

SPIDER OPTIMIZATION: PROBING THE SYSTEMATICS OF A LARGE-SCALE *B*-MODE EXPERIMENT

C. J. MAC TAVISH,¹ P. A. R. ADE,² E. S. BATTISTELLI,³ S. BENTON,⁴ R. BIHARY,⁵ J. J. BOCK,^{6,7} J. R. BOND,¹
J. BREVIK,⁶ S. BRYAN,⁵ C. R. CONTALDI,⁸ B. P. CRILL,^{1,9} O. DORÉ,¹ L. FISSEL,⁹ S. R. GOLWALA,⁷
M. HALPERN,³ G. HILTON,¹⁰ W. HOLMES,⁶ V. V. HRISTOV,⁷ K. IRWIN,¹⁰ W. C. JONES,^{6,7}
C. L. KUO,⁷ A. E. LANGE,⁷ C. LAWRIE,⁵ T. G. MARTIN,¹¹ P. MASON,⁷ T. E. MONTROY,⁵
C. B. NETTERFIELD,^{4,9} D. RILEY,⁵ J. E. RUHL,⁵ M. RUNYAN,⁷ A. TRANGSRUD,⁷
C. TUCKER,² A. TURNER,⁶ M. VIERO,⁹ AND D. WIEBE⁴

Received 2007 October 1; accepted 2008 July 28

ABSTRACT

Spider is a long-duration, balloon-borne polarimeter designed to measure large-scale cosmic microwave background (CMB) polarization with very high sensitivity and control of systematics. The instrument will map over half the sky with degree angular resolution in the *I*, *Q*, and *U* Stokes parameters in four frequency bands from 96 to 275 GHz. *Spider*'s ultimate goal is to detect the primordial gravity-wave signal imprinted on the CMB *B*-mode polarization. One of the challenges in achieving this goal is the minimization of the contamination of *B*-modes by systematic effects. This paper explores a number of instrument systematics and observing strategies in order to optimize *B*-mode sensitivity. This is done by injecting realistic-amplitude, time-varying systematics into a set of simulated time streams. Tests of the impact of detector noise characteristics, pointing jitter, payload pendulations, polarization angle offsets, beam systematics, and receiver gain drifts are shown. *Spider*'s default observing strategy is to spin continuously in azimuth, with polarization modulation achieved by either a rapidly spinning half-wave plate or a rapidly spinning goniometer and a slowly stepped half-wave plate. Although the latter is more susceptible to systematics, the results shown here indicate that either mode of operation can be used by *Spider*.

Subject headings: cosmic microwave background — gravitational waves —
methods: data analysis — polarization

Online material: color figures

1. INTRODUCTION

In the past decade, a wealth of data have pointed to a “standard model” of the universe, composed of ~5% ordinary matter, ~22% dark matter, and ~73% dark energy in a flat geometry (see, e.g., MacTavish et al. 2006). The flatness of the universe, the near isotropy of the cosmic microwave background (CMB), and the nearly scale-invariant nature of the primordial scalar perturbations from which structure grew support the existence of an early accelerating phase dubbed “inflation.” A necessary by-product of inflation is tensor perturbations from quantum fluctuations in gravity waves. Detection of this cosmological gravity-wave background (CGB) would give strong evidence of an inflationary period and determine its energy scale, while a powerful upper limit

would point to more radical inflationary scenarios, e.g., involving string theory, or some noninflationary explanation of the observations (see, e.g., Liddle & Lyth 2000).

The CGB imprints a unique signal in the curl-like, or *B*-mode, component of the polarization of the CMB; detection of a *B*-mode signal can be used to infer the presence of a CGB at the time of decoupling (Seljak & Zaldarriaga 1997). Direct detection of gravity waves is many decades off; the advanced *Big Bang Observer*, a successor to *LISA*, has been suggested as a way to achieve this (Phinney et al. 2005; Harry et al. 2006). Thus, a measurement of the primordial *B*-modes is the only feasible near-term way to detect the CGB and have a new window into the physics of the early universe (Bock et al. 2006).

A CGB with a potentially measurable amplitude is a by-product of the simplest models of single-field inflation which can reproduce the scalar spectral tilt observed in current combined CMB data (Spergel et al. 2007; MacTavish et al. 2006). Examples are chaotic inflation from power-law inflaton potentials (Linde 1983; Linde et al. 2005) or natural inflation from cosine inflaton potentials involving angular (axionic) degrees of freedom (Adams et al. 1993). The amplitude is usually parameterized in terms of the ratio of the tensor power spectrum to the scalar power spectrum, $r = \mathcal{P}_t(k_p)/\mathcal{P}_s(k_p)$, evaluated at a comoving wavenumber pivot k_p , typically taken to be 0.002 Mpc^{-1} . Chaotic inflation predicts $r \approx 0.13$ for a ϕ^2 potential and $r \approx 0.26$ for a ϕ^4 potential, and natural inflation predicts $r \approx 0.02$ – 0.05 .

The potential energy V driving inflation is related to r by $V \approx (10^{16} \text{ GeV})^4 r/0.1$ (Lyth 1997). Low-energy inflation models have low or negligible amplitudes for the CGB. To get the observed scalar slope and yet small r requires special tuning of the potential. These are often more complicated, multiple-field models, e.g., Linde (1994), or string-inspired brane or moduli models

¹ Canadian Institute for Theoretical Astrophysics (CITA), University of Toronto, Toronto, ON M5S 3H8, Canada.

² School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, Wales, UK.

³ Department of Physics and Astronomy, University of British Columbia, Vancouver, BC V6T 1Z1, Canada.

⁴ Department of Physics, University of Toronto, Toronto, ON M5S 3H8, Canada.

⁵ Department of Physics, Case Western Reserve University, Cleveland, OH 44106.

⁶ Jet Propulsion Laboratory, Pasadena, CA 91109.

⁷ Department of Physics, California Institute of Technology, Pasadena, CA 91125.

⁸ Theoretical Physics, Blackett Laboratory, Imperial College, London SW7 2A, UK.

⁹ Department of Astronomy and Astrophysics, University of Toronto, Toronto, ON M5S 3H4, Canada.

¹⁰ National Institute of Standards and Technology, Boulder, CO 80305.

¹¹ Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, ON M5S 3G8, Canada.

(Kallosh 2007). Given the collection of models it is difficult to predict a precise range for the expected tensor level, and the prior probability for r should be considered as wide open.

Recent CMB data have reached the sensitivity level required to constrain the amplitude of and possibly characterize the gradient-like, or E -mode, component of the polarization (Kovac et al. 2002; Hedman et al. 2002; Readhead et al. 2004; Montroy et al. 2006a; Page et al. 2007; Ade et al. 2008). A significant complication of the measurements is that the E -mode amplitude is an order of magnitude lower than the total intensity. In addition, galactic foregrounds such as synchrotron and dust are expected to be significant at these amplitudes (Kogut et al. 2007). Furthermore, the polarization properties of foregrounds are largely unknown. Constraining B -modes presents an even greater challenge, as it is a near certainty that polarized foregrounds will dominate the signal.

The next generation of CMB experiments will benefit from a revolution in detector fabrication in the form of arrays of antenna-coupled bolometers (Goldin et al. 2002; Myers et al. 2002; Kuo et al. 2006). The antenna-coupled design is entirely photolithographically fabricated, greatly simplifying detector production. In addition, the densely populated antennas allow very efficient use of the focal plane area.

Spider will make use of this technological advance in the form of 2624 polarization-sensitive detectors observing in four frequency bands from 96 to 275 GHz. A multifrequency observing strategy is a necessary requirement to allow for subtraction of the foreground signal. *Spider* will observe over a large fraction of the sky at degree-scale resolution, producing high signal-to-noise ratio polarization maps of the foregrounds at each frequency.

Extraordinarily precise understanding of systematic effects within the telescope will be required to measure the tiny B -mode signal. This paper presents a detailed investigation of experimental effects which may impact *Spider*'s measurement of B -modes. The strategy is to simulate a *Spider* flight time stream injecting systematic effects in the time domain. The aim is to determine the level of B -mode contamination at subsequent stages of the analysis. With these results one can set stringent requirements on experimental design criteria, in addition to optimizing the telescope's observing strategy.

