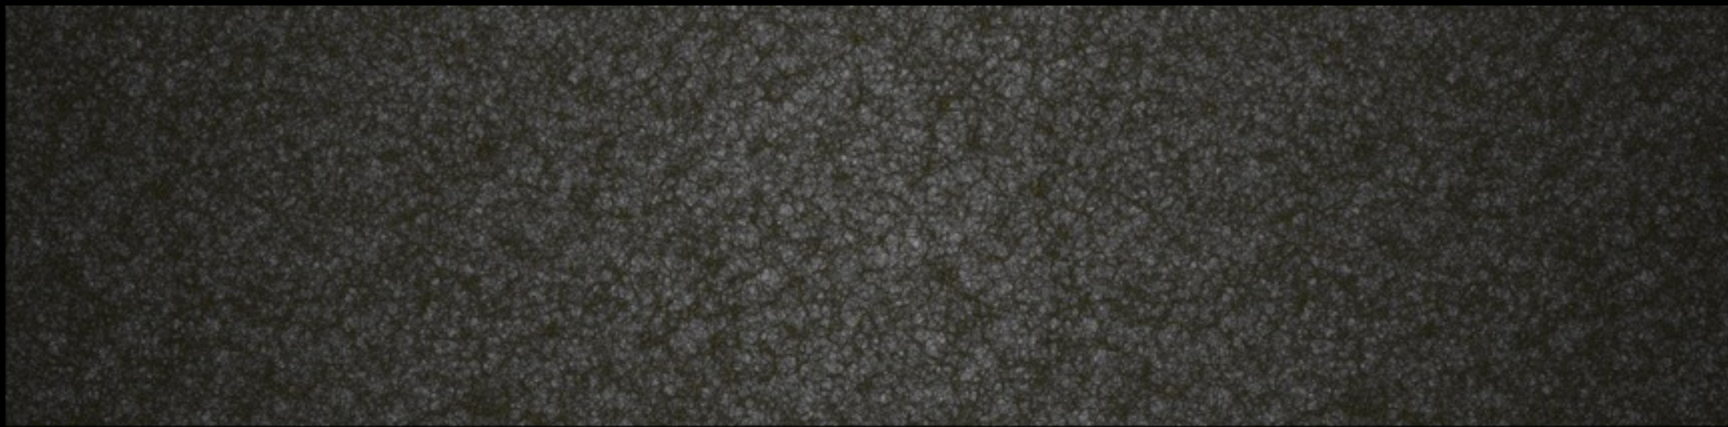

Simulating our Universe:

Leveraging present and future Supercomputers

Project:
2018184475

Joachim Stadel
Institute for Computational Science
University of Zurich

Universität Zürich[UZH]

COBE

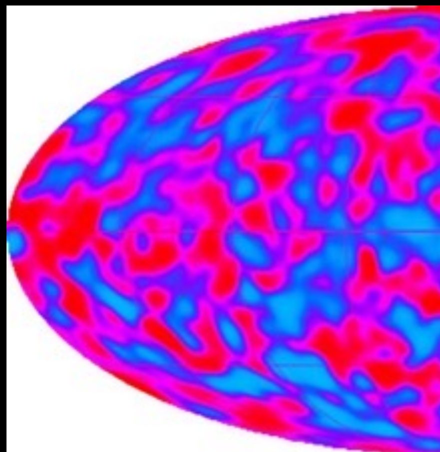

The Cosmic Microwave Background as seen by Planck and WMAP

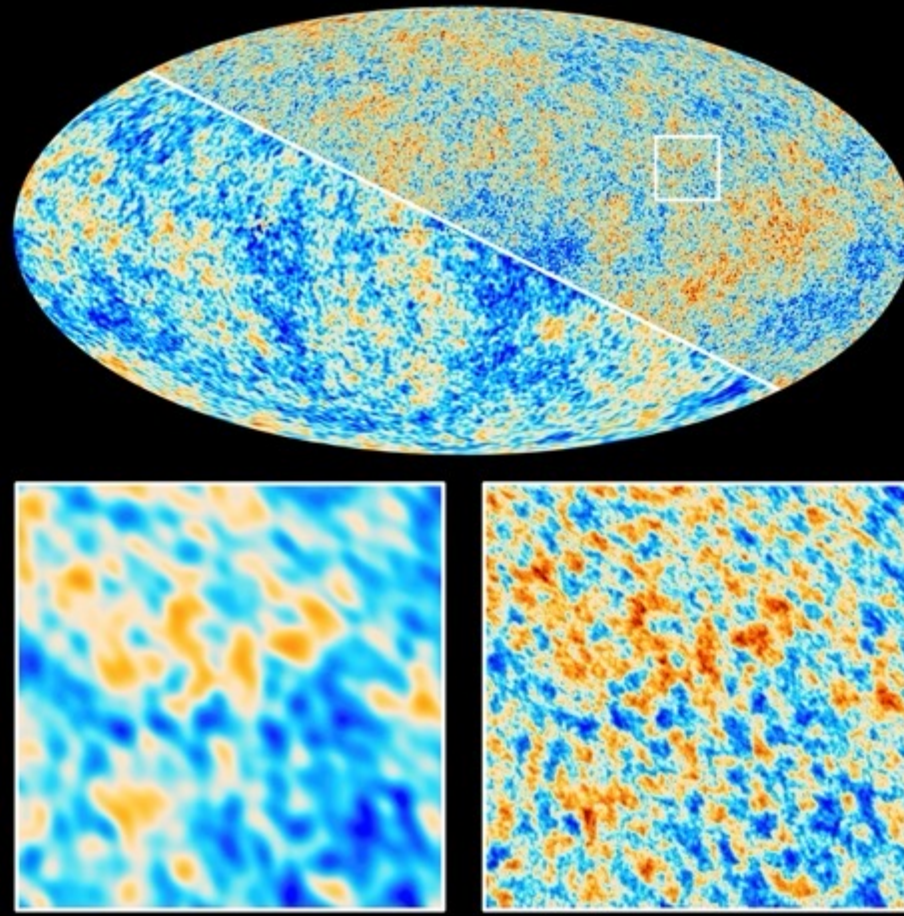

WMAP

Planck

The CMB fluctuations brought in the era of *precision cosmology*

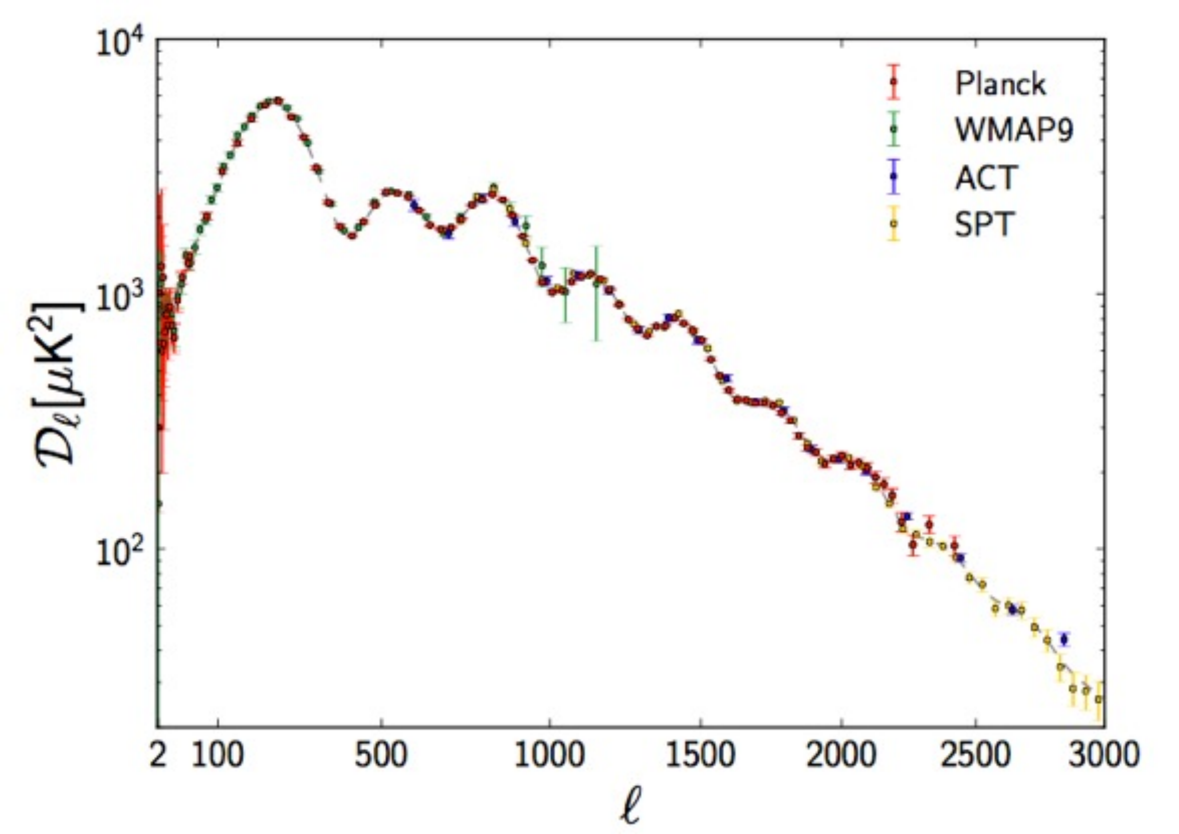

CMB fluctuations tell us about one early epoch.
(indirectly there are ways to get at other epochs too)

