‘Large’ Cosmological Simulations: Future Prospects (2020’s)

Previous activity: needs update, White Paper? Connect with Computing Frontier?

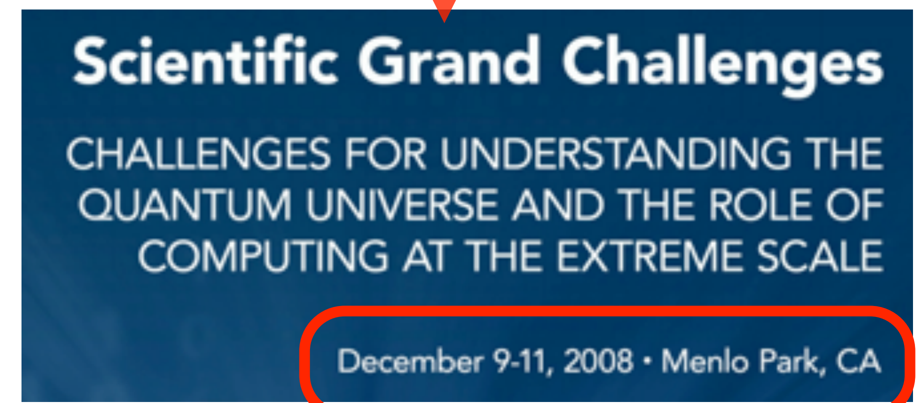

Salman Habib
HEP and MCS Divisions,
Argonne National Laboratory

Andy Connolly, Alex Szalay (CF)
George Fuller, Tony Mezzacappa (A&C)
Snowmass Computing Frontier

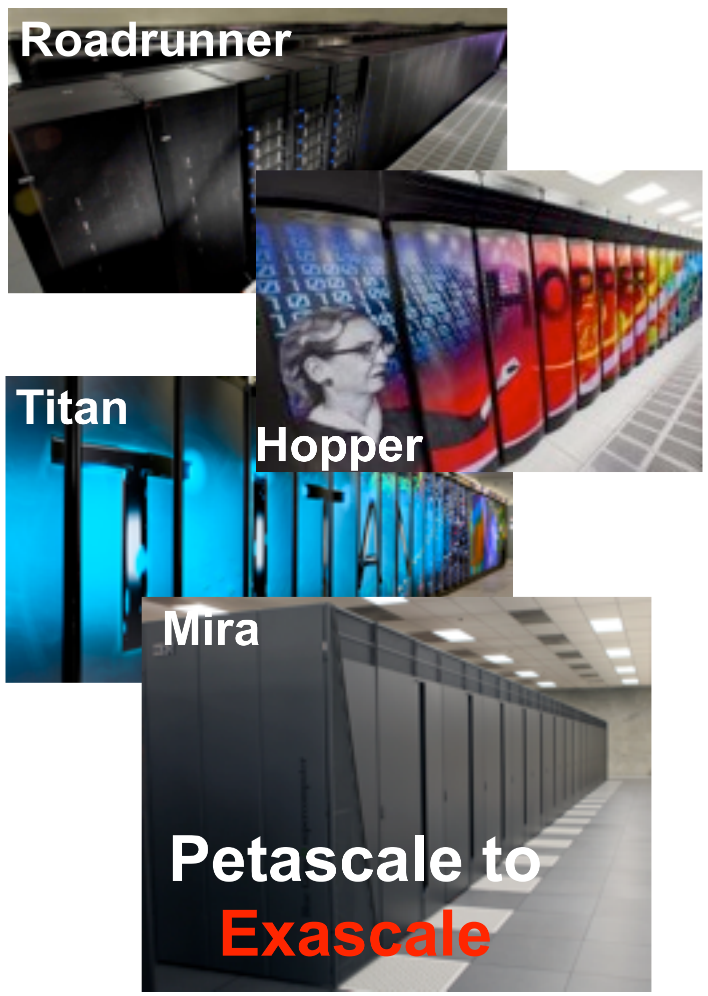

3 Gpc/h

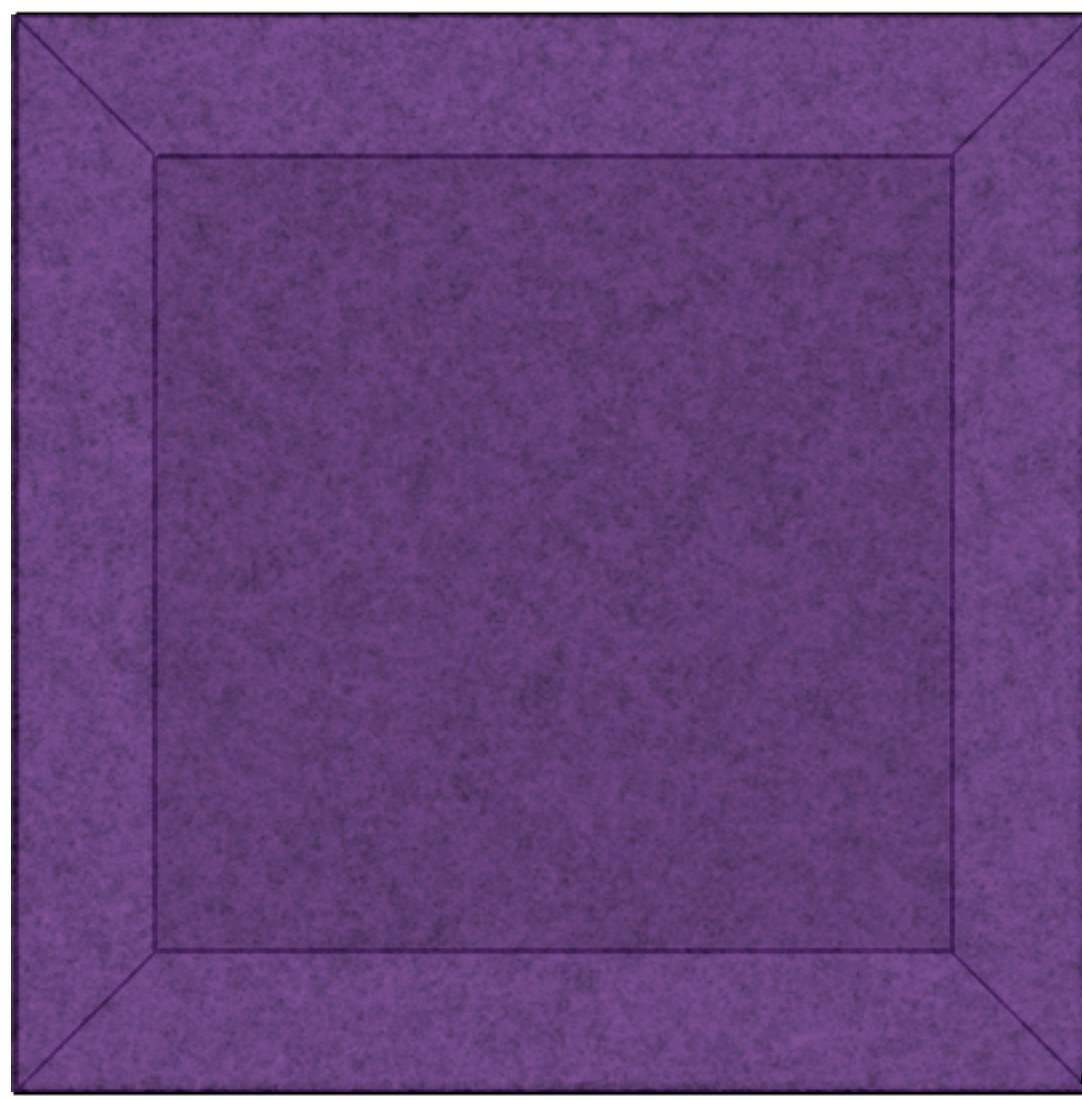

One out of 262,144 ranks; note force resolution is ~0.005 Mpc

47 Mpc/h

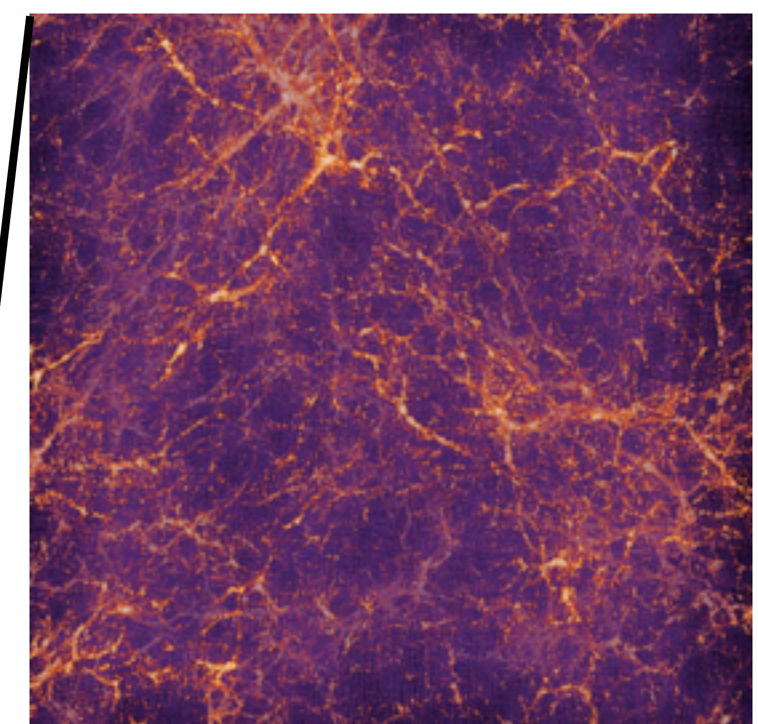

1.1 trillion particle HACC run at z=3

Cosmic Frontier Workshop
SLAC, March 6, 2013

Structure Formation: Key Issues

- **Inputs**
  - Basic theory mostly under control, some work always needed, but no show-stoppers (apparently)
  - Gastrophysics issues very serious, predictive modeling hopeless?
  - Development of simulation-based error-controlled phenomenology needed (multiple probe cross-cuts)
- **Drivers**
  - Precision cosmology from surveys (understand/control systematics)
  - 'Smoking gun' systems (indirect DM searches, extreme statistics)
- **Enabling 'technology'**
  - Next-generation HPC systems ('exascale')
  - Data-intensive computing
  - Statistical inverse problems
  - Analysis software frameworks

