TO BE SUBMITTED TO A JOURNAL TBD
Preprint typeset using LaTeX style emulateapj v. 04/17/13

BICEP2 I: DETECTION OF *B*-mode POLARIZATION AT DEGREE ANGULAR SCALES

BICEP2 COLLABORATION - P. A. R. ADE[1], R. W. AIKIN[2], D. BARKATS[3], S. J. BENTON[4], C. A. BISCHOFF[5], J. J. BOCK[2,6], J. A. BREVIK[2], I. BUDER[5], E. BULLOCK[7], C. D. DOWELL[6], L. DUBAND[8], J. P. FILIPPINI[2], S. FLIESCHER[9], S. R. GOLWALA[2], M. HALPERN[10], M. HASSELFIELD[10], S. R. HILDEBRANDT[2,6], G. C. HILTON[11], V. V. HRISTOV[2], K. D. IRWIN[12,13,11], K. S. KARKARE[5], J. P. KAUFMAN[14], B. G. KEATING[14], S. A. KERNASOVSKIY[12], J. M. KOVAC[5,16], C. L. KUO[12,13], E. M. LEITCH[15], M. LUEKER[2], P. MASON[2], C. B. NETTERFIELD[4], H. T. NGUYEN[6], R. O'BRIENT[6], R. W. OGBURN IV[12,13], A. ORLANDO[14], C. PRYKE[9,7,16], C. D. REINTSEMA[11], S. RICHTER[5], R. SCHWARZ[9], C. D. SHEEHY[9,15], Z. K. STANISZEWSKI[2,6], R. V. SUDIWALA[1], G. P. TEPLY[2], J. E. TOLAN[12], A. D. TURNER[6], A. G. VIEREGG[5,15], C. L. WONG[5], AND K. W. YOON[12,13]

to be submitted to a journal TBD

ABSTRACT

We report results from the BICEP2 experiment, a Cosmic Microwave Background (CMB) polarimeter specifically designed to search for the signal of inflationary gravitational waves in the *B*-mode power spectrum around $\ell \sim 80$. The telescope comprised a 26 cm aperture all-cold refracting optical system equipped with a focal plane of 512 antenna coupled transition edge sensor (TES) 150 GHz bolometers each with temperature sensitivity of $\approx 300\ \mu K_{CMB}\sqrt{s}$. BICEP2 observed from the South Pole for three seasons from 2010 to 2012. A low-foreground region of sky with an effective area of 380 square degrees was observed to a depth of 87 nK-degrees in Stokes Q and U. In this paper we describe the observations, data reduction, maps, simulations and results. We find an excess of *B*-mode power over the base lensed-ΛCDM expectation in the range $30 < \ell < 150$, inconsistent with the null hypothesis at a significance of $> 5\sigma$. Through jackknife tests and simulations based on detailed calibration measurements we show that systematic contamination is much smaller than the observed excess. We also estimate potential foreground signals and find that available models predict these to be considerably smaller than the observed signal. These foreground models possess no significant cross-correlation with our maps. Additionally, cross-correlating BICEP2 against 100 GHz maps from the BICEP1 experiment, the excess signal is confirmed with 3σ significance and its spectral index is found to be consistent with that of the CMB, disfavoring synchrotron or dust at 2.3σ and 2.2σ, respectively. The observed *B*-mode power spectrum is well-fit by a lensed-ΛCDM + tensor theoretical model with tensor/scalar ratio $r = 0.20^{+0.07}_{-0.05}$, with $r = 0$ disfavored at 7.0σ. Subtracting the best available estimate for foreground dust modifies the likelihood slightly so that $r = 0$ is disfavored at 5.9σ.

Subject headings: cosmic background radiation — cosmology: observations — gravitational waves — inflation — polarization

1. INTRODUCTION

The discovery of the Cosmic Microwave Background (CMB) by Penzias & Wilson (1965) confirmed the hot big bang paradigm and established the CMB as a central tool for the study of cosmology. In recent years, observations of its temperature anisotropies have helped establish and refine the "standard" cosmological model now known as ΛCDM, under which our universe is understood to be spatially flat, dominated by cold dark matter, and with a cosmological constant (Λ) driving accelerated expansion at late times. CMB temperature measurements have now reached remarkable precision over angular scales ranging from the whole sky to arcminute resolution, producing results in striking concordance with predictions of ΛCDM and constraining its key parameters to sub-percent precision (e.g. Bennett et al. 2013; Hinshaw et al. 2013; Story et al. 2013; Hou et al. 2014; Sievers et al. 2013; Das et al. 2013; Planck Collaboration XV 2013; Planck Collaboration XVI 2013).

Inflationary cosmology extends the standard model by postulating an early period of nearly exponential expansion which sets the initial conditions for the subsequent hot big bang. It was proposed and developed in the early 1980s to resolve mysteries for which the standard model offered no solution, including the flatness, horizon, smoothness, entropy, and monopole problems (Brout et al. 1978; Starobinsky 1980; Kazanas 1980; Sato 1981; Guth 1981; Linde 1982, 1983; Albrecht & Steinhardt 1982; see Planck Collaboration XXII 2013 for a review). Inflation also explains the universe's

[1] School of Physics and Astronomy, Cardiff University, Cardiff, CF24 3AA, UK
[2] Department of Physics, California Institute of Technology, Pasadena, CA 91125, USA
[3] Joint ALMA Observatory, ESO, Santiago, Chile
[4] Department of Physics, University of Toronto, Toronto, ON, Canada
[5] Harvard-Smithsonian Center for Astrophysics, 60 Garden Street MS 42, Cambridge, MA 02138, USA
[6] Jet Propulsion Laboratory, Pasadena, CA 91109, USA
[7] Minnesota Institute for Astrophysics, University of Minnesota, Minneapolis, MN 55455, USA
[8] SBT, Commissariat à l'Energie Atomique, Grenoble, France
[9] Department of Physics, University of Minnesota, Minneapolis, MN 55455, USA
[10] Department of Physics and Astronomy, University of British Columbia, Vancouver, BC, Canada
[11] National Institute of Standards and Technology, Boulder, CO 80305, USA
[12] Department of Physics, Stanford University, Stanford, CA 94305, USA
[13] Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, 2575 Sand Hill Rd, Menlo Park, CA 94025, USA
[14] Department of Physics, University of California at San Diego, La Jolla, CA 92093, USA
[15] University of Chicago, Chicago, IL 60637, USA
[16] Corresponding authors: pryke@physics.umn.edu and jmkovac@cfa.harvard.edu

