Computational Issues in CMB Analysis, Now & Then

(aka CMB Stuff with Bob & Doug eh!)

Dick Bond & all CMBers on McKenzie

Analysis = Theory + Simulation + Experiment + Phenomenology

CMB experimental timeline

CMB Pipelines, Now (e.g. CBI – some new results, Boomerang (~month), Acbar (~month), WMAP2/3 (sigh, ~month) & Then (QuAD/Bicep, ACT/SPT, Quiet, Planck, Spider)

Computational costs, now & projected then (e.g. the case for large HPC@CITA)

Theoretical Simulations & Monte Carlo analyses:

Early Universe – Acceleration Histories & the Inflation Landscape; Preheating to reheating; defects; topology

Nonlinear Secondary Anisotropies: inhomogeneous reionization; point sources- ULIRGs, radio galaxies ..; tSZ, kSZ (cluster/gp web); lensing - but homogeneous & isotropic

Galactic Foregrounds: template based; polarization frontier; CMB, IRAS/DIRBE, HI, HII, IGPS …; synchrotron, bremsstrahlung, dust - vibrating, spinning, in HVCs, local …

CMBers on McKenzie

CITA

- Bond
- Contaldi
- Lewis
- Pogosyan (U Alberta)
- Prunet (IAP – France)
- Sievers
- Myers (NRAO)
- Pen

UofT

- Netterfield
- MacTavish

Others

- Crill (Caltech)
- Hivon (Caltech)
- Jones (Caltech)
- Montroy (Case Western)
- Kisner (Case Western)

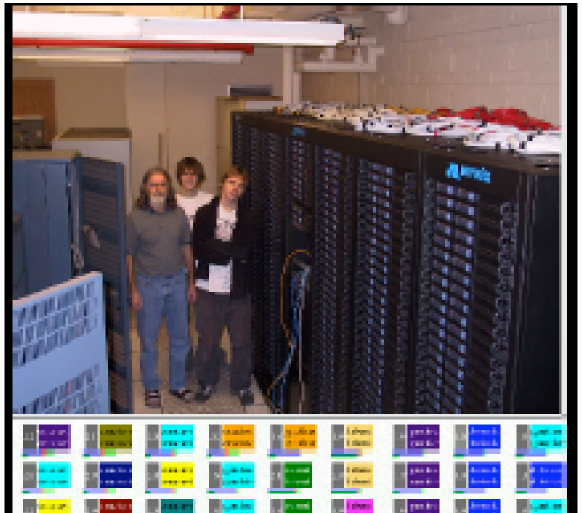

now@cita, Cdn$0.9M

CMB analysis ~ 25%

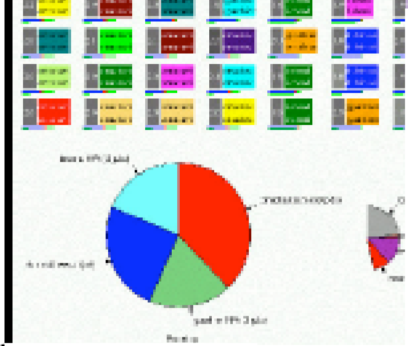

- 264 compute nodes
  - 528 CPUs (2.4GHz Xeons); 35TB of usable local scratch
  - 285 GB RAM
- nodes logically arranged in 2D mesh (See Fig 1)
  - each node connected to two switches
  - each node acts as a router
  - maximum one hop through a compute node for any message
- networking cost < 10% of entire budget ($900K)

then@cita ~ Cdn$30M?

8,100 dual-CPU nodes (2GB/node) and 1PB of storage

nodes arranged in 90x90, 2D-mesh as in McKenzie using GigE

CBI ongoing to Sept'05+

Acbar ongoing to Sept'06+

Bicep

QUaD

Quiet1

Quiet2
(1000 HEMTs)
Chile

APEX
(~400 bolometers)
Chile

SZA
(Interferometer)
California

SCUBA2
(12000 bolometers)
JCMT, Hawaii

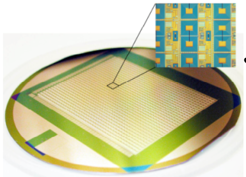

Spider
(1856 bolometer LDB)

ACT
(3000 bolometers)
Chile

Clover

2017
CMBpol

Boom03

2003 2004 2005 2006 2007 2008

WMAP ongoing to 2007+

DASI

CAPMAP

AMI

GBT

Polarbear
(300 bolometers)
California

SPT
(1000 bolometers)
South Pole

Planck
(84 bolometers)
HEMTs L2

ALMA
(Interferometer)
Chile

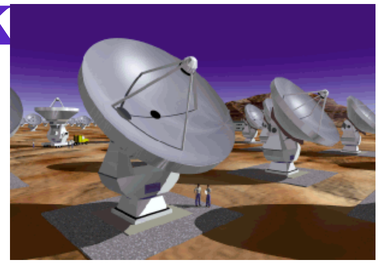

SN1: Oct04 ~100 @ z ~ .3-.7

LIGO1

LIGO2

LISA
2013

~ 10 @ z ~ 1-1.5

~30

2008-12

CFHT-Legacy ongoing to 08 (165 spec, 700 in can) ~400+ SN/5yr

ESSENCE ongoing to 06+ ~150 SN/5yr

WEAK LENSING:

Pan-STARRS

LSST

Oct04: RCS1 53 sq deg, Virmos-Descart 11 sq deg +

2003 2005 2007 2017

2004 2006 2008

JDEM
space

Deep Lens Survey ongoing 28 sq deg

CFHT-Legacy ongoing to 08 (first great results 05) 140 sq deg

RCS2 ongoing 1000 sq deg

KIDS (960 sq deg), UKIDS

SDSS ongoing

CLUSTER/GROUP system in the Cosmic Web

SZ/ PVmeasure: SZA, APEX, GBT, AMI, ACT. SPT, Planck, ALMA

Optical: RCS, RCS2, SDSS + ongoing

Large Optical Surveys for
SZ tomography (DES?, .)

forecast ~ 2008+

Planck1 +

WMAP4 +

SPT/ACT

if (PSB arrays

& 1000 sq deg)

+ QuAD

+ BiCEP

+ Quiet2

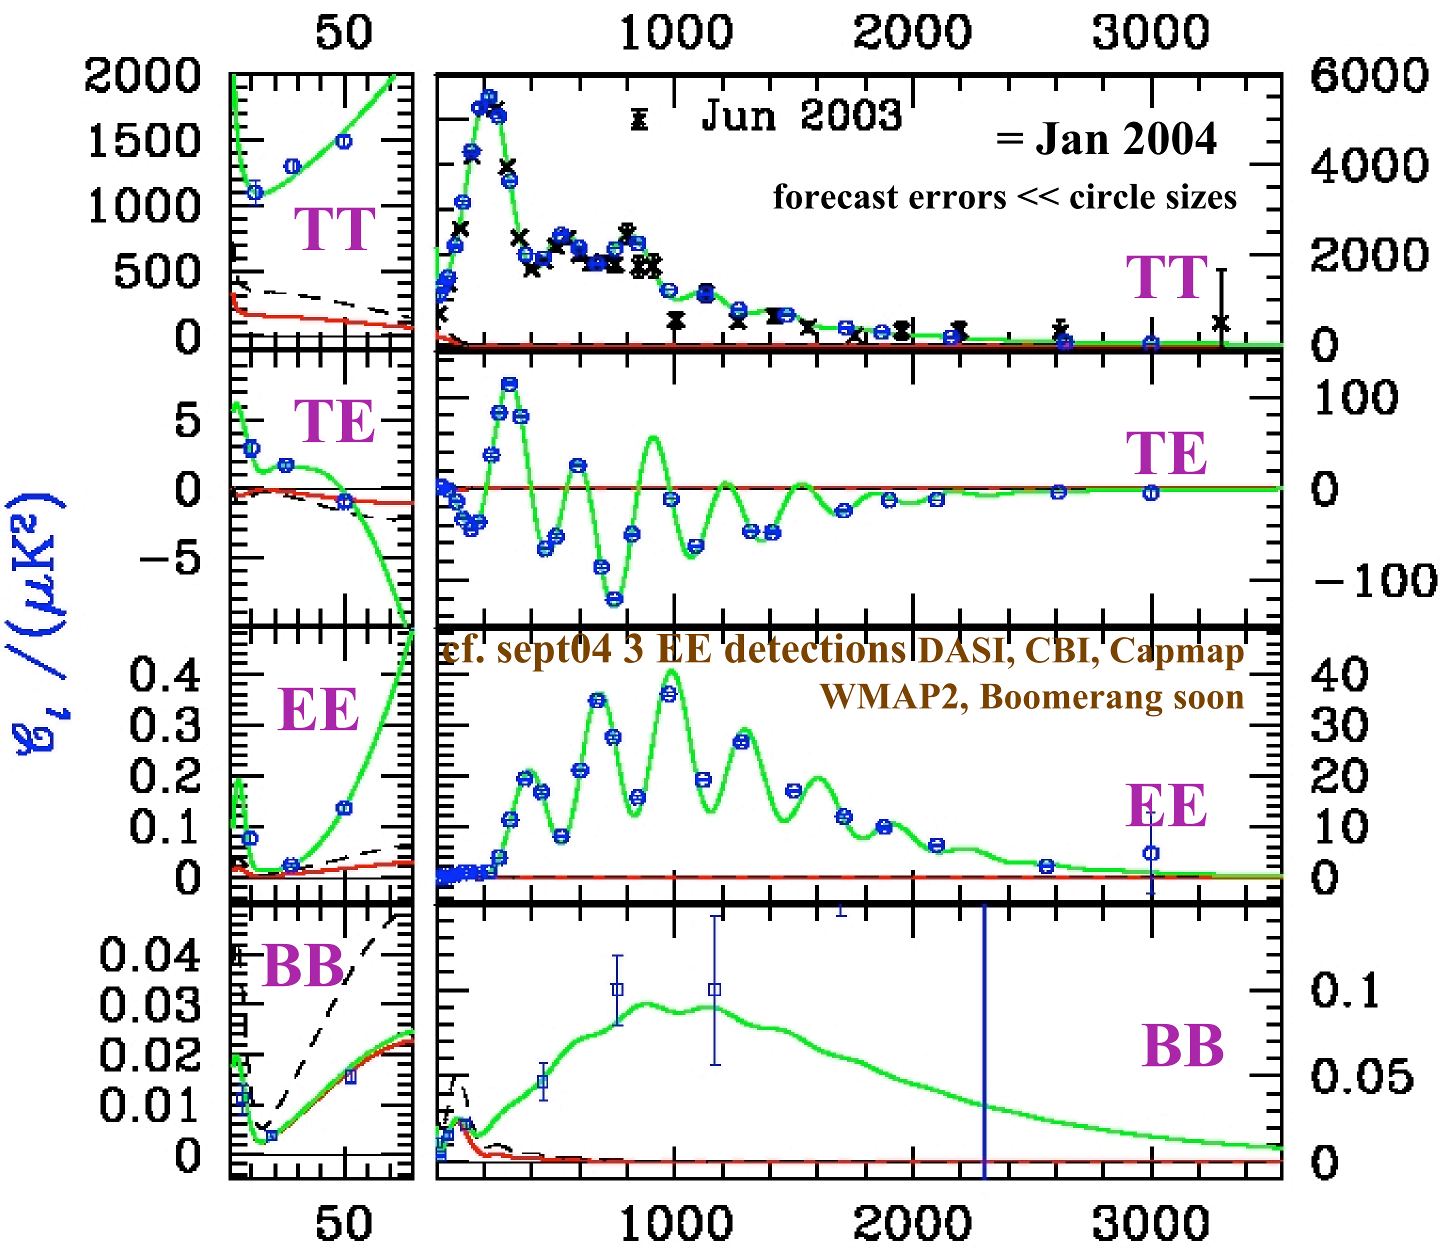

GW/scalar curvature: current from CMB+LSS: r < 0.7 or < 0.36 95% CL; good shot at **0.03** 95% CL with **BB polarization** BUT fgnds/systematics??

some CMB ANALYSIS PIPELINES

from Timestreams or Visibilities (time-chunked interferometers) through generalized maps to bandpowers++ to parameters (via Monte Carlo, MCMC)

Signal-noise separator of ToD & Maps: via mapcumba (B98), MADnes, MADmap; jiqu (B03), newsky (B03), sky2naive (B98,B03}, gridder of uv visibilities (CBI)

Bandpowers via near-optimal isotropized MC QUADest: MASTER (pseudo-CL); SPICE, polspice; FASTER/XFASTER (B98, B03); optimal quadest: BJK, MADCAP (B98), Mlikely/MPIlikely (CBI); hybrids;

Spice, XFaster, Gridder/Mlikely extended to polarization (B03, CBI)

banded higher point stats (B98, CBI, CBIpol, WMAP)

USE healpix or alternative (e.g. ice), fast spherical harmonic transform

Parameters of all sorts via Monte Carlo Markov Chain feasible e.g. COSMOMC (Lewis)**; fixed adaptive grids for some parameter mappings**

some CMB Analysis Actions

Compressing: time ordered data to generalized maps to bandpowers etal to cosmic parameters

Mocking: from naïve forecasts to full simulation end-to-end through the CMB pipelines

Forecasting: power spectrum errors, cosmic parameter errors, usual homogeneous sky coverage in the continuum limit

Constraining Theories: power spectra, parameters, non-Gaussian higher order statistics and pattern indicators. Feedback to early/late Universe, dark matter/energy theorizing, etc.

Cleaning & Separating: cleaning **systematics;** separating **foregrounds, secondary and primary backgrounds**, finding & understanding the **residuals** in the data. fully characterized separated maps (mean plus correlated errors).

Comparing: One data set or subset to another, often internal to an expt, with different pix, sky coverage, beams, frequencies, jackknifes of all sorts – data-halves, channel-splits, etc. sometimes “Interpolating theory” used, e.g. Gaussian best fit model. Are the sets compatible? If not, why not? The residual hunt.

some CMB PIPELINE COMPUTERS

B98,B03,CBI CITA: 538-CPU. 256-node xeons 1.5 Tflops ; (cf.16-CPU wildfire SMP) ~25% of all McKenzie cycles have gone to CMB projects

WMAP: 6 32-node origin 300s, 1 16-node origin 2000, 3 16-node altix itanium SMPs; 12-node linux dedicated to beam; use 3 map-makers (one Wright one)

Planck: **UK** Cosmos: 2 64-node altix itanium SMPs connected as 128-node; **France** ~ 100 nodes + ; **Germany (MPA++), Spain (++), …**

US NERSC: in 04-05, 0.5-1 ExaFlop = 0.2-0.35 Tflop/s; ExaFlop = 10^{18} flops

Planck USPDA estimate of 14 Tflop/s (41 ExaFlop) dedicated in 2009. ~ 3100 CPUs/yr

Spider (balloon borne CMB polarization on large scales) will need equivalent of 2700 CPUs/result/year (assuming small efficiency factor); ACT 6200 CPUs/ result/yr

"efficiency" fudge factor, 10 as a minimum cf. B03 ~ 200, CBI ~ 50

i.e., dedicated access to ~10000 CPUs needed for Planck, Spider, ACT etal. analyses

CMB Analysis Pipelines

•Step 1 : Time stream filtering ,cleaning, deconvolution, calibration, noise estimation, pointing determination etc…

$$\mathbf{d_t} = \mathbf{P_{tp}}\boldsymbol{\Delta}_\mathbf{p} + \mathbf{n_t}$$

•Step 2 : Map making

$$\chi^2 = (\mathbf{d} - \mathbf{P}\boldsymbol{\Delta})^\dagger \mathbf{N}^{-1} (\mathbf{d} - \mathbf{P}\boldsymbol{\Delta})$$

$$\tilde{\boldsymbol{\Delta}} = \left(\mathbf{P}^\dagger \mathbf{N}^{-1} \mathbf{P}\right)^{-1} \mathbf{P}^\dagger \mathbf{N}^{-1} \mathbf{d}$$

•Step 3 : Power Spectrum estimation

$$\ln L = -\frac{1}{2}\{\Delta^\dagger C^{-1} \Delta + \mathrm{Tr}[\ln C]\}$$

$$\mathbf{W}_{tot} = C^{-1} = (C_N + C_T + C_K + C_{res})^{-1}, \quad C_K = \gamma\ \mathbf{K}^{-1}\ \gamma^t$$

into time chunks to treat non-stationary and non-Gaussian & bad data

Gap filling because of non-whiteness

templates γ = modes

(temporal, spatial, frequency, polarization dependent) e.g. systematic modes in data, YLM patterns, measured foreground templates, source patterns to be projected out, pixels in "position space", in "momentum space" (interferometry), splines, …

extension to polarization: same algorithms, larger matrices

d_{cpt}, P_{cptx}, $N_{cpt,c'p't'}$,

c = channel, p= T,E,B pol, t=time-bit, x="pixel" aka mode/template

channel-channel cross-correlations

polarization cross-correlations TT, EE, BB, TE, TB, EB – leakage

Δ_{cpt} $C_{N(cpx,c'p'x')}$ $C_{S(cpx,c'p'x')}$ $C_{res(cpx,c'p'x')}$

$x=1, \ldots N_{pix/modes}$ $S=1, \ldots N_{sigs}$

channel-channel cross-correlations (single-frequency compression)

polarization cross-correlations TT, EE, BB, TE, TB, EB in theory

Cross terms imply much larger analysis demands

CMB Statistics: Beyond Isotropic Bandpowers

In compression staqes loss of information that is not essential for some is crucial to others. e.g. statistical anisotropy of foregrounds & topology cf. isotropized power spectra for inflation - highly reduced quadratic (data V data) space. Full pixel-pixel covariance for topology: e.g. SOCCER BALL Universe.

