Next-Generation Cosmology Simulations with HACC: Challenges from Baryons

Katrin Heitmann

Lunch Talk, Computational Performance Workshop

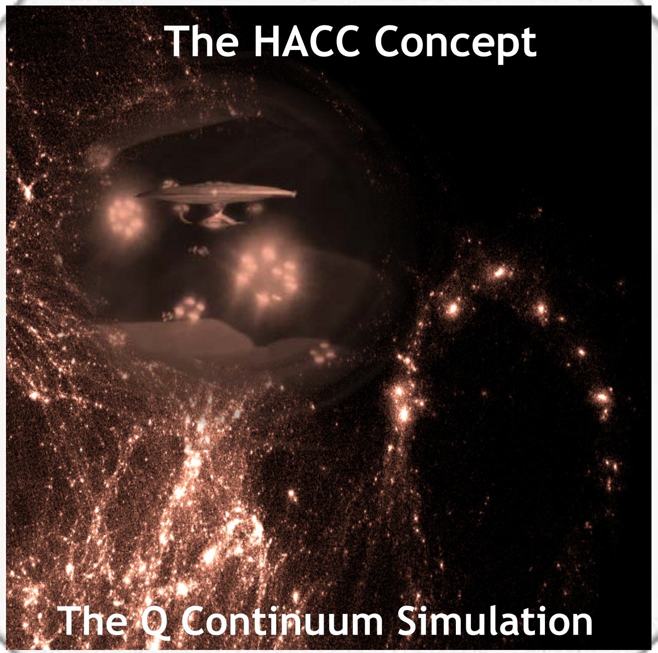

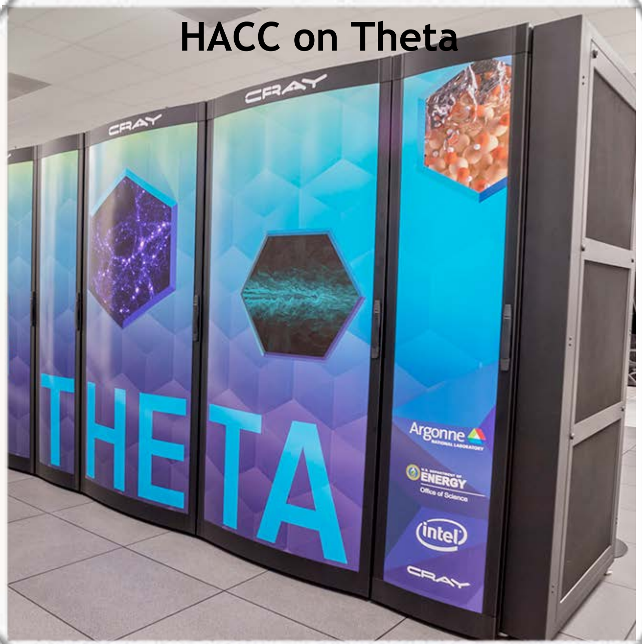

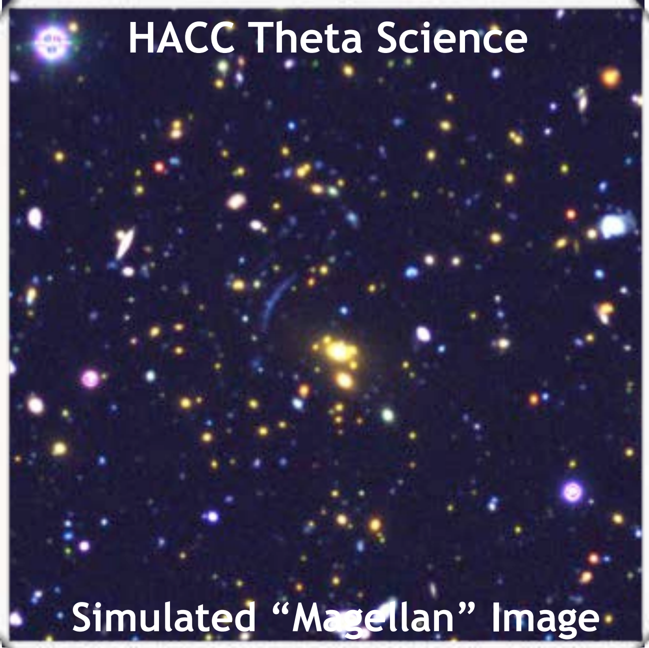

* HACC = Hardware/Hybrid Accelerated Cosmology Code

The Cosmologist's View on Baryons

The Cosmologist's View on Baryons

Next-Generation Cosmology Simulations with HACC*: Challenges from Baryons

Katrin Heitmann

Lunch Talk, Computational Performance Workshop

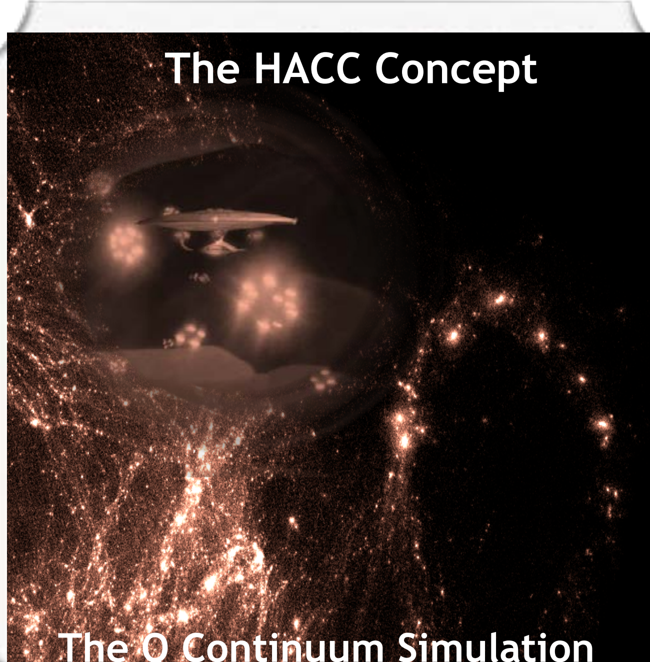

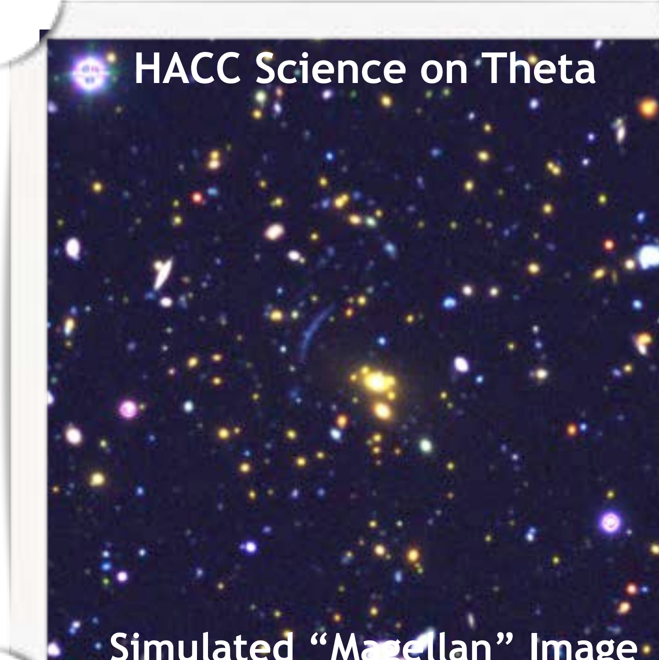

* HACC = Hardware/Hybrid Accelerated Cosmology Code

Our Group at Argonne and Close Collaborators

HEP Division

ALCF Division

External

● Core HACC Team ● HACC Science Collaborators

The Evolution of the Universe: Structure Formation

- **Solid understanding of structure formation; success underpins most cosmic discovery**
  - Initial conditions determined by primordial fluctuations, measured from the cosmic microwave background
  - Initial perturbations amplified by gravitational instability in a dark matter-dominated Universe
  - Relevant theory is gravity, field theory, and atomic physics ('first principles')
- **Early Universe:** Linear perturbation theory very successful (CMB)
- **Latter half of the history of the Universe:** Nonlinear domain of structure formation, impossible to treat without large-scale computing

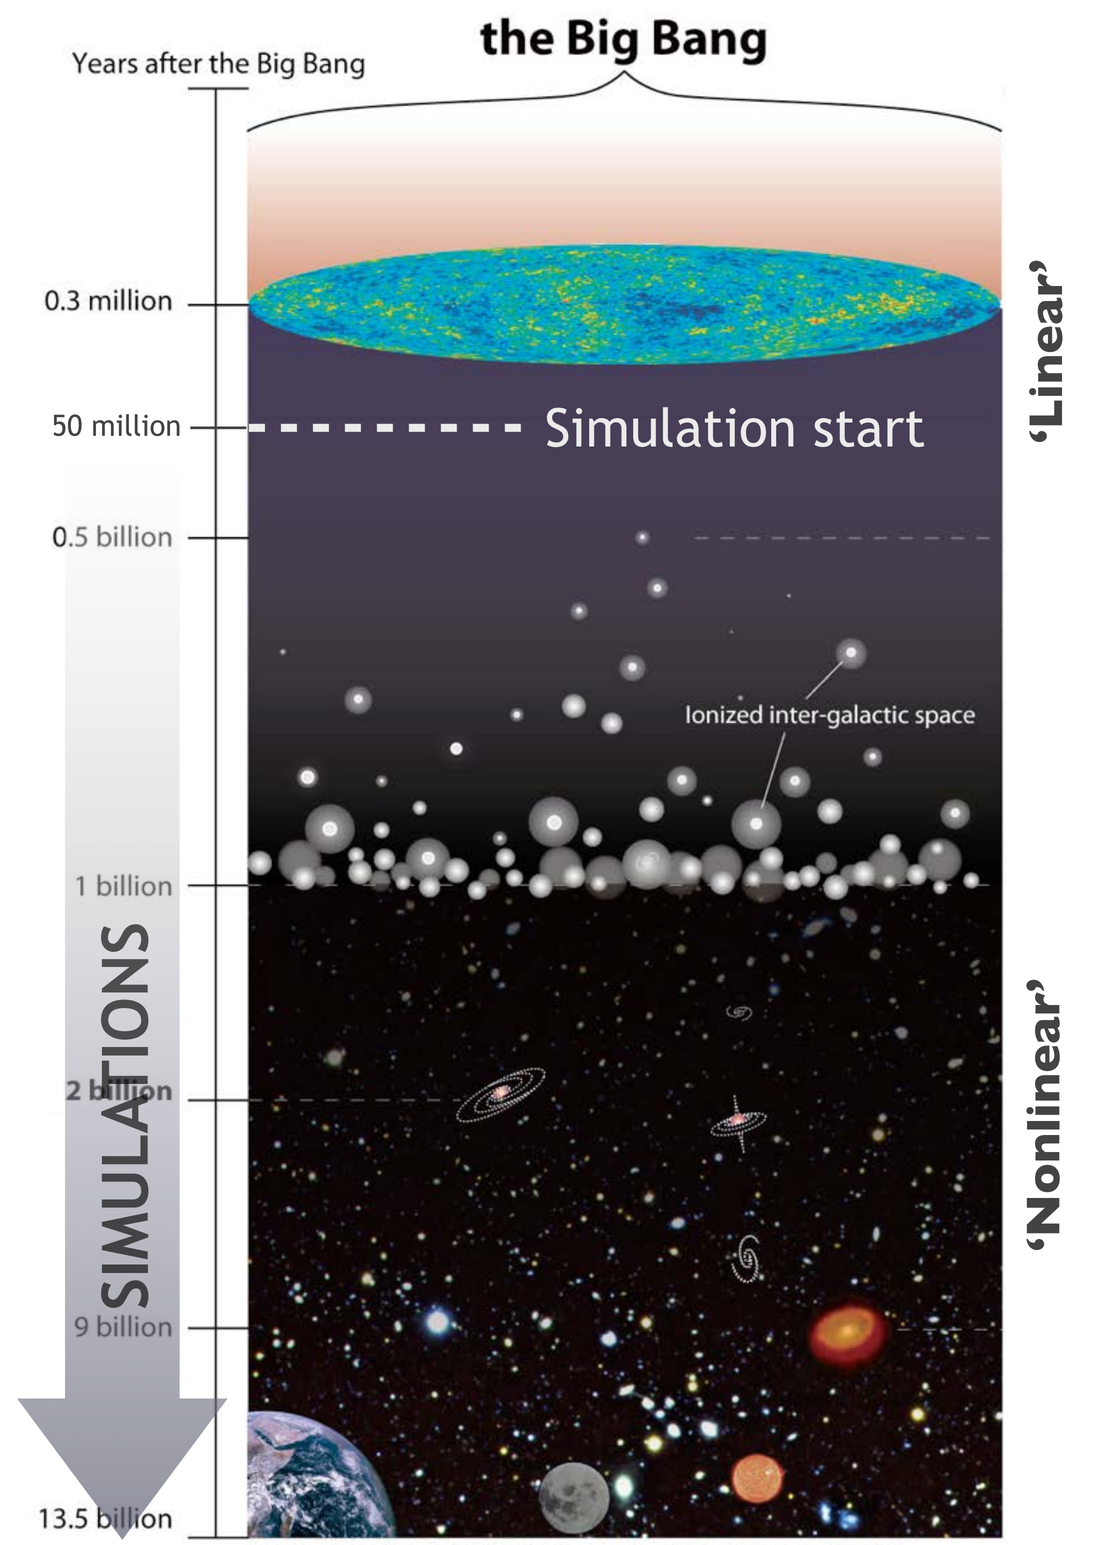

Computing the Universe: Gravity-only

- Gravity dominates **on large scales**, use Monte Carlo sampling of density with tracer particles
- Particles are tracers of the dark matter in the Universe, mass typically at least ~10^9 $M_\odot$

$$m_p \sim V/n_p$$

- Simulate galaxy size objects ($v^2/c^2 << 1$), Newtonian description accurate
- **At smaller scales, add gas physics, feedback etc.**, sub-grid modeling inevitable

> "The Universe is far too complicated a structure to be studied deductively, starting from initial conditions and solving the equations of motion."
> *Robert Dicke (Jayne Lectures, 1969)*

$$\ddot{\mathbf{x}} + 2\frac{\dot{a}}{a}\dot{\mathbf{x}} = -\frac{\nabla\Phi}{a^2}$$ Equation of motion for tracer particles in expanding Universe

$$\frac{\dot{a}}{a} = H = \frac{H_0}{a^{3/2}}\sqrt{\Omega_{tot} + a^3\Omega_\Lambda}$$ CDM + baryons + DE

$$\nabla^2\Phi(\mathbf{x}) = 4\pi G a^2[\rho(\mathbf{x},t) - \rho_b(t)]$$ Poisson equation

