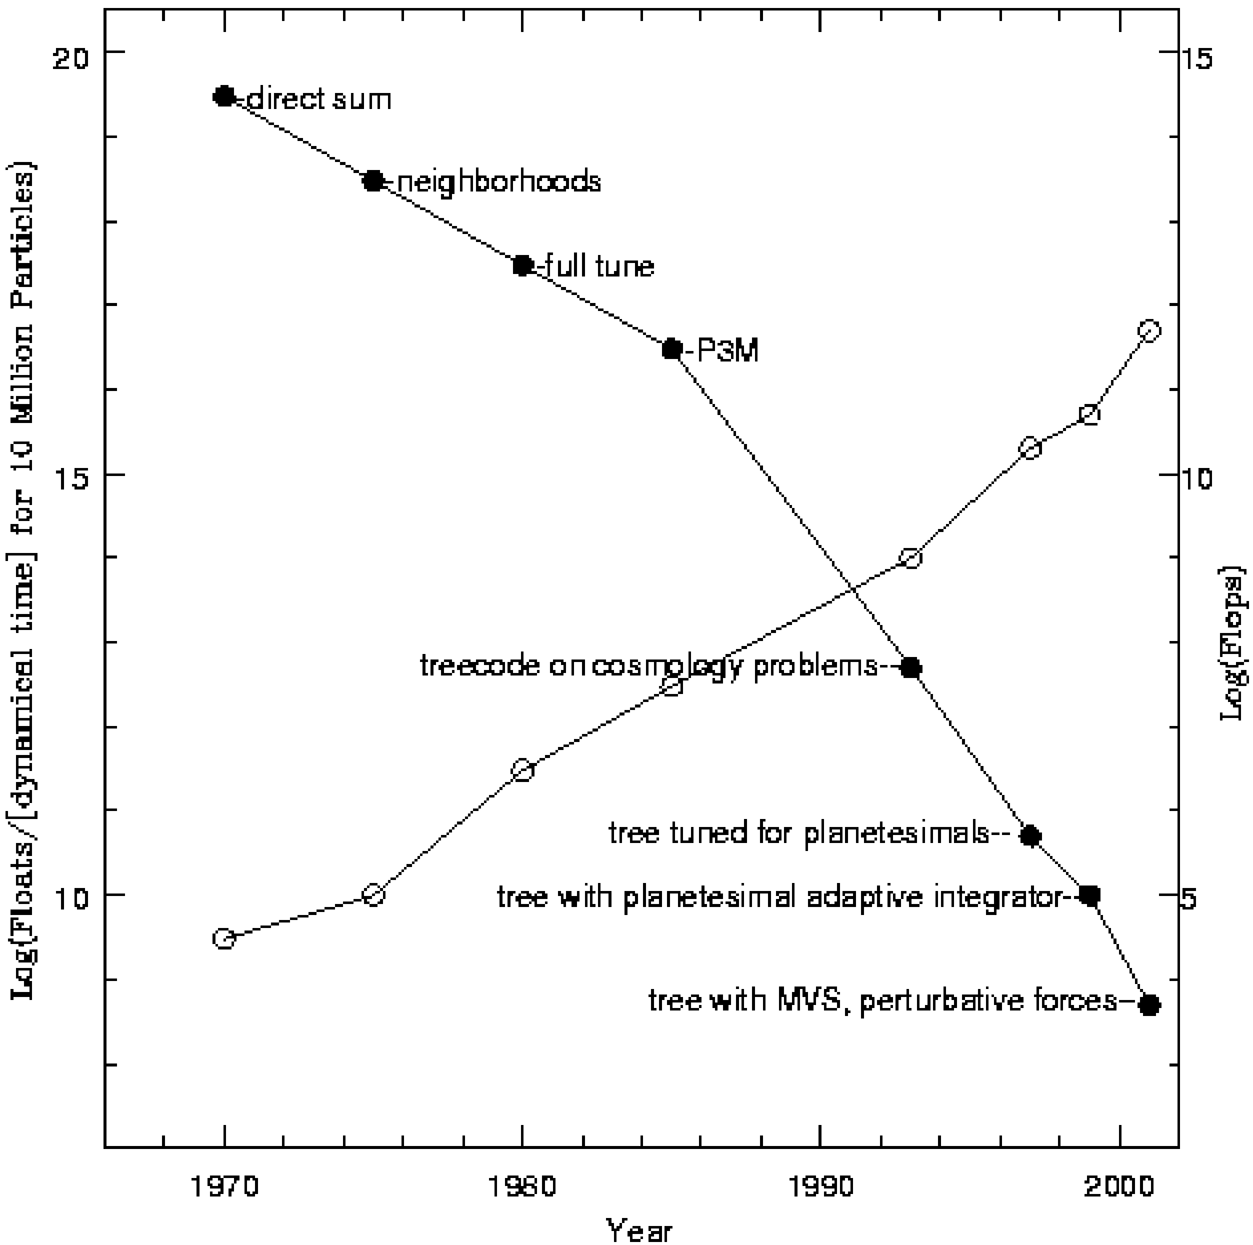

-direct sum
-neighborhoods
-full tune
-P3M
treecode on cosmology problems--
tree tuned for planetesimals--
tree with planetesimal adaptive integrator--
tree with MVS, perturbative forces--
20
15
10
15
10
5
1970
1980
1990
2000
Year
Log(Floats/[dynamical time] for 10 Million Particles)
Log(Flops)

Have We converged?

