MIT KAVLI INSTITUTE

Cosmological Simulations of Galaxy Formation

Ryan McKinnon

May 25, 2017

The Big Picture

The Big Picture

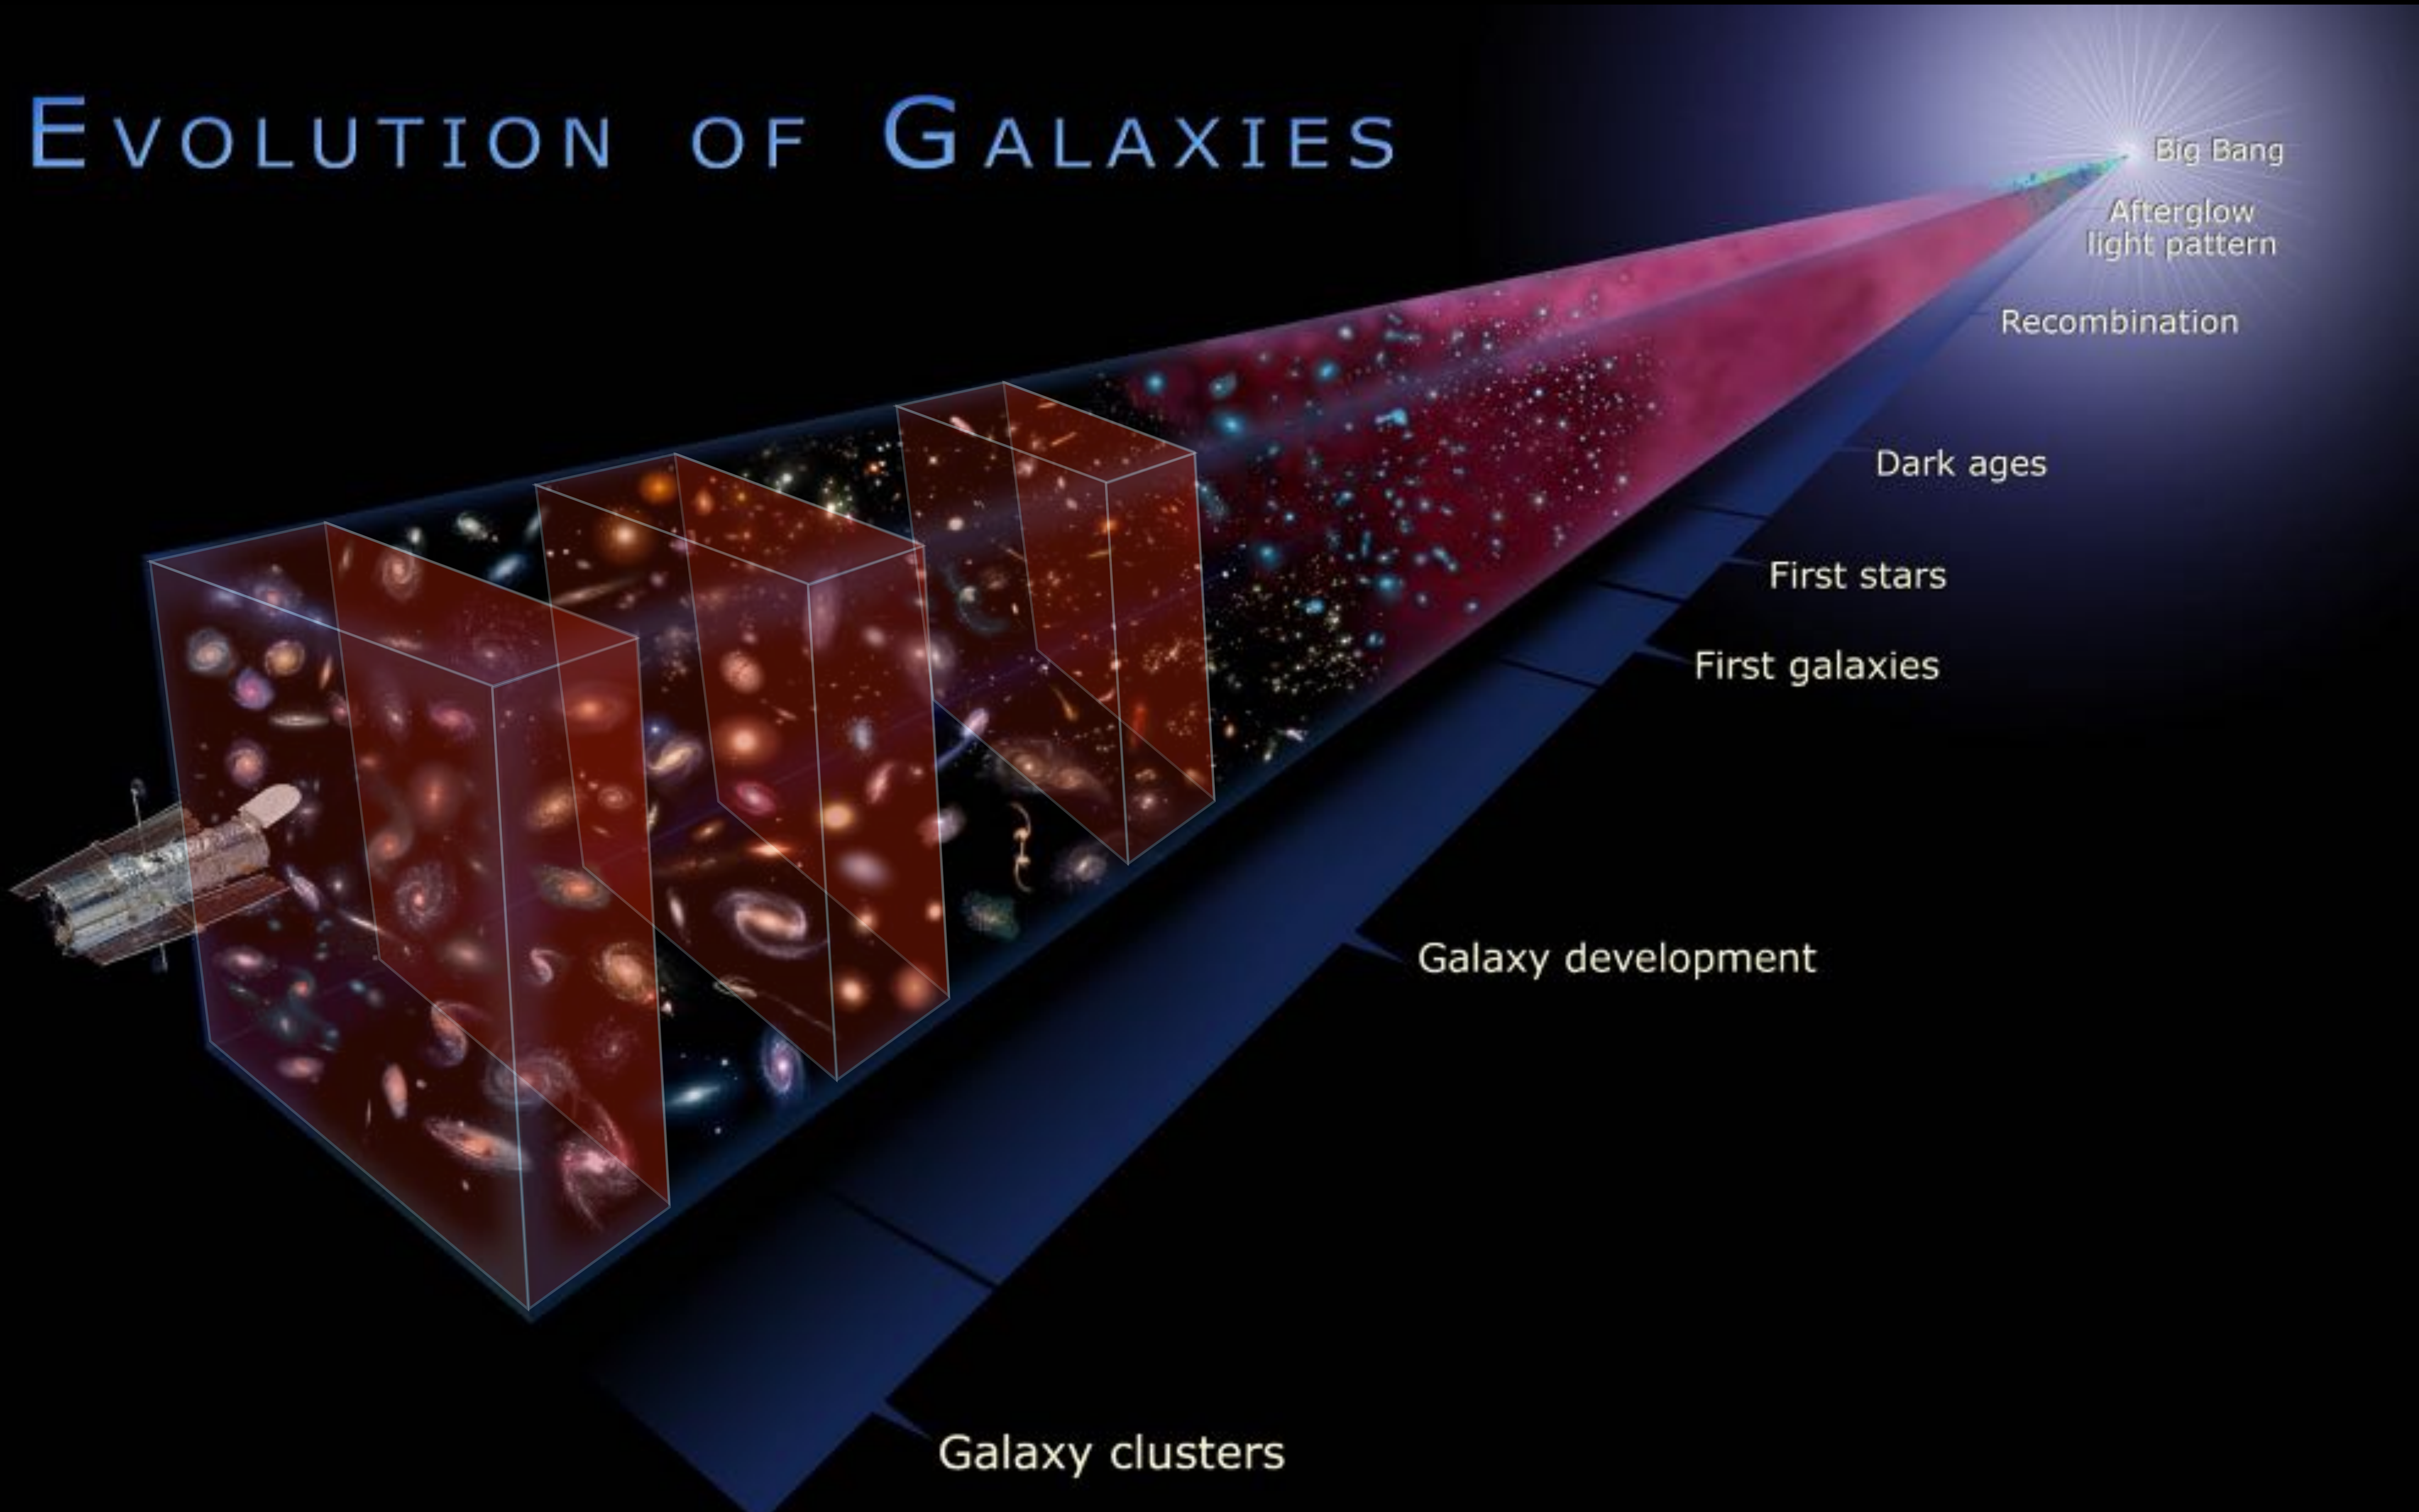

The Big Picture

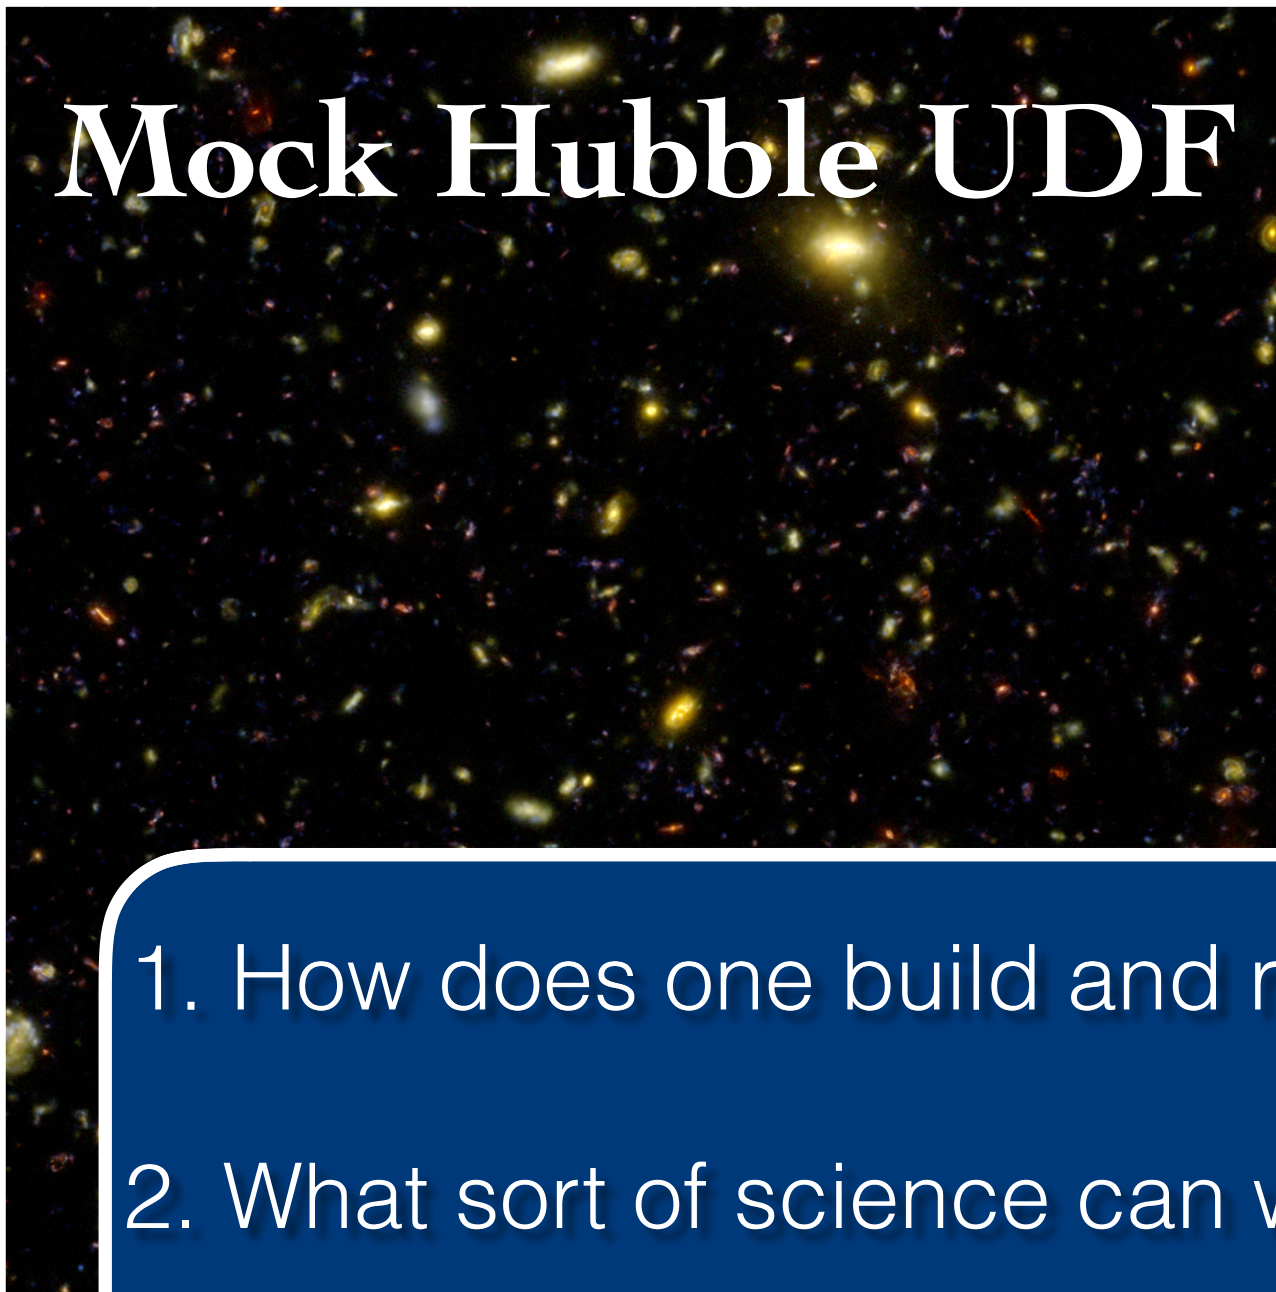

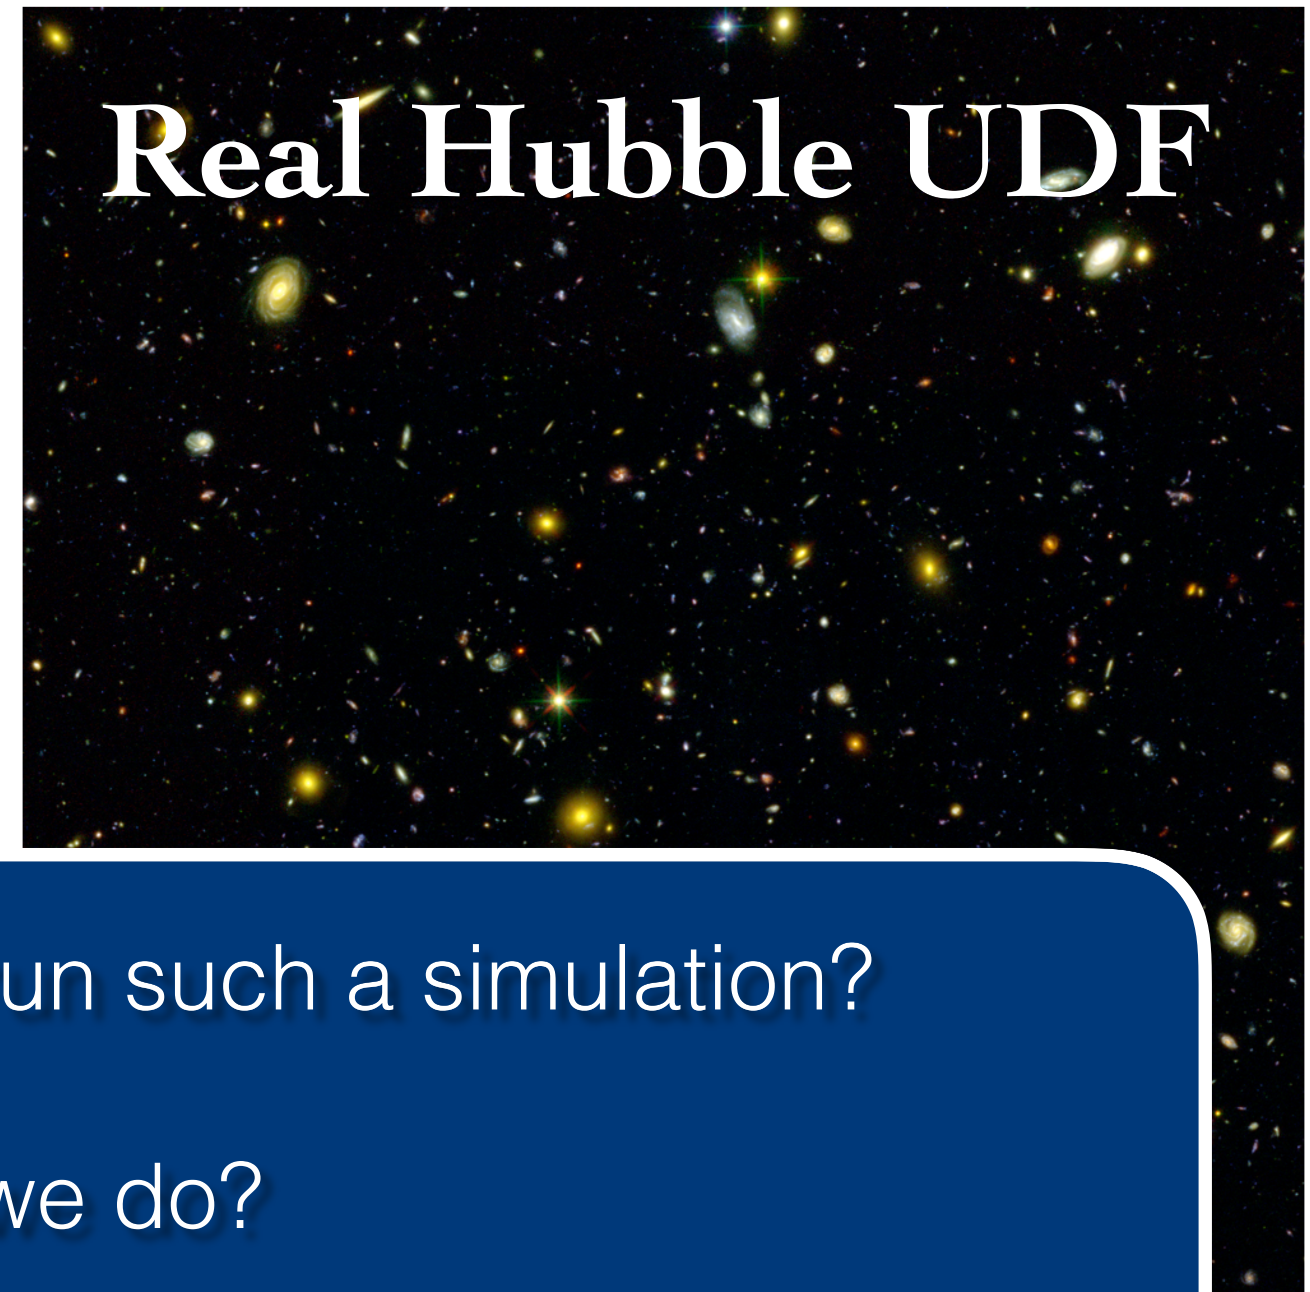

1. How does one build and run such a simulation?
2. What sort of science can we do?
3. What are the computational costs?

Image Credit: Greg Snyder

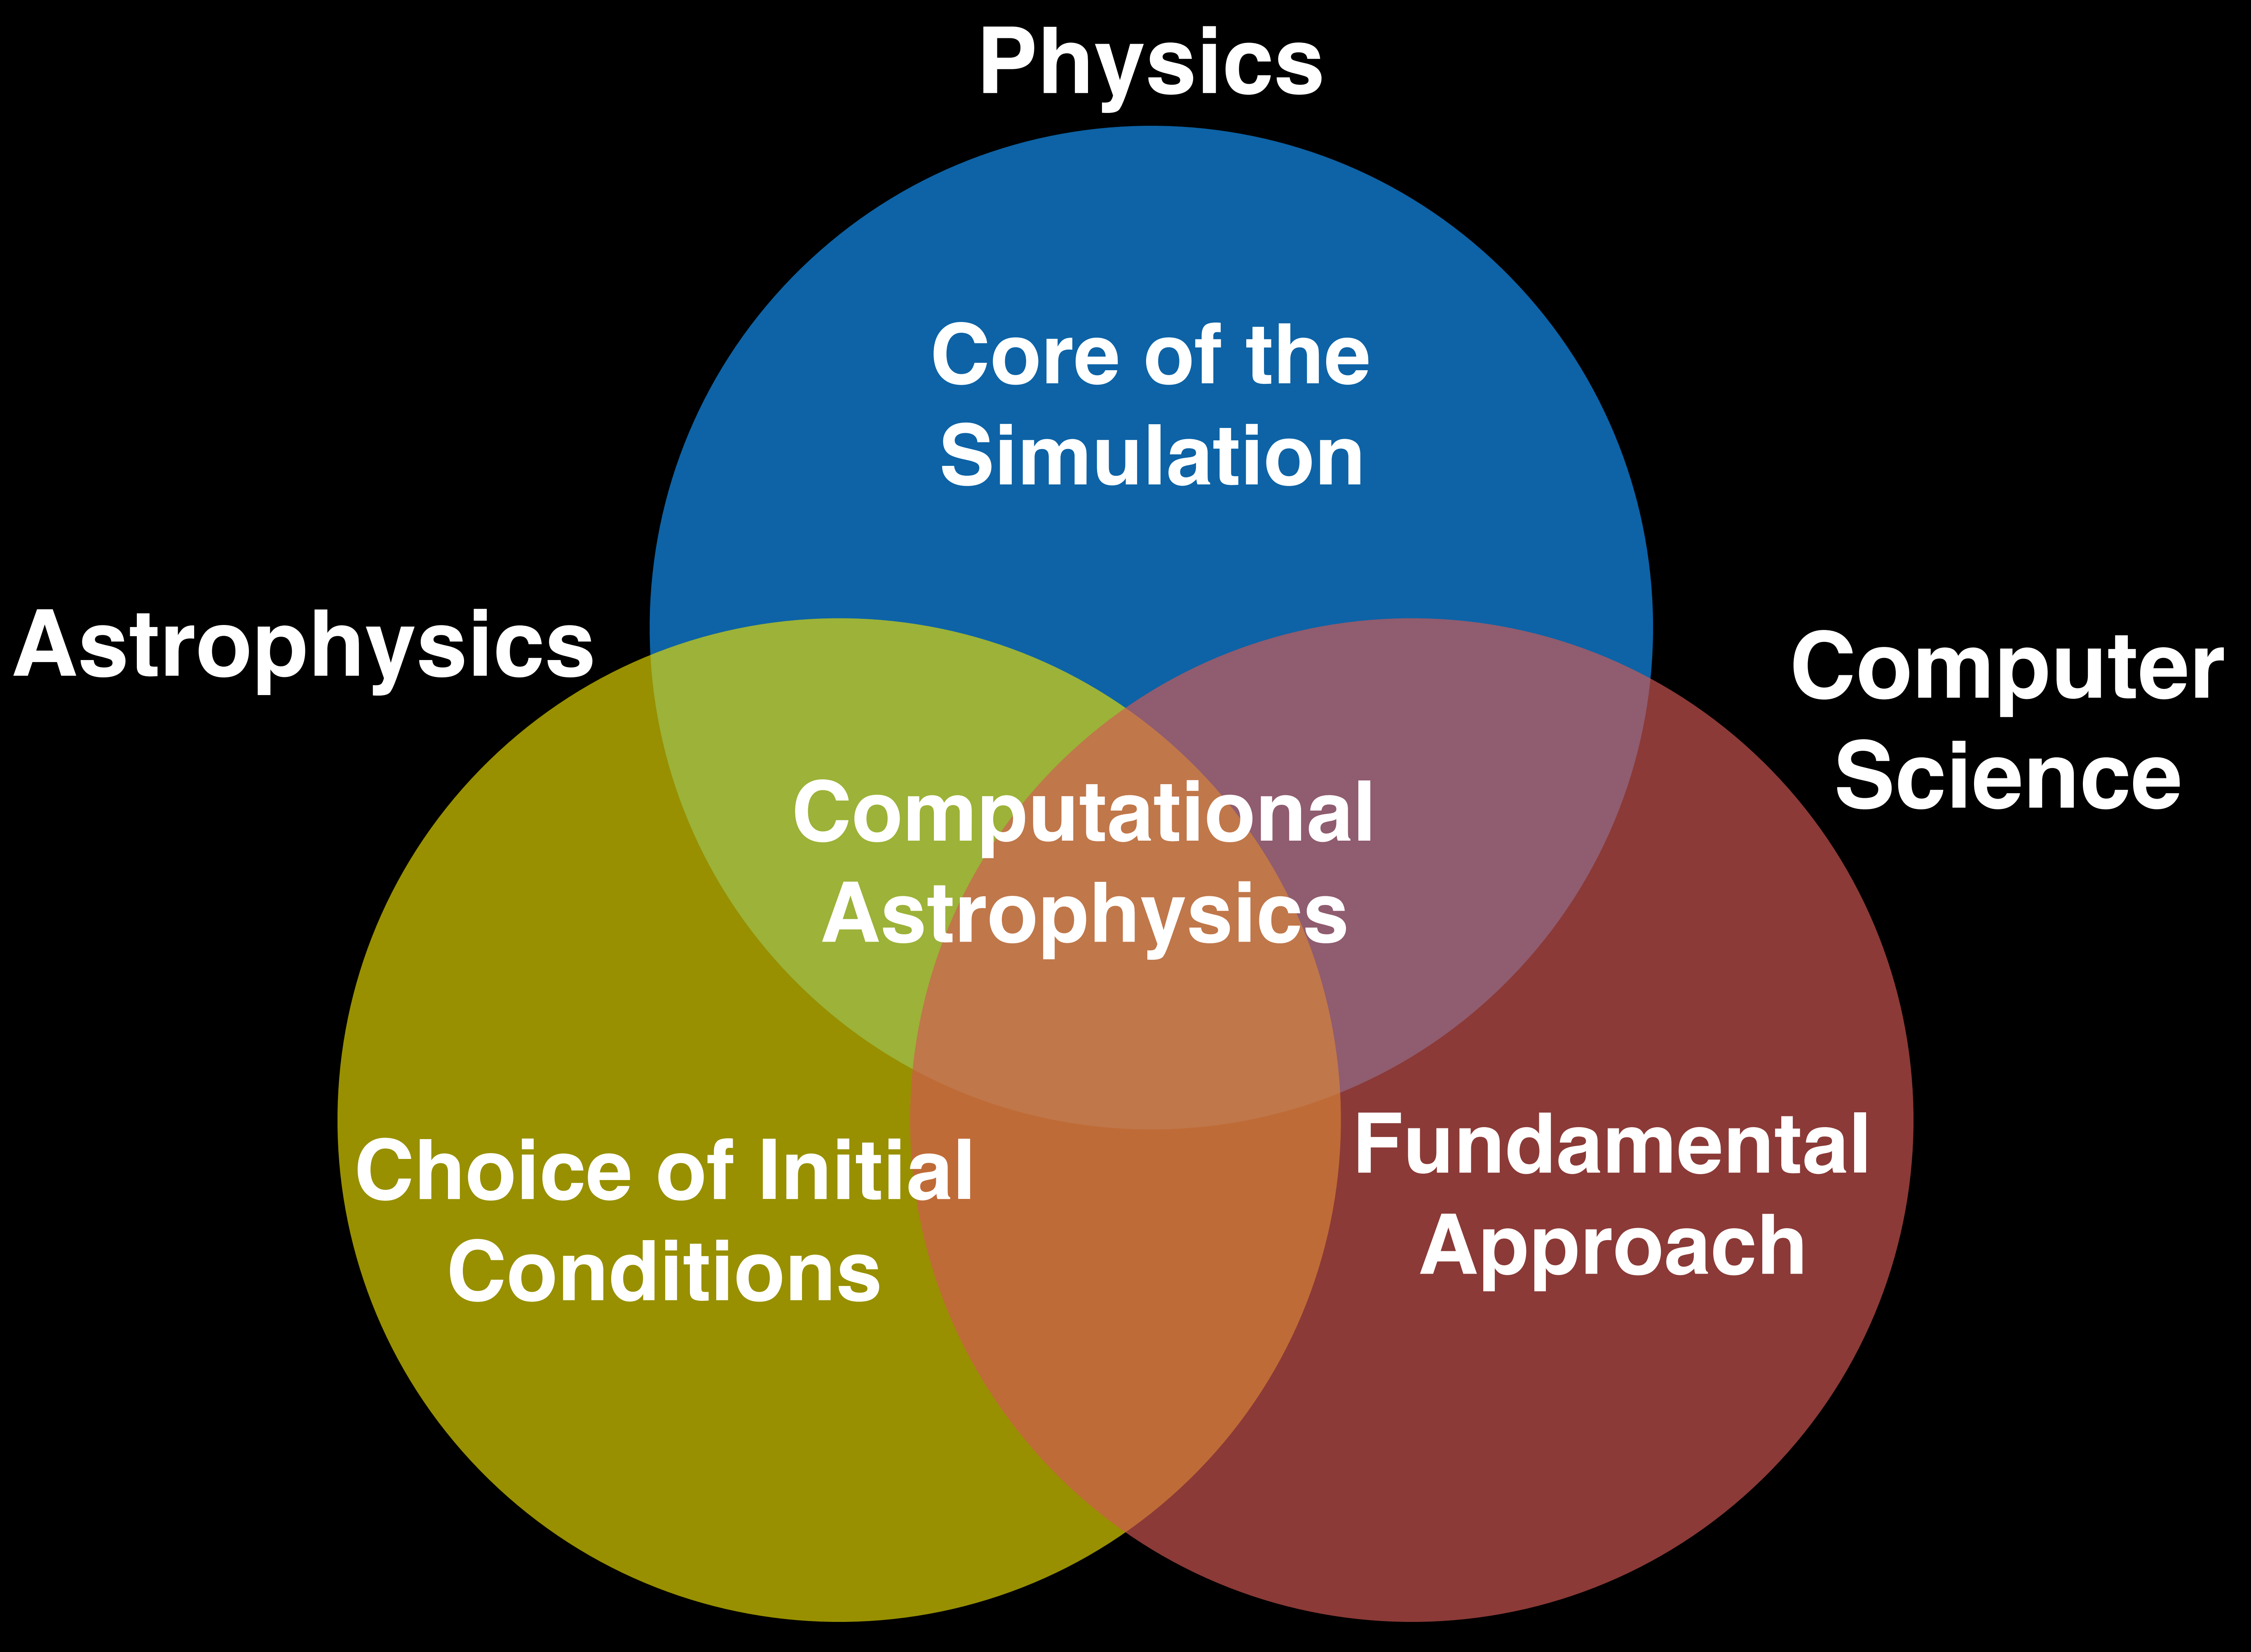
Physics
Core of the Simulation
Astrophysics
Computer Science
Computational Astrophysics
Choice of Initial Conditions
Fundamental Approach

Generic Simulation Approach

Discretize Medium and Equations of Motion

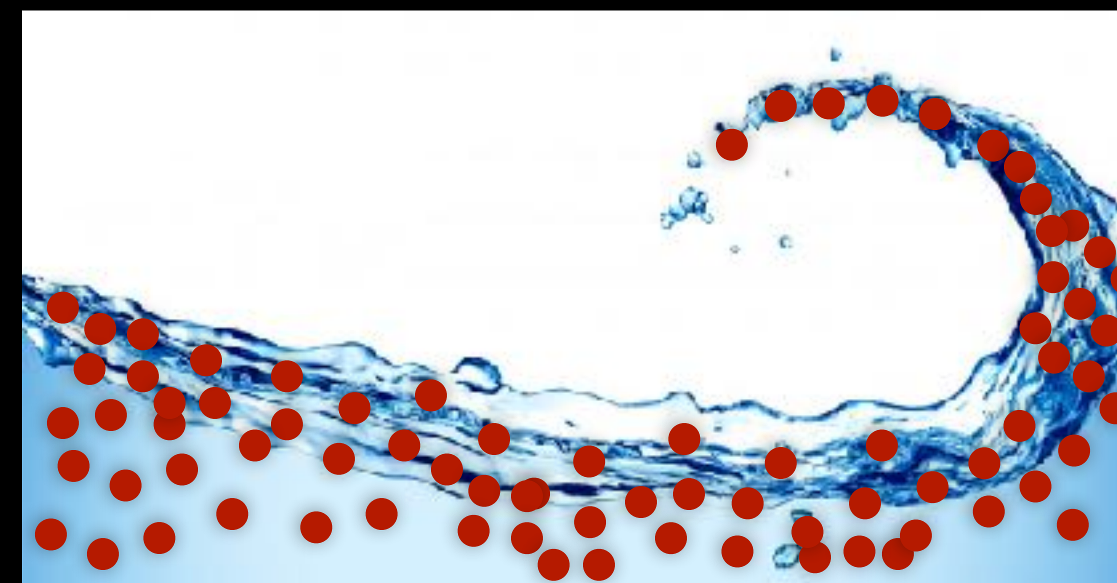

$$\rho = \rho(\vec{x})$$

$$\vec{v} = \vec{v}(\vec{x})$$

$$\vec{a} = \vec{a}(\vec{x})$$

Physical Input

Simulation Approach

$$\rho_i \approx \sum_j m_j W_{ij}$$

$$\vec{v}_i = \vec{v}_i$$

$$\vec{a}_i = \vec{a}_i(\vec{x}, \vec{v}, \rho, \ldots)$$

Cosmological Simulation Approach

Implement Relevant Physics

1. Gravity
    - Drives structure formation
    - Influences all matter on large scales

2. Hydrodynamics
    - Affects flow of gas within and around galaxies
    - Computationally more challenging than gravity

3. "Galaxy Formation Physics"
    - Various processes that shape internal structure of galaxies
    - Star formation, gas cooling, "feedback" onto surrounding gas, etc.