Focus of this talk

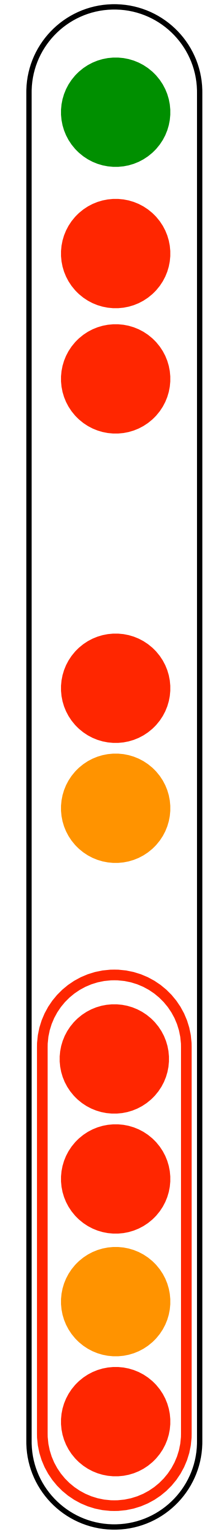

Role of Computational Cosmology

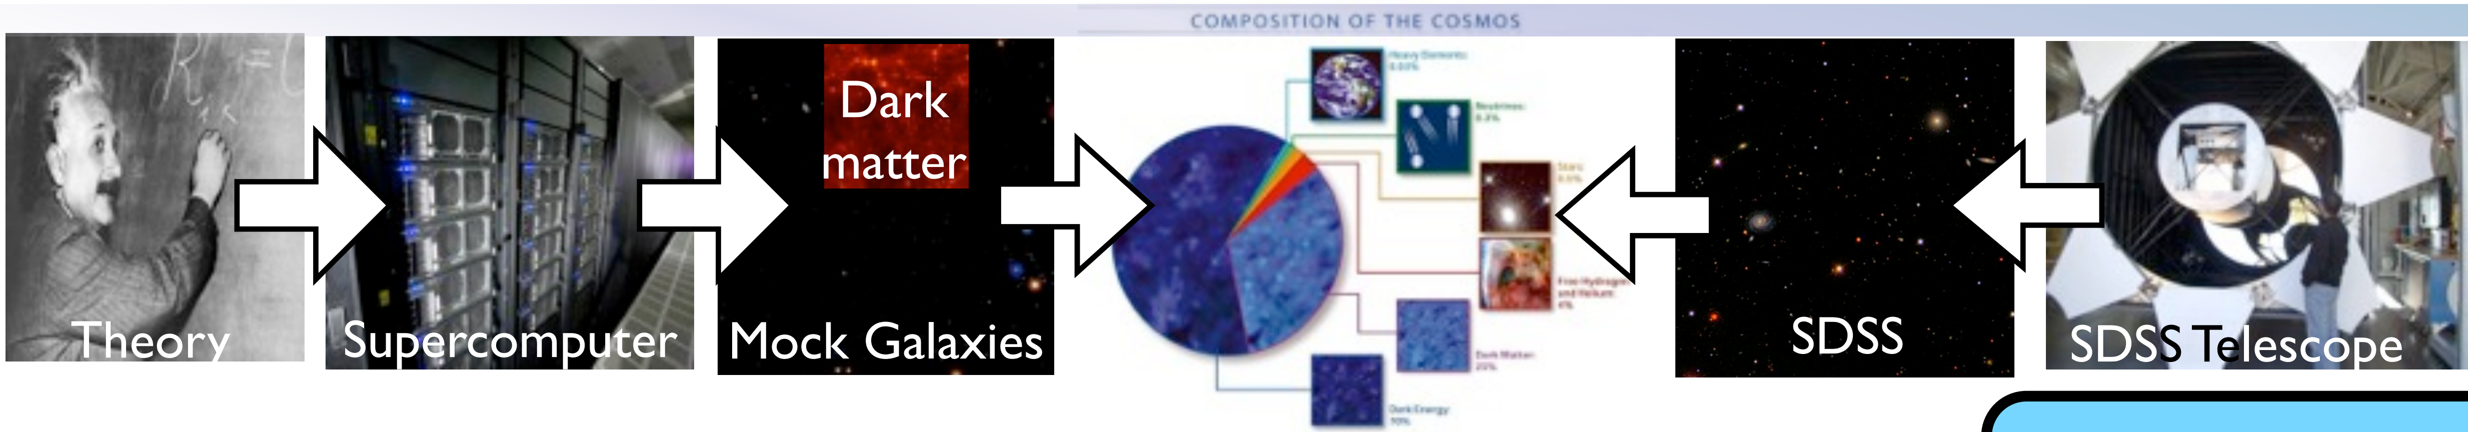

- **Three Roles of Cosmological Simulations**
  - Basic ‘theory’ of (nonlinear) cosmological probes, incl. errors
  - Production of high-fidelity ‘mock skys’ for end-to-end tests of the observation/analysis chain
  - Essential component of analysis toolkits
- **Simulation and Analysis Challenges**
  - Large dynamic range simulations; control of subgrid modeling and feedback mechanisms (combine with observations)
  - Design and implementation of complex analyses on large datasets; new fast (approximate) algorithms
  - Solution of large statistical inverse problems of scientific inference (many parameters, ~10-1000’s?)