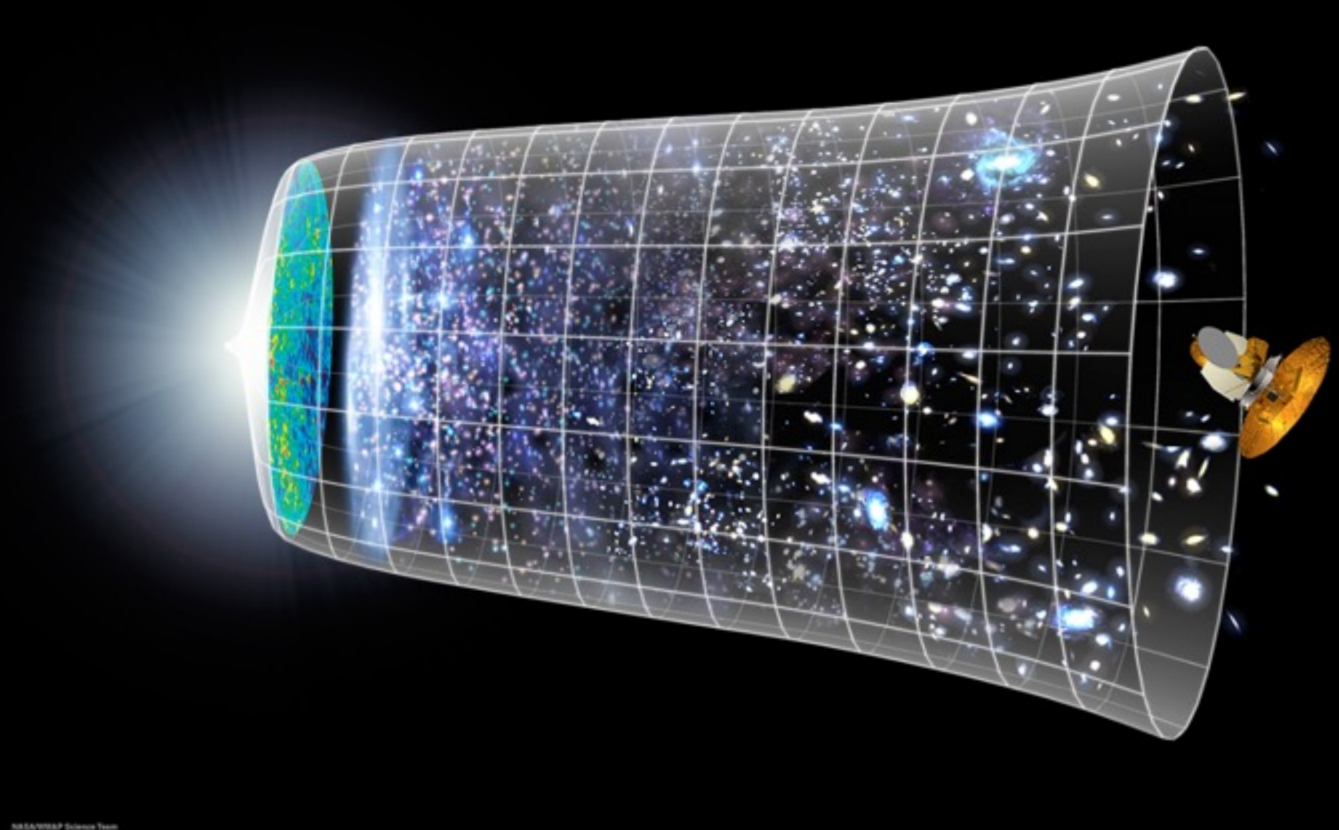

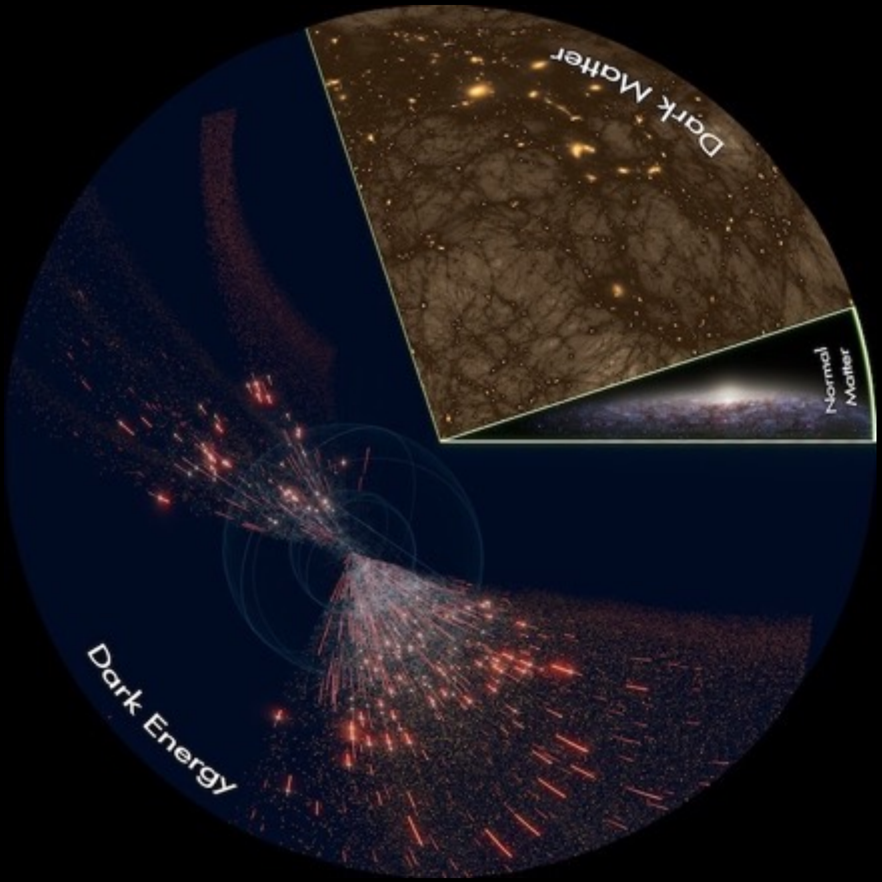

Dark Matter
Normal Matter
Dark Energy

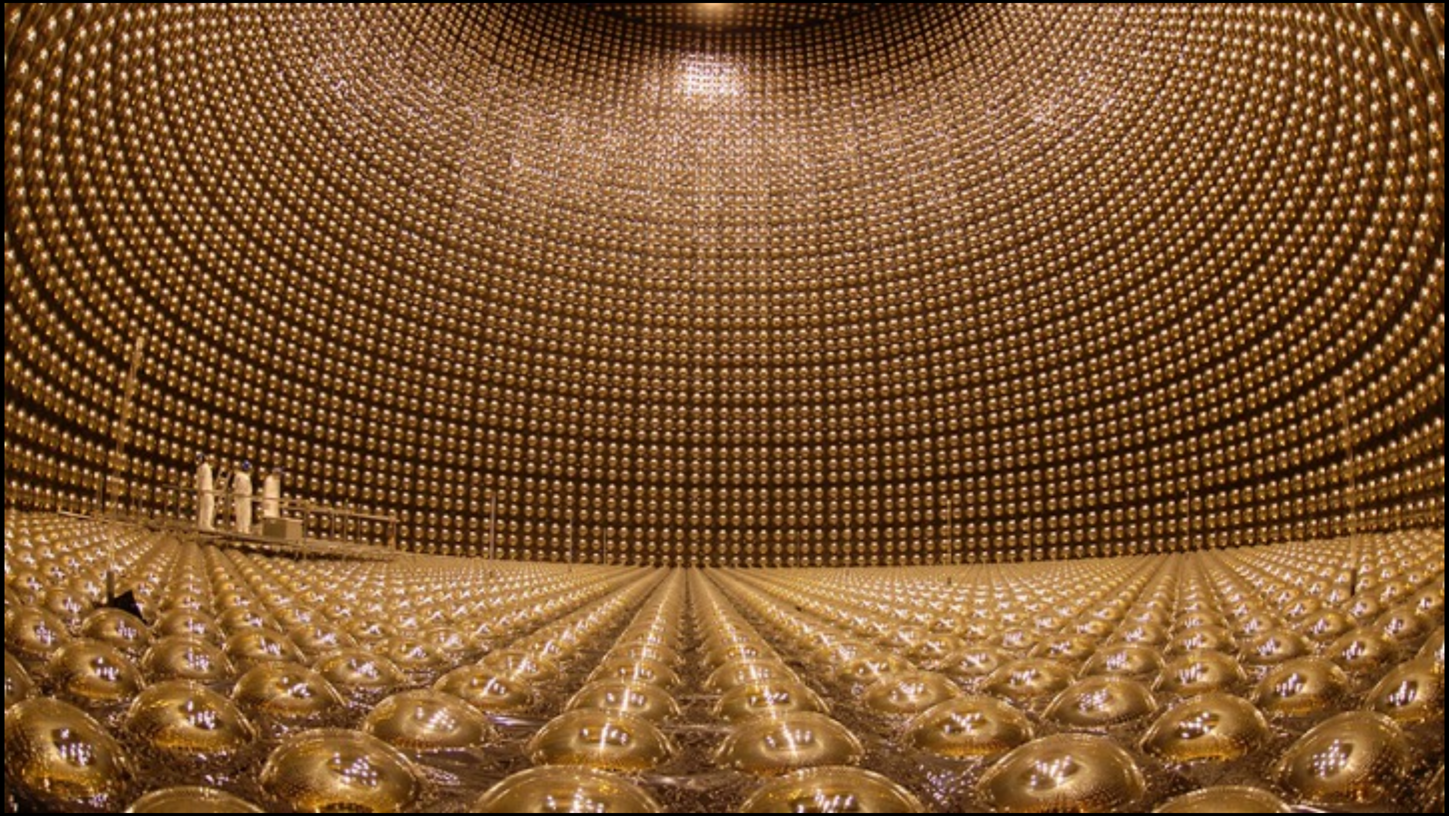

Euclid Mission 2021-2027 >10 billion galaxy images (photo-z) >10 million redshifts (Hα)

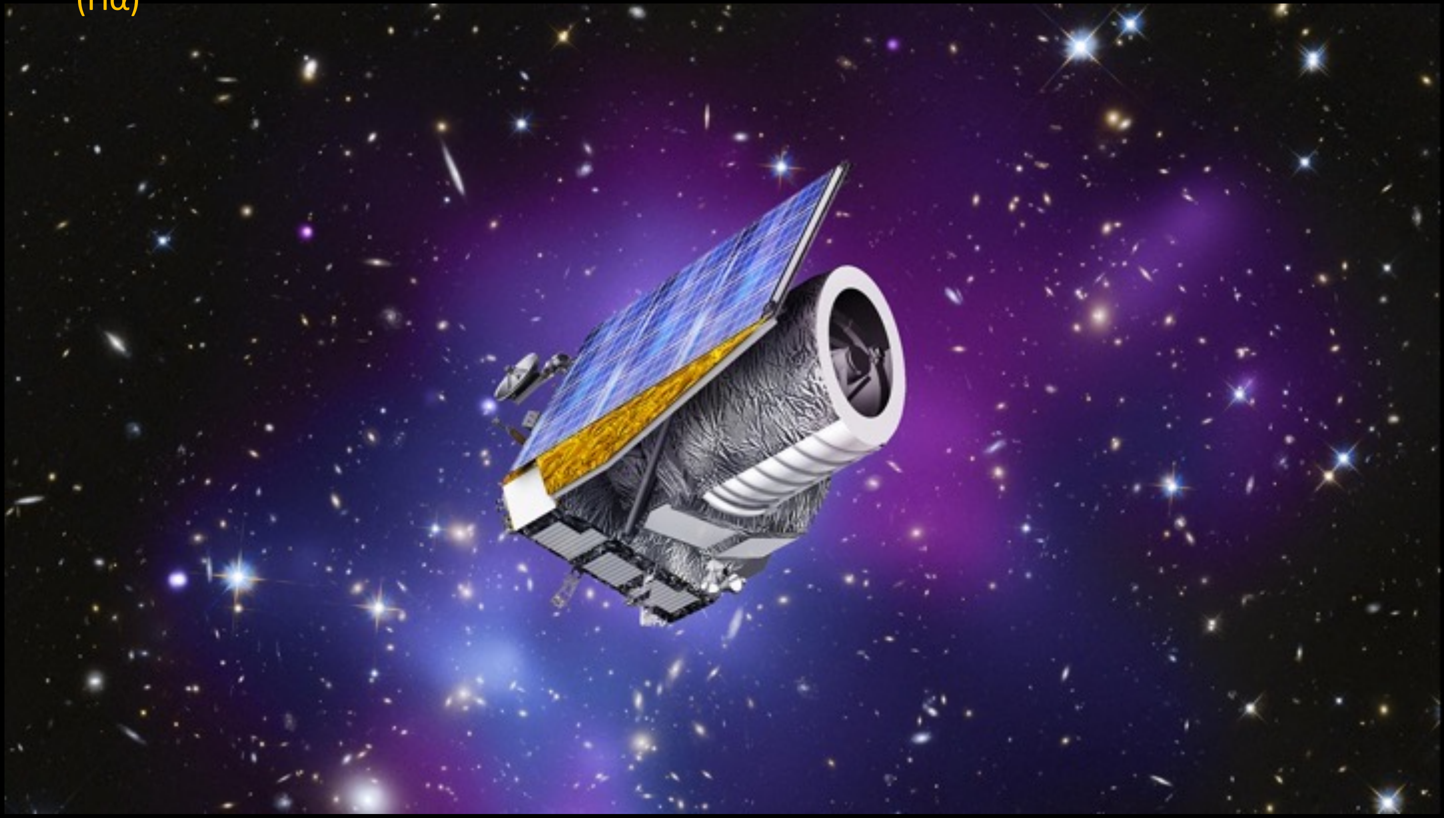

15 000 square degrees, most of the sky above the plane of the Milky Way

Hubble Deep Field image was taken over a very small fraction of the sky.

Euclid will cover all of the sky at something approaching this level of resolution!

Abel 370

DEFLECTION OF LIGHT RAYS CROSSING THE UNIVERSE, EMITTED BY DISTANT GALAXIES

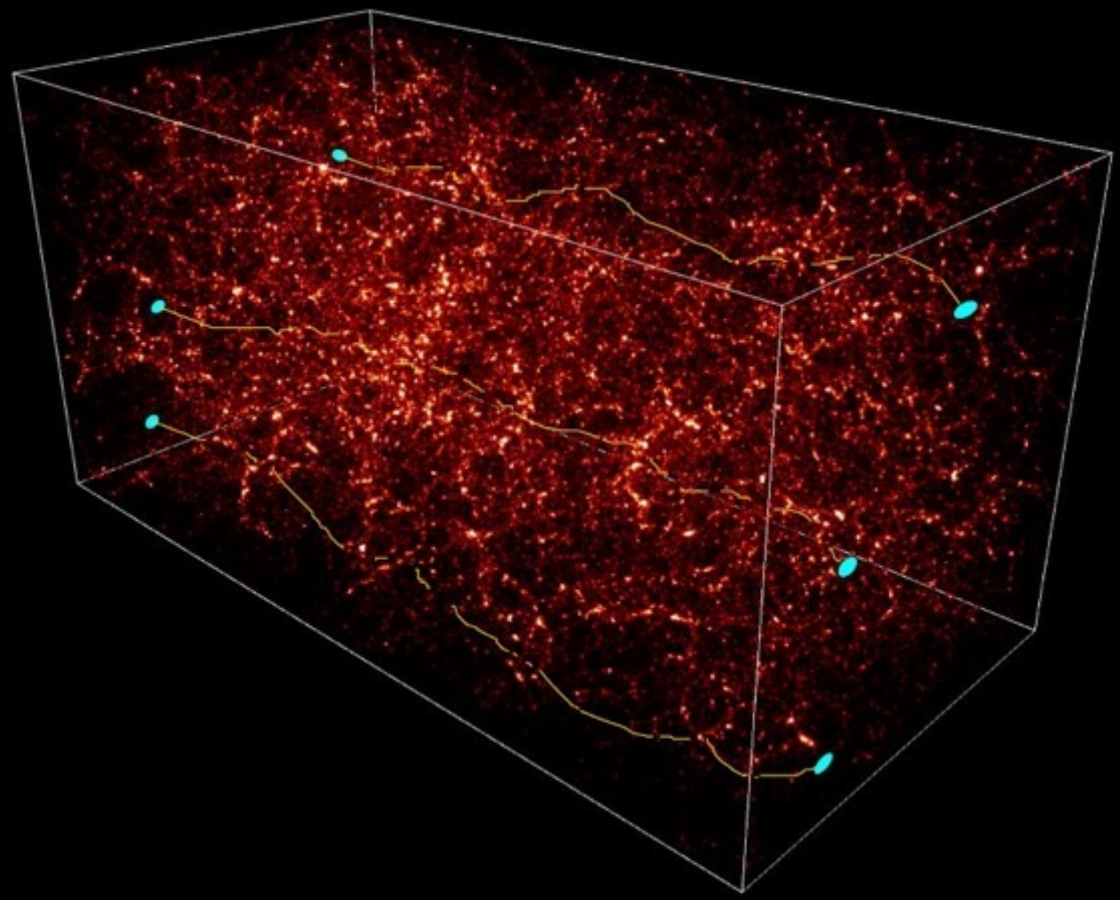

SIMULATION: COURTESY NIC GROUP, S. COLOMBI, IAP.