Several studies have been done which investigate the impact of systematics on CMB polarization measurements. Hu et al. (2003) and, more recently, O'Dea et al. (2007) use analytical methods to determine benchmark parameters for a variety of systematic effects. For these results there is no analysis in the map or time domains (as is done here); pixelization effects and time-varying systematics are ignored. Also, the noise (when included) is characterized as white; there is no accounting for the impact of $1/f$ noise. Much work has also been done that focuses primarily on the impact of systematics due to beam effects (Shimon et al. 2008; Rosset et al. 2007; Franco et al. 2003). The effects studied include optical aberration and various beam asymmetries (differential beam size, ellipticity, and cross-polar beam pattern). For *Spider*, the beam shape is determined not from feed horns but by the coherent interference of the detector antennas, and beams have been shown to be highly symmetric (see Fig. 2, discussed in § 2). The most worrisome beam effect comes from reflections in *Spider*'s refractive optics. Possible systematics from reflections are discussed in § 4.4.

The outline of this paper is as follows. Section 2 gives an overview of the instrument, flight, and observing strategy. Section 3 describes the details of the simulation methodology. Results for several systematic effects are presented in § 4. Section 5 concludes with a summary and discussion of the results.

2. THE INSTRUMENT

An initial description of *Spider* can be found in Montroy et al. (2006b). Since that publication, some of the telescope features have been changed in order to simplify design and further optimize the instrument. An overview of *Spider* instrumentation is given here.

Spider is a balloon-borne polarimeter designed to measure the polarization of the CMB at large angular scales. A clear advantage of a balloon platform is the increase in raw sensitivity, especially in the higher frequency channels, achievable above the Earth's atmosphere. For the long-duration balloon (LDB) flight, *Spider* will launch from Australia with an ~ 25 day, around-the-world, constant-latitude trajectory. The first test flight, of ~ 48 hr duration, is scheduled for fall 2009 from Alice Springs, Australia.

A schematic of the *Spider* payload is depicted in Figure 1 (*left*). The *Spider* gondola will spin in azimuth at a fixed elevation, observing only when the Sun is at least 10° below the horizon.¹² A constant-latitude 25 day flight, launching from Australia (with the optical axis tilted at 41° from the zenith), yields a sky coverage of $\sim 60\%$, as illustrated in Figure 4 (discussed in § 3.1).

Azimuthal attitude control is provided by a reaction wheel below the payload and a torque motor in the pivot located above the gondola. *Spider* will employ a number of sensors to obtain both short- and long-timescale pointing solutions. These include two star cameras, three gyroscopes, a GPS, and a three-axis magnetometer. The pointing system is based on proven BOOMERANG (Masi et al. 2006) and BLAST (Pascale et al. 2008) techniques. The pair of star cameras are mounted above the cryostat on a rotating platform, which will allow them to remain fixed on the sky, providing pointing reconstruction accurate to $\sim 6''$. Solar arrays pointing toward the Sun during daytime operation will recharge the batteries supplying payload power.

The instrument consists of six monochromatic telescopes operating from 96 to 275 GHz. All six telescope inserts are housed in a single LHe cryostat which provides >30 days of cooling power at 4 K (for the optics) and 1.5 K (for the sub-K cooler). The detectors are further cooled to 300 mK using simple ^3He closed-cycle sorption fridges, one per insert, which are cycled each day when the Sun prevents observations. Specifications for each of the six telescopes, including observing bands and detector sensitivities, are given in Table 1.

Spider uses antenna-coupled bolometer arrays cooled to 300 mK (Kuo et al. 2006). Figure 2 shows an image of a prototype detector and the measured beam response of a single dual-polarization antenna. The antenna arrays are intrinsically polarization-sensitive, with highly symmetrical beams on the sky and low sidelobes. Each spatial pixel consists of a phased array of 288 slot dipole antennas, with a radiation pattern defined by the coherent interference of the antenna elements. Each of the feed antennas provides an edge taper of roughly -13 ± 1 dB on the primary aperture. A single spatial pixel has orthogonally polarized antennas. The optical power incident on an antenna is transmitted to a bolometer and detected with a superconducting transition-edge sensor (TES) immediately adjacent to the spatial pixel.

The TESs will be read out using superconducting quantum interference device (SQUID) current amplifiers with time-domain multiplexing (Chervenak et al. 1999; de Korte et al. 2003; Reintsema et al. 2003; Irwin et al. 2004). Ambient-temperature multichannel electronics (Battistelli et al. 2008), initially developed

¹² An additional daytime (anti-Sun) scanning mode may be implemented but is not discussed here.




FIG. 1.—*Left:* *Spider* payload. Six independent monochromatic telescopes are housed in a single long hold time cryostat. Each telescope is fully baffled from radiation from the ground. Power is supplied by solar arrays. The baseline observing strategy is to spin the payload in azimuth at fixed elevation. *Spider* is designed to obtain maximum sky coverage during a 20–30 day, midlatitude, around-the-world flight. *Right:* *Spider* optical train. The telescope yields a flat and telecentric focal plane. The apodized Lyot stop, which is fixed with regard to the instrument, is maintained at 4 K. All dimensions are in millimeters. [See the electronic edition of the *Journal* for a color version of this figure.]

for SCUBA2 (Holland et al. 2006), will work in concert with the time-domain multiplexers.

The optical design for the inserts, shown in Figure 1 (*right*), is based on the BICEP/Robinson (Keating et al. 2003) optics. The monochromatic, telecentric refractor comprises two AR-coated polyethylene lenses and is cooled to 4 K in order to reduce the instrumental background to negligible levels. The primary optic is 302 mm in diameter, and the clear aperture of the Lyot stop is 264 mm, which produces a 45' beam at 145 GHz.

A cryogenic half-wave plate is located in front of the Lyot stop of each telescope. It consists of a single birefringent sapphire plate coated with a single layer of Herasil quartz on each side. Rotating or stepping the half-wave plate provides increased Q and U redundancy or improved polarization cross-linking in a pixel. In addition, with the half-wave plate located on the sky side of the primary, the beam remains fixed while the polarization rotates. This will eliminate false polarization signals due to

beam differences. Finally, rotated rapidly, the half-wave plate will modulate the polarization signal above detector $1/f$ noise.

Initially, polarization modulation was to be achieved via a *continuously spinning half-wave plate* (Montroy et al. 2006b). This work examines the viability of a *fast-spinning gondola* modulating the incoming signal with the half-wave plate stepping $22.5^\circ \text{ day}^{-1}$. Section 4.1 illustrates that either of these modes of operation can be used for *Spider*. The latter mode, stepping the half-wave plate, is easier to design mechanically and more robust to operate, and it is therefore preferred.

3. SIMULATION METHODOLOGY

The simulation pipeline is based largely on the analysis pipeline described in Jones et al. (2007), which was developed for the analysis of the data obtained from observations made with the BOOMERANG telescope during the 2003 LDB Antarctic flight (Montroy et al. 2006a; Piacentini et al. 2006; Jones et al. 2006;

TABLE 1
Spider CHANNEL SPECIFICATIONS

Obs. Band (GHz)	Orientation	Bandwidth (GHz)	Number of Detectors	NET ($\mu\text{K}\sqrt{\text{s}}$)	Beam FWHM (arcmin)
96.....	×	24	288	100	58
96.....	+	24	288	100	58
145.....	×	35	512	100	40
145.....	+	35	512	100	40
225.....	×	54	512	204	26
275.....	+	66	512	351	21

NOTES.—This table includes instrument orientation, observing bands, detector counts, sensitivities, and beams. Instrument orientation indicates the orientation of the polarization sensitivity of the orthogonal antenna pairs in a detector array for a given telescope insert. A total of 2624 detectors is distributed between six telescopes, with two operating at 96 GHz and two operating at 145 GHz.




FIG. 2.—*Left*: Single pixel of a 145 GHz antenna-coupled bolometer, comprising a 288 element phased array of dual-polarization slot antennas coupled to a matched load by a superconducting microstrip network. Microstrip filters, which determine the spectral response, and TES detectors, which measure the power dissipated in the load, are visible at bottom. *Right*: Measured beam pattern of the horizontal component of the dual-polarization antenna (the vertical beam pattern is similar). The upper limit on differential beam ellipticity is 1%, limited by the testbed. The polarization efficiency is greater than 98%. It is important to note that the beam pattern here is the *feed* beam pattern. The beam on the sky is influenced by the telescope. While the *Spider* telescope edge taper is modest, the beam on the sky will be more symmetric than the feed pattern shown here. In particular, the visibly large and asymmetric lobes above will not propagate to the sky. [See the electronic edition of the *Journal* for a color version of this figure.]