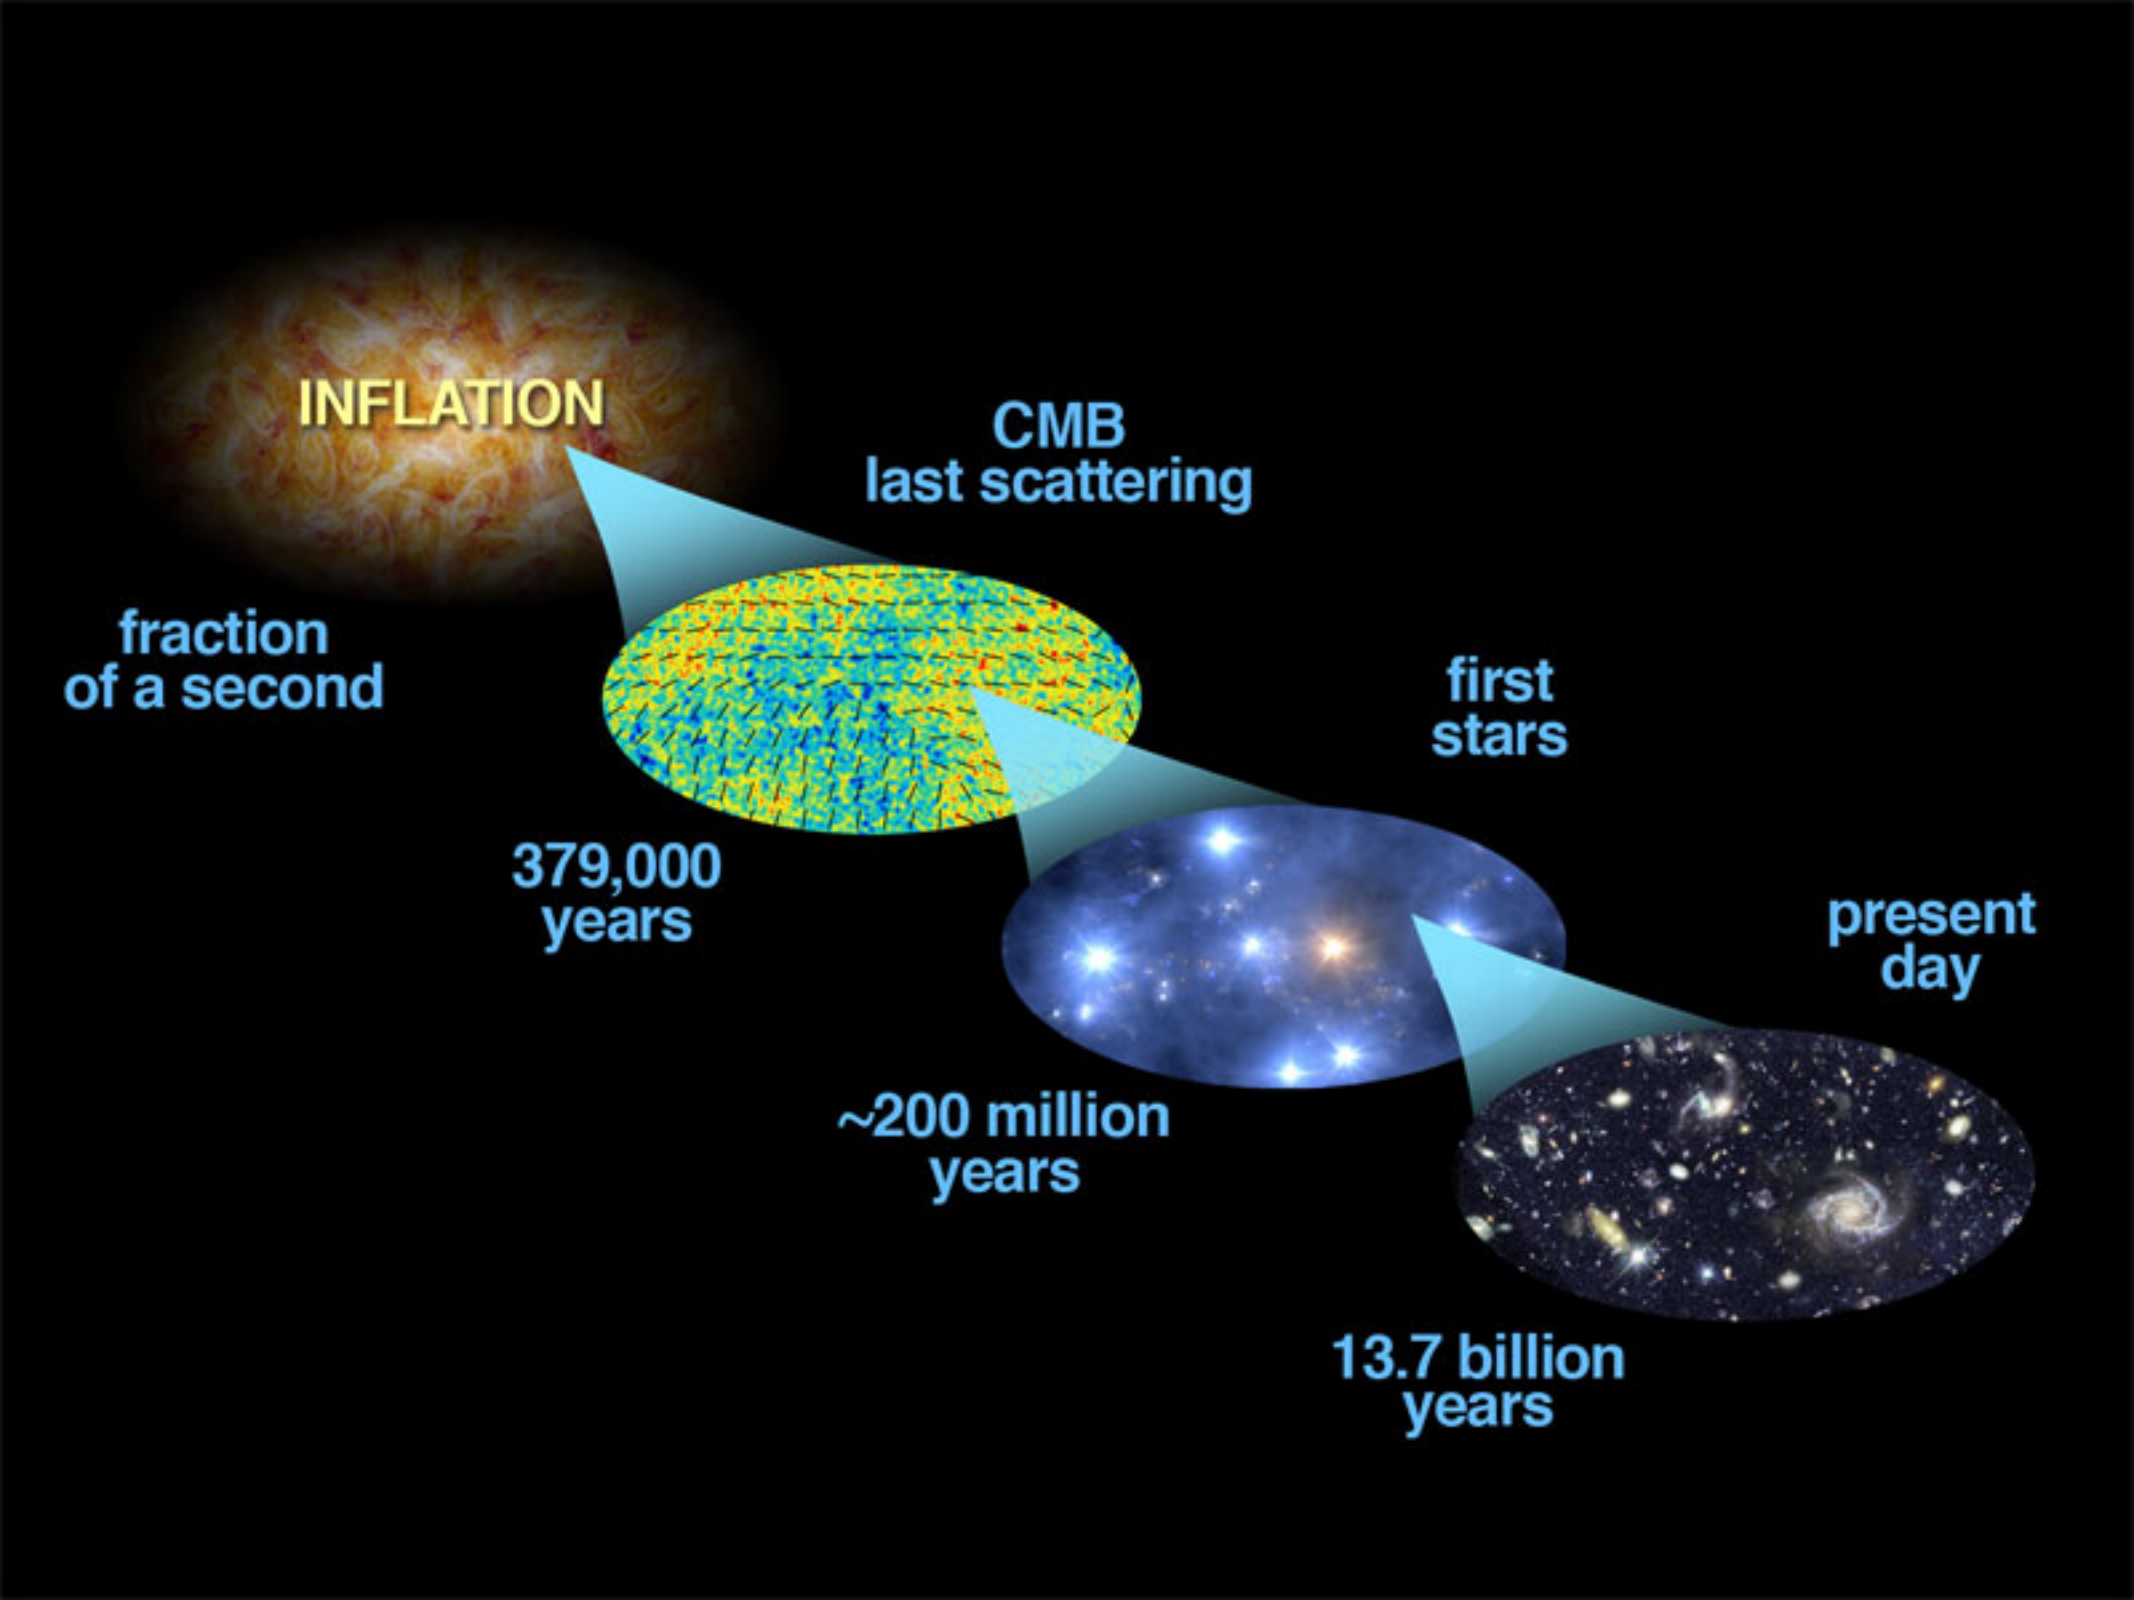
INFLATION
CMB
last scattering
fraction
of a second
379,000
years
first
stars
~200 million
years
present
day
13.7 billion
years

Properties of "standard “ inflation measurable with the CMB

- Spectrum of spatial temperature fluctuations
  - Intrinsic quantum fluctuations both scalar and tensor
  - Scale invariant spectrum k^n, n ~ 1
- Gaussian probability distribution of fluctuations
- Tensor gravitational waves generated due to acceleration during inflation, related to energy scale of inflation

$$r = \frac{\Delta T_{tensor}}{\Delta T_{scalar}} \leq 0.3$$

  - cause quadrupolar density and temperature fluctuations in the primeval plasmas
  - uniquely measurable in the patterns of the CMB polarization due to Thomson scattering of the CMB by electrons (B modes)

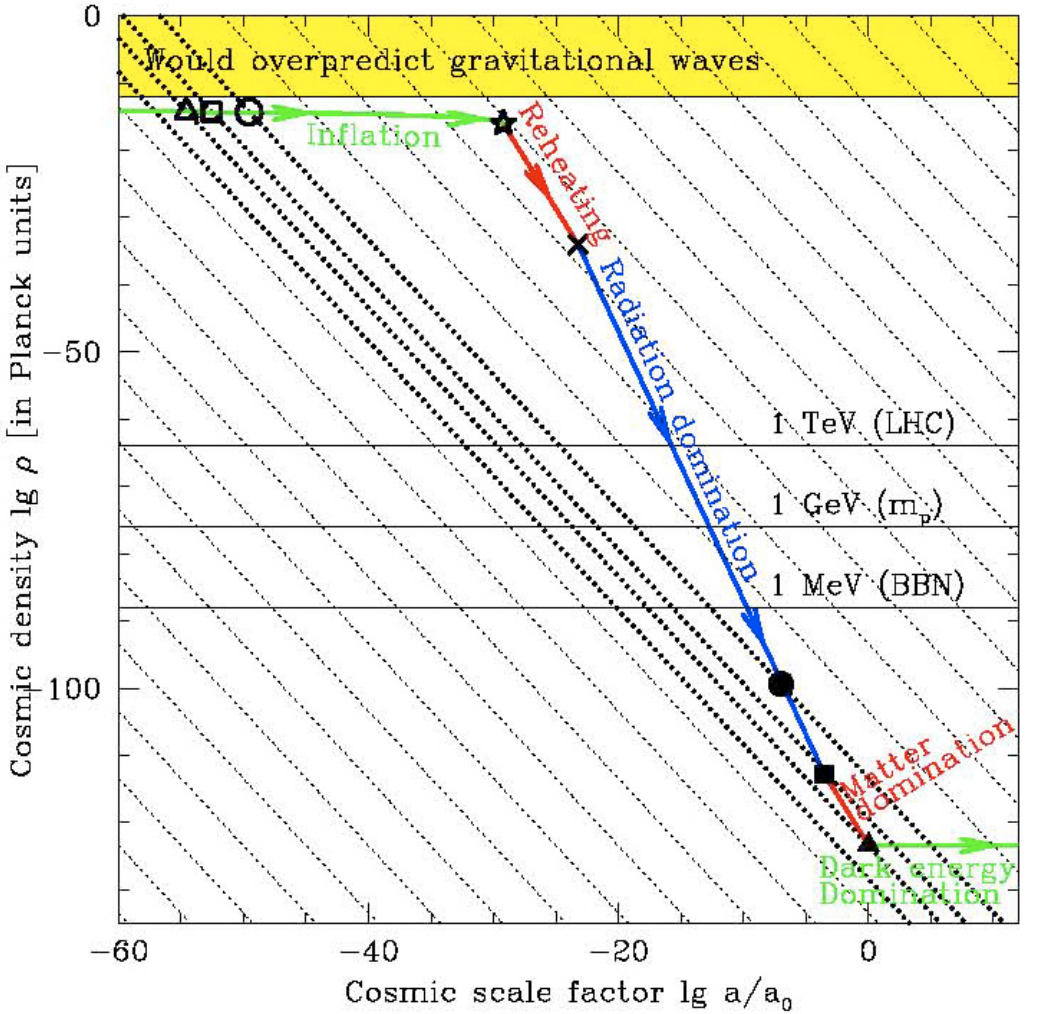

Figure 1: The cosmic mean density (solid curve) is diluted as the Universe expands. Inflation is a period when there is almost no dilution, causing the expansion to accelerate. The result is that the curve decreases slower than the dotted diagonal lines of slope −2. The two triangles lie on the same diagonal, which means that quantum fluctuations generated during inflation at the open triangle have been stretched into the horizon-scale fluctuations that we observe today at the filled triangle in the CMB. Detecting inflationary gravitational waves with CMB polarization would directly measure the shape of this cosmic density curve in the upper left corner of the plot, just as the proposed Joint Dark Energy mission would directly measure the same curve in the lower right corner.