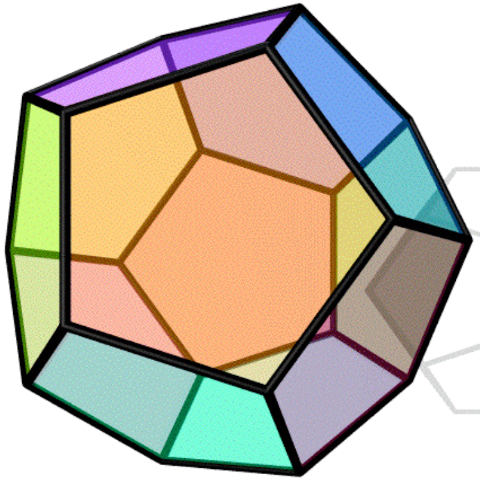

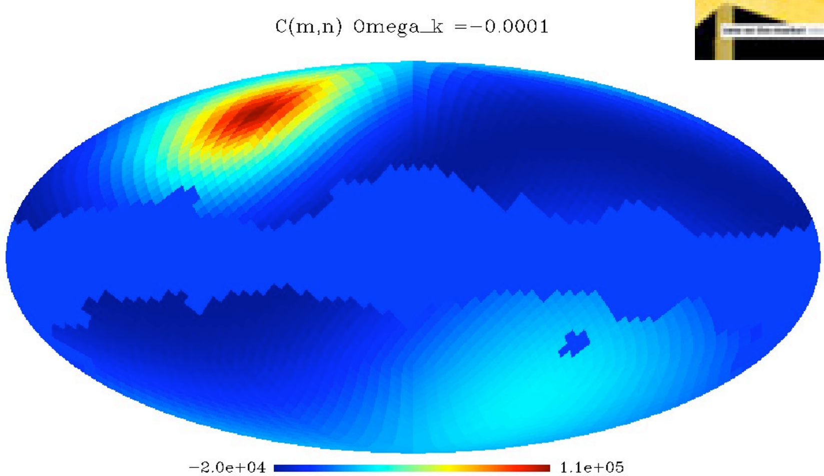

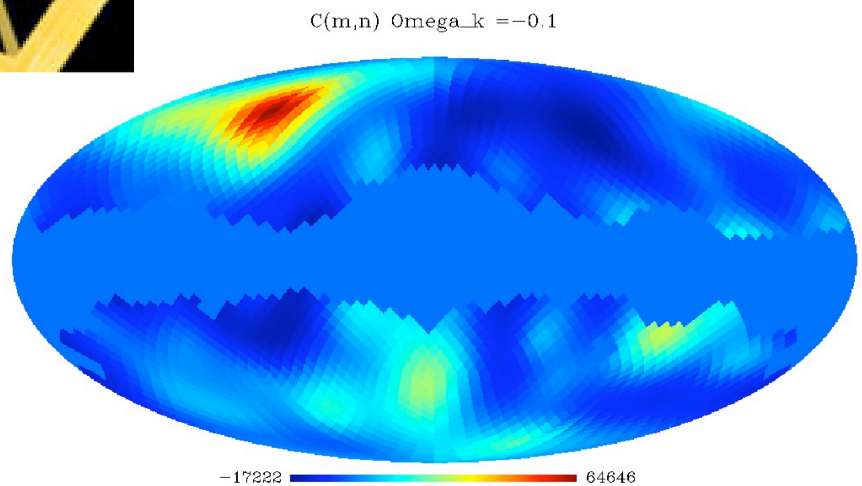

Cosmic Background Imager Polarization

- 13 element interferometer @ 30GHz
- 5000m Atacama Plateau, Chile
- Polarizers in Oct02. HEMTs RL pol
- 2+ yrs of data collected (to Jan05 +40%) Compact array optimized L~600-800

WMAP1 “synchrotron” map

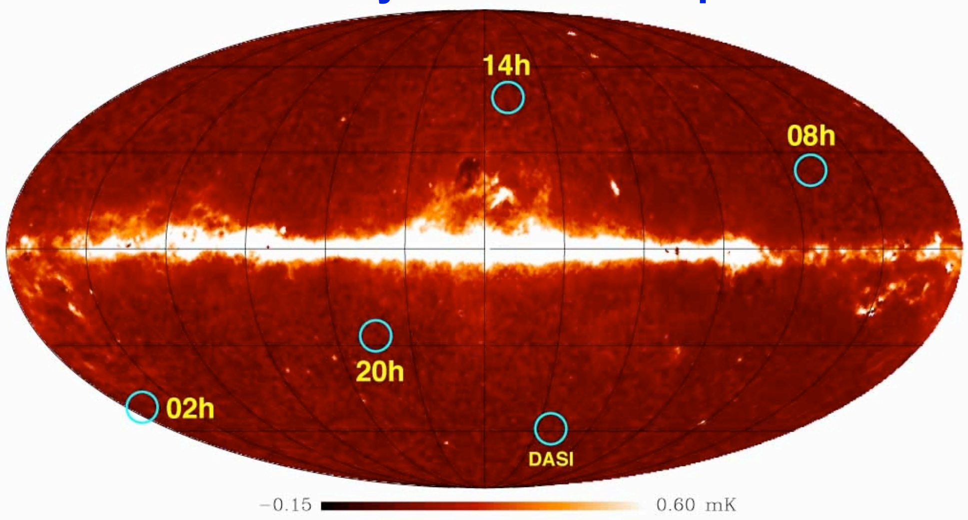

CITA ICAT
CIAR

interferometry primer

Measures visibilities = intensities in baseline-dependent Fourier mode convolved with the dish antenna-pattern.

$$V(\mathbf{u}) \approx \tilde{A}(\mathbf{u}) \otimes \tilde{T}(\mathbf{u})$$

Primary beam transform

CMB intensity transform

$$\ell \approx 2\pi|\mathbf{u}|$$
$$C(|\mathbf{u}|) \approx C_\ell$$

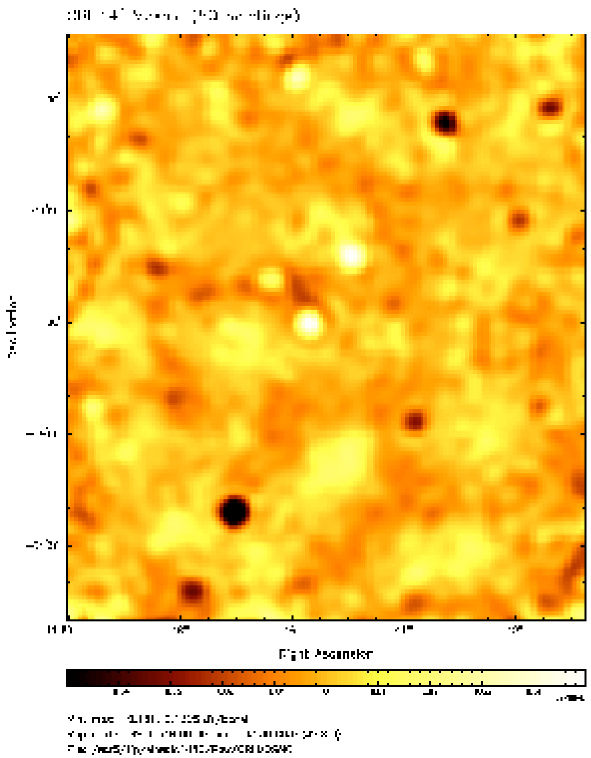

uncorrelated noise

$$C_{kk'} = \langle V_k V_{k'}^* \rangle + 2\sigma_k^2 \delta_{kk'}$$

We compress onto a coarse-grained (u,v)-plane lattice.

Polarization – Stokes parameters

- **CBI receivers can observe either RCP or LCP**
  - cross-correlate RR, RL, LR, or LL from antenna pair

$$\begin{pmatrix} \langle e_R e_R^* \rangle \\ \langle e_R e_L^* \rangle \\ \langle e_L e_R^* \rangle \\ \langle e_L e_L^* \rangle \end{pmatrix} = \begin{pmatrix} I+V \\ (Q+iU)e^{-i2\psi} \\ (Q-iU)e^{i2\psi} \\ I-V \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & e^{-i2\psi} & ie^{-i2\psi} & 0 \\ 0 & e^{i2\psi} & -ie^{i2\psi} & 0 \\ 1 & 0 & 0 & -1 \end{pmatrix} \begin{pmatrix} I \\ Q \\ U \\ V \end{pmatrix}$$

  - CMB not circularly polarized, ignore V (RR = LL = I)
  - parallel hands RR, LL measure intensity I
  - cross-hands RL, LR measure complex polarization R-L phase gives electric vector position angle
   ψ = _ tan-1 (U/Q)
  - rotates with parallactic angle of detector ψ on sky

Decompose polarization signal into "gradient" and "curl modes" – E and B

$$\widetilde{Q}(\mathbf{v}) + i\,\widetilde{U}(\mathbf{v}) = \left[\widetilde{E}(\mathbf{v}) + i\,\widetilde{B}(\mathbf{v})\right] e^{i2\div_{\mathbf{v}}}$$

$$\chi_{\mathbf{v}} = \tan^{-1}\left(v/u\right)$$

E & B response smeared by phase variation over aperture A

$$V_{ij\nu}^{RL}(\mathbf{u}_{ij\nu}) = \int d^2\mathbf{v}\, P_{ij\nu}(\mathbf{v})[\widetilde{E}(\mathbf{v}) + i\,\widetilde{B}(\mathbf{v})]\, e^{i2(\chi_{\mathbf{v}} - \psi_{ij})} + \mathrm{e}_{ij\nu}^{RL}$$

interferometer "directly" measures (Fourier transforms of) E & B!

CBI 2004 Polarization results

• 2nd measurement of the E-type CMB polarization spectrum, best so far (DASI02, CBI04, DASI04, CAPmap04 @ COSMO04) & WMAP1 ’03 TE

• Now 40% more data analyzed – cbi9

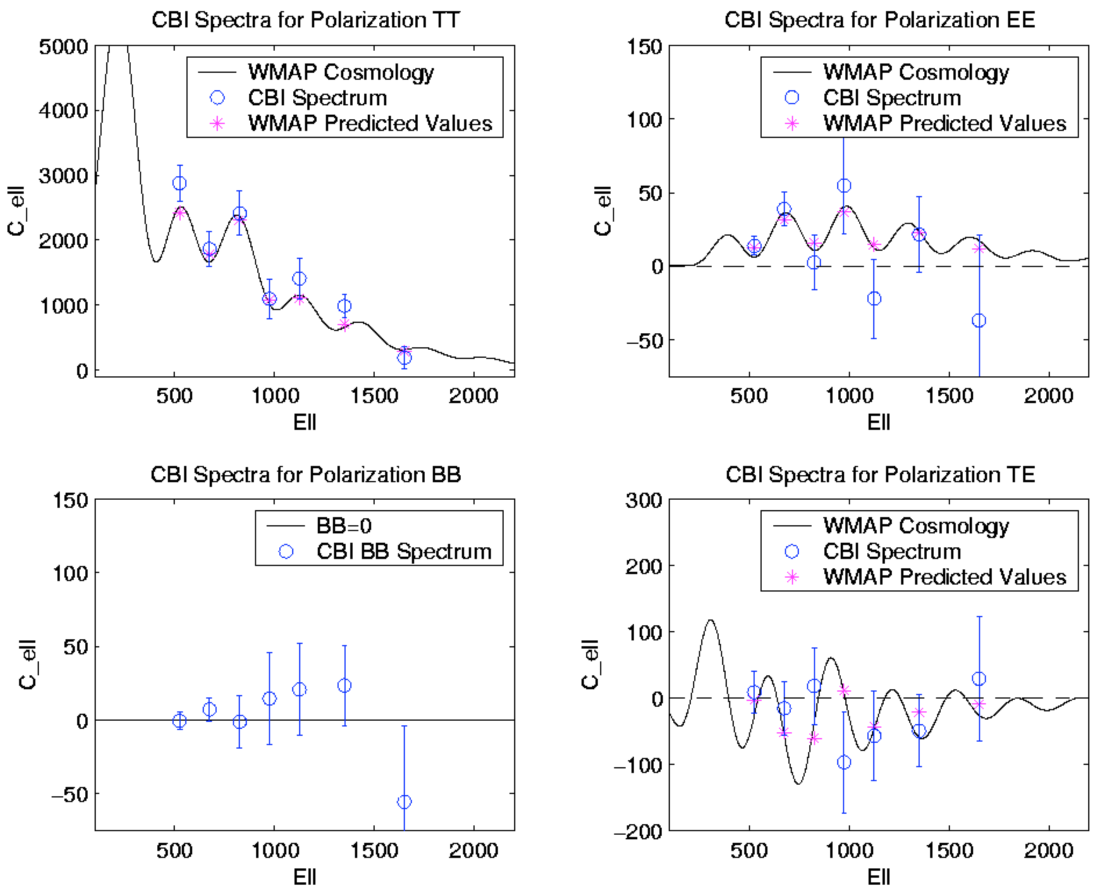

[Readhead et al. Science Nov 2004, , v306]

First Year Polarization Results EE

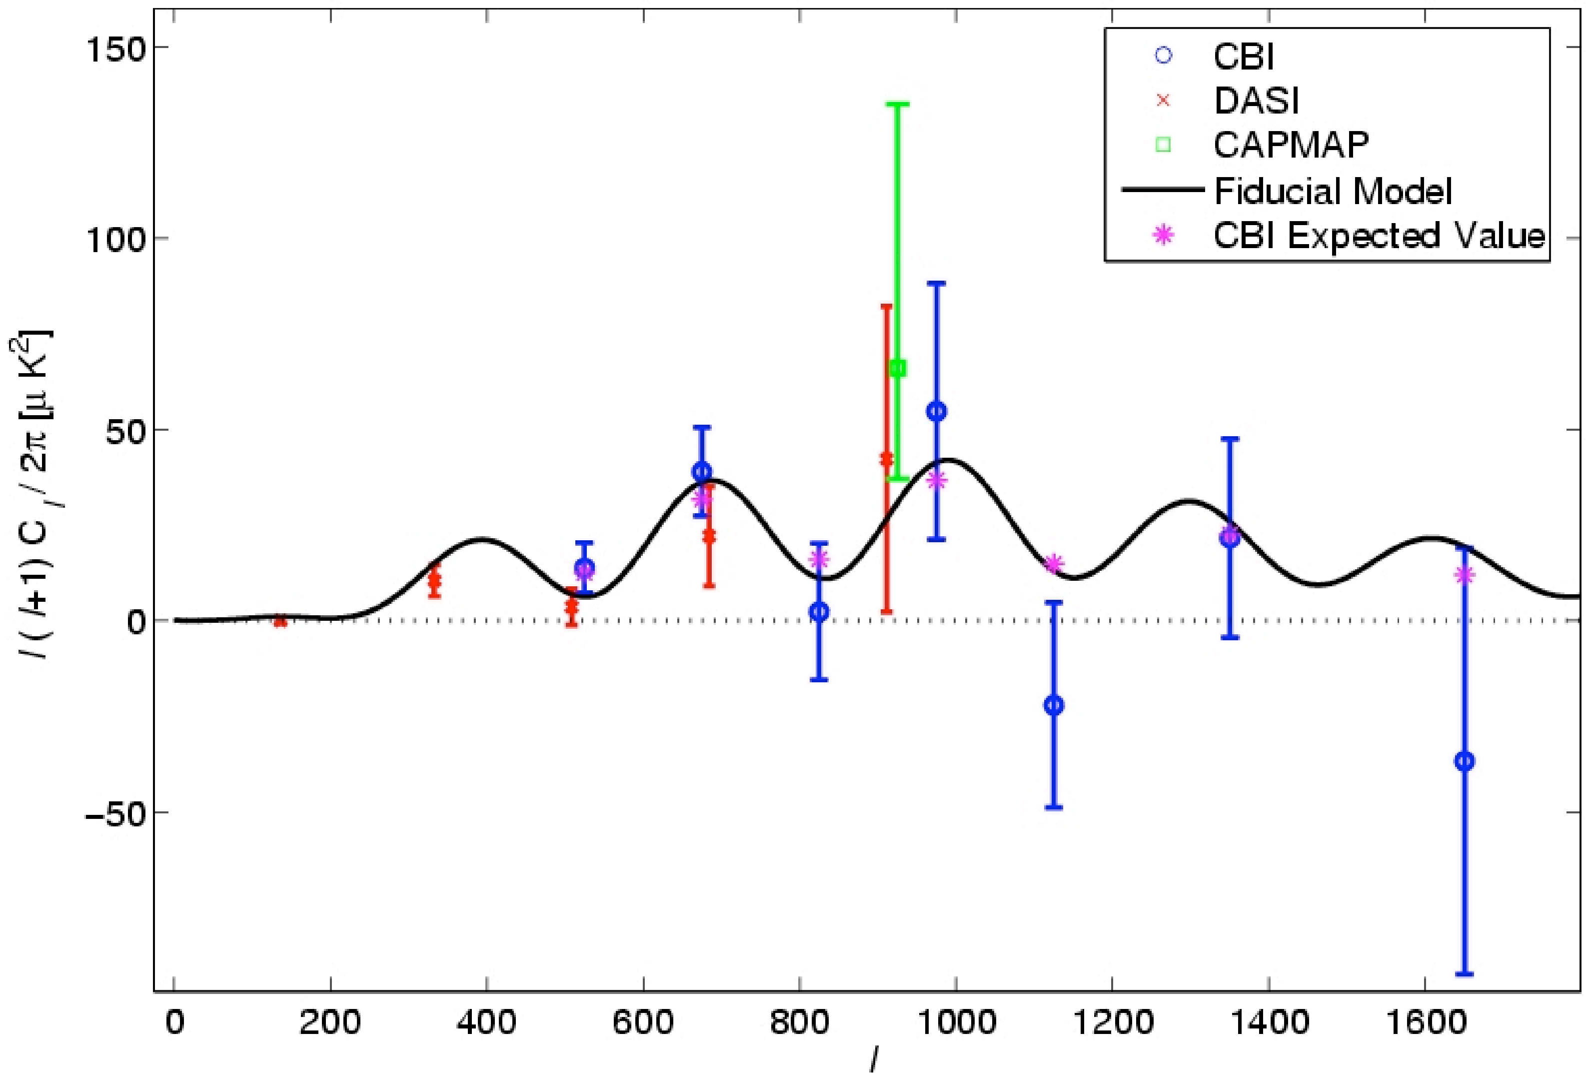

[Readhead et al. astro-ph/0409569]

CBI ~300 nights fall 2002 – Jan 2005 in three 4deg by 4deg mosaics and a 4deg by 45' deep strip.

measure spectrum using a maximum likelihood estimator (CITA McKenzie cluster).

EE detected with high significance (10.7-_), most to date; TE with moderate significance (3.6-_).

TT is consistent with previous measurements. BB consistent with zero, as expected.

- **7-band spectra**
- **(D*l* = 150 for 600<*l*<1200)**
- **consistent with WMAPext (TT from WMAP, ACBAR, 2000 + 2001 CBI)**
- **& with Science'04 pol'n data, 40% more data, errors smaller**