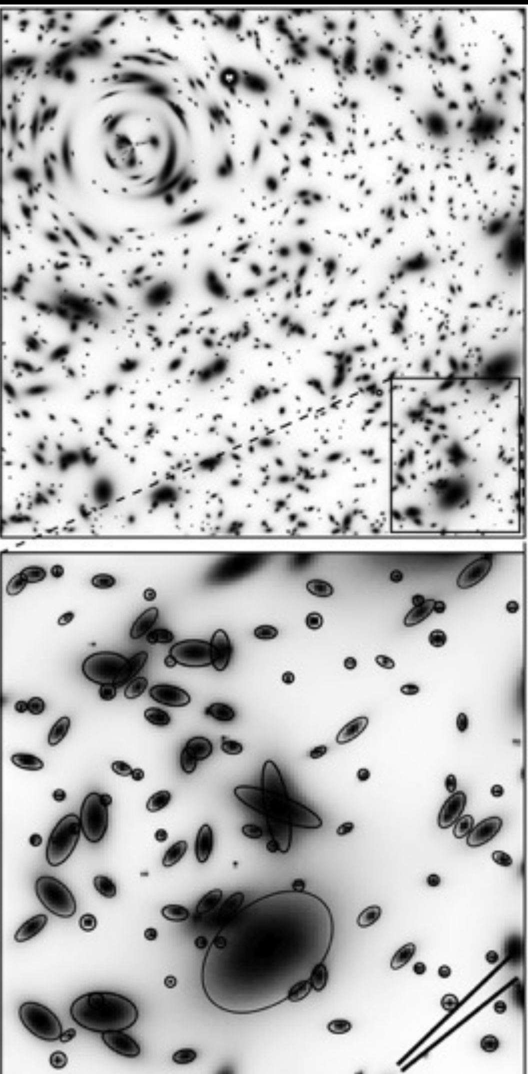

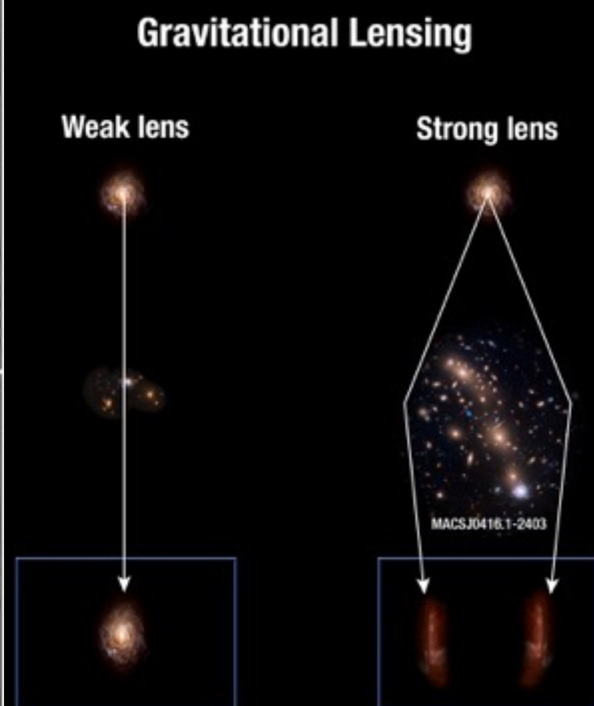
Gravitational Lensing
Weak lens
Strong lens
MACSJ0416.1-2403


Abel 370

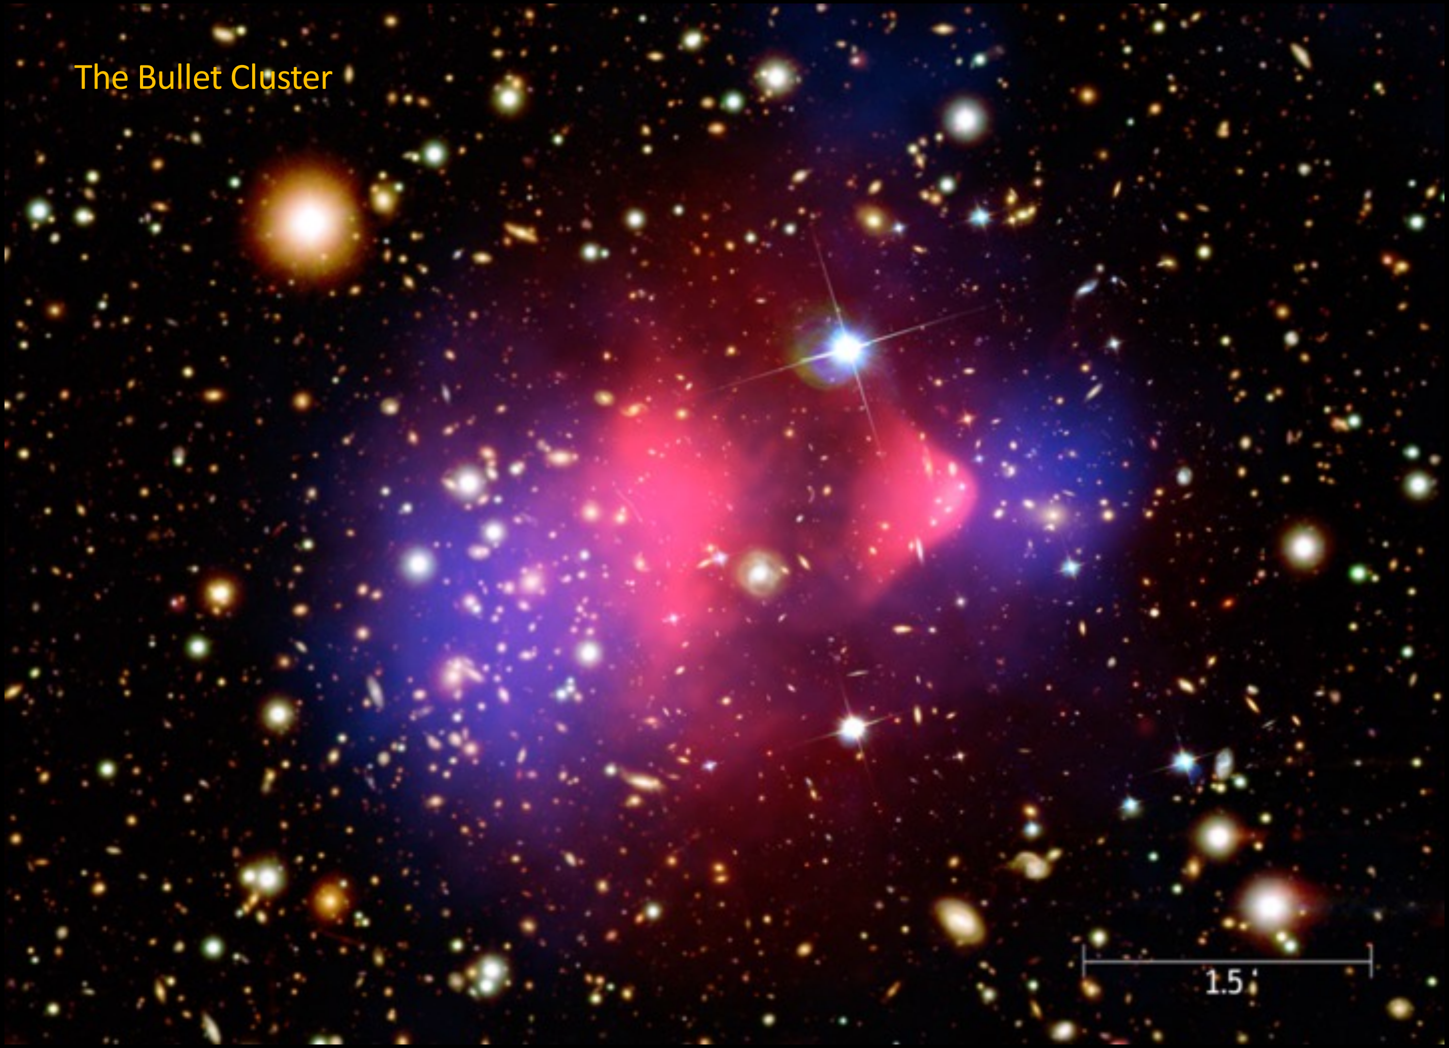
The Bullet Cluster
1.5'

Theory of LSS

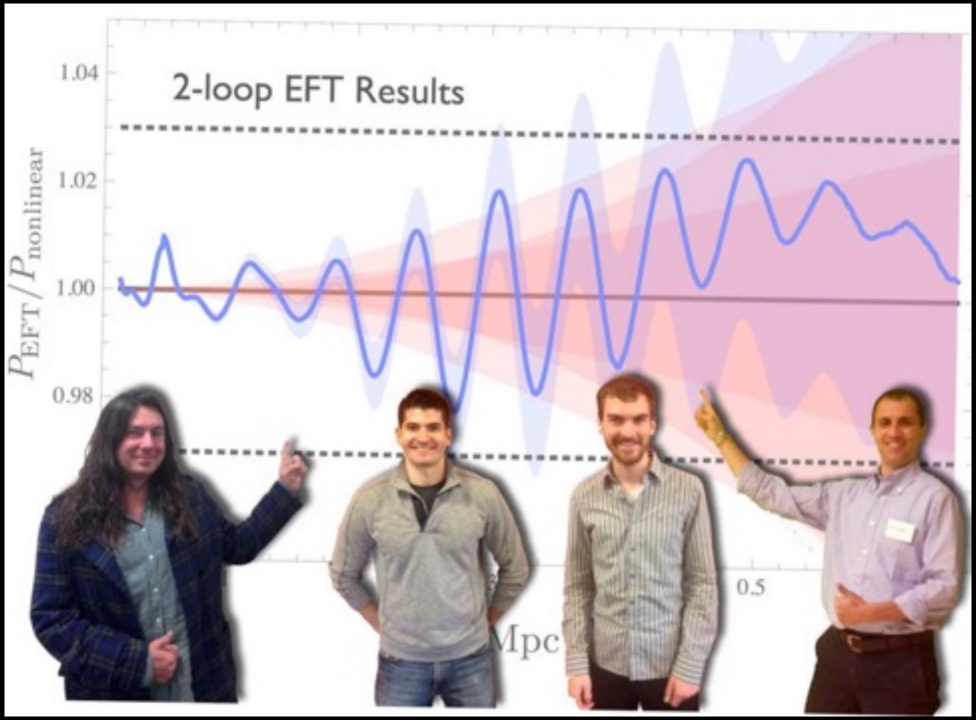

“Accurate” Theory? Simulations as T

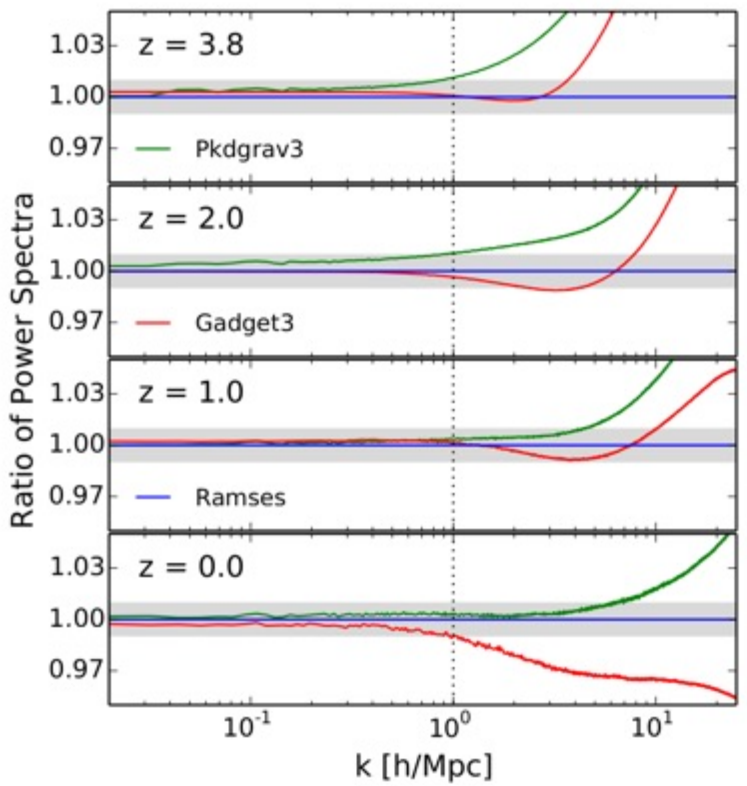

Figure 1. Comparison of the Power Spectra from the three different N-body codes at different redshifts. Green lines correspond to **Pkdgrav3**, red lines to **Gadget3**, and blue lines to **Ramses** (reference lines). One percent accuracy is obtained for $k \leqslant 1$ h/Mpc (dotted vertical line).

Simulations now agree to better than 1% at k ~ 10

The amount of information scales at k^3 there are 1’000’000x more “k-modes” of information to be used here.

Euclid wants to mine the deeply non-linear regime of data at k>0.1.

But do we have an accurate theory for describing this highly non-linear regime?

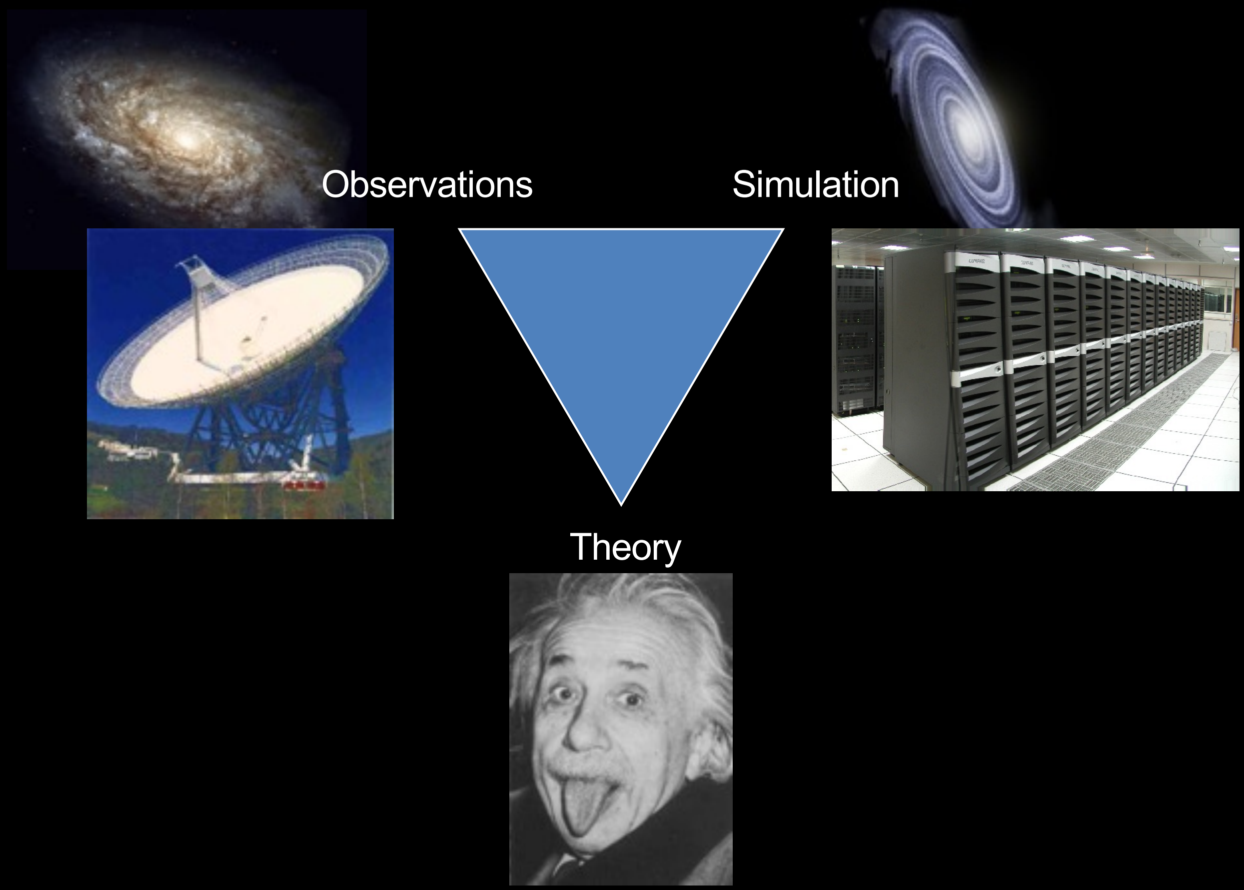
Observations
Simulation
Theory

Where there are galaxies, there are dark matter halos (and visa versa)

- Our Milky Way galaxy lives in a $10^{12}M_{\odot}$ dark matter halo
- To simulate such halos we need *at least* 1000 “particles”
- Our particles should be about $10^{9}M_{\odot}$
- The “piece” of the Universe that Euclid will see requires simulating a volume with side length of 12 billion lightyears
- These 2 factors lead to a *minimum* simulation size of 4 trillion particles!