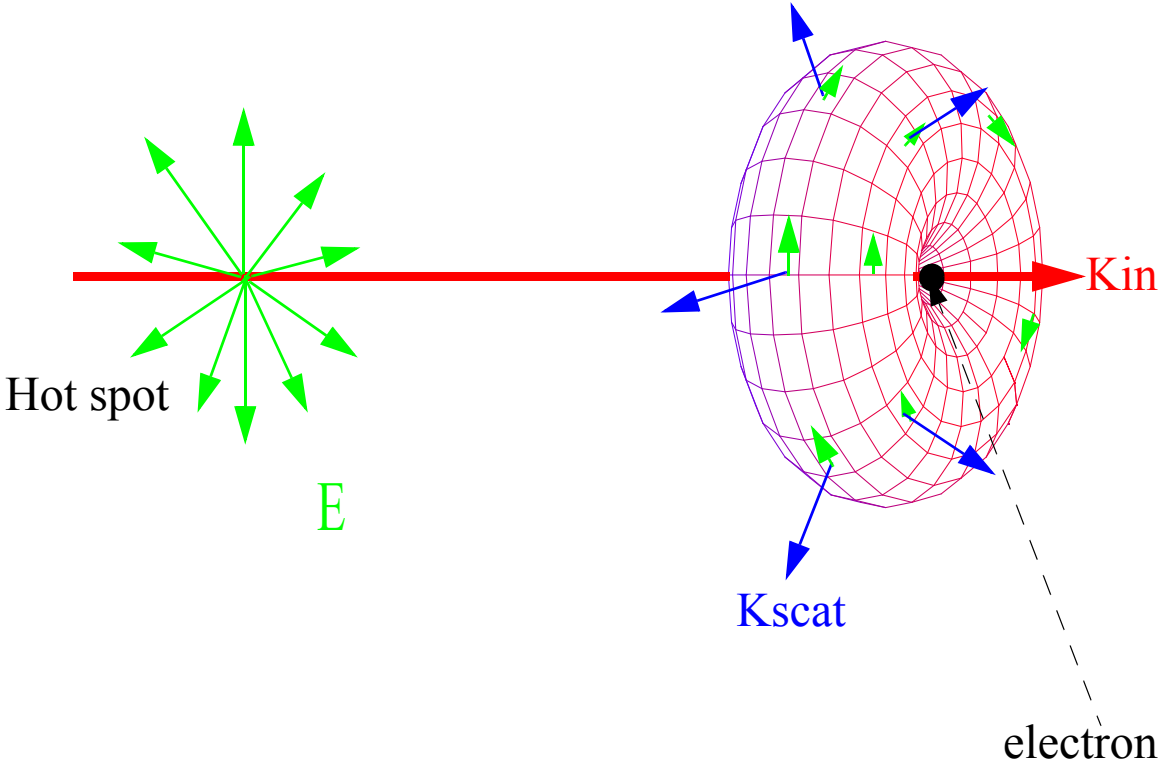
Hot spot
E
Kin
Kscat
electron

H. Edens

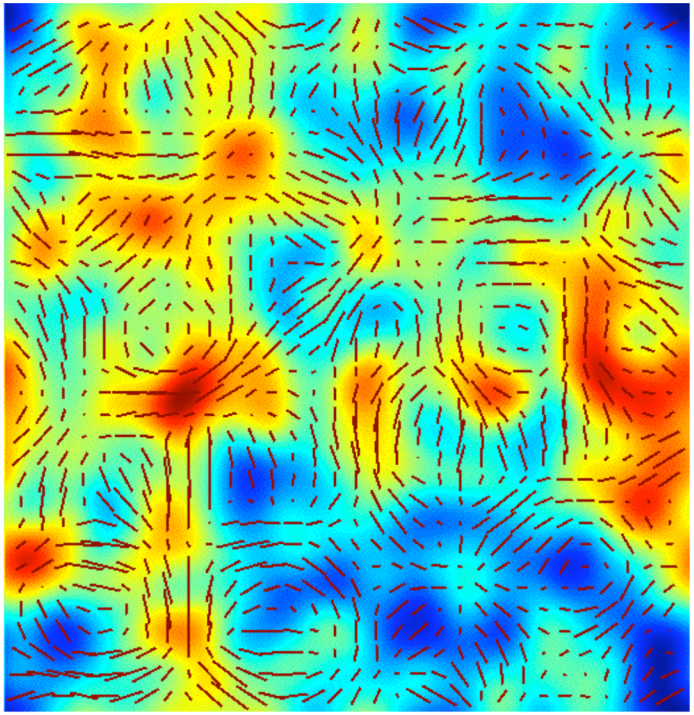

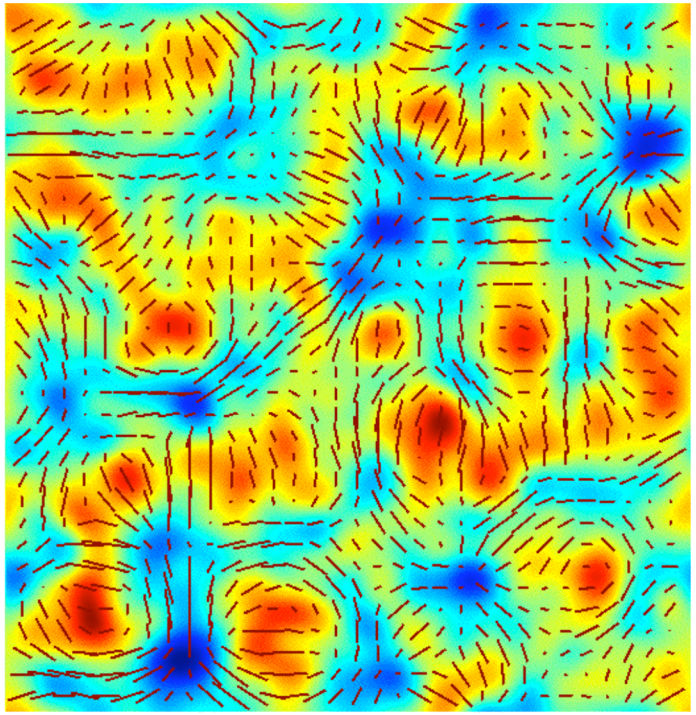

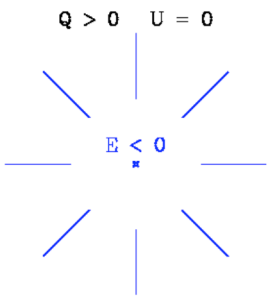

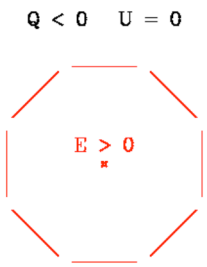

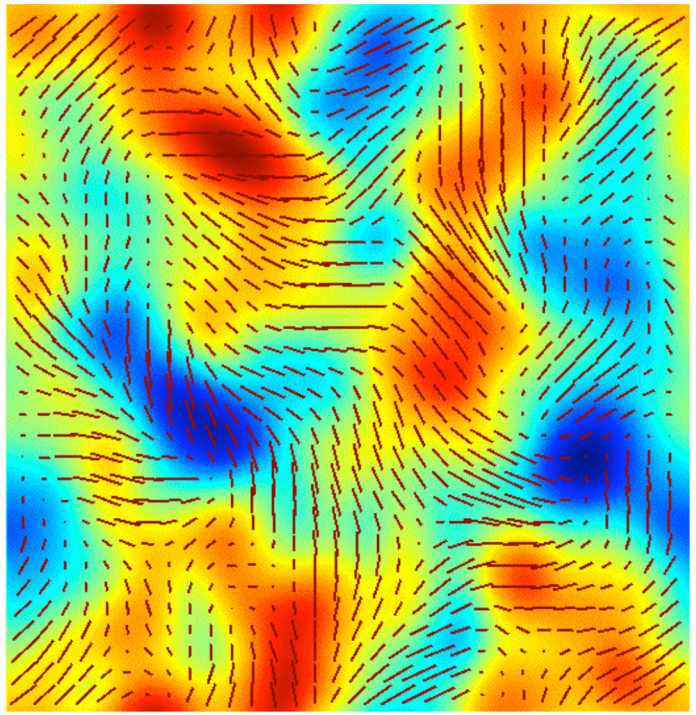

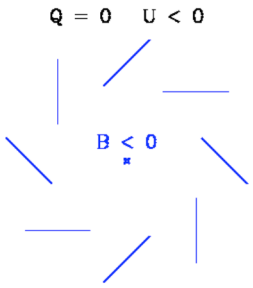

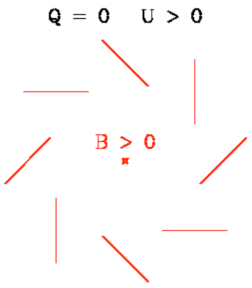

M. Zaldarriaga

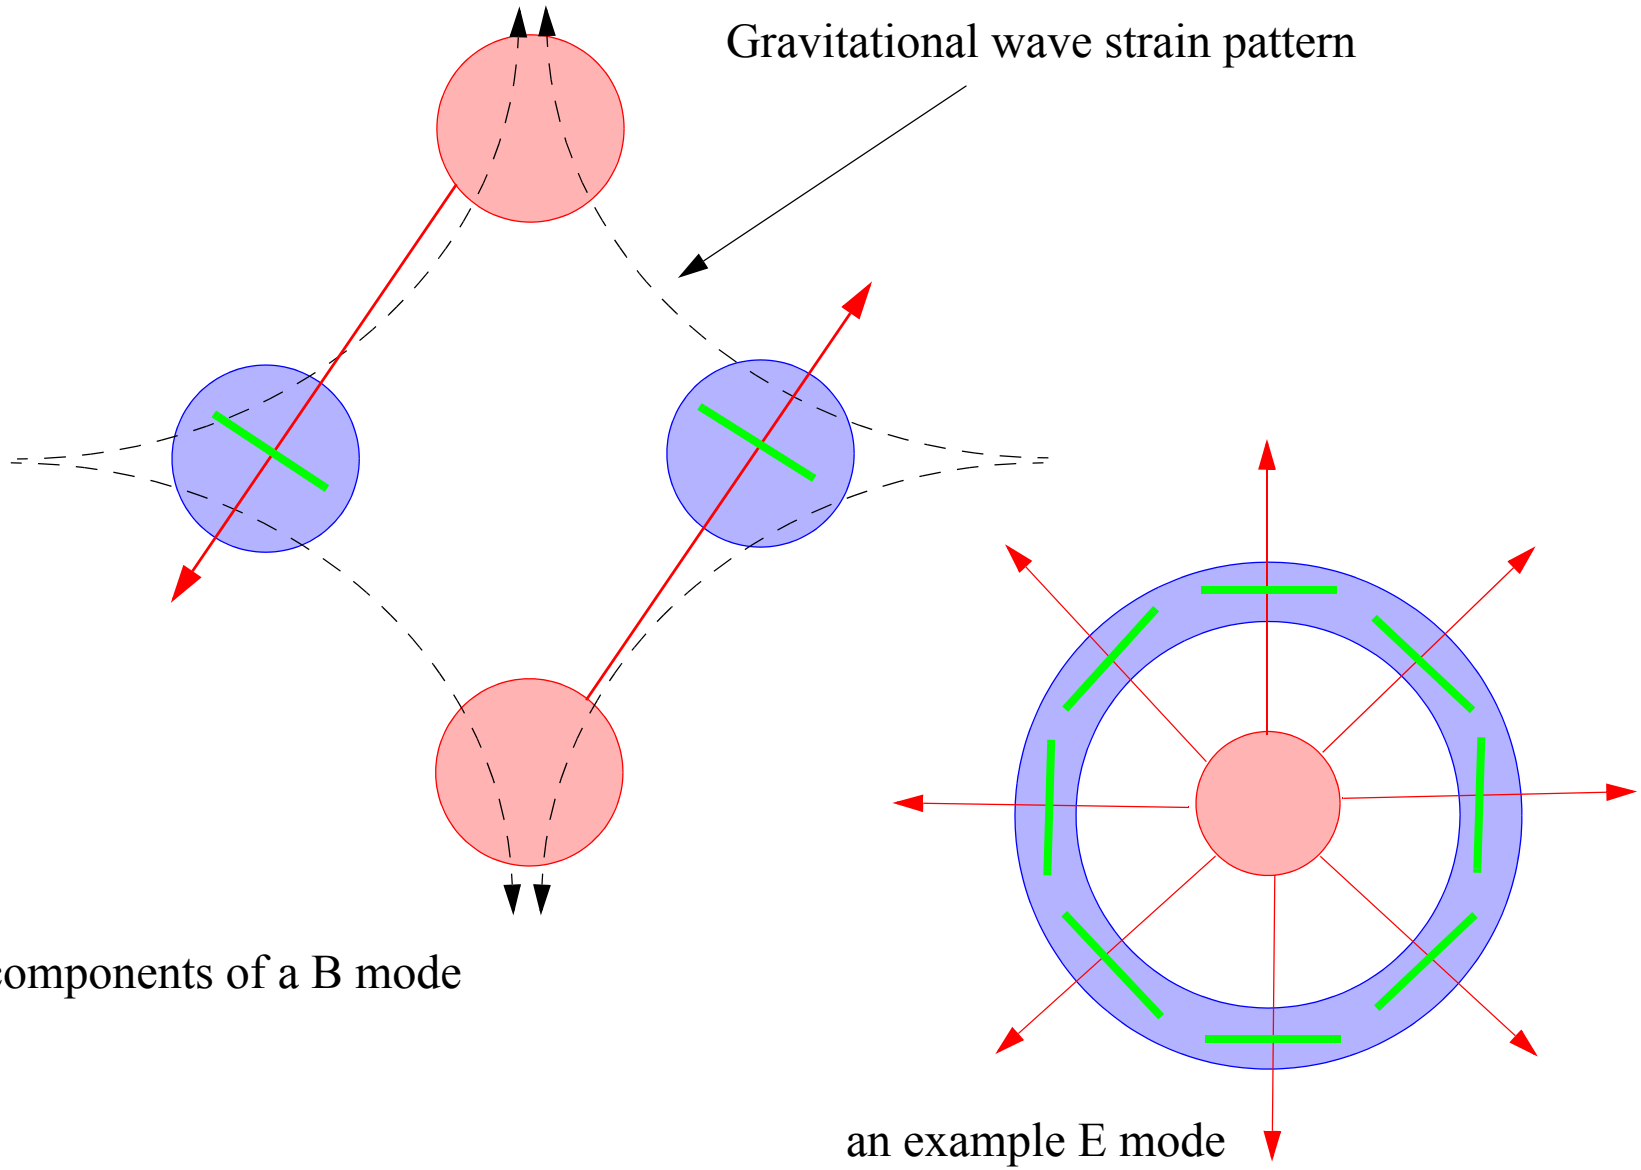
Gravitational wave strain pattern
components of a B mode
an example E mode