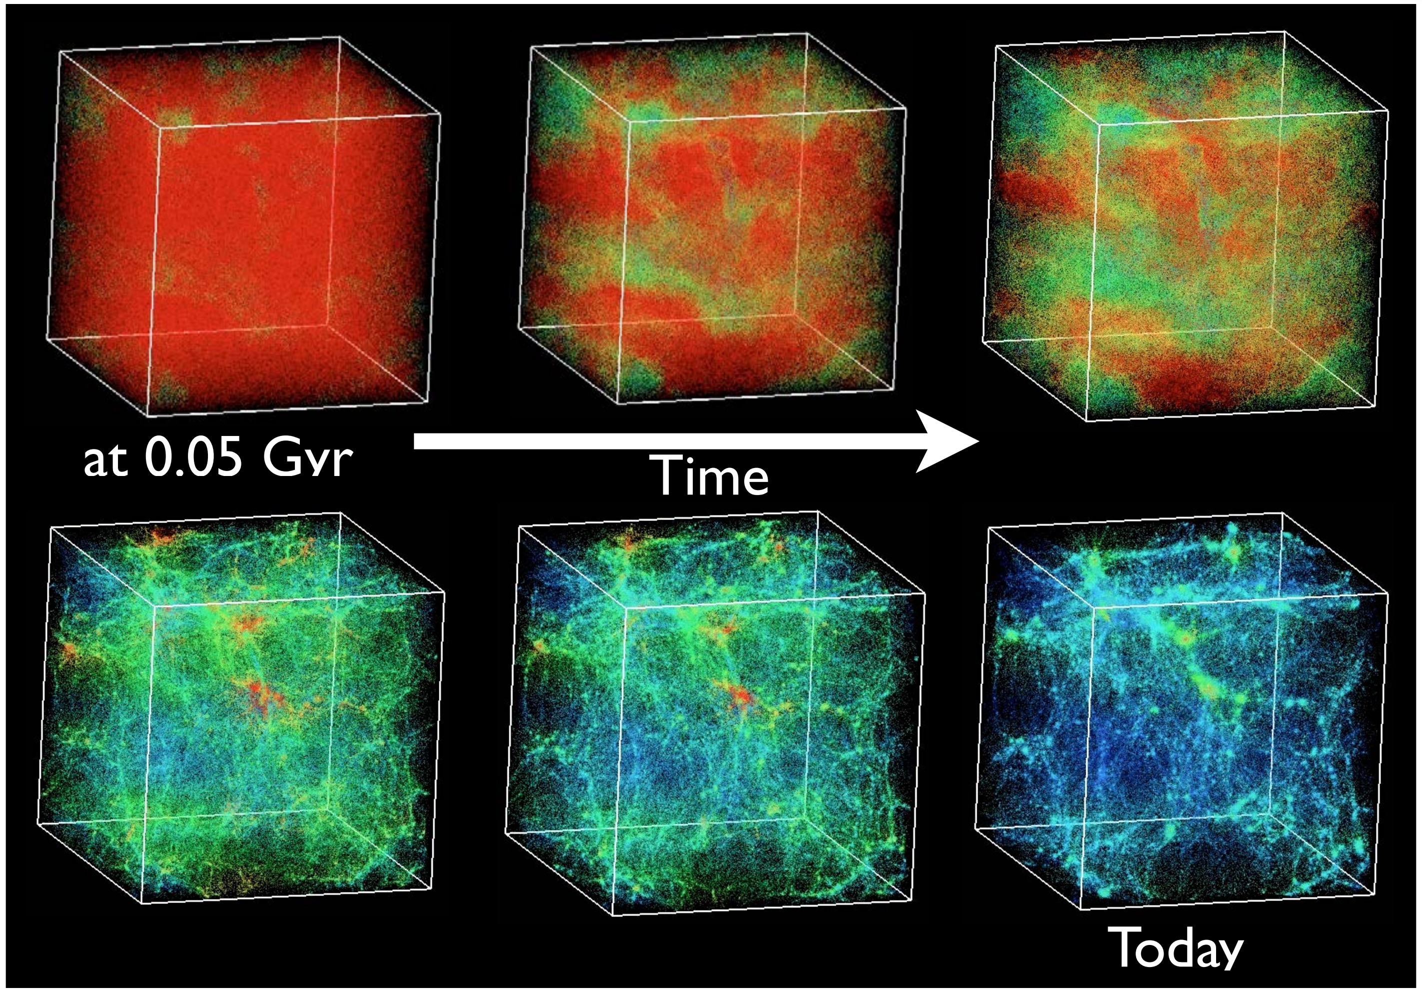

HACC in a Nutshell

S. Habib et al. 2016, New Astronomy

- **Long-range/short range force splitting:**
  - **Long-range:** Particle-Mesh solver, C/C++/MPI, **unchanged for different architectures,** FFT performance dictates scaling (custom pencil decomposed FFT)
  - **Short-range:** **Depending on node architecture** switch between tree and particle-particle algorithm; tree needs "thinking" (building, walking) but computationally less demanding (BG/Q, X86, KNL), PP easier but computationally more expensive (GPU)
- **Overload concept to allow for easy swap of short-range solver and minimization of communication (reassignment of passive/active in regular intervals)**
- **Adaptive time stepping, analysis on the fly, mixed precision, custom I/O, ...**

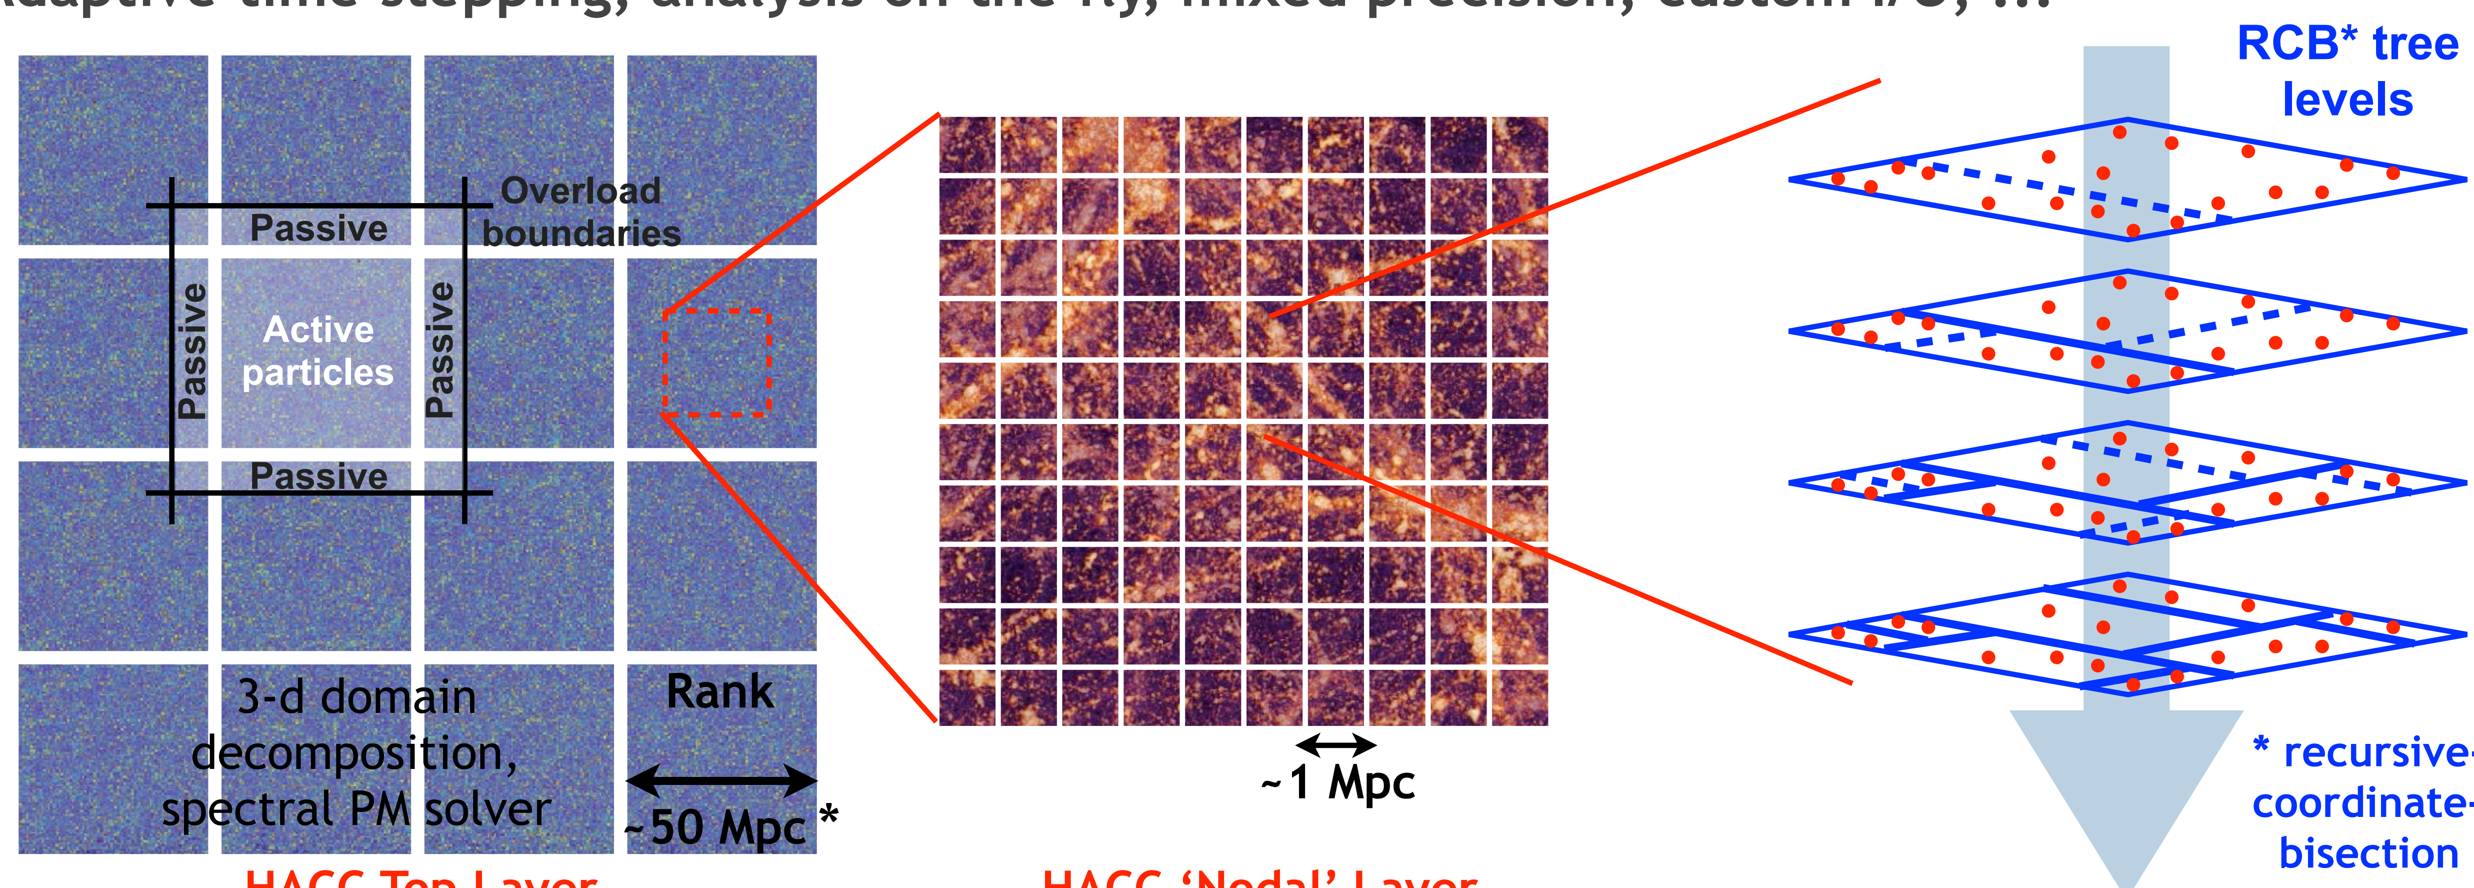

HACC Top Layer

HACC 'Nodal' Layer

* Mpc~3M light years, distance between 2 galaxies

Tree Portability

- Each rank divides its volume into chaining mesh cells of width CM
- An RCB tree us built in each CM cell. Tree bisection stops once the number of particles on a leaf drops below PPN
- Customize PP force kernel for specific architecture
- Data structure arranged for easy vectorization of the PP kernel and tree walk parallelized using dynamic OpenMP threads
- Flexibility in CM and PPN allows for easy portability to new machines
    - For Mira (BG/Q): PPN ~ 128 optimal
    - For Edison (X86): PPN ~ 20 optimal
    - For Theta (KNL): PPN again higher

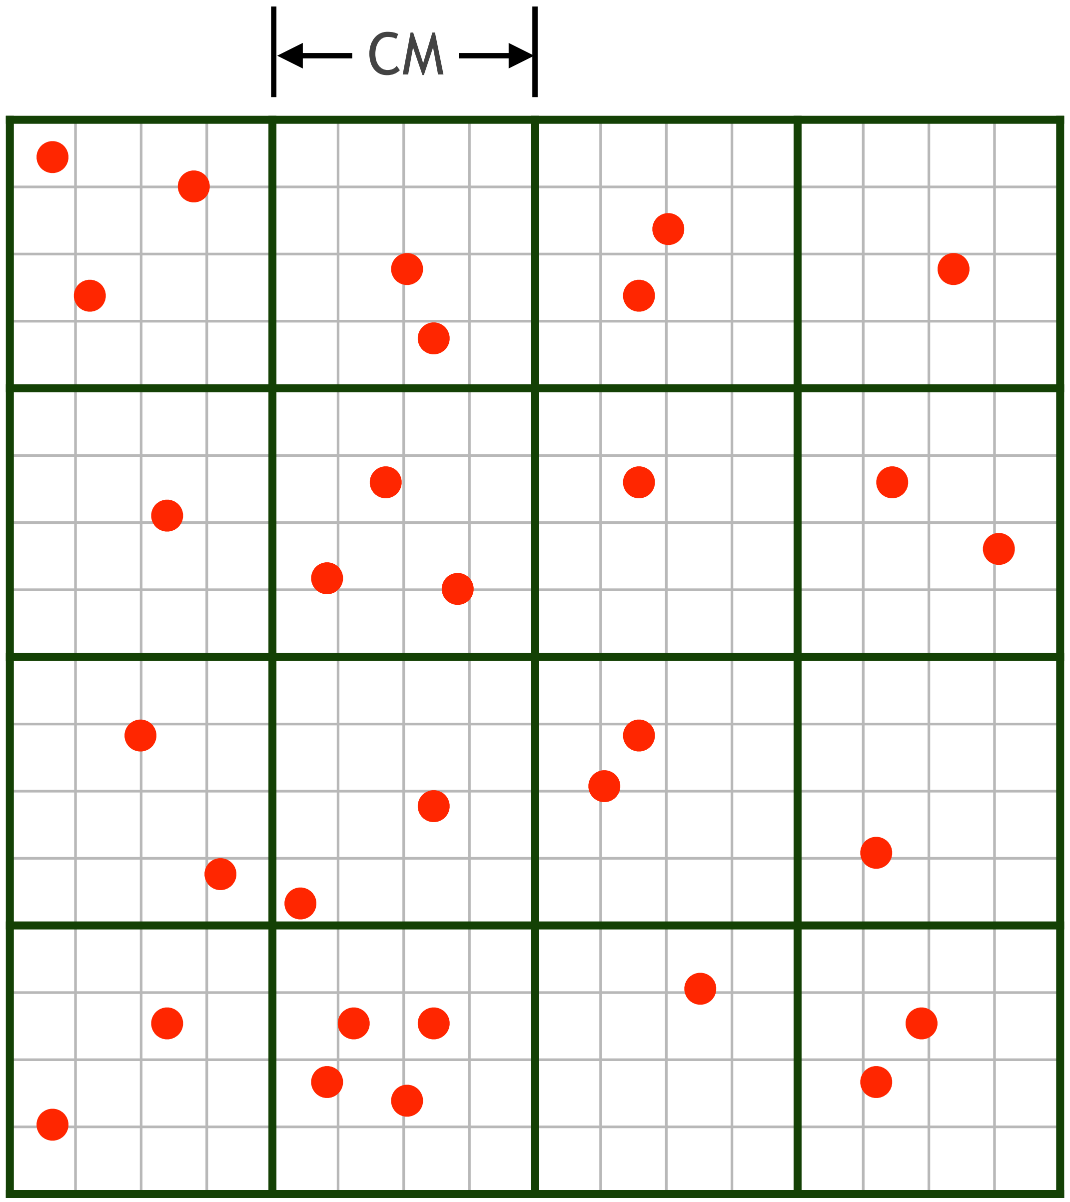

Single rank tree structure

HACC Timing: One BG/Q Node versus 1/4 KNL Node

Cores	RPN	OMP	TH	BG/Q [Time/s]	KNL [Time/sec]	Ratio
16	4	4	16	4297	616.33	6.98
16	4	8	32	2677	543.73	4.92
16	4	16	64	2504	530.23	4.72
16	8	2	16	4362	544.75	8.00
16	8	4	32	2571	459.53	5.59
16	8	8	64	**2278**	**437.21**	5.18
16	16	4	64	2581	468.50	5.50