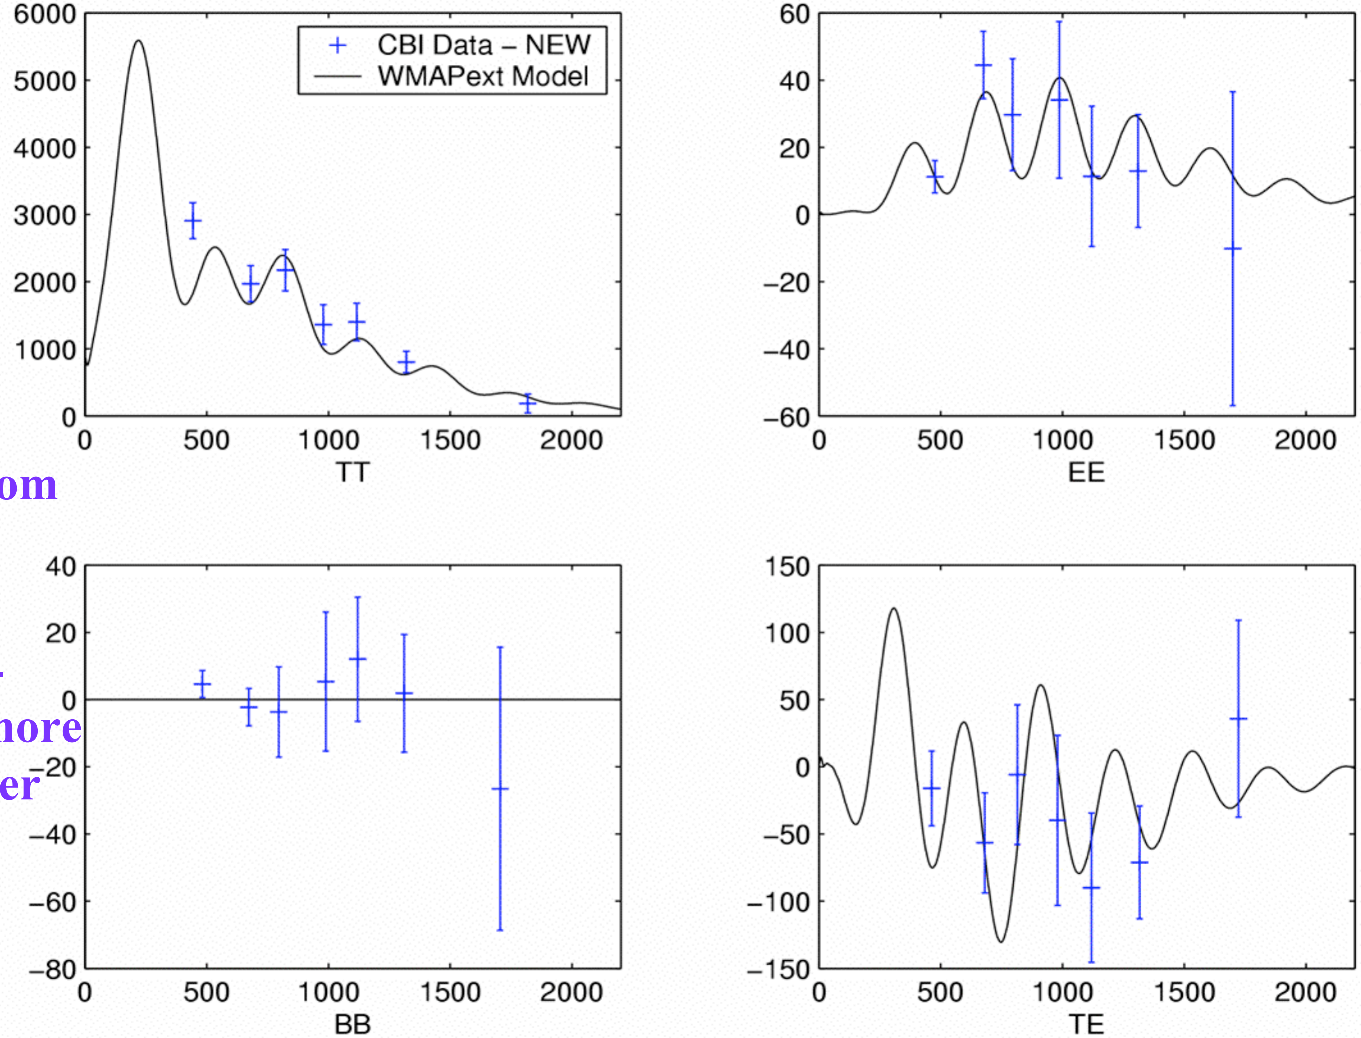

forecast
CBI05

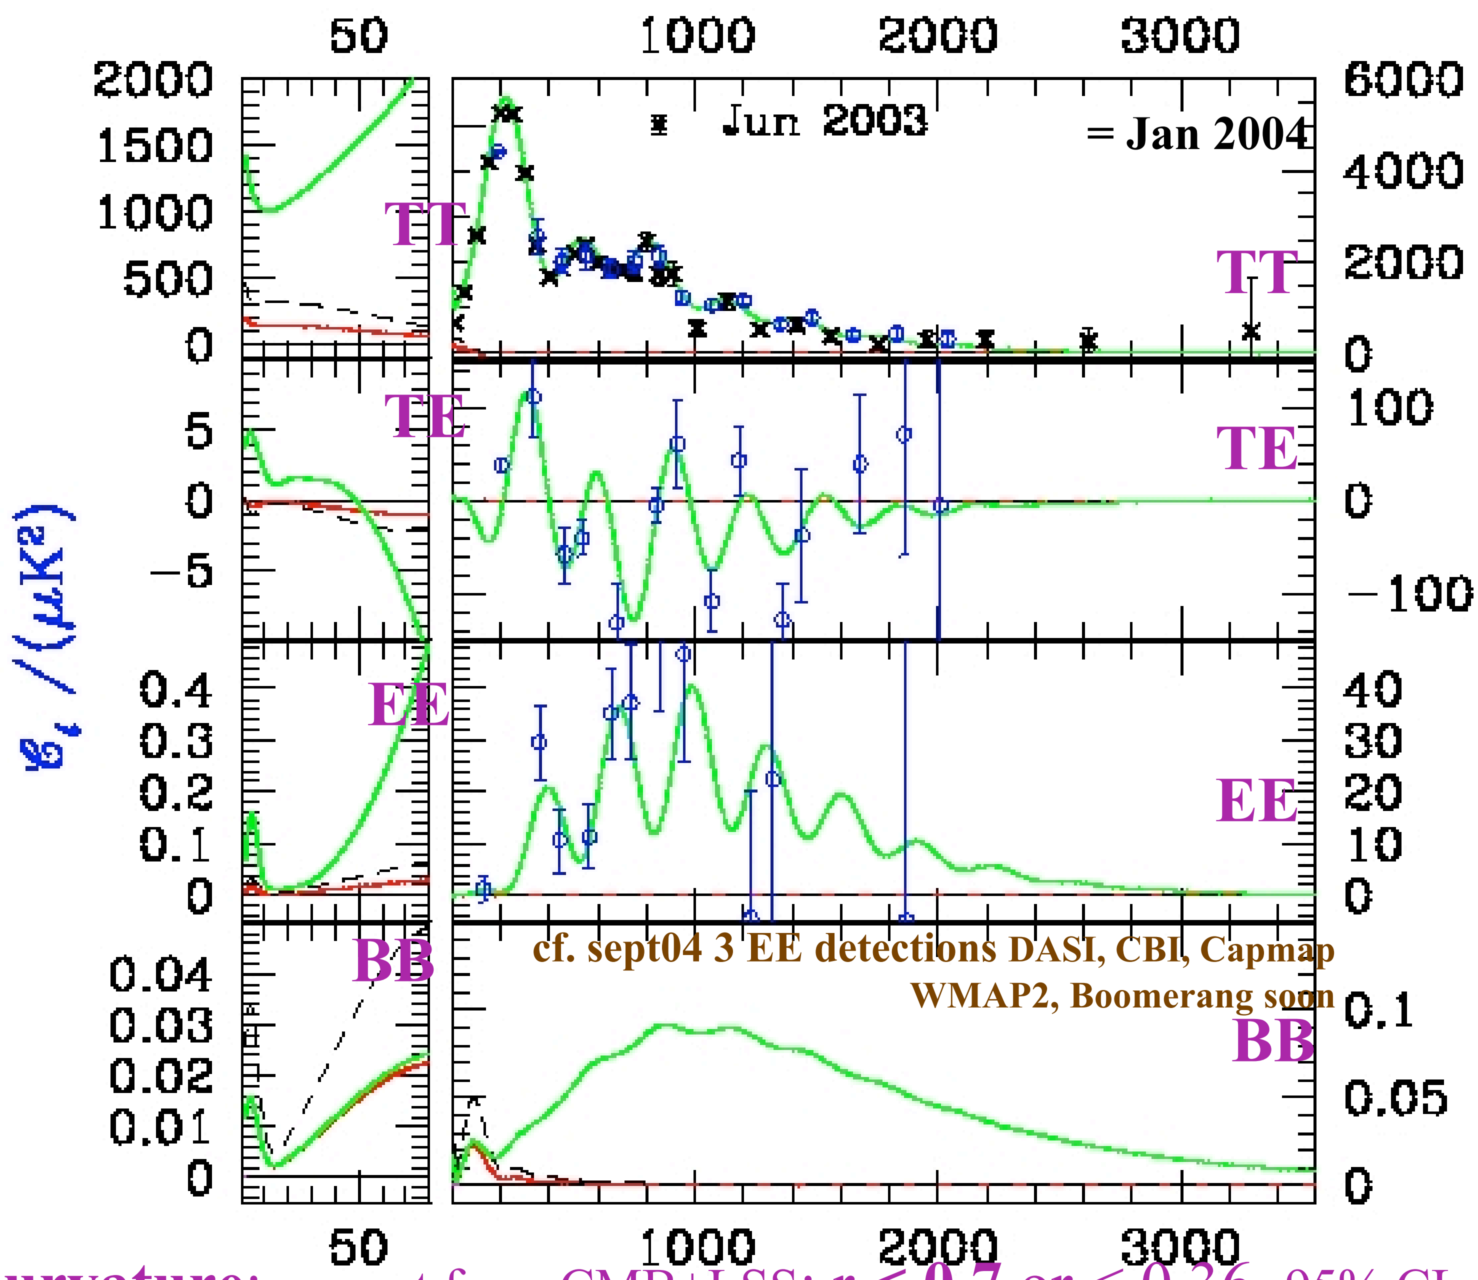

GW/scalar curvature: current from CMB+LSS: r **< 0.7** or < 0.36 95% CL; good shot at **0.03** 95% CL with **BB polarization** BUT fgnds/systematics??

CBI : Compress onto coarse Q-grid & Brute force Power Spectrum

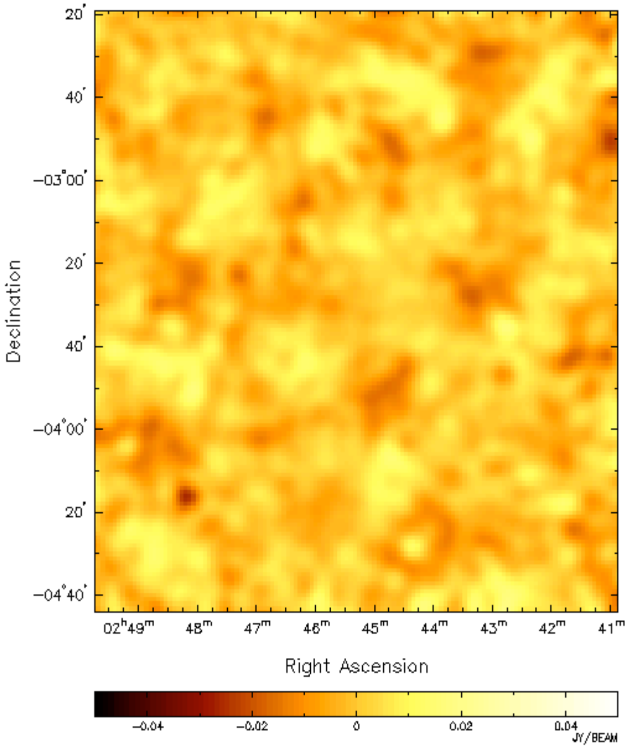

- N_{vis} ~ 10000K (was 100K, finer chunks): reduced onto a gridded set of uv estimators. ~4K sources [MPIGridder]
- N_{pix} ~ 10K (coarse-grid), 40K fine-grid: Sufficiently small to allow brute force search for maximum likelihood using a full iterative (~10) quadratic estimator [MPILikely]
- Storage : $(N_{pix} \times 3)^2 \times N_b / 2 \sim 30$ Gb
- Scaling : $N_{run} \sim (N_{pix} \times 3)^3 \times N_b \sim 2560$ cpu hours
- Codes are parallelized using MPI and use the scaLAPACK (MPI) linear algebra library to solve for $x = M^{-1} y$ [www.netlib.org]

hours + hours in both at 32 McKenzie nodes per field - scale as Area^3

'Efficiency' prefactor ~ 2% → total production time ~ 128,000 cpu hours

A_s n_s $dn_s/d\ln k$

$dn_t/d\ln k$

ω_b ω_c

$dw_{DE}/d\ln a$

Ω_{DE} τ_C A_t n_t w_{DE}

$d^2w_{DE}/d\ln a^2$

Ω_K ω_ν

isocurvature & other subdominant

6 + 1 + 2 + 1 + (1+1) + 1 + (1,2) + 1 + (3+1) + many many more parameters

Any acceleration trajectory for early & late inflation is a-priori allowed, restricted only by the observed data (including "anthropic data" – heat/light, life)

e.g. "blind" search for patterns in the primordial power spectrum : 1+q(ln a), H

e.g. "blind" search for evolution of the dark energy equation of state w(z) : q(ln a)

cf. "guided" searches with theory priors: the cost of baroqueness

CMB futures ~2008++: Planck1+WMAP4+SPT/ACT/Quiet+Bicep/QuAD/Quiet; Planck2.5+Spider

parameter eigenmodes: 6/9 to 1%, rest to 10%

+ Blind-ish search for primordial patterns: 10/35 to 1%, 10/35 to 2%, 9/35 to 10%

Polarization is fundamental to the blind pattern search: T >> E >> B modes

Phase recognition in EE

- *In the standard, scale invariant, pure adiabatic model the phase of the scalar EE spectrum is fully correlated with that of the TT*
- *Polarization sourced by the velocity term at last scattering → peaks are in phase with dips (doppler contribution) in TT*
- *CBI polarization results have begun to test this prediction*

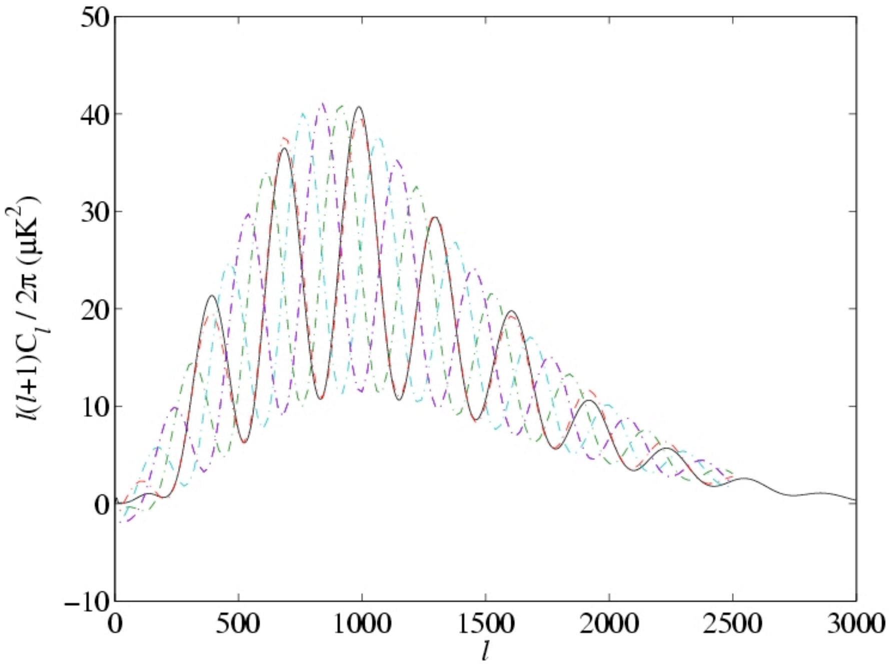

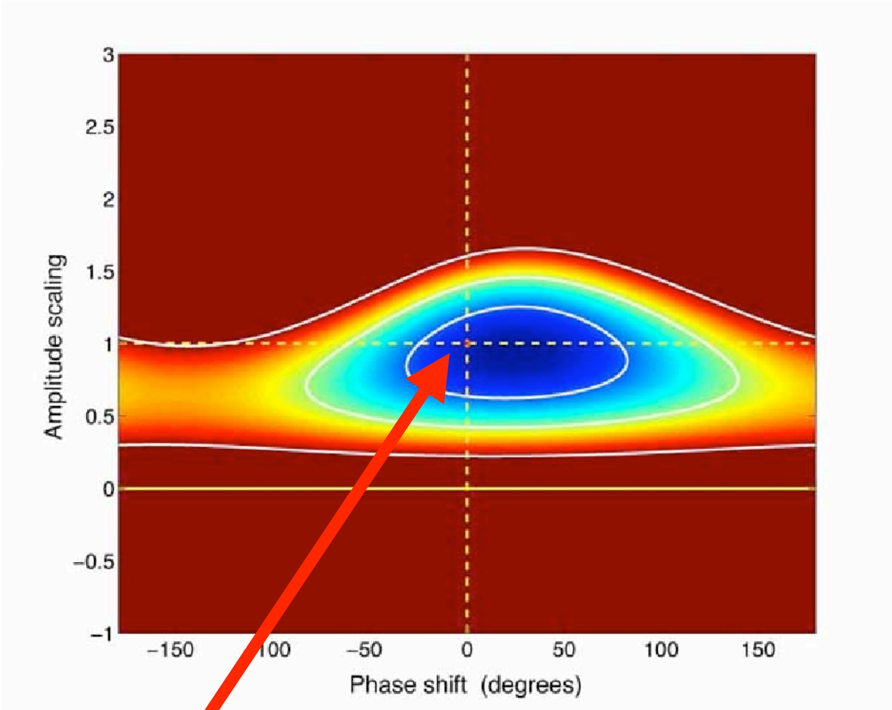

From TT concordance model

ò
ò0 - *An Independent test of origin of perturbations*

DASI→ CBI→ DASI+CBI→ CMB TT

Current CMB constraint = 0.998 +/- 0.005 (WMAP1+CBIpol TT/TE/EE)

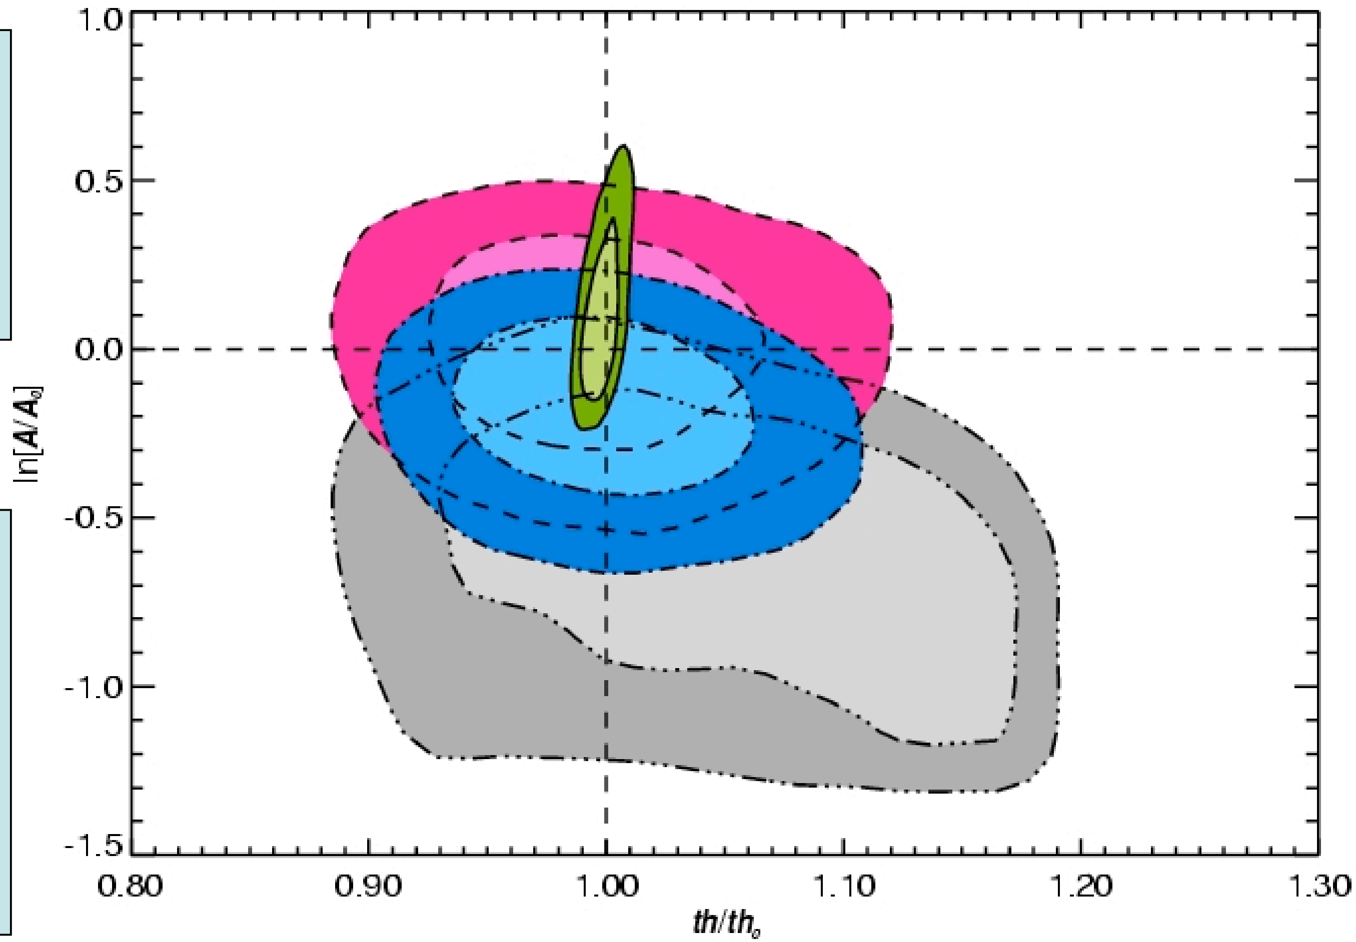

iff TT and EE agreement = no radically broken scale invariance, significant isocurvature modes, etc…

