LiteBIRD TOAST Simulations for Scan and Instrument Optimization

(Internal White Paper)

T. Kisner[1,2], J. Borrill[1,2], S. Kernasovskiy[3,4], and R. Keskitalo[1,2]

[1] Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, CA, 94707, USA

[2] Space Science Laboratory, University of California, Berkeley, CA, 94707, USA

[3] Kavli Institute for Particle Astrophysics and Cosmology, SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

[4] Department of Physics, Stanford University, Stanford, CA 94305, USA

January 23, 2017

ABSTRACT

Aims. We use several metrics and constraints based on previous experience with CMB satellite data processing to optimize the LiteBIRD scan strategy and evaluate the necessity of a Half Wave Plate (HWP).

Methods. Time domain pointing simulations with one full frequency (140GHz) of the proposed LiteBIRD focalplane are used to evaluate the parameter space of precession and spin angles and rates. Constraints are applied to this parameter space based on minimizing the deformation of the effective beam and sufficient sampling of the beam on the sky. The best scan strategy both with and without the HWP is chosen based on uniformity of coverage, pixel condition number, and the ability to calibrate the instrument using the solar sytem and orbital dipole signals. After selecting the best scan strategy based on these criteria, an ensemble of noise simulations are performed to estimate the level of noise present on the measured angular power spectra.

Results. We find that regardless of HWP rotation rate, the best precession opening angles are 35° and 55°. Using noise Monte Carlos we show that the case without a HWP has significantly more noise power at low-ℓ in $\langle BB \rangle$ than all cases with a perfect HWP.

Key words. LiteBIRD – CMB – TOAST

1. Overview

Many factors impact the final science achievable by LiteBIRD. Previous forecasts used approximate numbers for array sensitivity at each frequency, white noise angular power spectra, and mock component separation performance to arrive at high-level estimates for LiteBIRD's final measurements of the cosmological B-mode signal. In this study, we assume the focalplane sensitivity and layout used in these earlier estimates, but we now investigate how the choice of scan strategy impacts our ability to make frequency maps of the sky which have uniform coverage, good separation of temperature and polarization, stable effective beams along the scan direction, and sufficient dipole signal to perform relative calibration between detectors.

The focalplane properties and geometry used here is the nominal one from the CSR. Table 1 lists the detector properties for all frequencies.

Figure 1 shows the configuration of the Low Frequency Telescope (LFT) and Figure 2 shows the configuration of the High Frequency Telescope (HFT).

Table 1. Detector properties assumed for this study, listed by central frequency in GHz. The slope parameter of the noise model was 1.5 for all detectors.

Property	40	50	60	68	78	89	100	119	140	166	195	235	280	337	402
NET (μK$\cdot\sqrt{Hz}$)	266.9	170.8	140.9	115.3	96.1	83.3	104.9	76.0	67.1	64.1	67.1	76.0	99.6	147.1	282.8
f_{knee} (Hz)	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05
FWHM (arcmin)	69.0	56.0	48.0	43.0	39.0	35.0	29.0	25.0	23.0	21.0	20.0	19.0	24.0	20.0	17.0

Fig. 1. Full LiteBIRD LFT focalplane. The blue shaded wafers have trichroic pixels which include the 140GHz detectors used in this study.

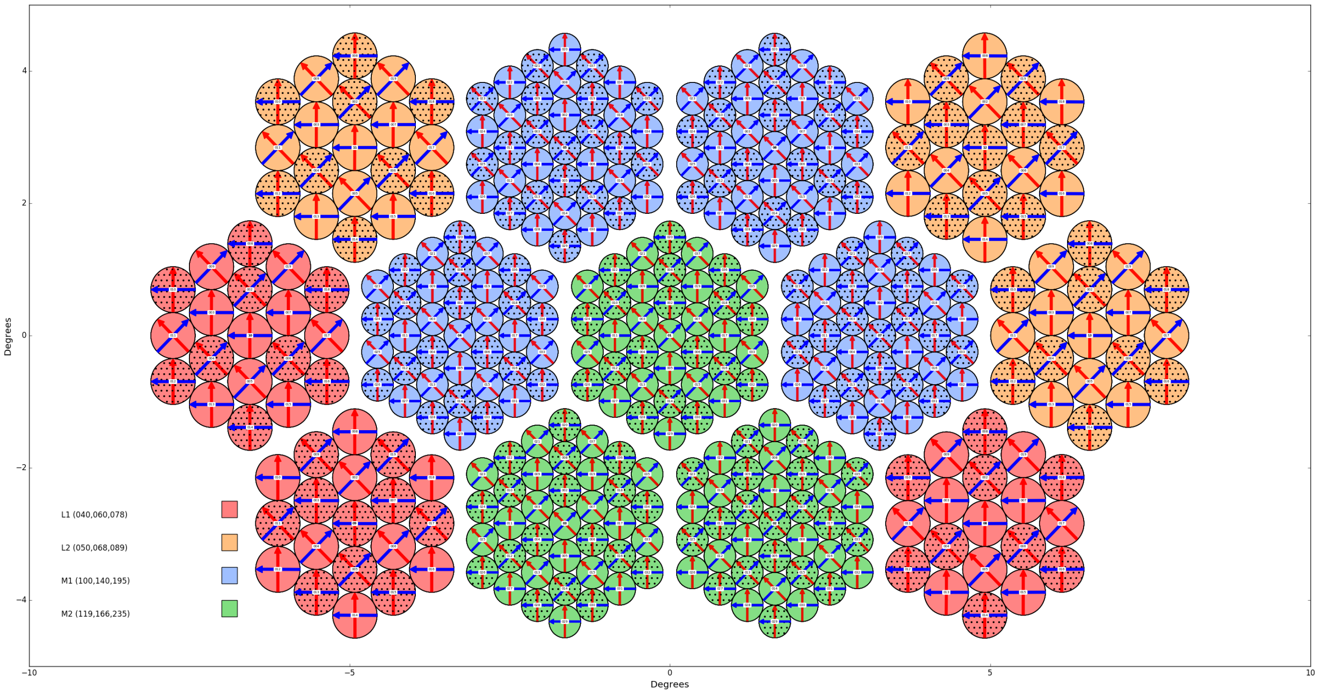

We carry out these studies using only a single frequency (140GHz). We expect all frequencies to be similarly impacted by the scan strategy. However, when using the beam FWHM to set constraints on the scan parameters, we use the smallest beam on the focalplane (17”). Our basic approach to scan strategy optimization consists of several steps:

1. Sweep across the range of allowed precession angles and use contraints driven by the HWP rate, samplerate, and effective beam deformation to set the precession rate and the spin angle and rate.
2. Evaluate pointing-related quantities such as the uniformity of the hit map, pixel condition number and dipole amplitude. Determine the best configurations both with and without the HWP.
3. For these best cases, generate an ensemble of noise Monte-Carlo maps in order to estimate the noise angular power spectra.