RPN: MPI ranks for each run; **OMP:** OpenMP threads per MPI rank, **TH:** total number of threads = RPN*OMP

- Problem set up: 320^3 particles, 320^3 grid, 3 full time steps, 5 sub-cycles
- Results for flat mode and cache mode are essentially identical

HACC Weak Scaling on Theta

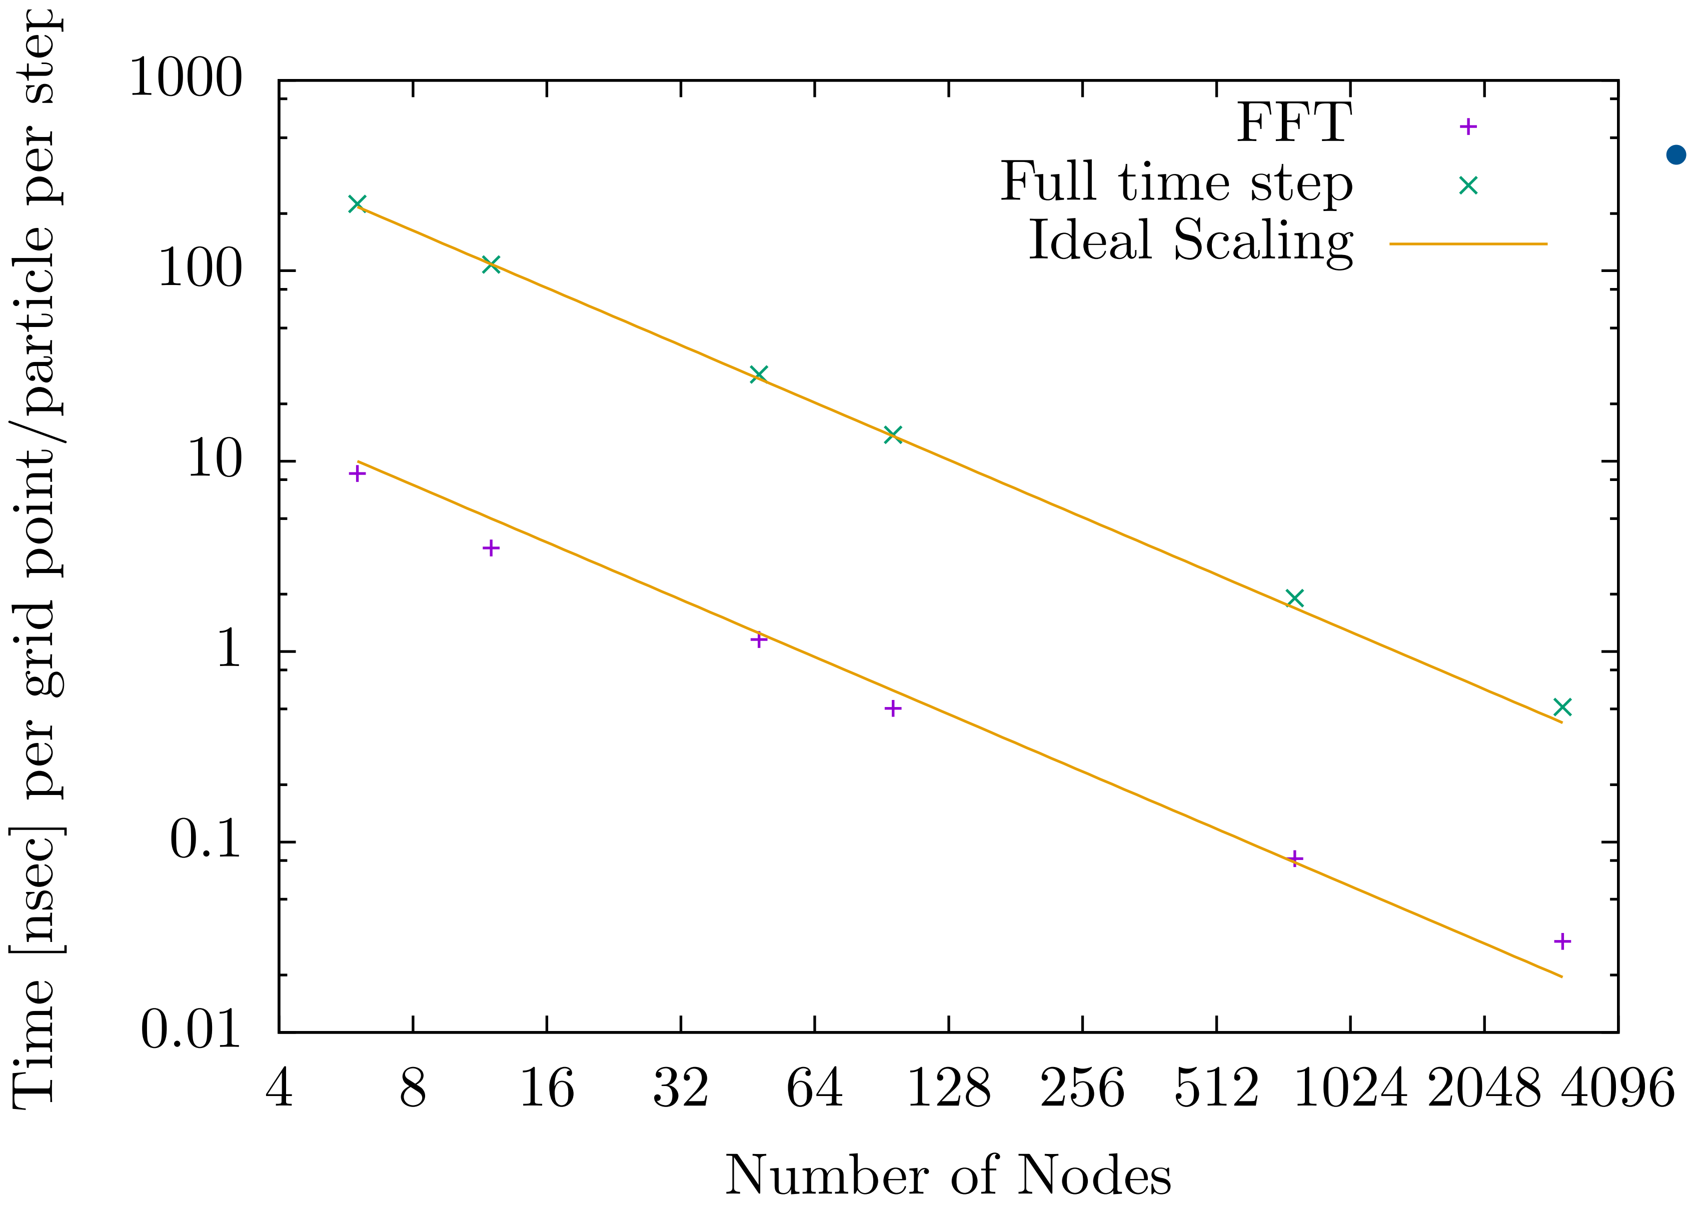

- **Problem set up:**
  - 6 nodes: 1152^3 particles in $(312.5\ \mathrm{Mpc/h})^3$ volume
  - 3072 nodes: 9216^3 particles in $(2500\ \mathrm{Mpc/h})^3$ volume
  - 8 ranks per nodes, 16 threads
  - Cache-quad mode

HACC Strong Scaling on Theta

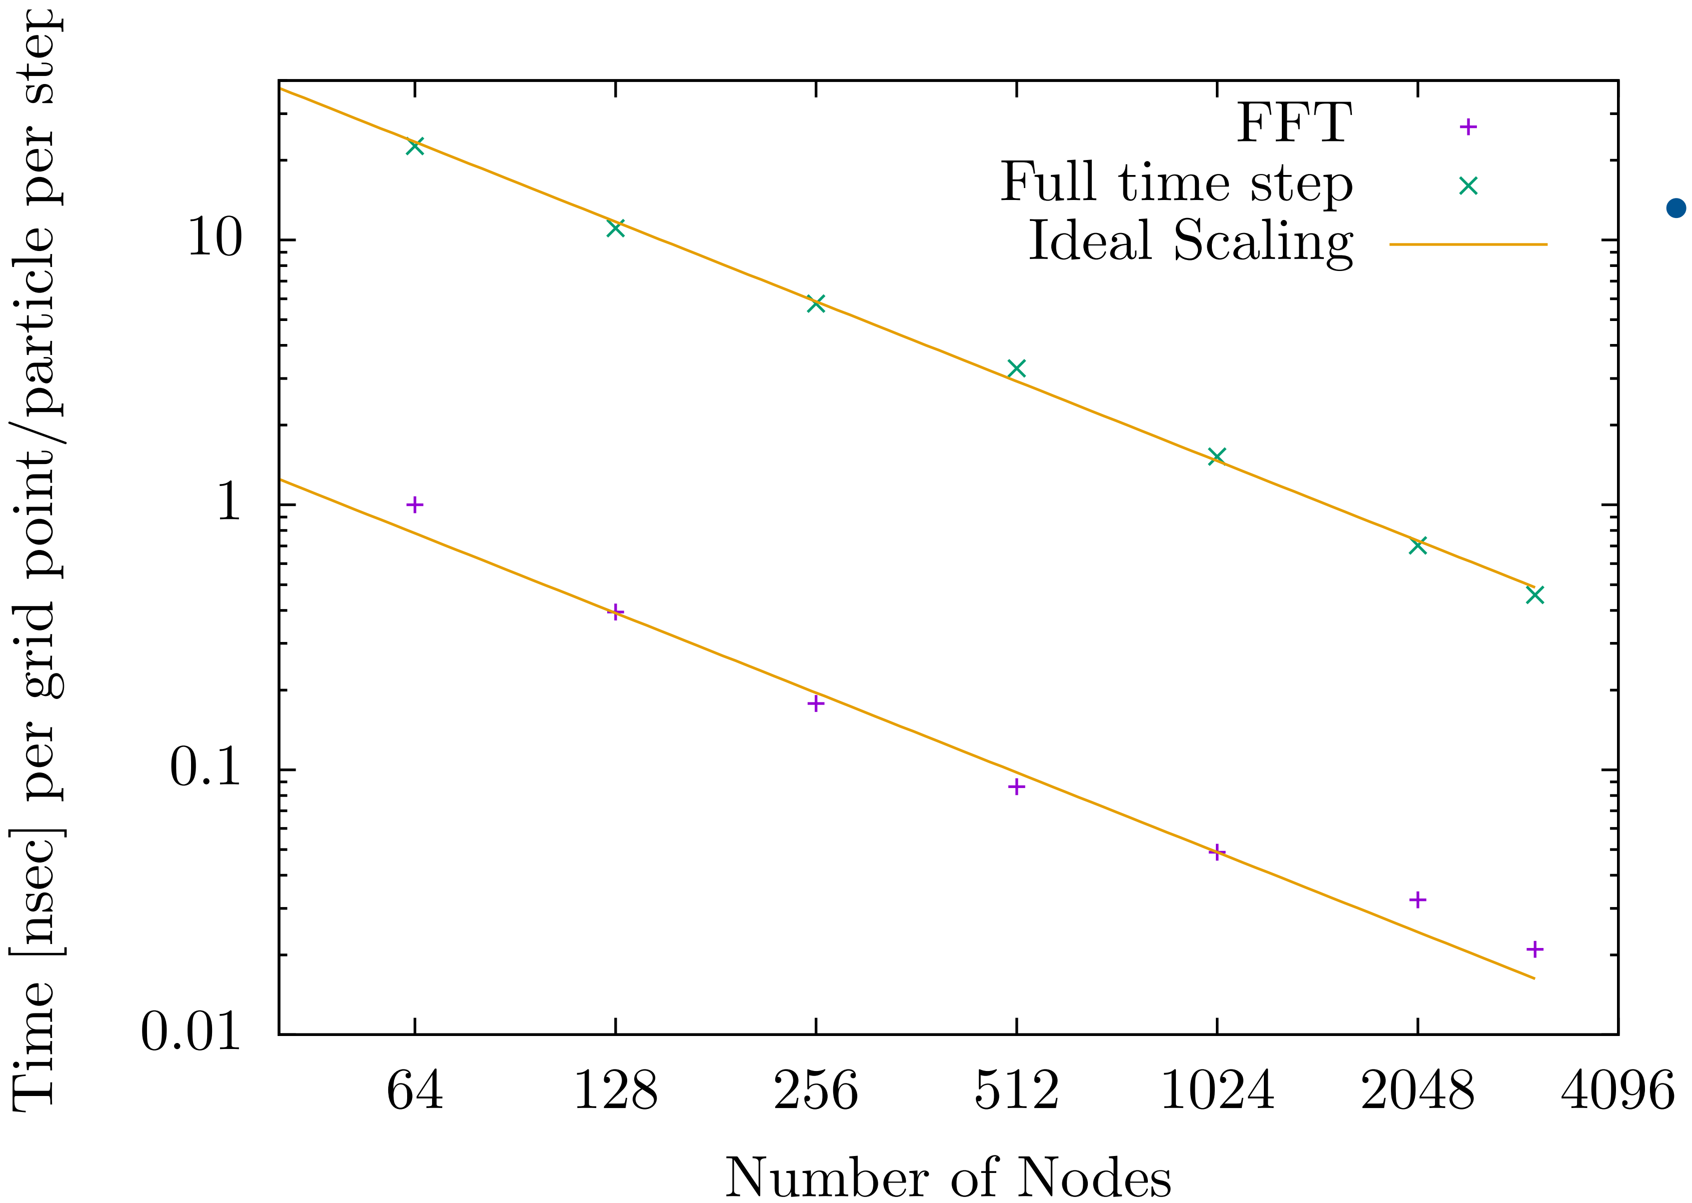

- **Problem set up:**
  - 2048^3 particles on 64 to 2048 nodes (and 2304^3 particles on 3072 nodes)
  - 8 ranks per nodes, 16 threads
  - Cache-quad mode

The Borg Cube Simulation

- **Science target on Theta**
  - (2x)2304^3 particles in $(800\ \mathrm{Mpc/h})^3$ volume
  - Study of maps relevant to cross-correlation between cosmic microwave background and large-scale structure
  - Gravity-only simulation + modeling of baryons in post-processing vs. hydro simulation
  - Gravity-only simulation finished