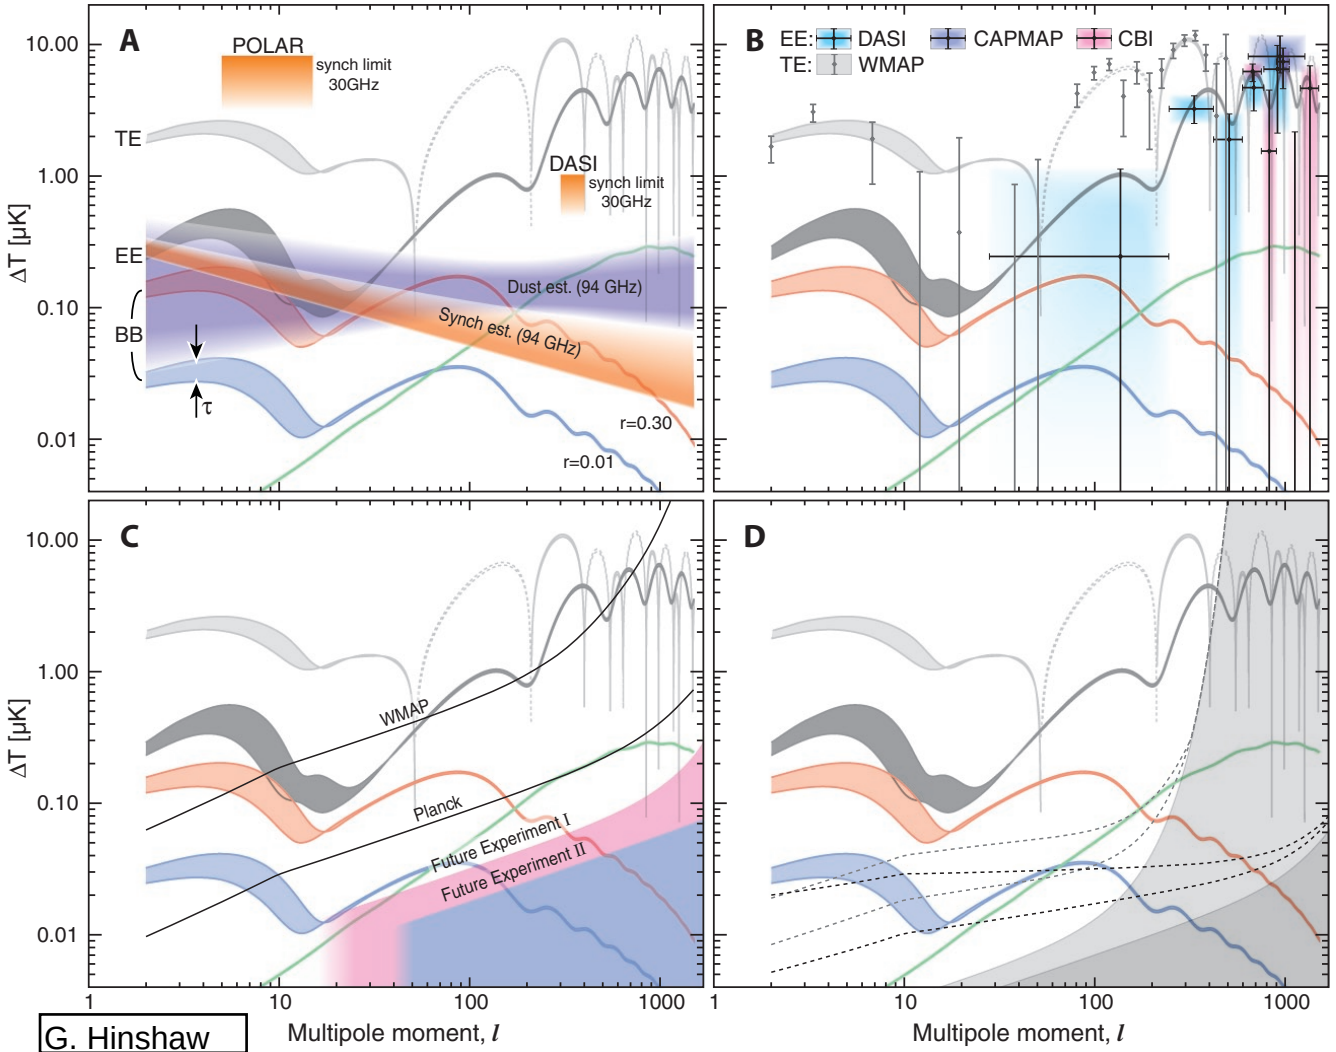
A
POLAR
synch limit
30GHz
TE
DASI
synch limit
30GHz
EE
BB
τ
Dust est. (94 GHz)
Synch est. (94 GHz)
r=0.30
r=0.01
B
EE:
TE:
DASI
WMAP
CAPMAP
CBI
C
WMAP
Planck
Future Experiment I
Future Experiment II
D
ΔT [μK]
10.00
1.00
0.10
0.01
1
10
100
1000
Multipole moment, l
G. Hinshaw

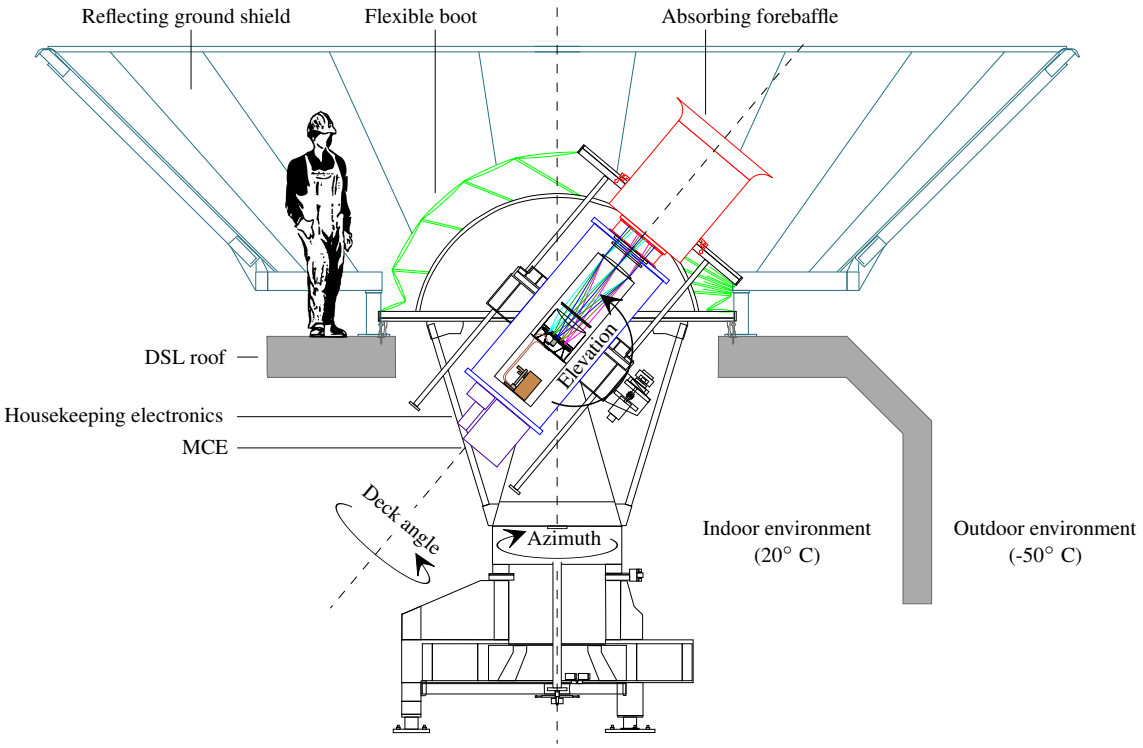

FIG. 1.— The BICEP2 telescope in the mount, looking out through the roof of the Dark Sector Laboratory (DSL) located 800 m from the geographic South Pole. The three-axis mount allows for motion in azimuth, elevation, and boresight rotation (also called "deck rotation"). An absorbing forebaffle and reflective ground screen prevent sidelobes from coupling to nearby objects on the ground. A flexible environmental seal or "boot" maintains a room temperature environment around the cryostat and mount. The telescope forms an insert within the liquid helium cryostat. The focal plane with polarization-sensitive TES bolometers is cooled to 270 mK by a ^{4}He/^{3}He/^{3}He sorption refrigerator. The housekeeping electronics (§8.4) and Multi-Channel Electronics (MCE, §9.2) attach to the lower bulkhead of the cryostat.

The South Pole is an excellent site for millimeter-wave observation from the ground, with a record of successful polarimetry experiments including DASI, BICEP1, QUAD and SPT. Situated on the Antarctic Plateau, it has exceptionally low precipitable water vapor (Chamberlin et al. 1997), reducing atmospheric noise due to the absorption and emission of water near the 150 GHz observing band. The South Pole site also has very stable weather, especially during the dark winter months, so that the majority of the data are taken under clear-sky conditions of very low atmospheric $1/f$ noise and low loading. The consistently low atmospheric loading is crucially important because the sensitivity of the experiment is limited by photon noise, so that low atmospheric emission is a key to high CMB mapping speed.

Finally, the Amundsen-Scott South Pole Station has hosted scientific research continuously since 1958. The station offers well-developed facilities with year-round staff and an established transportation infrastructure. BICEP1 and BICEP2 were housed in the Dark Sector Laboratory (DSL), which was built to support radio and millimeter-wave observatories in an area 1 km from the main station buildings and isolated from possible sources of electromagnetic interference.

3.2. *Telescope mount and drive*

The telescope sits in a three-axis mount (Fig. 1) supported on a steel and wood platform attached to the structural beams of the DSL building. The mount was originally built for BICEP1 by Vertex-RSI[17] along with a second, identical mount that has remained in North America for pre-deployment testing The mount attaches to a flexible environmental shield or "boot" (Fig. 2) attached to the roof of the building, so that the cryostat, electronics, and drive hardware are kept inside a climate-controlled, room temperature environment.