  Use of Advanced Statistical methods/Machine Learning

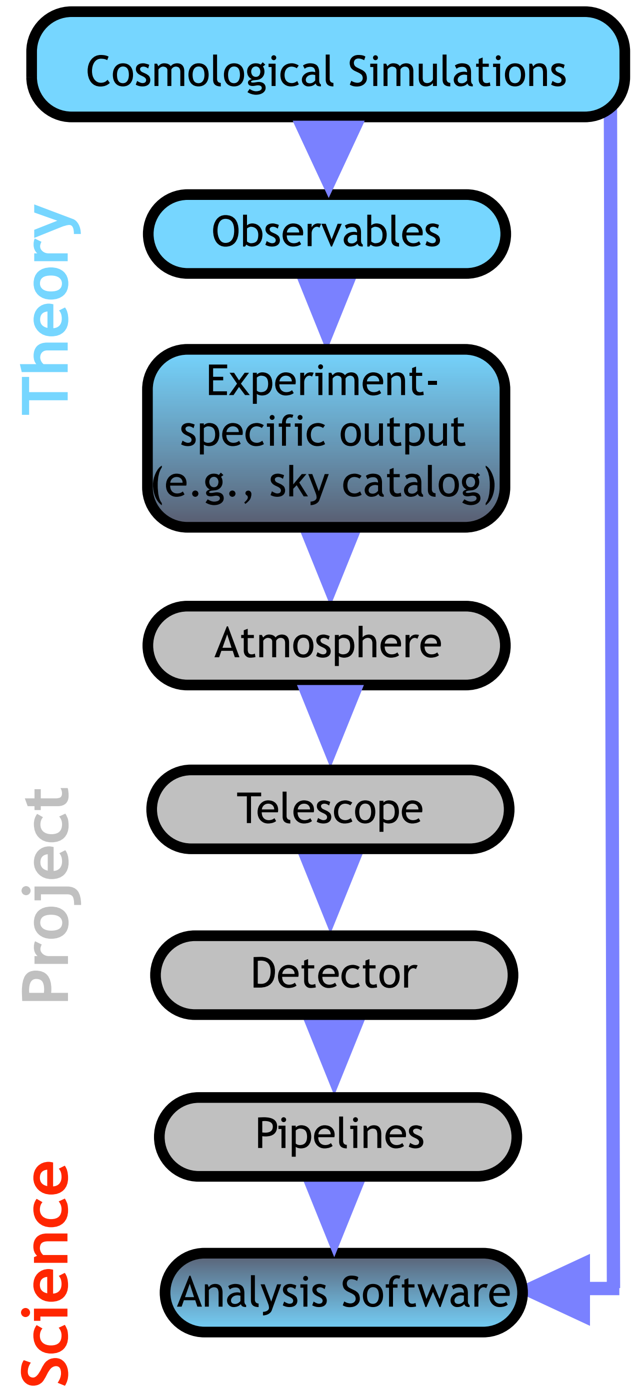

Precision Cosmology: Big Data Meets Supercomputing

Supercomputer Simulation Campaigns

Mapping the Sky with Survey Instruments

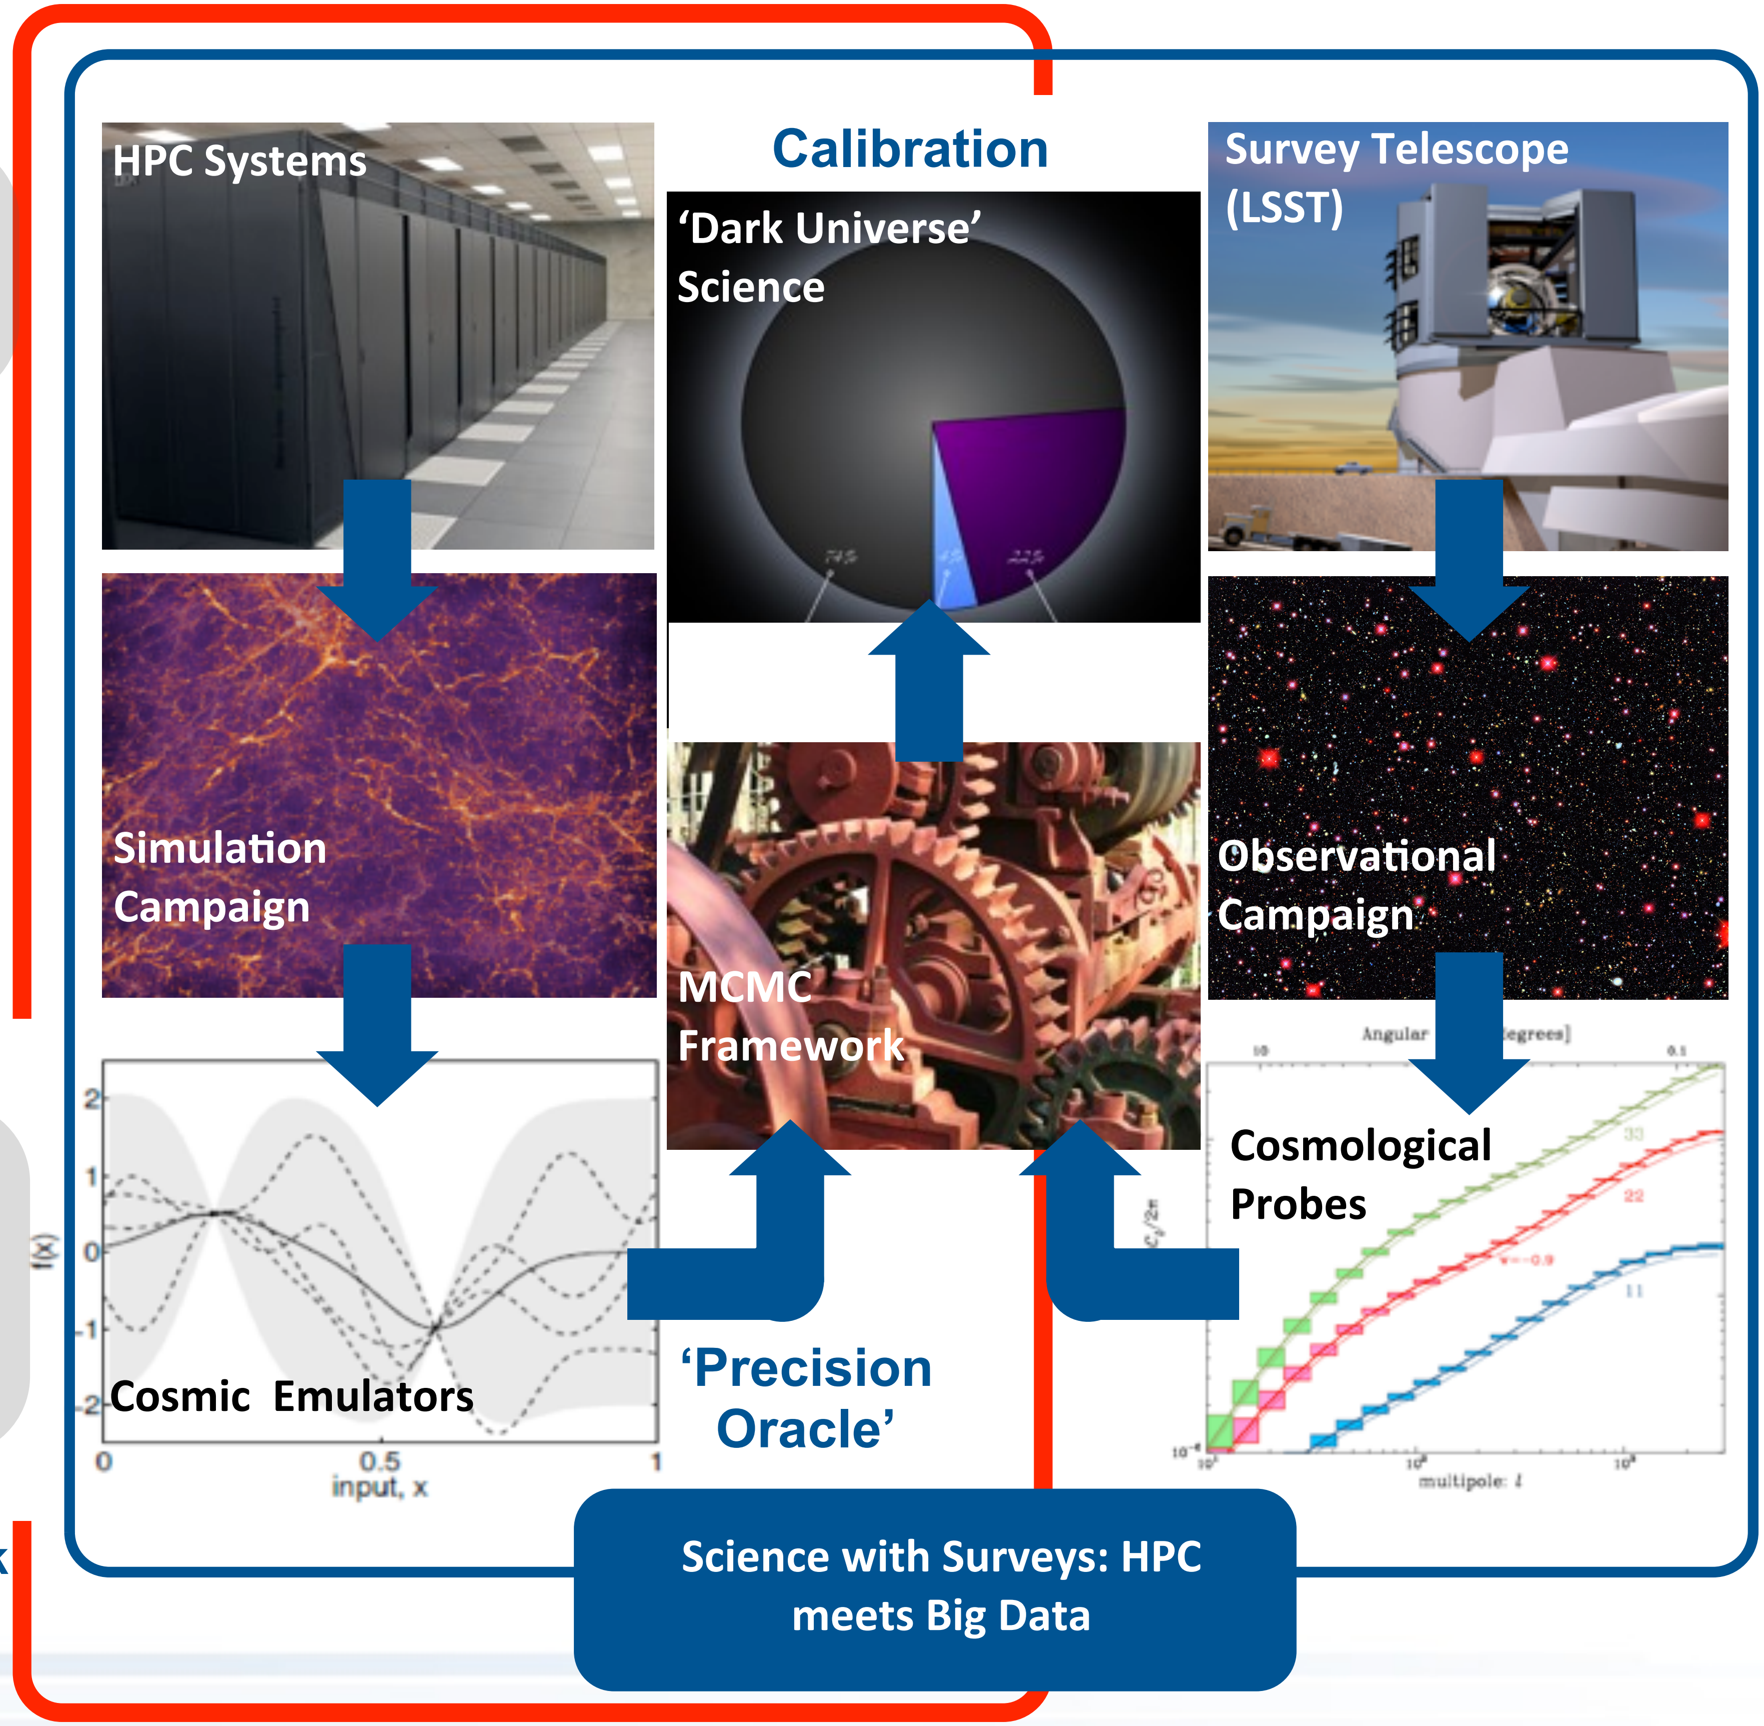

Simulations + CCF

Analysis pipelines

Emulator based on Gaussian Process Interpolation in High-Dimensional Spaces

CCF= Cosmic Calibration Framework

Simulation and Analysis Requirements

- **Resolution:**
  - To the extent that very small scales will not be treated in large simulations, the physics will need to be resolved to ~kpc scales (or a little less), global dynamic range is likely to be about 10 million to 1 (~exaflop performance and ~100 PB 'RAM', many exa/zettabytes of storage)
  - Mass resolutions in large simulations can reach as low as 10^6-10^7 solar masses (but time evolution will be a mess), simulations in the ~100 trillion -1 quadrillion particle class, (equivalent of) multi-billion core-hours per simulation (both N-body and hydro)
  - Accuracy requirements have not been treated seriously enough so far, need to be addressed
- **Simulation modes:**
  - Smaller number of 'extreme' runs for specific goals
  - Large number of 'workhorse' runs for simulation campaigns (systematics studies, inference engines, covariance estimation, --)
  - High-resolution, physics-rich simulations to build sub-grid models (will this require special hardware?)
- **Analysis:**
  - Multiple probes and cross-correlations, need to build integrated analysis framework that has HPC and data-intensive computing components, how to best incorporate legacy components?
  - Next-generation tools for statistical inverse problems and optimization
  - Critical issue: development of human resources (find/support people with multiple talents)

Where is Conventional Supercomputing Headed?