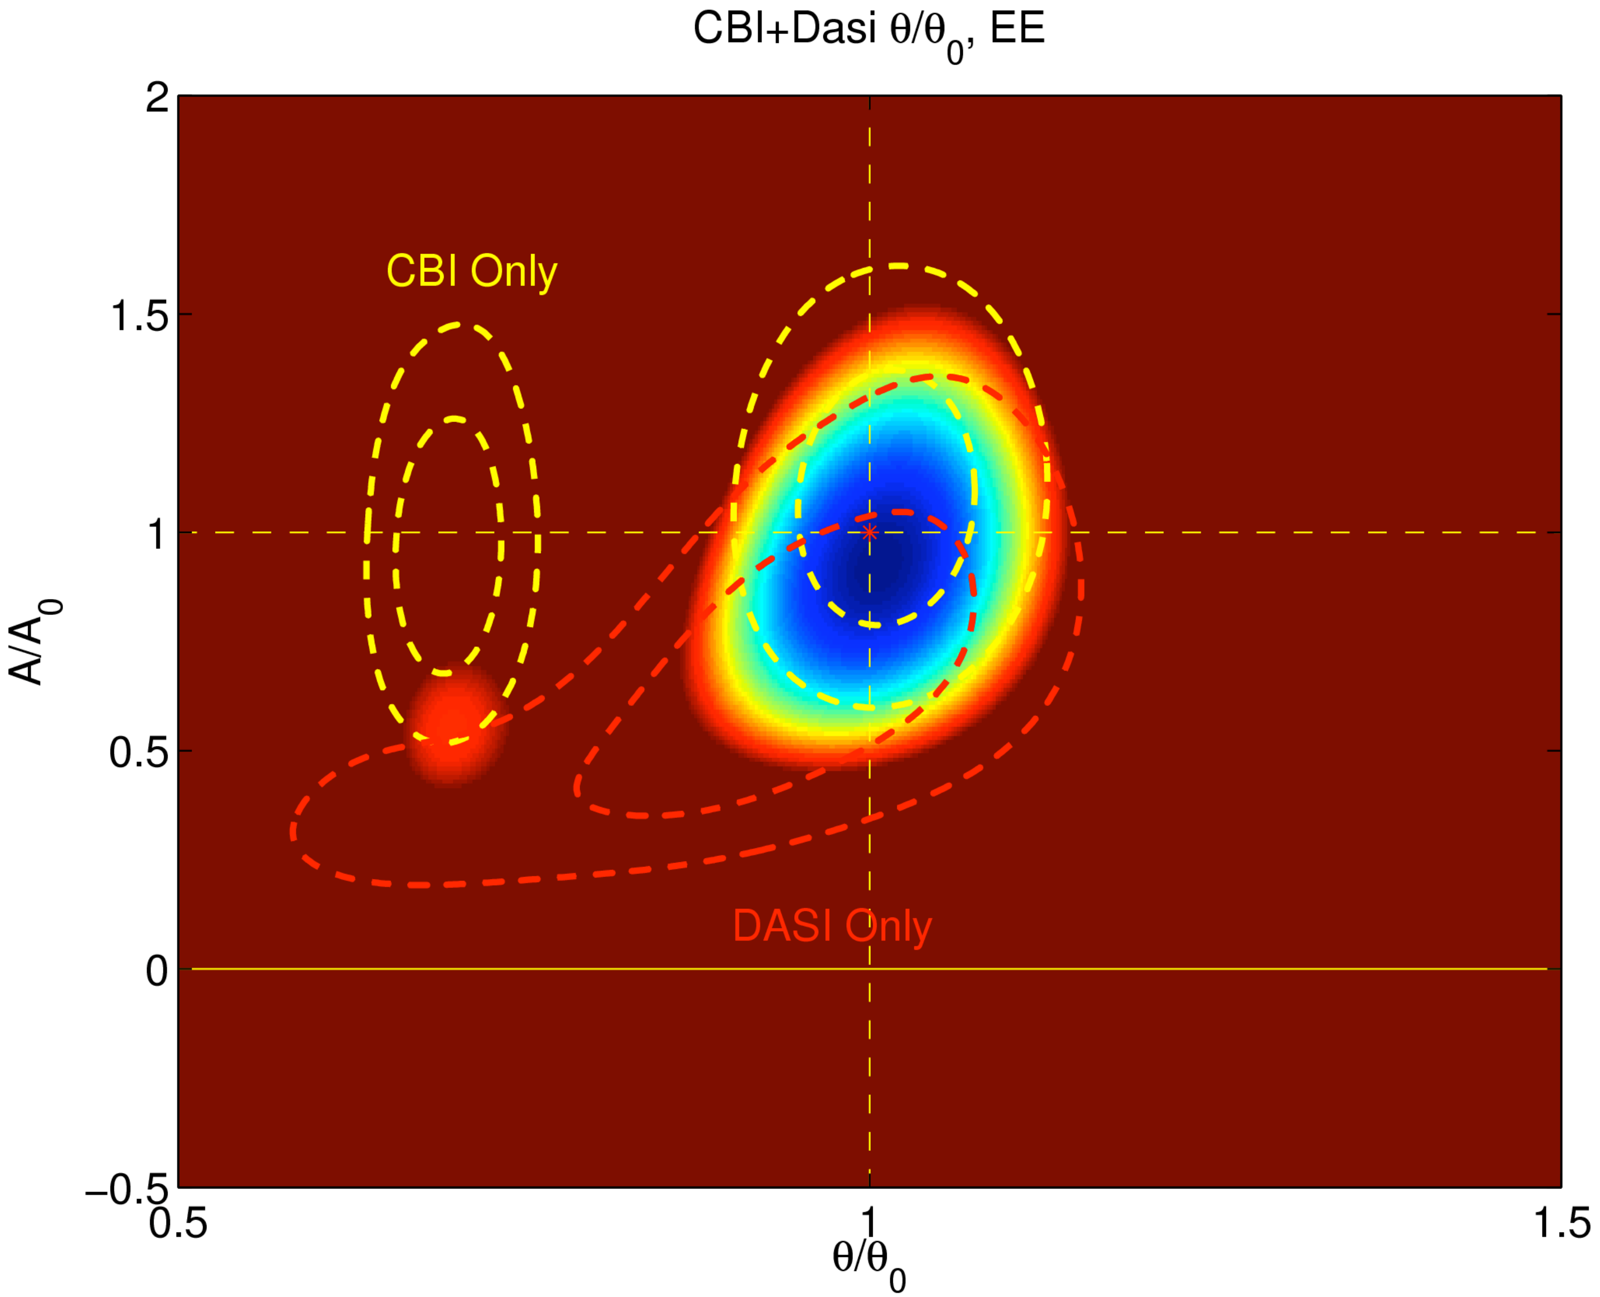

CBI+Dasi θ/θ0, EE
CBI Only
DASI Only
A/A0
θ/θ0
2
1.5
1
0.5
0
−0.5
0.5
1
1.5

Are there any isocurvature modes?

- ***Perturbation of the entropy (in one or more species e.g. baryons, CDM, photons etc…) as opposed to perturbation in the curvature***

$$\hat{\imath} = \hat{\imath}_R + \hat{\imath}_S$$

- ***Lots of models allow for isocurvature modes e.g. multiple-field inflation, curvaton models etc. Data does not allow for too much isocurvature contribution however a subdominant component will bias standard model parameters***

- ***Overall contribution to even/odd peaks depends on species perturbed. Isocurvature modes and adiabatic modes can be correlated!!***

$$C_\ell^{tot} = C_\ell^{SS} + C_\ell^{RR} + C_\ell^{RS}$$

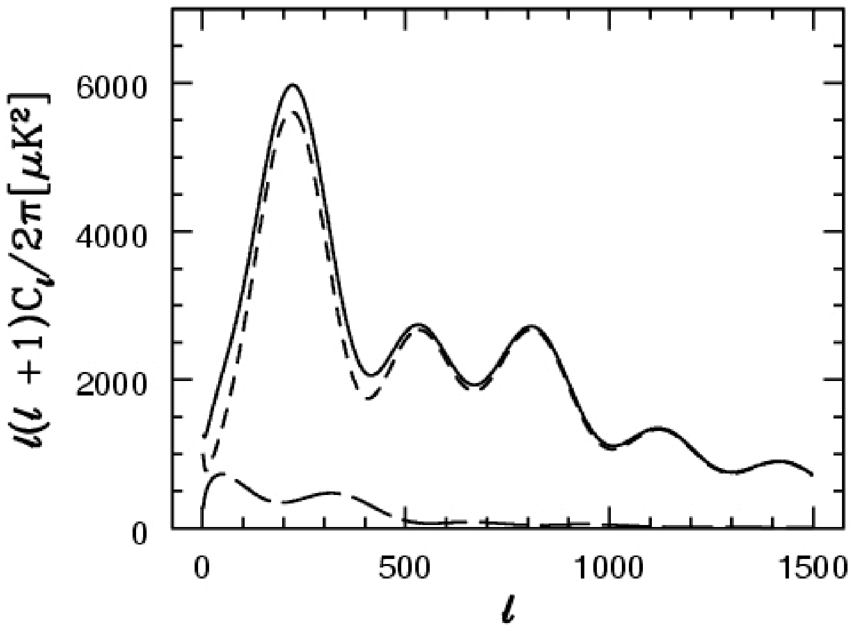

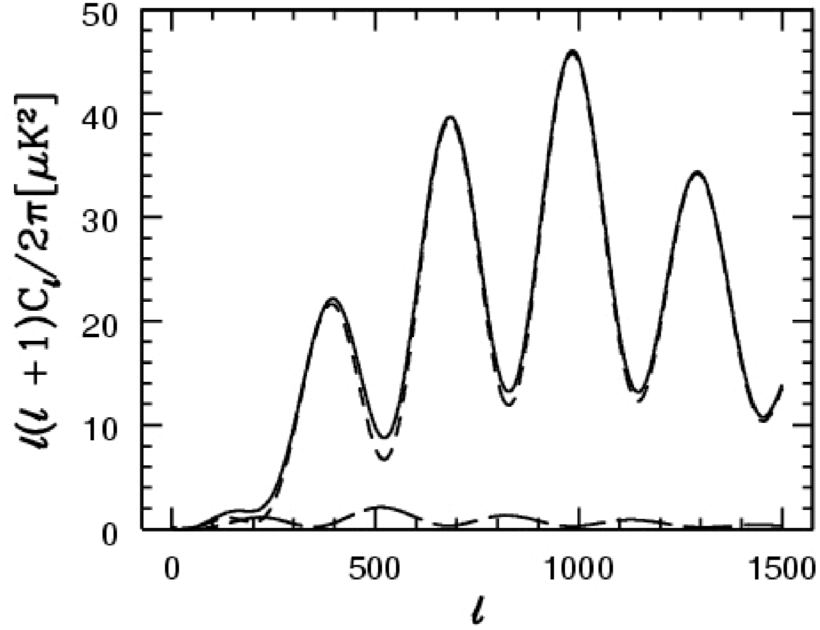

Sample CBI results: Subdominance of isocurvature mode cf. inflationary curvature modes even with just the polarization data

adiabatic amplitude to best-fit inflation

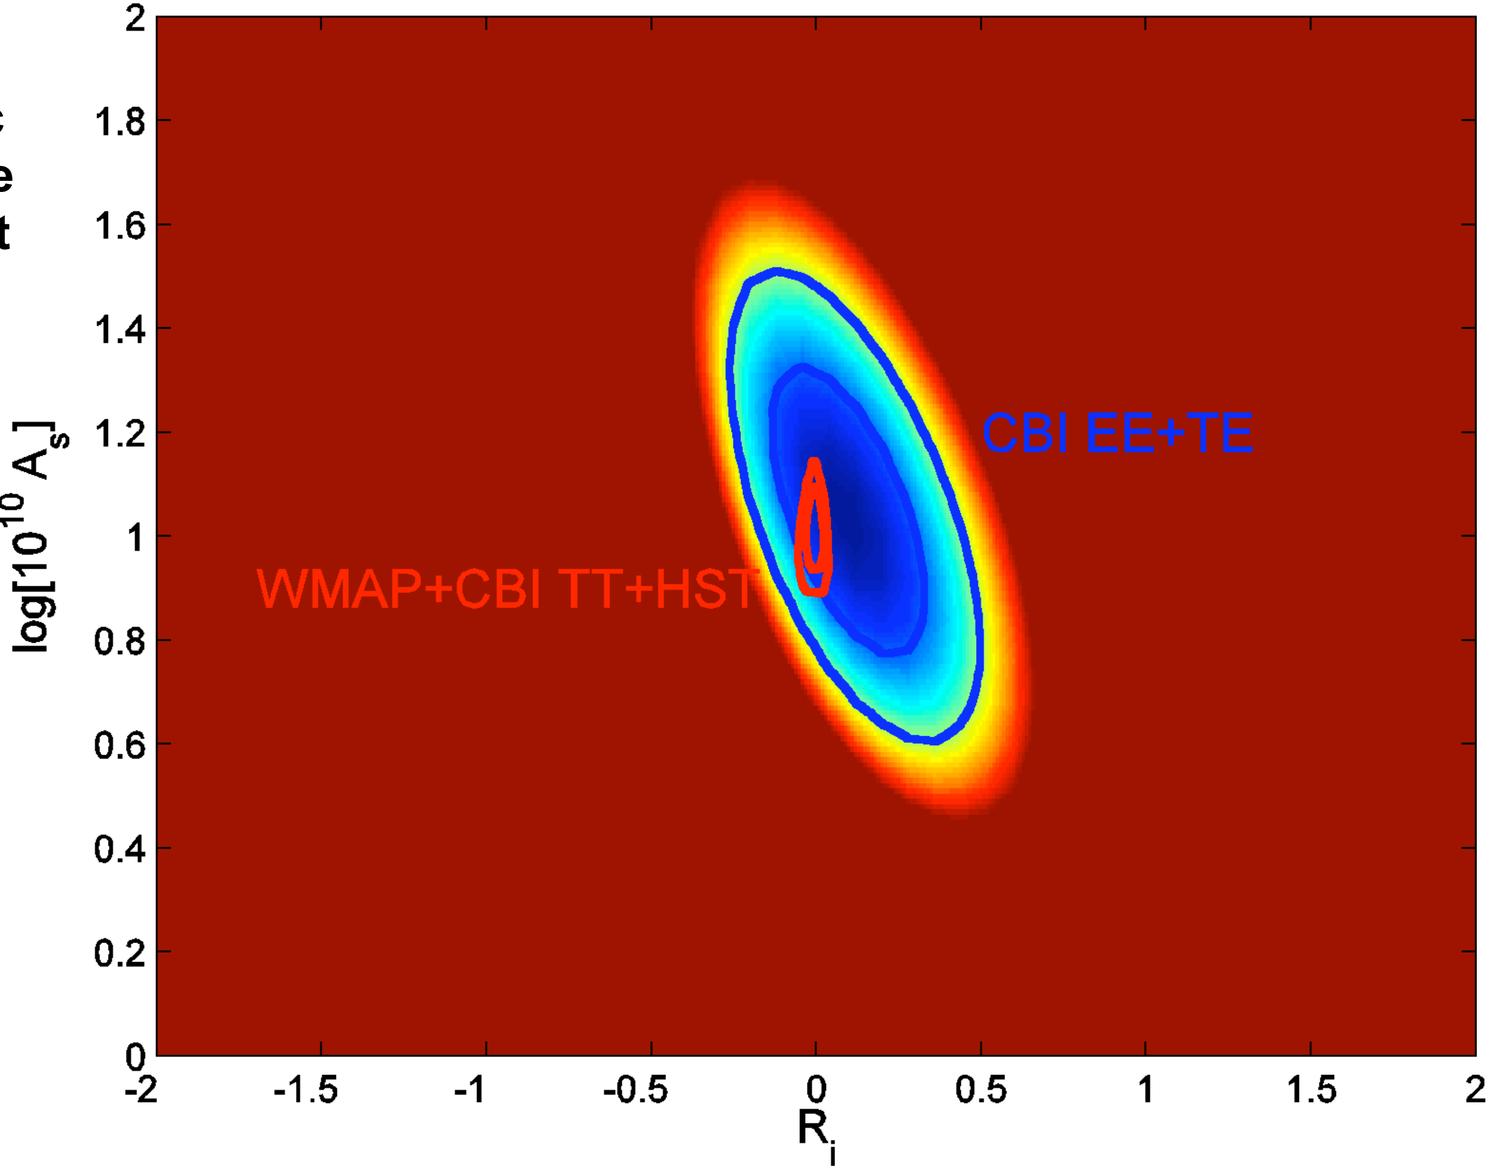

Isocurvature to adiabatic amplitude ratio

CBI 2000+2001, WMAP, ACBAR, BIMA

Readhead et al. ApJ, 609, 498 (2004)

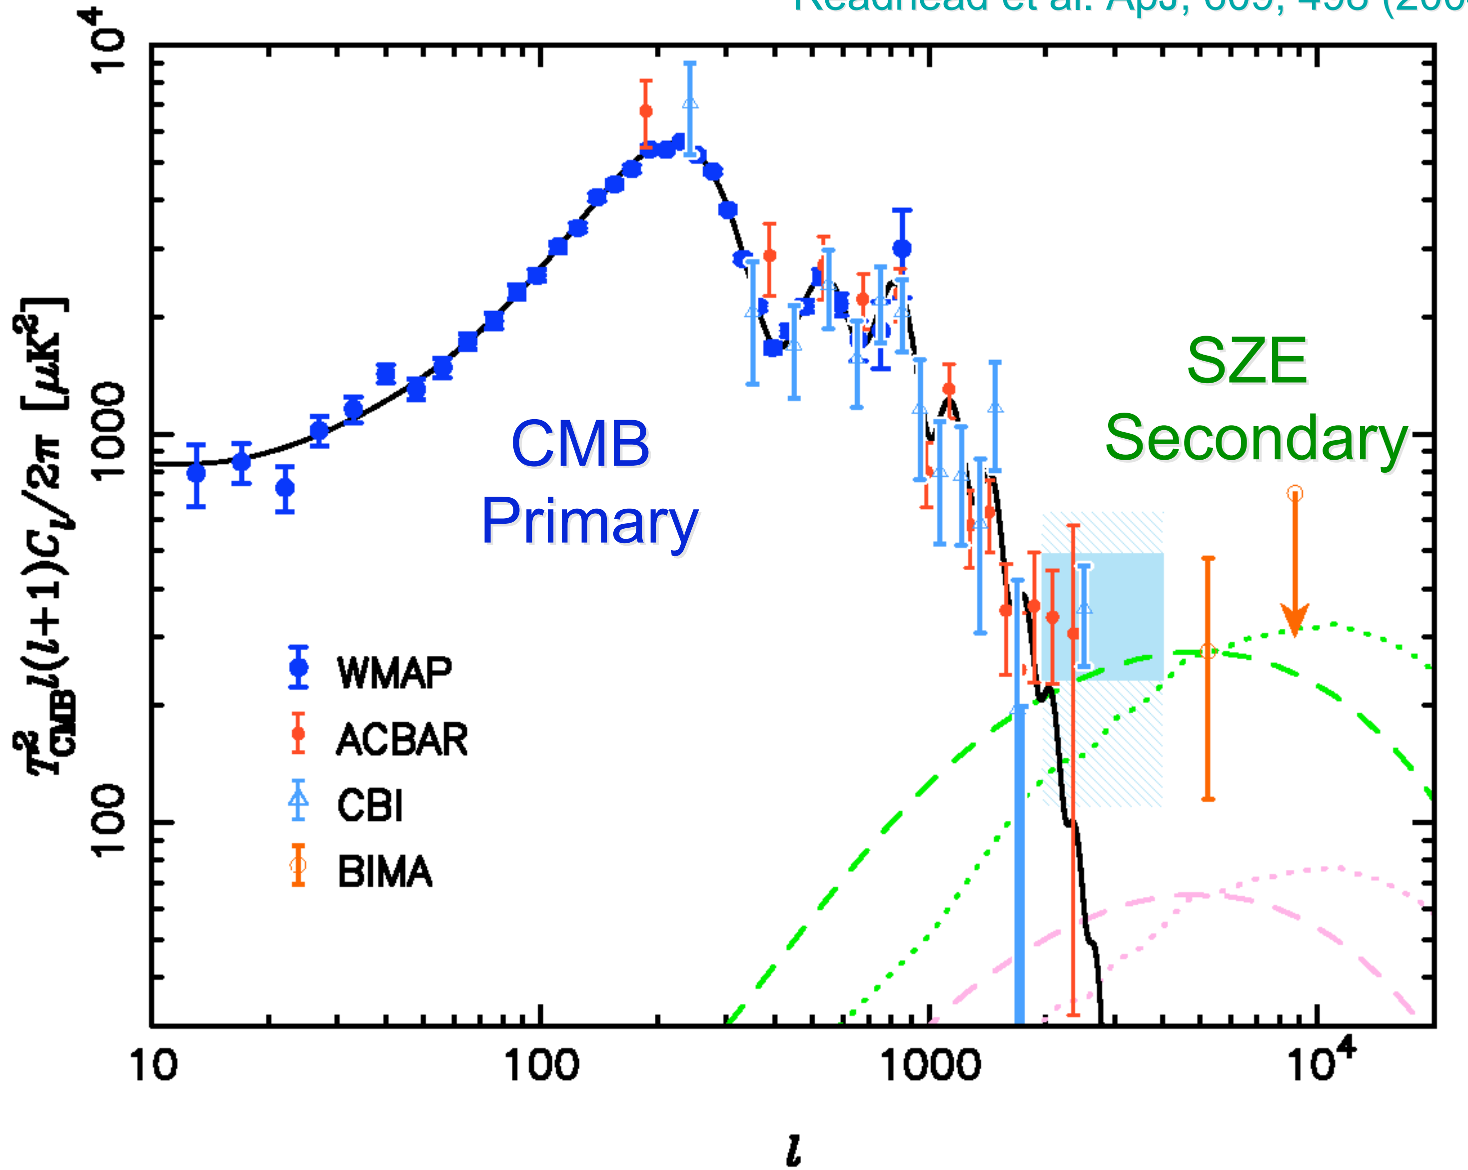

Acbar05: very nice TT, release soon05. parameters & new excess analysis as SZ

Non-Gaussianity

- **Decompose data into uncorrelated S/N eigenmodes for each bin.**
- **Pick out modes expected to have signal**
- **Check distribution for non-Gaussianity**
- **Keep total of 5500 modes TT, 3800 EE – everything consistent with Gaussian**
- **First check of EE Gaussianity**