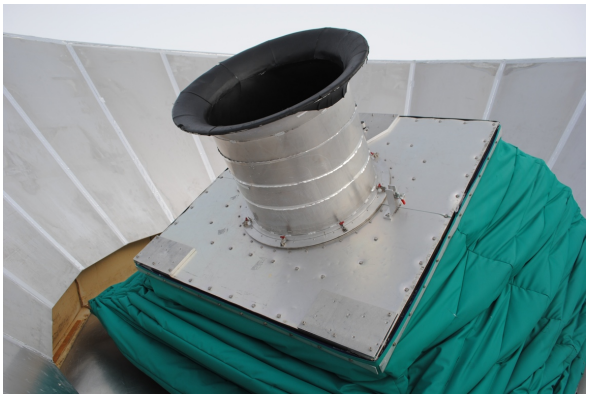

FIG. 2.— BICEP2 absorbing forebaffle, flexible environmental seal (the "boot"), and ground shield. The telescope and mount sit below the boot inside the Dark Sector Laboratory.

The mount moves in azimuth and elevation (which closely approximate right ascension and declination when observing from the South Pole). Its third axis is a rotation about the boresight of the telescope, also known as the "deck angle". When installed in DSL its range of motion is 50°–90° in elevation

[17]Now General Dynamics Satcom Technologies, 1500 Prodelin Drive, Newton, North Carolina 28658, 1-828-464-4141, http://www.gdsatcom.com/vertexrsi.php

short transitions into a superconducting state from the center outwards so that it does not trap flux as is possible with type-II superconductors.

6.3. *Printed circuit board*

An FR-4 printed circuit board (PCB) carries superconducting Al electrical traces and serves as a base for wire-bonding the tiles and the SQUID chips. Between the Cu plate and the PCB we place sheets of Metglas 2714A to create a low-field environment around the SQUID chips. The planar geometry between the Cu and Nb plates is especially effective in lowering the normal field component to which the SQUID chips are most sensitive. The SQUIDs sit on alumina carriers on the PCB, giving sufficient separation from the Metglas sheet to prevent magnetic coupling that could cause increased readout noise.

6.4. *Assembly*

Each detector tile is stacked with a high-conductivity z-cut crystal quartz anti-reflection (AR) wafer. We attach the detector tiles and AR wafers to the Cu plate in a way that provides precision alignment, allows for differential thermal contraction, and ensures sufficient heat-sinking. First precision-drilled holes and slots are made in the detector tiles and AR wafers. These register to pins that are press-fit in the Cu adjacent to each window. The detector tile and AR wafer stacks are clamped to the plate with machined tile clamps that allow slipping under thermal contraction. The weak clamping force is insufficient to effectively heat-sink the tiles, so we further connect a gold "picture frame" around the tile edges with gold wire bonds that make direct contact with the gold-plated Cu frame. The thermal conductivity (limited by the Kapitza resistance between the silicon substrate and the gold) is sufficient to prevent tile heating under thermal loading.

The detector tiles have Nb pads on their back edges to be connected to the PCB traces with superconducting Al wire bonds. The Nyquist and MUX chips are similarly wire-bonded to the PCB, as are NTD thermometers and heaters mounted directly on the detector tiles.

7. DETECTORS

The focal plane was populated with integrated arrays of antenna-coupled bolometers. This technology combines beam-defining planar slot antennas, inline frequency-selective filters, and TES detectors into a single monolithic package. The JPL Microdevices Laboratory produces these devices in the form of square silicon tiles, each containing an 8×8 array of dual-polarization spatial pixels (64 detector pairs or 128 individual bolometers). The BICEP2 focal plane had four of these tiles, for a total of 500 optically coupled detectors and 12 dark (no antenna) TES detectors. The detector tiles were characterized at Caltech and JPL during 2008–2009. The rapid fabrication cycle of the Caltech-JPL detectors made it possible to incorporate results of pre-deployment testing into the final set of four tiles deployed in BICEP2. Further details of the detector design and fabrication will be presented in the Detector Paper, which will report on improvements to the detector tiles leading up to BICEP2 as well as further developments in subsequent generations informed by BICEP2 testing.

7.1. *Antenna networks*

Optical power couples to each detector through an integrated planar phased-array antenna. The sub-radiators of

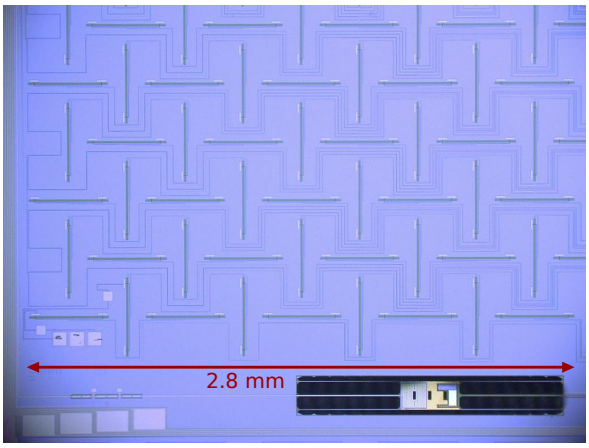

FIG. 8.— Partial view of one BICEP2 dual-polarization pixel, showing the band-defining filter (lower left), TES island (lower right), and part of the antenna network and summing tree. The vertically oriented slots are sensitive to horizontal polarization and form the antenna network for the A detector, while the horizontally oriented slots receive vertical polarization and are fed into the B detector. In this way the A and B detectors have orthogonal polarizations but are spatially co-located and form beams that are coincident on the sky. This view corresponds to a boresight angle of 90°. At boresight angle of 0° the A detectors receive vertical polarization and the B detectors receive horizontal polarization.

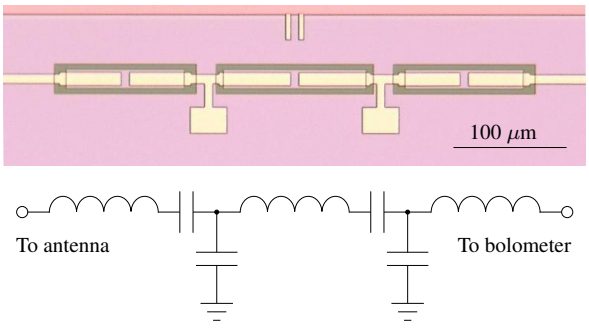

FIG. 9.— 150 GHz band-defining filter and equivalent circuit. Each filter consisted of three inductors in series, coupled to each other through a T-network of capacitors.

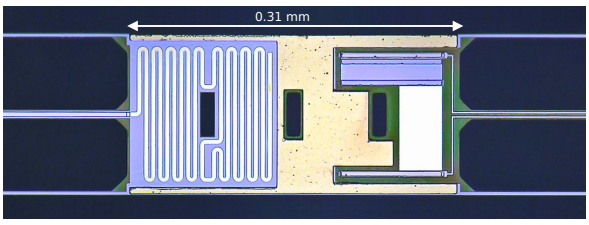

FIG. 10.— TES island for a single BICEP2 detector. The island was supported by six lithographically etched legs. Microwave power, entering from the left, terminated into a resistive meander. The deposited heat is measured as a decrease in electrical power (or current) dissipated in the titanium TES, which appears as a blue rectangle on the right of the island. The TES voltage bias was provided by two micro-strip lines at right. To increase the dynamic range of the device, an aluminum TES (seen as a white rectangle below the titanium film) was deposited in series with the titanium TES, providing linear response across a wide range of background loading conditions. The heat capacity of the island was tuned by adding 2.5 μm thick evaporated gold, which is distributed across the remaining real estate of the island. This made the detector time constants (§10.6) slow enough for stable operation.

that this procedure gives beam centers accurate to 2′ rms. We have simulated the effect of cosmological TE correlations as a bias on the beam centers and find it well below 5″. The same beam-fitting procedure has been repeated with *Planck* 143 GHz maps (Planck Collaboration et al. 2013; Planck HFI Core Team et al. 2011) instead of WMAP templates. The results are identical to within 15″ for all BICEP2 detectors.

When we compare the beam centers as fit from CMB maps at different boresight rotation angles, we detect an offset in the elevation direction of an average of 1′. We interpret this offset as an internal flexure of the focal plane assembly relative to the cryostat shell and the telescope mount.

12. OBSERVING STRATEGY

The BICEP2 observing strategy is based on deep integration in the region of the sky least contaminated by polarized foregrounds. The telescope spends 90% of its observing time on this CMB field, and the other 10% on a secondary Galactic field. These observations are grouped in schedules of three sidereal days, followed by a six-hour cryogenic service period. Within one three-day schedule the telescope scans in azimuth at a fixed boresight angle—the orientation of the telescope about its own axis. The details of the observing schedule have been chosen to allow for control of possible systematics such as drift in detector gain and ground-fixed signals.

12.1. *Observing fields*

BICEP2 spends most of its time observing the primary CMB field centered at (RA = 0 hr, dec = −57.5°). This 1000-degree^2 field (2% of the sky) lies well away from the Galactic plane, within a larger region known as the "Southern Hole" where polarized foregrounds are expected to be especially low. The BICEP2 field is the same one observed by BICEP1. It was selected for its very low level of expected Galactic dust emission, less than 1% of the sky median (Finkbeiner et al. 1999) as shown in Fig. 20. If the dust signal is polarized at 5%, the resulting contamination of the B-mode signal at 150 GHz will be below $r = 0.02$. The faint synchrotron signal within the Southern Hole has not been well measured, but a scaling of WMAP data at 23 GHz implies that the B-mode contamination at 150 GHz is at a level similar to or below that from dust (Nguyen et al. 2008).

The secondary BICEP2 field covers a part of the Galactic plane centered at (RA = 15 : 42 hr, dec = −55.0°). Observations of this field are used for Galactic science objectives (Bierman et al. 2011) and as a bright, partially polarized source for use in instrument characterization.

These same two fields have also been observed by BICEP1[29] and the *Keck Array*. Coverage of the same fields by the three experiments allows for consistency tests, cross-calibrations on the bright Galactic signal, and the possibility of achieving greater map depth by stacking CMB maps across multiple experiments. The additional frequencies of BICEP1 and the *Keck Array* (beginning in the 2014 season) also give spectral information needed to separate any foreground signals from the CMB.

12.2. *Scan pattern*

The telescope scans at 2.8°/s in azimuth, so that at an elevation of 57.5° a signal with frequency f (in Hz) corresponds to a multipole $\ell = 240f$. This sets the science band for the experiment: 0.05–1 Hz for $20 \leq \ell \leq 200$, or 2.6 Hz for $\ell = 500$.

[29] BICEP1 also observed a third field in a different part of the Galactic plane. This field has not been covered by BICEP2 or the *Keck Array*.

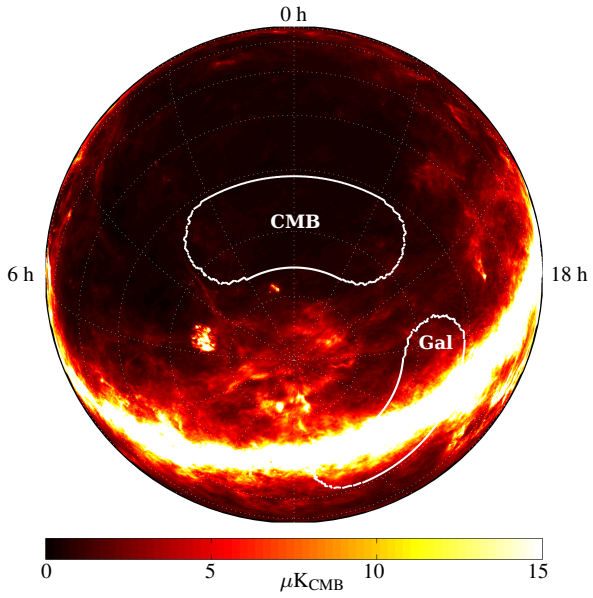

FIG. 20.— BICEP2 observing fields relative to the polarization amplitude predicted from FDS (Finkbeiner et al. 1999) model 8, assuming a 5% polarization fraction.