Masi et al. 2006). A schematic outlining the components of the simulation pipeline and the various inputs and outputs is given in Figure 3.

The flight simulator generates time-ordered pointing data in the form of right ascension, declination, and polarization angle for each detector. Data are simulated for 16 TES detectors (8 dual-polarization pixels), arranged in evenly spaced pairs in a single column which extends the full height of the focal plane. Detectors in a single pixel pair are sensitive to orthogonal polarizations. Since signal-only simulations are used in this work, this is sufficient to test most of the systematic effects and observing strategies considered here. The small number of time streams also significantly reduces the data storage and computation requirements, which themselves present a unique challenge for the actual analysis.

It is assumed that the telescope is fixed in elevation 41° from the zenith and that the payload is moving in longitude (beginning at 128.5° east) at a speed of 3.76×10^{-4} degrees per second (dps) at constant latitude (25.5° south). Data are simulated for 4 days of operation, assuming a mid-November launch. Four days of operation allows for one complete observing cycle for the stepped half-wave plate operating mode, after which the cycle is repeated. This is also the minimum required to ensure sufficient coverage for polarization reconstruction of the entire observed area.

Full-sky intensity and polarization maps are simulated and smoothed with the *Spider* beam using the *synfast* program, which is part of the HEALPix software (Górski et al. 2005). In order to ensure that signal variation within a pixel is negligibly small, the full-sky maps are pixelized at a resolution which corresponds to a pixel size of $\sim 3.4'$.

Full-sky maps are then converted into time-ordered data (TOD) using pointing information from the flight simulator. Thus, the time-stream generator constructs d_t for each detector from the equation

$$d_t = G \left\{ I_{\text{pix}} + \frac{\rho}{2 - \rho} [Q_{\text{pix}} \cos(2\psi_t) + U_{\text{pix}} \sin(2\psi_t)] \right\}. \quad (1)$$

Here I_{pix} , Q_{pix} , and U_{pix} are the Stokes values in the pixel observed at time t , ρ parameterizes the polarization efficiency, ψ is the final projection of the orientation of a detector on the sky, and G is the detector gain or responsivity.

All time streams are high-pass-filtered at 10 mHz during the map-making stage. This is done to test the impact of the filtering that is required in the case of real data, which are affected by long-timescale systematics. Particular systematics of concern are knowledge of system transfer functions (or, equivalently, knowledge of the gains) and knowledge of the noise amplitude/statistics on long timescales.

During the time-stream generation, the (stepped or spinning) half-wave plate polarization angle is added to the intrinsic polarization angle of the individual detectors. In addition, polarization angle systematics, beam offsets, and gain drift are also applied during time-stream generation. Additional pointing jitter and pendulation systematics are added to the pointing time streams during flight simulation. For the case of simulation of optical ghosting in the refractive optics, a pointing time stream is produced for each beam (the nominal beam and the reflection, or “ghost” beam), and the full-sky map is observed by each beam/pointing. TODs produced for each beam are combined with various weighting schemes in order to gauge the impact of the ghost-beam contamination.

Finally, *Spider* maps are constructed in terms of the observed Stokes parameters, I^{obs} , Q^{obs} , and U^{obs} , with an iterative map maker, an adaptation of the Jacobi method (described in detail in Jones et al. 2007) solving for the signal according to

$$\tilde{m} = (A^T N^{-1} A)^{-1} A^T N^{-1} d. \quad (2)$$

Here \tilde{m} is the general least-squares map triplet (I^{obs} , Q^{obs} , and U^{obs}), A is the pointing matrix which maps time-domain samples to pixels on the sky, and N^{-1} is the inverse noise filter. For *Spider* simulations the noise kernel is defined by a uniform white noise level, with a $1/f$ knee component at low frequencies. Although the simulations are for pure signal, the map-maker algorithm performs inverse noise filtering of the time streams. This filtering

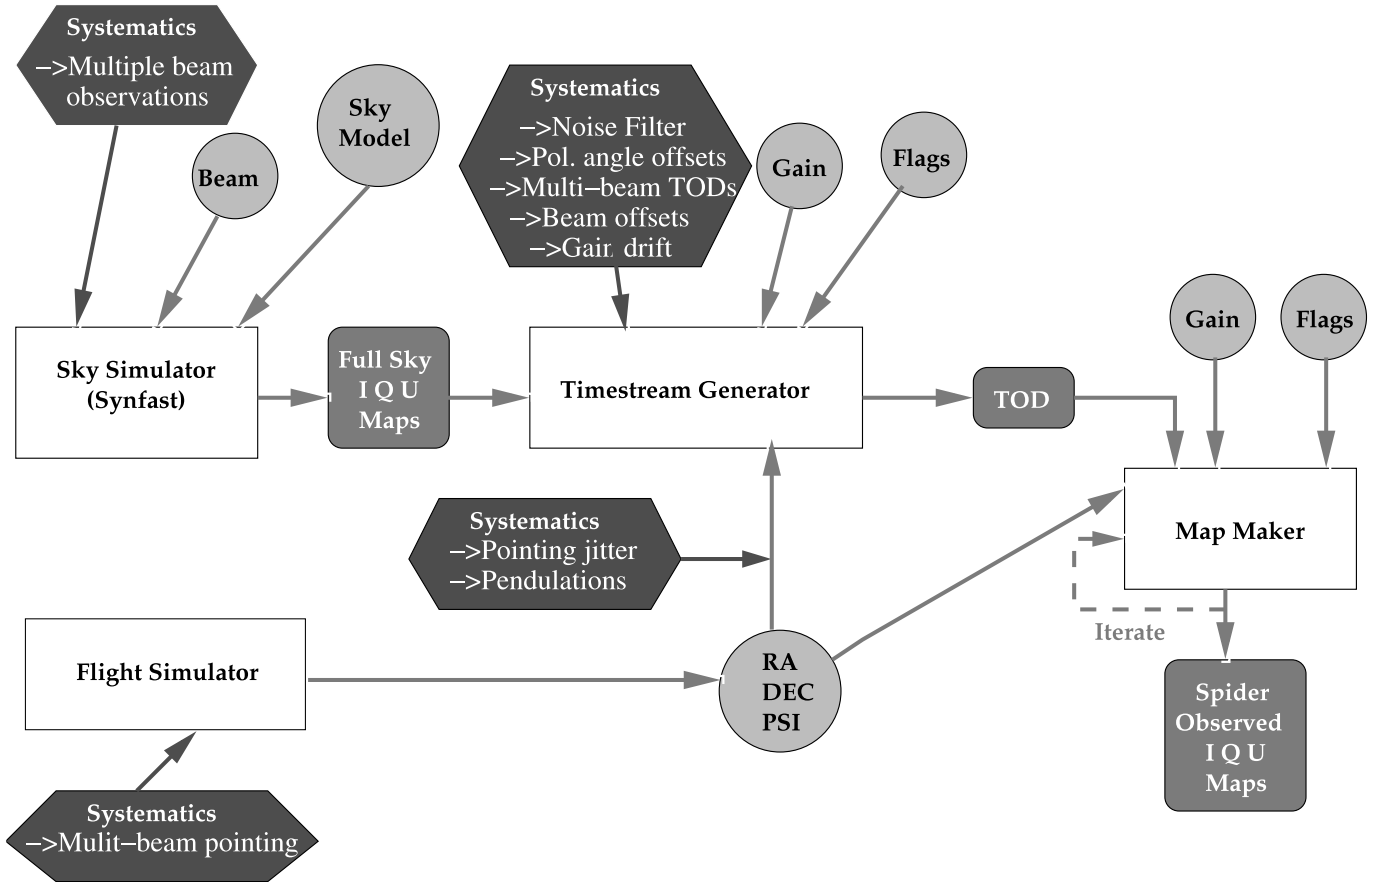


FIG. 3.— Schematic representation of the simulation pipeline. During the time-stream generation the (stepped or spinning) half-wave plate polarization angle is added to the intrinsic polarization angle (PSI) of the individual detectors. In addition, polarization-angle systematics, beam offsets, and gain drift are also applied during time-stream generation. Pointing jitter and pendulation systematics are added to the pointing time streams during flight simulation. For the case of simulation of optical ghosting in the refractive optics, a pointing time stream (multibeam pointing in this figure) is produced for each beam (the nominal beam and the reflection, or ghost, beam), and the full-sky map is observed by each beam/pointing (multibeam observation in the sky simulator). Finally, the time-stream generator produces time-ordered data (multibeam TODs) for each beam. The TODs are combined with various weighting schemes in order to simulate the ghost-beam contamination. [See the electronic edition of the Journal for a color version of this figure.]

would be included in noisy time streams to reduce the strongest effects of $1/f$ noise, which can significantly reduce map-making efficiency. In this case it is also included to make the simulations used here accurate representations of the full pipeline. To reduce computation time, output maps are pixelized at a resolution which corresponds to a pixel size of $\sim 13.7'$ (about one-fourth of the beam size).