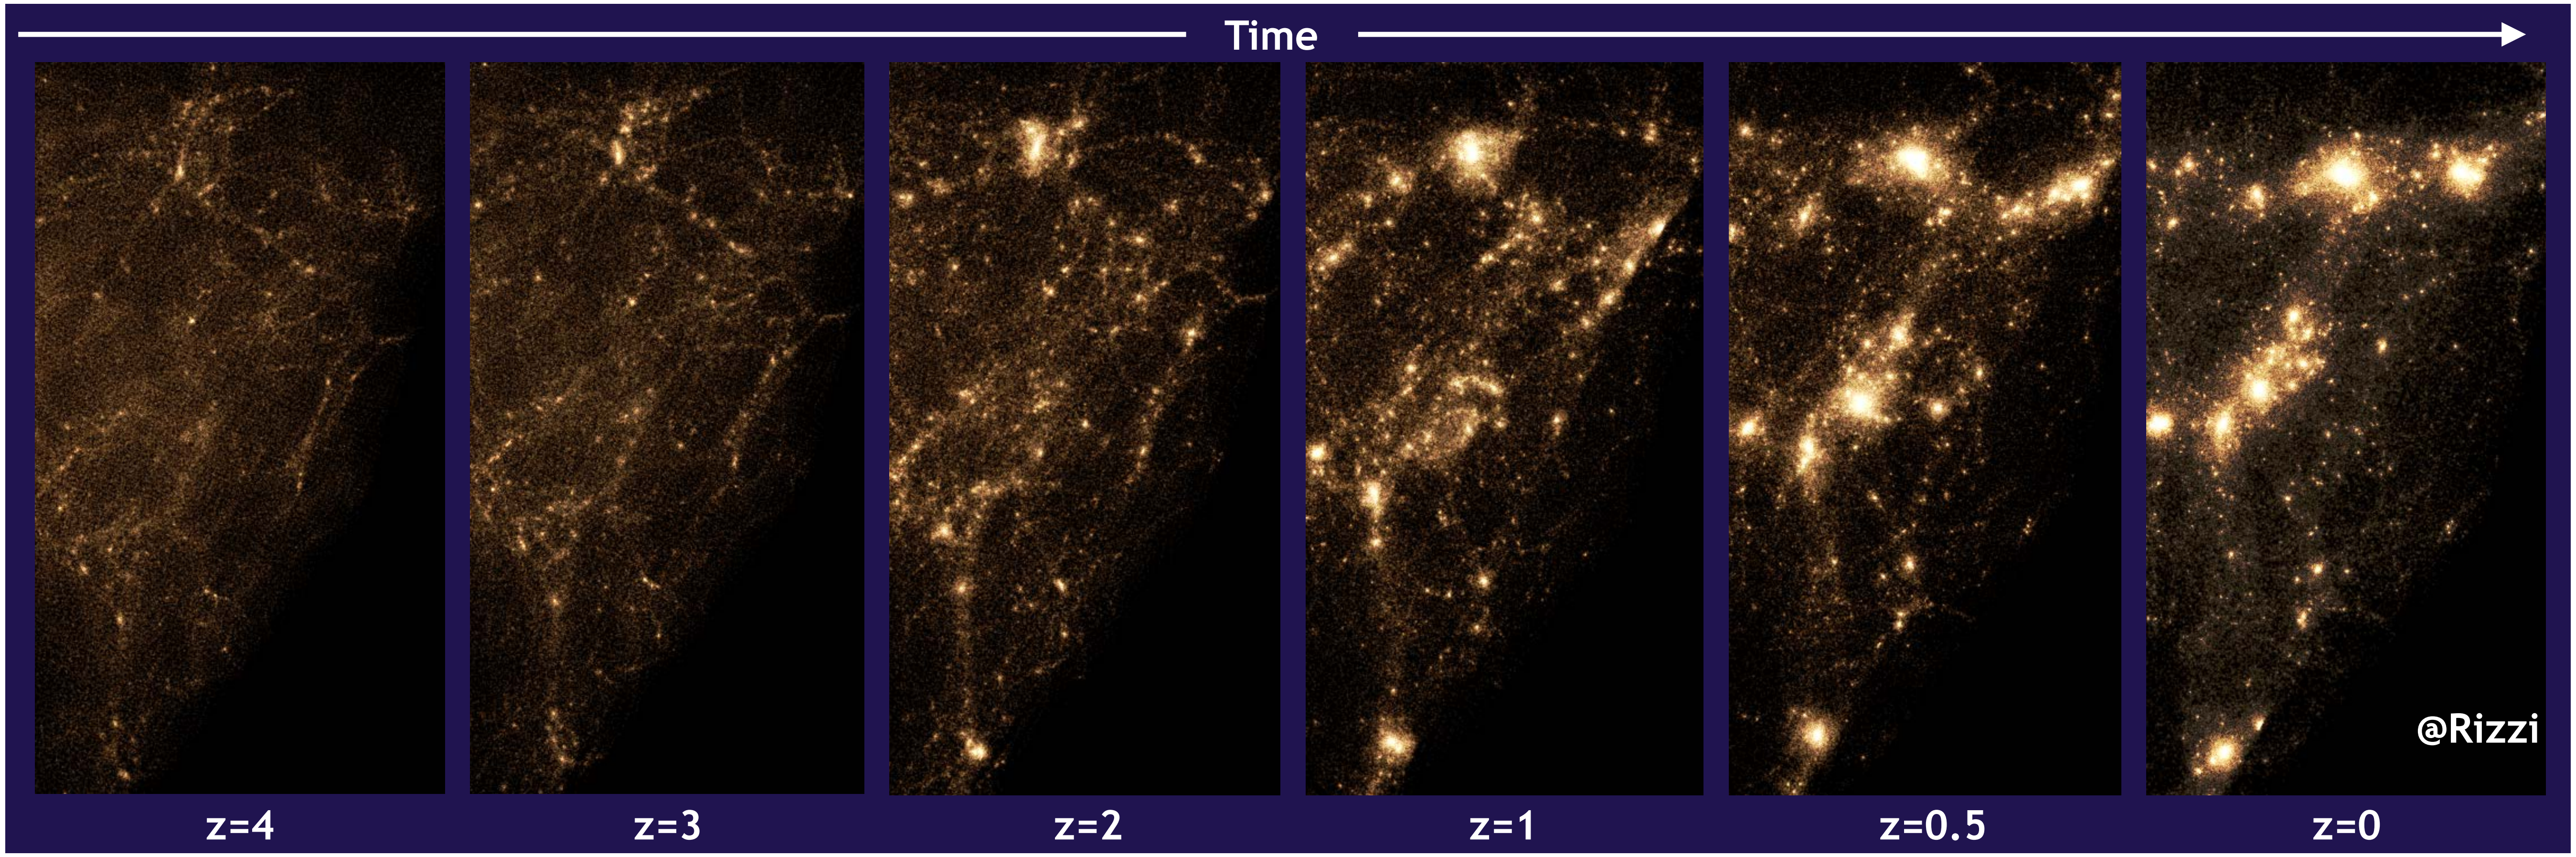

Portable Analysis Tools with Thrust

- Halo finder identifies groups of particles within a certain distance
- Center finder is expensive piece, calculate local potential minimum
- Thrust: NVIDIA library of parallel algorithms and data structures to perform sorts, scans, transforms, etc.
- Works on GPU and KNL
- Test on 64 KNL nodes

