High Performance Computing For B-Mode Data Analysis

Julian Borrill
Computational Cosmology Center, Berkeley Lab
& Space Sciences Laboratory, UC Berkeley

CMB Power Spectra

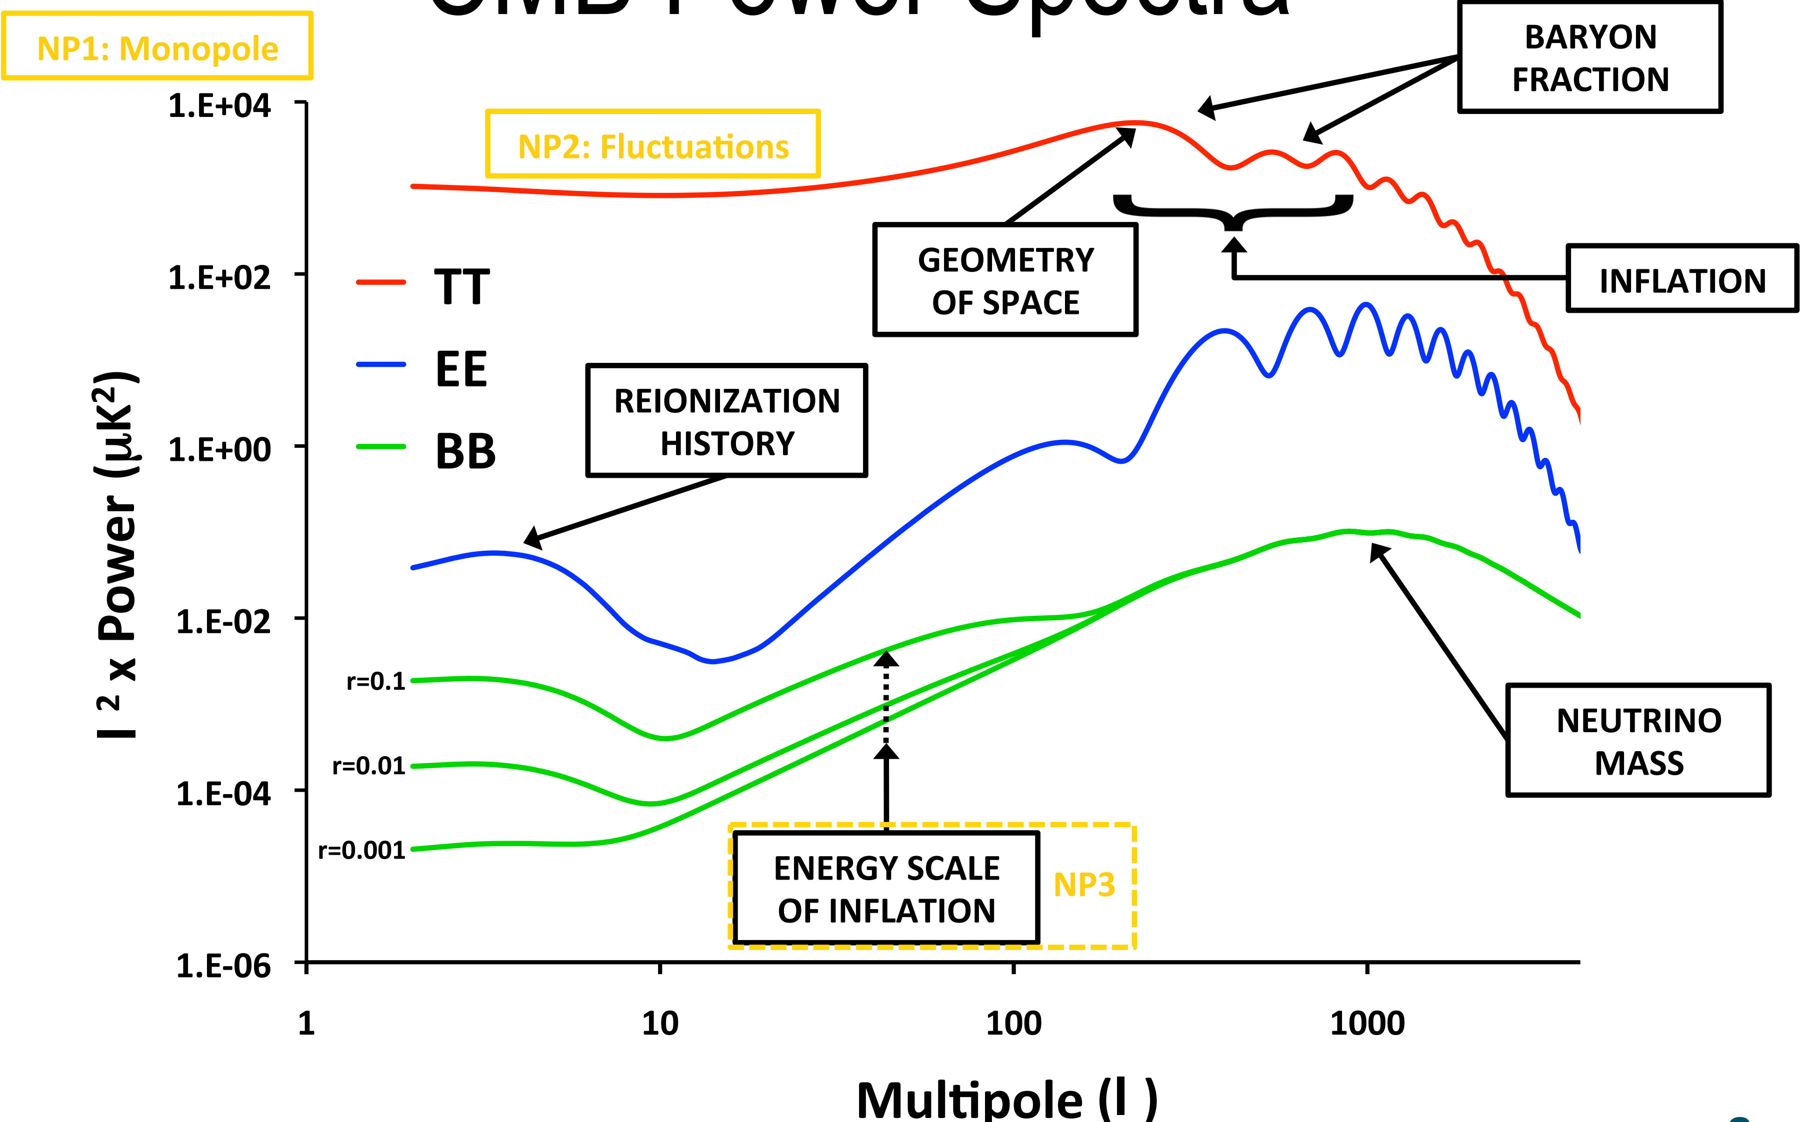

Another View …

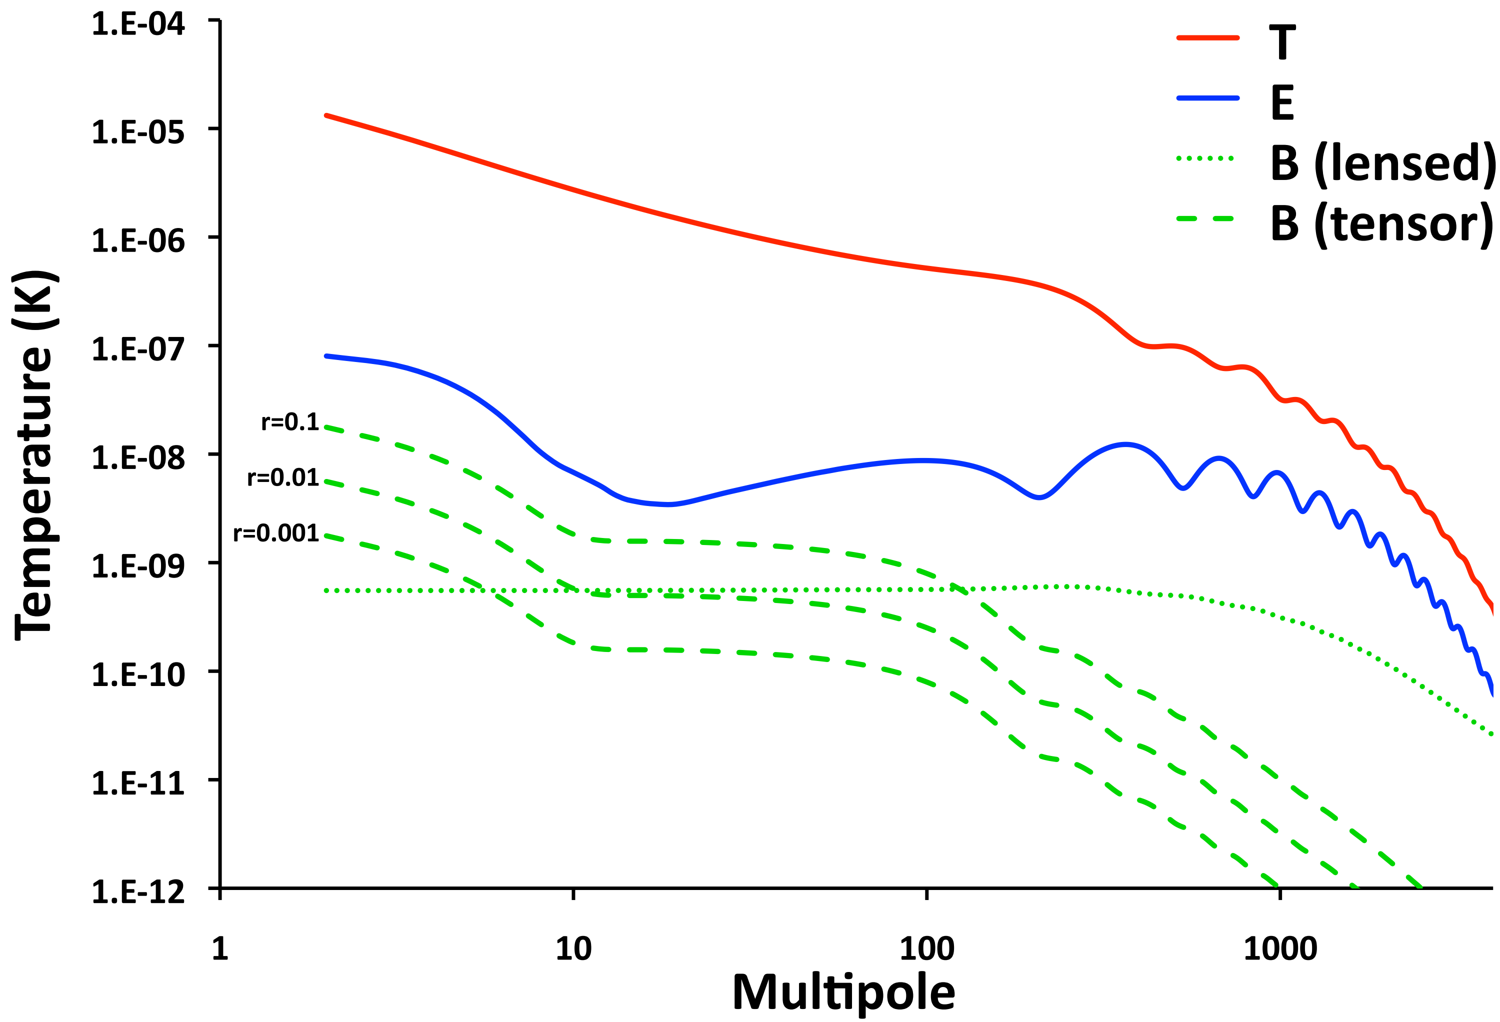

CMB Signals

COMPONENT	AMPLITUDE (K)	ERA
TT : Monopole	1	1968 (Penzias & Wilson)
TT : Anisotropy	10^{-5}	1990 (COBE)
TT : 1^{st} Peak	10^{-6}	2000 (BOOMERanG, MAXIMA)
EE : Reionization	10^{-7}	2005 (DASI)
BB : Lensing	10^{-9}	2015 (SPT, POLARBEAR)
BB : Gravity Waves	$< 10^{-9}$	**(AdvACTpol, BICEP3, CLASS, Simons Array, SPT-3G, CMB-S4; EBEX-10K, PIPER, SPIDER-2; LiteBIRD, PIXIE, COrE+)**

CMB Data Evolution

- Detecting fainter signals requires
  - Larger data sets to beat down noise
  - Tighter control of systematics
    - Environment
      - Foregrounds, Atmosphere
    - Instrument
      - Beams, Bandpasses, Noise, (HWP)
    - Analysis
      - Approximations, Biases, Uncertainties