3.1. Residual Measure

The aim is to quantify the contamination of the B -mode angular power spectrum (BB) from systematics which induce either $I \rightarrow Q, U$ or $Q \leftrightarrow U$ mixing. To assess the impact of the various systematics on BB, the following procedure is implemented:

1. Generate full-sky input I, Q , and U maps with $C_l^{\text{BB}} = 0$.
2. “Observe” the maps with the simulation pipeline including a chosen systematic (but no noise) giving signal-only $I^{\text{obs}}, Q^{\text{obs}}$, and U^{obs} .
3. Take the difference between the input and output maps over the survey area,

$$\begin{aligned} I^{\text{res}} &= I^{\text{obs}} - I, \\ Q^{\text{res}} &= Q^{\text{obs}} - Q, \\ U^{\text{res}} &= U^{\text{obs}} - U. \end{aligned} \quad (3)$$

4. Apply a mask with pixel weighting determined by the number of observations per pixel to the residual maps.

5. Spherical harmonic transform the weighted maps to obtain pseudo- C_l spectra of the BB residual, or

$$\tilde{C}_l^{\text{BBres (noBB)}} \quad (4)$$

6. Compute residual measure R_l^{BB} , defined below.

The residual map obtained in this way does not have uniform signal variance due to the filtering of the time stream and scanning of the input map. The reconstruction of polarization is also scan-dependent and breaks down at the edges of the map, where Q and U are degenerate because of a single crossing. The residual map is weighted by a nonuniform mask to reduce the effect in the final pseudo- C_l . The increased variance would not bias the ensemble averages of the pseudo- C_l but would dominate the scatter in the single realization if the mask were not applied. In actual data (or signal plus noise simulations), the mask weighting would also take into account the noise variance per pixel.

The mask that is applied to all of the simulations for which the gondola spin rate is 36 dps is shown in Figure 4. The dark band represents a Galactic cut at $\pm 10^\circ$ in Galactic latitude. The white region represents the portion of the sky that cannot be observed

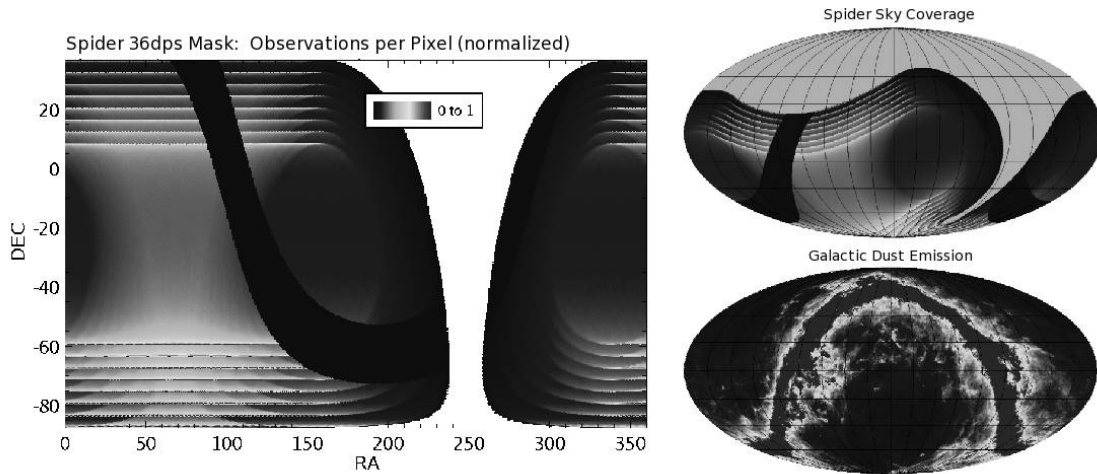


FIG. 4.—*Left:* Spider mask for simulations with a 36 dps gondola spin rate. The dark band represents a galactic cut at $\pm 10^\circ$ in galactic latitude, and the white region represents Sun flagging. With these regions flagged the fraction of the sky covered for this observing strategy is $\sim 60\%$. Pixel values are the number of observations in the pixel divided by the maximum hit value. The most obvious features are constant declination lines where scan circles on the sky for each detector overlap and the coverage is deepest. The pixel weighting is applied in order to reduce the effect of badly sampled pixels at the edge of the map. *Top right:* Spider coverage projected into equatorial coordinates. *Bottom right:* IRAS $100\ \mu\text{m}$ map (Schlegel et al. 1998) of galactic dust shown for comparison. [See the electronic edition of the Journal for a color version of this figure.]

because of the Sun. With these regions flagged, the fraction of the sky covered for this observing strategy is $\sim 60\%$. Each pixel value in the mask is the number of observations in the pixel divided by the value in the pixel with the maximum hits.

The spectra obtained from the method above are raw cut-sky, or pseudo- C_l , power spectra (Hivon et al. 2002). Since no B -mode power is present in the original full-sky simulation, any B -mode power in the final maps will be due to the mixing of modes from either systematics or cut-sky effects which mix E - and B -modes (Lewis et al. 2002; Bunn et al. 2003) or, indeed, T - and B -modes. The desired measure quantifies the size of the residual in the B -modes with respect to the expected original signal. To do this, the resulting power spectrum (eq. [4]) is divided by the spectrum obtained by a simulation with no input EE signal, $\tilde{C}_l^{\text{BB (noEE)}}$. This ratio is approximately independent of overall cut-sky effects, since both spectra are obtained from geometrically identical maps. The ratio is then multiplied by the original full-sky B -mode spectrum to obtain an approximate measure of the raw, experiment-induced B -mode power on the sky in μK^2 ,

$$R_l^{\text{BB}} = \frac{\tilde{C}_l^{\text{BBres (noBB)}}}{\tilde{C}_l^{\text{BB (noEE)}}} C_l^{\text{BB}}. \quad (5)$$

The residual measure defined above is not designed to give a complete picture of how well the original BB signal can be reconstructed from the observations. A complete treatment would require a full unbiased power spectrum estimation method, which is beyond the scope of this work. Instead, equation (5) isolates the impact of the systematics under study on the observed signal by minimizing the impact of the $E \rightarrow B$ mixing from cut-sky effects.

Note that for all of the input maps the same initial seed value is used to generate the full CMB sky; i.e., the sample scatter is the same for all simulations.

4. SIMULATION RESULTS

The presentation of results begins by illustrating the base residuals for two basic modes of half-wave plate operation: stepped and continuously rotating. For the remaining subsections, the B -mode residuals from experimental systematic effects for the

stepped half-wave plate case are examined. All simulations are for signal only (with no noise), but time streams are inverse noise-filtered during the map-making phase according to equation (2). For these simulations the noise kernel is defined by a uniform white noise level, with a $1/f$ knee component at low frequencies. Aside from § 4.2, which explores two knee-frequency values, the $1/f$ knee for the noise filter is 100 mHz for all simulations. In all plots the case labeled *nominal* is a 36 dps gondola spin rate, with the half-wave plate stepping 22.5° once per day, with 10 iterations (sufficient to recover the residual levels of the continuously rotating half-wave plate case) of the map maker, a Jacobi iterative solver (Jones et al. 2007).