- **Not faster, but 'bigger':**
  - Much more parallelism but at sharply reduced memory/compute core; complex/heterogeneous nodes; nontrivial memory/network hierarchy; serious I/O issues (e.g., local checkpoints)
  - End of 'weak scaling' as we know it; major penalty for complex data structures
  - Time evolution codes will suffer as more physics is added, new exploits of concurrency are needed

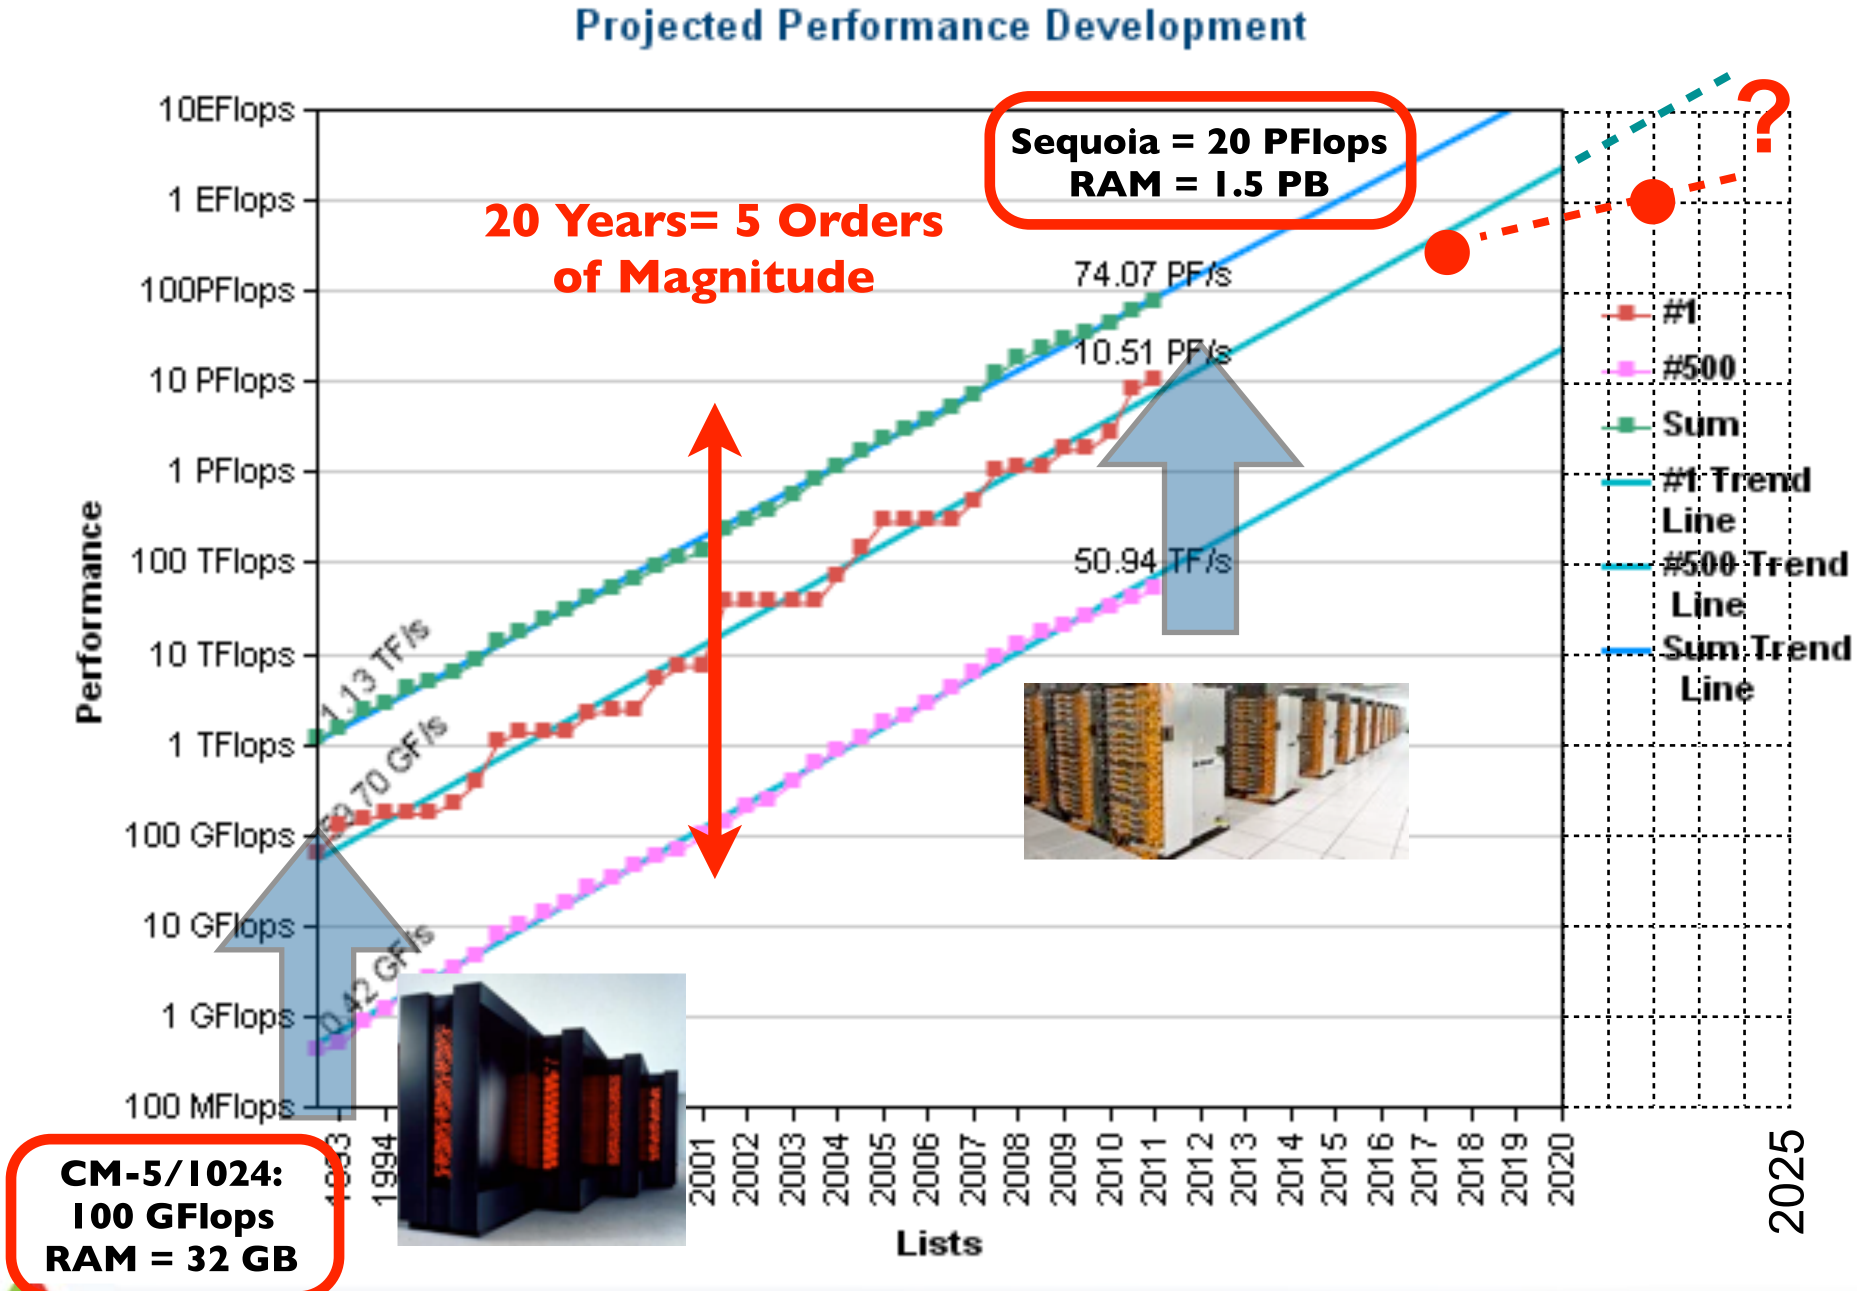

- **Other important things**
  - Programming model space is very murky
  - A "data analysis cloud" with dynamically allocatable resources will complement the supercomputer
  - Ability to run very large simulation campaigns on supercomputers
  - Power-awareness, fault-tolerance, whose responsibility?
  - To design codes for the future **must know soon** what codes have to do! It will not just happen --