CMB Data Growth

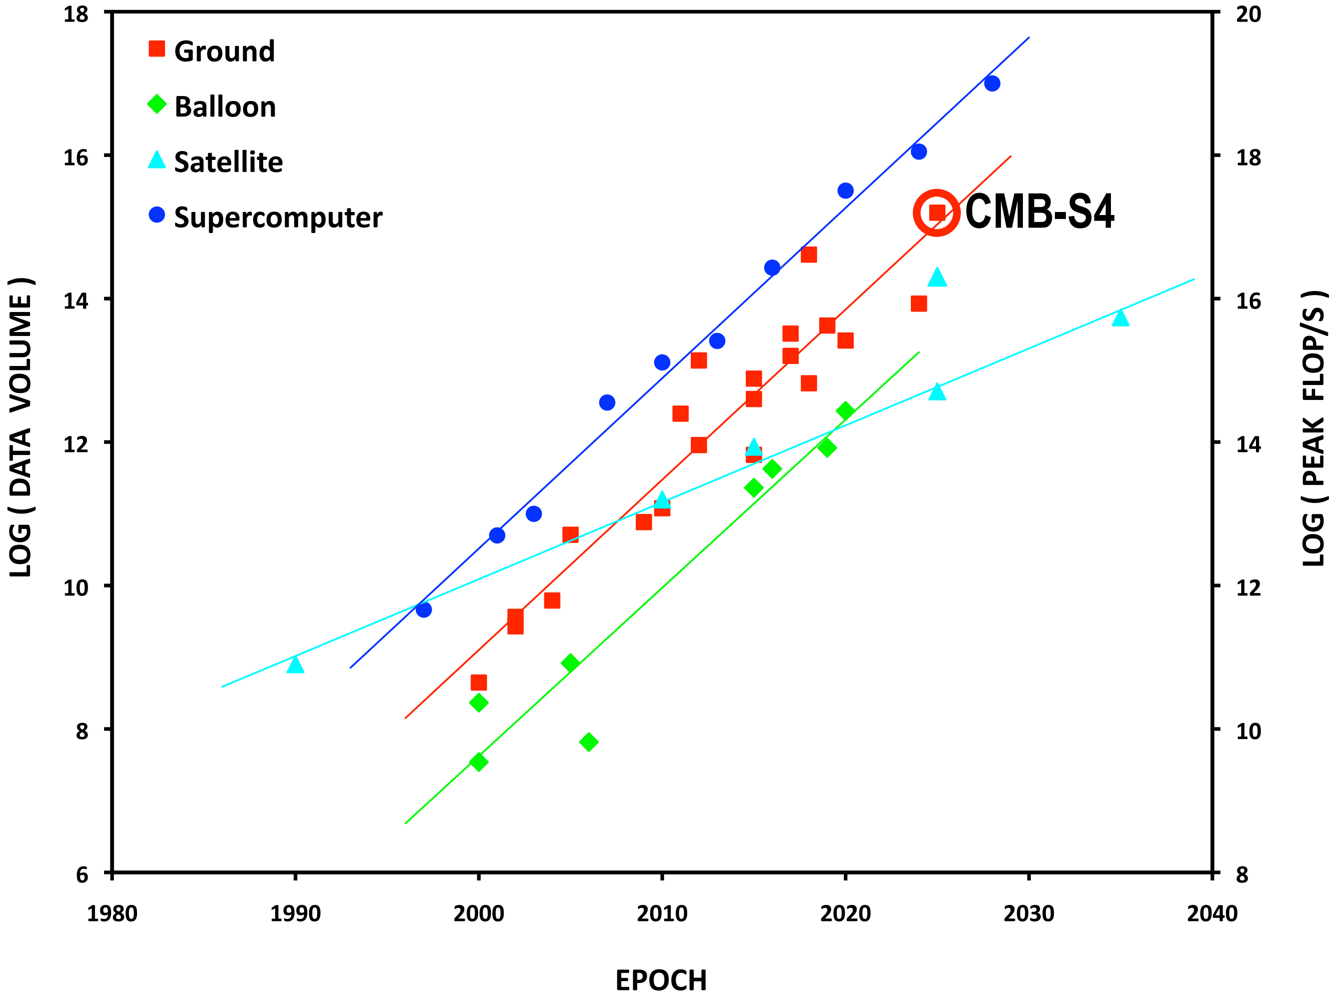

CMB Data Flow

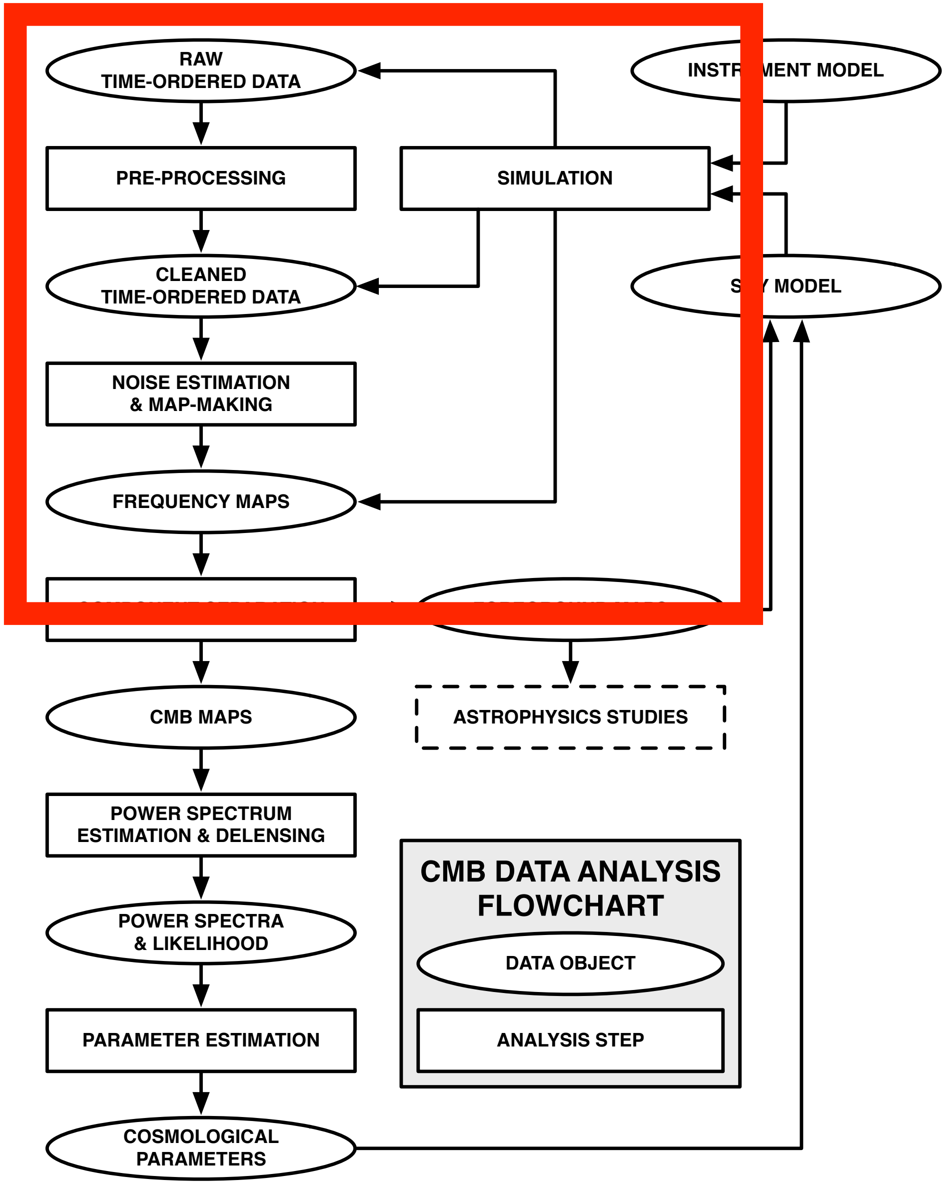

CMB Data Algorithms

EXACT

- Maximize Gaussian likelihoods
- Scales as $O(\mathcal{N}_p^3)$
- High computational efficiency
- $\mathcal{N}_p$ driven up by
  - high resolution
  - large sky coverage
  - many frequencies
  - polarization
- Intractable for
  - $\mathcal{N}_p > 10^6$ in 2000
  - $\mathcal{N}_p > 10^7$ in 2015
  - $\mathcal{N}_p > 10^8$ in 2030