Cosmological Simulation Approach

Set up initial conditions

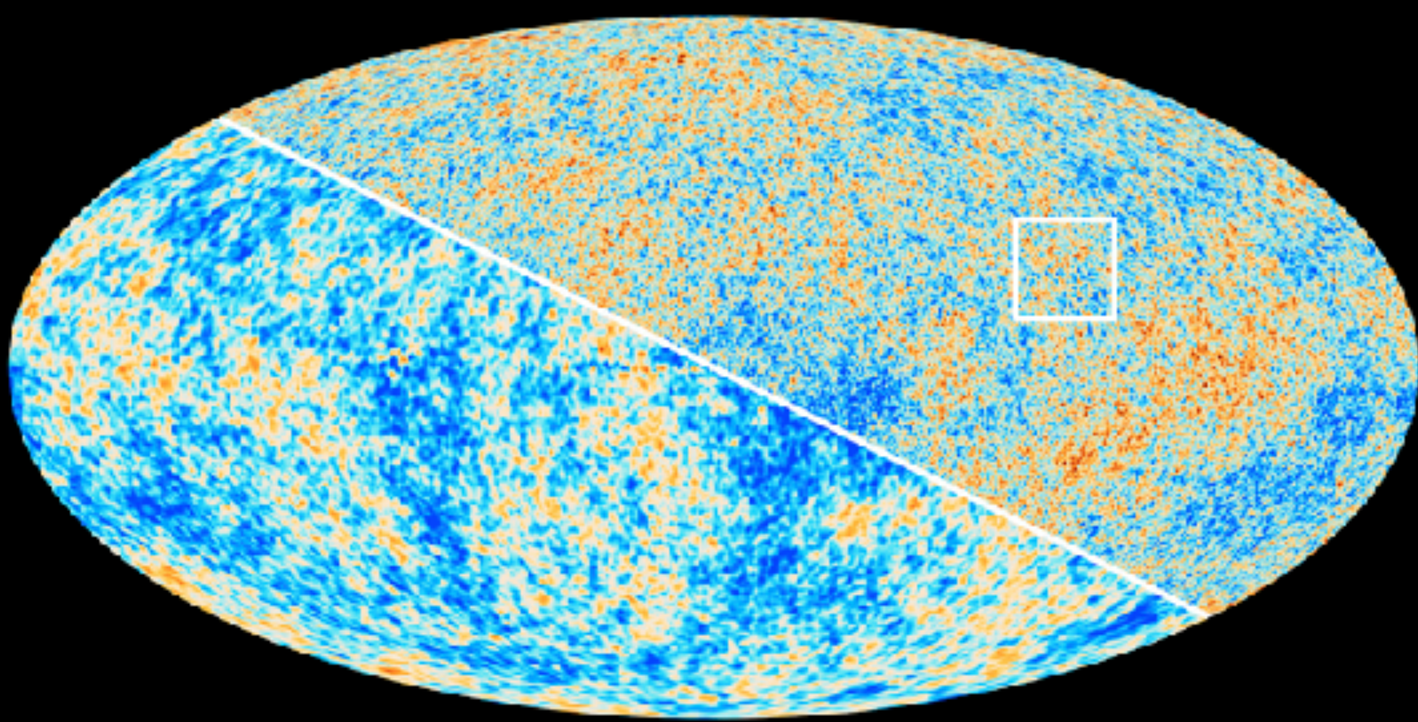

Use early universe cosmology to set initial **distribution** of matter

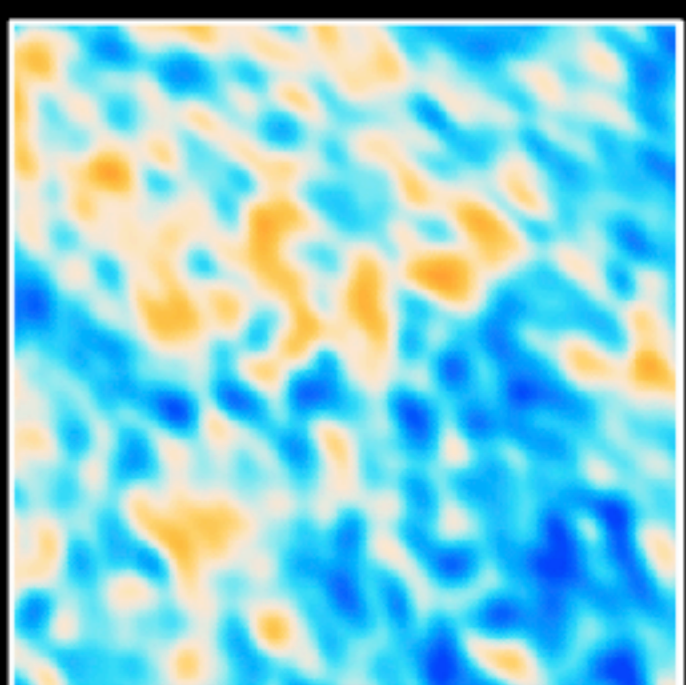

WMAP

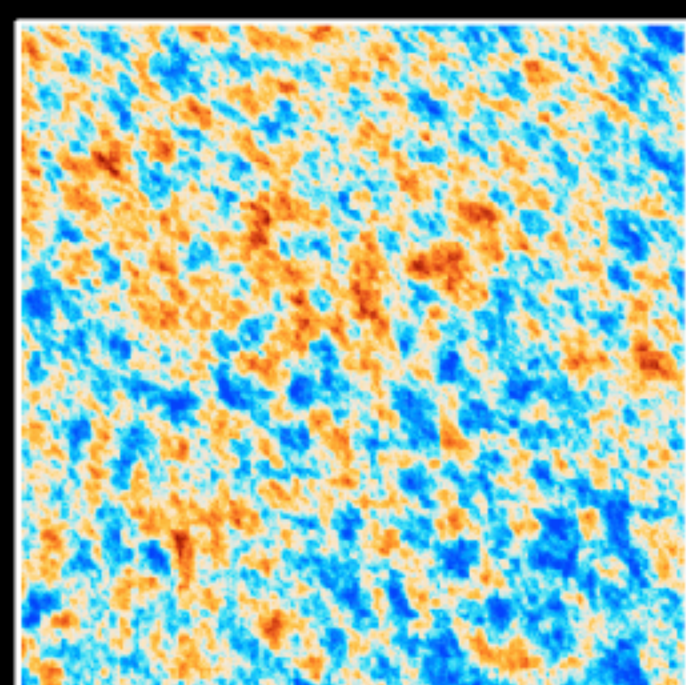

Planck

Early universe almost uniform, with ~10 parts per million fluctuations

Millennium-II Simulation

http://www.mpa-garching.mpg.de/galform/millennium-II

Mike Boylan-Kolchin

Max Planck Institute for Astrophysics

Cosmological Simulation Approach

Implement Relevant Physics

1. Gravity
   - Drives structure formation
   - Influences all matter on large scales

2. Hydrodynamics
   - Affects flow of gas within and around galaxies
   - Computationally more challenging than gravity

3. "Galaxy Formation Physics"
   - Various processes that shape internal structure of galaxies
   - Star formation, gas cooling, "feedback" onto surrounding gas, etc.

Cosmological Simulation Approach

Hydrodynamical Methods

Lagrangian Methods

(SPH) e.g., GADGET

Eulerian Methods

(AMR) e.g., ENZO

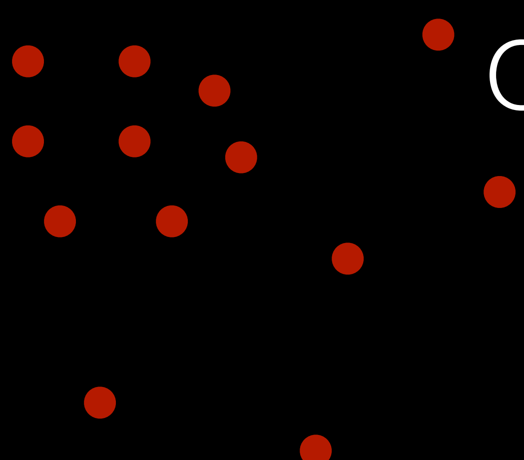

Quasi-Lagrangian Finite-Volume

(moving mesh) e.g., AREPO

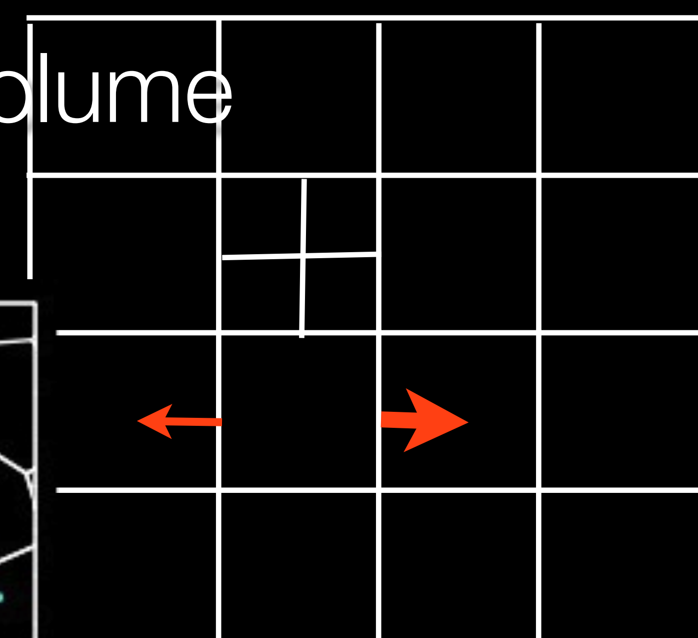

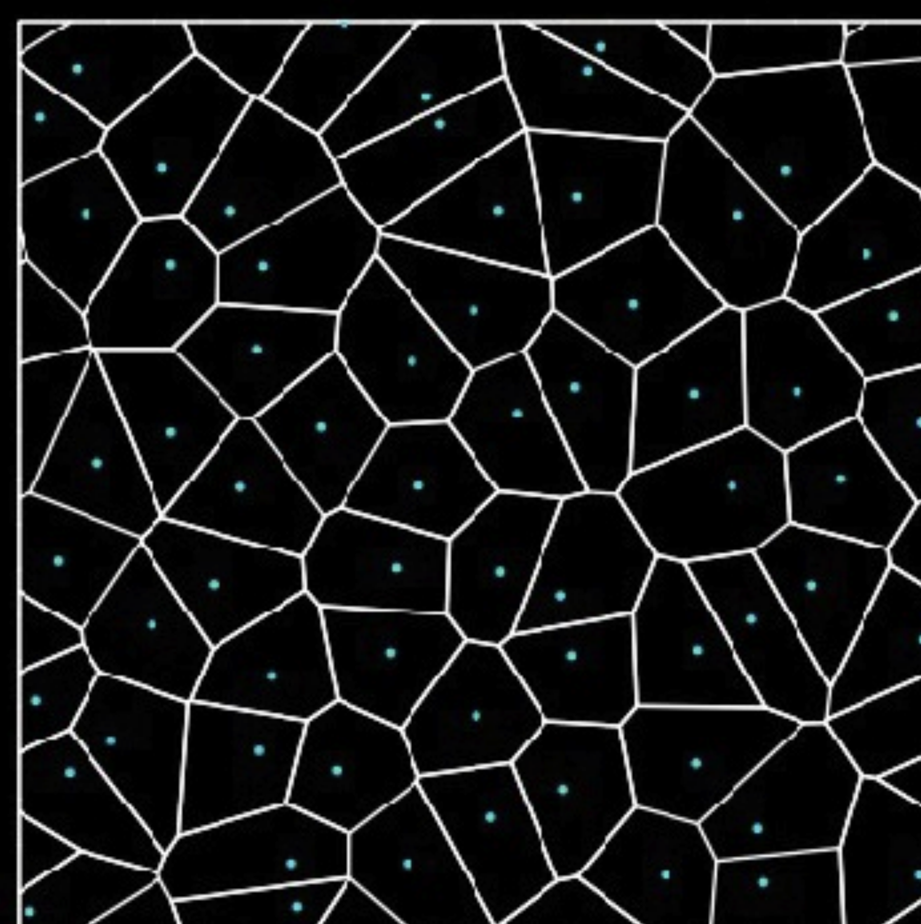

- Geometrically flexible
- Naturally adaptive resolu
- Galilean invariant

stability handling

atural shock capturing

hase boundary resolution

Springel (2010)

Cosmological Simulation Approach

Example of Kelvin-Helmholtz Instability

Earth

Jupiter

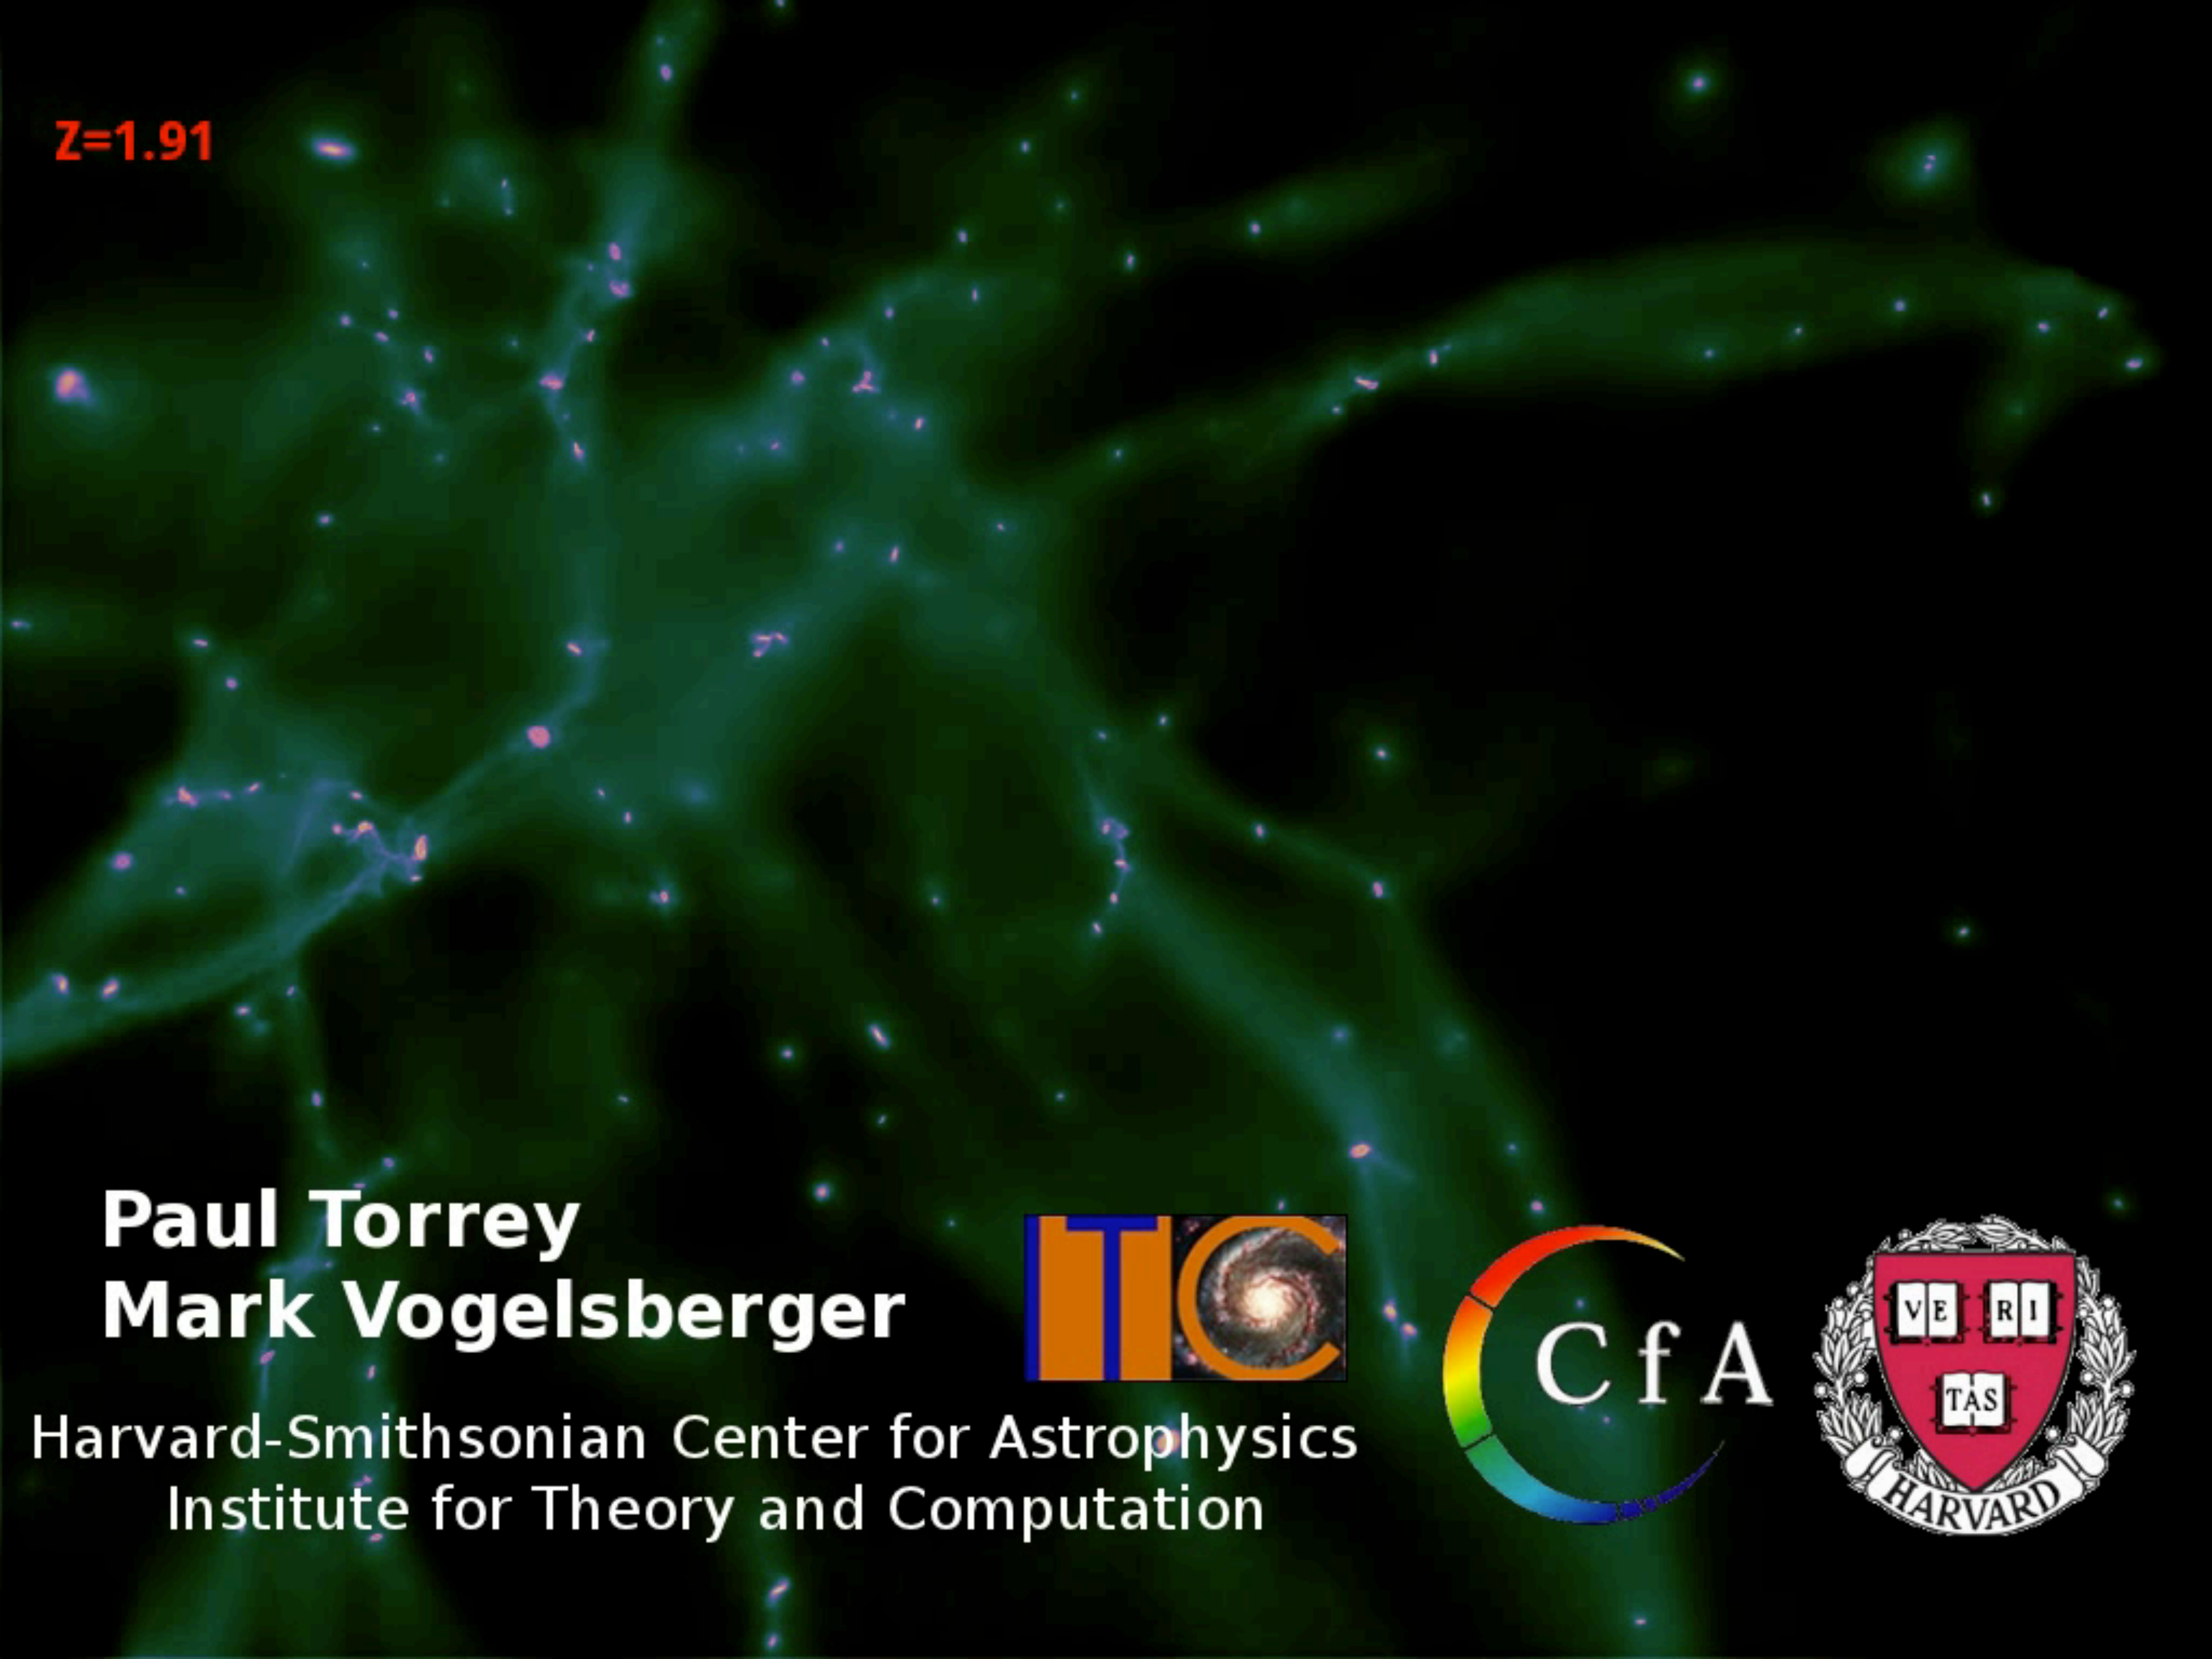
Z=1.91
Paul Torrey
Mark Vogelsberger
Harvard-Smithsonian Center for Astrophysics
Institute for Theory and Computation