2. Pointing Simulations

For the first set of tests, we use the 140GHz nominal focalplane and the NET values listed in table 1. We assume a sample rate of 23.0Hz, since that was used for earlier studies and is presumably set by the downlink bandwidth and compression. Once the sample rate is fixed, we scan the range of precession opening angles from 10-85° (avoiding the extrema of the allowed range 0-95). The sum of the spin and precession angles must total 95°, so now the spin angle is also fixed. To set the precession rate, we must precess slow enough that the effective beam in the scan direction does not change much over the course of the scan. The speed of a detector on the sky depends on both the spin and precession motions. When the spin angle is aligned with the precession angle, then these add

Fig. 2. Full LiteBIRD HFT focalplane.

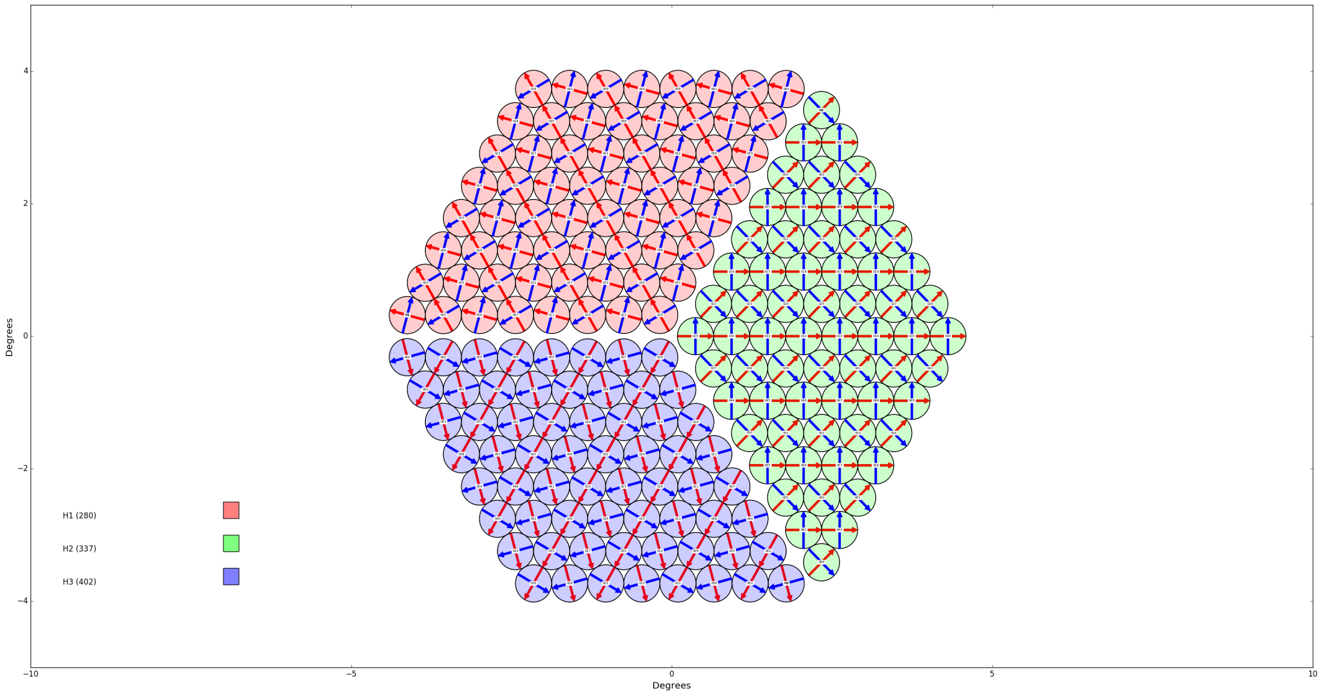

constructively. The minimum and maximum speeds on the sky of the focalplane boresight are

$$
\begin{aligned}
\alpha &= \text{Precession opening angle (rad)} \\
p &= \text{Precession rate (rad/s)} \\
\beta &= \text{Spin opening angle (rad)} \\
s &= \text{Spin rate (rad/s)} \\
v_{\text{min}} &= \beta s - \alpha p \qquad (1) \\
v_{\text{max}} &= \beta s + \alpha p \qquad (2)
\end{aligned}
$$

and so the maximum variation in speed for a particular scan strategy is

$$v_{\text{var}} = v_{max} - v_{min} = 2\alpha p. \qquad (3)$$

We require that this speed variation integrated over one sample be a small fraction of the beam FWHM:

$$
\begin{aligned}
B &= \text{Beam FWHM (rad)} \qquad (4) \\
\frac{2\alpha p}{f_{\text{sample}}} &= X \cdot B \qquad (5)
\end{aligned}
$$

The original nominal scan strategy corresponded to $X \approx 2.2\%$, so that is what we use for this study. This could be adjusted based on further beam studies. This sets an upper limit on the precession rate:

$$p = \frac{X \cdot B \cdot f_{\text{sample}}}{2\alpha} \qquad (6)$$

Without a HWP, we want to cover the sky as quickly as possible to gain as many attack angles as possible on each pixel. However, the spin rate is always constrained by the sample rate and the requirement that we want to have on average 3 hits per pixel. Typically we use a pixelization that is about 1/3 of the beam FWHM. So for these tests we require 9 samples per beam FWHM on the sky:

$$s_{(\text{No HWP})} \cdot \beta = \frac{B \cdot f_{\text{sample}}}{9} \qquad (7)$$

For studies *with* a HWP, we additionally want to ensure that the HWP angle is modulated by 45° over the course of one pixel crossing. This requirement allows us to resolve the polarization in the pixel with one crossing and to (hopefully) disentangle any HWP-induced systematics. This requirement sets an upper limit on the spin rate that gets smaller as the HWP is slowed down:

$$\begin{aligned} H &= HWP_{\text{rate}}\ (\text{rad/s}) = \frac{2\pi\,(HWP_{\text{rpm}})}{60} \\ s_{(\text{With HWP})} \cdot \beta &= \frac{B \cdot H}{3} \end{aligned} \tag{8}$$

With a HWP, we choose the slower of the two spin rate requirements above. For this study, we sweep over the precession angle in 5° increments. For each angle, we fix the other parameters using the constraints just described. For the simulations, we keep the precession axis pointed in the anti-sun direction in a constant slew. There are no discrete steps of the precession axis direction. We also simulate 28 hour observations containing 24 hours of science data collection and 4 hours of unusable cooler cycling. Table 2 shows the resulting parameters for the case without a HWP, and also for the case of spinning the HWP at 88 and 44 RPM.

Table 2. Scan properties used for this study. Precession and spin periods are in minutes. The free parameters for each case are the HWP rotation speed and the precession opening angle (α). Note that the case of an 88 RPM HWP rotation is fast enough that its spin rate is limited by equation (7) rather than equation (8).