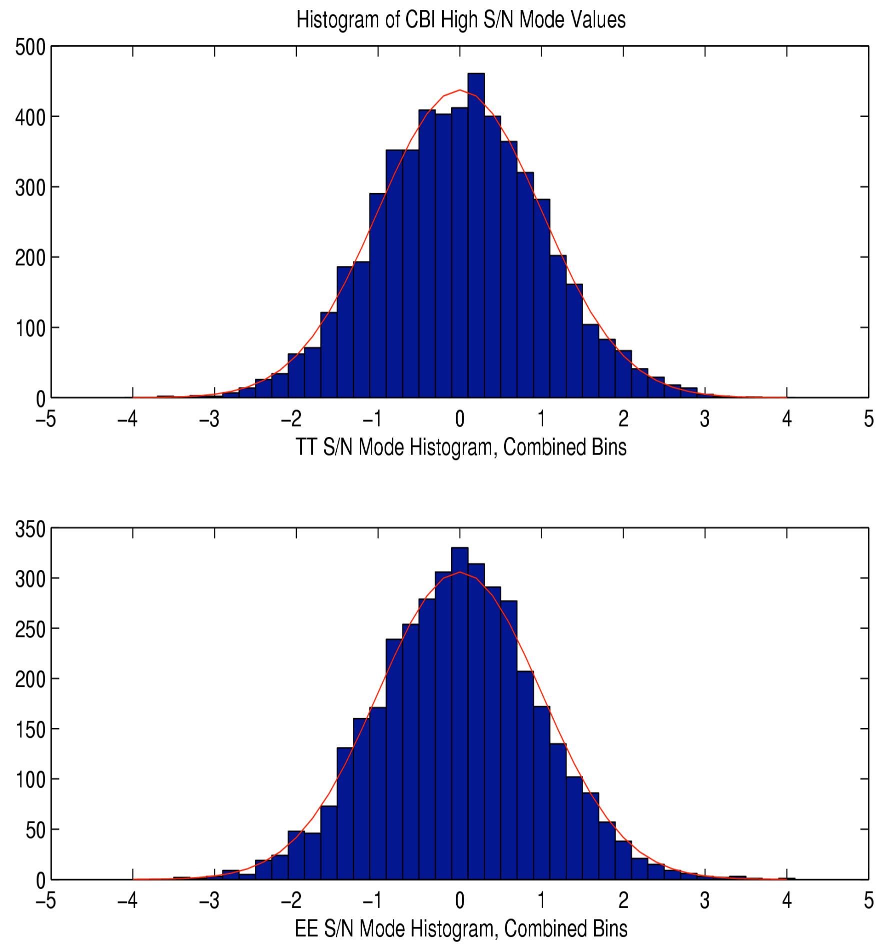

Non-Gaussianity *cont.*

- Check non-Gaussianity in each bin
- Might show l-dependent effect (such as foreground)
- Individual bins consistent with Gaussian.

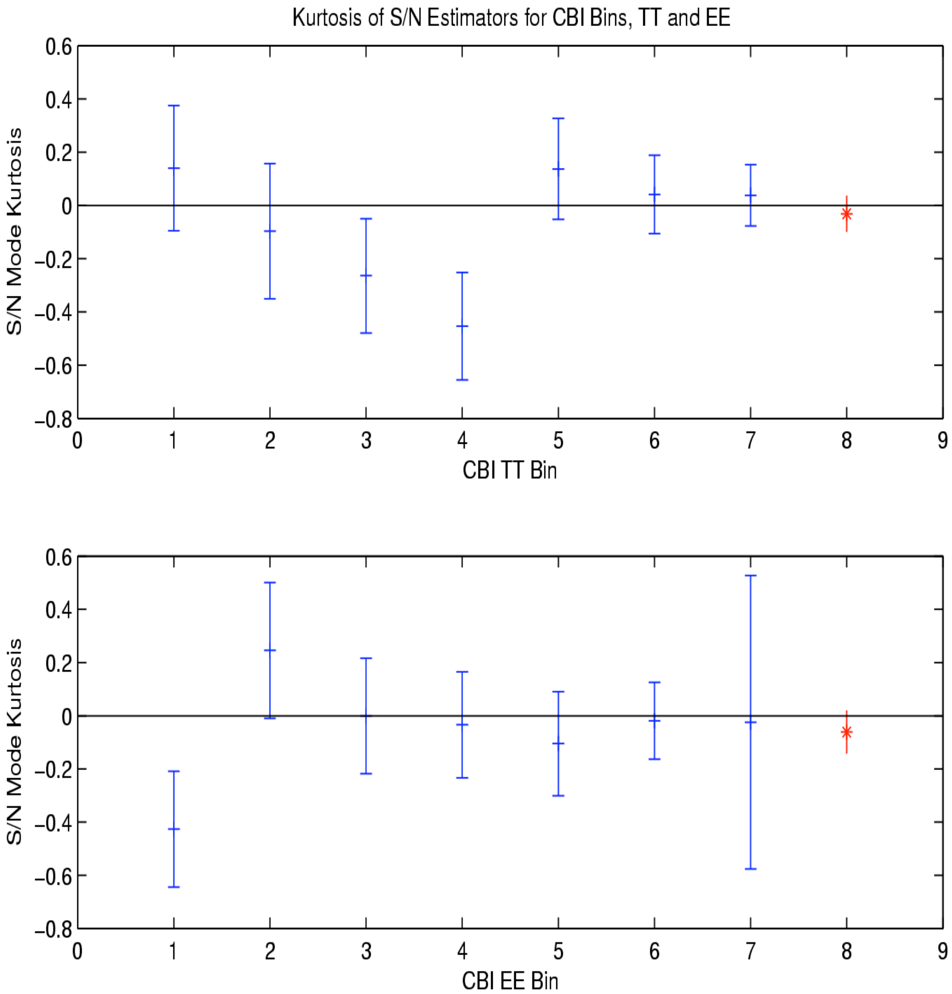

Foregrounds – CBI Radio Sources

Project ~3500 sources in TT,

~550 in polarization

Located in NVSS at 1.4 GHz,

VLA at 8.4 GHz

Predominant on long baselines

No evidence for contribution of sources in polarization – very conservative approach

“masking” out much of sky – need GBT measurements to reduce the number of sources projected