Are we done?

What other physics do we need to include?

Science with Simulations

Importance of Feedback on Galaxy Growth

Star formation is **too efficient** unless we account for physics that can regulate growth of galaxies!

Simulated Sky

Real Sky

Cosmological Simulation Approach

Implement Relevant Physics

1. Gravity
    - Drives structure formation
    - Influences all matter on large scales

2. Hydrodynamics
    - Affects flow of gas within and around galaxies
    - Computationally more challenging than gravity

3. "Galaxy Formation Physics"
    - Various processes that shape internal structure of galaxies
    - Star formation, gas cooling, "feedback" onto surrounding gas, etc.

Science with Simulations

The Illustris Collaboration

Core Collaboration Members:

- Mark Vogelsberger (MIT)
- Paul Torrey (MIT)
- Shy Genel (Columbia)
- Debora Sijacki (Cambridge)
- Volker Springel (HITS)
- Lars Hernquist (Harvard)

MIT KAVLI INSTITUTE

Heidelberg Institute for Theoretical Studies HITS

HARVARD UNIVERSITY

UNIVERSITY OF CAMBRIDGE

COLUMBIA UNIVERSITY

Current Collaboration Status:

- ~30 Active Members
- ~10 Institutions
- Wide range of expertise and interests

Science with Simulations

The Illustris Simulation

Science with Simulations

Testing Various Physics Models

Simulated Sky (1)

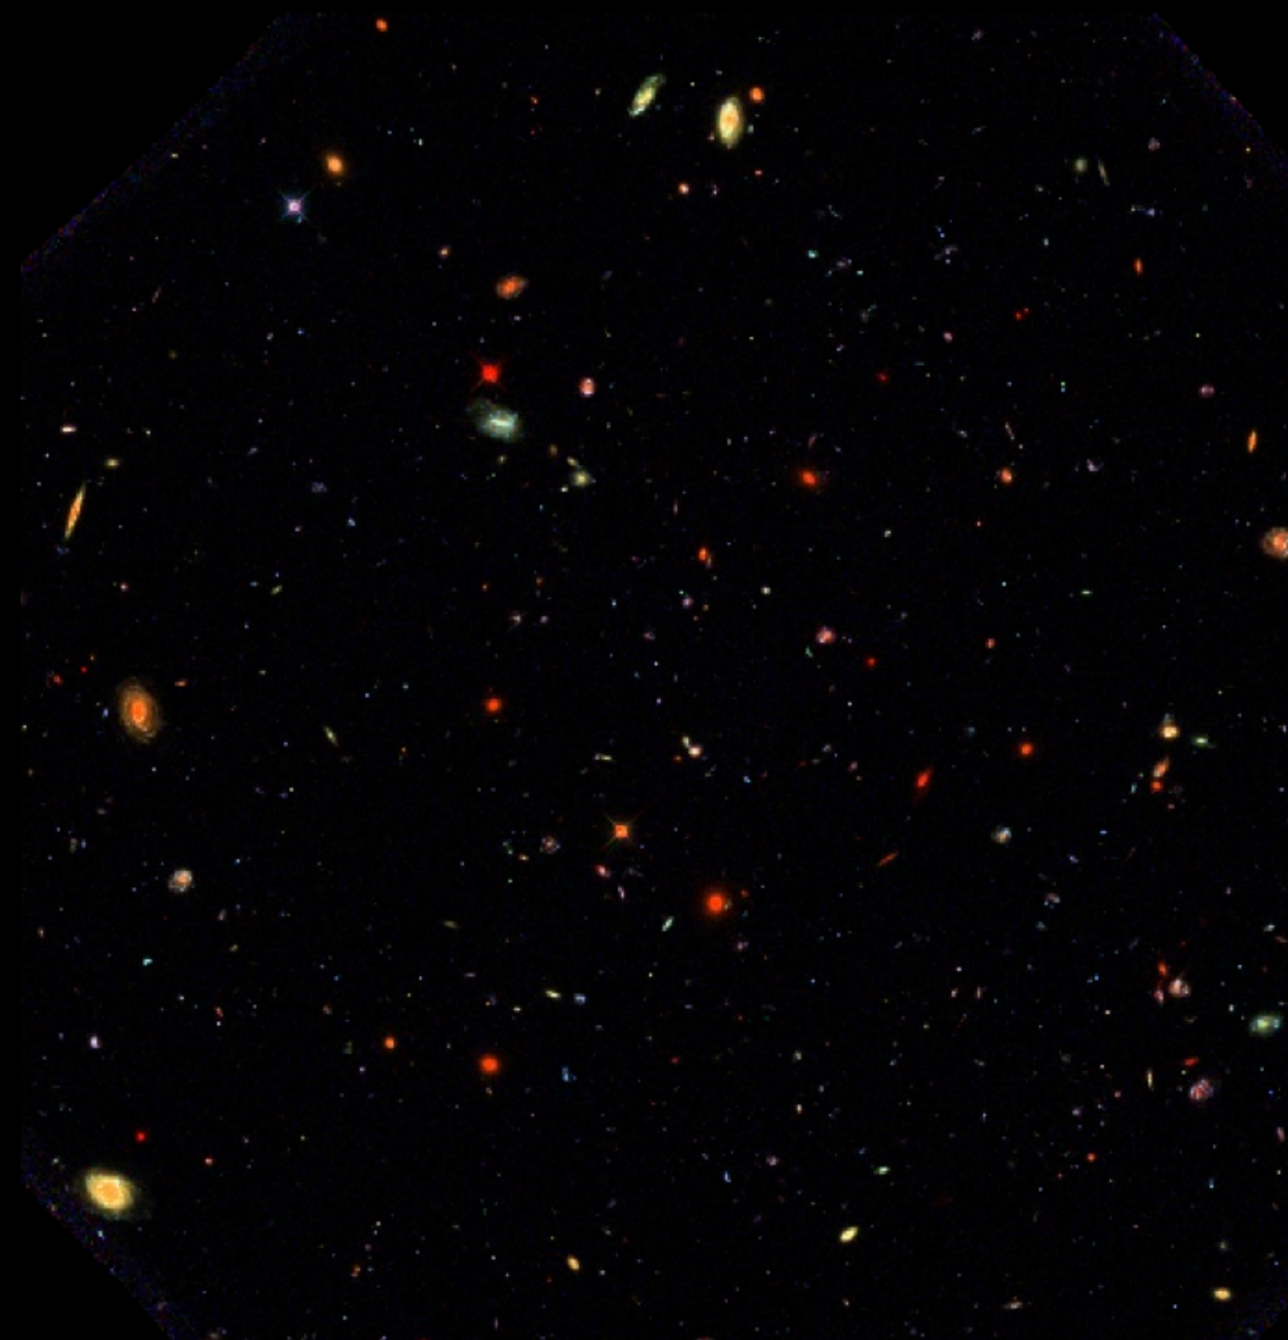

Real Sky

Science with Simulations

Testing Various Physics Models

Simulated Sky (2)

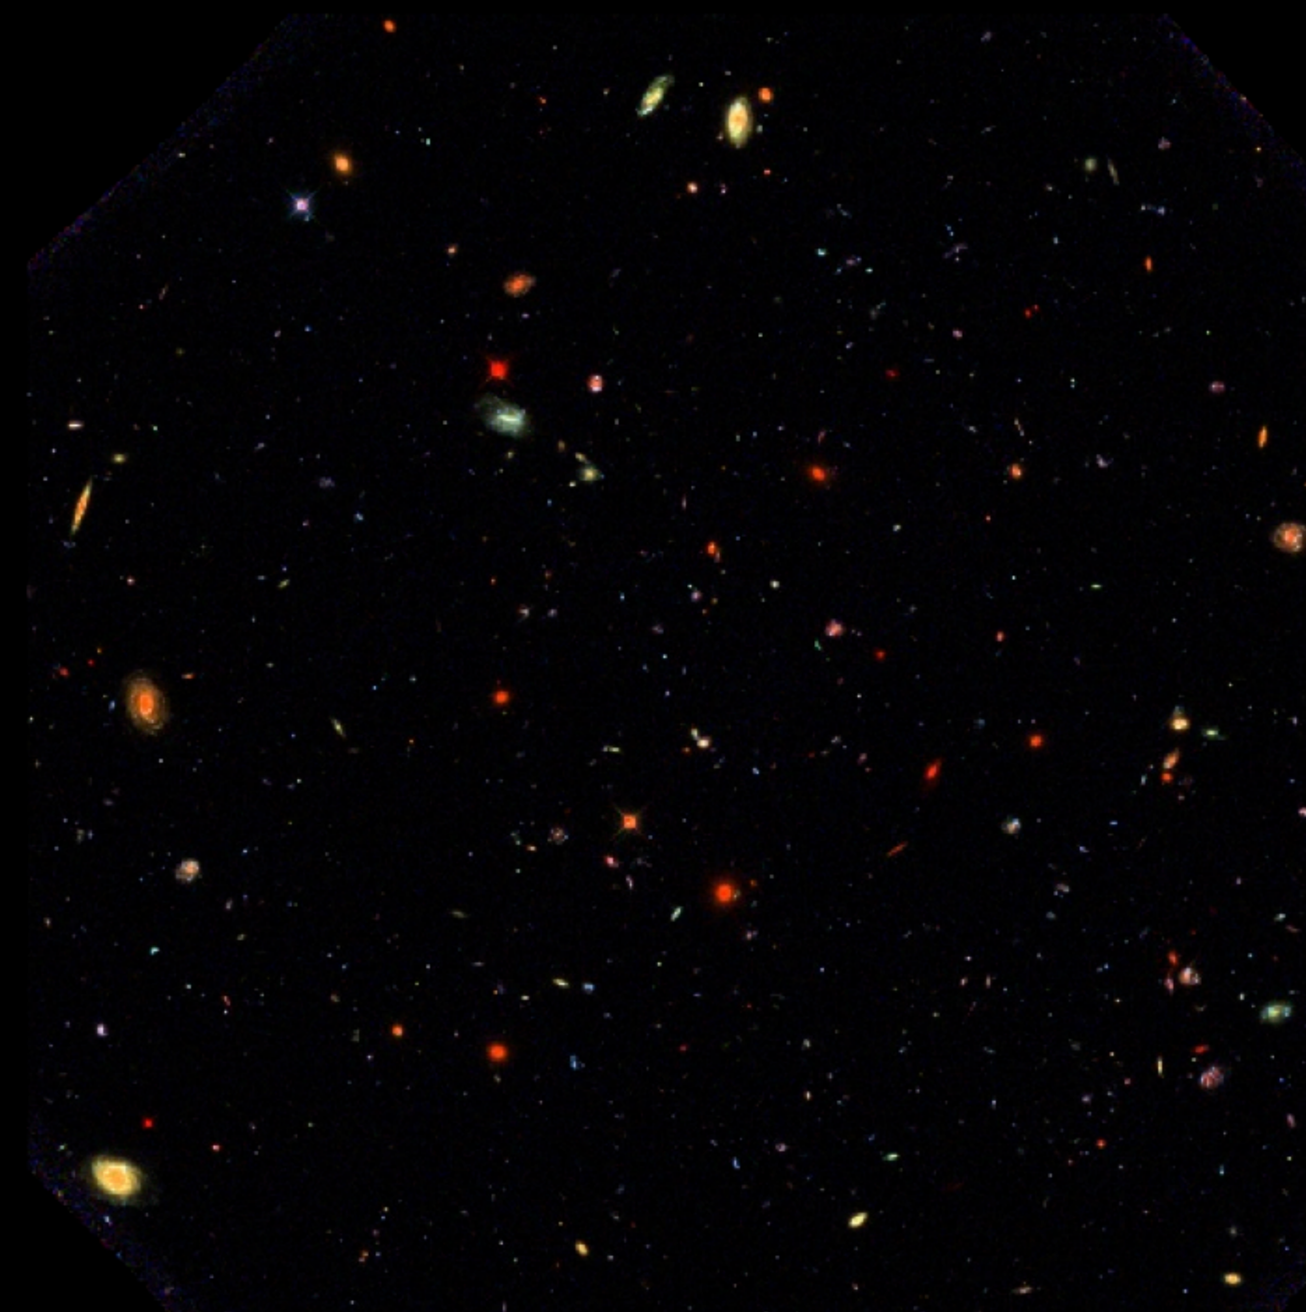

Real Sky

Science with Simulations

Testing Various Physics Models

Simulated Sky (3)