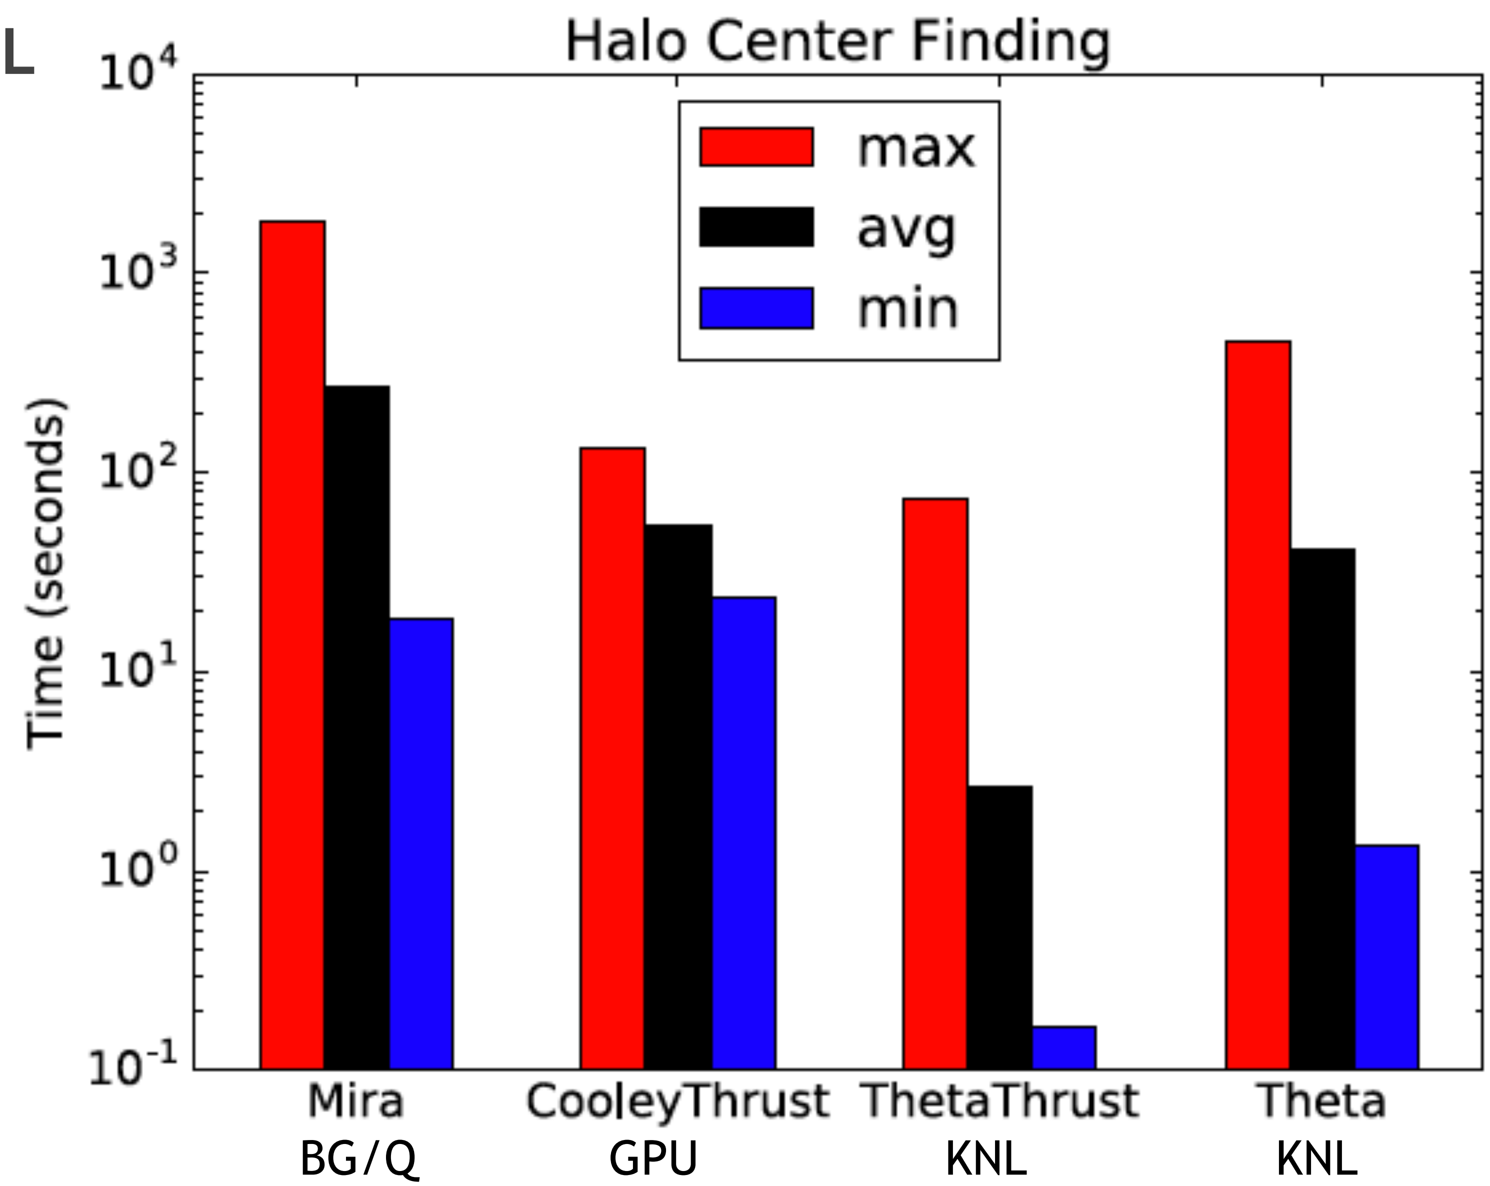

Current State-of-the-Art Gravity-Only Simulations

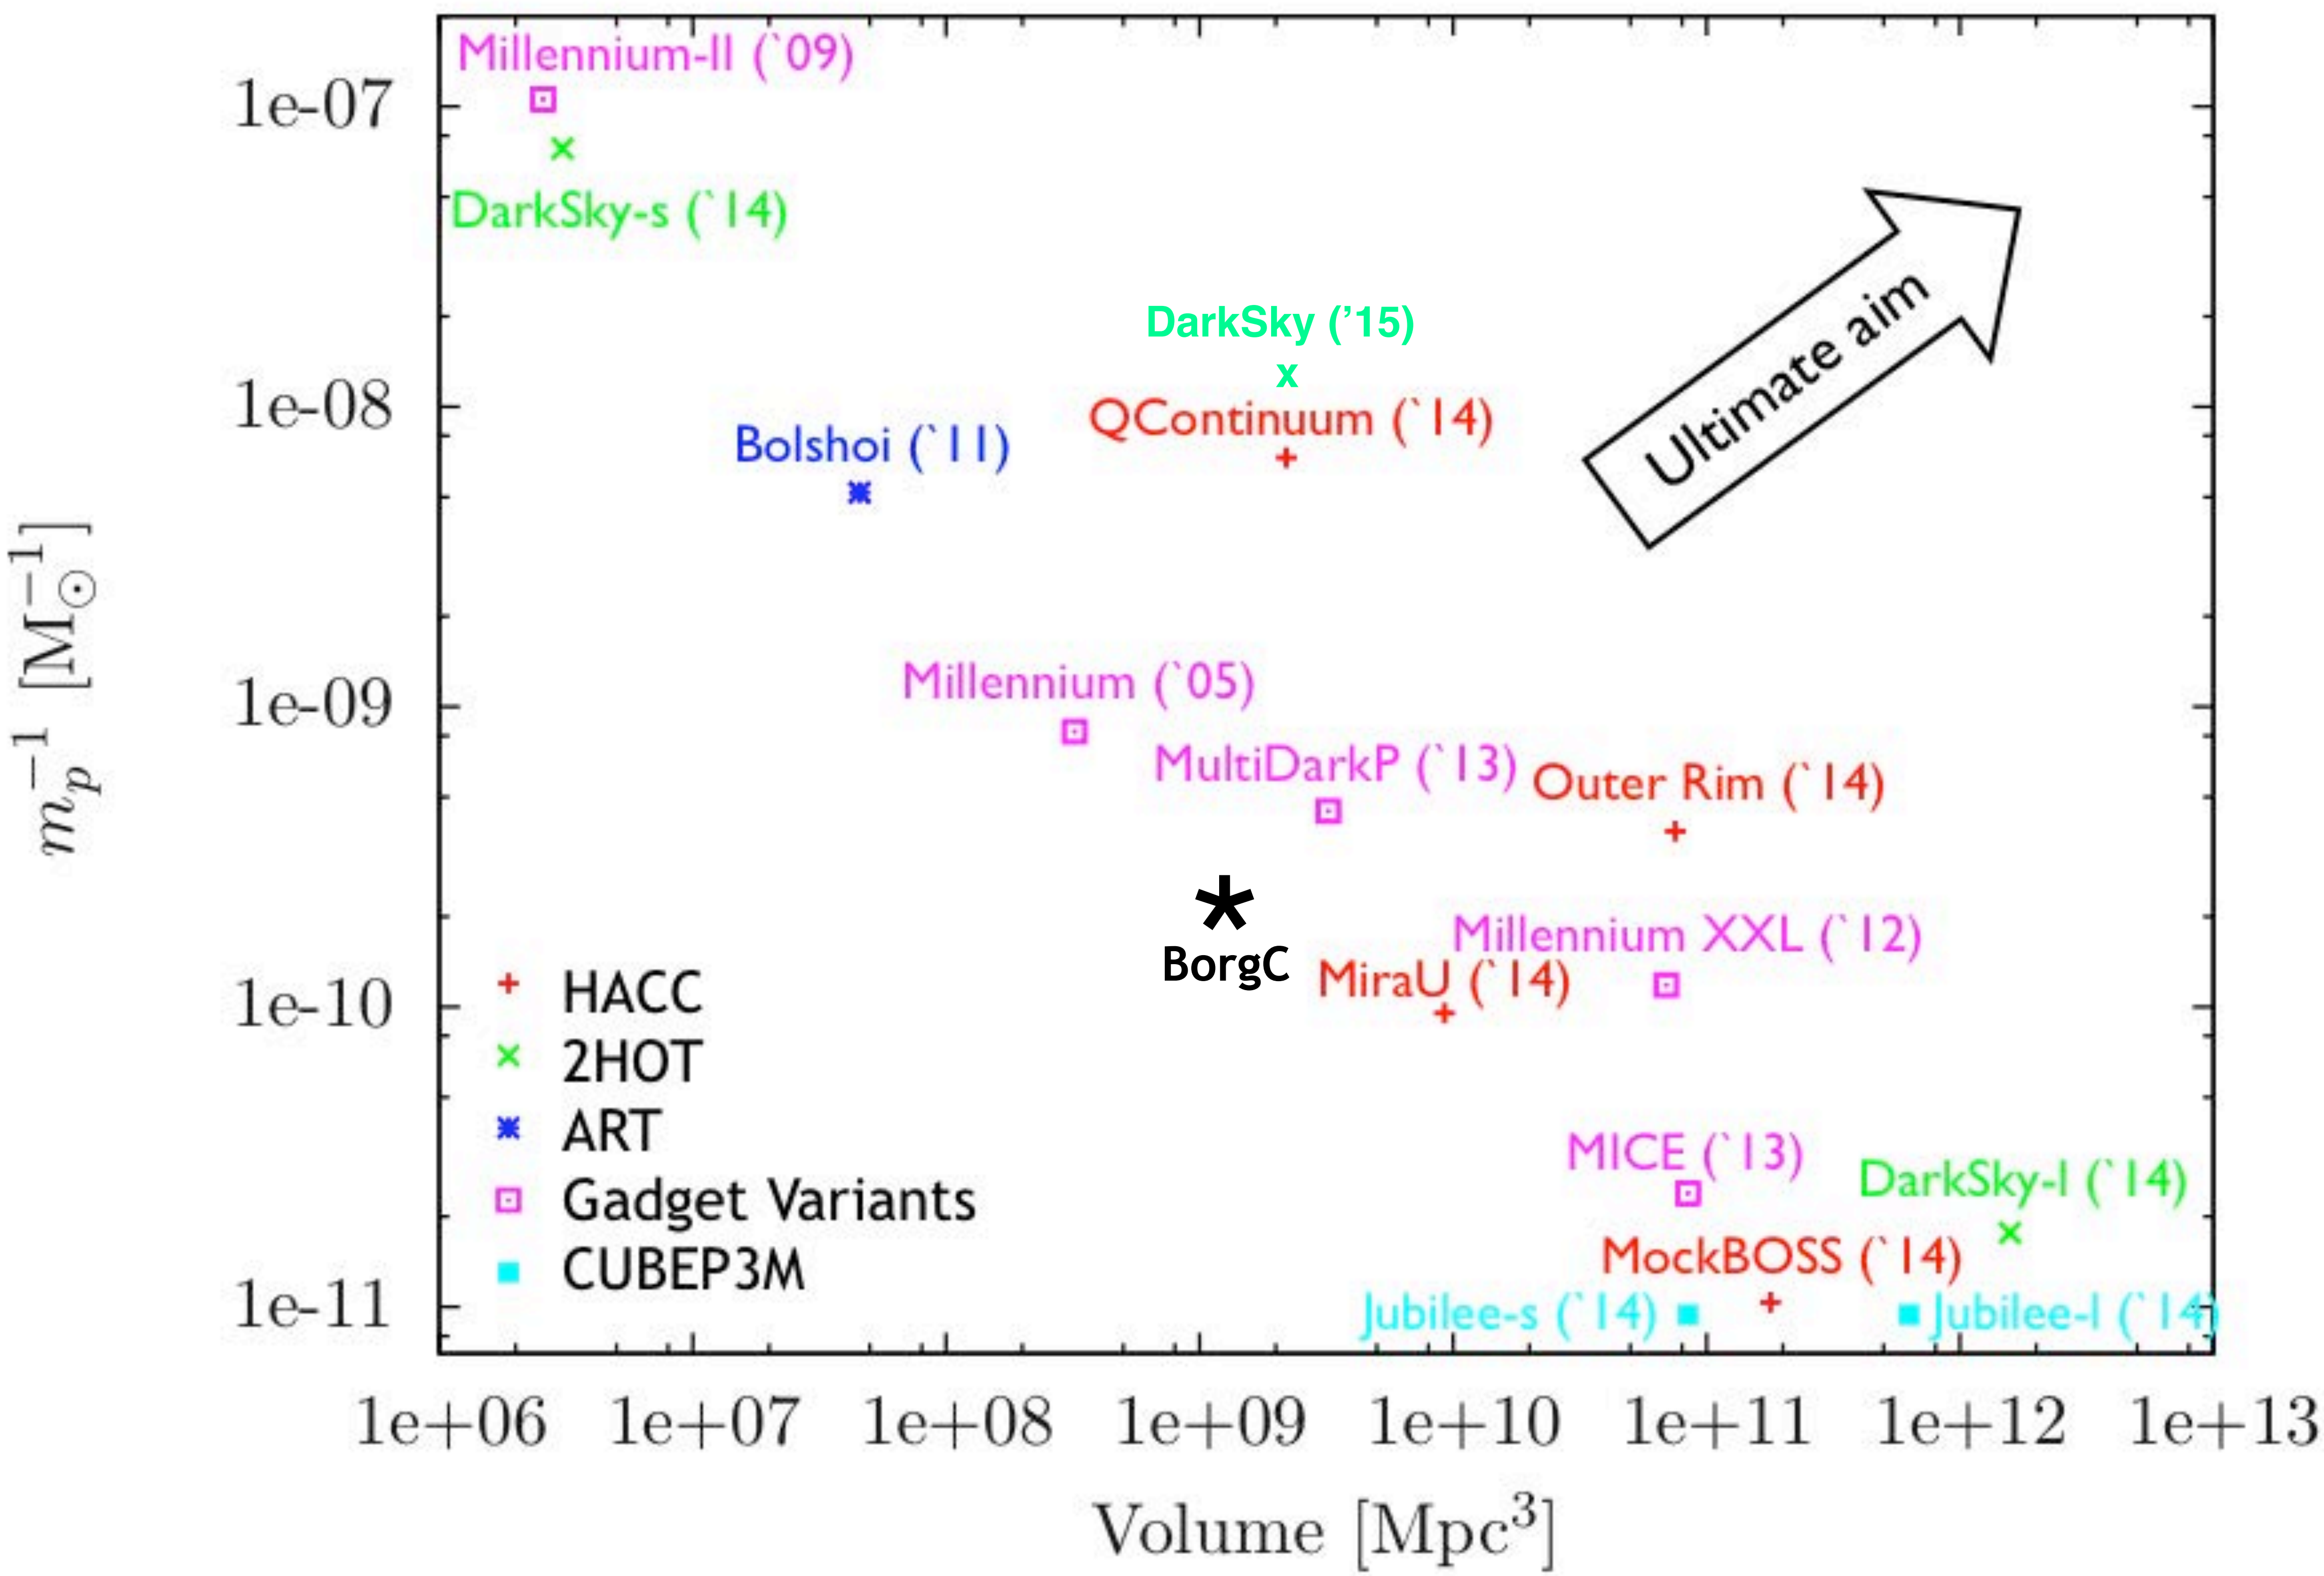

Current State-of-the-Art Simulations

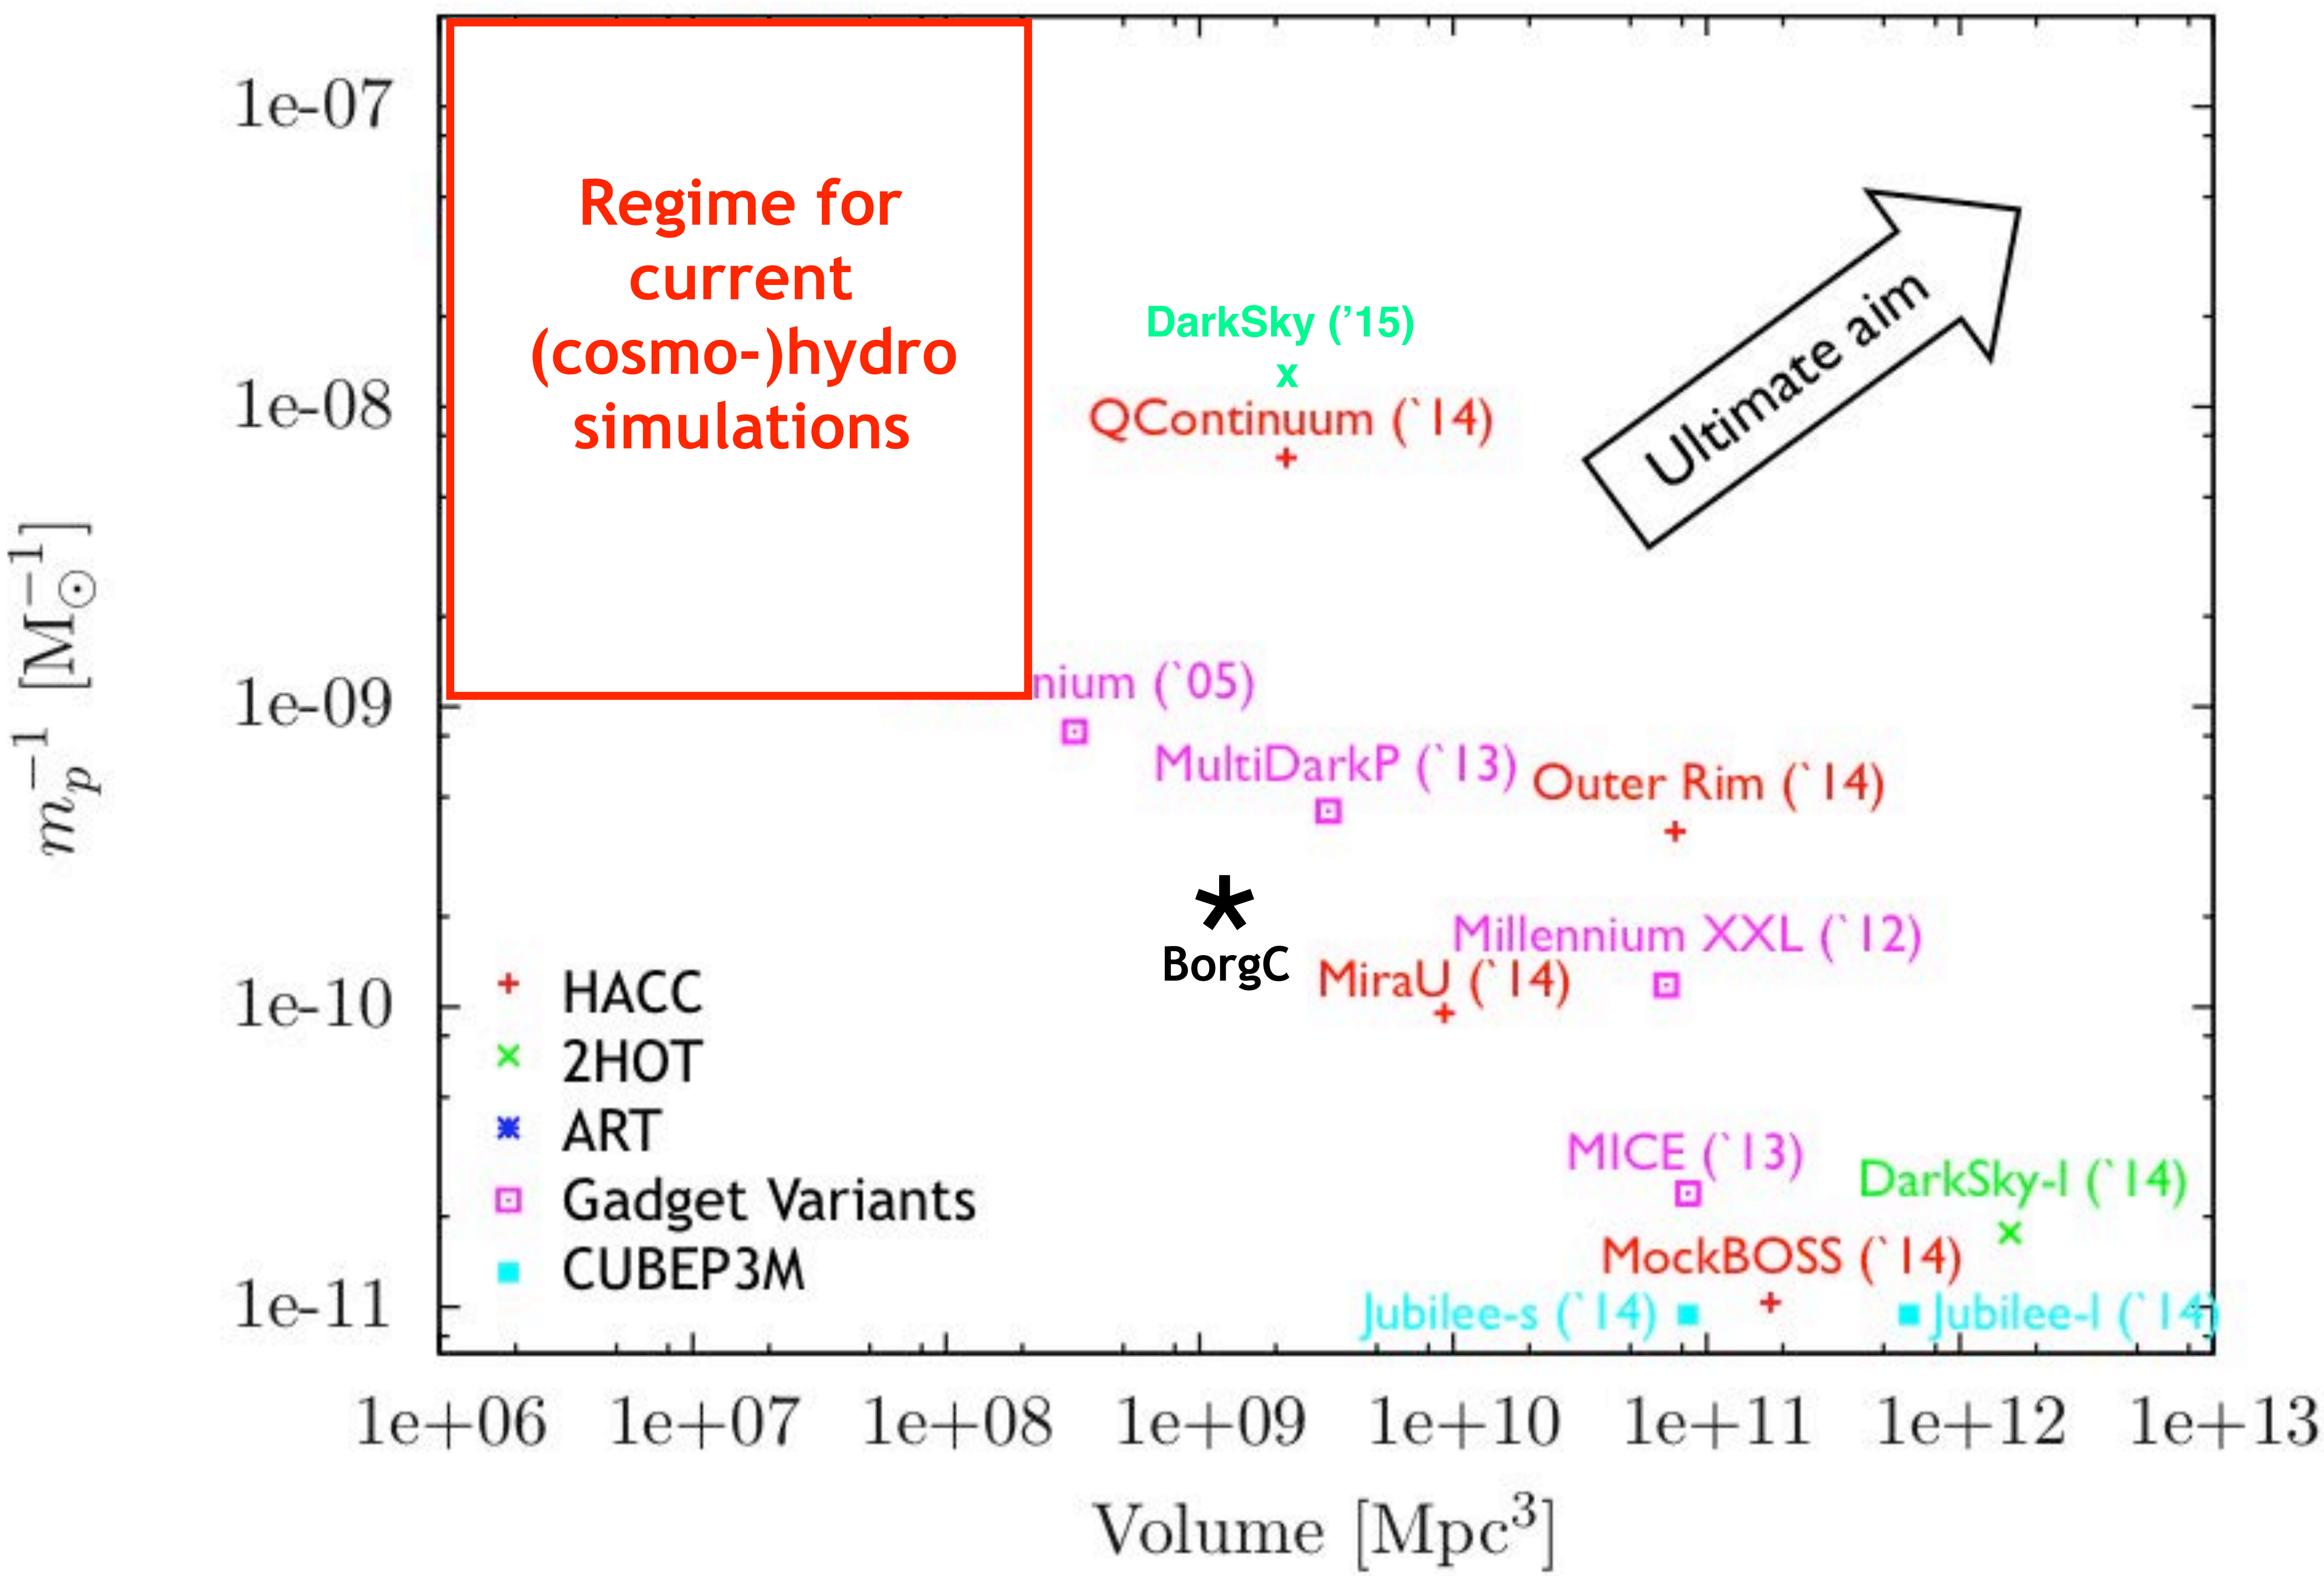

Current State-of-the-Art Simulations

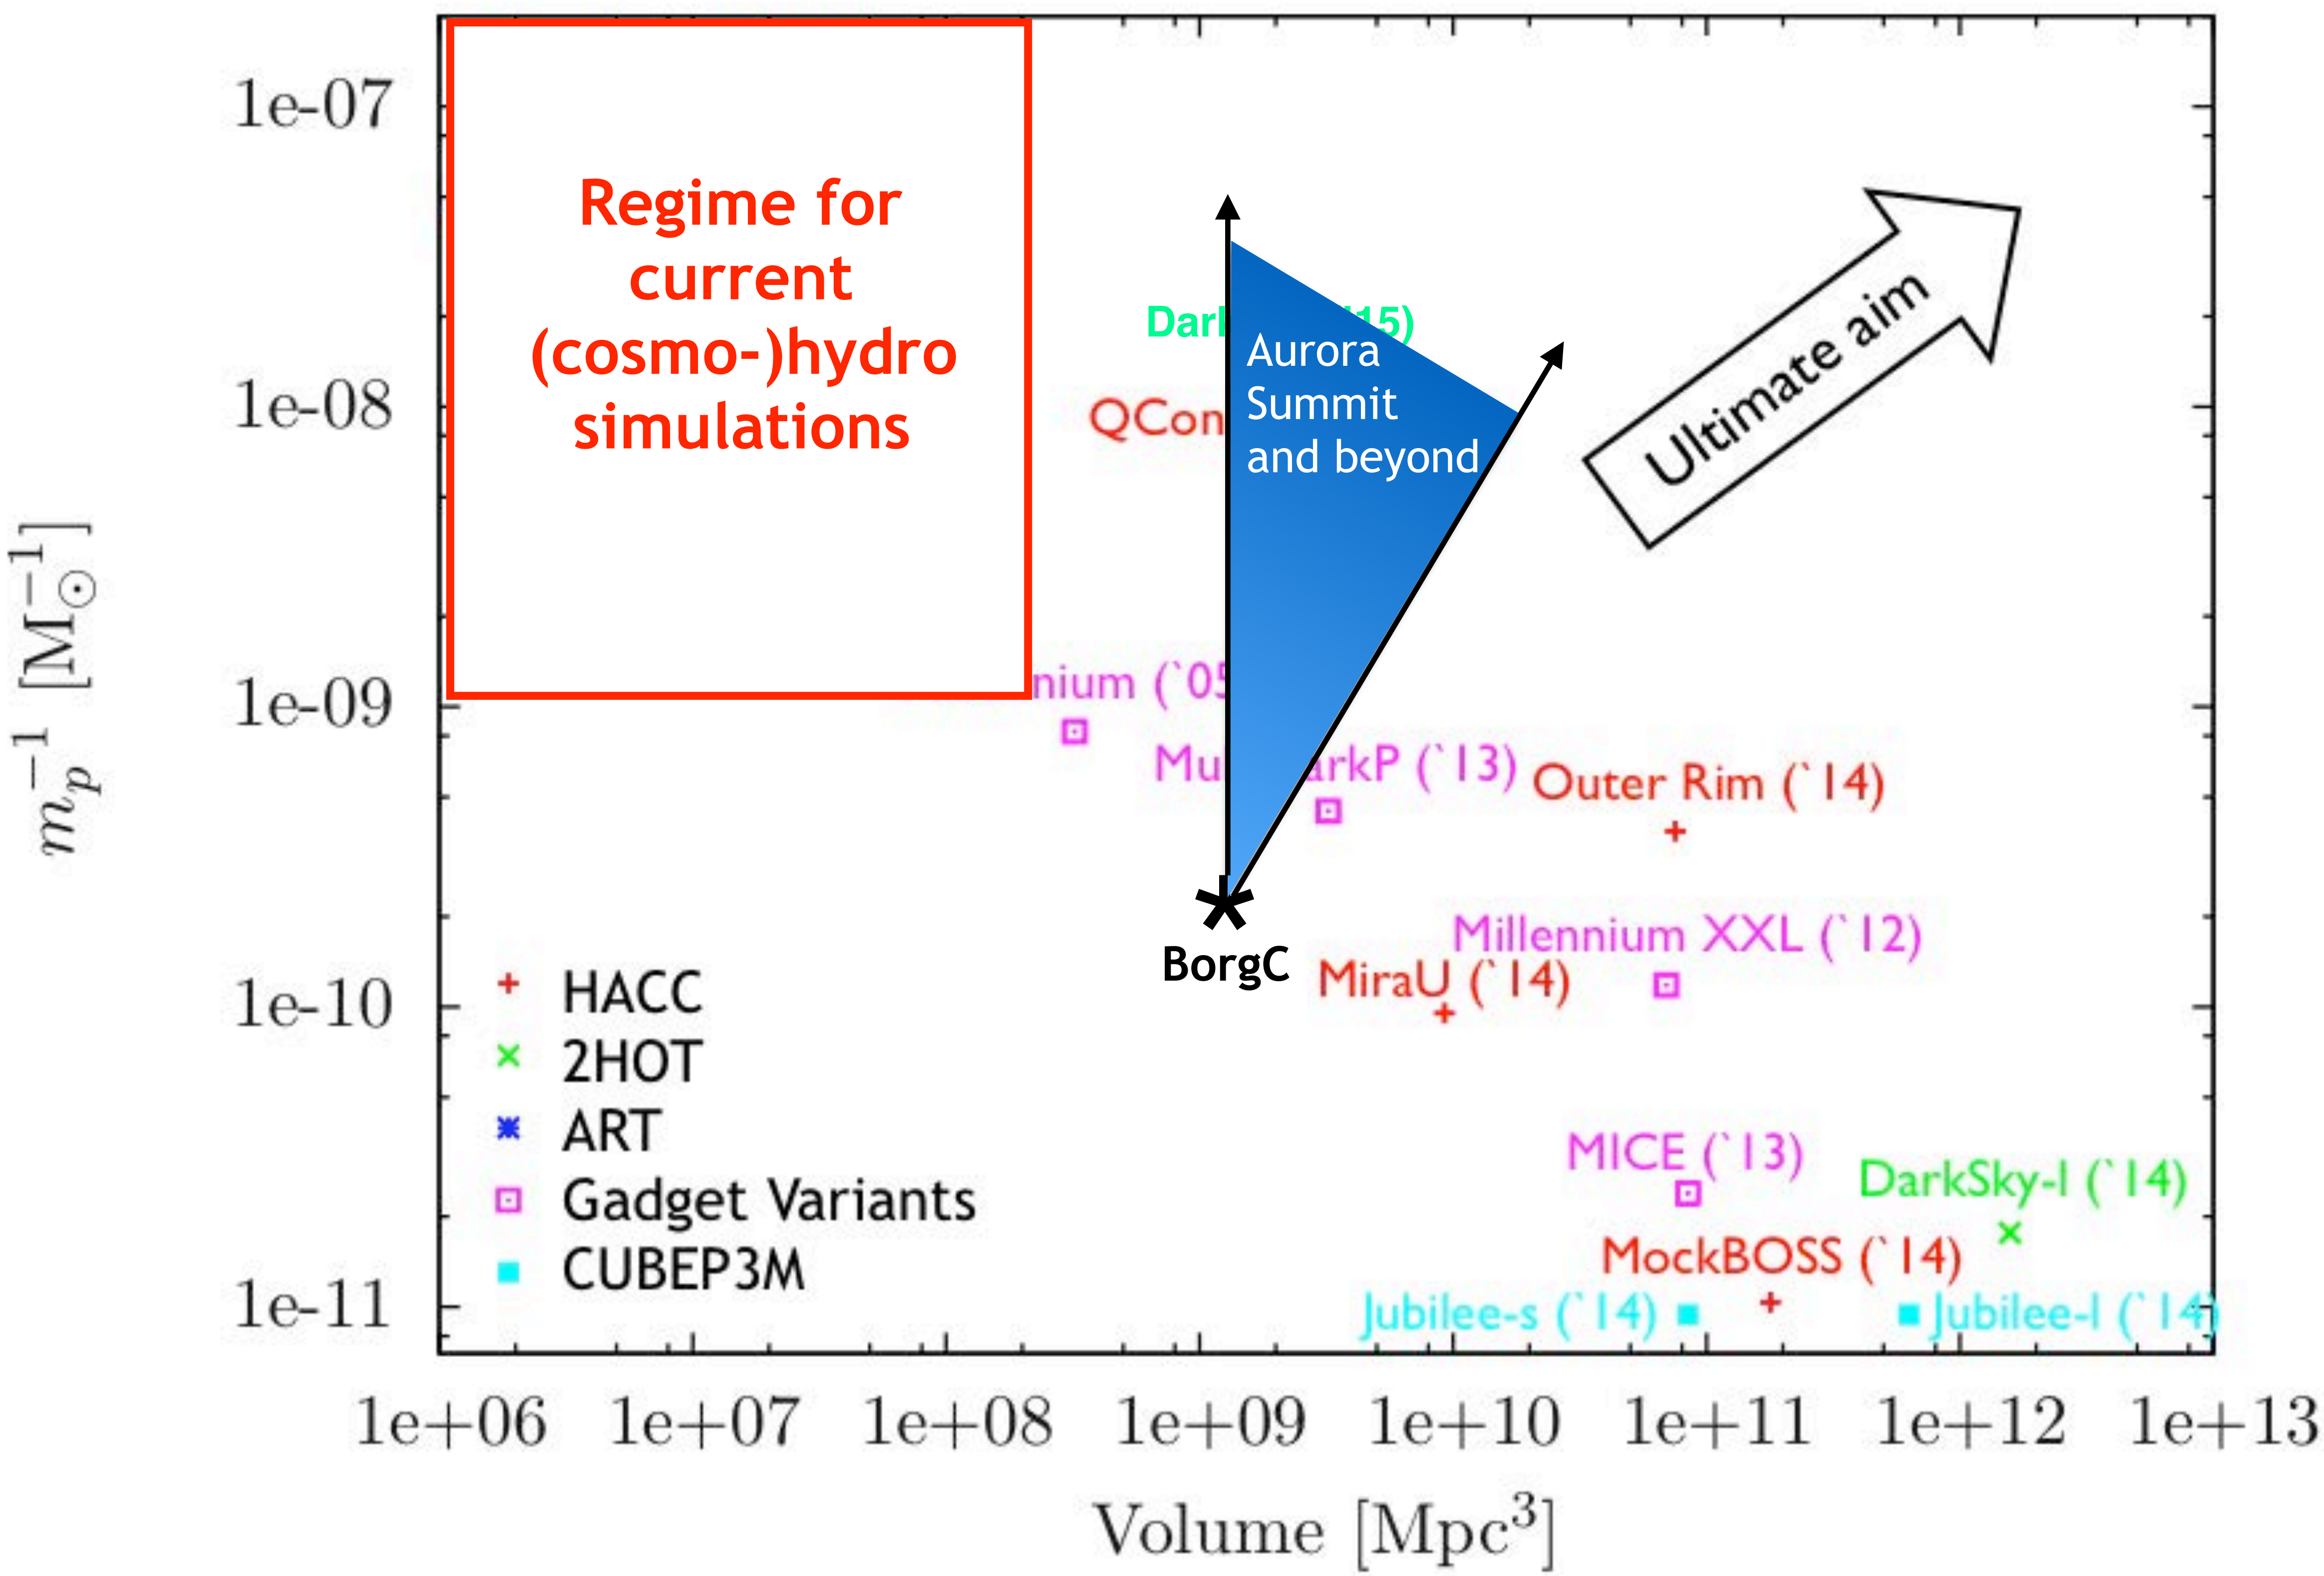

CRK-HACC: Tackling the Baryons

- HACC: Assume that the baryons follow the dark matter and the physics of the baryons is negligible on the scales of interest
- Now: Model the baryons separately as adiabatic, compressible fluid described by the Euler fluid equations; describe physics on very small scales via sub-grid models
- Range of new variables (density, pressure, temperature, …) increases memory requirements considerably as well as computational cost
- HACC is particle-based code, therefore use Smoothed Particle Hydrodynamics (SPH) approach with new developments (Conservative reproducing kernel) to overcome traditional SPH short-comings