4.1. Polarization Modulation

Since *Spider*'s default observing strategy is to spin continuously in azimuth, two modes of polarization modulation are explored. For the first case the half-wave plate spins continuously at 10 Hz, while the gondola rotates at 6 dps. For the second case the half-wave plate is stepped by $22.5^\circ \text{ day}^{-1}$, while the gondola rotates at 36 dps or more. Therefore, in the first case the half-wave plate is modulating the incoming polarization signal, and in the second case the gondola itself is used to modulate the signal.

Modulation by the gondola spin has a number of design advantages over the inclusion, in the optical train, of a continuously rotating half-wave plate. The half-wave plate adds a degree of complexity in the design with a subsequent impact on the robustness of the instrument. In addition, it is a potential source of a number of systematic effects, for example, microphonics, thermalization effects, magnetic pickup, and higher power dissipation at 4 K. It is therefore preferable (and nearly equivalent, as will be shown) to step the half-wave plate once per day, in order to increase Q and U redundancy in a single pixel while rapidly spinning the gondola in order to move the signal above the detector $1/f$ knee frequency.

Figure 5 shows Q residual maps for the two modulation modes. Maps of the U residuals are not shown here but are of similar amplitude. The top panel shows the residuals for the first case (continuous half-wave plate rotation at 10 Hz). For this case a “naive,” or zero-iteration, map is shown. The naive map is equivalent to a

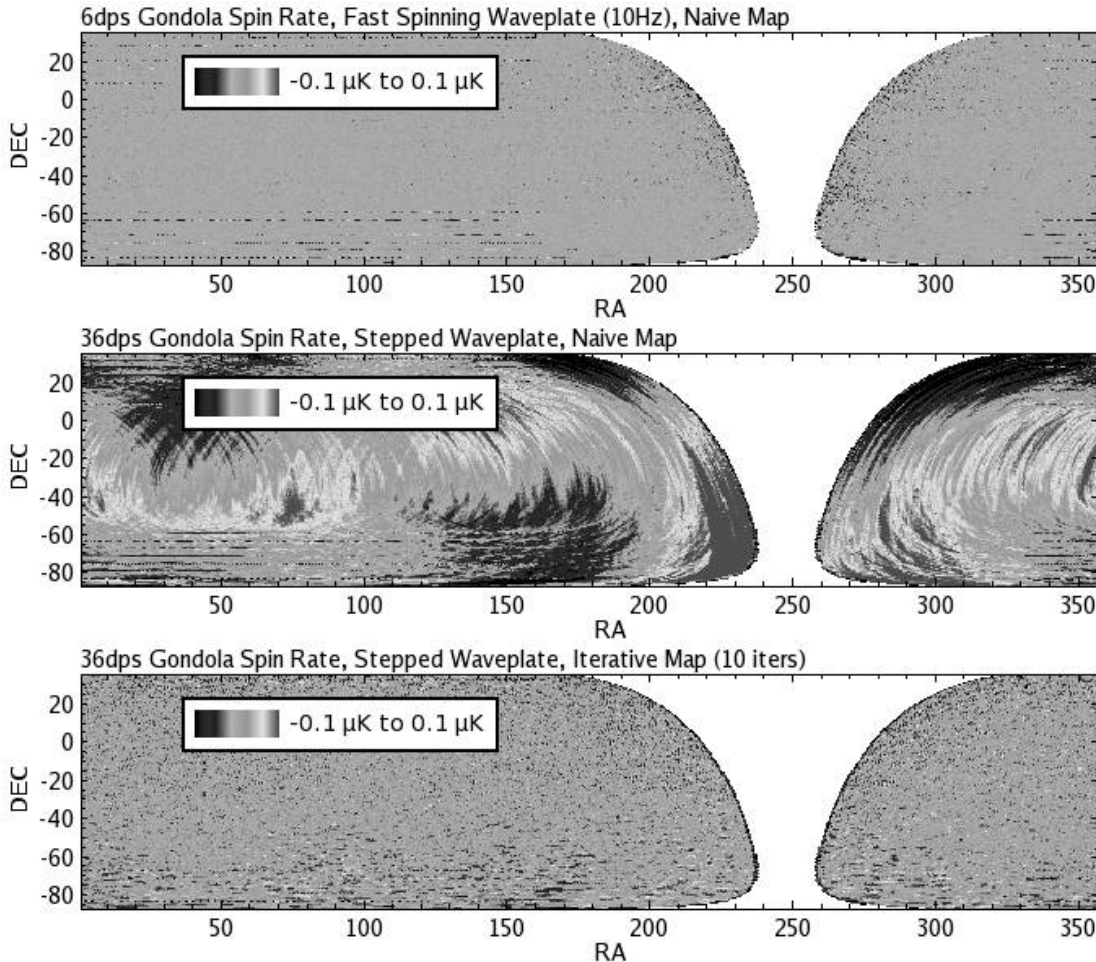


FIG. 5.—Maps of the residuals in the Q Stokes parameters for the two polarization modulation strategies. The top panel shows the residuals for the continuous half-wave plate rotation case. For this case a “naive,” or zero-iteration, map is shown. The middle panel shows the naive map for the stepped half-wave plate mode with the gondola spinning at 36 dps. In this case significant striping is present due to the loss of low-frequency modes. This is caused by the inverse noise filtering of the time streams during the map-making phase, which uses a noise kernel with a realistic 100 mHz $1/f$ knee. For the stepped case the polarization modulation is not sufficient. However, iterated map-making reduces the impact of the striping, as shown in the bottom panel, and 10 iterations of the map maker are sufficient to recover most of the lost modes. [See the electronic edition of the *Journal* for a color version of this figure.]

simple (pixel hit-weighted) binning of the time stream into pixels. Iterations of the map-making step are not required in this case, since the signal is modulated to frequencies higher than the expected $1/f$ knee; there is no signal degradation at low frequencies from the inverse noise filtering of time streams. This is one of the benefits of a design which includes a continuously rotating half-wave plate.

The middle panel of Figure 5 shows the naive map for the stepped half-wave plate mode with the gondola spinning at 36 dps. In this case significant striping is present because of the impact of inverse noise filtering of the time streams. Hence, for this case the real signal will be degraded on large scales because of the low-frequency $1/f$ response of the detectors. These translate to large-scale modes along the individual scans and result in the striping obvious in the maps. Iterating the map maker in this case reduces the effect of striping as the large-scale modes are recovered. After 10 iterations the striping is significantly reduced, as shown in the bottom panel of Figure 5. One possible way to reduce the computational load of a map-making stage with many iterations is to spin the gondola faster to modulate the signal into higher frequencies.

The power spectra for the BB residual measure (eq. [5]) are shown in Figure 6. The optimal solution is the continuous half-wave plate modulation scheme. This yields the lowest residual

compared to an $r = 0.01$ fiducial BB model. In the stepped 36 dps, noniterated case the residuals have the same amplitude as the model on the largest scales. The residuals are reduced to $<6\%$ levels for multipoles $l < 100$ when the map maker is iterated 10 times. The stepped 70 dps spin case with 10 iterations yields even smaller residuals at the largest scales.

Given the design and implementation advantages, the simple stepped half-wave plate system appears to be a feasible choice for the *Spider* scan strategy, albeit with significant additional computational costs.¹³ The remainder of this section is restricted to the stepped half-wave plate case. In particular, the focus is on probing whether any other systematic effects invalidate this choice of modulation scheme.

4.2. Noise

To examine the impact of different $1/f$ profiles on the stepped-mode residuals, a number of different cases were run:

1. 36 dps gondola spin rate with 100 mHz detector $1/f$ knee.

¹³ A full exploration of faster, suboptimal map-making algorithms is left for future work. In particular, destriping algorithms (see, e.g., Ashdown et al. 2007) may provide a much faster alternative, although it is still not clear that these can be applied to a *Spider* observing strategy and polarization-sensitivity requirements.




FIG. 6.—Comparison of the BB residual from continuously spinning and stepped half-wave plate polarization modulation schemes. For the first case the half-wave plate spins continuously at 10 Hz, while the gondola rotates at 6 dps. For the second case the half-wave plate is stepped by $22.5^\circ \text{ day}^{-1}$, while the gondola rotates at 36 dps or more. Operation in stepped half-wave plate mode, with the modulation provided by the spinning gondola, requires iterated map-making, as the signal is not modulated as far from the $1/f$ knee as in the continuously rotating half-wave plate mode. For $l > 20$, 10 iterations of the map maker are sufficient to recover the residual levels of the continuously rotating half-wave plate. Fewer iterations may be required if the gondola spin rate is even higher (70 dps). The input BB spectrum for $r = 0.01$ is plotted for comparison. [See the electronic edition of the Journal for a color version of this figure.]