e.g., lead-trail radio sources in CBI mosaic field cf. TT image

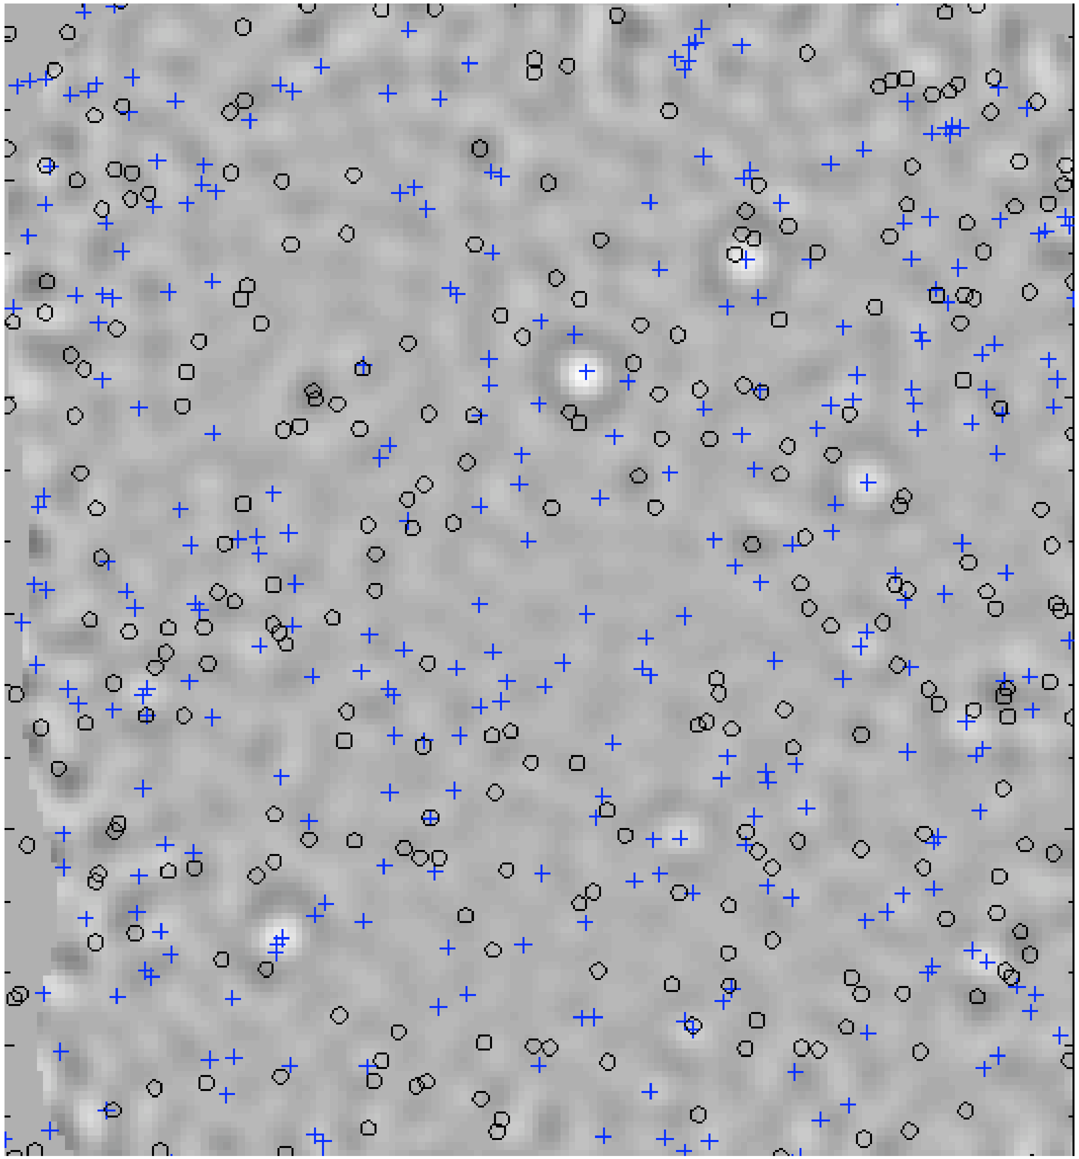

B03: BOOMERANG Jan03 flight

- Caltech
- Rome - La Sapienza
- U of T, CITA
- Case Western
- many others…

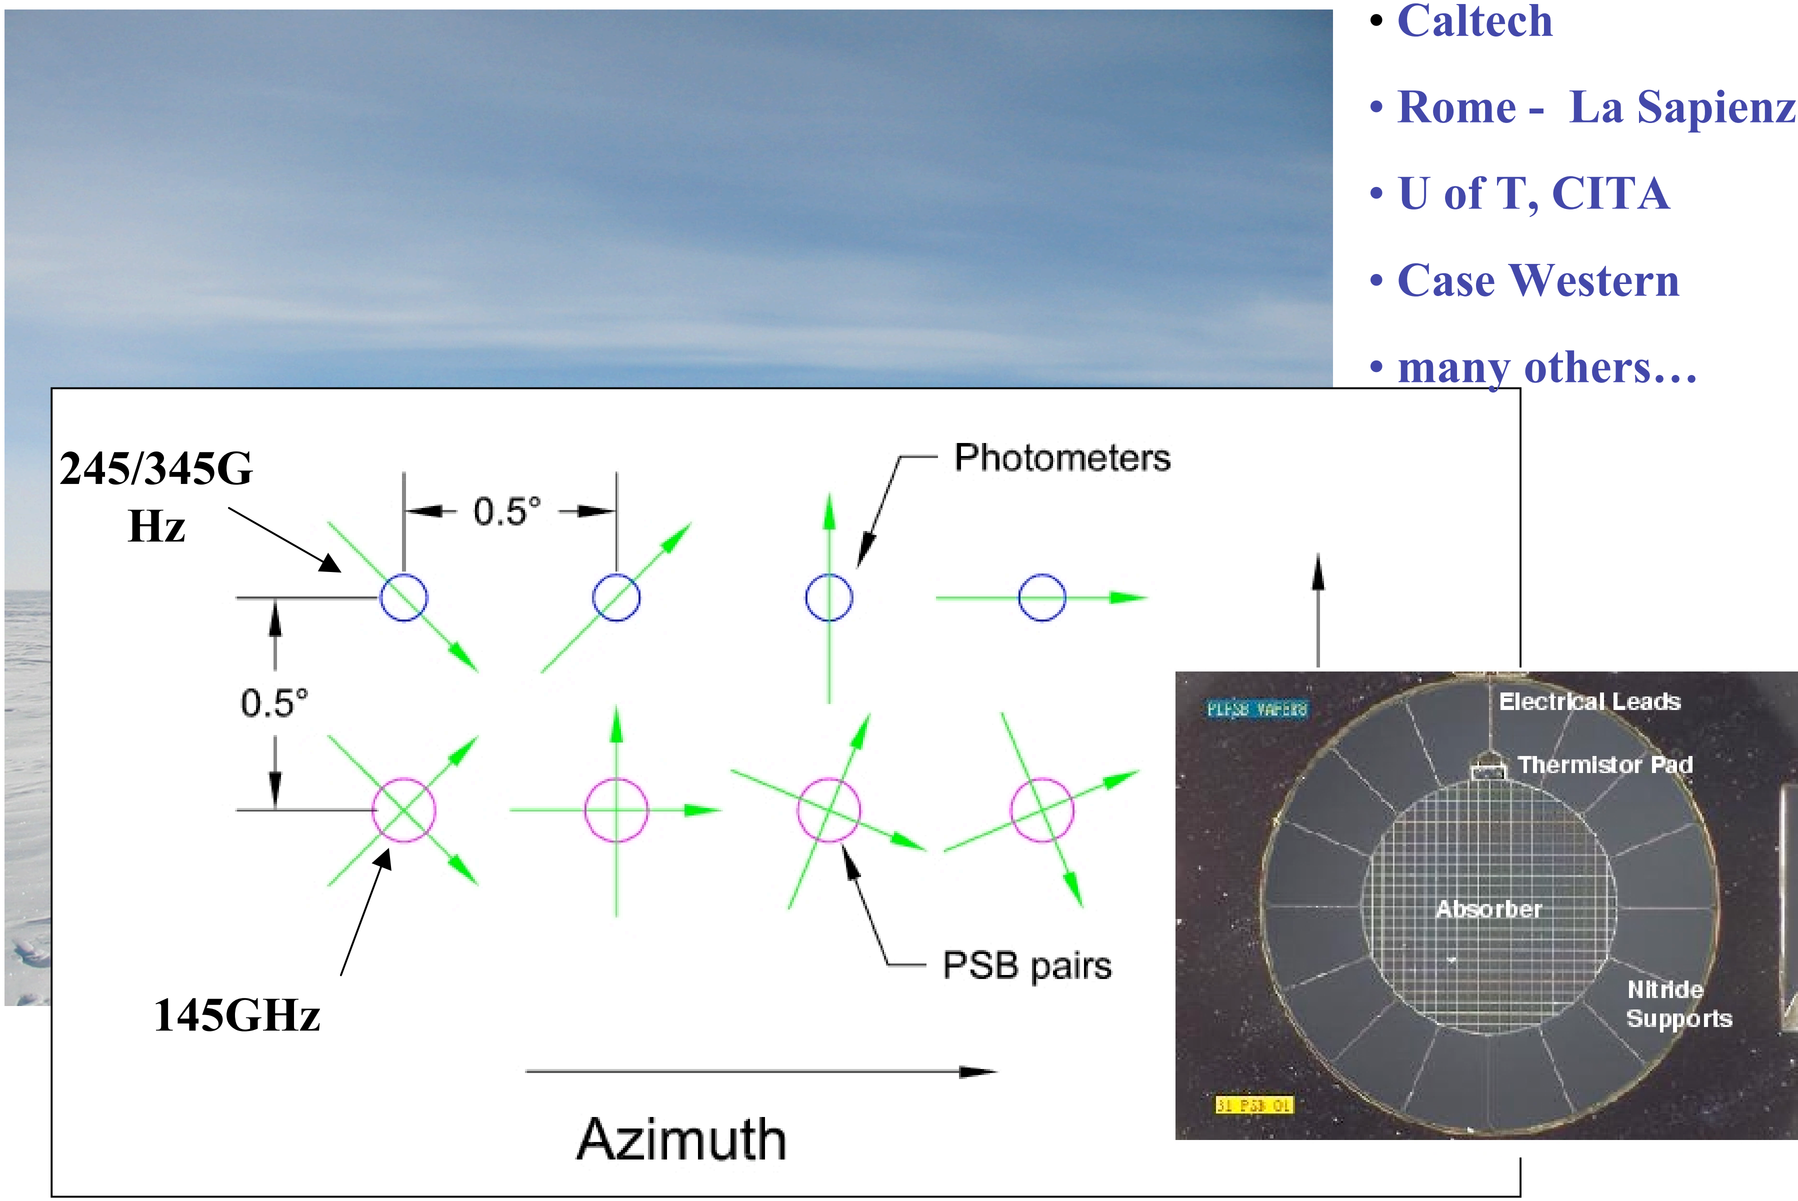

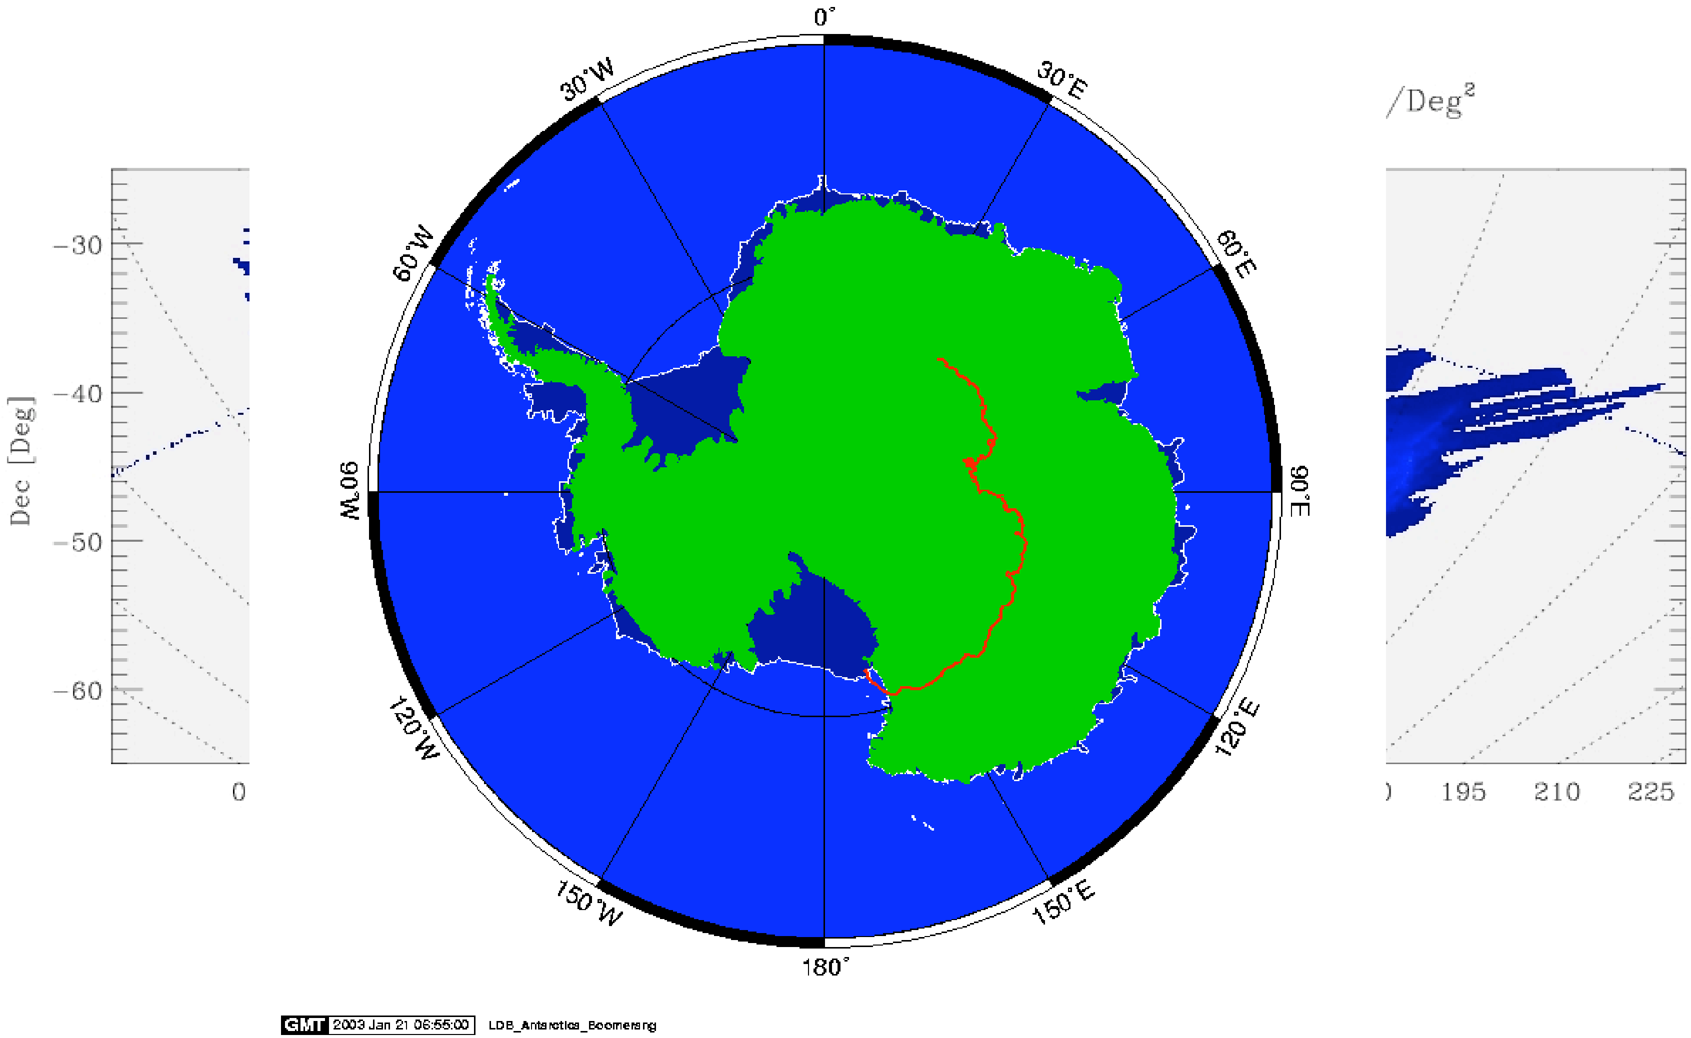

0°
30°W
30°E
60°W
60°E
90°W
90°E
120°W
120°E
150°W
150°E
180°
/Deg²
Dec [Deg]
−30
−40
−50
−60
0
195
210
225
GMT 2003 Jan 21 06:55:00
LDB_Antarctica_Boomerang

B03: PSBs for Polarization

B03: TT very good, TE, EE good detections. Release Jun05

Masi etal 05, Montroy etal 05, Piacentini etal 05, Jones etal 05, MacTavish etal 05

Contaldi etal 05 XFASTER

Boomerang pipeline

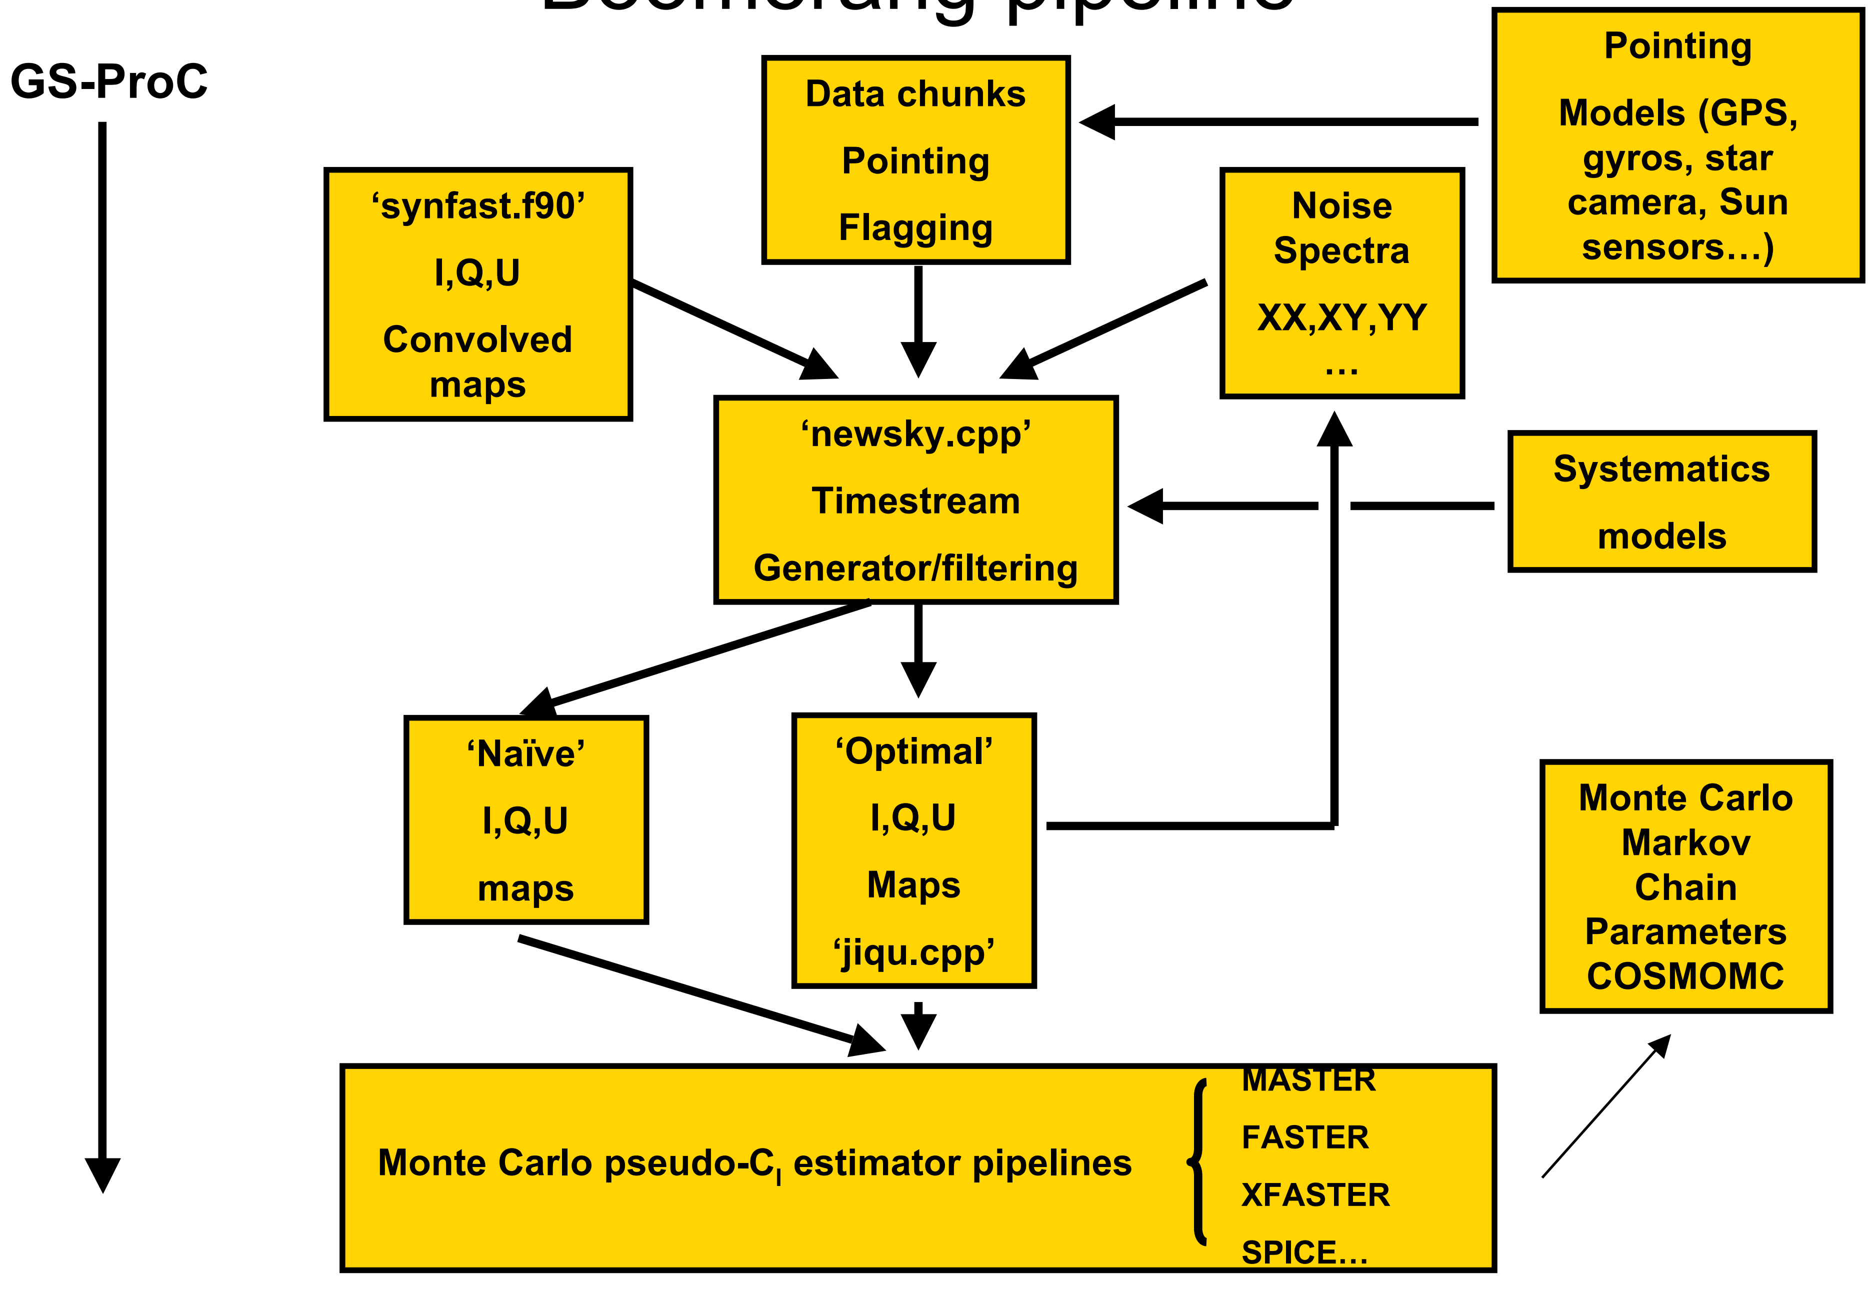

Boomerang : Monte Carlo Methods

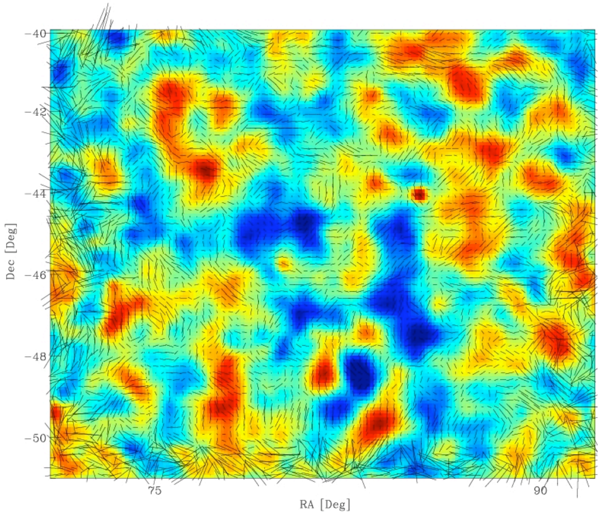

- $N_t \sim 3x10^7$: reduced onto a map pixelized at 3.5' pixels (HEALPIX nside=1024) with $N_{pix} \sim 200{,}000$ x 3. (JIQU : CG linear iteration) ~ 2 hours, single node run

- Storage : single precision pixels ~ 25 Mb/map

- Monte Carlo the full scan strategy to estimate the biases of pseudo-C_l due to noise and filtering. Requires ~ 1000 simulations of the experimental time stream and runs of the iterative map-maker [MASTER, XFASTER, …]

- Scaling : N_{maps} x 2 hrs x 2 ~ 1600 cpu hours, for a standard test run

'Efficiency' prefactor ~ 0.2 % → <u>total development time ~ 800,000 cpu hours</u>

(Jan04-Jan05: PBS recorded 360,000 cpu hrs on Boomerang jobs, more since)

CBI ongoing to Sept'05+

Acbar ongoing to Sept'06+

Bicep

QUaD

Quiet1

Quiet2
(1000 HEMTs)
Chile

APEX
(~400 bolometers)
Chile

SZA
(Interferometer)
California

SCUBA2
(12000 bolometers)
JCMT, Hawaii

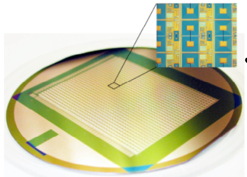

Spider
(1856 bolometer LDB)

ACT
(3000 bolometers)
Chile

Clover

2017
CMBpol

Boom03

2003 2005 2007

2004 2006 2008

WMAP ongoing to 2007+

DASI

CAPMAP

AMI

GBT

Polarbear
(300 bolometers)
California

SPT
(1000 bolometers)
South Pole

Planck
(84 bolometers)
HEMTs L2

ALMA
(Interferometer)
Chile

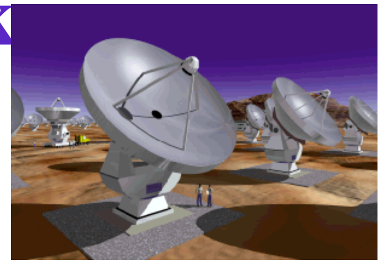

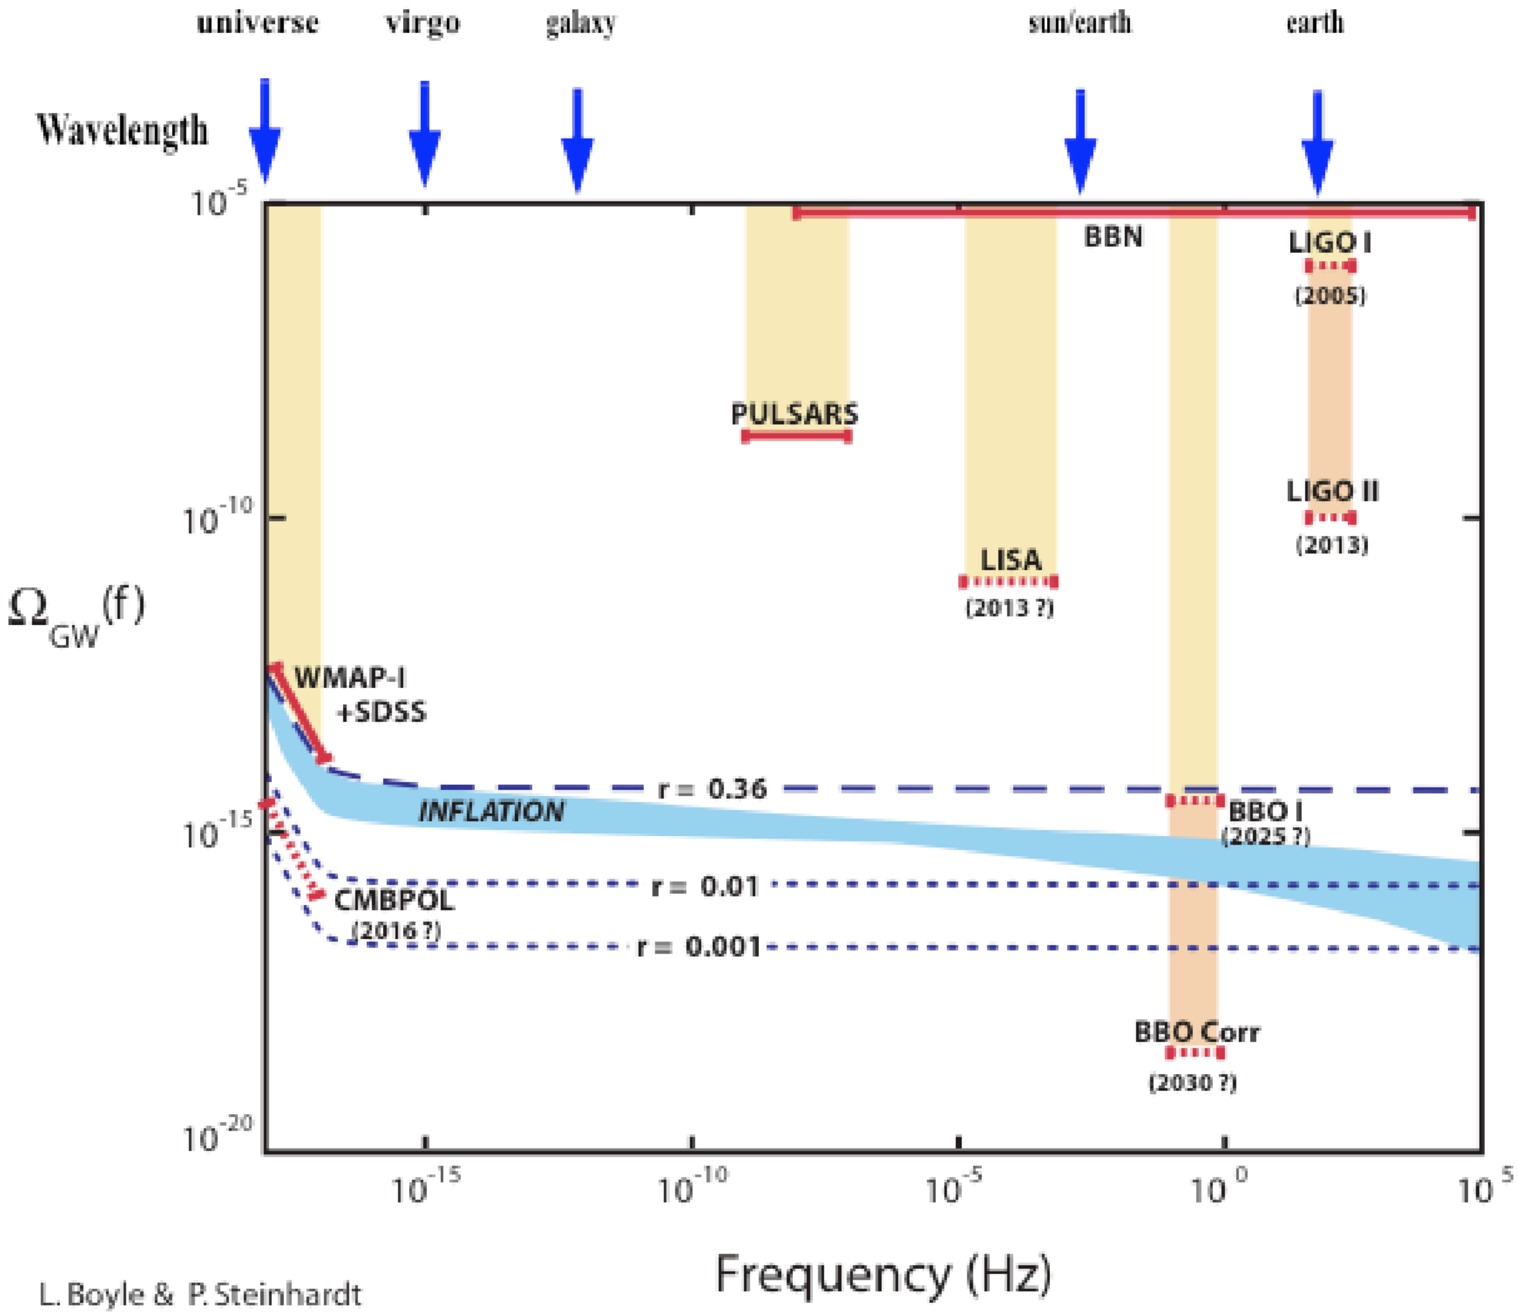

L. Boyle & P. Steinhardt

tensor (gravity wave) power to curvature power, a direct measure of (q+1), q=deceleration parameter during inflation

q may be highly complex (scanning inflation trajectories)

many inflaton potentials give the same curvature power spectrum, but the degeneracy is broken if gravity waves are measured

(q+1) =~ 0 is possible - low scale inflation – upper limit only

Very very difficult to get at this with direct gravity wave detectors – even in our dreams

Response of the CMB photons to the gravitational wave background leads to a unique signature within the CMB at large angular scales of these GW and at a detectable level. Detecting these B-modes is the new "holy grail" of CMB science.