		No HWP		HWP = 44 RPM		HWP = 88 RPM	
α	β	Prec. Period	Spin Period	Prec. Period	Spin Period	Prec. Period	Spin Period
10	85	14.61	12.29	14.61	20.45	14.61	12.29
15	80	21.91	11.57	21.91	19.25	21.91	11.57
20	75	29.22	10.85	29.22	18.05	29.22	10.85
25	70	36.52	10.12	36.52	16.84	36.52	10.12
30	65	43.83	9.40	43.83	15.64	43.83	9.40
35	60	51.13	8.68	51.13	14.44	51.13	8.68
40	55	58.44	7.95	58.44	13.24	58.44	7.95
45	50	65.74	7.23	65.74	12.03	65.74	7.23
50	45	73.04	6.51	73.04	10.83	73.04	6.51
55	40	80.35	5.79	80.35	9.63	80.35	5.79
60	35	87.65	5.06	87.65	8.42	87.65	5.06
65	30	94.96	4.34	94.96	7.22	94.96	4.34
70	25	102.26	3.62	102.26	6.02	102.26	3.62
75	20	109.57	2.89	109.57	4.81	109.57	2.89
80	15	116.87	2.17	116.87	3.61	116.87	2.17
85	10	124.17	1.45	124.17	2.41	124.17	1.45

When simulating LiteBIRD scannning strategies, there are several metrics we can use to evaluate the relative effectiveness of a particular configuration:

- the overall uniformity of coverage (RMS of the hitmap).
- the condition number of the 3x3 I/Q/U noise covariance at each pixel (which depends only on the pointing and the NET).
- the amplitude of the CMB dipole in the time domain, which we use to determine the relative gains between detectors.

The first two items are straightforward, but the last item is often not appreciated. Even with active temperature control, every time we cycle the cooler, re-tune detectors, or otherwise change the state of the system, we may have small changes in the relative gains. The metric we use in this study is the maximum variation in the dipole signal during one cooler cycle. This number is a proxy for how well we can determine the relative gains over one cooler cycle where we expect the system to be stable. Figure 3 shows how the hit map RMS changes with scan strategy. Note that since the precession rates are the same for all cases of HWP and the spin rates are similar, the overall hit map standard deviations for all cases are very close. We can see that there are two minima, at 35 and 55 degrees. Figures 4, 5, and 6 shows how much dipole variation we have across a year of data at the focalplane boresight. Each datapoint is the value for a single cooler cycle.

Fig. 3. Hit map RMS as a function of precession opening angle. All cases use the same angles and same precession rates, and the 0 and 88 RPM cases also have identical spin rates. At 44 RPM, the difference in the spin rate has only a slight impact on the hit RMS.

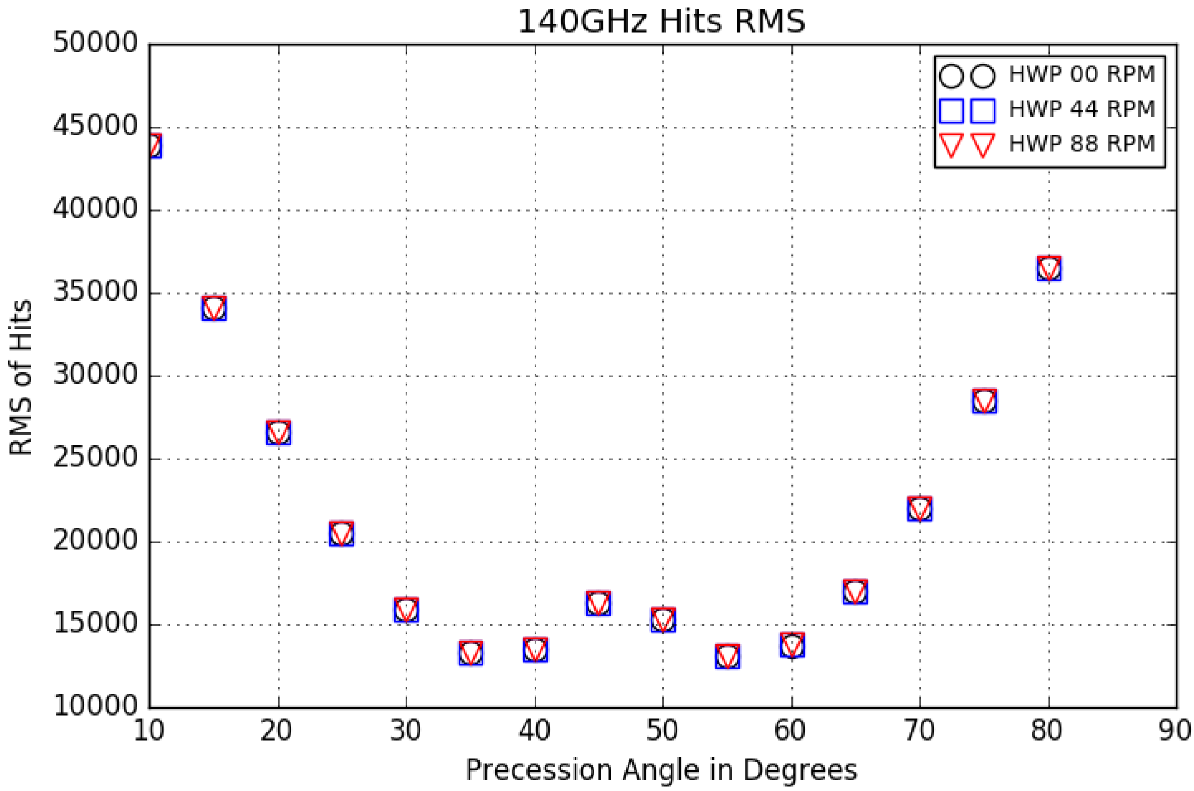

Fig. 4. Difference in CMB dipole signal extrema within a single ADR cycle without a HWP.