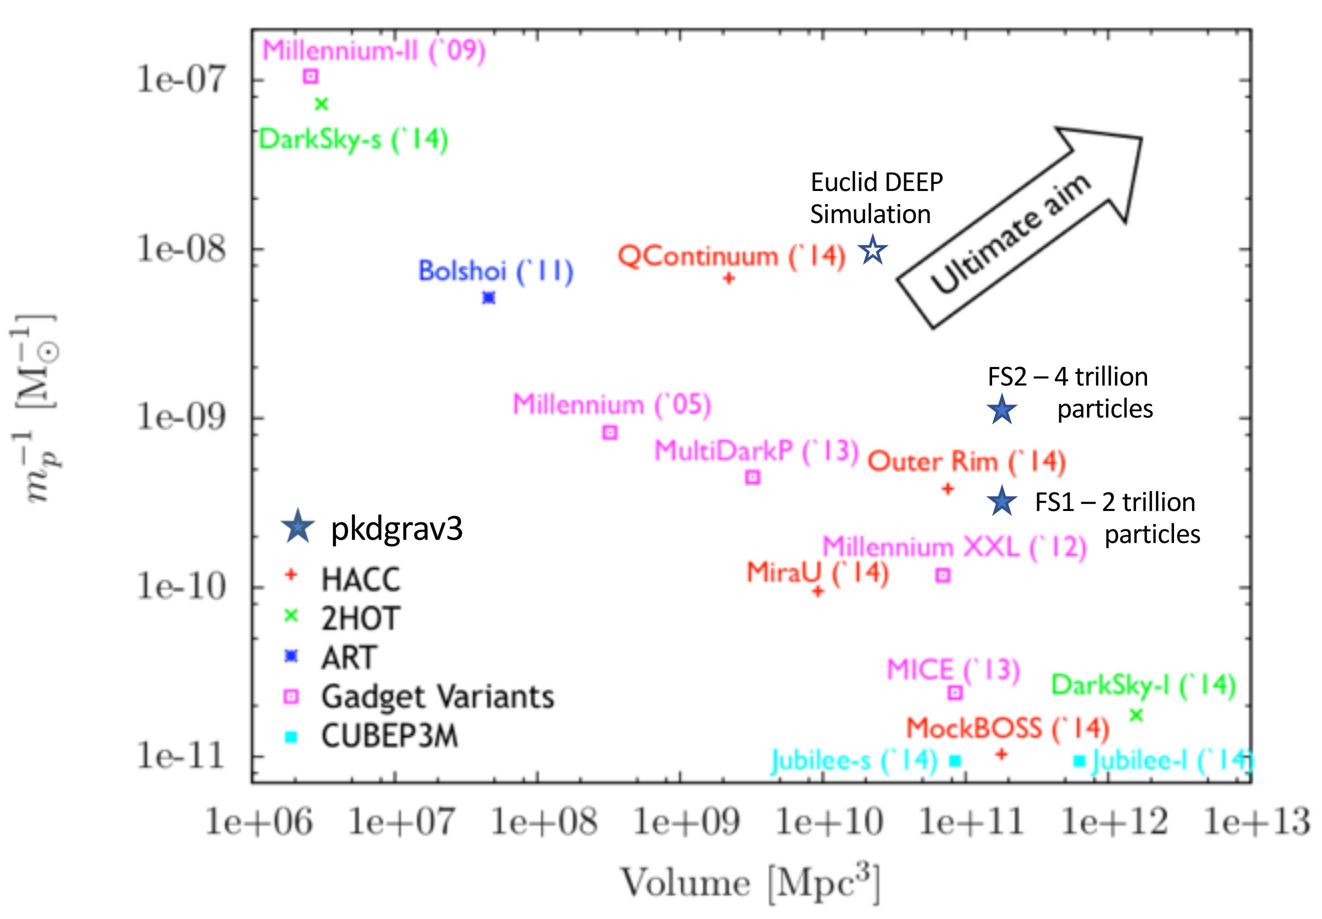

Millennium-II (`09)
DarkSky-s (`14)
Euclid DEEP Simulation
Ultimate aim
Bolshoi (`11)
QContinuum (`14)
FS2 – 4 trillion particles
Millennium (`05)
MultiDarkP (`13)
Outer Rim (`14)
FS1 – 2 trillion particles
pkdgrav3
HACC
2HOT
ART
Gadget Variants
CUBEP3M
Millennium XXL (`12)
MiraU (`14)
MICE (`13)
DarkSky-l (`14)
MockBOSS (`14)
Jubilee-s (`14)
Jubilee-l (`14)
1e-07
1e-08
1e-09
1e-10
1e-11
m_p^{-1} [$M_\odot^{-1}$]
1e+06
1e+07
1e+08
1e+09
1e+10
1e+11
1e+12
1e+13
Volume [Mpc^3]

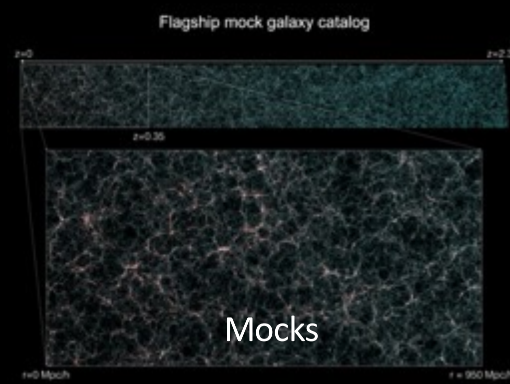

N-Body

Klypin & Prada 17, Blot+

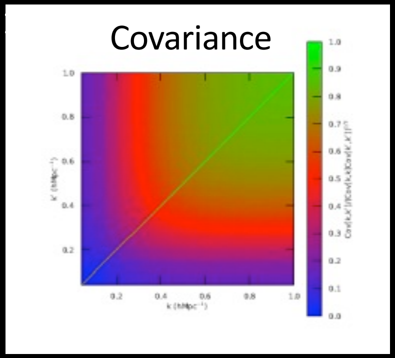

Knabenhans+ (19,20)