APPROXIMATE

- Monte Carlo replaces covariance
- Scales as $O(\mathcal{N}_t)$
- Low computational efficiency
- $\mathcal{N}_t$ driven up by
  - quest for S/N
  - 10^4 MC realizations
  - many maps per realization
- Intractable for
  - $\mathcal{N}_t > 10^{12}$ in 2000
  - $\mathcal{N}_t > 10^{15}$ in 2015
  - $\mathcal{N}_t > 10^{18}$ in 2030

Planck Monte Carlo

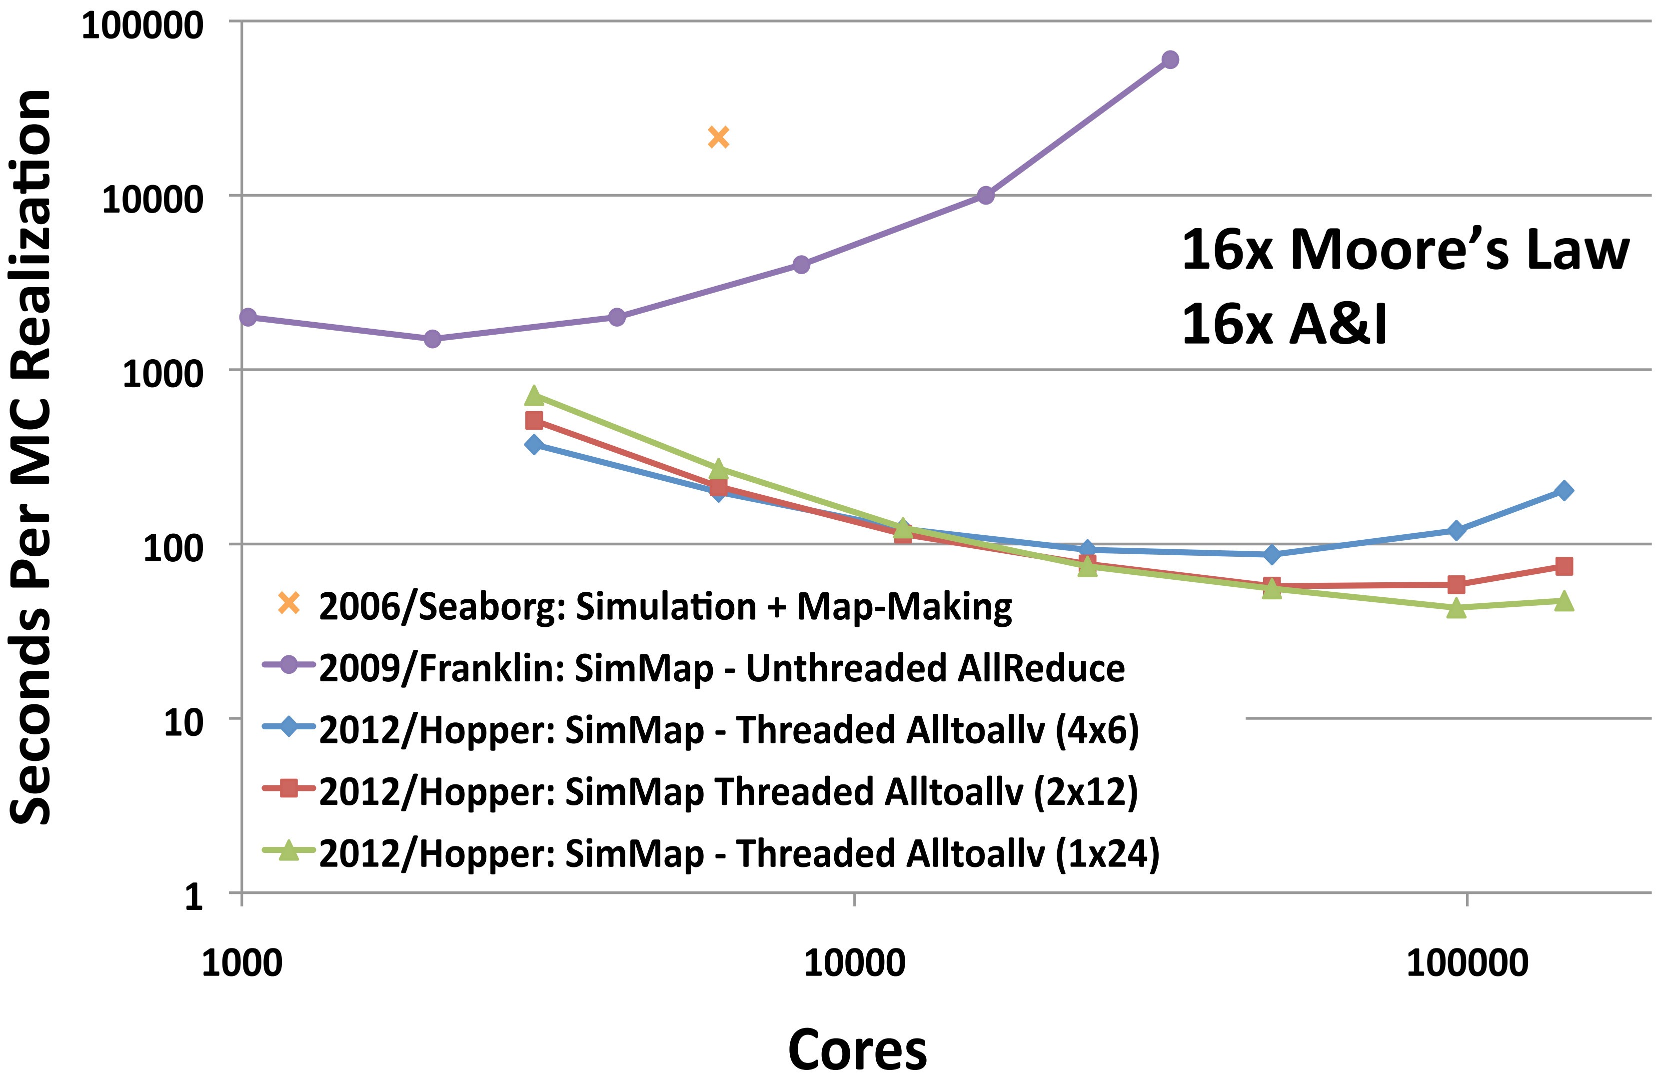

BERKELEY LAB