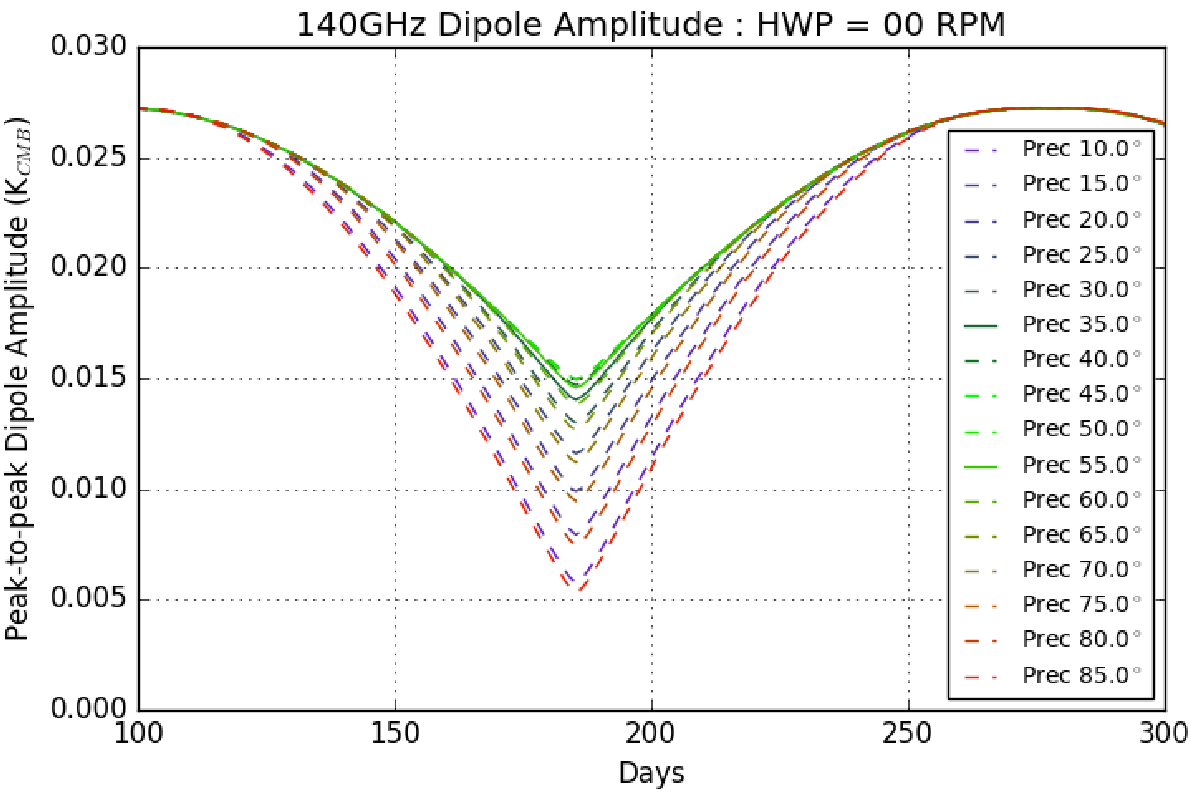

Fig. 5. Difference in CMB dipole signal extrema within a single ADR cycle with a HWP at 44 RPM.

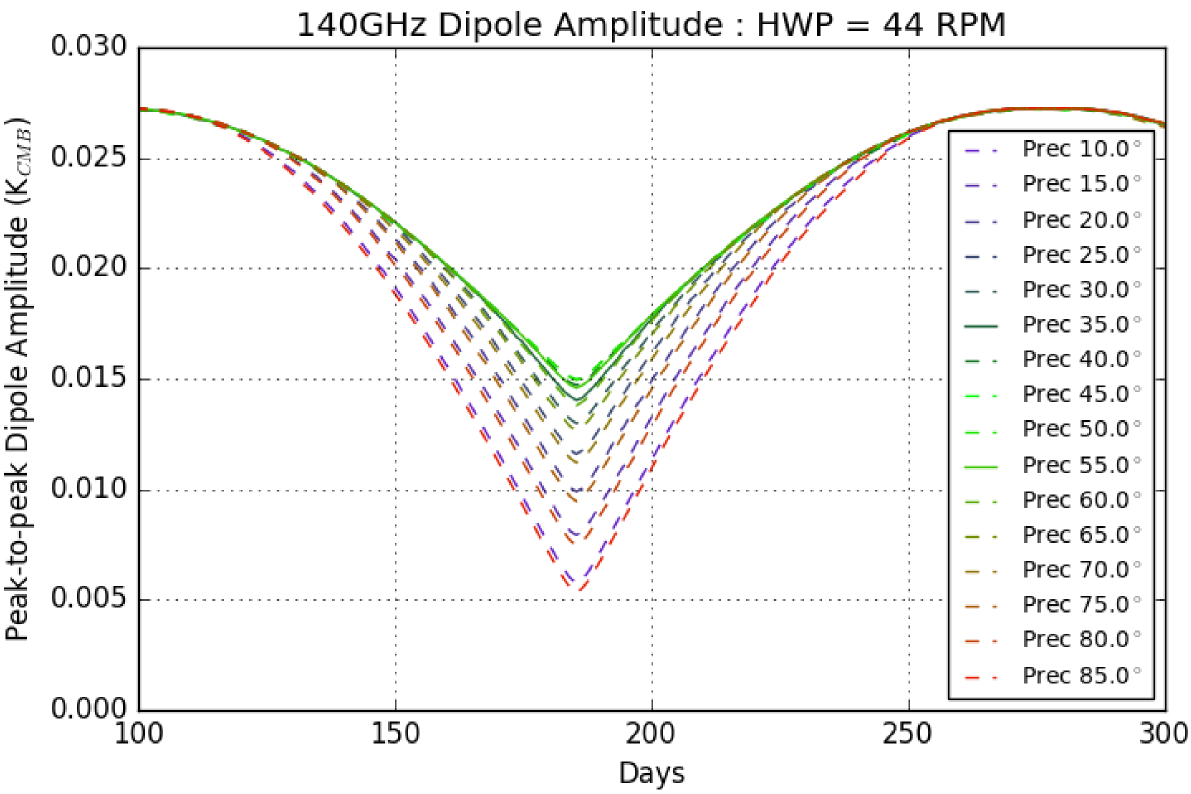

Fig. 6. Difference in CMB dipole signal extrema within a single ADR cycle with a HWP at 88 RPM.