Each scan spans 64.2° in azimuth, of which the central 56.4° (77.7% of the duration of the scan) is covered at uniform speed and is used for mapmaking. The region around each turn-around is excluded from CMB analysis. The trajectory of each scan was optimized at the time of BICEP2 deployment for a gain of 4% in the usable, central part of the scan relative to BICEP1. The elevation is kept fixed as the telescope executes 53 round-trip scans over a period of 46 minutes. During this single "scan set" the telescope scans back and forth within fixed limits in azimuth, rather than continuously tracking the sky. Each scan set is preceded and followed by bracketing calibrations as described in §12.4, bringing the total duration of each scan set up to 50 minutes.

At the end of each 50-minute scan set, the telescope steps up by 0.25° in elevation and shifts the azimuth of the scan center to follow the apparent motion of the field on the sky before beginning the next scan set.

This scan pattern deliberately scans across a fixed range in azimuth within each 50-minute observing block, rather than a fixed range in right ascension. After 50 minutes the CMB has drifted by 12.5° relative to the ground. Therefore, any pickup of ground-fixed optical power, the magnetic field of the Earth or nearby structures, scan-fixed thermal fluctuations, or scan-fixed vibrational noise will all appear in the same locations from scan to scan. This allows us to remove these signals using a simple ground-subtraction algorithm (§13.6).

12.3. *Schedules and boresight angles*

A three-day schedule is divided into groups of 50-minute scan sets. Each of these groups, called an observing phase, contains ten scan sets (nine hours total) or seven scan sets (six hours total) along with the accompanying calibrations. During one full three-day schedule the telescope completes one six-hour cryogenic service phase, six 9-hour phases and one 6-hour phase on the CMB field, and one 6-hour phase on the Galactic field, as listed in Table 6. The azimuth/elevation pattern of a typical observing schedule is shown in Fig. 21. This

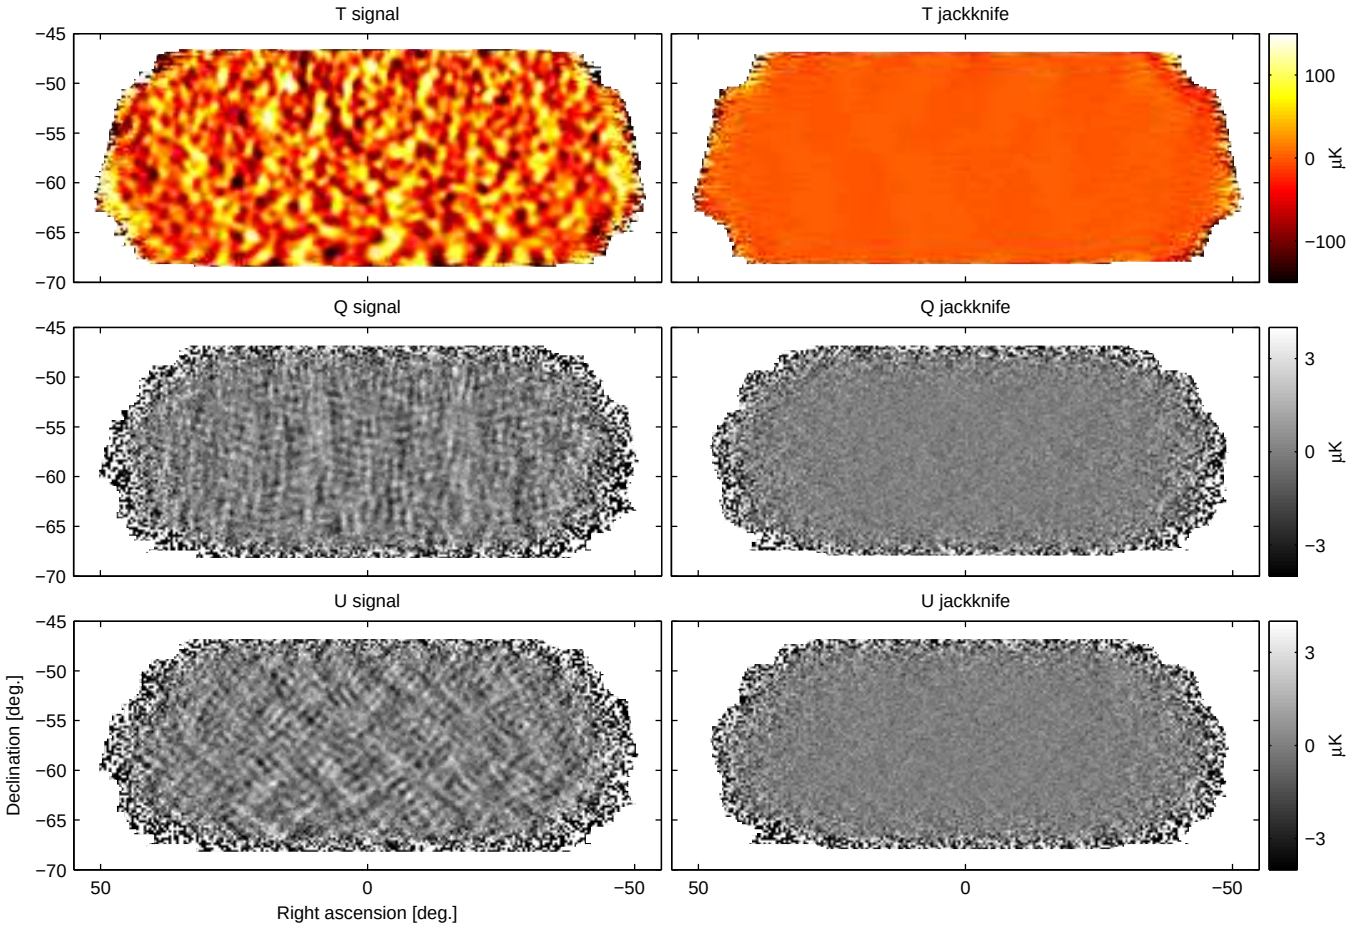

FIG. 1.— BICEP2 T, Q, U maps. The left column shows the basic signal maps with 0.25° pixelization as output by the reduction pipeline. The right column shows difference (jackknife) maps made with the first and second halves of the data set. No additional filtering other than that imposed by the instrument beam (FWHM 0.5°) has been done. Note that the structure seen in the Q&U signal maps is as expected for an E-mode dominated sky.

where the C_ℓ's are ΛCDM spectra from CAMB[22] with cosmological parameters taken from Planck Collaboration XVI (2013), and the $n_{\ell m}$ are normally distributed complex random numbers. For C_ℓ^{TT} we use a lensed-ΛCDM spectrum since the $a_{\ell m}^T$ from *Planck* NILC inherently contain lensing. We have found the noise level in the *Planck* NILC maps for our region of observation and multipole range to be low enough that it can be ignored.

Using the $a_{\ell m}^E$ we generate Nside=2048 maps using `synfast`. We substitute in the $a_{\ell m}^T$ from *Planck* 143 GHz so that the T map more closely resembles the T sky we expect to see with BICEP2. (This is also the map that is used in §4.6 to construct deprojection templates.) Additionally, we add in noise to the T map at the level predicted by the noise covariance in the *Planck* 143 GHz map, which allows us to simulate any deprojection residual due to noise in the *Planck* 143 GHz map.

5.1.2. *Lensing of input maps*

Lensing is added to the unlensed-ΛCDM maps using the `LensPix`[23] software (Lewis 2011). We use this software to generate lensed versions of the constrained CMB input $a_{\ell m}$'s described in §5.1.1. Input to the lensing operation are deflection angle spectra that are generated with CAMB as part of the standard computation of ΛCDM spectra. The lensing operation is performed before the beam smoothing is applied to form the final map products. We do not apply lensing to the 143 GHz temperature $a_{\ell m}^T$ from *Planck* since these inherently contain lensing. Our simulations hence approximate lensed CMB maps ignoring the lensing correlation between T and E.

[22] `http://camb.info/`
[23] `http://cosmologist.info/lenspix/`

5.2. *Noise Pseudo Simulations*

The previous BICEP1 and QUAD pipelines used a Fourier based procedure to make simulated noise timestreams, maintaining correlations between all channels (Pryke et al. 2009). For the increased channel count in BICEP2 this is computationally very expensive, so we have switched to an alternate procedure adapted from van Engelen et al. (2012). We perform additional coadds of the real pairmaps randomly flipping the sign of each scanset. The sign-flip sequences are constructed such that the total weight of positively and negatively weighted maps is equal. We have checked this technique against the older technique, and against another technique which constructs map noise covariance matrices, and have found them all to be equivalent to the relevant level of accuracy. By default we use the sign flipping technique and refer to these realizations as "noise pseudo simulations."

We add the noise maps to the lensed-ΛCDM realizations to form signal plus noise simulations — hereafter referred to as lensed-ΛCDM+noise.

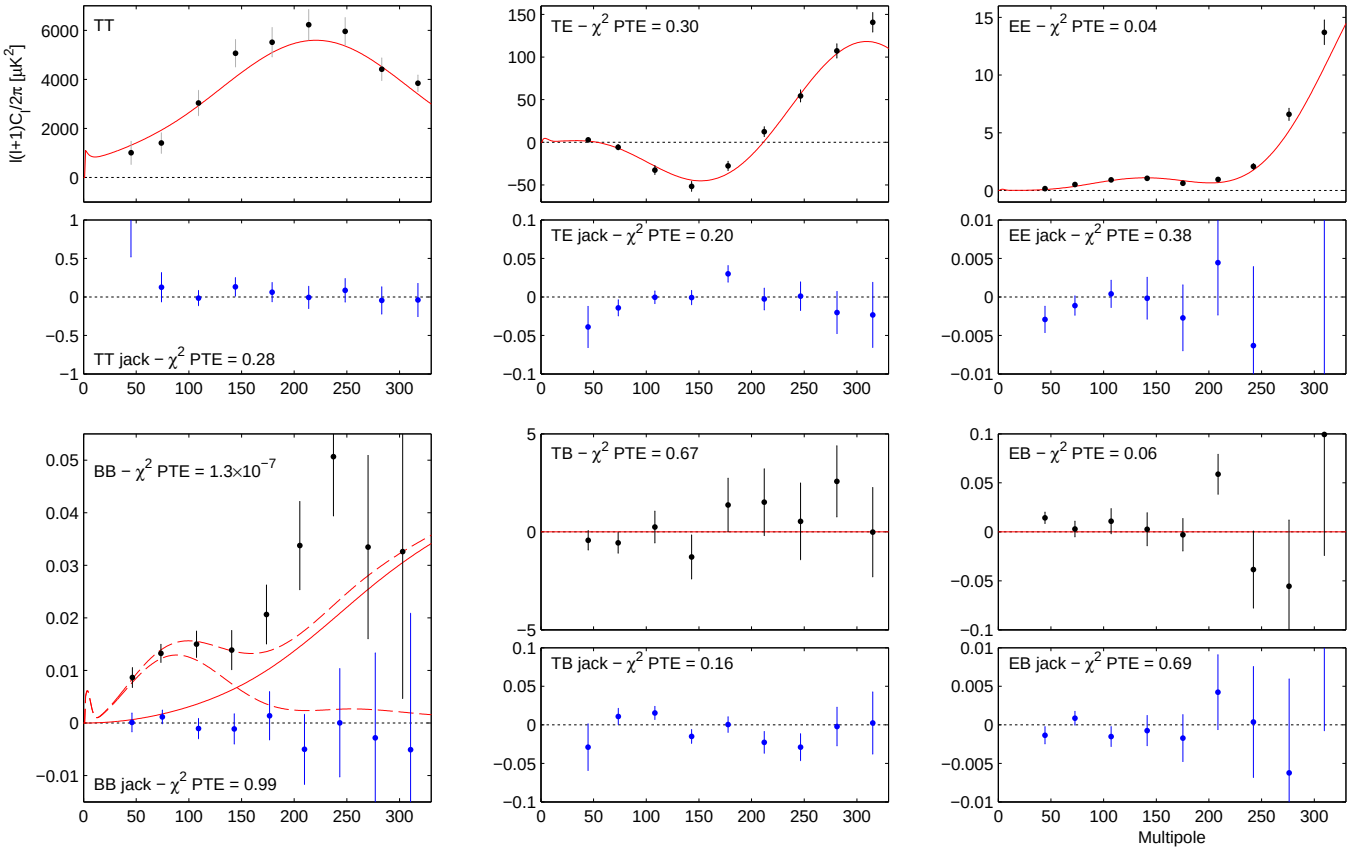

FIG. 2.— BICEP2 power spectrum results for signal (black points) and temporal-split jackknife (blue points). The red curves show the lensed-ΛCDM theory expectations — in the case of *BB* an $r = 0.2$ spectrum is also shown. The error bars are the standard deviations of the lensed-ΛCDM+noise simulations. The probability to exceed (PTE) the observed value of a simple χ^2 statistic is given (as evaluated against the simulations). Note the very different y-axis scales for the jackknife spectra (other than *BB*). See the text for additional discussion of the *BB* spectrum.