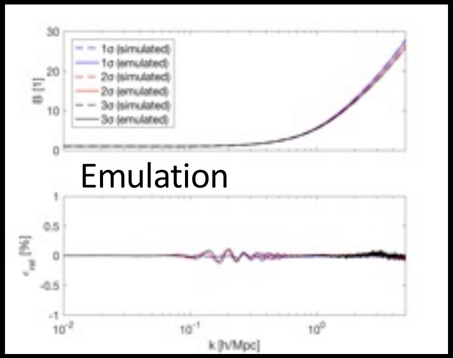

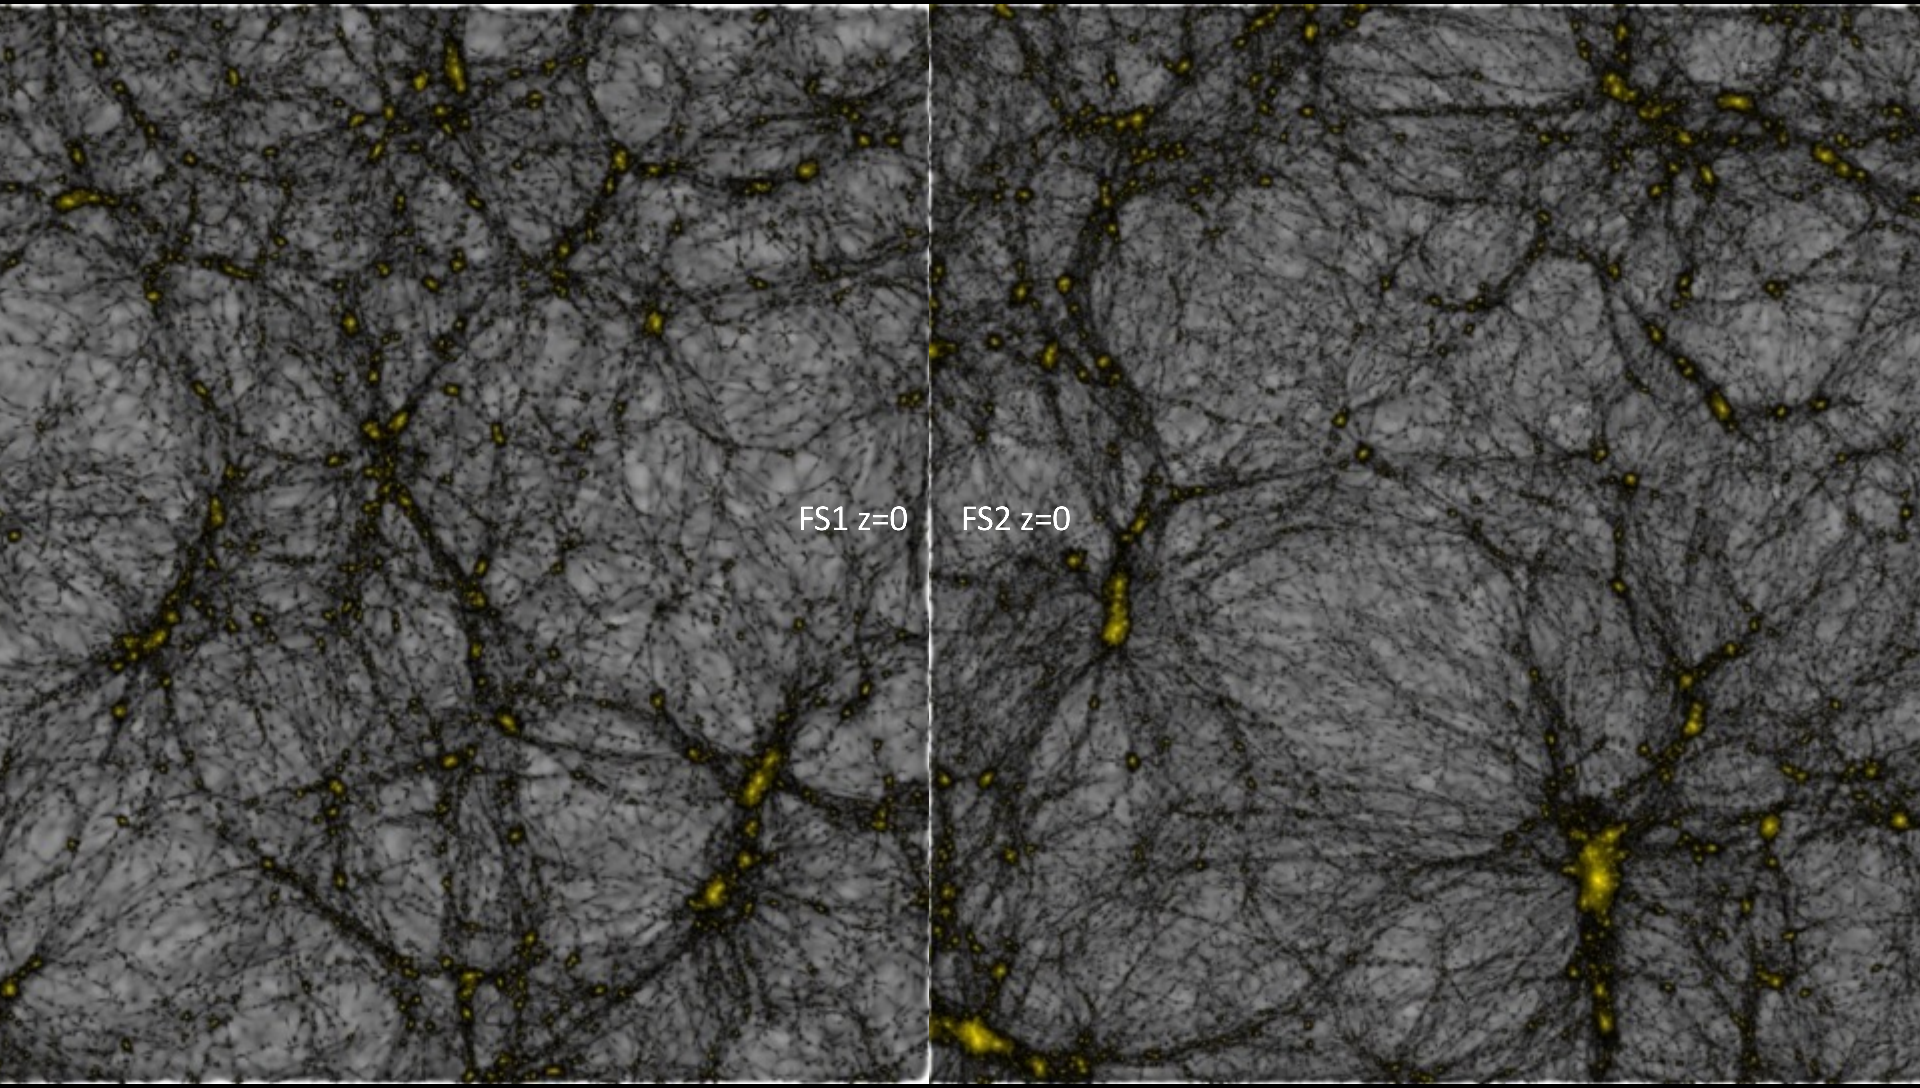
FS1 z=0
FS2 z=0

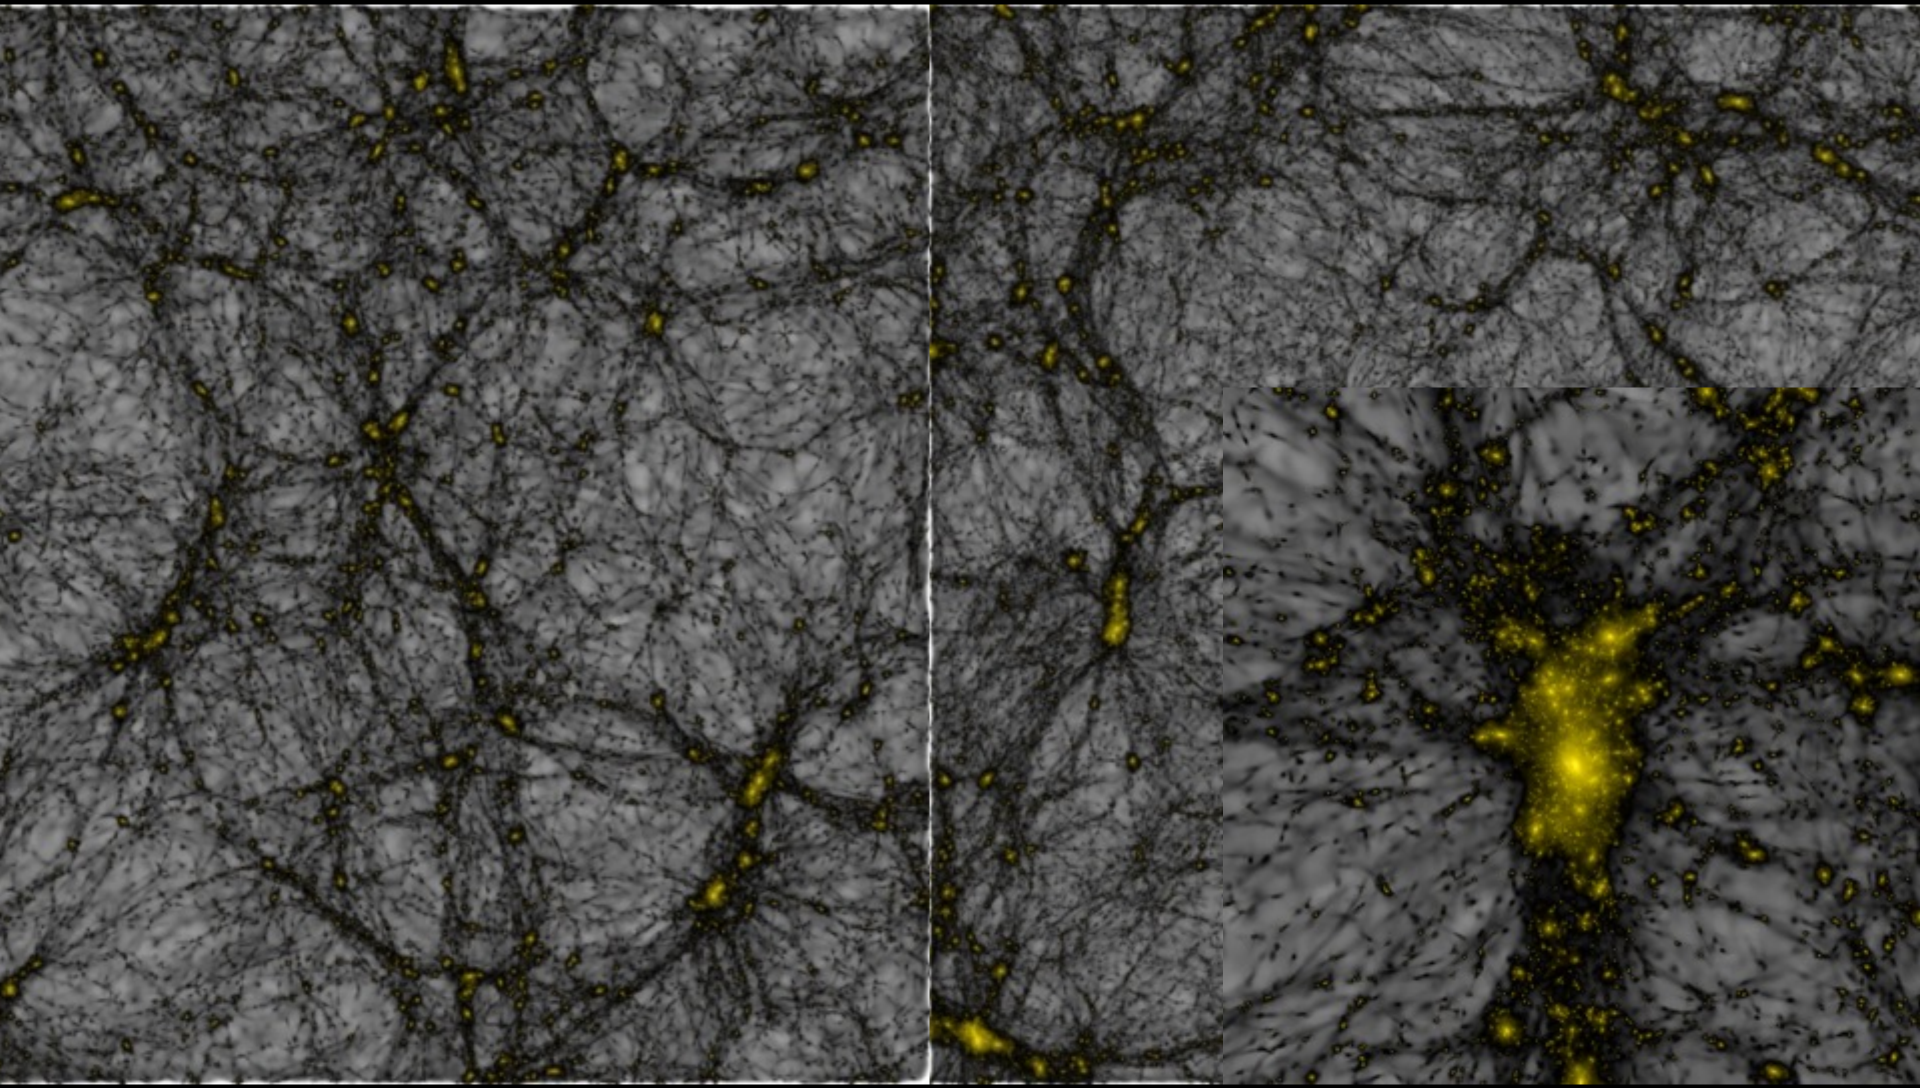

Flagship v2.0 Big Numbers

1. $(16'000)^3$ = 4.1 trillion particles
2. $(3'600\ h^{-1}\mathrm{Mpc})^3$ volume
3. $10^9\ h^{-1}\ M_{\odot}$ particle mass (Millenium res)
4. 835’000 node hours (12 core CPU + P100 GPU)
5. 1.3 Pbytes of on-the-fly data
   - 31 trillion particle light cone (700 TB)
   - z=10, 1.35, 1.00, 0.78, 0.54, 0 full particle snapshots (112 TB each)
   - 50x 8000^3 delta(**k**) grids from z=50 to 0 (100 TB)
   - 200 Healpix Maps (nSide=16384) (2 TB)
6. ≈150 billion Rockstar halos with particle subset for placing satellites (in progress) (??? TB)

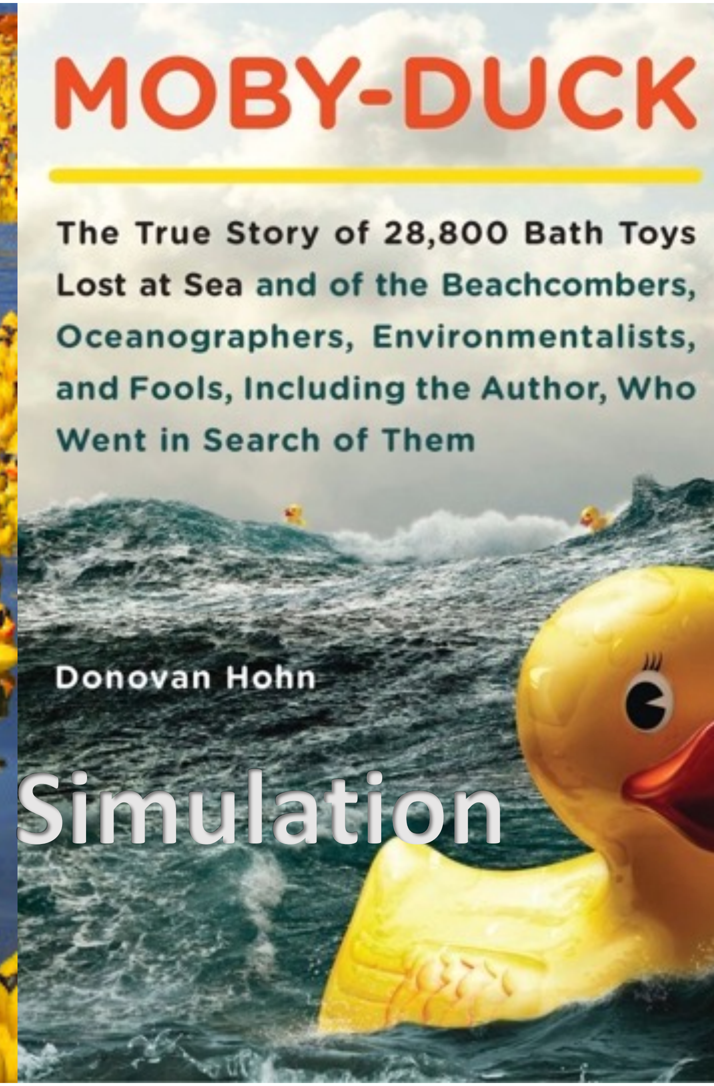

Collisionless N-body Simulation

Collisions can be critical!

Real stars and planets don't always ignore each other!

This requires both very precise forces and very good integration of the orbits at close approach.

The N-Body Solution of a 6-D Fluid Collisionless Boltzmann Equation

$$\ddot{x}_i = \sum_{j\neq i}^{N} -\nabla G m_j / |x_i - x_j|$$

$$\frac{\partial f}{\partial t} + \mathbf{v} \bullet \frac{\partial f}{\partial \mathbf{r}} - \nabla \Phi \bullet \frac{\partial f}{\partial \mathbf{v}} = 0$$

The density in 6-D phase space is conserved. Where the spatial density is high, the spread in velocity space is high (lower velocity space density).

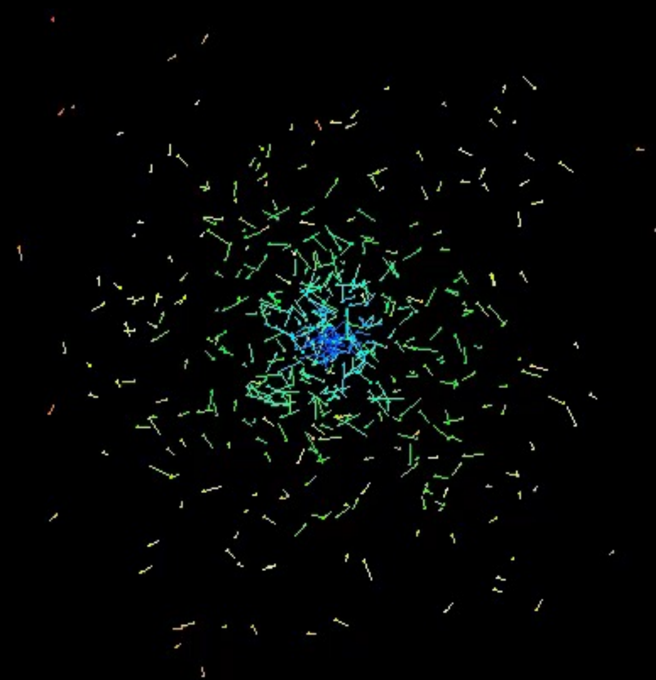

Simulating with particles

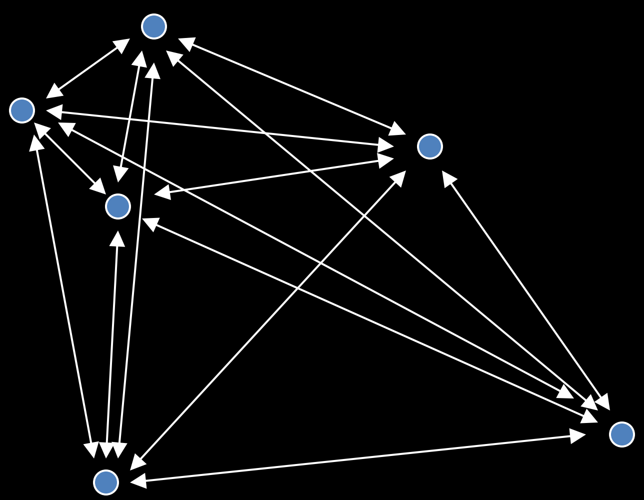

For N particles – N(N-1)/2 Forces!

Today we use > 1 000 000 000 000 = 10^24 Forces!

pkdgrav3 and Fast Multipole

Quick explanation of FMM

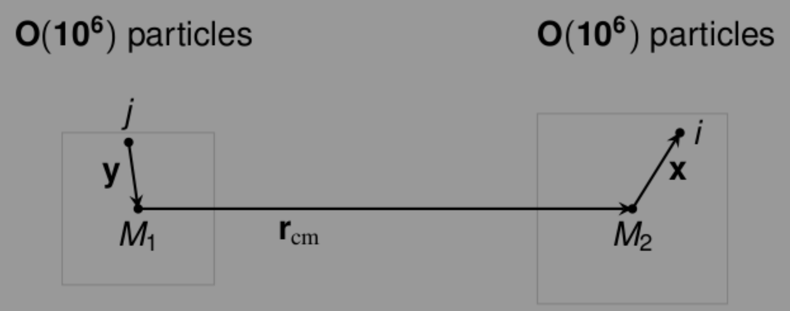

Direct $\mathbf{O}(\mathbf{10^{12}})$ interactions to calculate! $O(N^2)$ code.