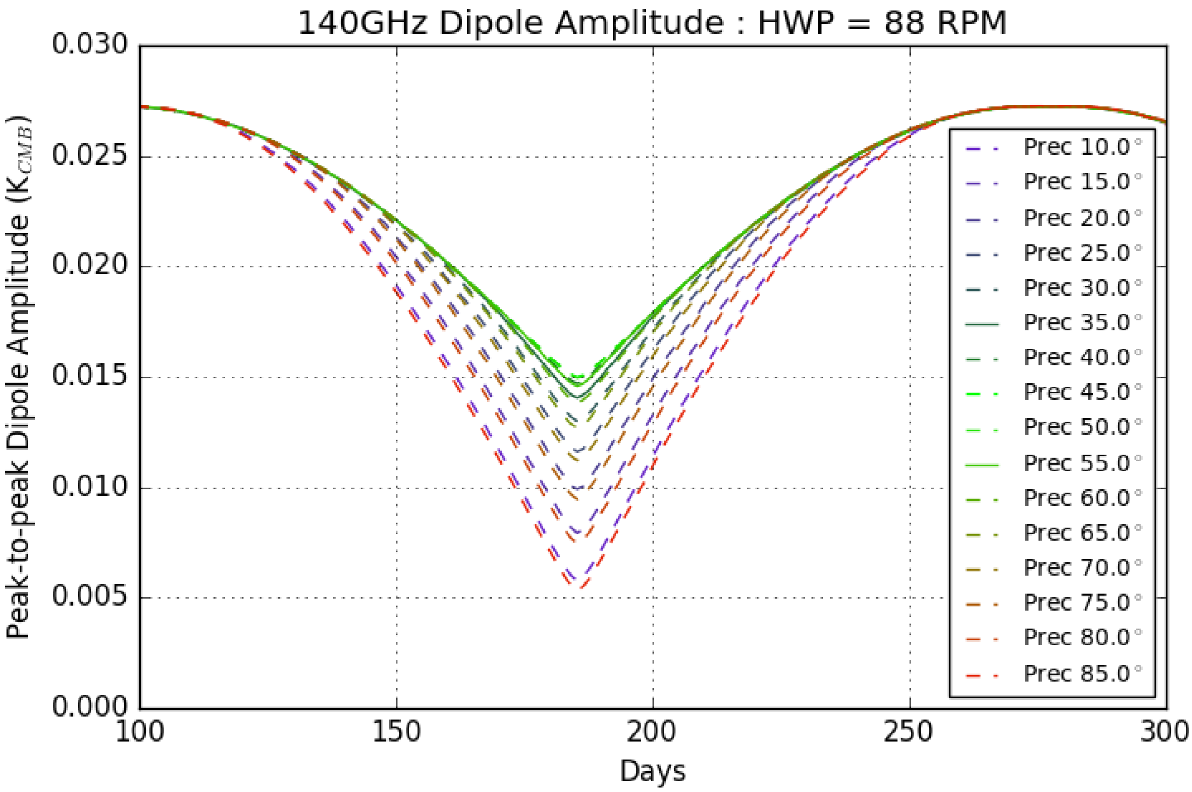

Fig. 7. Maps of the inverse pixel condition number for the three HWP states and two best precession opening angles.

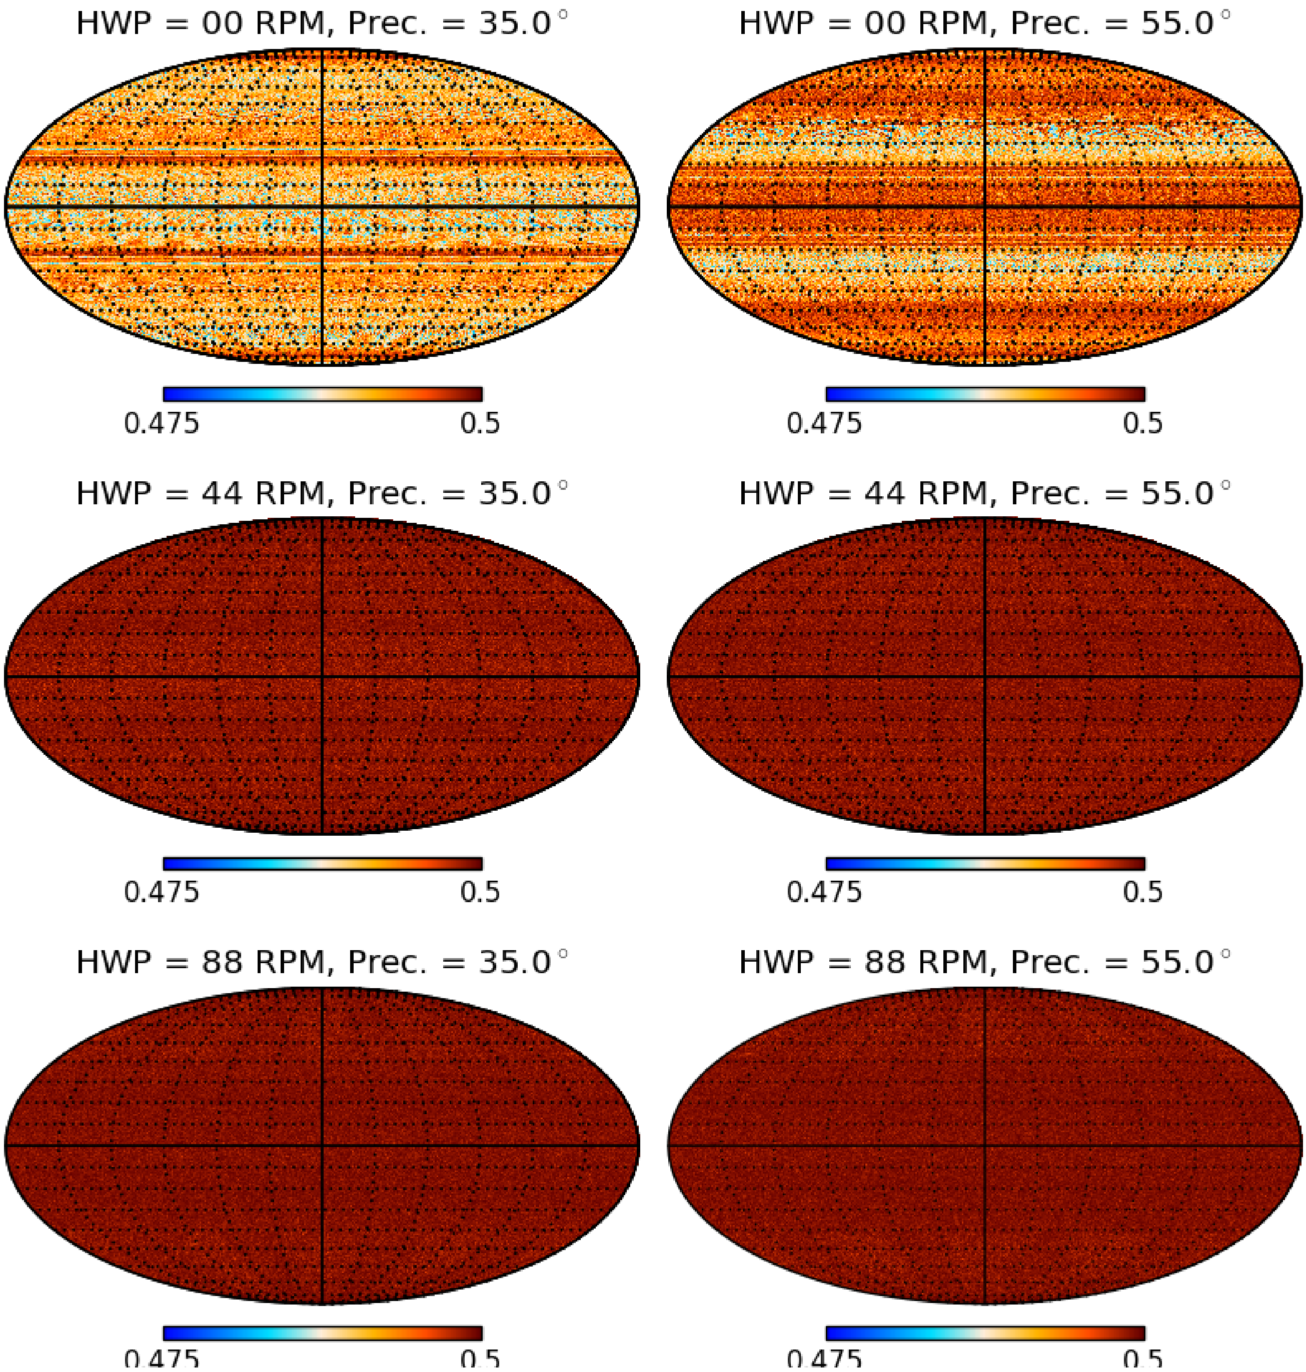

Fig. 8. Histograms of the inverse pixel condition numbers for the case with no HWP.