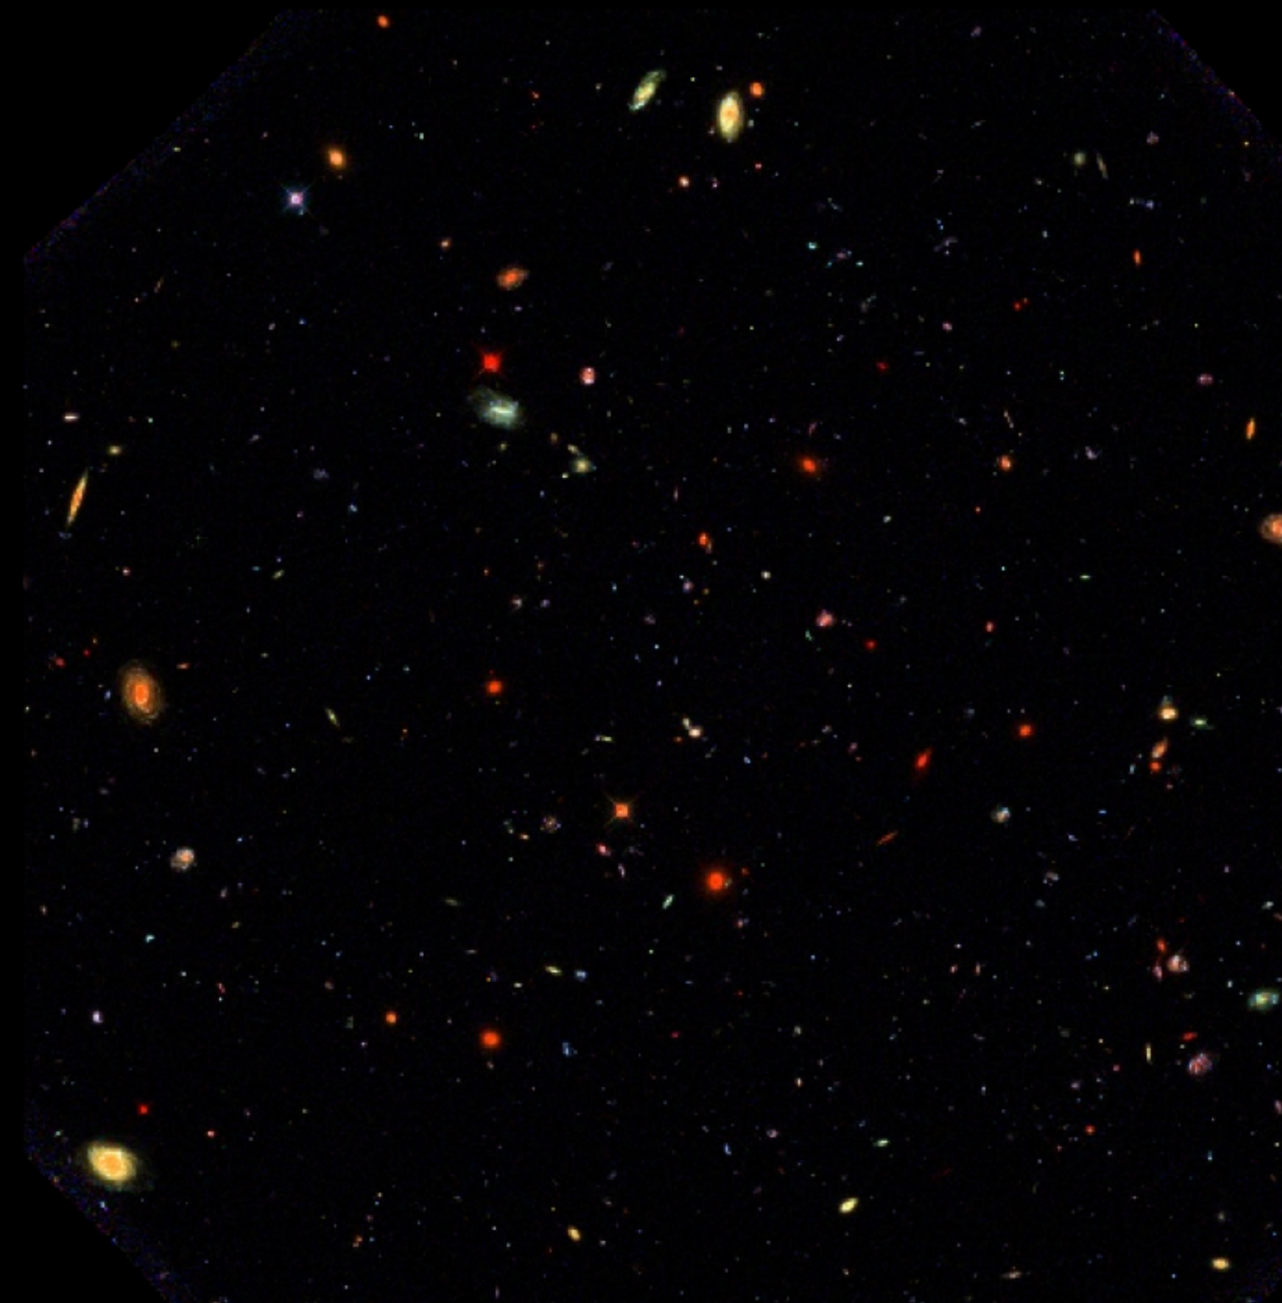

Real Sky

Science with Simulations

Testing Various Physics Models

Simulated Sky (4)

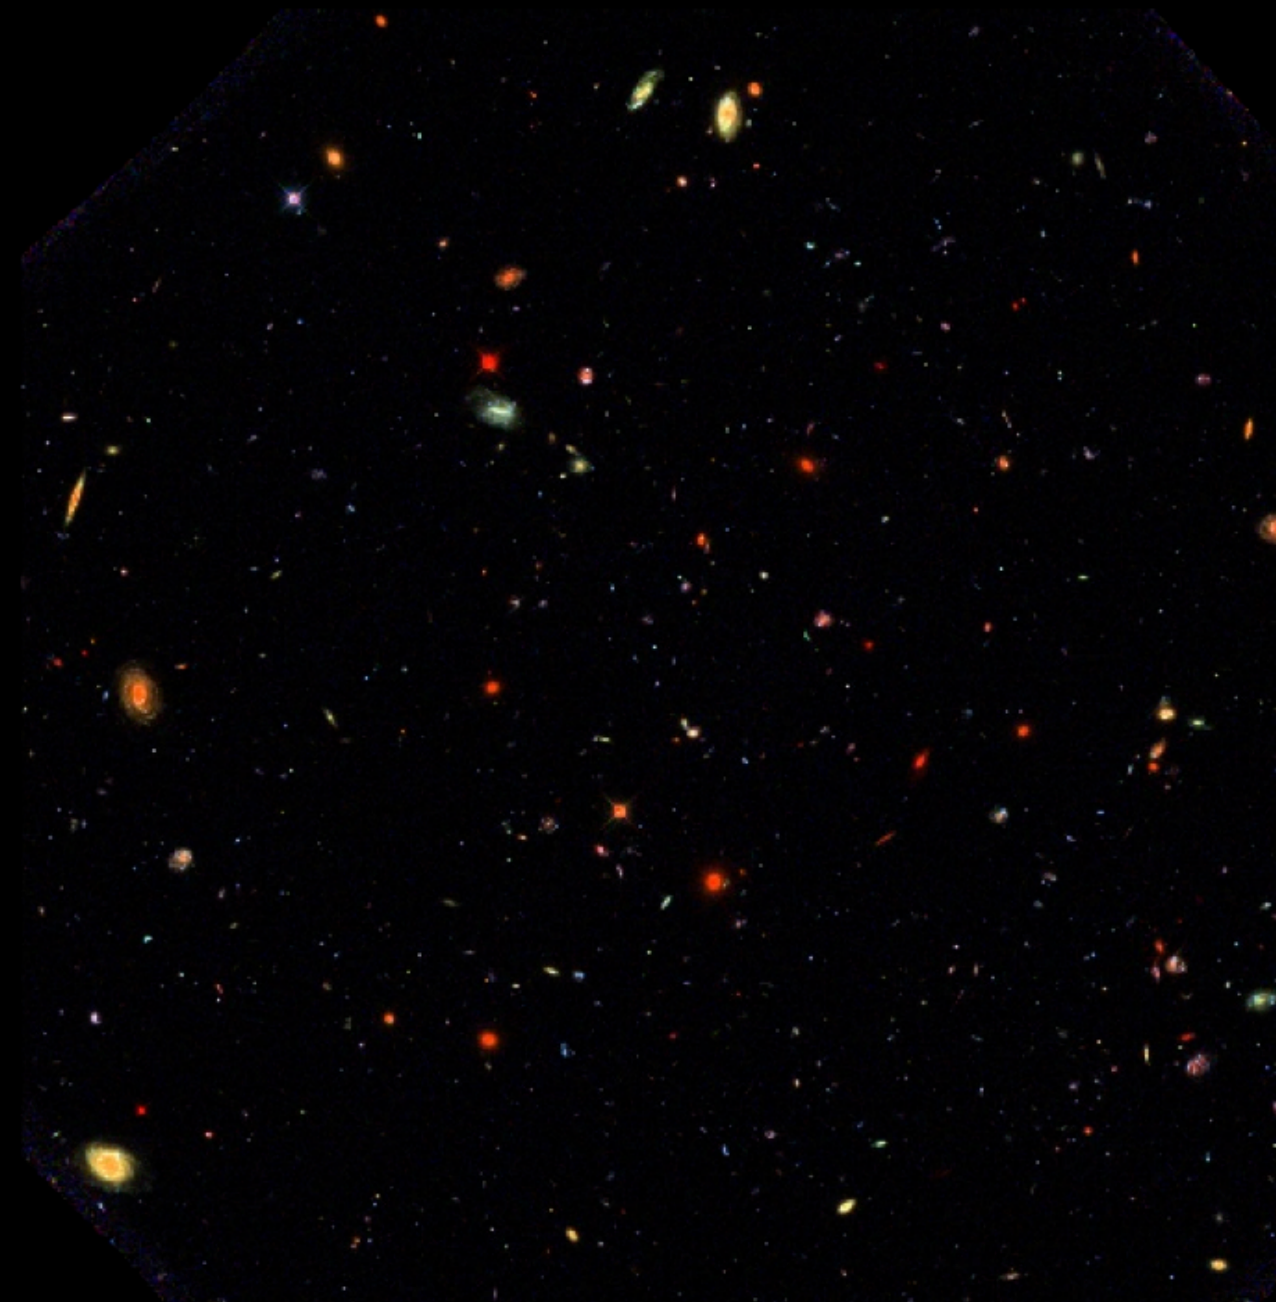

Real Sky

Computing Trends and Projections

How Big is Illustris?