Tree Use a multipole approximation for the mass at M_2 to calculate the force at each j: $\mathbf{O}(\mathbf{10^6})$ interactions to calculate. $O(N \log N)$ code.

FMM Use a multipole approx for the mass at M_2 to approximate the "potential landscape" at M_1 (n^{th} order gradients of the potential): $\mathbf{O}(\mathbf{1})$ interaction to calculate. $O(N)$ code!

FMM: memory balance = compute balance when all N particles are computed!

Calculating Forces – Direct summation $\mathcal{O}(N^2)$

- There are $\frac{N(N-1)}{2}$ individual forces $\boldsymbol{F}_{ij}$ to calculate at each step of the integration of the equations of motion.
- Each of these forces requires about 20 floating point operations.
- The fastest computers today can theoretically do about 10^{17} flop/s.
- How long would it take to calculate the forces 10^{12} particles **once**?
- $= 10^{25}$ flop $= 10^{8}$s $= 3.2$ years!
- Typically we need to calculate forces several hundreds to thousands of times per cosmology simulation, so this is a big problem.
- Accuracy in an N-body simulation comes primarily from N, the number of particles used, *not from the accuracy of the force calculation!*

Calculating Forces – Multipole approximation

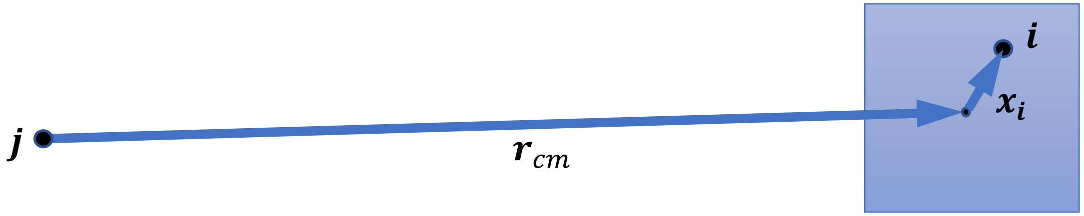

- The gravitational potential at a point in space due to a mass distribution over the volume V is given by,

$$\Phi = \int_V \gamma(|\boldsymbol{r}|)\rho(\boldsymbol{r})d^3\boldsymbol{r},$$

where $\gamma(|\boldsymbol{r}|)$ denotes the Green's function; for *unsoftened* gravity this is given by $\gamma(r) = -1/r$ (setting G=1). The mass density for a distribution of particles,

$\rho(\boldsymbol{r}) = \sum_{i\in V}\delta(\boldsymbol{r}_i - \boldsymbol{r})m_i$ this then leads to,

$$\Phi = \sum_{i\in V} m_i\gamma(|\boldsymbol{r}_i|) = \sum_{i\in V} m_i\gamma(|\boldsymbol{r}_{cm} + \boldsymbol{x}_i|)$$

Fast Multipole Method (FMM): $\mathcal{O}(N)$

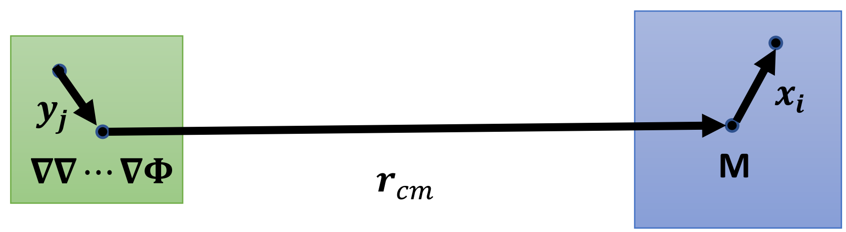

- Same as before but this time assuming $|\boldsymbol{x}_i + \boldsymbol{y}_j| \ll |\boldsymbol{r}_{cm}|$,

$$\Phi_j = \sum_{i \in V} m_i \sum_{n=0}^{5} \frac{1}{n!} [\partial_{\vec{n}} \gamma(\boldsymbol{r}_{cm})] : (\boldsymbol{x}_i + \boldsymbol{y}_j)^{\vec{n}}$$

- Now we need to use the binomial theorem for the tensor at the end,

$$\Phi_j = \sum_{i \in V} m_i \sum_{n=0}^{5} \frac{1}{n!} [\partial_{\vec{n}} \gamma(\boldsymbol{r}_{cm})] : \sum_{l=0}^{n} \frac{1}{l!} \boldsymbol{x}_i^{\vec{l}} \boldsymbol{y}_j^{\overrightarrow{n-l}}$$

collect multipole!

Fast Multipole Method (FMM): $\mathcal{O}(N)$

- Each M-L interaction takes about 450 flop at 5^{th} order.
- In reality the M-L interaction is more complex than outlined.
- We use *trace-free* moments for M and L.
- We use unit vectors to avoid very high powers of r that can occur. This is also very important for *mixed precision calculation*.

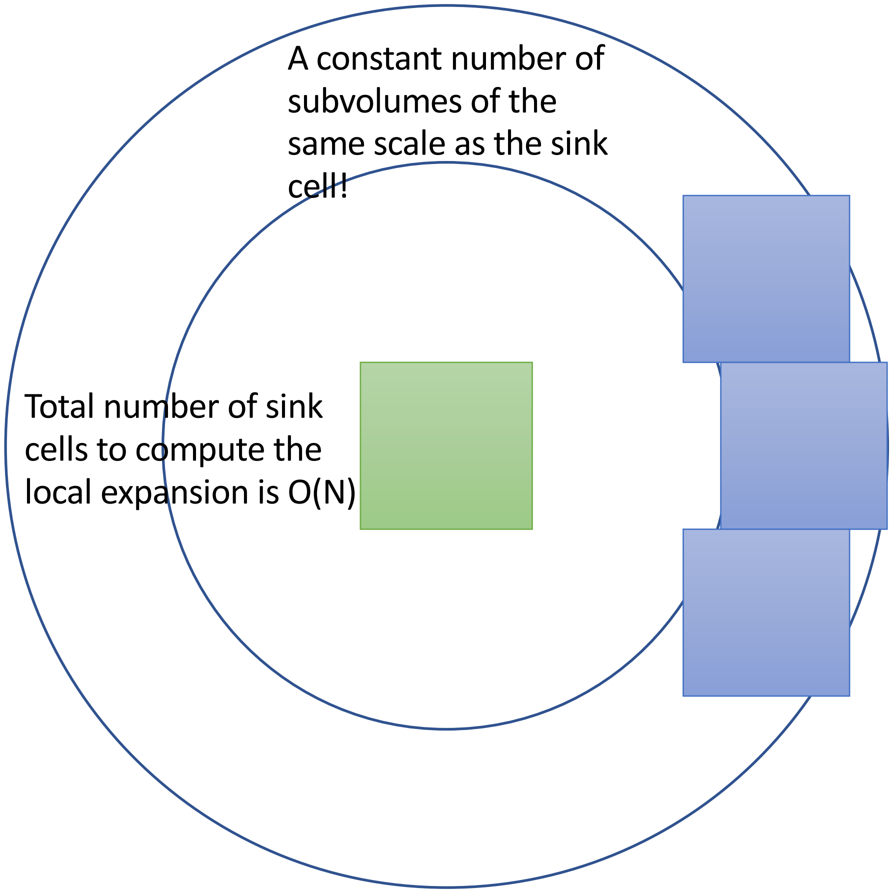

Time to calculate all forces on *a trillion* particles using FMM is only a few seconds vs. 3.2 years!

Computation…

My first N-body machine (very ugly)

i486 DX (had a math coprocessor, bought with my own money)

N = 32’000 particles while studying in Toronto 1992

Piz Daint – over 5000 GPU Nodes
4000 Nodes were used

Swiss National Computing Center (CSCS) in Lugano, Switzerland

The pkdgrav3 N-Body Code

- Started development in 1992 (NASA HPCC)
- Fast Multipole Method, O(N), 5th order in Φ
- Open source, available at: www.pkdgrav.org

Douglas Potter →

System Complexity

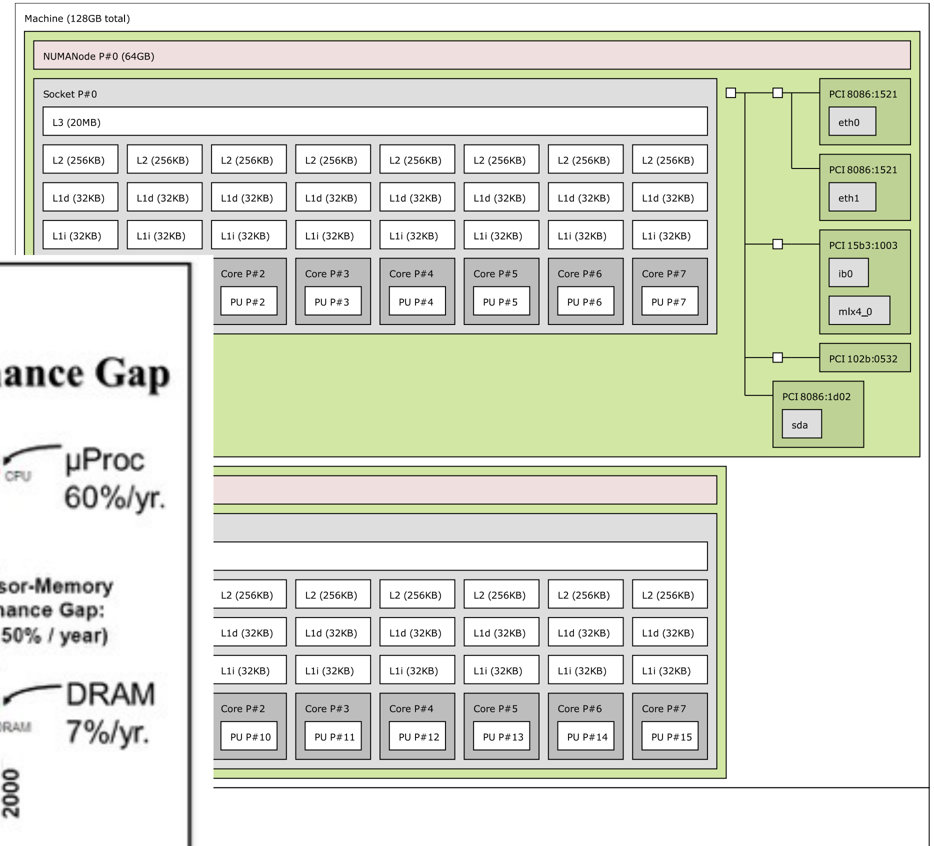

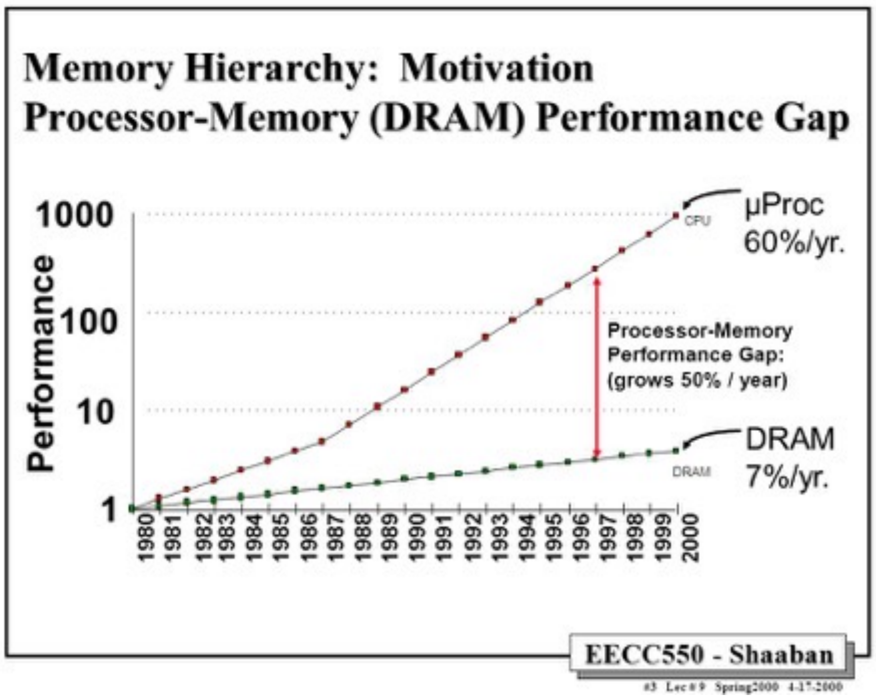

Memory Usage in pkdgrav3

0.5 billion particles can fit on a 32 Gbyte Node like Piz Daint

28 bytes persistent

6 bits: old rung 24: group id	
pos[0]	int32_t
pos[1]	int32_t
pos[2]	int32_t
vel[0]	float
vel[1]	float
vel[2]	float

<28 bytes / particle

Tree Cells
Binary Tree

4th order
Multipoles
(float prec)

~5 bytes / particle

Cache/Buffers

0-8 bytes ephemeral

Group finding
Other analysis

ClAoS is used for the particle and cell memory which makes moving particles around simple

AoSoA is used for all interaction lists which are built by the TreeWalk algorithm.

Reducing memory usage increases the capability of existing machines, but also increases performance somewhat. Simulations are limited more by memory footprint.