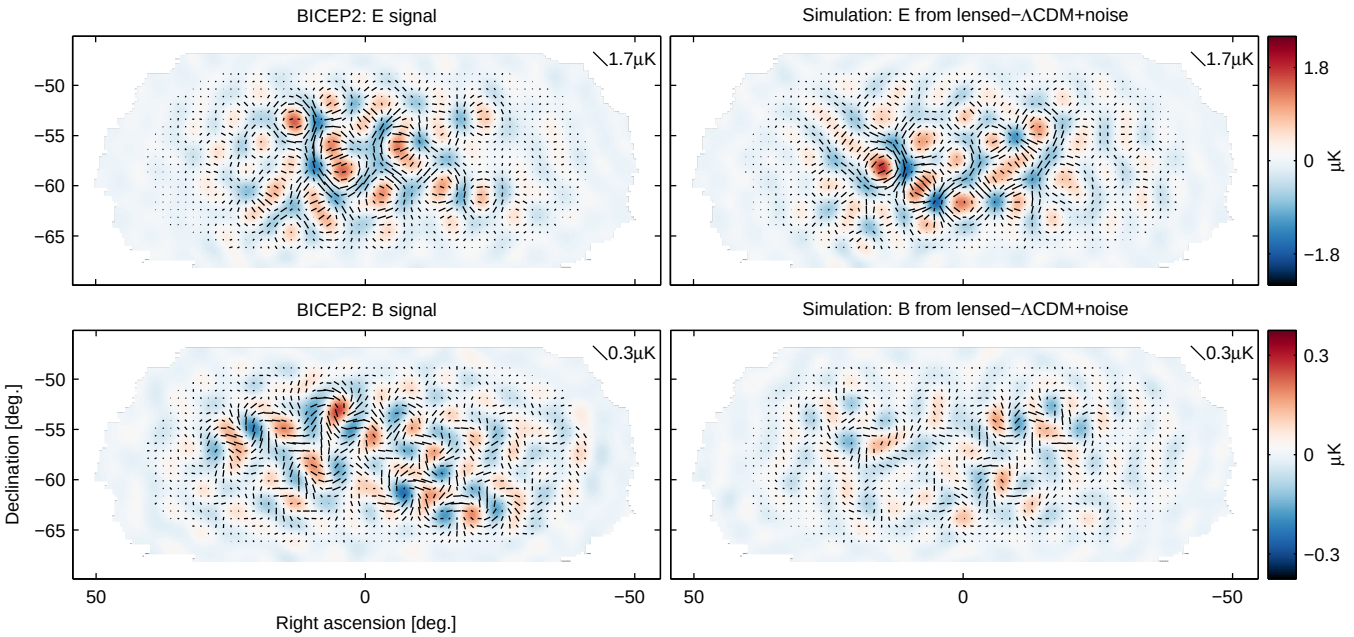

FIG. 3.— *Left:* BICEP2 apodized *E*-mode and *B*-mode maps filtered to $50 < \ell < 120$. *Right:* The equivalent maps for the first of the lensed-ΛCDM+noise simulations. The color scale displays the *E*-mode scalar and *B*-mode pseudoscalar patterns while the lines display the equivalent magnitude and orientation of linear polarization. Note that excess *B*-mode is detected over lensing+noise with high signal-to-noise ratio in the map ($s/n > 2$ per map mode at $\ell \approx 70$). (Also note that the *E*-mode and *B*-mode maps use different color/length scales.)

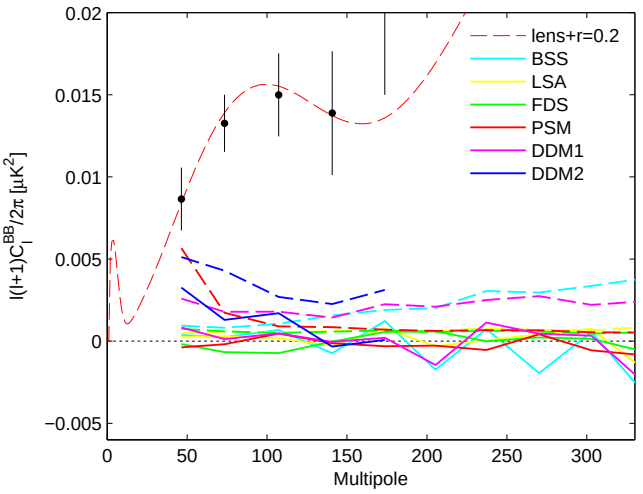

FIG. 6.— Polarized dust foreground projections for our field using various models available in the literature, and two new ones formulated using publically available information from *Planck*. Dashed lines show autospectra of the models, while solid lines show cross spectra between the models and the BICEP2 maps. The cross spectra are consistent with zero, and the DDM2 auto spectrum (at least) is noise biased high (and is hence truncated to $\ell < 200$). The BICEP2 auto spectrum from Figure 2 is also shown with the lensed-ΛCDM+$r = 0.2$ spectrum.

stant emissivity value of 1.6 and a constant temperature of 19.6 K. In our field these values agree both with the mean values shown by the *Planck* Collaboration in dust polarization[31], and with the median values of the recently delivered *Planck* dust model (Planck Collaboration et al. 2013a). A uniform 5% sky polarization fraction is assumed in agreement with the first all-sky images of dust polarization shown by the *Planck* Collaboration[32]. The polarization angles are taken from the PSM.

DDM2: "Data Driven Model 2" (DDM2) constructed using all publicly available information from *Planck*. Uses the same dust model temperature map as DDM1, with polarization fractions and angles matching those shown by the *Planck* Collaboration[32].

All of the the models except FDS make explicit predictions of the actual polarized dust pattern in our field — presumably with varying probabilities of success. We can therefore search for a correlation between the models and our signal by taking cross spectra against the BICEP2 maps. Figure 6 shows the resulting *BB* auto and cross spectra — note that the autospectra are all well below the level of our observed signal and that the cross spectra are consistent with zero[33]. We also note that the DDM2 model auto spectrum (which is the highest) contains uncorrected noise bias from the polarization fraction and angle maps (which is why this curve in Figure 6 is truncated to $\ell < 200$).

9.2. *Synchrotron*

In our field and at angular scales of $\ell > 30$ the WMAP K-band (23 GHz) maps are noise dominated. Extrapolating them to our observing frequency using a spectral index of $\beta = -3.3$ derived from WMAP foreground products results in an upper limit to synchrotron contamination equivalent to $r = 0.003$. Taking the cross spectrum against our observed map indicates that the true value is lower.

9.3. *Point Sources*

Extragalactic point sources might also potentially be a concern. Using the 143 GHz fluxes for the sources in our field from the *Planck* catalog (Planck Collaboration et al. 2013b), together with polarization information from ATCA (Massardi et al. 2011) we find that the contribution to the *BB* spectrum is equivalent to $r \approx 0.001$. This is consistent with the projections of Battye et al. (2011).

10. CROSS SPECTRA

10.1. *Cross Spectra with* BICEP1

BICEP1 observed essentially the same field as BICEP2 from 2006 to 2008. While a very similar instrument in many ways the focal plane technology of BICEP1 was entirely different, employing horn fed PSBs read out via neutron transmutation-doped (NTD) germanium thermistors (see T10 for details). The high-impedance NTD devices and readouts have different susceptibility to microphonic pickup and magnetic fields, and the shielding of unwanted RFI/EMI was significantly different from that of BICEP2. The beam systematics were also quite different with a more conservative edge taper and a more complex pattern of observed pair centroid offsets. BICEP1 had detectors at both 100 and 150 GHz.

Figure 7 compares the BICEP2 *EE* and *BB* auto spectra with cross spectra taken against the 100 and 150 GHz maps from BICEP1. For *EE* the correlation is extremely strong, which simply confirms that the mechanics of the process are working as expected. For *BB* the signal-to-noise is of course much lower, but there appear to be positive correlations. To test the compatibility of the *BB* auto and cross spectra we take the differences and compare to the differences of lensed-ΛCDM+noise+$r = 0.2$ simulations (which share common input skies)[34]. Using bandpowers 1–5 the χ^2 and χ PTEs are mid-range indicating that the spectra are compatible to within the noise. (This is also true for *EE*.)

Calculating the *BB* χ^2 and χ statistics against the lensed-ΛCDM model the BICEP2$\times$BICEP1$_{150}$ spectrum has PTEs of 0.37 and 0.05 respectively. However, BICEP2$\times$BICEP1$_{100}$ has PTEs of 0.005 and 0.001 corresponding to $\approx 3\sigma$ detection of power in the cross spectrum. While it may seem surprising that one cross spectrum shows a stronger detection than the other, it turns out not to be unusual in lensed-ΛCDM+noise+$r = 0.2$ simulations.

10.2. *Spectral Index Constraint*

We can use the BICEP2 auto and BICEP2$\times$BICEP1$_{100}$ spectra shown in Figure 7 to constrain the frequency dependence of the nominal signal. If the signal at 150 GHz were due to synchrotron we would expect the frequency cross spectrum to be much larger in amplitude than the BICEP2 auto spectrum. Conversely if the 150 GHz power were due to polarized dust emission we would not expect to see a significant correlation with the 100 GHz sky pattern.

[31] http://www.rssd.esa.int/SA/PLANCK/docs/eslab47/Session09_Data_Processing/47ESLAB_April_04_10_30_Aumont.pdf

[32] http://www.rssd.esa.int/SA/PLANCK/docs/eslab47/Session07_Galactic_Science/47ESLAB_April_04_11_25_Bernard.pdf

[33] The cross spectra between each model and real data are consistent with the cross spectra between that model and (uncorrelated) lensed-LCDM+noise simulations.

[34] For all spectral difference tests we compare against lensed-ΛCDM+noise+$r = 0.2$ simulations as the cross terms between signal and noise increase the variance even for perfectly common sky coverage.