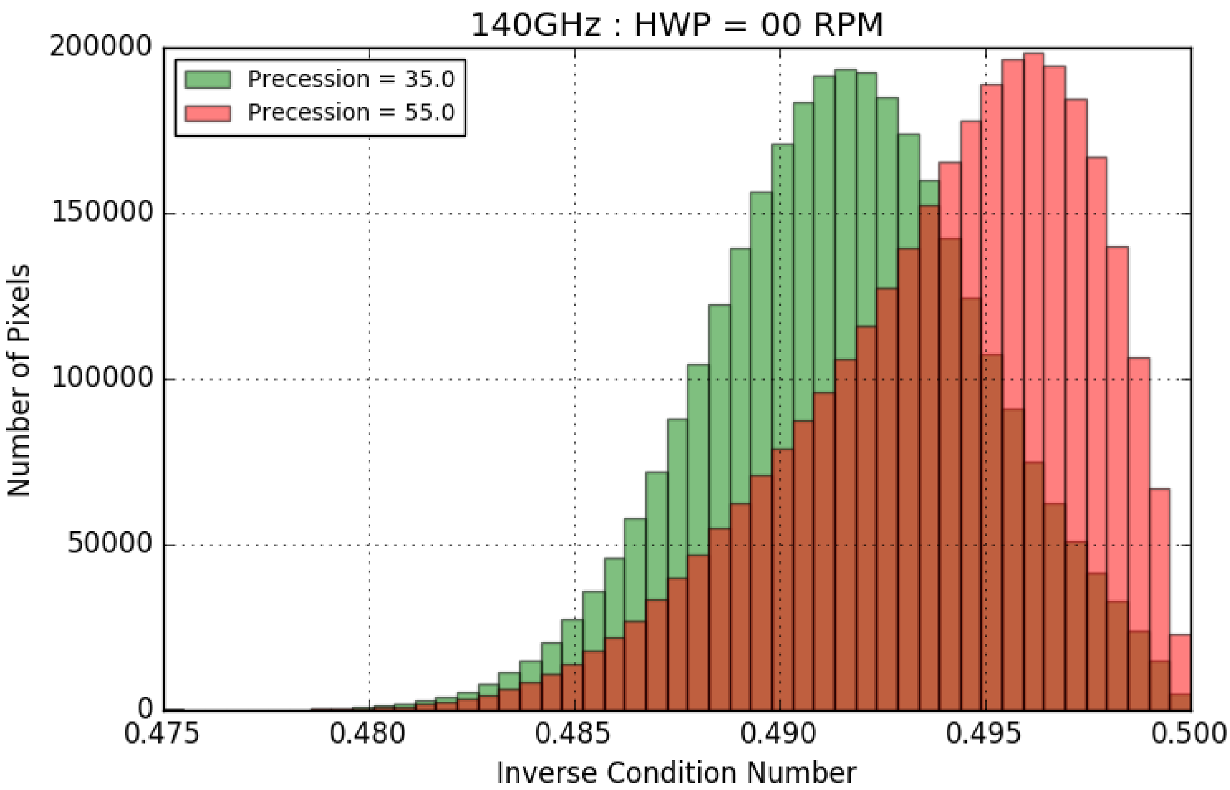

Fig. 9. Histograms of the inverse pixel condition numbers for the case of a HWP at 44 RPM.

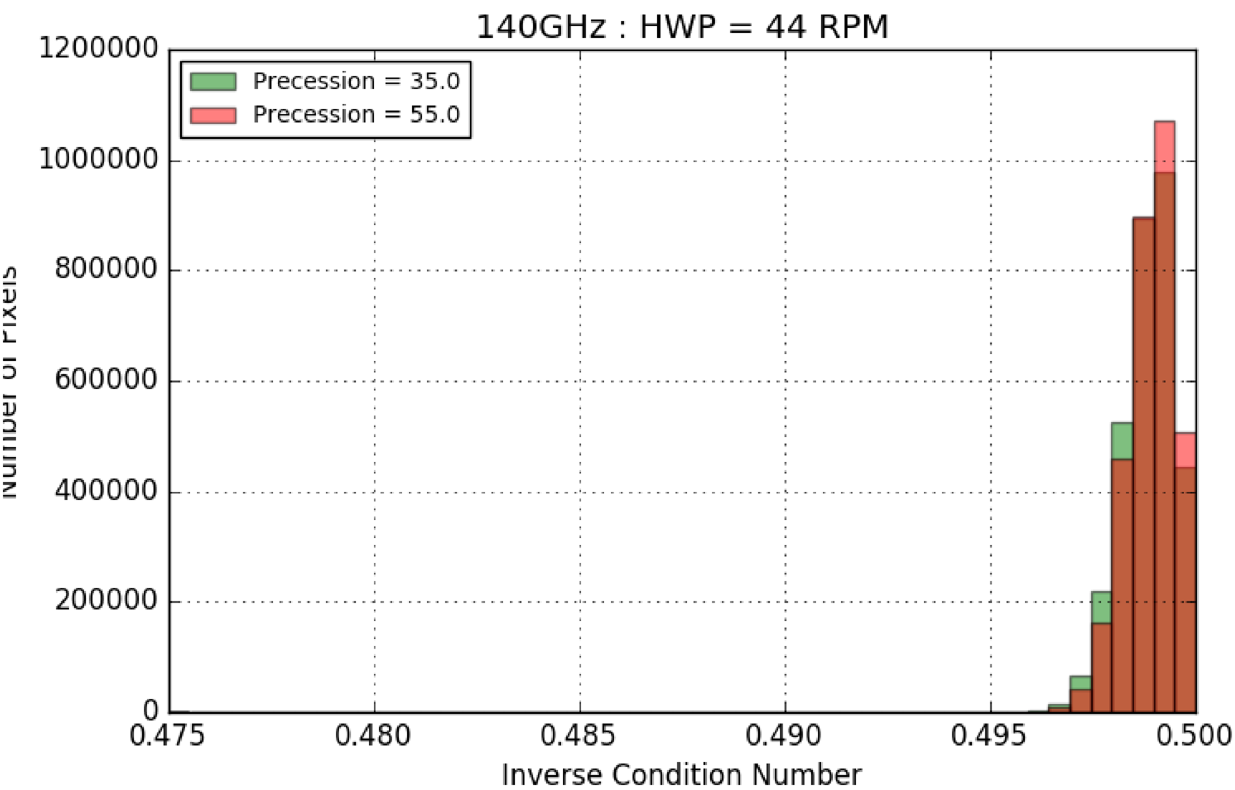

Fig. 10. Histograms of the inverse pixel condition numbers for the case of a HWP at 88 RPM.