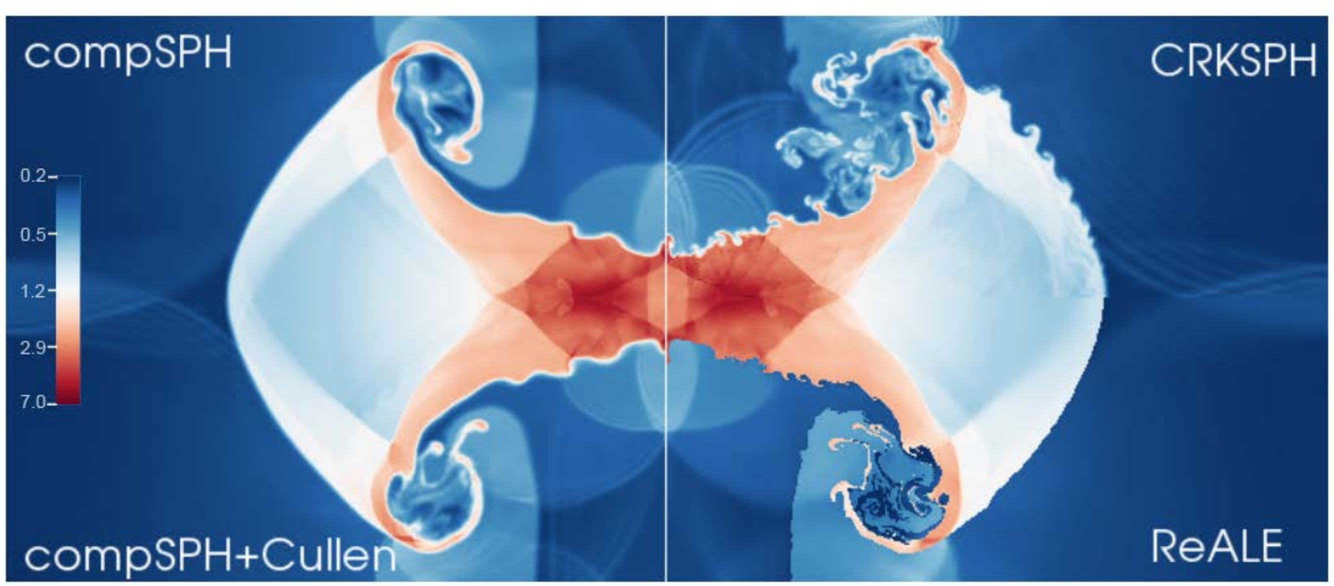

Triple-point shock test with CRKSPH

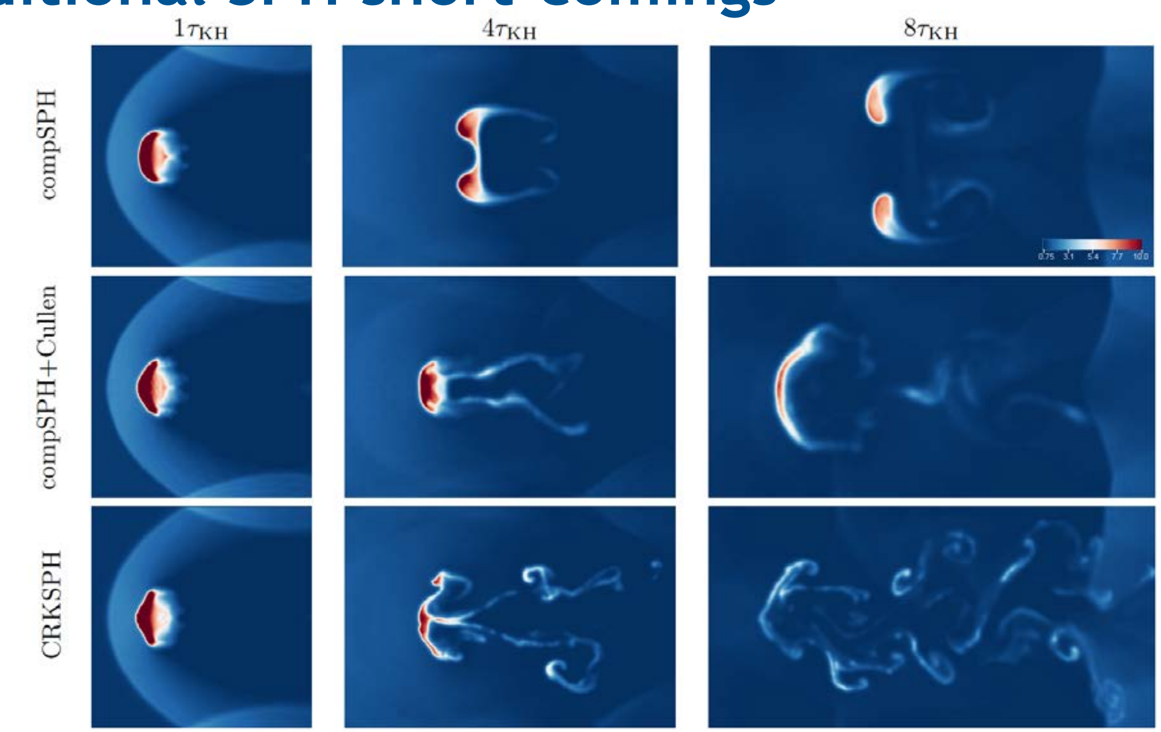

Blob test with CRKSPH

CRK-HACC: Performance on the KNL, PRELIMINARY!

Cores	RPN	OMP	TH	Gravity [Time/s]	CRK-HACC [Time/s]	Ratio
16	4	4	16	616.33	21881	35.5
16	4	8	32	543.73	18598	34.2
16	4	16	64	530.23	19885	37.1
16	8	2	16	544.75	21750	39.9
16	8	4	32	459.53	**18548**	40.3
16	8	8	64	**437.21**	20026	45.8
16	8	4	64	468.50	21339	45.4

RPN: MPI ranks for each run; **OMP:** OpenMP threads per MPI rank, **TH:** total number of threads = RPN*OMP

- **Problem set up: 2x320^3 particles, 320^3 grid, 3 full time steps, 5 sub-cycles**
- **Results for flat mode and cache mode are essentially identical**

CRK-HACC: Strong Scaling on the KNL, PRELIMINARY!

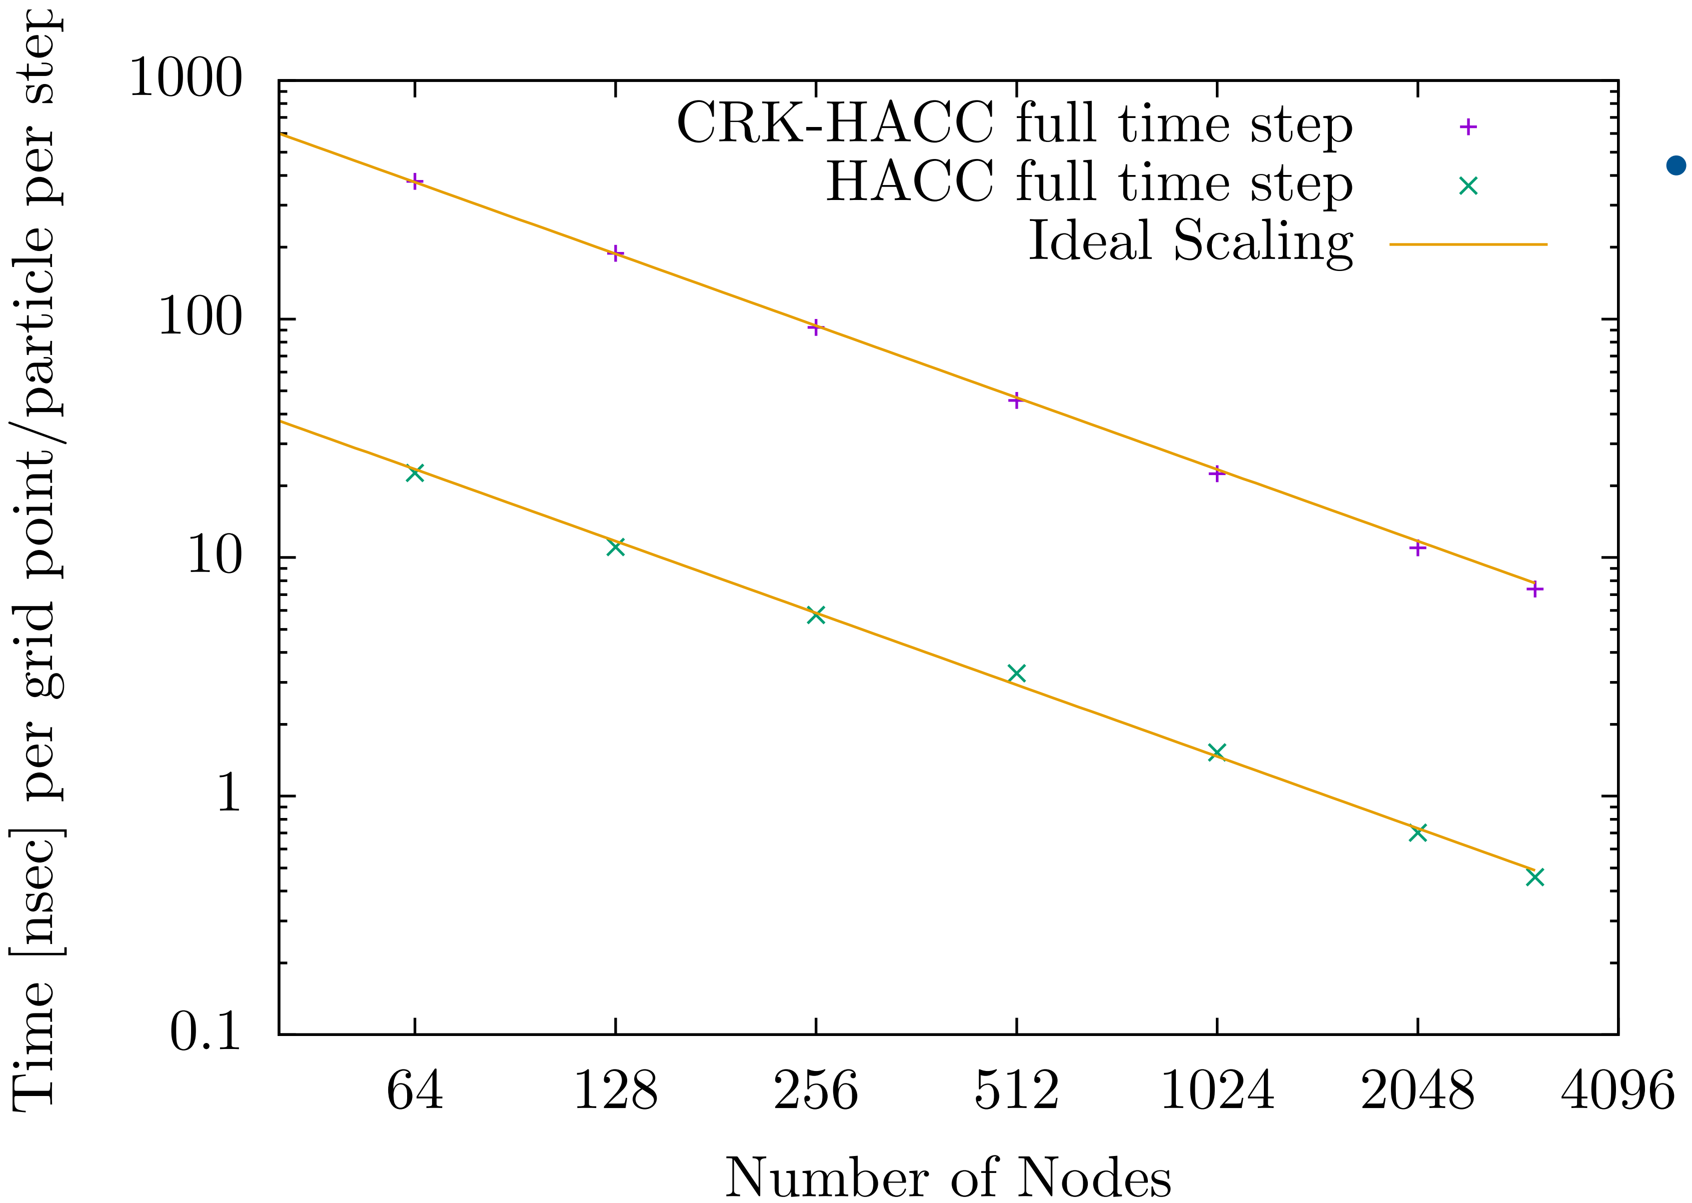

- **Problem set up:**
  - 2048^3 particles on 64 to 2048 nodes (and 2304^3 particles on 3072 nodes)
  - 8 ranks per nodes, 16 threads
  - Cache-quad mode
  - Per particle CRK-HACC ~15 times slower

CRK-HACC: The Santa Barbara Cluster

- Standard test for cosmological hydro codes, first carried out in 1999 (Frenck et al.)