2. 36 dps gondola spin rate with 500 mHz $1/f$ knee.
3. 110 dps gondola spin rate with 500 mHz $1/f$ knee.

The time-stream high-pass filter cutoff is kept at 10 mHz in all cases. A comparison of BB signal residuals varying the detector knee frequency is shown in Figure 7. Although the simulations are signal-only, the time streams are inverse noise-filtered, as would be done for the real data. This reveals the impact of the detector noise characteristics in terms of the degradation of the polarization signal on the largest scales. The effect of a 500 mHz knee is clearly seen on the largest angular scales. Even for 10 iterations of the map maker, the residuals are at the level of the BB model for $r = 0.01$. Increasing the spin rate to 110 dps reduces the impact of the higher $1/f$ knee and approaches the nominal 36 dps, 100 mHz knee case. A 500 mHz knee frequency for the detectors and readout electronics is pessimistic but would not be catastrophic, since polarization modulation can still be achieved by the faster spinning gondola. *Spider*'s high-frequency response is limited by the noise and response time of the detectors, the combination of which sets the maximum gondola spin rate. With 5 ms optical time-constant detectors, *Spider* can spin up to 110 dps before being affected by the detector noise and time constants. Thus, for the stepped half-wave plate case, the limit for the $1/f$ knee frequency is ~ 500 mHz.

4.3. Pointing Systematics

One of the more challenging aspects of balloon-borne telescope observations is pointing reconstruction, one of the limiting systematics in the interpretation of CMB data obtained from balloon-borne telescopes. For *Spider* the pointing system is based on largely on BOOMERANG (Masi et al. 2006) and BLAST (Pascale et al. 2008) instrument pointing. The former achieved $\sim 1.5'$ pointing, and the latter reconstructed pointing to a few arcseconds. The requirement for *Spider* will be much less demanding given the low resolution. Nonetheless, the design goal for *Spider* is subarcminute pointing reconstruction, and establishing a precise requirement is still important given the particular sensitivity of polarization measurements to offsets in the pointing.




FIG. 7.—Impact of a higher $1/f$ detector knee on the BB residuals. Simulations are for signal only (with no noise), but time streams are inverse noise-filtered in the map maker, giving a realistic estimate of modes that would be lost from $1/f$ effects in the real data. With the higher knee frequency of 500 mHz the gondola spin rate must be increased to reduce the residuals over the target range of multipoles, $l < 100$. The maximum gondola spin rate of ~ 110 dps is limited by the time response and noise characteristics of the *Spider* detectors. The input BB spectrum for $r = 0.01$ is plotted for comparison. [See the electronic edition of the Journal for a color version of this figure.]

The pair of star cameras will provide absolute pointing on ~ 1 Hz timescales. Integrated velocity data from fiber optic rate gyros will provide shorter timescale pointing information.

Two main pointing offsets are explored in this work. The first is a random pointing jitter added to the original pointing solution. The offsets are added to the right ascension (RA) and declination (DEC) value of each sample in the solution. The offsets are constant for 6 s, after which new random values are drawn. The 6 s bandwidth regime is chosen, since this will mimic pointing systematics which occur within the timescale of one gondola spin.¹⁴ A nominal run with $1'$ rms and a worst-case scenario with $10'$ rms for the instrument jitter are considered. To translate from instrument jitter to true jitter on the sky, a factor of $\sec(\text{DEC})$ is applied to the RA offsets.

The second systematic consists of a sinusoidal oscillation with an amplitude of $6'$ and a 20 minute period. This effect simulates the pendulation of the gondola. The pointing offsets examined are typical of in-flight conditions, albeit the $10'$ jitter is extremely pessimistic.

The results for the residual measure for the three cases are shown in Figure 8. The pendulation case also includes a long-timescale (1 day period), $1'$ rms jitter. This is intended to simulate a diurnal pointing offset caused by thermal flexure in the star camera mount. All cases except the large $10'$ rms jitter result in negligible contributions to the residual measure at $l < 100$ compared to the nominal stepped half-wave plate/36 dps spin mode without any systematic.

An additional pointing systematic affecting polarization measurements is the requirement to reconstruct the angle ψ (in eq. [1]) of the detector polarization relative to the fixed, local Q and U frame of reference on the sky. The systematic can arise in two distinct ways. The first is a relative offset between the ψ angles of different detectors. The second is an overall offset in the focal plane reference frame and the frame on the sky. The latter is generated by any error in the calibration of the polarization angle of the instrument.

¹⁴ White noise pointing jitter of similar amplitude that varied from sample to sample was also explored and was found to have an insignificant effect.

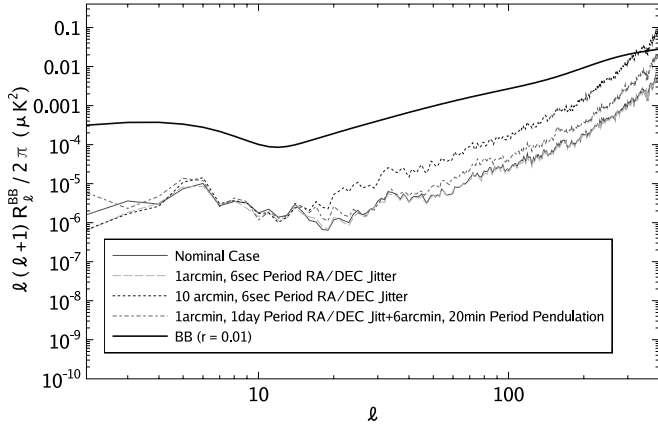


FIG. 8.—Three cases of pointing jitter. In the first two cases, offsets are added to the RA and DEC, with new offset values (with rms amplitudes $1'$ or $10'$) every 6 s. In the third case, a sinusoidal oscillation is implemented with an amplitude of $6'$ and a 20 minute period. This effect simulates the pendulation of the gondola. In addition, for the latter case, a long-timescale (1 day period) $1'$ rms jitter is added. Both types of pointing error, reconstruction error and in-flight pendulations, have negligible effects when compared with the BB $r = 0.01$ model. [See the electronic edition of the Journal for a color version of this figure.]

Results from simulation of the ψ systematics are shown in Figure 9. In the first two cases random 0.5° and 1° rms ψ errors are added to each detector. This simulates fixed, random offsets in the relative polarization angles of the detectors. The results show that the relative offsets contribute a comparable amount to the residuals as the nominal stepped-mode case.

In the remaining cases the same offset (0.25° , 0.5° , and 1°) is applied to all channels. This simulates an overall calibration error in the half-wave plate ψ angle. The results show that this systematic gives a much larger contribution to the residual. The 0.25° offsets produce an acceptable residual level, well below the model. Again, simulations consider only eight pairs in a single column, or 16 detectors total. The rms result for the full focal plane should average down as \sqrt{N} , where N is the number of detectors. This factor has not been applied to the result. The calibration of the half-wave plate ψ angles will be performed preflight on the ground. The level of subpercent precision required is not a difficult measurement and is made much easier by the compact optics and




FIG. 9.—Residuals from absolute and relative offsets of detector polarization angles. In the first two cases random 0.5° and 1° rms ψ errors are added to each detector. Residuals for these relative offsets are negligible. In the remaining cases the same offset (0.25° , 0.5° , and 1°) is applied to all channels, simulating an overall calibration error in the half-wave plate ψ angle. The 0.25° offsets bring the residual level a factor of ~ 10 below the BB $r = 0.01$ model. [See the electronic edition of the Journal for a color version of this figure.]