GPU Hybrid Computing

Piz Daint example

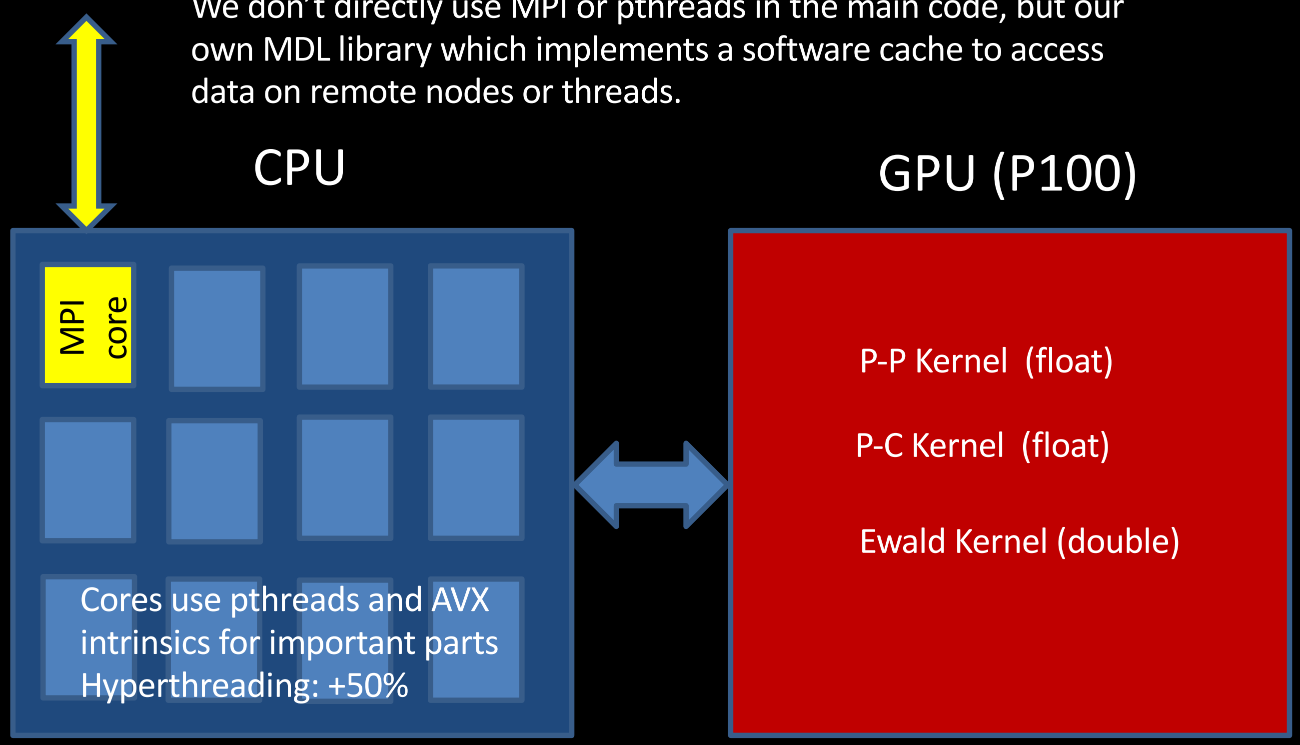

AVX instructions, what are they?

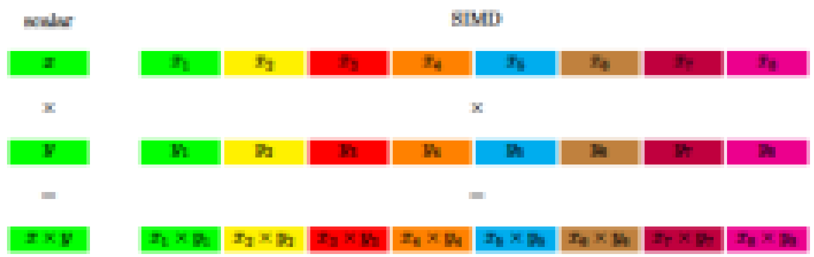

- SSE: 4 x float (128 bit), AVX: 8 x float (256 bit), AVX-512: 16 x float
- Arrays of structures of arrays. Reorganizing data before computing!

particle structure: each field is a vector of 32

x0-x31	x32-x63						
y0-y31	y32-y63	…					
z0-z31	z32-z63						

GPU

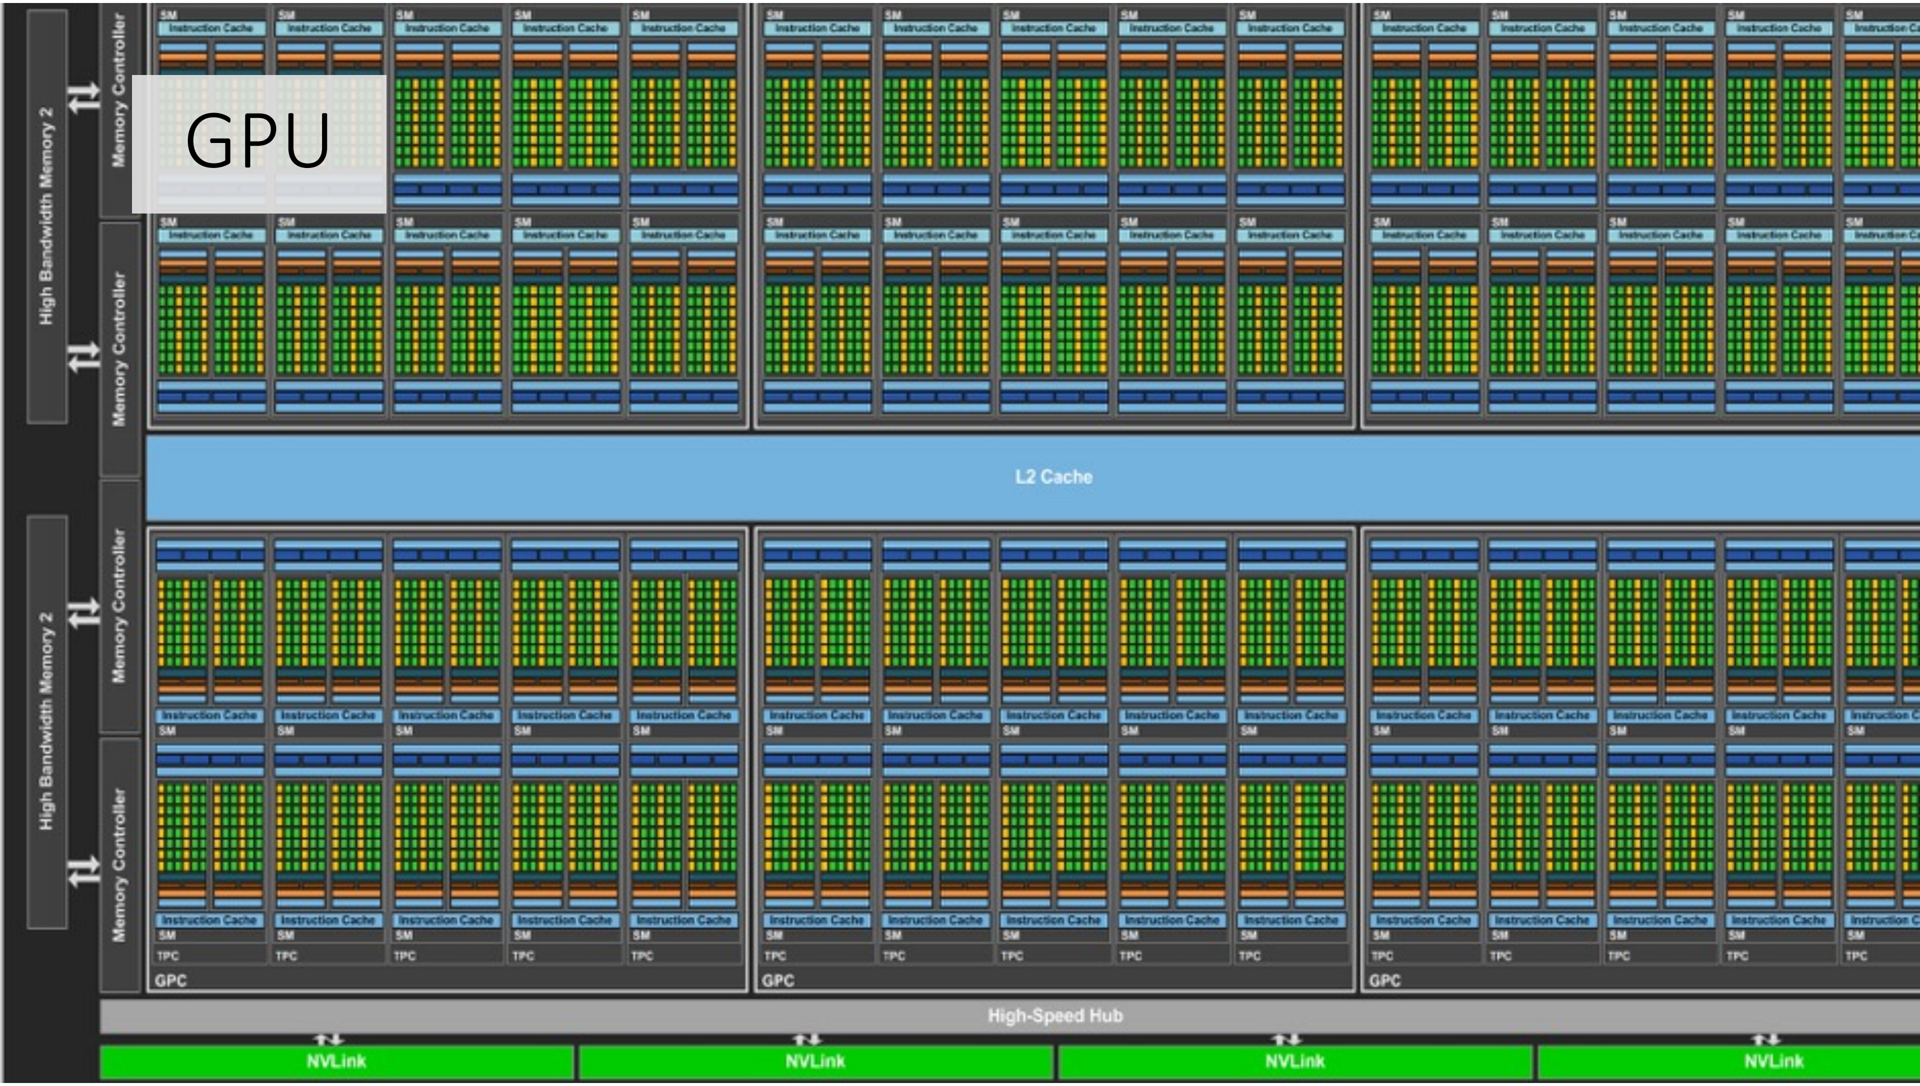
High Bandwidth Memory 2
Memory Controller
SM
Instruction Cache
L2 Cache
TPC
GPC
High-Speed Hub
NVLink

Mixed Precision (and Tensor Ops?)

- Scaled multipoles: Instead of $M^{kl} = \sum_{i=1}^{n_{cell}} m_i \, x_i^{kl}$, use

$$M^{kl} = \sum_{i=1}^{n_{cell}} \frac{m_i}{M_{cell}} \left(\frac{\boldsymbol{x}_i}{r_{cell}}\right)^{\boldsymbol{kl}}$$

- Each expansion can be calculated in single precision (*fp32*) with the force being $M/r^2(1 + \theta^2 + \theta^3 + \theta^4 + \cdots)$ with most of the flop being in the ( ). This can probably also be done with *fp16*. M/r^2 can still be calculated in double precision (*fp64*).

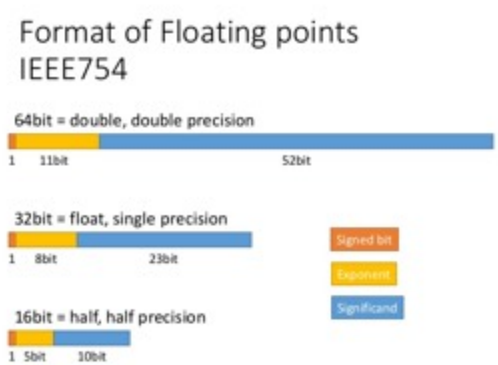

Feature	Tesla V100 SXM2 16GB/32GB	Tesla V100 PCI-E 16GB/32GB	Tesla V100S PCI-E 32GB	Quadro GV100 32GB
GPU Chip(s)	Volta GV100			
TensorFLOPS	125 TFLOPS	112 TFLOPS	130 TFLOPS	118.5 TFLOPS
Integer Operations (INT8)*	62.8 TOPS	56.0 TOPS	65 TOPS	59.3 TOPS
Half Precision (FP16)*	31.4 TFLOPS	28 TFLOPS	32.8 TFLOPS	29.6 TFLOPS
Single Precision (FP32)*	15.7 TFLOPS	14.0 TFLOPS	16.4 TFLOPS	14.8 TFLOPS
Double Precision (FP64)*	7.8 TFLOPS	7.0 TFLOPS	8.2 TFLOPS	7.4 TFLOPS

Profile of 4.1 trillion particle simulation (Piz Daint)
O(N) and everything else matters!