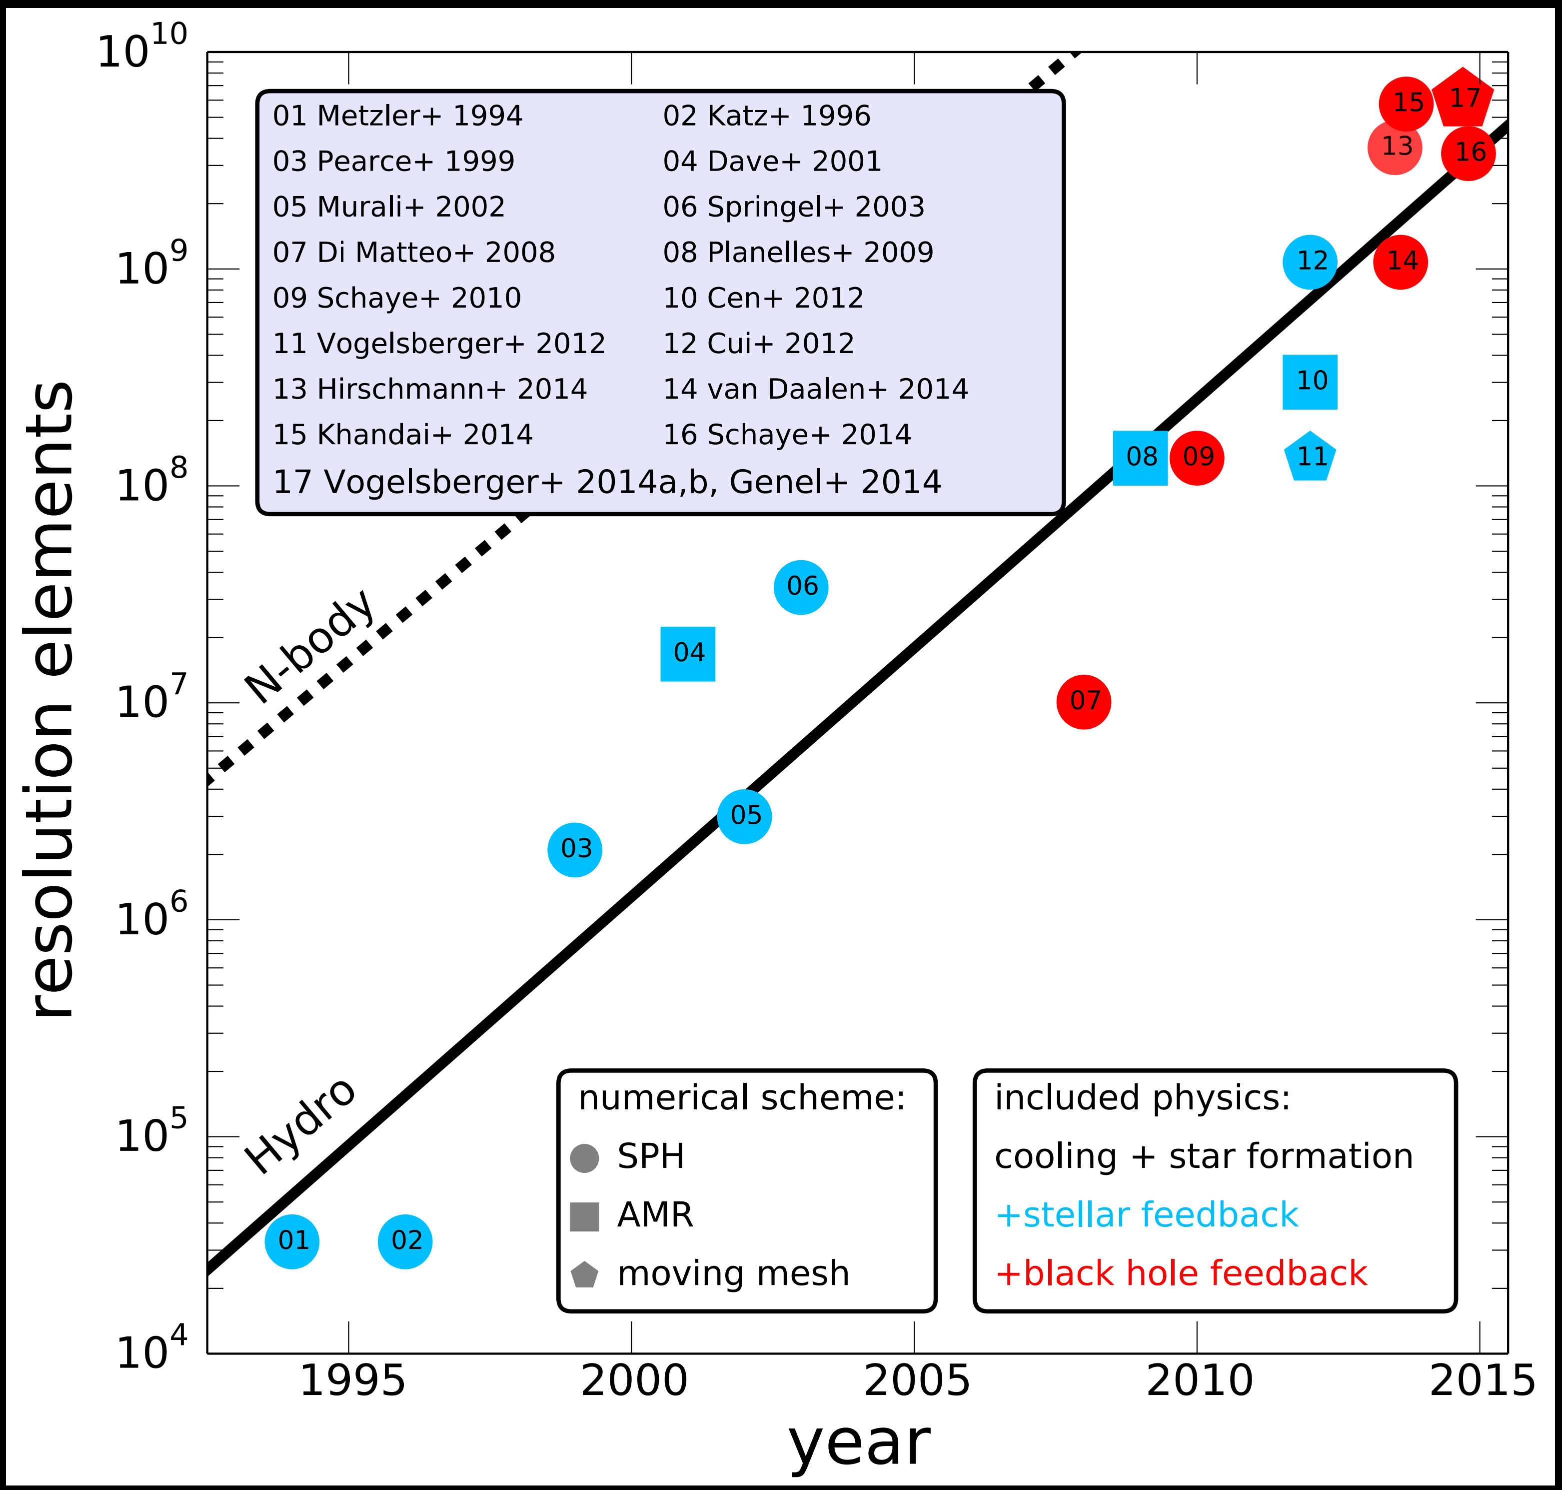

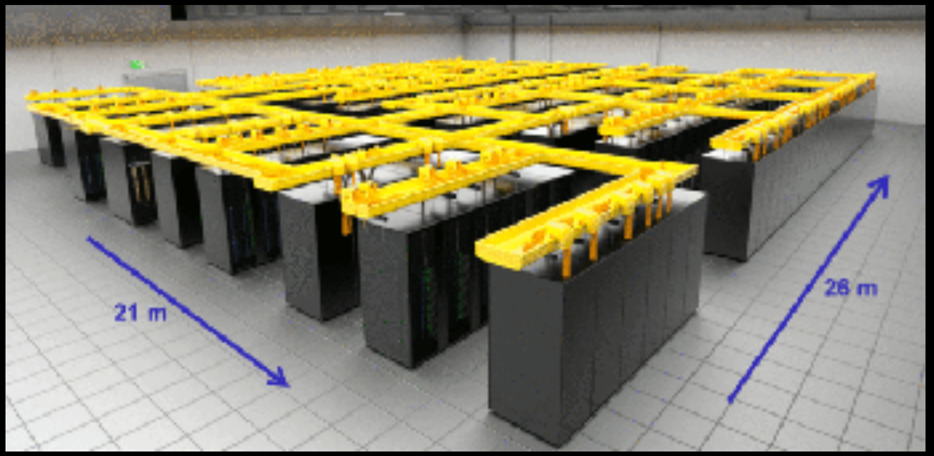

Computing Trends and Projections

Moore's Law: The Good

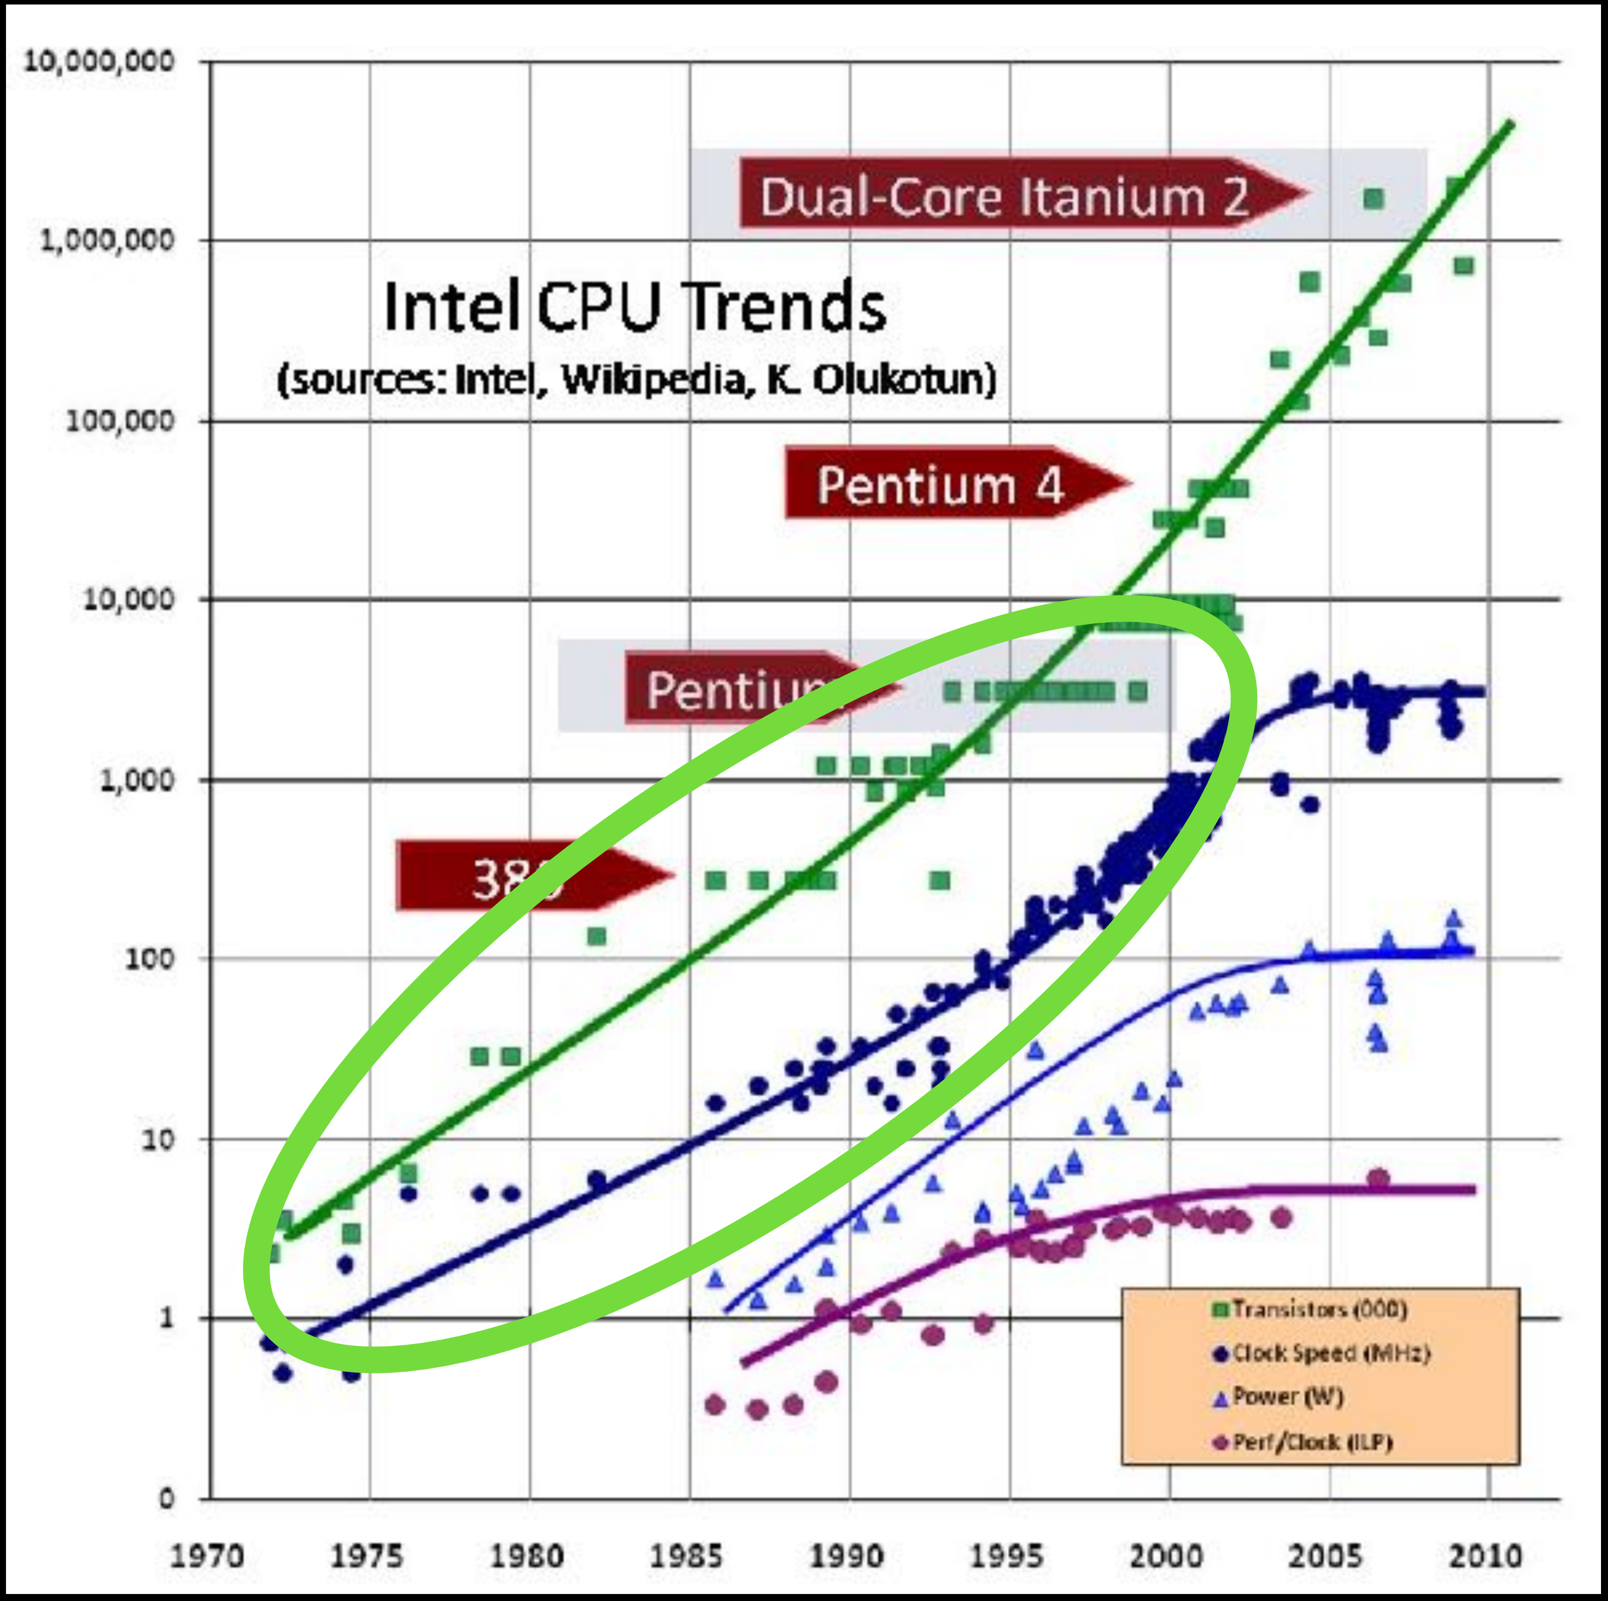

Computing Trends and Projections

Moore's Law: The Good

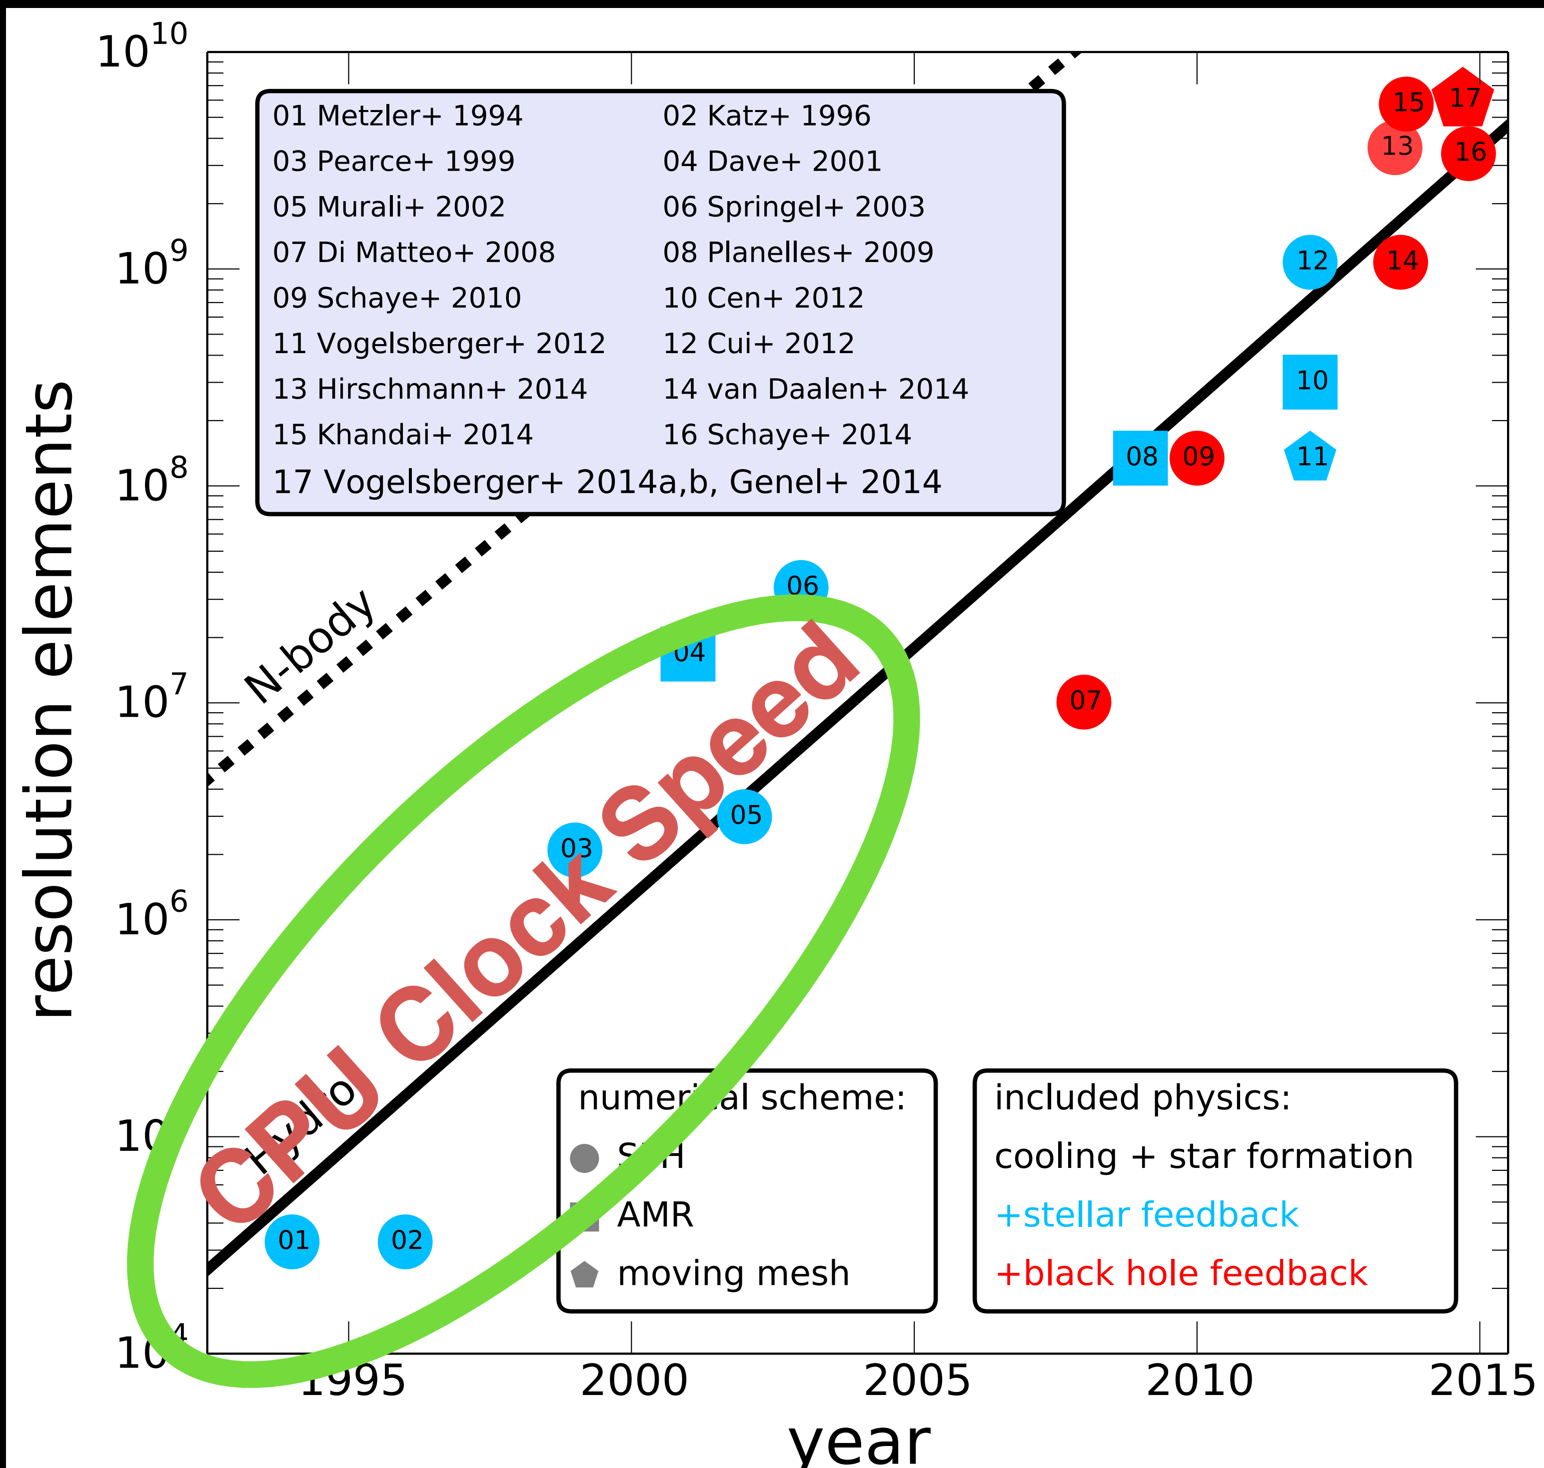

Computing Trends and Projections

Moore's Law: The Bad

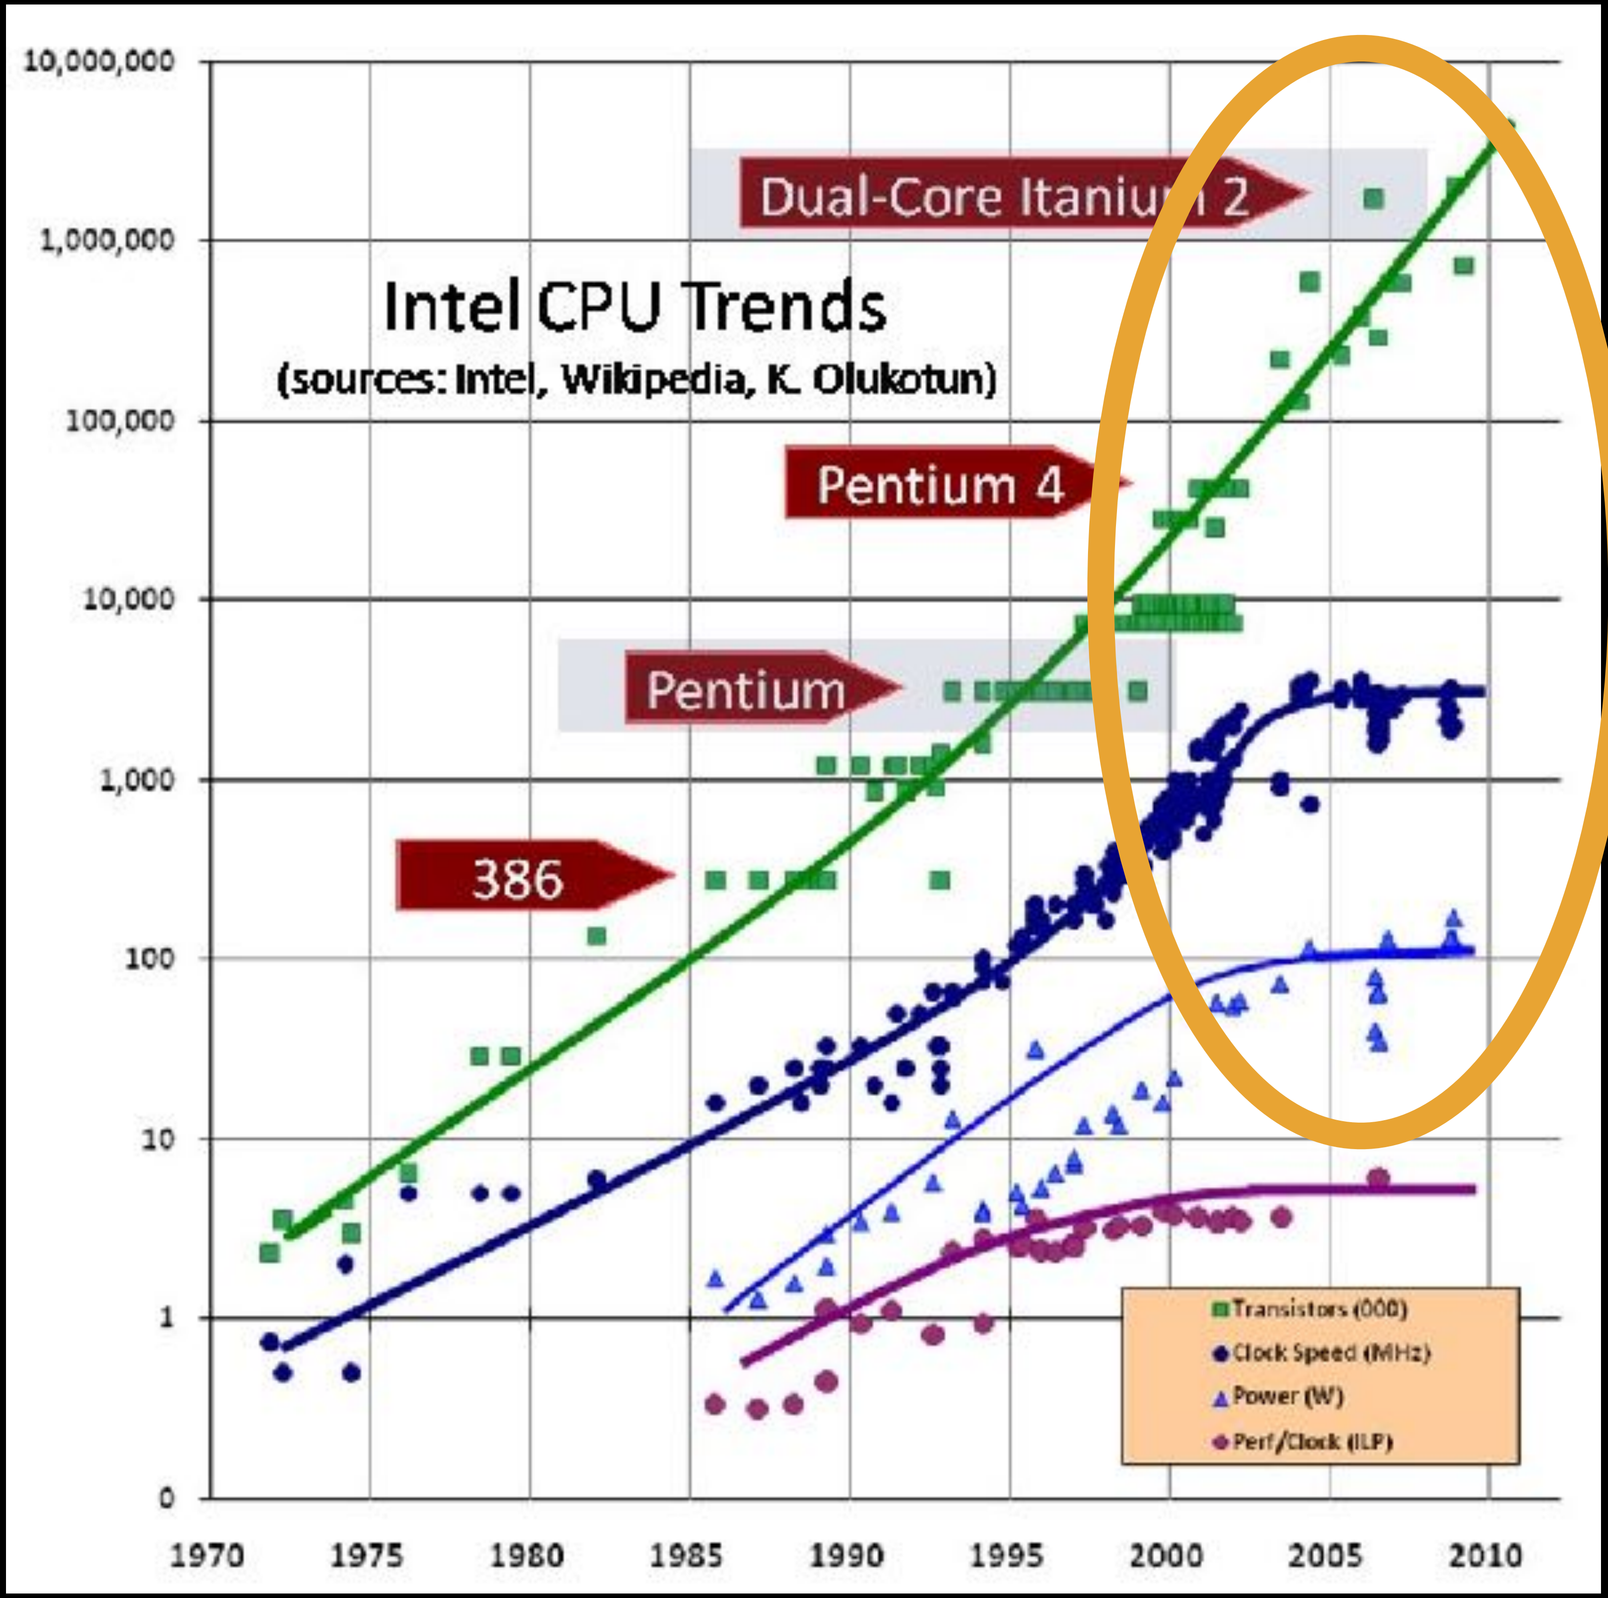

Computing Trends and Projections

Moore's Law: The Bad

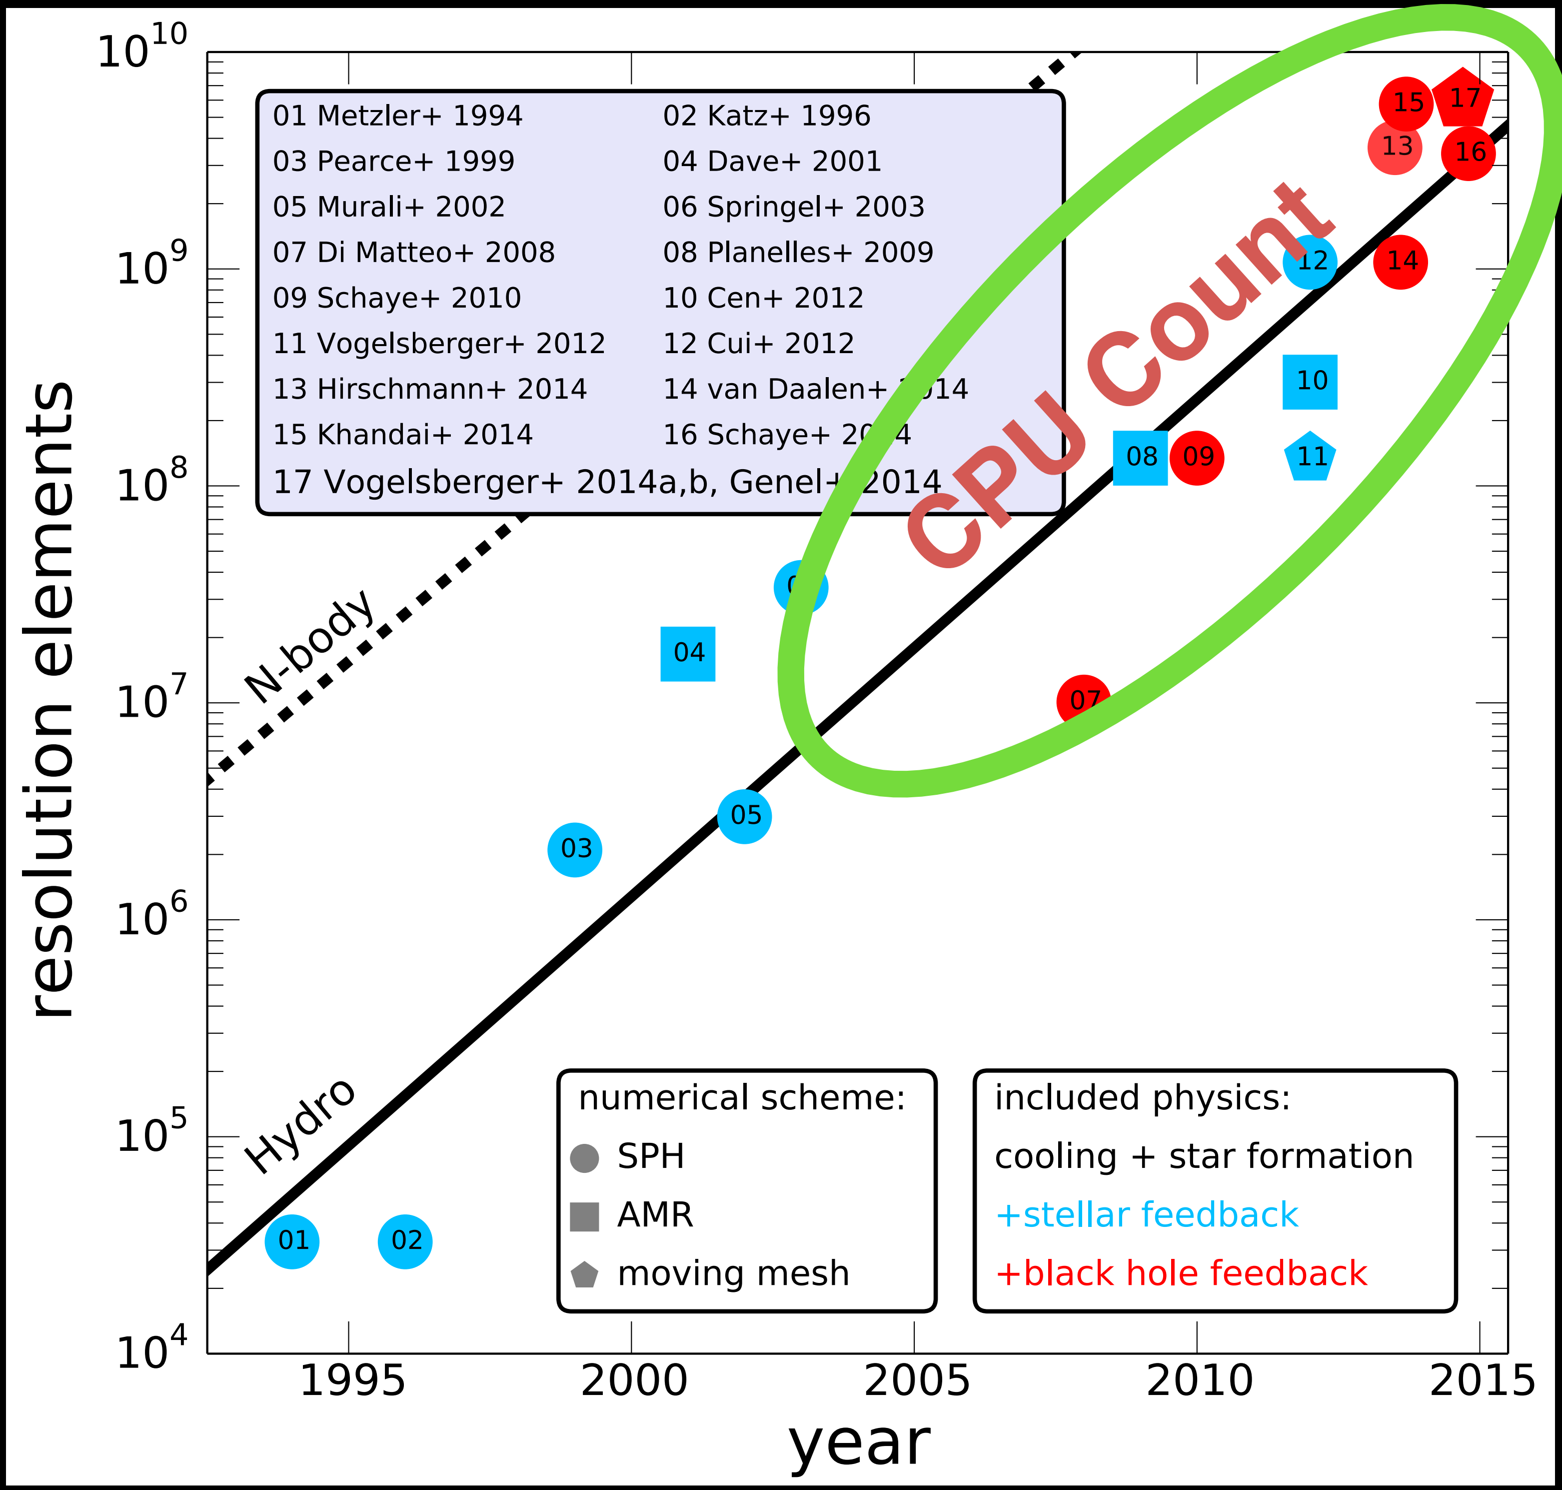

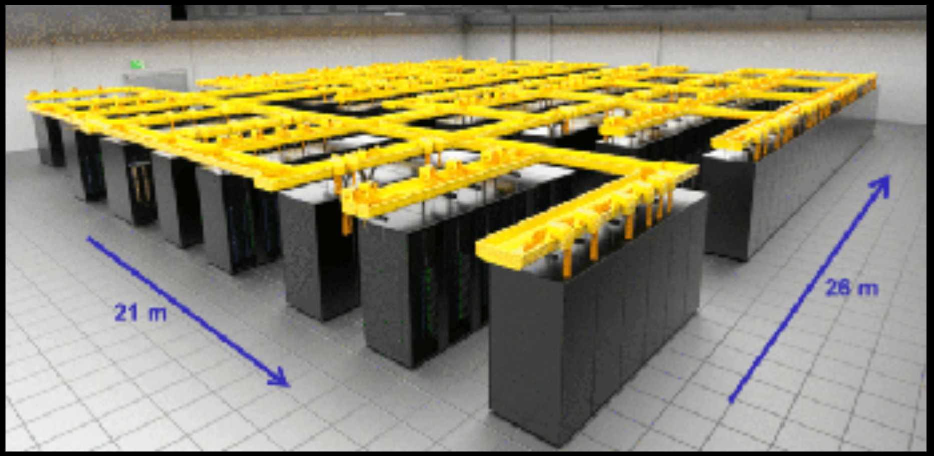

Computing Trends and Projections

Moore's Law: The Bad

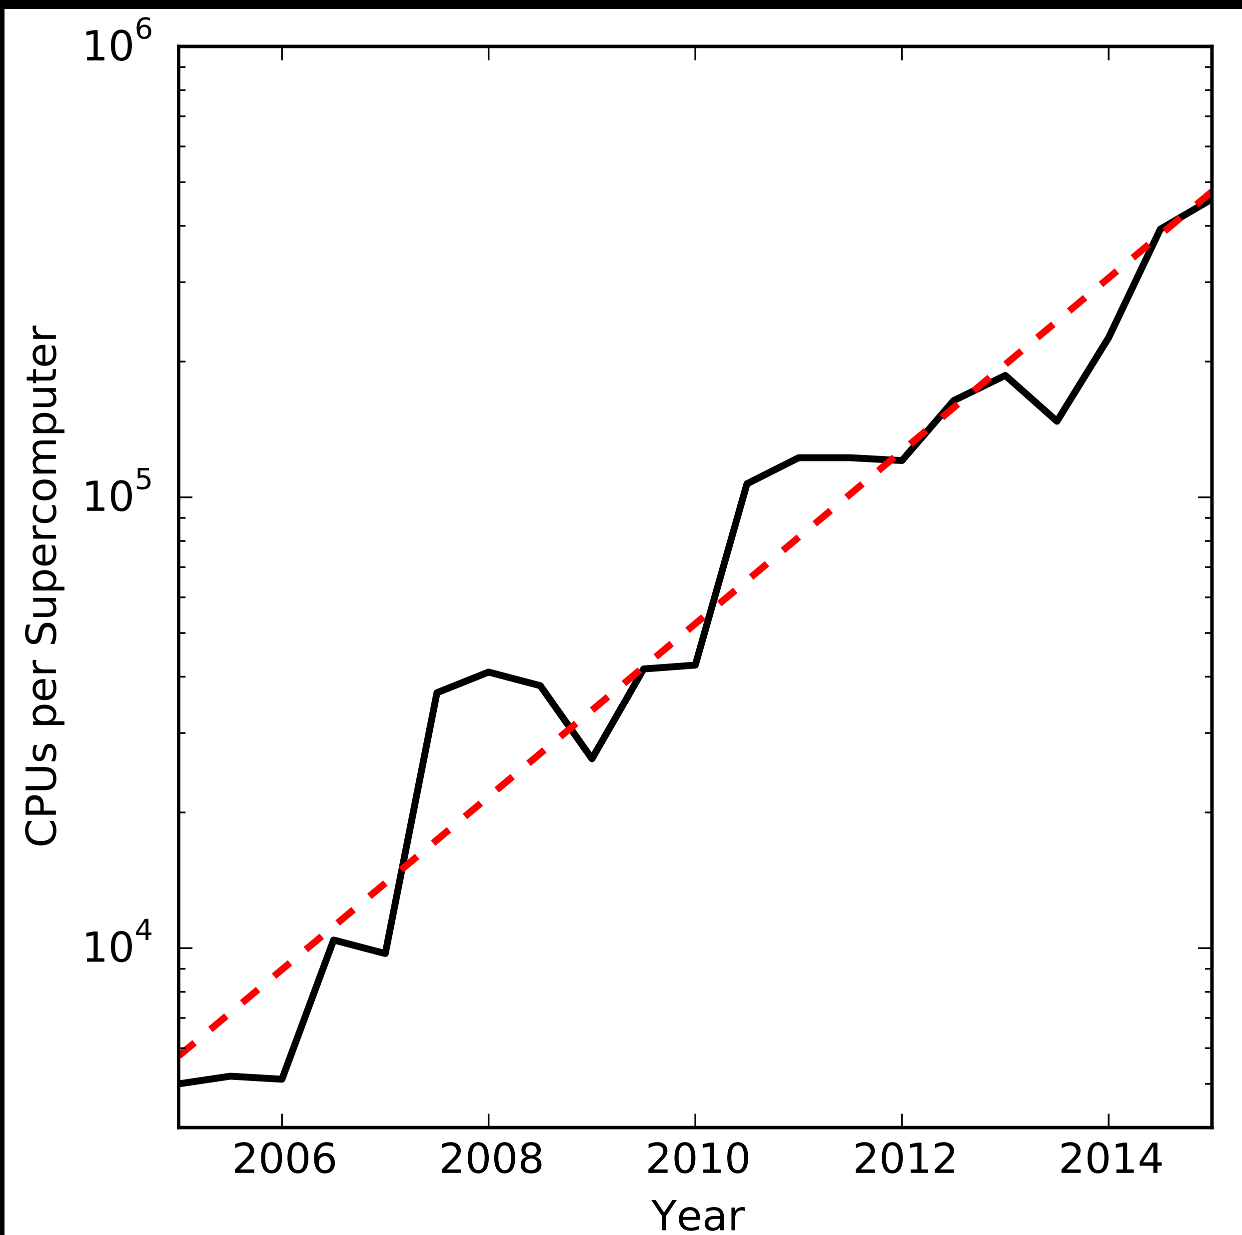

TOP 500 The List.

www.top500.org/lists/2015/06/

RANK	SITE	CORES	RMAX (TFLOP/S)	RPEAK (TFLOP/S)	POWER (KW)
1	National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
2	DOE/SC/Oak Ridge National Laboratory United States	560,640	17,590.0	27,112.5	8,209
3	DOE/NNSA/LLNL United States	1,572,864	17,173.2	20,132.7	7,890
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	705,024	10,510.0	11,280.4	12,660
5	DOE/SC/Argonne National Laboratory United States	786,432	8,586.6	10,066.3	3,945
6	Swiss National Supercomputing Centre (CSCS) Switzerland	115,984	6,271.0	7,788.9	2,325
7	King Abdullah University of Science and Technology Saudi Arabia	196,608	5,537.0	7,235.2	2,834
8	Texas Advanced Computing Center/Univ. of Texas United States	462,462	5,168.1	8,520.1	4,510