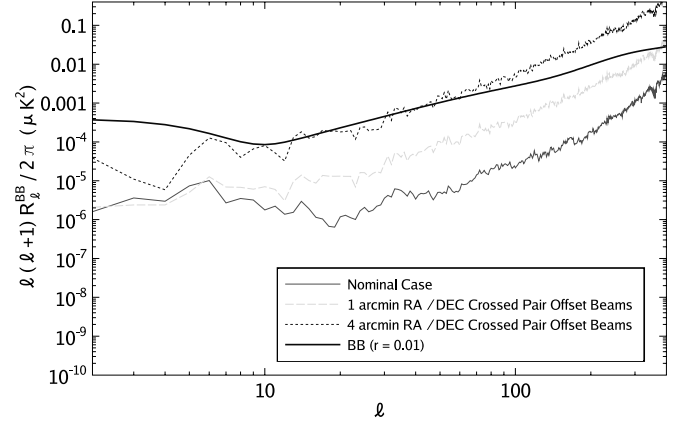


FIG. 10.—For these simulations one detector in a crossed pair (of the eight pairs) has a constant offset in RA and DEC of $1'$ or $4'$. For *Spider* $40'$ beams, the corresponding A-B amplitudes are 3.6% and 14% for $1'$ and $4'$ offsets. Residuals for beam offsets are plotted in the figure. The effect is well below the $r = 0.01$ model for the $1'$ offset case. [See the electronic edition of the Journal for a color version of this figure.]

correspondingly close “far field” of *Spider*. For the BICEP experiment (Yoon et al. 2006), polarization angles are determined within 1° . This systematic uncertainty is based on the way the polarizer is mounted during measurements, which can be improved. Preliminary measurements in the lab show that polarization angles can be measured to $< 0.125^\circ$.

Another systematic that is tested concerns pointing offsets of crossed pairs. This occurs if the E -field distribution is not identical at the feed or if there are polarization-dependent properties in the optics. For these simulations one TES in a single-pixel orthogonal polarization pair (of the eight pairs) has a constant offset in RA and DEC of $1'$ or $4'$. Again, for the case of RA a factor of $\sec(\text{DEC})$ is applied to the offset and hence represents the true offset on the sky. For *Spider* $40'$ beams, the corresponding A-B amplitudes are 3.6% and 14% for $1'$ and $4'$ offsets. Residuals for beam offsets are plotted in Figure 10. The requirement of subarcminute knowledge of beam centroids should be sufficient for $l < 30$ and sub- $0.5'$ for a larger range of l . This result can be compared with the result in Hu et al. (2003), who found that pointing errors between orthogonal polarization detectors should be less than a percent of the Gaussian beam width.

4.4. Beam Systematics

The *Spider* antenna array and optics define highly symmetric beams on the sky. The beam pattern shown in Figure 2 is the *feed* beam pattern. The beam on the sky is influenced by the telescope. While the *Spider* telescope edge taper is modest, the beam on the sky will be more symmetric than the feed pattern shown here. In particular, the visibly large and asymmetric lobes above will not propagate to the sky. The largest amplitude beam effect expected comes from reflections, or “ghosting,” in the *Spider* optics. Ghosting is common in refractive optics and results from unintended multiple reflections in the optics. The effect is a smaller amplitude beam image which is mirrored with respect to the pixel position from the center of the focal plane.

Ghosting is simulated by summing two time streams from two beams. One of the time streams is constructed using offset pointing from the ghost beam. The offset pointing of the ghost is determined in instrument coordinates (elevation and azimuth), hence the ψ angle for the ghost pointing is calculated appropriately as the final projection of the orientation of a detector on the sky. The second time stream is constructed from the pointing from the




FIG. 11.—Residuals from beam ghosting in the *Spider* optics. The ghosting is added in terms of a fractional contamination of the nominal time stream. Final maps are then reconstructed assuming no reflection contamination with no attempt to correct for the image distortion. A 10% contamination in the TOD yields a BB fractional residual higher than the model at some values of l . For 5% ghosting in the TOD this effect is already down by more than half, and it is negligible for 1% contamination in the TOD. [See the electronic edition of the *Journal* for a color version of this figure.]

nominal beam. The two time streams are summed with various weighting schemes depending on the ghost-beam contamination.¹⁵

The residuals from beam ghosting are summarized in Figure 11. For each case the ghosting effect is added to the time stream in terms of a fractional contamination added to the nominal time stream. Final maps are then reconstructed assuming no reflection contamination; there is no attempt to correct for image distortion. With 10% contamination in the TOD, the effect in the residual is as high as the BB $r = 0.01$ model. For 5% ghosting in the TOD, this effect is already down by more than half, and it is negligible for 1% in the TOD. Since simulations are done only for a single row (of eight pairs of detectors), the distance between the original and ghost images is, on average, smaller for this row than for any other. It is therefore worth noting that the full focal plane may show a larger effect than that simulated here.

4.5. Calibration Drift

Diurnal variations in the detector sensitivity will occur due to altitude-induced changes in background loading. These sensitivity changes will be tracked using 4 K semiconductor emitters (fired intermittently) similar to those used in BOOMERANG flights. For the BOOMERANG 2003 flight, the gain drift for each individual detector was measured and removed. Residual drift in the corrected data is less than the limit of uncertainty on an individual calibration pulse, 0.05%. With an improved version of the calibration lamp and *Spider*'s higher sensitivity, the expected uncertainty is $\sim 0.01\%$. Knowledge of the detector model (determined in preflight testing) will also allow the calculation of the sensitivity for any given operating point.

Figure 12 shows the effect of uncorrected calibration drift. Two cases are considered. For the first case the calibration drift is the same for all detectors, changing on a diurnal timescale with a maximum amplitude of 3%. For this case the resulting residual is small, less than 10% at all scales.

For the second case the gain drift for all detectors is of the same amplitude and has a 24 hr period, but each of the 16 detectors has a gain drift with a different phase. Thus, for any given

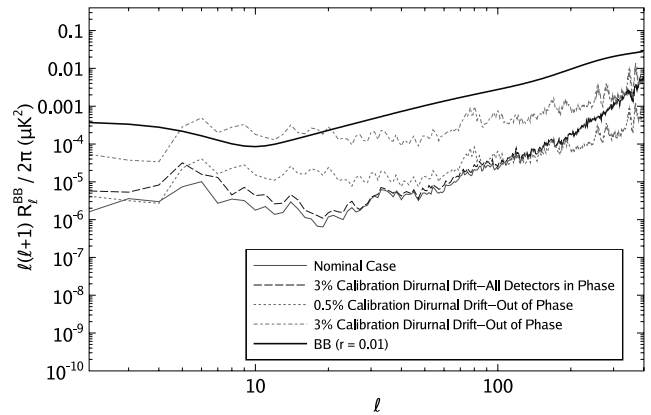


FIG. 12.—Effect of uncorrected calibration drift. In the first case the calibration drift is the same for all detectors, changing on a diurnal timescale with a maximum amplitude of 3%. For the second case the gain drift for all detectors is of the same amplitude and diurnal but out of phase. For this case the BB fractional residual is as much as 50% for a gain drift amplitude of 3% and drops below 5% for a 0.5% amplitude gain drift. Simulations consider only eight pairs in a single column, or 16 detectors total. The rms result for the full focal plane will average down as \sqrt{N} , where N is the number of detectors. For the second case (out-of-phase gain drifts), a factor of $\sqrt{16}/\sqrt{1024}$ (assuming 1024 CMB science channels) has been applied to each residual result. [See the electronic edition of the *Journal* for a color version of this figure.]

time sample the calibration factor within a pair of detectors will be different. Simulations consider only eight pairs in a single column, or 16 detectors total. Again, the rms result for the full focal plane should average down as \sqrt{N} , where N is the number of detectors. For this case a factor of $\sqrt{16}/\sqrt{1024}$ (assuming 1024 CMB science channels) has been applied to each residual result. The BB fractional residual is as high as the model for a gain drift amplitude of 3%. For a 0.5% amplitude drift the residual drops below the 30% level. As with all simulations in this work, there has been no attempt to correct for the gain systematic. For BOOMERANG 2003, the final relative calibration uncertainty was 0.4% (Masi et al. 2006). *Spider* is expected to achieve an uncertainty of 0.1% or less, which will be more than adequate to meet science goals.

5. CONCLUSIONS

The results from § 4 are summarized in terms of experimental specifications in Table 2. While results are *Spider*-specific, the order of magnitude of various effects can be translated to other CMB polarization experiments. The rms B -mode signal for $r = 0.01$ is roughly 1000 nK^2 and scales linearly with r . Experimental specifications are set by limiting the allowed systematic residual level to a factor of ~ 10 smaller than the B -mode signal for $r = 0.01$.