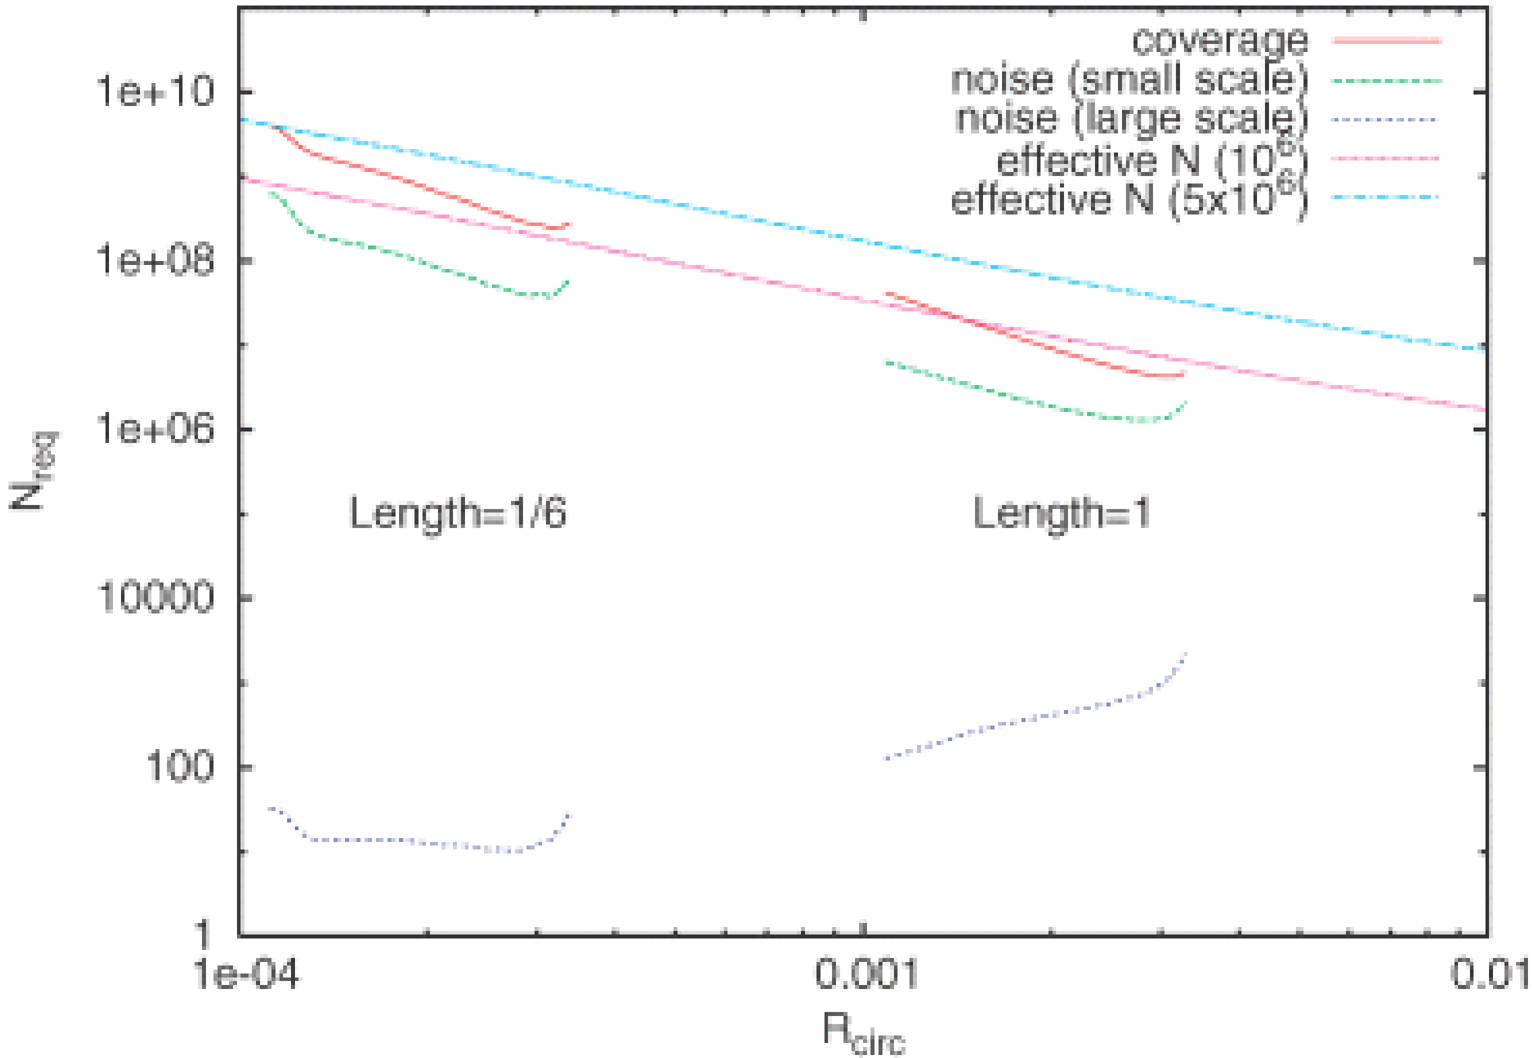

Weinberg & Katz (2007)

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.