C3 COMPUTATIONAL COSMOLOGY CENTER

10^4 MC: Planck Noise

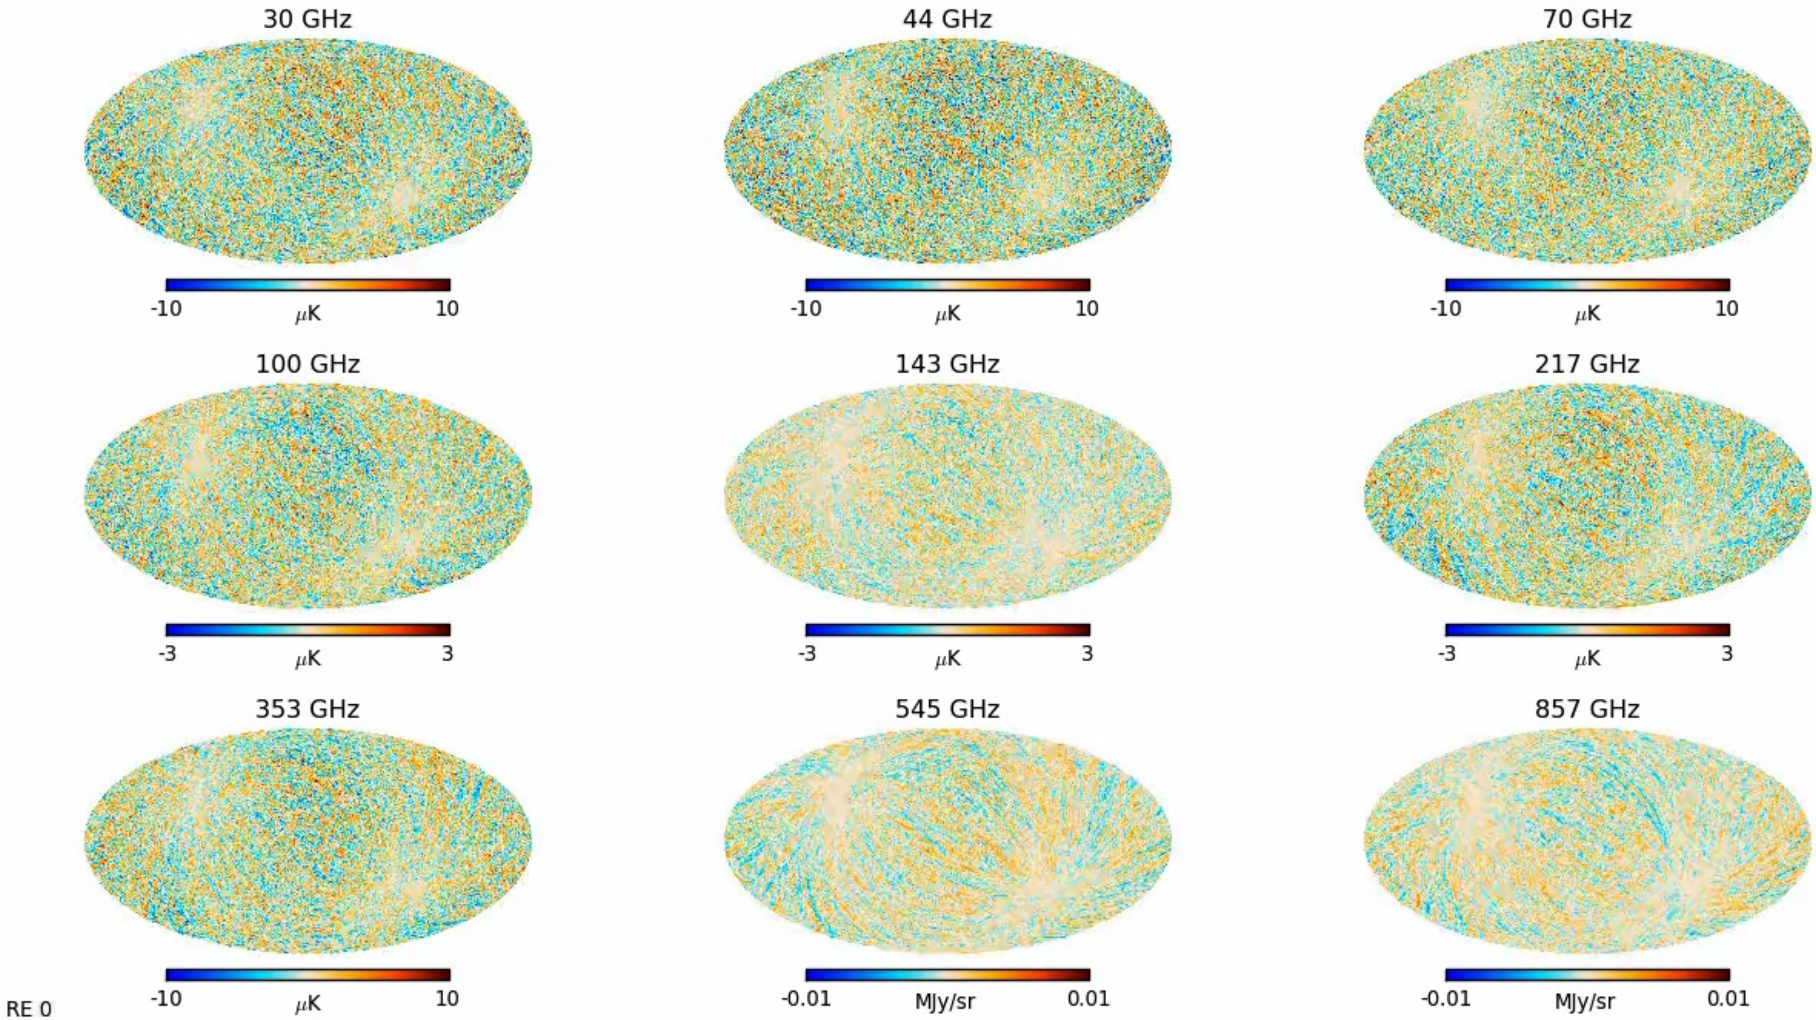

~50 million CPU-hours!

Future Challenges

- Precise analysis of forthcoming B-mode experiments requires our algorithms and implementations to track Moore's Law for
  - 10+ years.
  - 3+ generations of supercomputer.
- The means of achieving Moore's Law also evolves
  - Clock speed, node count, accelerators/many-core, ???
- HPC is increasingly energy-constrained
  - Energy-aware algorithms, dark silicon
- Overall: increasing heterogeneity, complexity, imbalance.

Conclusion

Realizing the enormous scientific promise of the coming generations of CMB B-mode experiments will critically depend on our ability to exploit the capabilities of the coming generations of supercomputer, whatever challenges their architectures and energy-constraints may pose.