A Flavor of the Hardware Complexity Challenge

- **HPC Myths**
  - The magic compiler
  - The magic programming model/language (DSL)
  - Special-purpose hardware
  - Co-Design?
- **Dealing with (current and future) HPC Reality**
  - Follow the architecture
  - Know the boundary conditions
  - There is no such thing as a 'code port'
  - Think out of the box
  - Put the best team together
  - Coordinate efforts

BQC:
- 16 cores
- 205 GFlops, 16 GB
- 32 MB L2, crossbar at 400 GB/s (memory connection is 40 GB/s)
- 5-D torus at 40 GB/s

Xeon Phi:
- 60 cores
- 1 TFlops, 8 GB
- 32 MB L2, ring at 300 GB/s (connects to cores and memory)
- 8 GB/s to host CPU

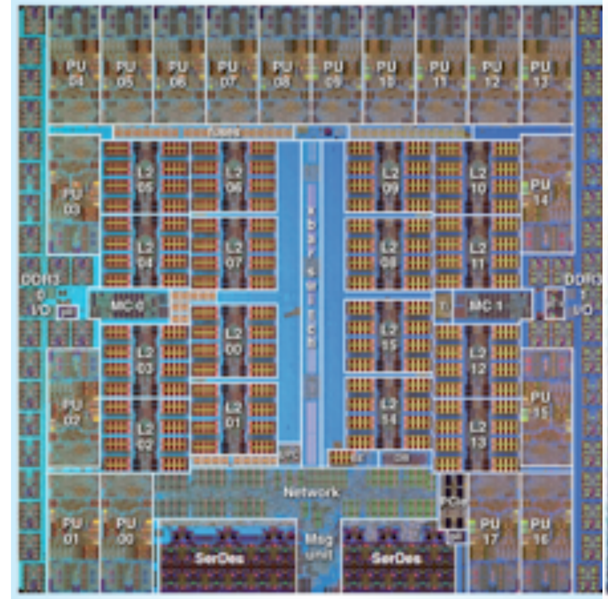

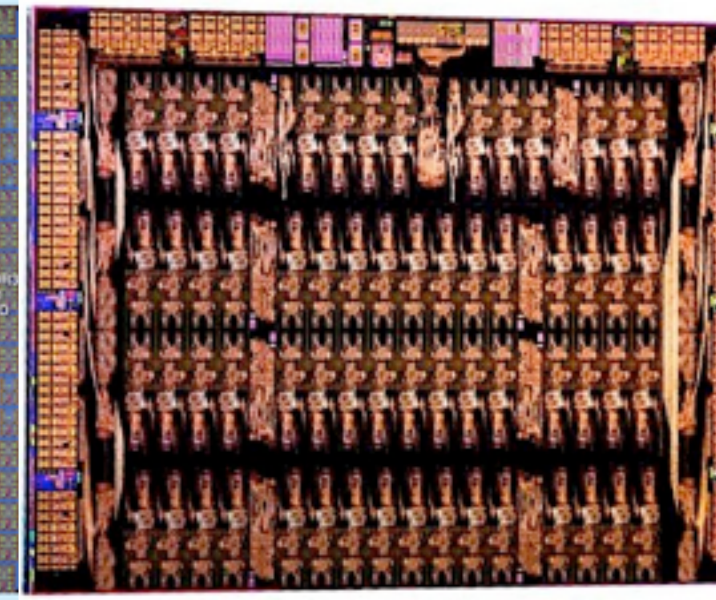

Average performance speed-up on ~10 applications codes on Titan is ~2 (ranging from 1.few to 7), but of Titan's 27 PFlops, only 2.5 PFlops are in the CPU! What can we learn from this?

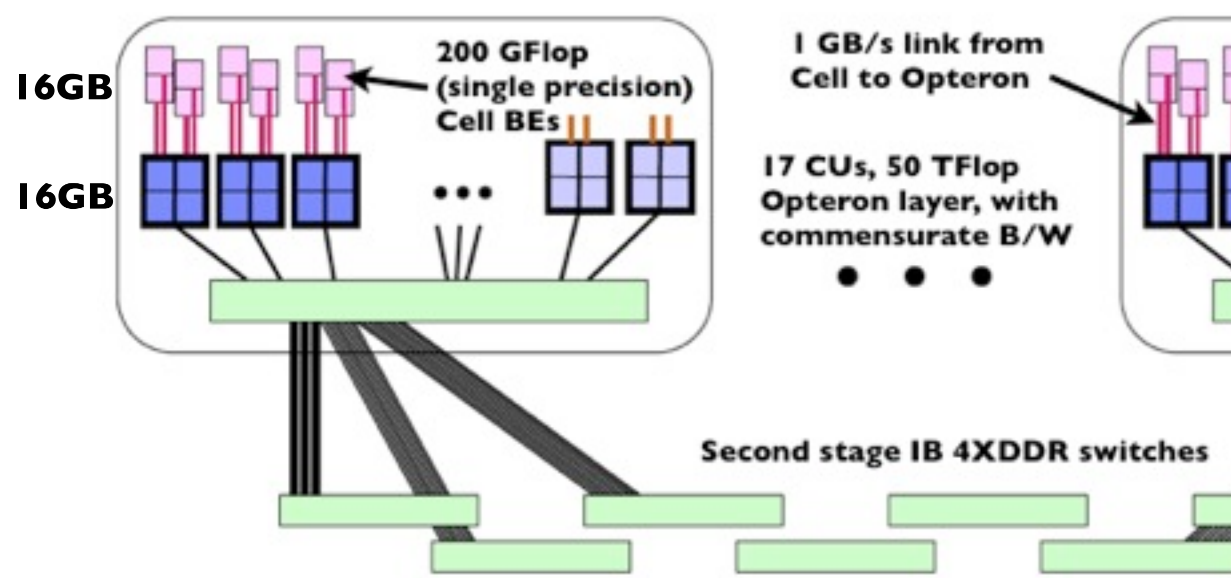

Roadrunner

Meeting the Challenge: HACC on the BG/Q

- New Cosmological N-Body Framework
  - Designed for extreme performance AND portability, including heterogeneous systems
  - Supports multiple programming models
  - Memory efficient
  - In situ analysis framework
  - Production science code