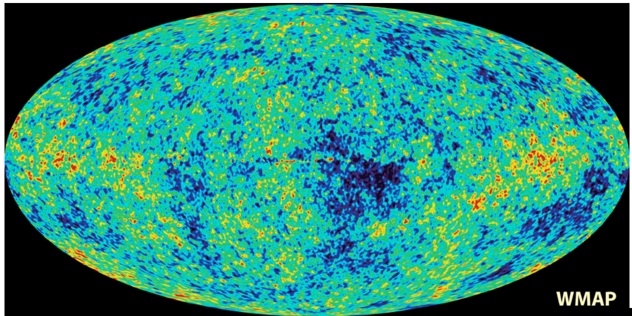

the **Terrain** for **Planck** in the **CMB Landscape**

Synchrotron
Free-Free
Thermal Dust

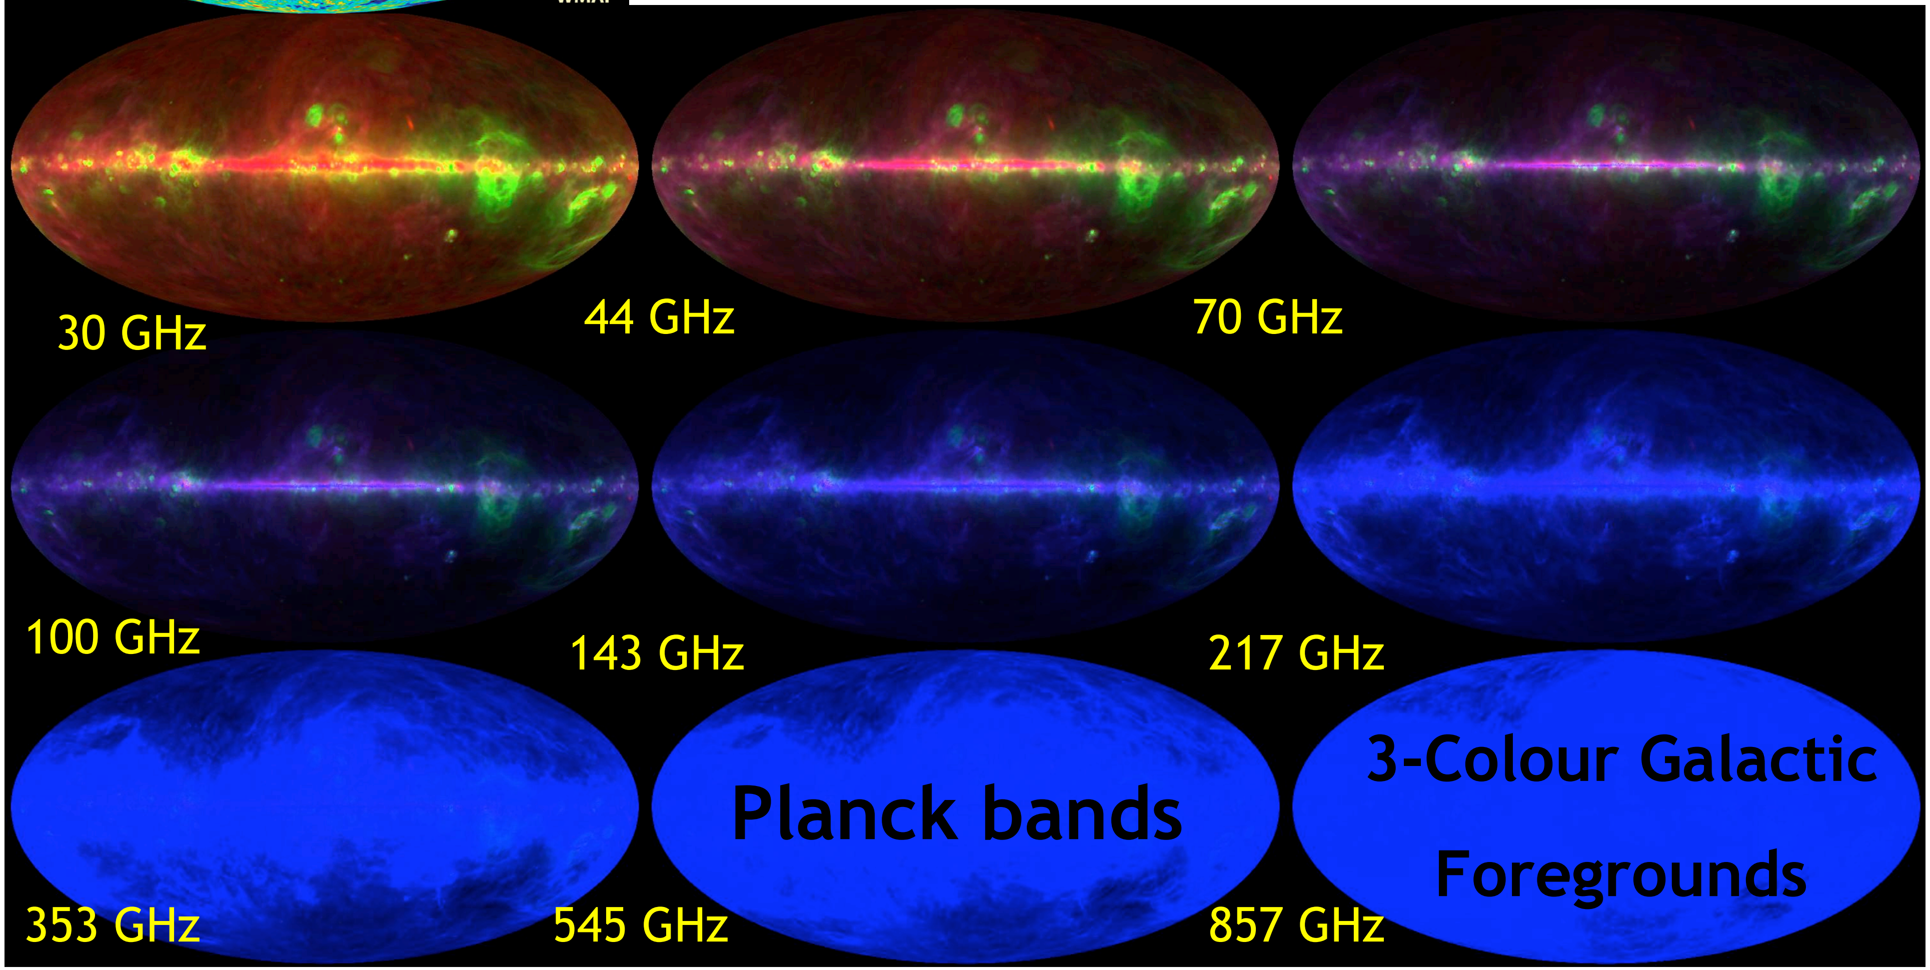

$\Delta T = \delta f/(df_{cmb}/dT)$ in deg K, linear in sqrt(ΔT), 1K threshold

forecast Planck1yr

2007.8+n,

n ~ 2

Planck2.5

is possible

Synchrotron pol'n

< .004 ??

Dust pol'n

< 0.1 ??

Template removals from multi-frequency data

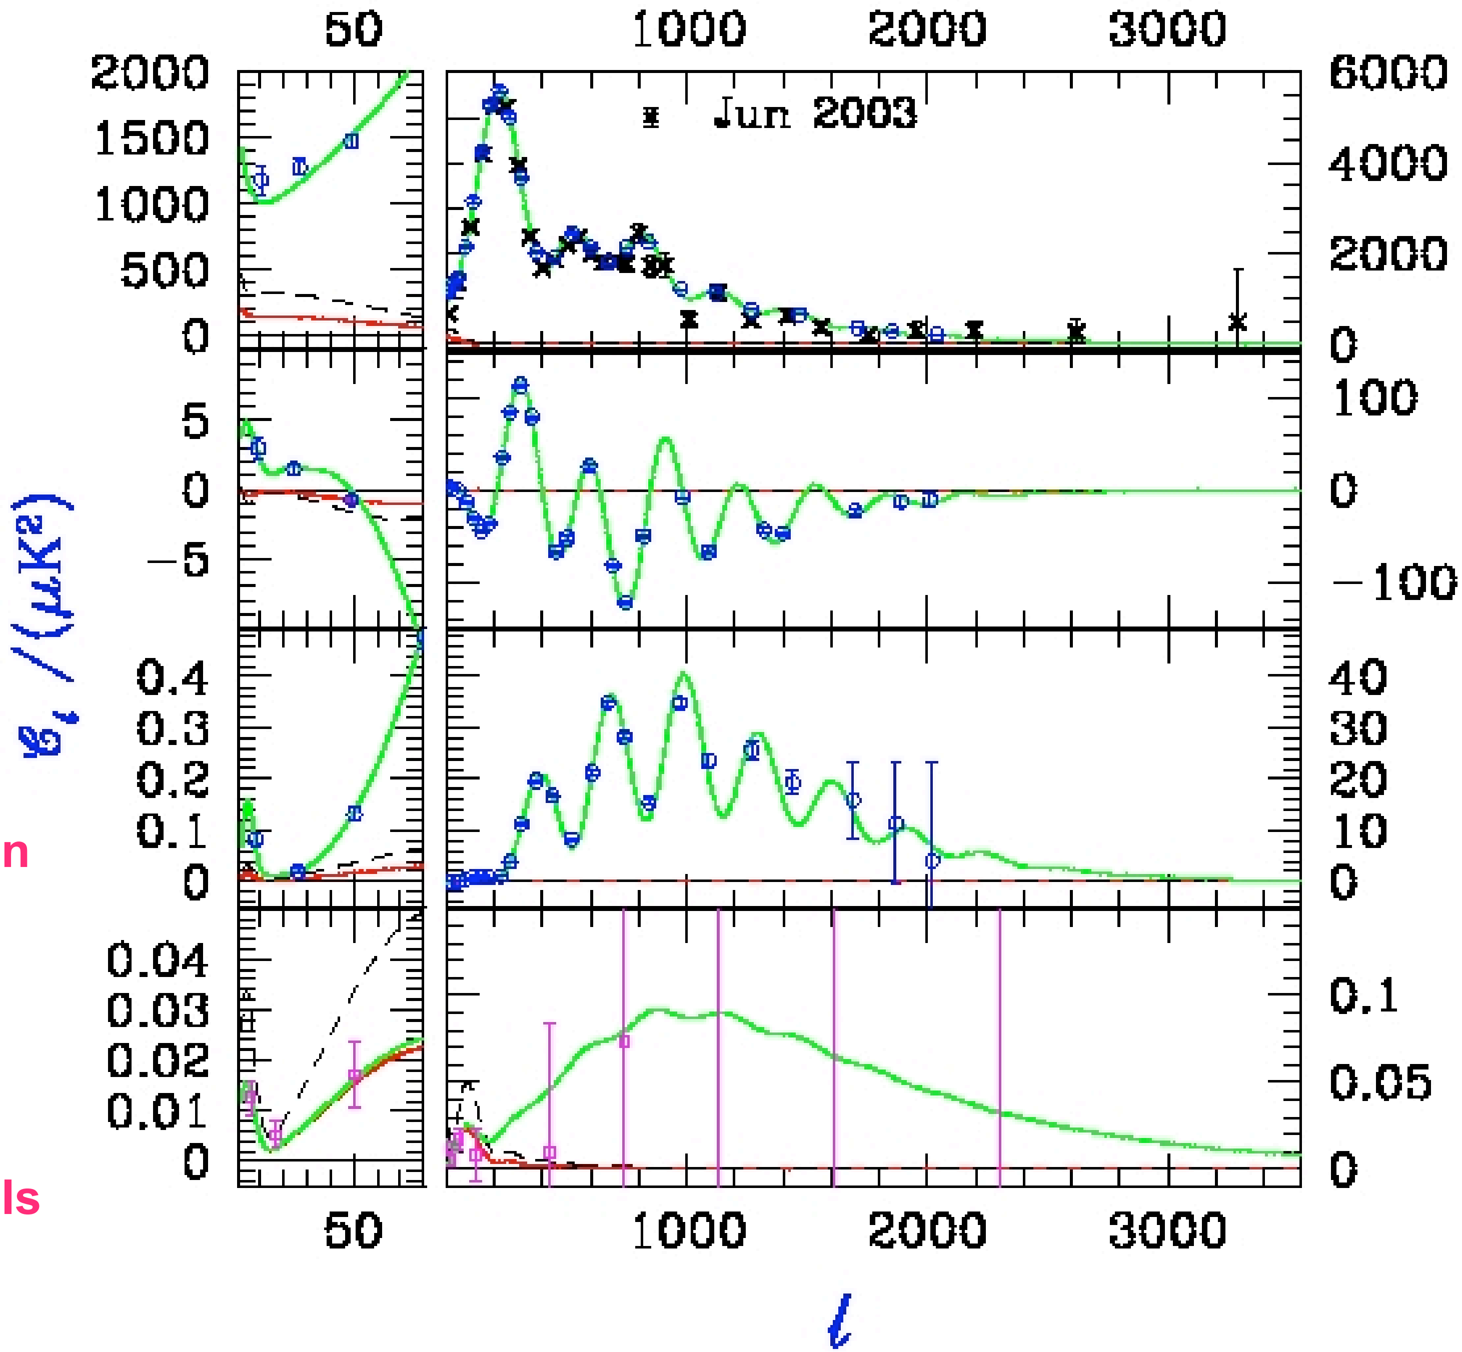

forecast
Planck2.5

2007.8+n,

n ~ 2

100&143

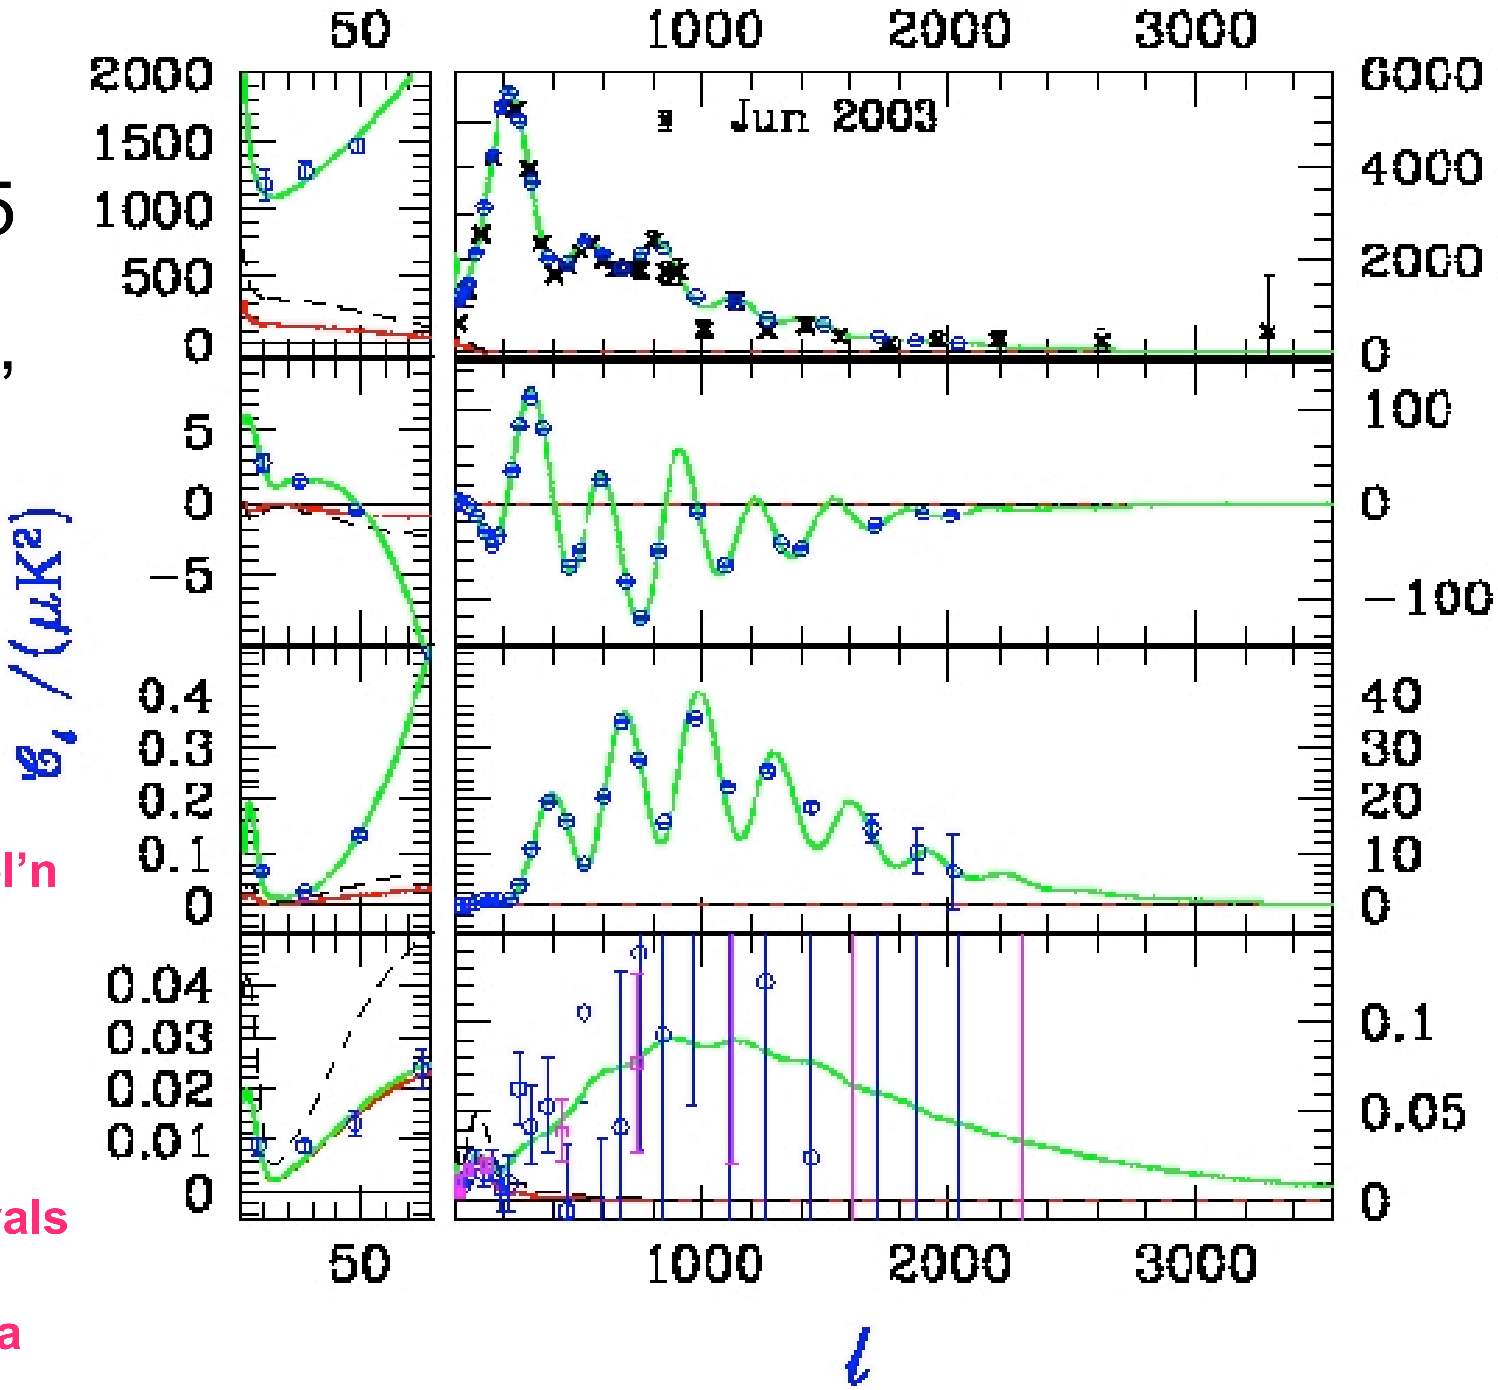

Synchrotron pol'n

< .004 ??