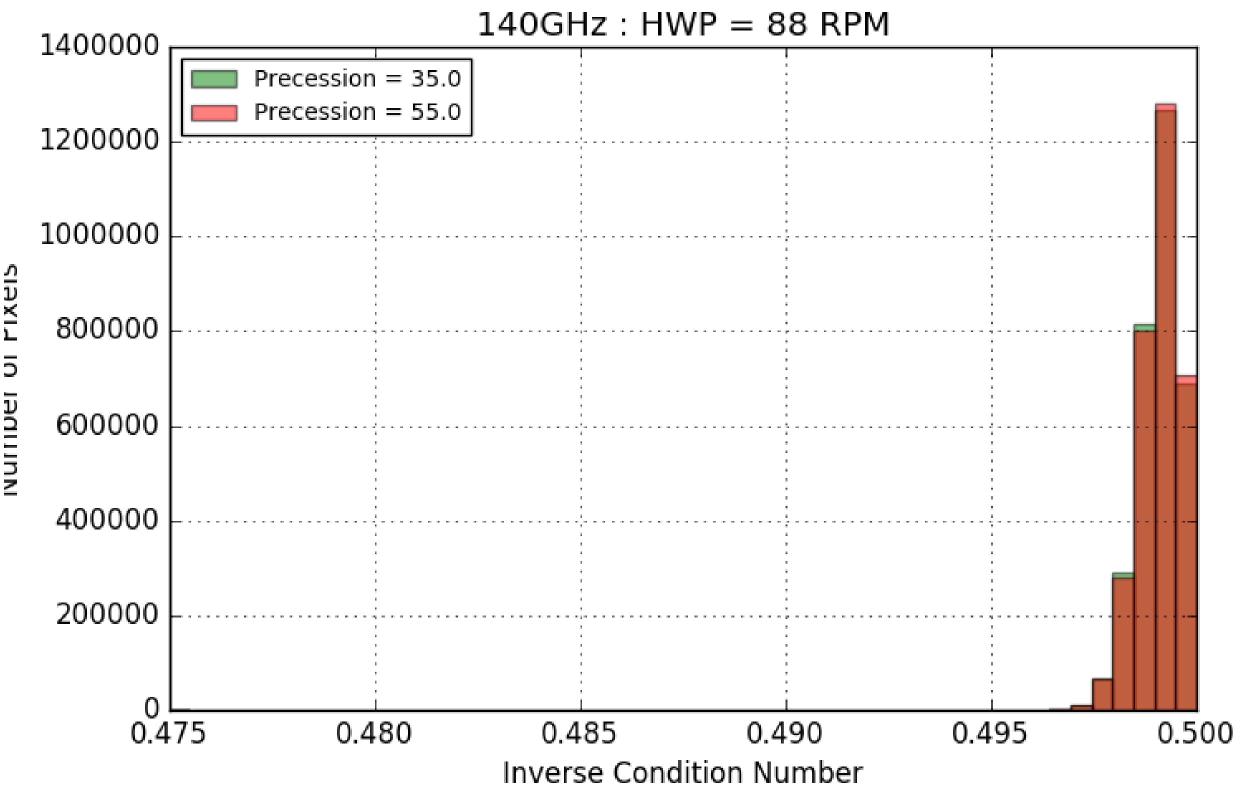

Based on the hit map RMS, the best case scanning parameters for all cases is close to either 35° or 55° precession opening angle. Both of those cases are also close to the maximum in terms of dipole signal amplitude. Now we compare the pixel condition numbers for these two "best" values of the precession opening angle for each HWP case. Figure 7 shows a comparison of the condition number maps for the three HWP states and two best precession angles. Histograms of the condition number maps are shown in figure 8, 9, and 10. Without a HWP, the larger precession opening angle of 55° looks better. However, all of these condition numbers seem "good".

3. Noise Angular Power Spectra

Now that we have narrowed down the best scan strategies for several HWP rotation speeds, we can estimate the expected noise level on our single-frequency angular power spectra. The TOAST software stack (see appendix B) is used to simulate noise timestreams from the detector parameters in table 1. Recall that for a simple simulation containing only sky signal and noise, there is **exactly one** maximum likelihood map given our input data. We use the MADAM destriping map maker configured in a way that will produce results very close to the maximum likelihood map. Most modern map making codes use iterative techniques to solve the map making equation, since the computational cost of the explicit solution scales as the cube of the number of pixels. This means that although we can make a map that is close to the maximum likelihood result, we have no estimate of the errors on this result. Instead, we generate an ensemble of noise simulations which are all consistent with the timestream noise PSD. We propagate these datasets through the map making operation and compute the angular power spectra. See appendix A for a detailed description of map making techniques used for this study.

For each HWP state and scan parameters, we construct our ensemble of noise Monte-Carlo maps and then compute the six angular autospectra. We use a relatively low NSIDE value (128) for the map making, since we are primarily interested in low-ℓ features and we wish to reduce the overall cost of these simulations. Figures 11, 12 and 13 show the resulting spectra for the three different HWP cases using the case of a 55° precession opening angle. Figures 14, 15, and 16, show the same for the case of a 35° precession opening angle. Figure 17 shows a comparison of the ensemble mean spectral values for all cases compared to some fiducial $\langle BB \rangle$ total spectra for different values of r. These figures do not use many realizations, but are binned to improve the sample variance.

Fig. 11. Angular power spectra from an ensemble of noise timestream realizations for the case of no HWP and a 55° precession opening angle.