13.94 PFlops, 69.2% peak, 90% parallel efficiency on 1,572,864 cores/MPI ranks, 6.3M-way concurrency

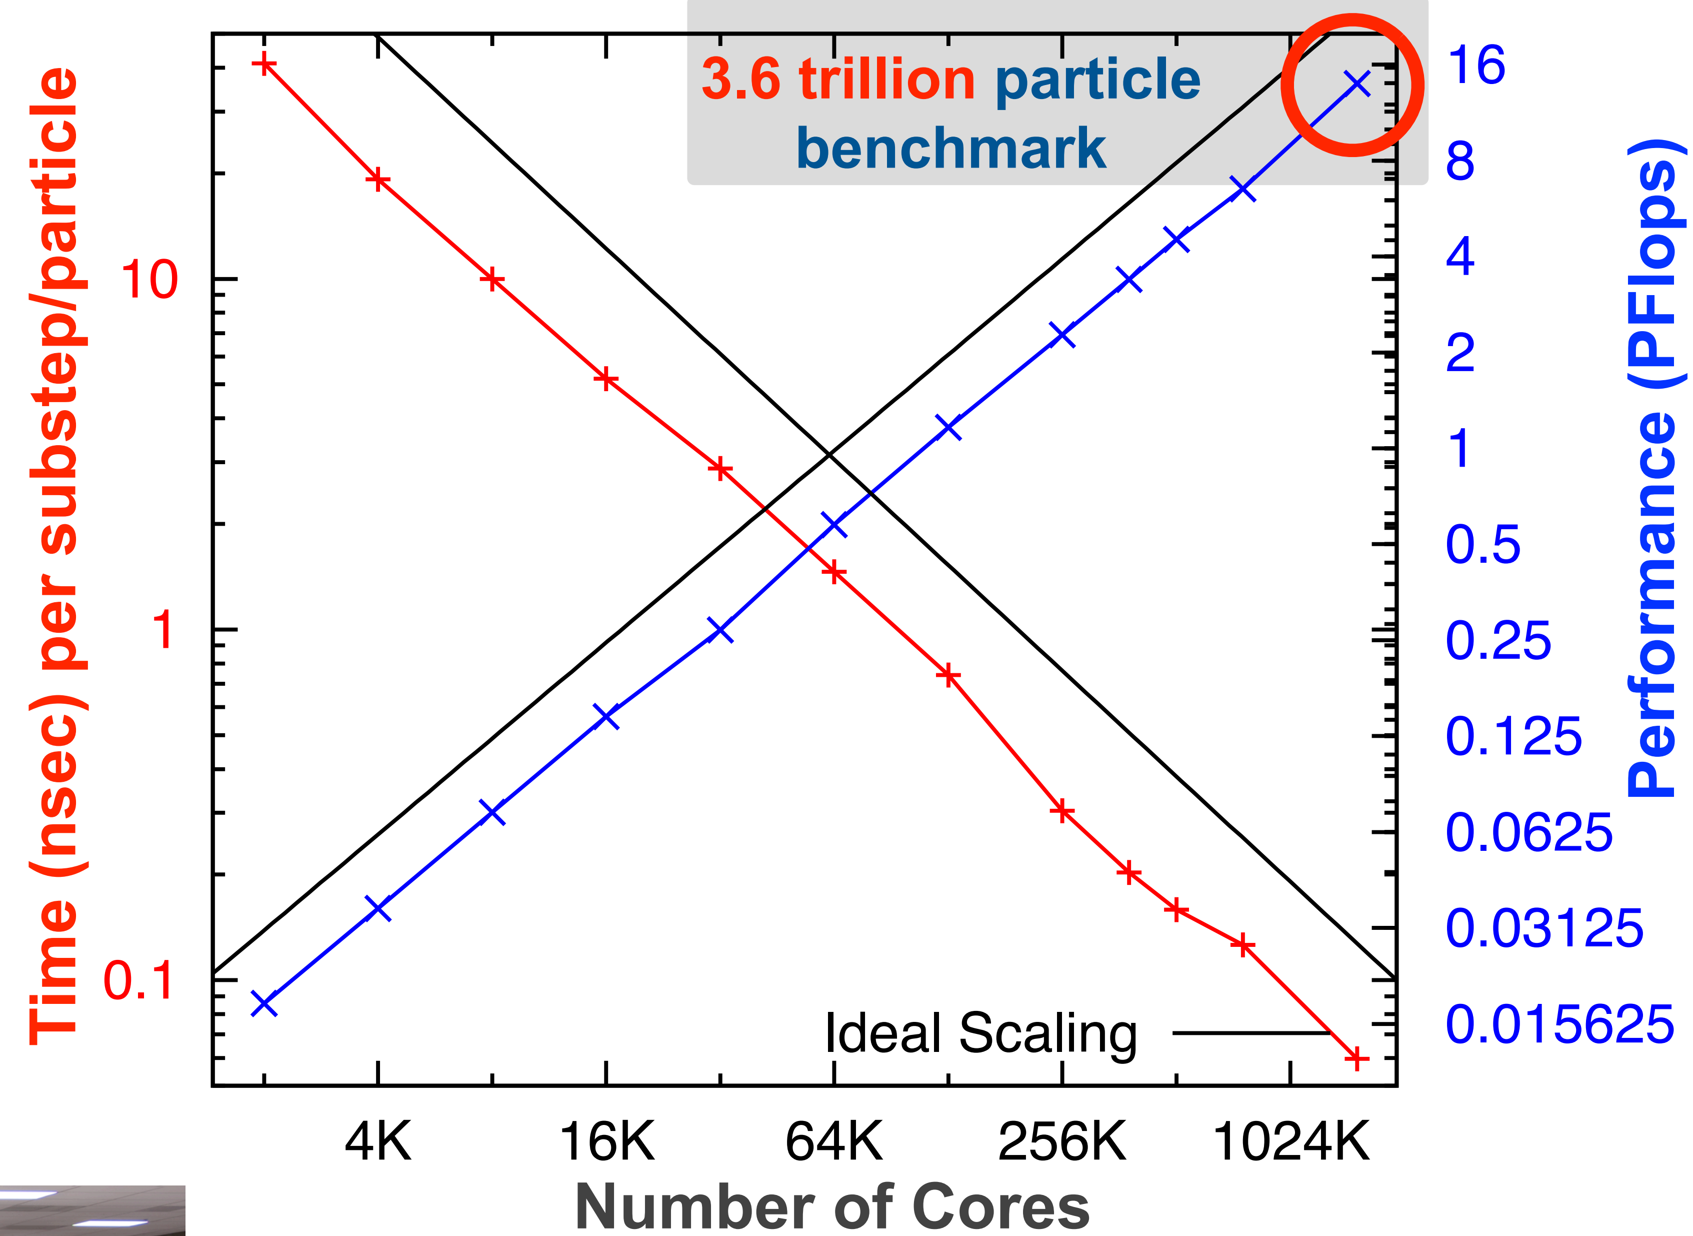

HACC weak scaling on the IBM BG/Q (MPI/OpenMP)

Fast In Situ Analysis

- **Data Reduction:** A trillion particle simulation with 100 analysis steps has a storage requirement of ~4 PB -- in situ analysis reduces it to ~200 TB
- **I/O Chokepoints:** Large data analyses difficult because I/O time > analysis time, plus scheduling overhead
- **Fast Algorithms:** Analysis time is only a fraction of a full simulation timestep
- **Ease of Workflow:** Large analyses difficult to manage in post-processing