Computing Trends and Projections

Moore's Law: The Ugly

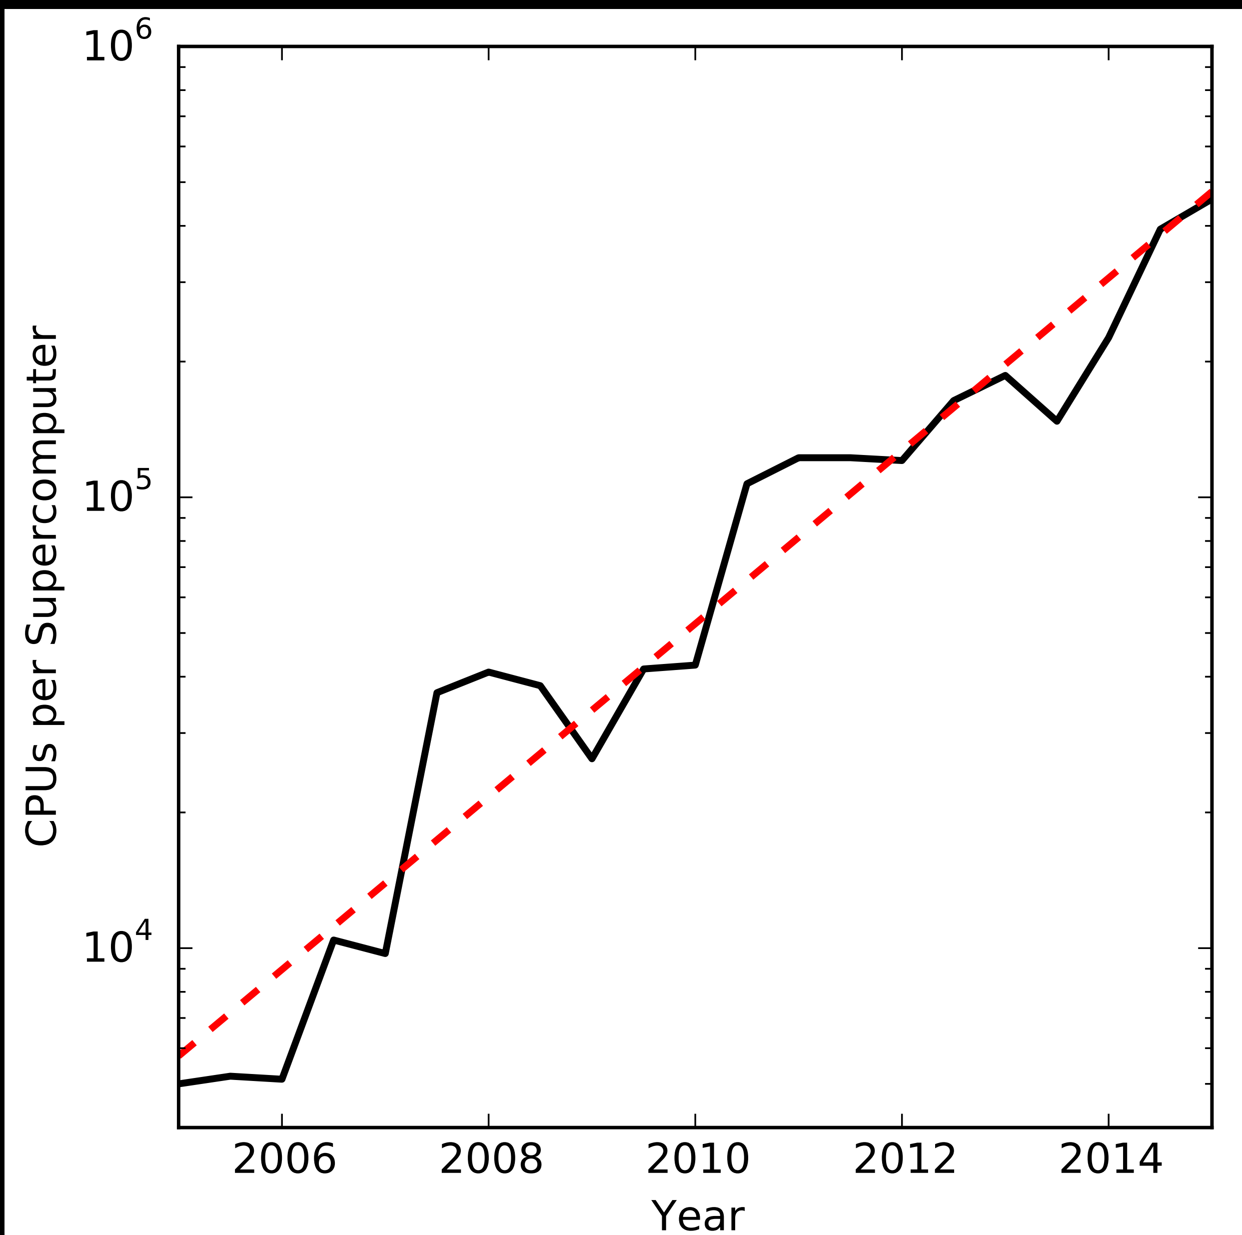

TOP 500
The List.

www.top500.org/lists/2015/06/

RANK	SITE	CORES	RMAX (TFLOP/S)	RPEAK (TFLOP/S)	POWER (KW)
1	National Super Computer Center in Guangzhou China	3,120,000	33,862.7	54,902.4	17,808
2	DOE/SC/Oak Ridge National Laboratory United States	560,640	17,590.0	27,112.5	8,209
3	DOE/NNSA/LLNL United States	1,572,864	17,173.2	20,132.7	7,890
4	RIKEN Advanced Institute for Computational Science (AICS) Japan	705,024	10,510.0	11,280.4	12,660
5	DOE/SC/Argonne National Laboratory United States	786,432	8,586.6	10,066.3	3,945
6	Swiss National Supercomputing Centre (CSCS) Switzerland	115,984	6,271.0	7,788.9	2,325
7	King Abdullah University of Science and Technology Saudi Arabia	196,608	5,537.0	7,235.2	2,834
8	Texas Advanced Computing Center/Univ. of Texas United States	462,462	5,168.1	8,520.1	4,510

Computing Trends and Projections

Moore's Law: The Ugly

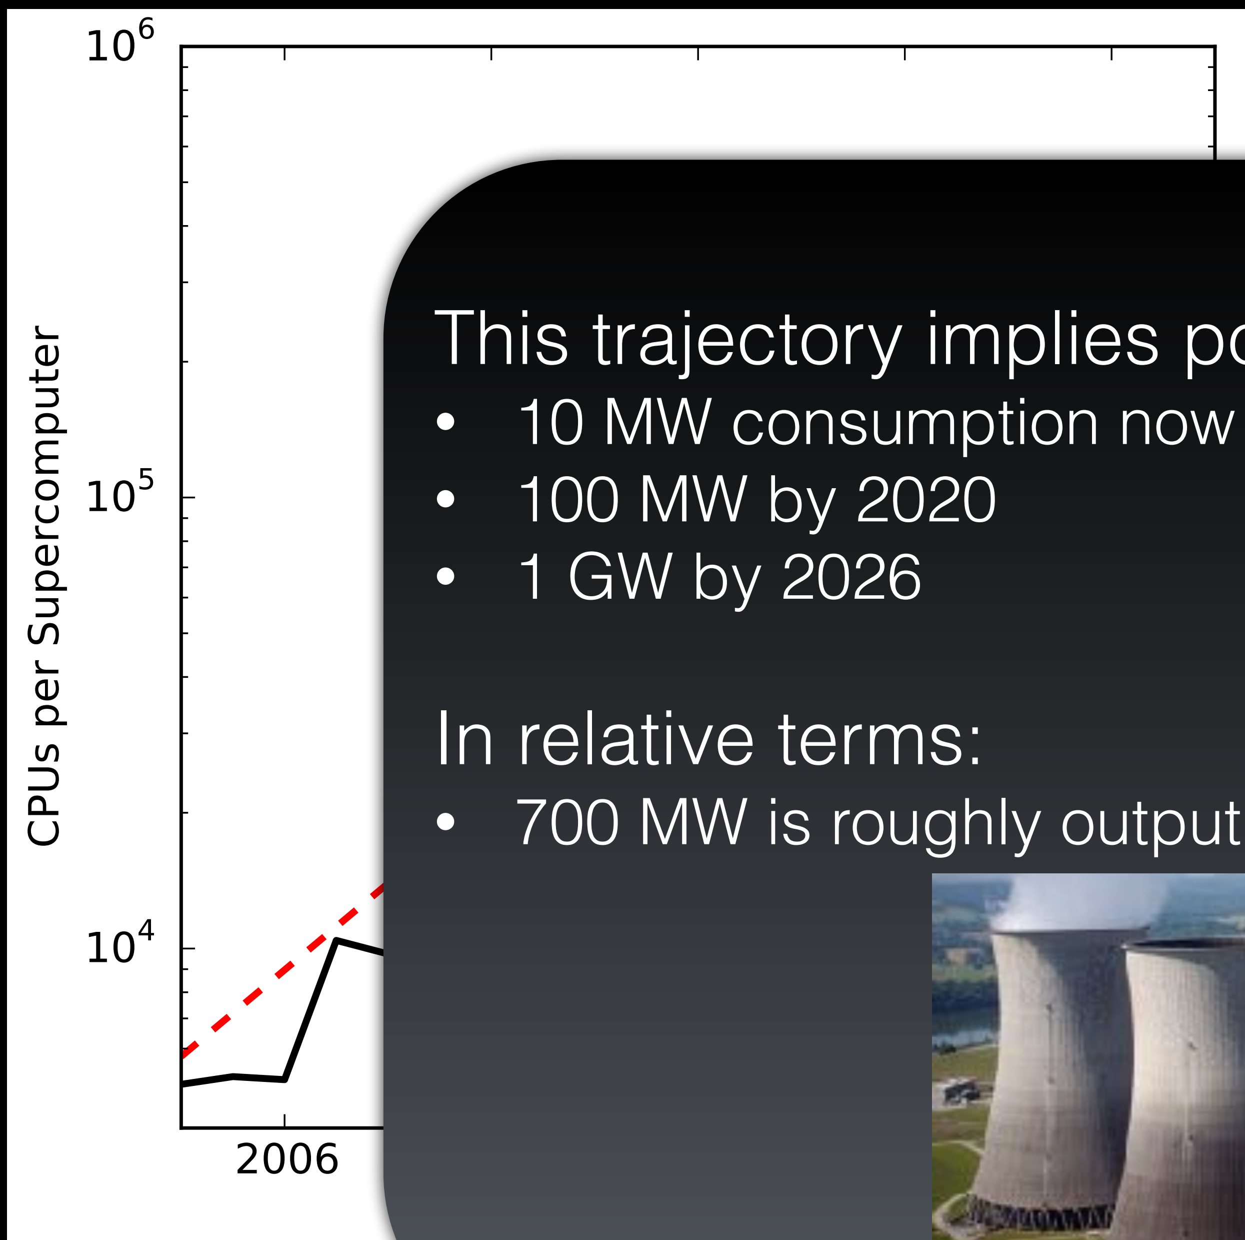

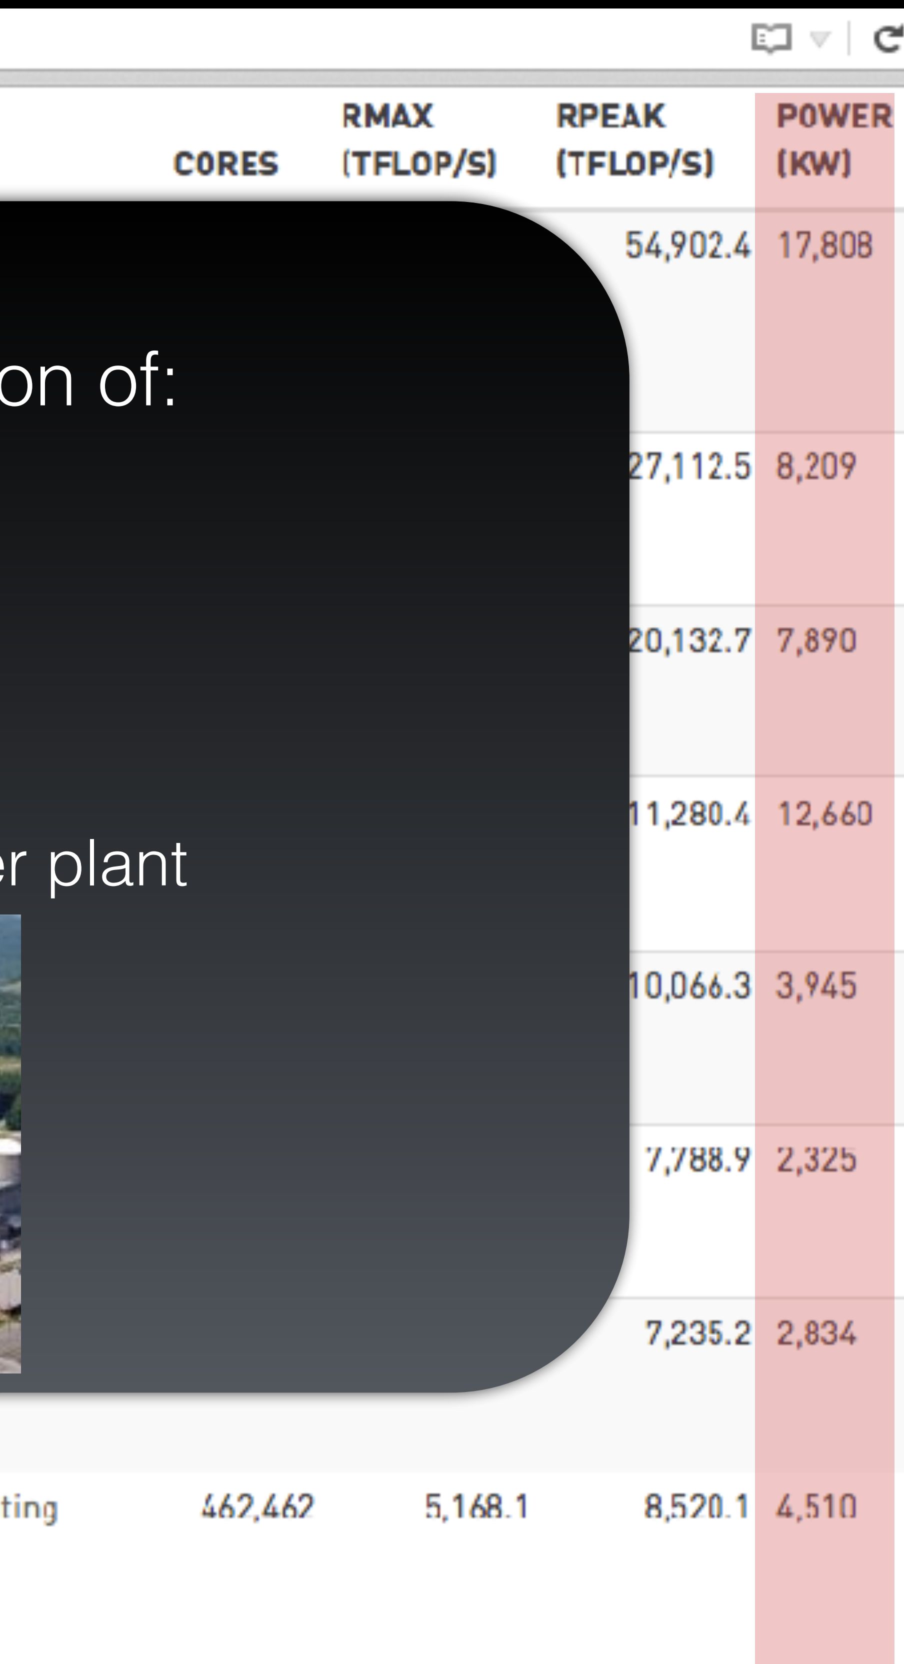

TOP 500
The List.

This trajectory implies power consumption of:

- 10 MW consumption now
- 100 MW by 2020
- 1 GW by 2026

In relative terms:

- 700 MW is roughly output of a nuclear power plant

Computing Trends and Projections

Moore's Law: A More Realistic Future

Better metrics of computing efficiency?

How to handle complexity of heterogeneous architectures (e.g. CPUs + GPUs)?

Complexity of hybrid parallelism (e.g. MPI + threading)?

The Green500 List

Listed below are the June 2015 Thcomputers ranked from 1 to 10.

Green500 Rank	MFLOPS/W	Computer*	Total Power (kW)
1	7,031.58	Shoubu - ExaScaler-1.4 80Brick, Xeon E5-2618Lv3 8C 2.3GHz, Infiniband FDR, PEZY-SC	50.32
2	6,842.31	Suiren Blue - ExaScaler-1.4 16Brick, Xeon E5-2618Lv3 8C 2.3GHz, Infiniband, PEZY-SC	28.25
3	6,217.04	Suiren - ExaScaler 32U256SC Cluster, Intel Xeon E5-2660v2 10C 2.2GHz, Infiniband FDR, PEZY-SC	32.59
4	5,271.81	ASUS ESC4000 FDR/G2S, Intel Xeon E5-2690v2 10C 3GHz, Infiniband FDR, AMD FirePro S9150	57.15
5	4,257.88	TSUBAME-KFC - LX 1U-4GPU/104Re-1G Cluster, Intel Xeon E5-2620v2 6C 2.100GHz, Infiniband FDR, NVIDIA K20x	39.83
6	4,112.11	XStream - Cray CS-Storm, Intel Xeon E5-2680v2 10C 2.8GHz, Infiniband FDR, Nvidia K80	190.00
7	3,962.73	Storm1 - Cray CS-Storm, Intel Xeon E5-2660v2 10C 2.2GHz, Infiniband FDR, Nvidia K40m	44.54
8	3,631.70	Wilkes - Dell T620 Cluster, Intel Xeon E5-2630v2 6C 2.600GHz, Infiniband FDR, NVIDIA K20	52.62
9	3,614.71	Taurus GPUs - Bull bullx R400, Xeon E5-2680v3 12C 2.5GHz, Infiniband FDR, Nvidia K80	58.01
10	3,543.32	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband, NVIDIA K20x	54.60

Summary

- Galaxy formation is complicated!
- Need to account for gravity, hydrodynamics, and small-scale physics within galaxies
- Can use simulations to “test out” various physics models, see what processes drive galaxy formation
- HPC advances are needed in the future to enable larger simulations, easier parallelization