Dust pol'n

< 0.1 ??

Template removals from multi-frequency data

Planck Software Development (huge effort at many centres)

Level 0, Level 1, Level 2, Level 3, Level S, Level 4

Planck HFI & LFI DPCs

Quick Look Analysis. rapid ToD; PITOU rapid visualizer of maps on spheres

MADnes/ MADmap & MAPCUMBA similar algorithms. Outgrowth of original B98 & Maxima optimal separator. cf. JIQU/Newsky - also similar.

PolSpice (cut sky via corr fns) cf. XFASTER

MCMC methods .. Quadratic .. Hybrids (Maxima, Hanson, Gorski, Hivon 03, Efstathiou 04, WMAP)

Stompor, Borrill estimate of cost for one full end-to-end analysis of Planck: 3.3 ExaFlop times 15 or so. Needs of 40+ ExaFlop by 09

SPIDER collaboration (NASA/CSA)

Institute	Responsibilities
Caltech-JPL	detector arrays, optics, receiver assembly/testing
Cardiff University	filters, optics
Case Western Reserve University	cooled _ wave plates and rotating mechanisms, optics
CEA (Grenoble)	He3 refrigerator
Imperial College	data analysis, theory
NIST	SQUID Multiplexers
University of Toronto-CITA	Gondola, tracking, data analysis
University of British Columbia	Readout electronics

SPIDER LDB 09: **Antenna-Coupled bolometer array + rotating half-wave plate**

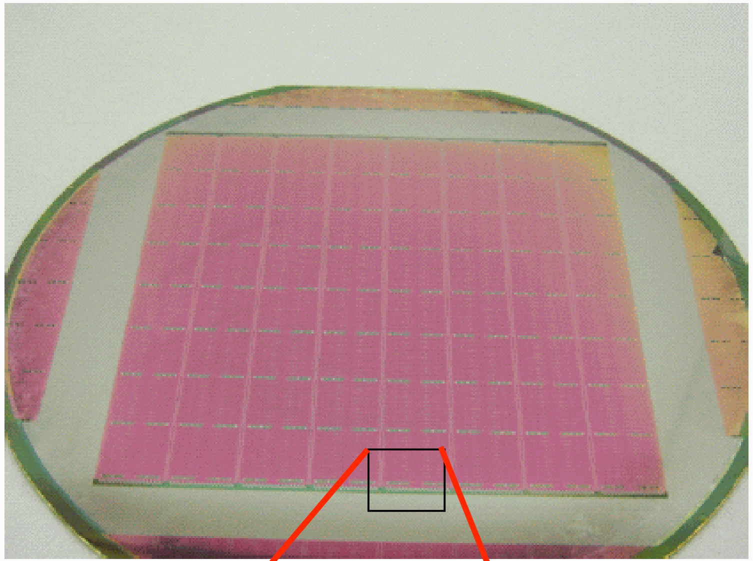

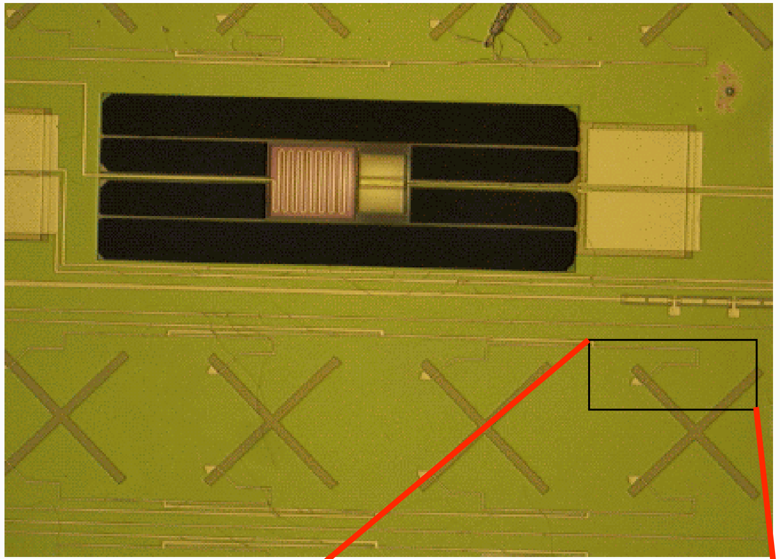

Figure 12: 4-inch-diameter wafer with 8×8 spatial pixels (left) and a closeup on a released TES and four antenna pairs at $50\times$ magnification (right).

Observing Band (GHz)	Bandwidth (GHz)	Beam FWHM (arcmin)	Number of Spatial Pixels	Number of Detectors	Single-Detector Sensitivity (μK_{CMB} $s^{1/2}$)	Instrument Sensitivity (μK_{CMB} $s^{1/2}$)
40	10	145	32	64	130	16.3
84 92 100	33	69 63 58	128 128 128	256 256 256	60	3.8
145 220	32 40	40 26	256 256	512 512	80 150	3.5 6.6

Table 1: Observing bands, pixel and detector counts, and single-detector and instrument sensitivities. The latter is obtained by dividing the single-detector sensitivity by $\sqrt{N_{det}}$. A total of 1856 detectors are distributed between the six telescopes.

forecast

Spider 10d

95&150

~ 50% sky
@ 40' res

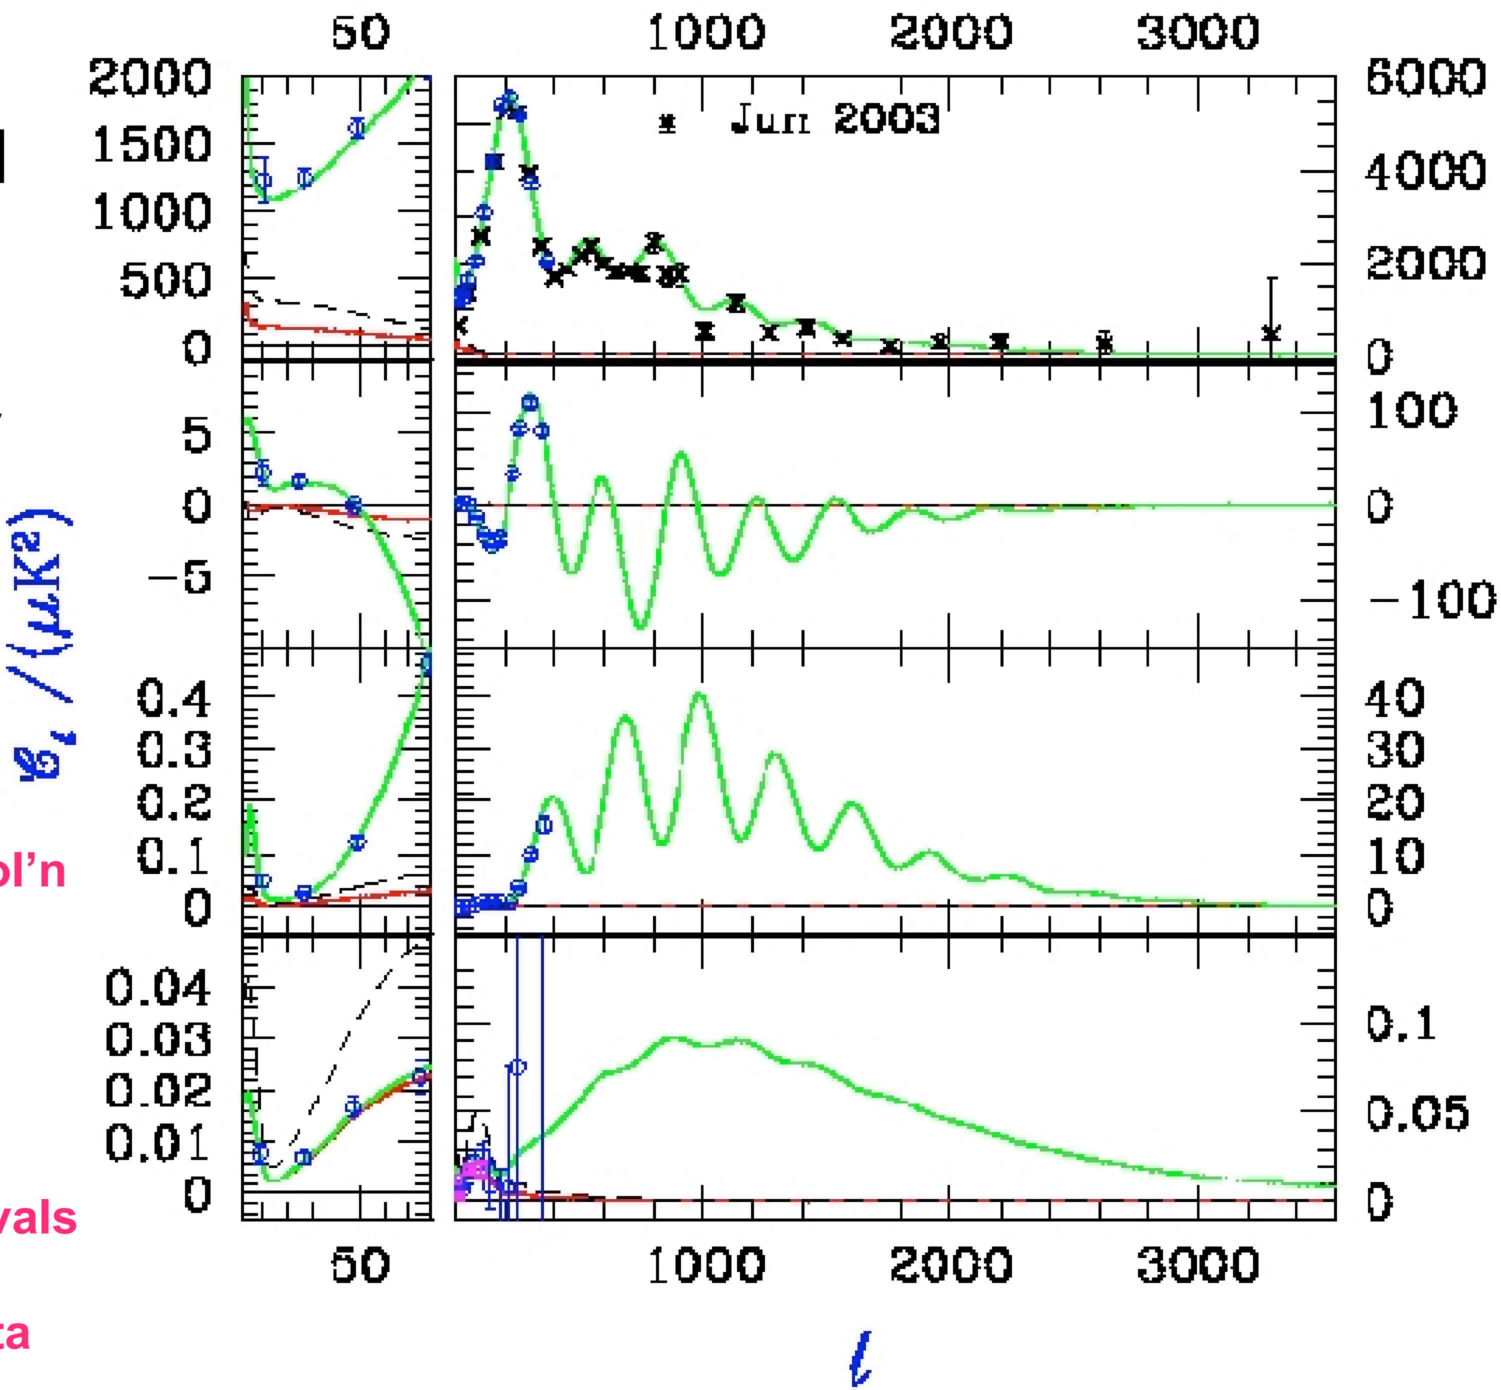

Synchrotron pol'n

< .004 ??

Dust pol'n

< 0.1 ??

Template removals from multi-frequency data

forecast
Planck2.5

100&143

Spider10d

95&150

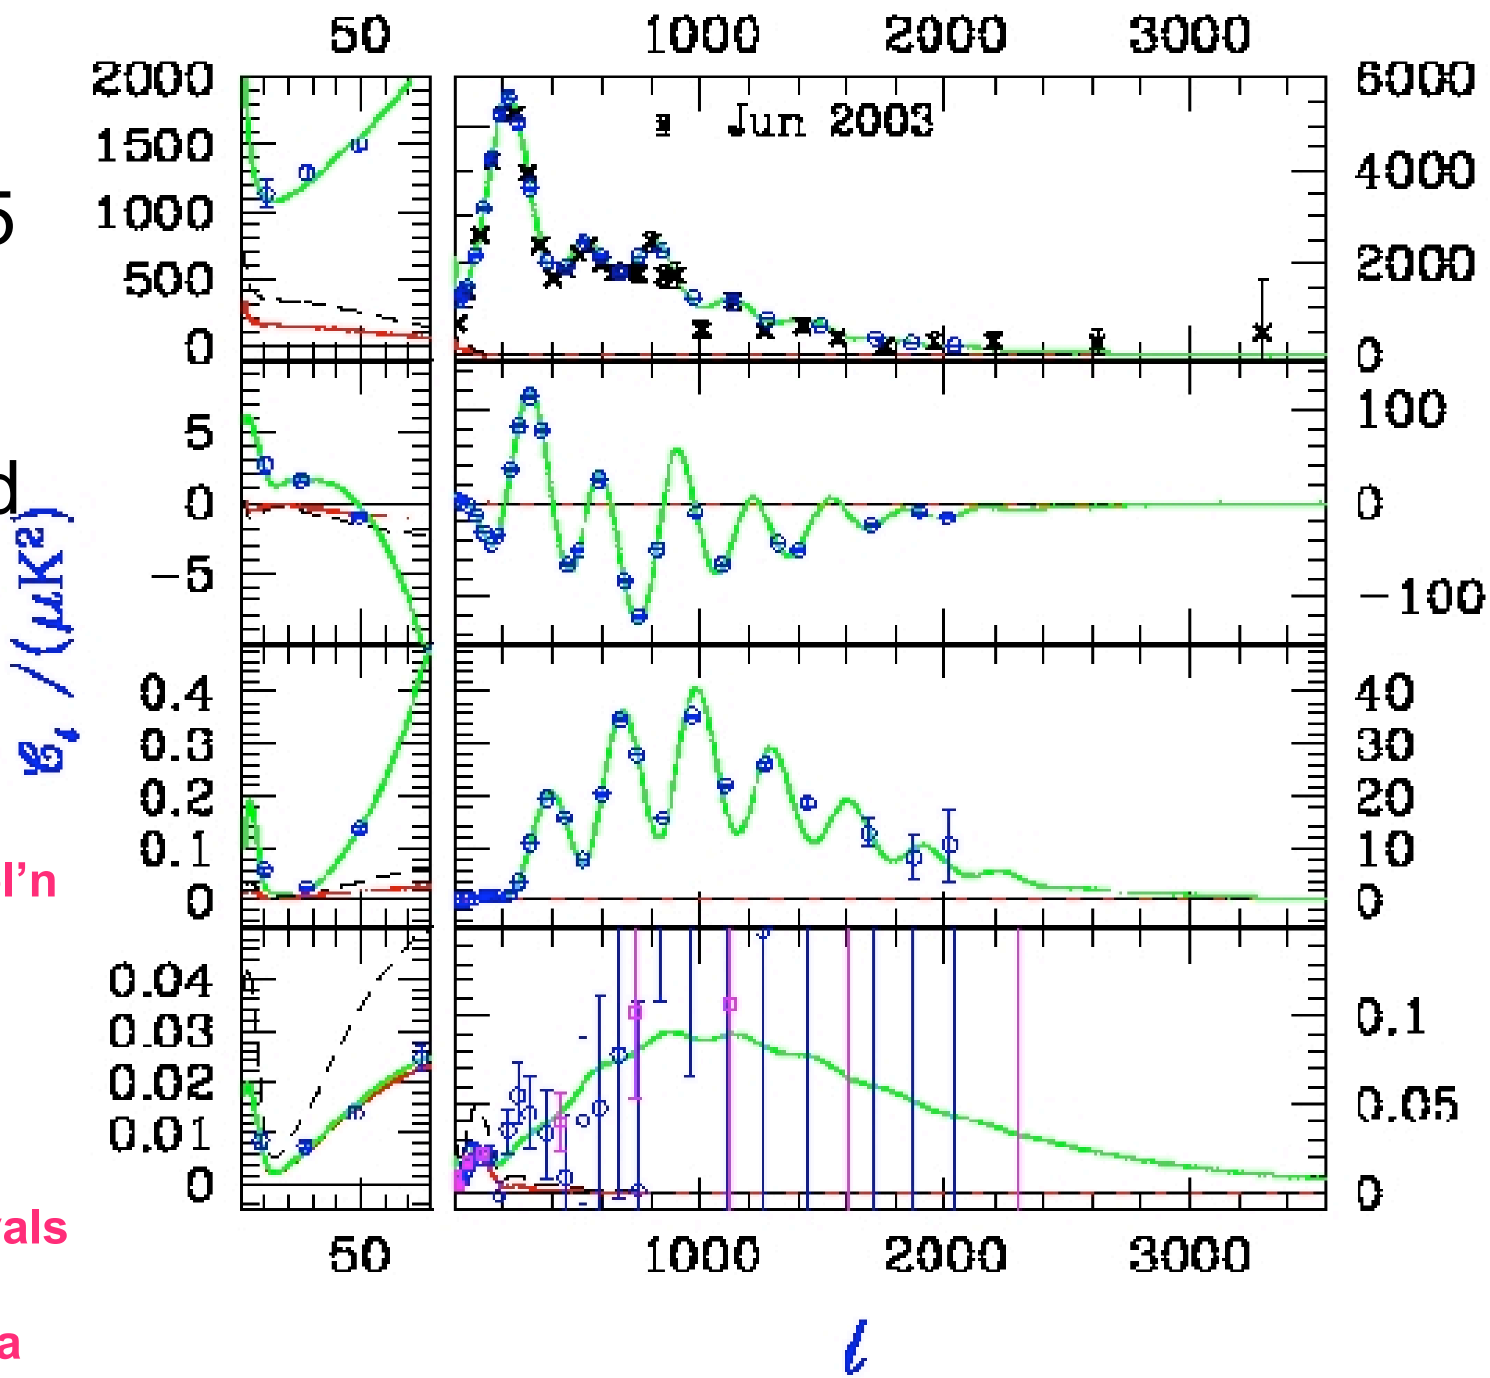

Synchrotron pol'n

< .004 ??

Dust pol'n

< 0.1 ??

Template removals from multi-frequency data

SUMMARY of CMB Computing Challenges, with Current Algorithms

McKenzie 1.5 Tflops - 201 on the Nov04 list (Current top 3 are 71, 52 and 36 Tflops)

20-25% of all McKenzie cycles have gone to CMB projects

CITA-HEP "plan" 8100 dual-CPU nodes. with 2.8GHz CPUs 90 Tflops.

Planck USPDA estimate of 14 Tflop/s (41 ExaFlop) dedicated in 2009. ~ 3100 CPUs/yr

Spider (balloon borne CMB polarization on large scales) will need equivalent of 2700 CPUs/result/year (assuming small efficiency factor); ACT 6200 CPUs/ result/yr

"efficiency" fudge factor, 10 as a minimum cf. B03 ~ 200, CBI ~ 50

i.e., dedicated access to ~10000 CPUs needed for Planck, Spider, ACT etal. analyses