While the simulations were signal-only, the impact of large, low-frequency detector noise ($1/f$ noise) is reflected in the large-scale degradation of the B -mode signal for the stepped half-wave plate mode of operation. Rapid, continuous half-wave plate modulation mitigates this effect entirely. Rapid, continuous gondola rotation also works, but only with iterative map-making which accurately recovers the larger scale signal.

It is important to note that the effects studied in §§ 4.1 and 4.2 (naive instead of iterated maps, spinning more slowly, and stepping half-wave plate instead of spinning half-wave plate) will degrade the signal-to-noise ratio achieved on the bandpowers. These effects differ from the systematics studied in §§ 4.3–4.5 (pointing reconstruction errors, polarization angle uncertainty, uncorrected ghosting, and uncorrected gain drifts), which will ultimately bias

¹⁵ In order to completely model the polarization of the ghost, a half-wave plate angle dependency should be included in the ratio of ghost amplitude to main-beam response. This is not done here but is left for future work.

TABLE 2
SUMMARY OF EXPERIMENTAL SPECIFICATIONS BASED ON SIMULATION RESULTS

Systematic	Experimental Spec.	Comments
Receiver $1/f$ knee.....	<200 mHz	For 110 dps gondola spin
Receiver $1/f$ knee.....	<100 mHz	For 36 dps gondola spin
Pointing jitter.....	<10'	Sufficient for $l < 50$
Absolute pol. angle offset.....	<0.25°	...
Relative pol. angle offset.....	<1°	...
Knowledge of beam centroids.....	<1'	Sufficient for $l < 30$
Optical ghosting.....	<2%	Contamination
Calibration drift.....	<3.0%	In phase
Calibration drift.....	<0.1%	Out of phase

NOTES.—Realistic-amplitude, time-varying systematics are injected into the simulated time streams. Maps are reconstructed without any attempt to correct for the systematic errors. Experimental specifications are set by limiting the allowed systematic residual level to a factor of ~ 10 smaller than the B -mode signal for $r = 0.01$. The nominal operating mode is a 36 dps gondola spin rate, with the half-wave plate stepping 22.5° once per day, with 10 iterations (sufficient to recover the residual levels of the continuously rotating half-wave plate case) of the map maker, a Jacobi iterative solver (Jones et al. 2007).

the final result. The requirements of the biasing effects are more difficult to treat than the signal-to-noise ratio issues.

The impact of systematics on B -mode polarimeter experiments is also discussed in Hu et al. (2003) and, more recently, O'Dea et al. (2007), where analytical methods are used for calculating the B -mode spectrum bias. The results are useful for setting experimental “benchmark parameters” at the very earliest phases of instrument design. This work goes a step further by considering the impact of systematics in the map/time domain, a necessary

step in the evolution of an experiment which aims to measure the tiny, primordial gravity-wave signal.

This research used the McKenzie cluster at CITA, funded by the Canada Foundation for Innovation. Some of the results in this paper have been derived using the HEALPix package (Górski et al. 2005), as well as the FFTW package (Frigo & Johnson 2005). We thank the anonymous referee for helpful comments and corrections.

REFERENCES

- Adams, F., Bond, J. R., Freese, K., Frieman, J., & Olinto, A. 1993, *Phys. Rev. D*, 47, 426
- Ade, P., et al. 2008, *ApJ*, 674, 22
- Ashdown, M. A. J., et al. 2007, *A&A*, 467, 761
- Battistelli, E., et al. 2008, *J. Low Temperature Phys.*, 151, 908
- Bock, J., et al. 2006, preprint (astro-ph/0604101)
- Bunn, E. F., Zaldarriaga, M., Tegmark, M., & de Oliveira-Costa, A. 2003, *Phys. Rev. D*, 67, 023501
- Chervenak, J. A., Irwin, K. D., Grossman, E. N., Martinis, J. M., Reintsema, C. D., & Huber, M. E. 1999, *Appl. Phys. Lett.*, 74, 4043
- de Korte, P. A. J., et al. 2003, *Rev. Sci. Instrum.*, 74, 3807
- Franco, G., Fosalba, P., & Tauber, J. A. 2003, *A&A*, 405, 349
- Frigo, M., & Johnson, S. G. 2005, *Proc. IEEE*, 93, 216
- Goldin, A., Bock, J. J., Hunt, C., Lange, A. E., Leduc, H., Vayonakis, A., & Zmuidzinas, J. 2002, in *AIP Conf. Proc. 605, Low Temperature Detectors*, ed. F. S. Porter et al. (Melville: AIP), 251
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *ApJ*, 622, 759
- Harry, G. M., Fritschel, P., Shaddock, D. A., Folkner, W., & Phinney, E. S. 2006, *Classical Quantum Gravity*, 23, C01
- Hedman, M. M., Barkats, D., Gundersen, J. O., McMahon, J. J., Staggs, S. T., & Winstein, B. 2002, *ApJ*, 573, L73
- Hivon, E., Górski, K. M., Netterfield, C. B., Crill, B. P., Prunet, S., & Hansen, F. 2002, *ApJ*, 567, 2
- Holland, W., et al. 2006, *Proc. SPIE*, 6275, 62751E
- Hu, W., Hedman, M. M., & Zaldarriaga, M. 2003, *Phys. Rev. D*, 67, 043004
- Irwin, K. D., et al. 2004, *Nucl. Instrum. Methods Phys. Res. A*, 520, 544
- Jones, W. C., et al. 2006, *ApJ*, 647, 823
- . 2007, *A&A*, 470, 771
- Kallos, R. 2007, preprint (hep-th/0702059)
- Keating, B. G., Ade, P. A. R., Bock, J. J., Hivon, E., Holzapfel, W. L., Lange, A. E., Nguyen, H., & Yoon, K. W. 2003, *Proc. SPIE*, 4843, 284
- Kogut, A., et al. 2007, *ApJ*, 665, 355
- Kovac, J., et al. 2002, *Nature*, 420, 772
- Kuo, C. L., et al. 2006, *Proc. SPIE*, 6275, 62751M
- Lewis, A., Challinor, A., & Turok, N. 2002, *Phys. Rev. D*, 65, 023505
- Liddle, A. R., & Lyth, D. H. 2000, *Cosmological Inflation and Large-Scale Structure* (Cambridge: Cambridge Univ. Press)
- Linde, A. D. 1983, *Phys. Lett. B*, 129, 177
- . 1994, *Phys. Rev. D*, 49, 748
- Linde, A., Mukhanov, V., & Sasaki, M. 2005, *J. Cosmol. Astropart. Phys.*, 10, 2
- Lyth, D. H. 1997, *Phys. Rev. Lett.*, 78, 1861
- MacTavish, C. J., et al. 2006, *ApJ*, 647, 799
- Masi, S., et al. 2006, *A&A*, 458, 687
- Montroy, T. E., et al. 2006a, *ApJ*, 647, 813
- . 2006b, *Proc. SPIE*, 6267, 62670R
- Myers, M. J., Lee, A. T., Richards, P. L., Schwan, D., Skidmore, J. T., Smith, A. D., Spieler, H., & Yoon, J. 2002, in *AIP Conf. Proc. 605, Low Temperature Detectors*, ed. F. S. Porter et al. (Melville: AIP), 247
- O'Dea, D., Challinor, A., & Johnson, B. R. 2007, *MNRAS*, 376, 1767
- Page, L., et al. 2007, *ApJS*, 170, 335
- Pascale, E., et al. 2008, *ApJ*, 681, 400
- Phinney, E. S., et al. 2005, *Phys. Rev. D*, 72, 3005
- Piacentini, F., et al. 2006, *ApJ*, 647, 833
- Readhead, A. C. S., et al. 2004, *Science*, 306, 836
- Reintsema, C. D., et al. 2003, *Rev. Sci. Instrum.*, 74, 4500
- Rosset, C., Yurchenko, V., Delabrouille, J., Kaplan, J., Giraud-Héraud, Y., Lamarre, J.-M., & Murphy, A. J. 2007, *A&A*, 464, 405
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
- Seljak, U., & Zaldarriaga, M. 1997, *Phys. Rev. Lett.*, 78, 2054
- Shimon, M., Keating, B., Ponthieu, N., & Hivon, E. 2008, *Phys. Rev. D*, 77, 083003
- Spergel, D. N., et al. 2007, *ApJS*, 170, 377
- Yoon, K. W., et al. 2006, *Proc. SPIE*, 6275, 62751K