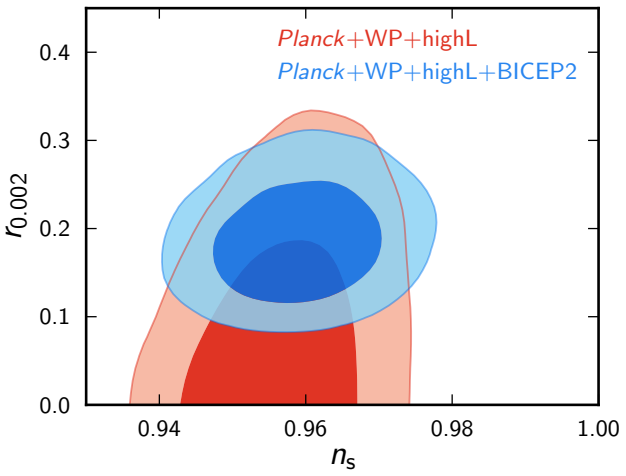

FIG. 13.— Indirect constraints on r from CMB temperature spectrum measurements relax in the context of various model extensions. Shown here is one example, following Planck Collaboration XVI (2013) Figure 23, where tensors and running of the scalar spectral index are added to the base ΛCDM model. The contours show the resulting 68% and 95% confidence regions for r and the scalar spectral index n_s when also allowing running. The red contours are for the "Planck+WP+highL" data combination, which for this model extension gives a 95% bound $r < 0.26$ (Planck Collaboration XVI 2013). The blue contours add the BICEP2 constraint on r shown in the center panel of Figure 10. See the text for further details.

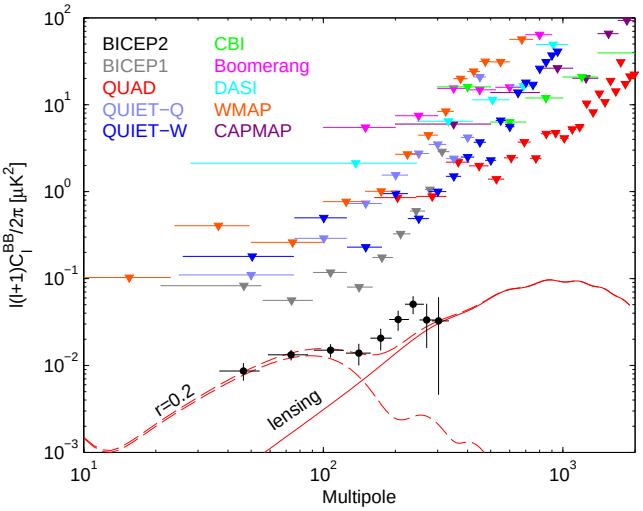

FIG. 14.— BICEP2 *BB* auto spectra and 95% upper limits from several previous experiments (Leitch et al. 2005; Montroy et al. 2006; Sievers et al. 2007; Bischoff et al. 2008; Brown et al. 2009; QUIET Collaboration et al. 2011, 2012; Bennett et al. 2013; Barkats et al. 2014). The curves show the theory expectations for $r = 0.2$ and lensed-ΛCDM.

To fully exploit this unprecedented sensitivity we have expanded our analysis pipeline in several ways. We have added an additional filtering of the timestream using a template temperature map (from *Planck*) to render the results insensitive to temperature to polarization leakage caused by leading order beam systematics. In addition we have implemented a map purification step that eliminates ambiguous modes prior to *B*-mode estimation. These deprojection and purification steps are both straightforward extensions of the kinds of linear filtering operations that are now common in CMB data analysis.

The power spectrum results are perfectly consistent with lensed-ΛCDM with one striking exception: the detection of a large excess in the *BB* spectrum in exactly the ℓ range where an inflationary gravitational wave signal is expected to peak. This excess represents a 5.2σ excursion from the base lensed-ΛCDM model. We have conducted a wide selection of jackknife tests which indicate that the *B*-mode signal is common on the sky in all data subsets. These tests offer very strong empirical evidence against a systematic origin for the signal.

In addition we have conducted extensive simulations using high fidelity per channel beam maps. These confirm our understanding of the beam effects, and that after deprojection of the two leading order modes, the residual is far below the level of the signal which we observe.

Having demonstrated that the signal is real and "on the sky" we proceeded to investigate if it may be due to foreground contamination. Polarized synchrotron emission from our galaxy is easily ruled out using low frequency polarized maps from WMAP. For polarized dust emission public maps are not yet available. We therefore investigate a range of models including new ones which use all of the information which is currently available from *Planck*. These models all predict auto spectrum power well below our observed level. In addition none of them show any significant cross correlation with our maps.

Taking cross spectra against 100 GHz maps from BICEP1 we find significant correlation and set a constraint on the spectral index of the signal consistent with CMB, and disfavoring synchrotron and dust by 2.3σ and 2.2σ respectively. The fact that the BICEP1 and *Keck Array* maps cross correlate is powerful further evidence against systematics.

The simplest and most economical remaining interpretation of the *B*-mode signal which we have detected is that it is due to tensor modes — the IGW template is an excellent fit to the observed excess. We therefore proceed to set a constraint on the tensor-to-scalar ratio and find $r = 0.20^{+0.07}_{-0.05}$ with $r = 0$ ruled out at a significance of 7.0σ. Multiple lines of evidence have been presented that foregrounds are a subdominant contribution: i) direct projection of the best available foreground models, ii) lack of strong cross correlation of those models against the observed sky pattern (Figure 6), iii) the frequency spectral index of the signal as constrained using BICEP1 data at 100 GHz (Figure 8), and iv) the spatial and power spectral form of the signal (Figures 3 and 10).

Subtracting the various dust models and re-deriving the r constraint still results in high significance of detection. For the model which is perhaps the most likely to be close to reality (DDM2 cross) the maximum likelihood value shifts to $r = 0.16^{+0.06}_{-0.05}$ with $r = 0$ disfavored at 5.9σ. These high values of r are in apparent tension with previous indirect limits based on temperature measurements and we have discussed some possible resolutions including modifications of the initial scalar perturbation spectrum such as running. However we emphasize that we do not claim to know what the resolution is.

Figure 14 shows the BICEP2 results compared to previous upper limits. The long search for tensor *B*-modes is apparently over, and a new era of *B*-mode cosmology has begun.

BICEP2 was supported by the US National Science Foundation under grants ANT-0742818 and ANT-1044978 (Caltech/Harvard) and ANT-0742592 and ANT-1110087 (Chicago/Minnesota). The development of antenna-coupled detector technology was supported by the JPL Research and Technology Development Fund and grants 06-ARPA206-0040 and 10-SAT10-0017 from the NASA APRA and SAT

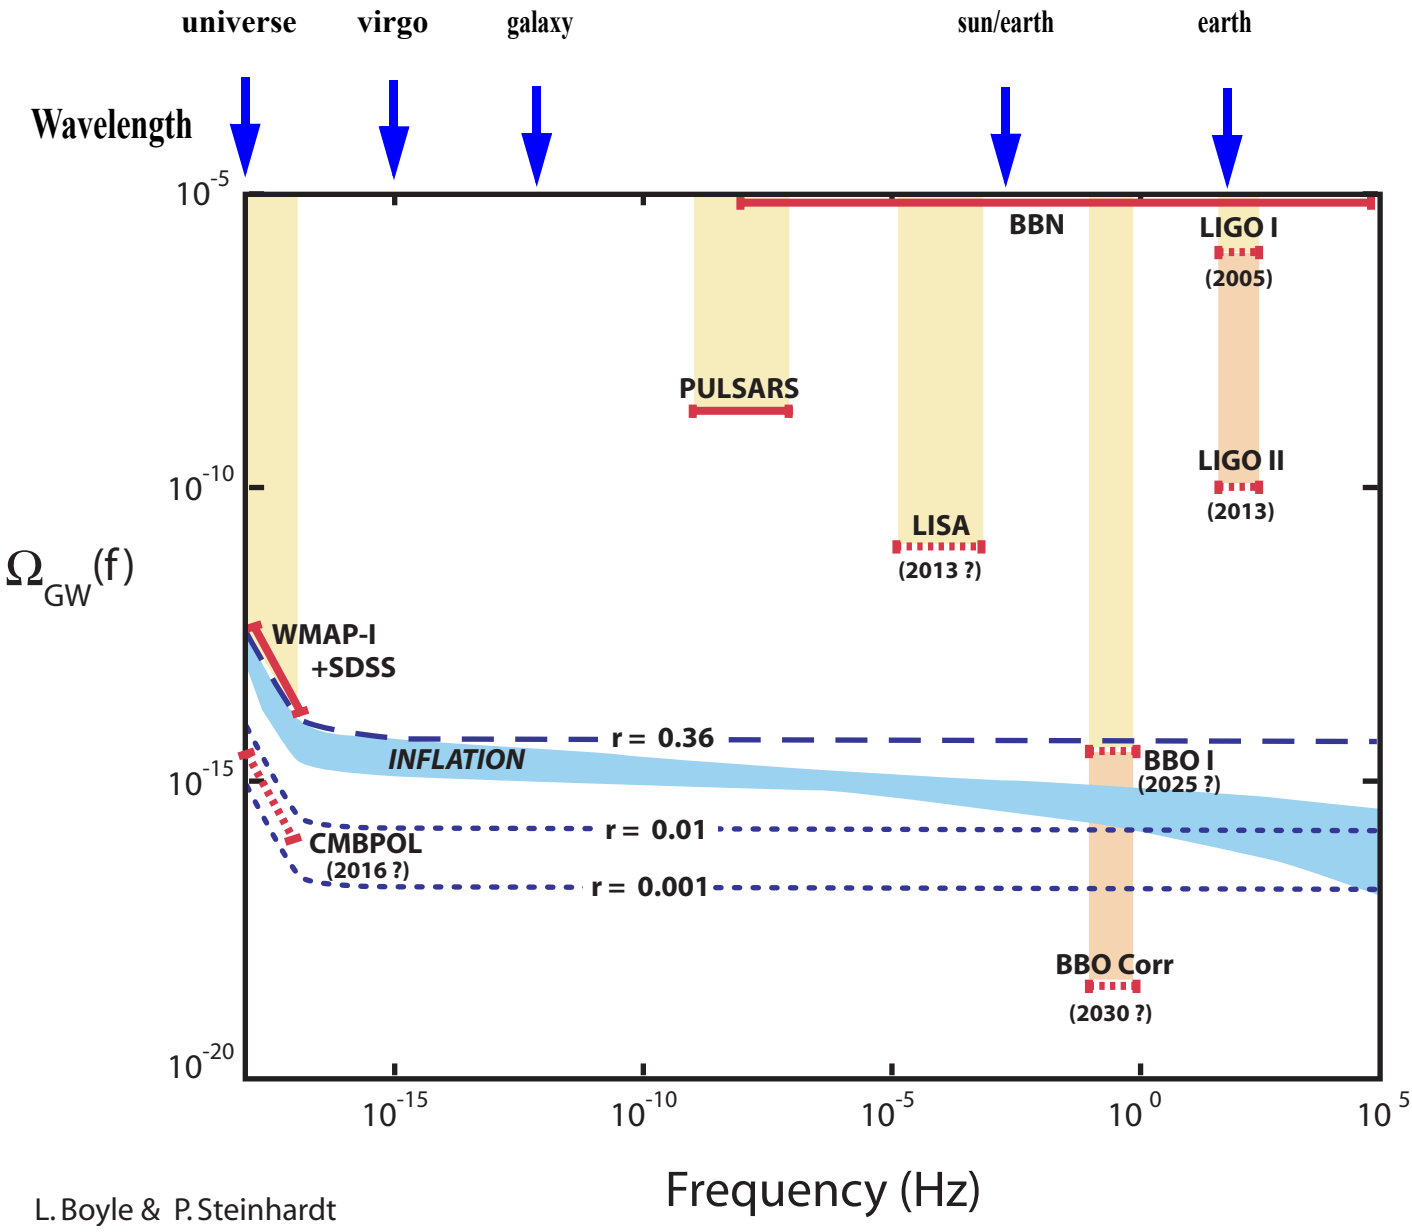

L. Boyle & P. Steinhardt