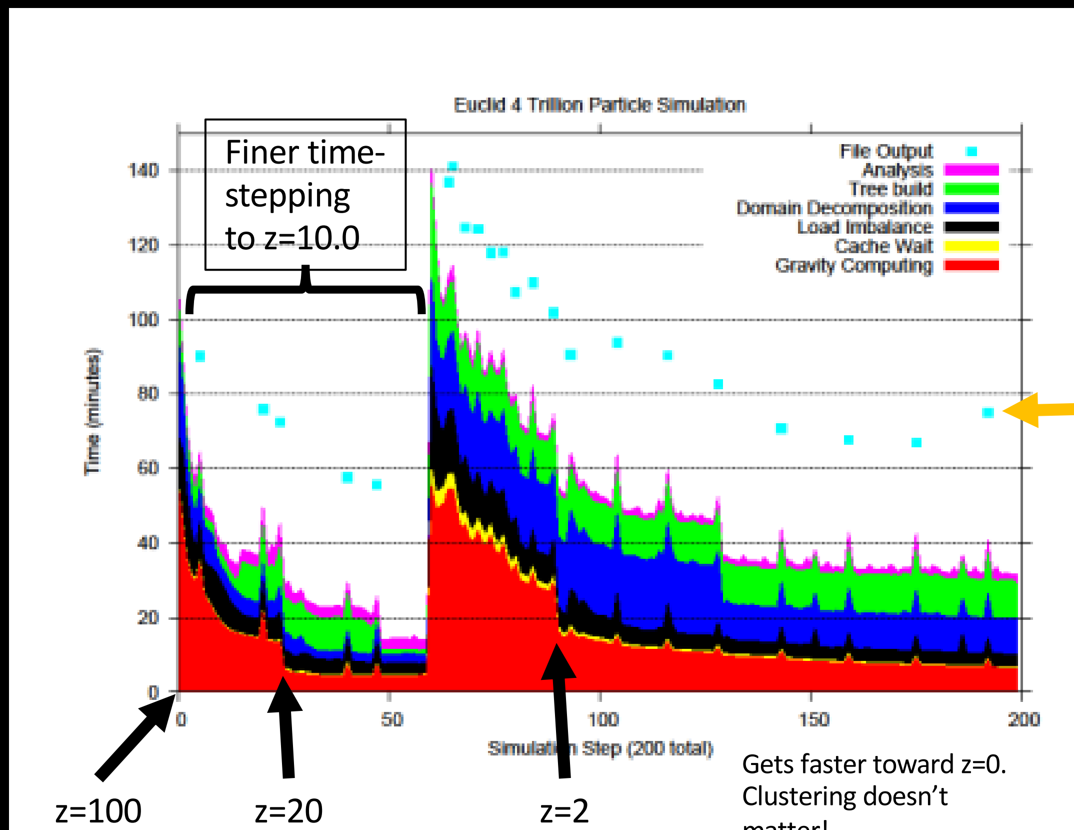

12 hour "large" Queue was used

Load balancing: domain decomposition

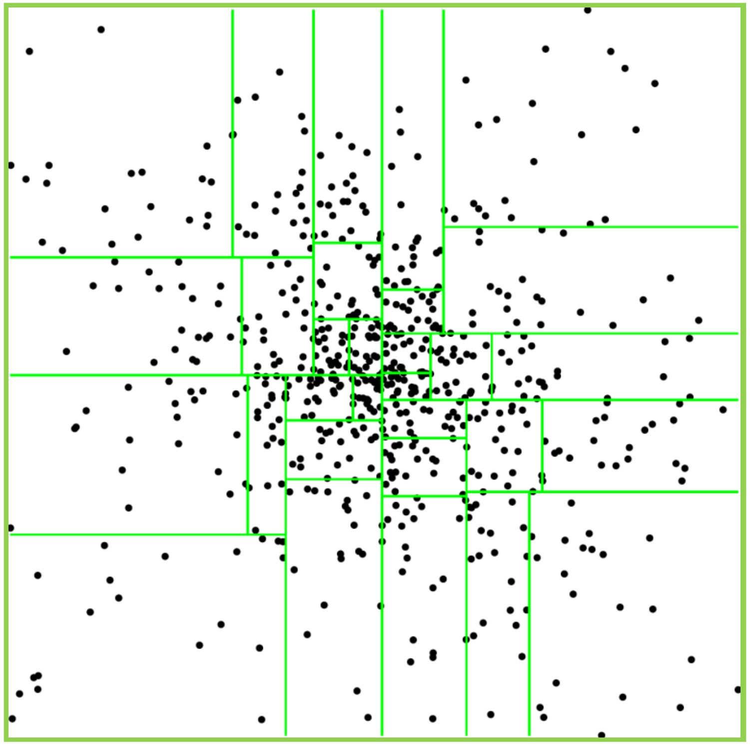

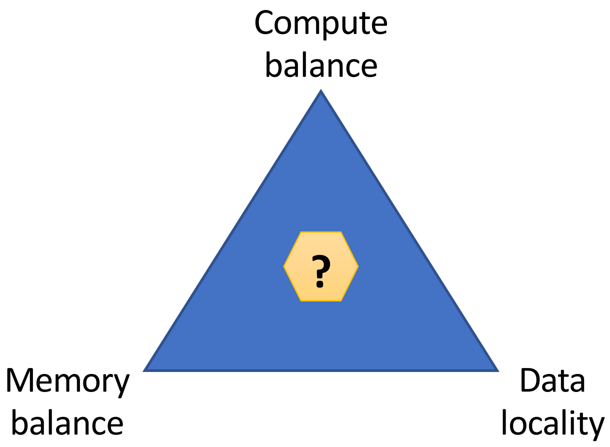

Is it possible to achieve all three?

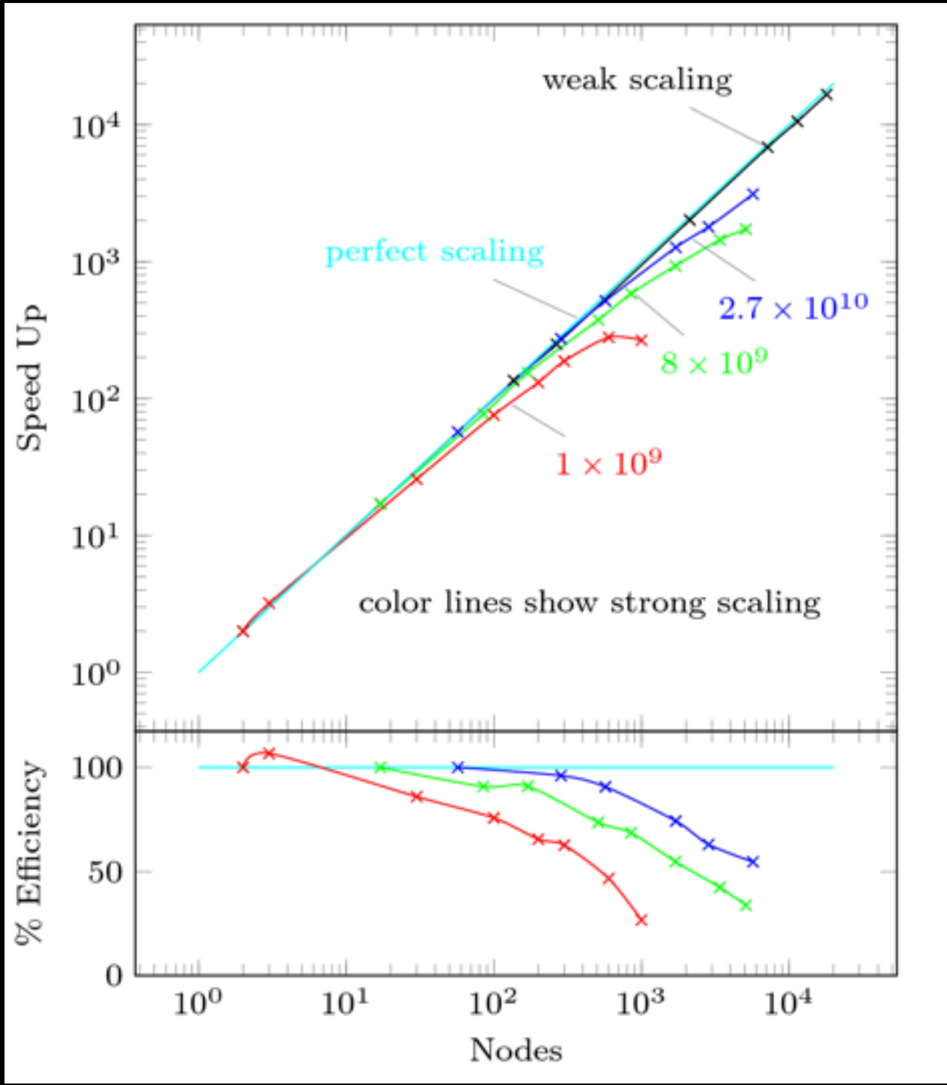

weak scaling
perfect scaling
2.7×10^{10}
8×10^9
1×10^9
color lines show strong scaling
Speed Up
10^0
10^1
10^2
10^3
10^4
% Efficiency
100
50
0
Nodes

Frontier: Simulating the Baryons as well as the dark stuff!

Zooms

Volumes

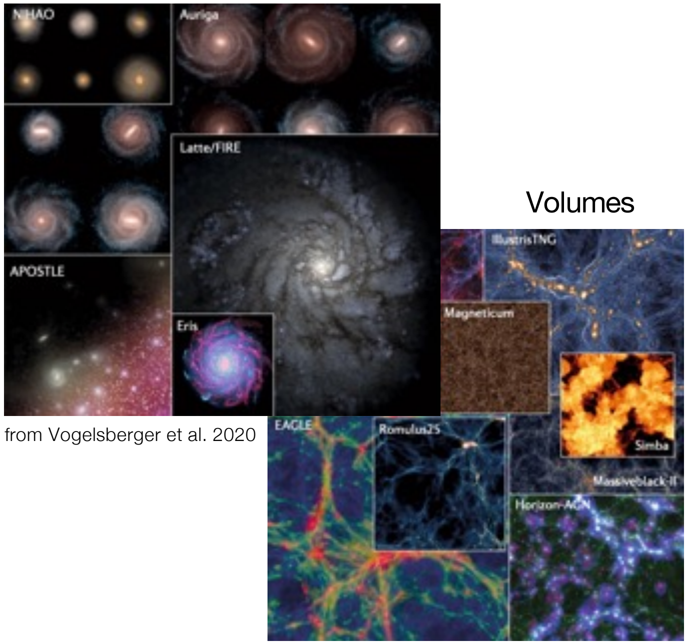

from Vogelsberger et al. 2020

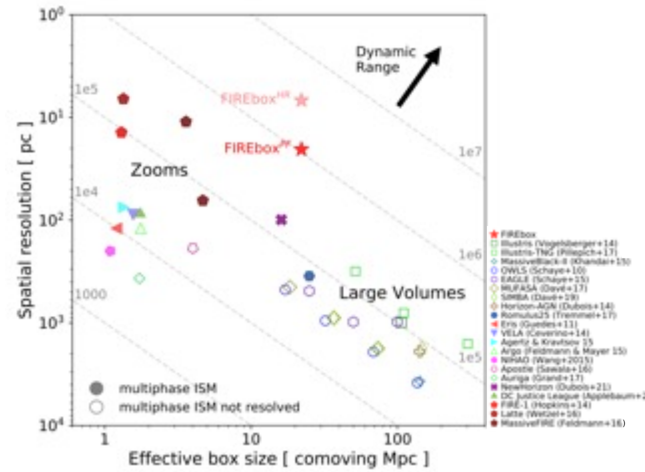

RF et al.

Modeling Challenges

- Galaxy formation highly complex
- range of scale problem
- model degeneracies

Complex numerical models required ➜ HPC & ML needed!

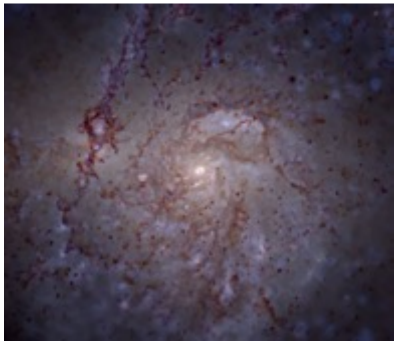

Astrophysicists solve a dark matter puzzle

By Liz Fuller-Wright, Princeton University, and Lucas Van Wyk Joel, University of California-Irvine
Feb. 14, 2022, 11 a.m.

In a new Nature Astronomy study, an international team of astrophysicists report how, when tiny galaxies collide with bigger ones, the bigger galaxies can strip the smaller galaxies of their dark matter — matter that we can't see directly, but which astrophysicists think must exist because, without its gravitational effects, they couldn't explain things like the motions of a galaxy's stars.

It's a mechanism that has the potential to explain how galaxies might be able to exist without dark matter — something once thought impossible.

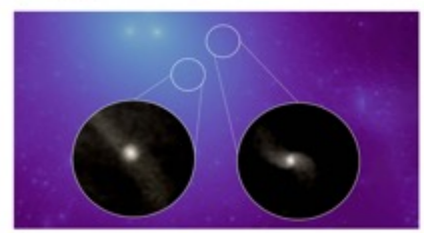

FIREbox

Cosmological volume simulation with FIRE physics (Feldmann)

Mitigates range of scale problem:

- FIREbox first simulation of its kind to reach a dynamic range > 10^6
- Resolve galaxy structure & multi-phase ISM in fully cosmological context down to z=0

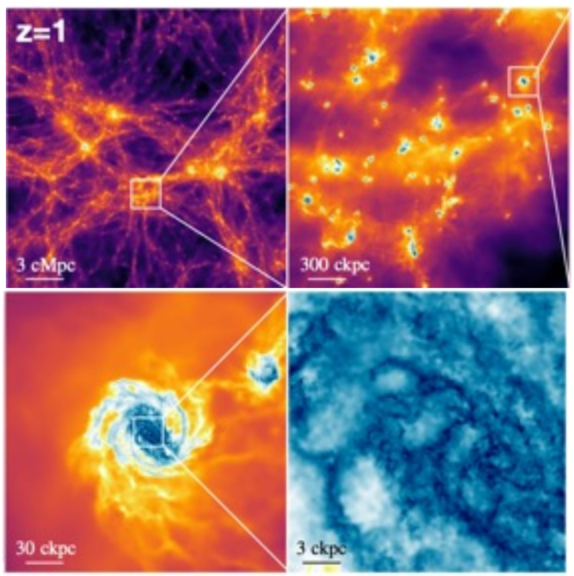

RF et al.

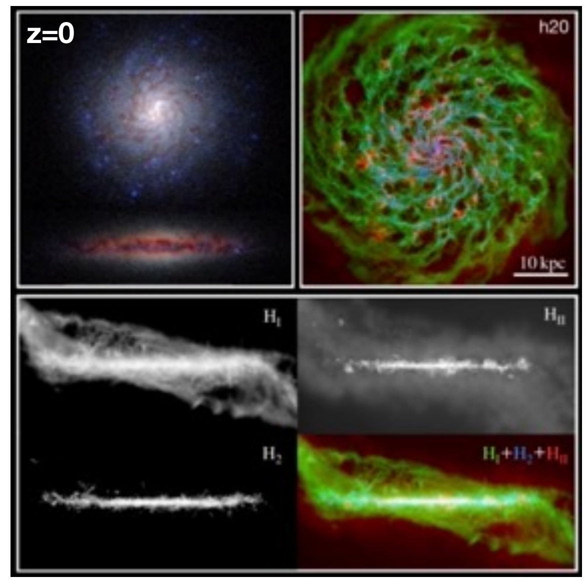

RF et al.

Summary

- Simulations are the only tool able to reliably calculate observables in the highly non-linear regime. They are effectively “the Theory” for upcoming observational surveys (Euclid, SKA).
- To reach the required precision, very large simulations and simulation campaigns are mandatory.
- Modern simulation codes, like PKDGRAV3, need to continually adapt to new supercomputing architectures! Soon simulations will reach >10 trillion particles, a big challenge to the data processing!
- Machine learning provides a way of “replacing” simulations, or parts of simulations, at a vastly reduced cost. Could be used to perform Galaxy formation over a large volume.