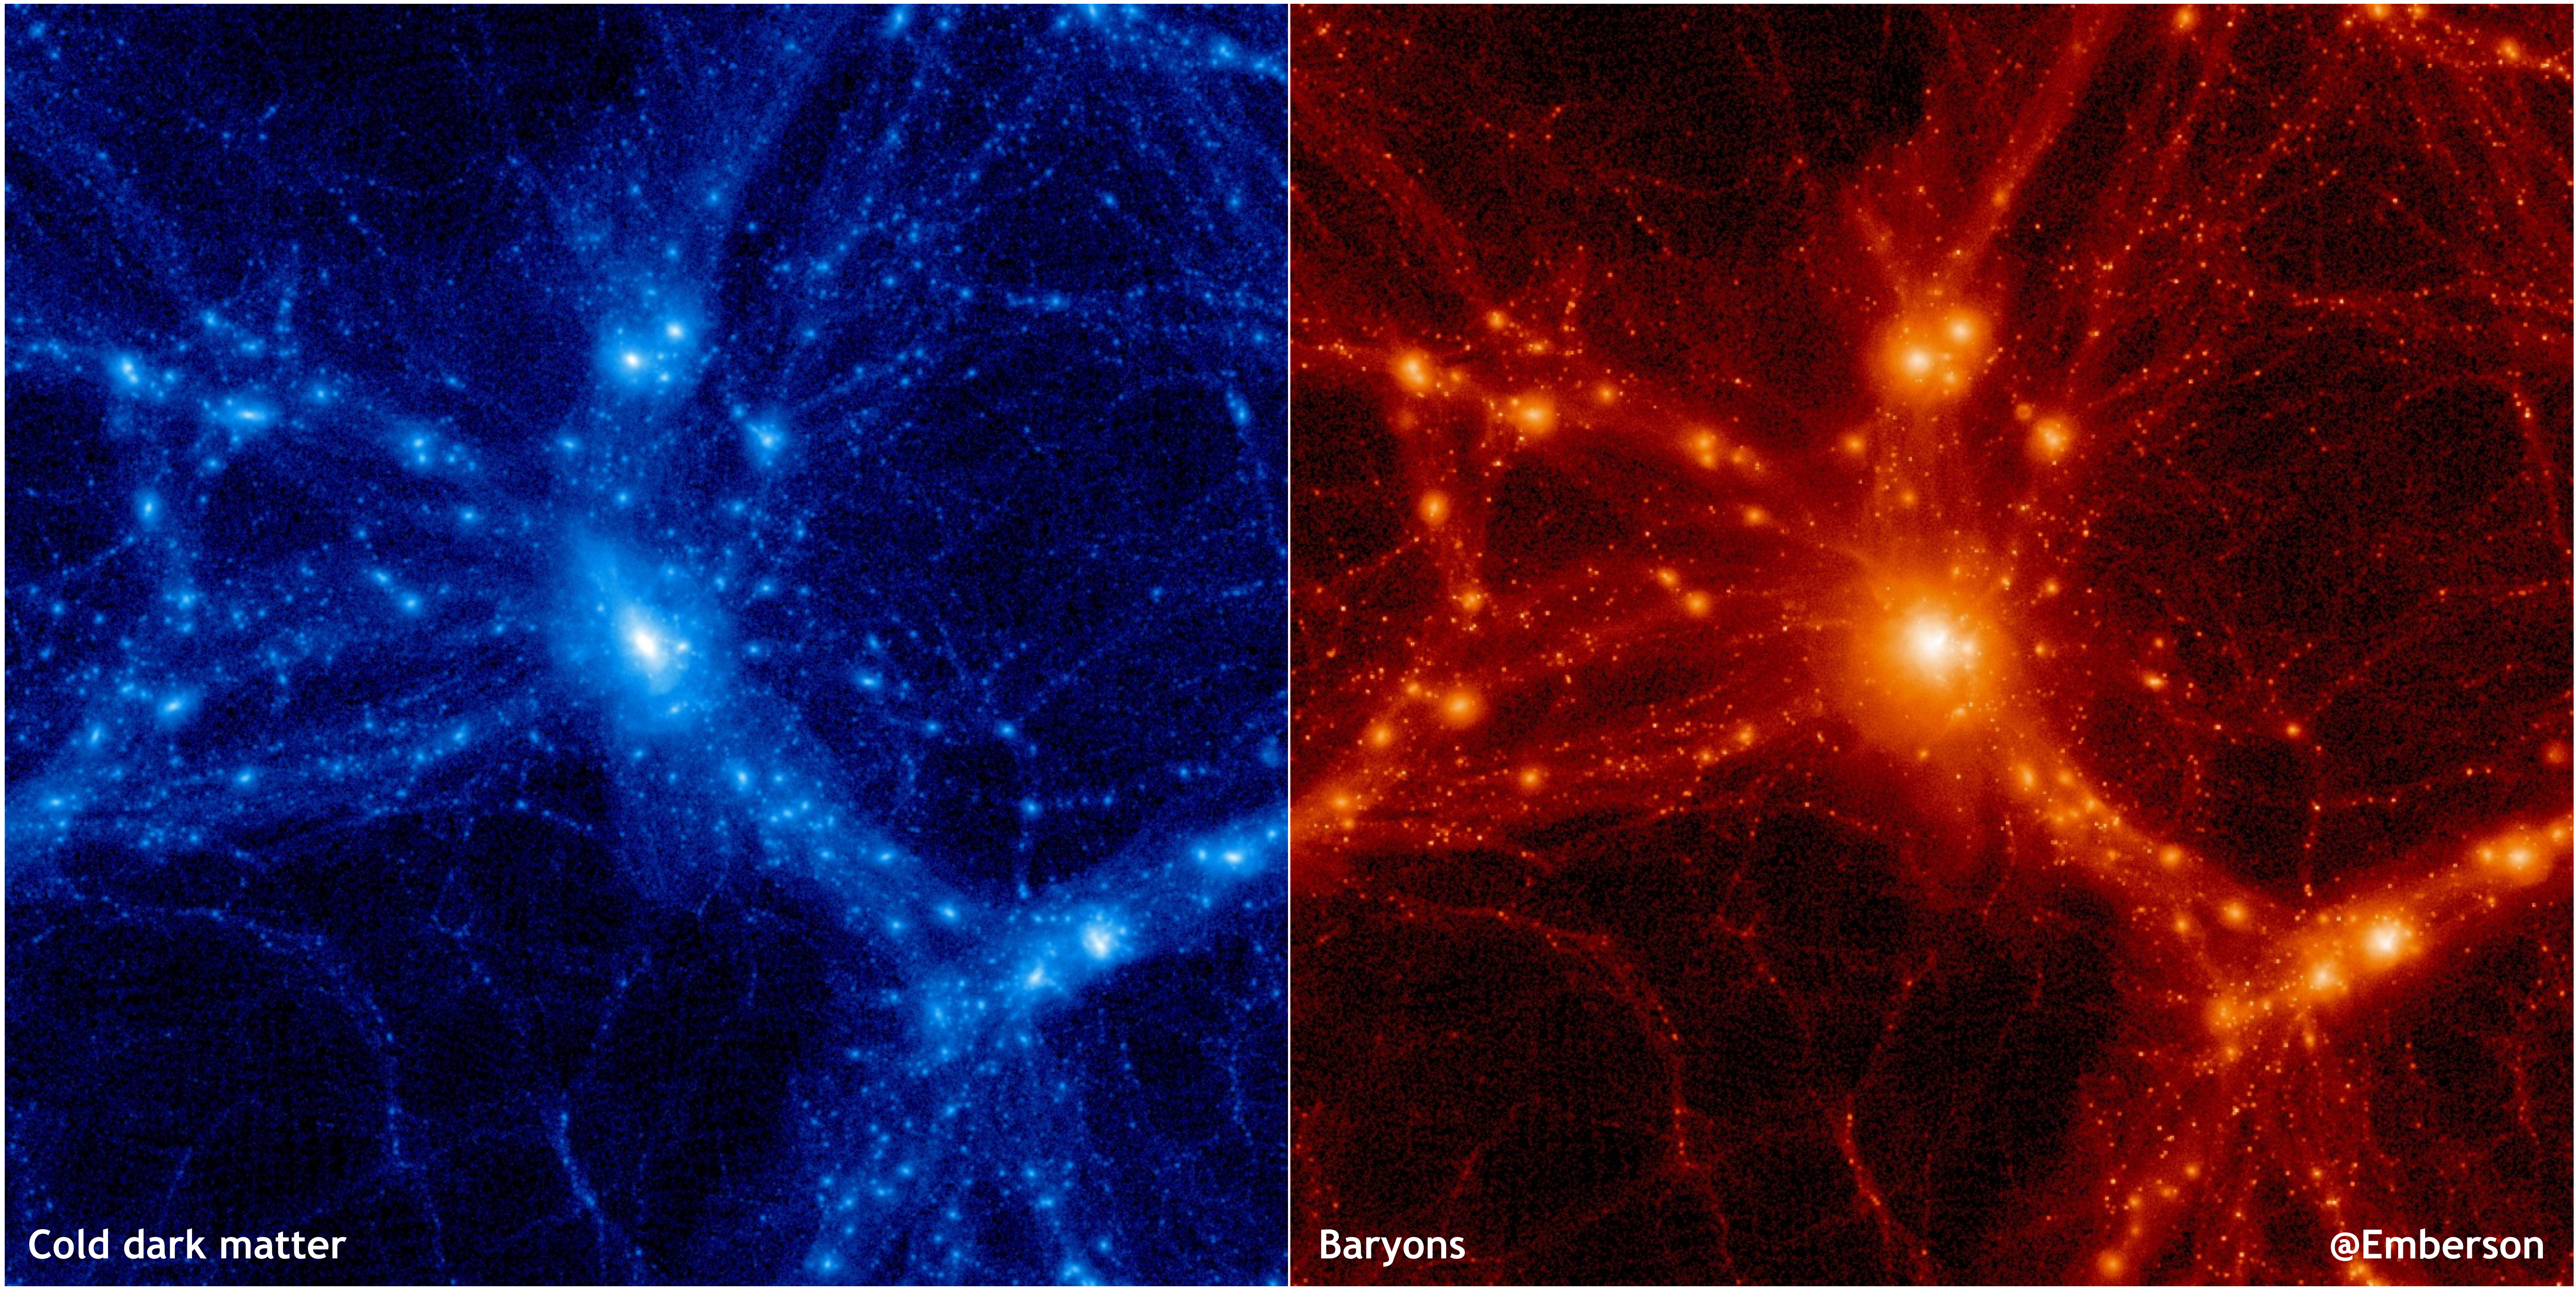

Density slices through the full simulation volume

Final Thoughts

- HACC was designed from the start to allow easy portability between a diverse set of architectures
- Large portions of the code base (~95%) of HACC remain unchanged when moving from one machine to another; 5% are optimized in several ways (algorithms, implementation, parameters) to achieve high performance on all architectures
- HACC had been optimized for Mira already, port to Theta was straightforward
- HACC weak and strong scales on the full machine, I/O could be better
- First science run successfully carried out, currently being analyzed
- With new generation of supercomputers around the corner, we started adding new baryonic physics to HACC, now have CRK-HACC
- CRK-HACC scales also well, still needs to be optimized (implemented adaptive time stepping) and some small issues ironed out
- One specific “user” comment: the submission system could be improved (many different variables, all ALCF machines have different syntax now, and it’s different from OLCF and NERSC … it’s really difficult to keep track of all of them!)

Computing the Sky at Extreme Scales

- **Why SPH?**
  – Galilean Invariant, intrinsically adaptive, agnostic to surface geometries, conservative, easily parallelizable/ scalable, particle/interpolant connectivity is mutable, enabling more accurate modeling of extreme material deformations
- **Why not?**
  – E0 Error: cannot reproduce even a constant field, becomes increasingly troubling over density discontinuities.
  – Overly diffusive Artificial Viscosity (AV)
- **CRKSPH (Conservative Reproducing Kernel SPH)**
  – Utilizes Linear Reproducing Kernels (improves E0)
  – Novel Limited Viscosity (improves AV diffusivity)
  – Conservative Formalism (maintains conservation even though the kernels are no longer symmetric

Triple-point shock test with CRKSPH

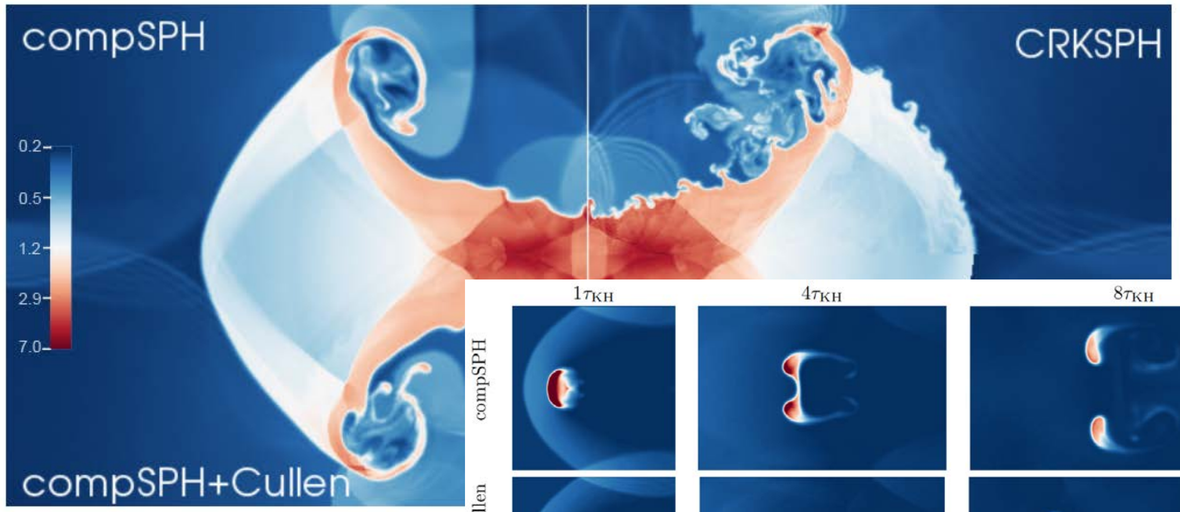

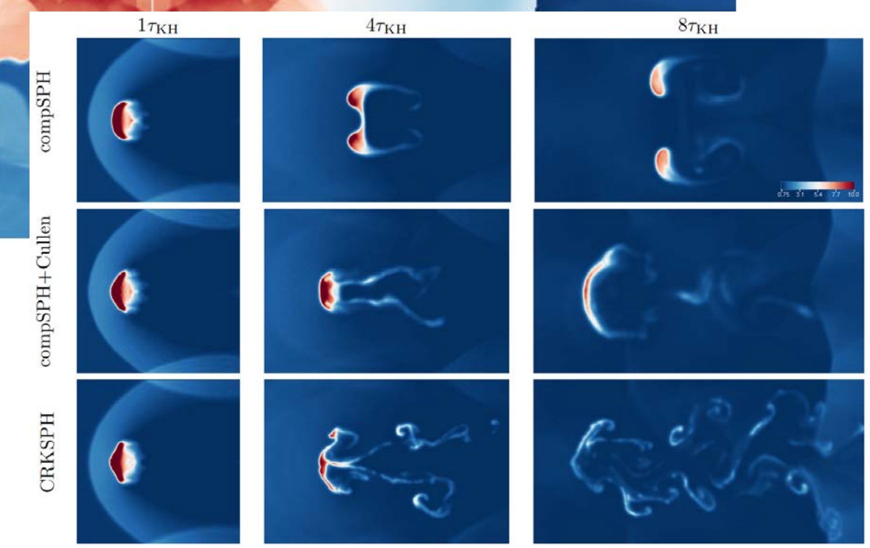

Near-term ECP project

Blob test with

Frontiere, Raskin, Owen, J. Comp. Phys. (2017); arXiv:1605.00725 [physics.comp-ph]