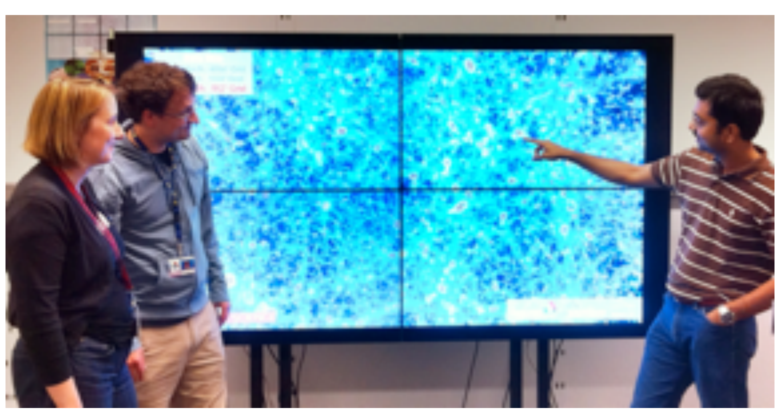

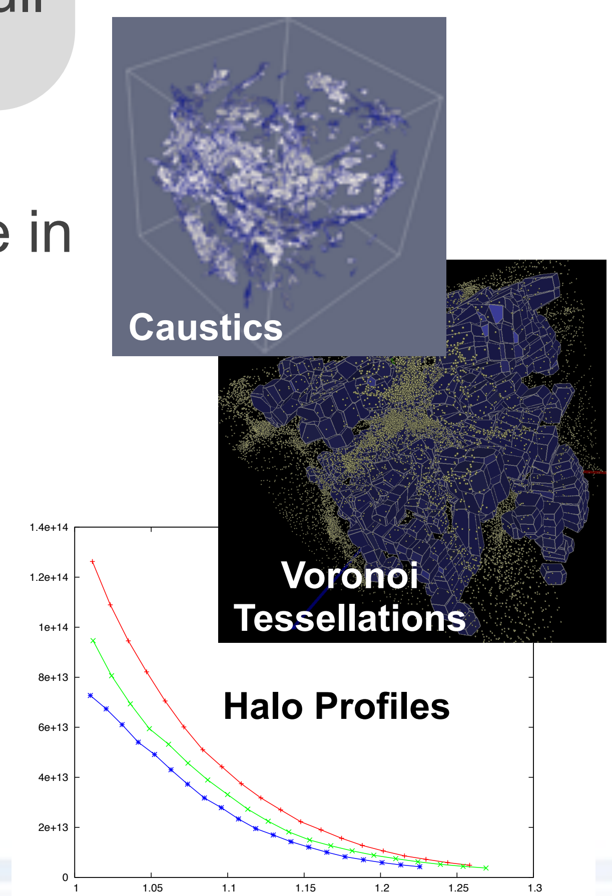

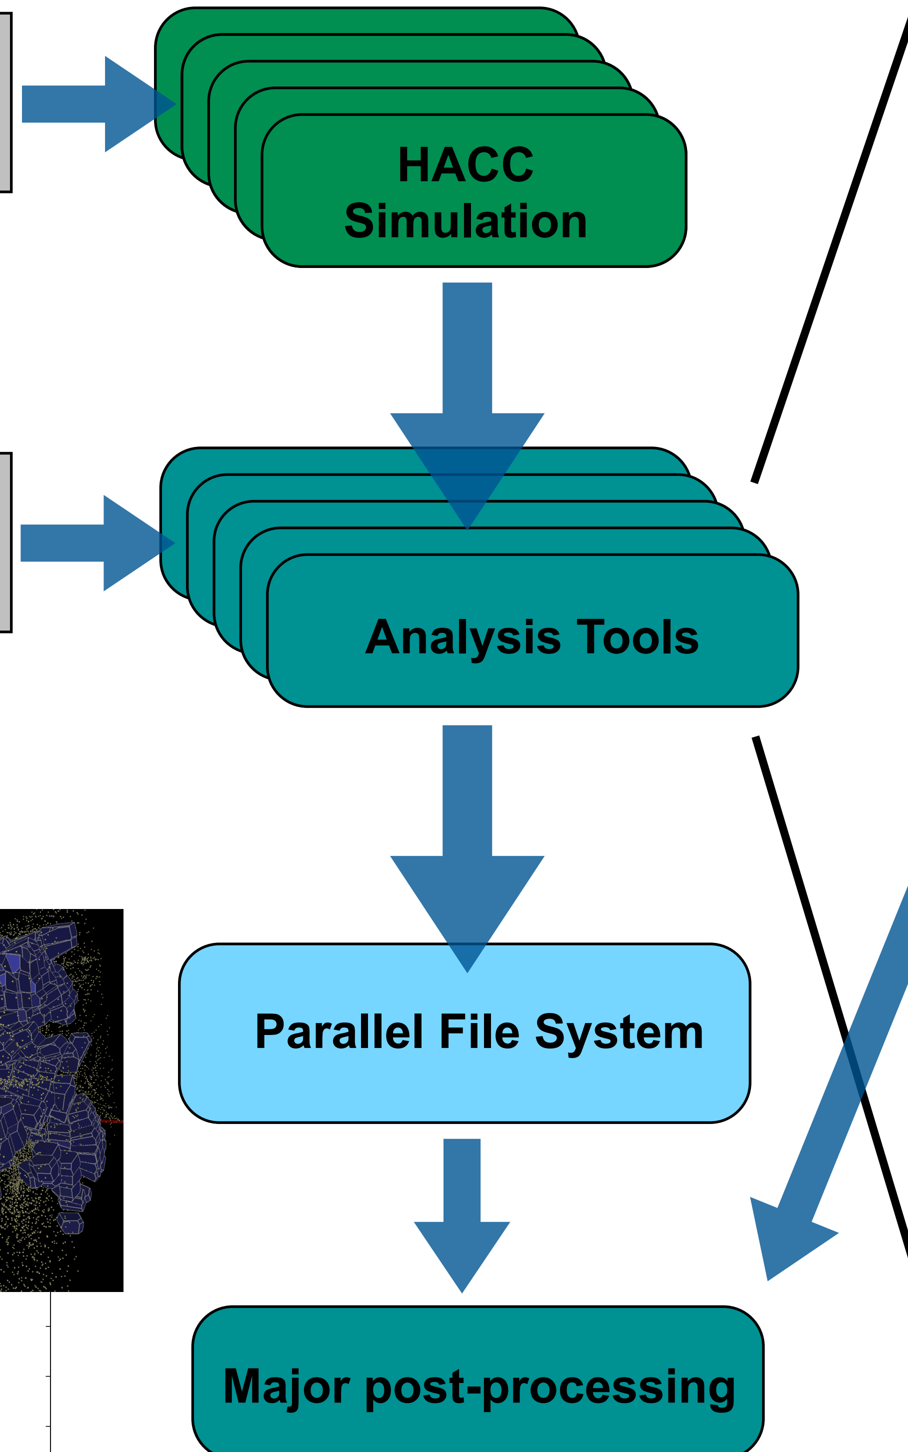

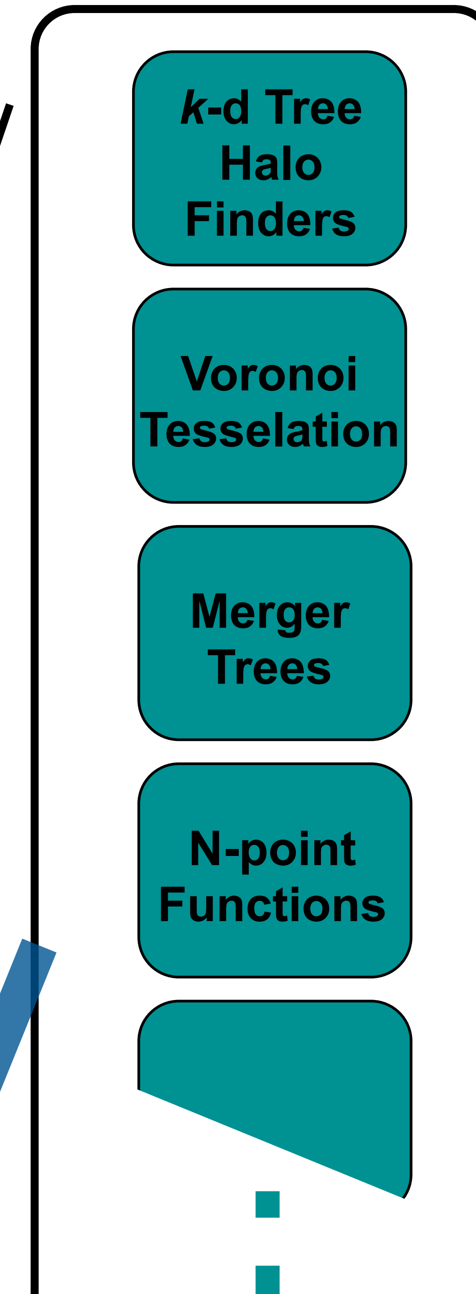

Predictions for solving cosmic inverse problems

Cosmic Calibration: Solving the Inverse Problem

- **Challenge:** To extract cosmological constraints from observations in the nonlinear regime, need to run Markov Chain Monte Carlo; input: 10,000 - 100,000 different models
- **Brute Force:** Hopeless (‘curse of dimensionality’ -- n^d)
- **Current Strategy:** Fitting functions, e.g. for P(k), accurate at 10% level, not good enough!
- **CS/ML Solution:** Precision emulators and discrepancy functions (systematic uncertainties)

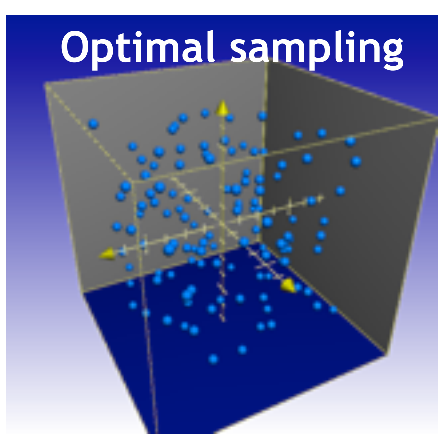

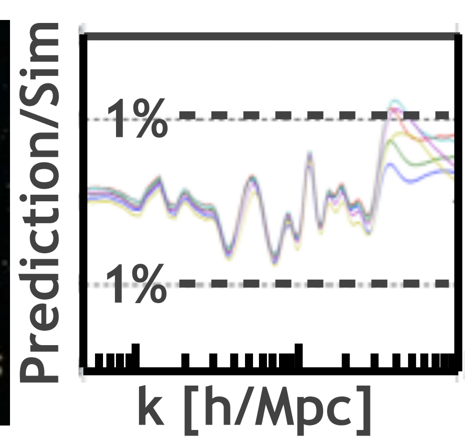

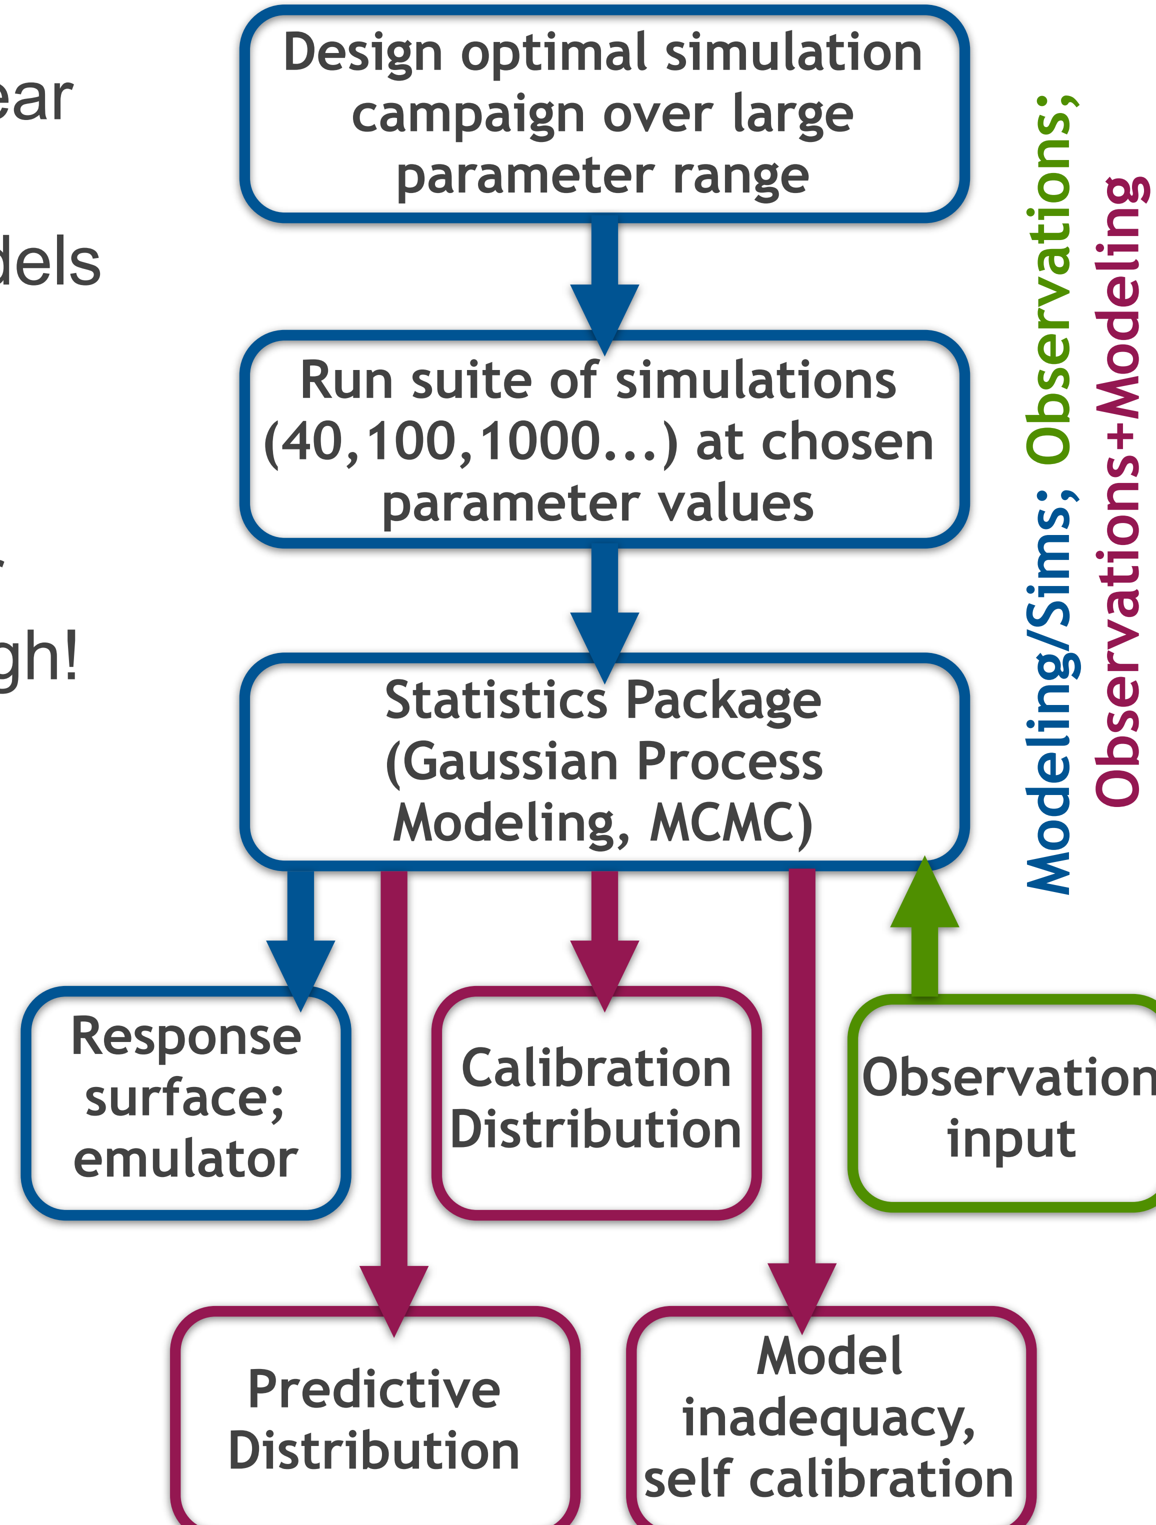

Summary: Computational Cosmology Ecosystem

- **Simulation Campaigns:**
  - Hundreds to thousands of large-scale simulations needed to build emulators and understand error covariances
  - Large numbers of parameters, combination of multiple simulation strategies
  - UQ issues: discrepancy functions, new Monte Carlo implementations, covariance across probes
- **Data/Analytics Issues:**
  - Development of high-performance *in situ* and post-processing viz/analysis frameworks
  - Joint analysis of data and simulations
- **Analysis Services (e.g., PDACS):**
  - Community interaction with large-scale simulation data archives/databases (w/ universities and Labs)
  - Repository of data and tools; access to computational resources and data transfer services, carry out complex analysis workflows

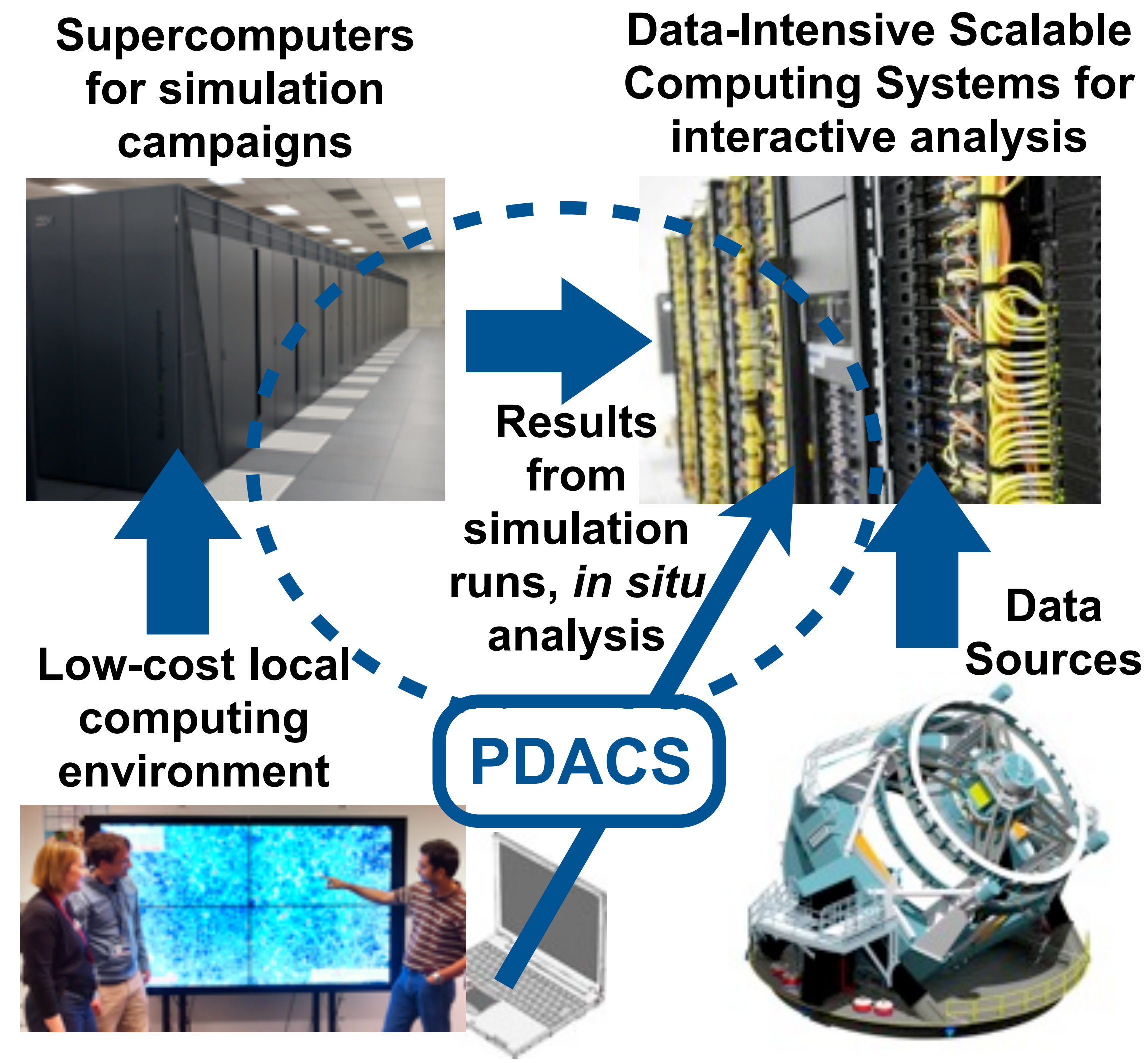

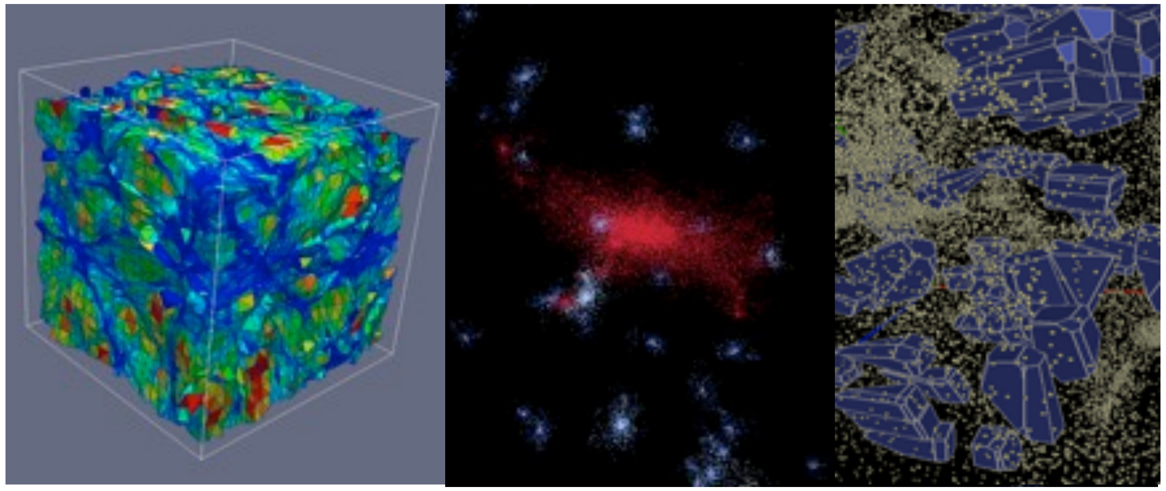

Snowmass Action Items?

- **Simulation/Theory Requirements:**
  - What are the science cases (need to collect/organize)? What sorts of (large) simulations are needed? What are the requirements imposed by the science? Is there a prioritized timeline? What are the things we may not have thought of?
  - What is required of the community for progress? How large and interactive is the research community? Who will help build/test the analysis framework(s)?
  - Resource estimates (people, underlying science, hardware, software, how do we do this?)
- **Simulation Methods Development:**
  - Are pure gravity codes almost 'done'? How to add a systematic subgrid component (needed or not)?
  - What should we do in 'hydro' space (cosmic frontier meets traditional astronomy/astrophysics -- e.g., clusters/groups, CMB, galaxy clustering, lensing, etc.)?
  - Build in inputs, including observational inputs, to achieve a more formal component for quantifying modeling errors (historically, errors have always been underestimated, sometimes seriously so)
- **Infrastructure**
  - Integrated repositories of data and tools; access to computational resources and data transfer services, carry out complex analysis workflows
  - Plan for exploits of future computational system/network ecology