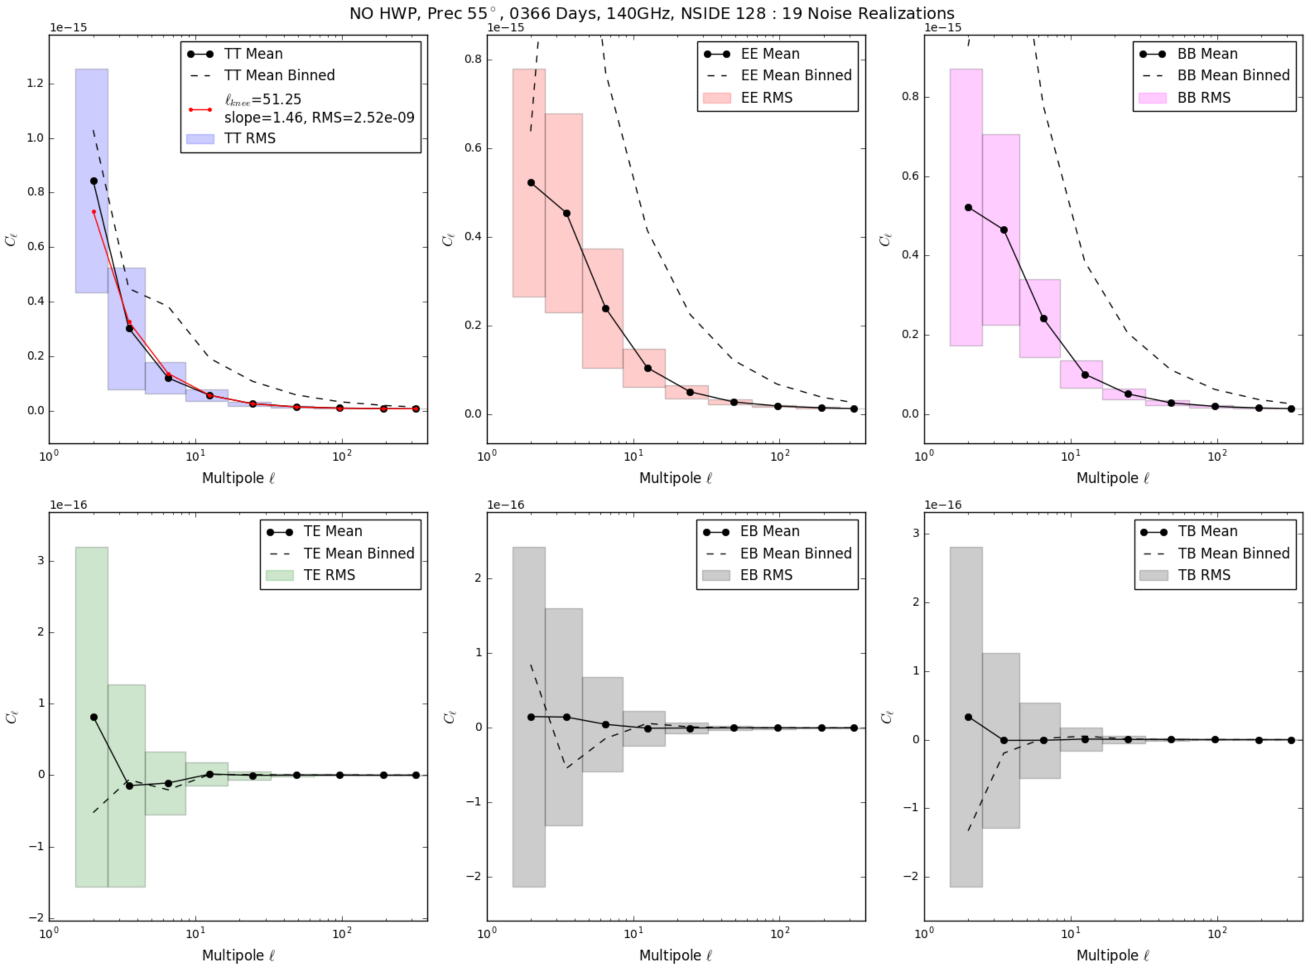

Fig. 12. Angular power spectra from an ensemble of noise timestream realizations for the case of a HWP rotating at 44 RPM and a 55° precession opening angle.

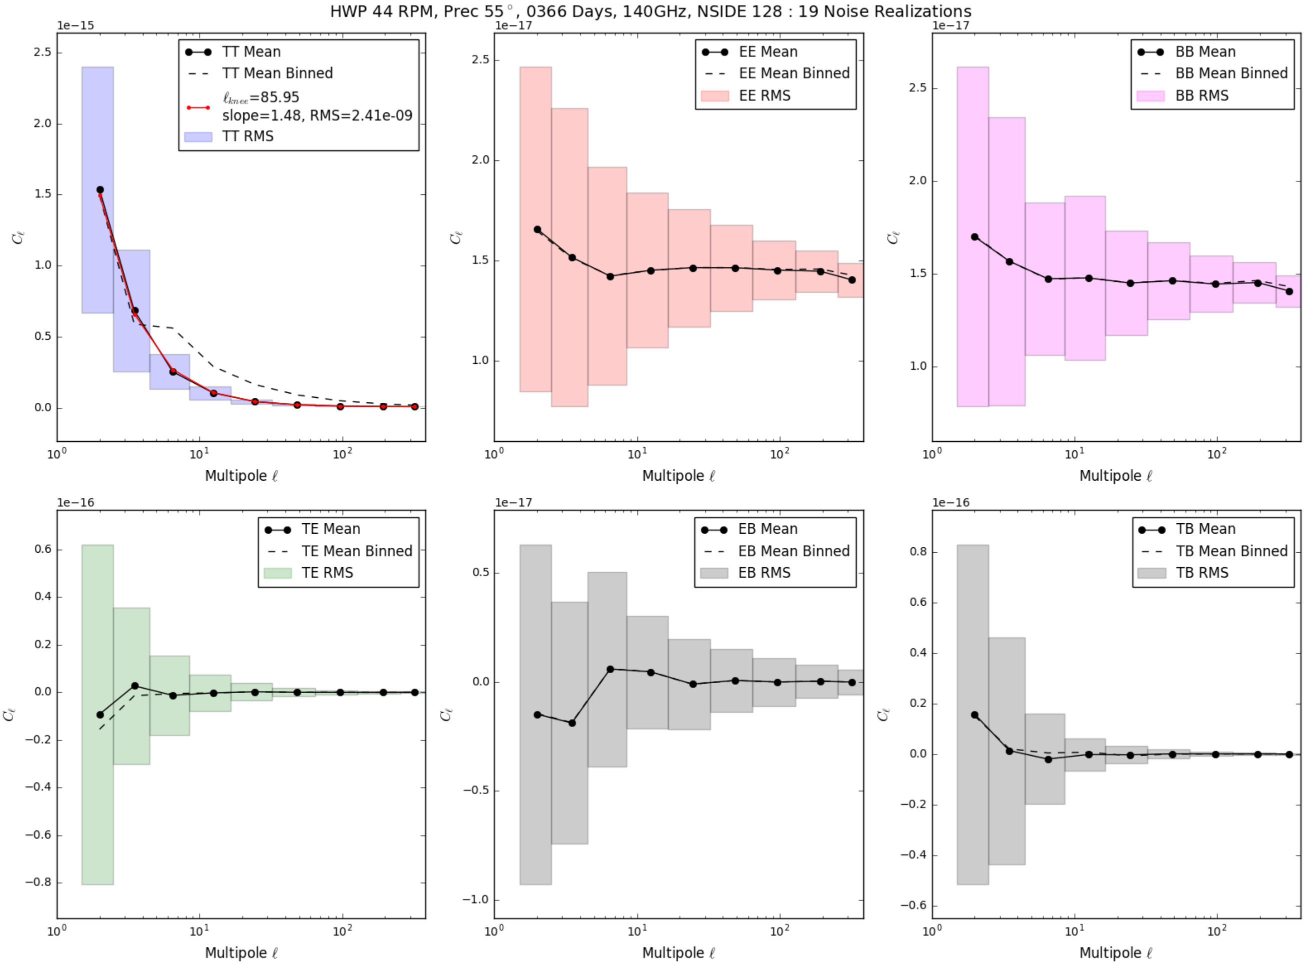

Fig. 13. Angular power spectra from an ensemble of noise timestream realizations for the case of a HWP rotating at 88 RPM and a 55° precession opening angle.

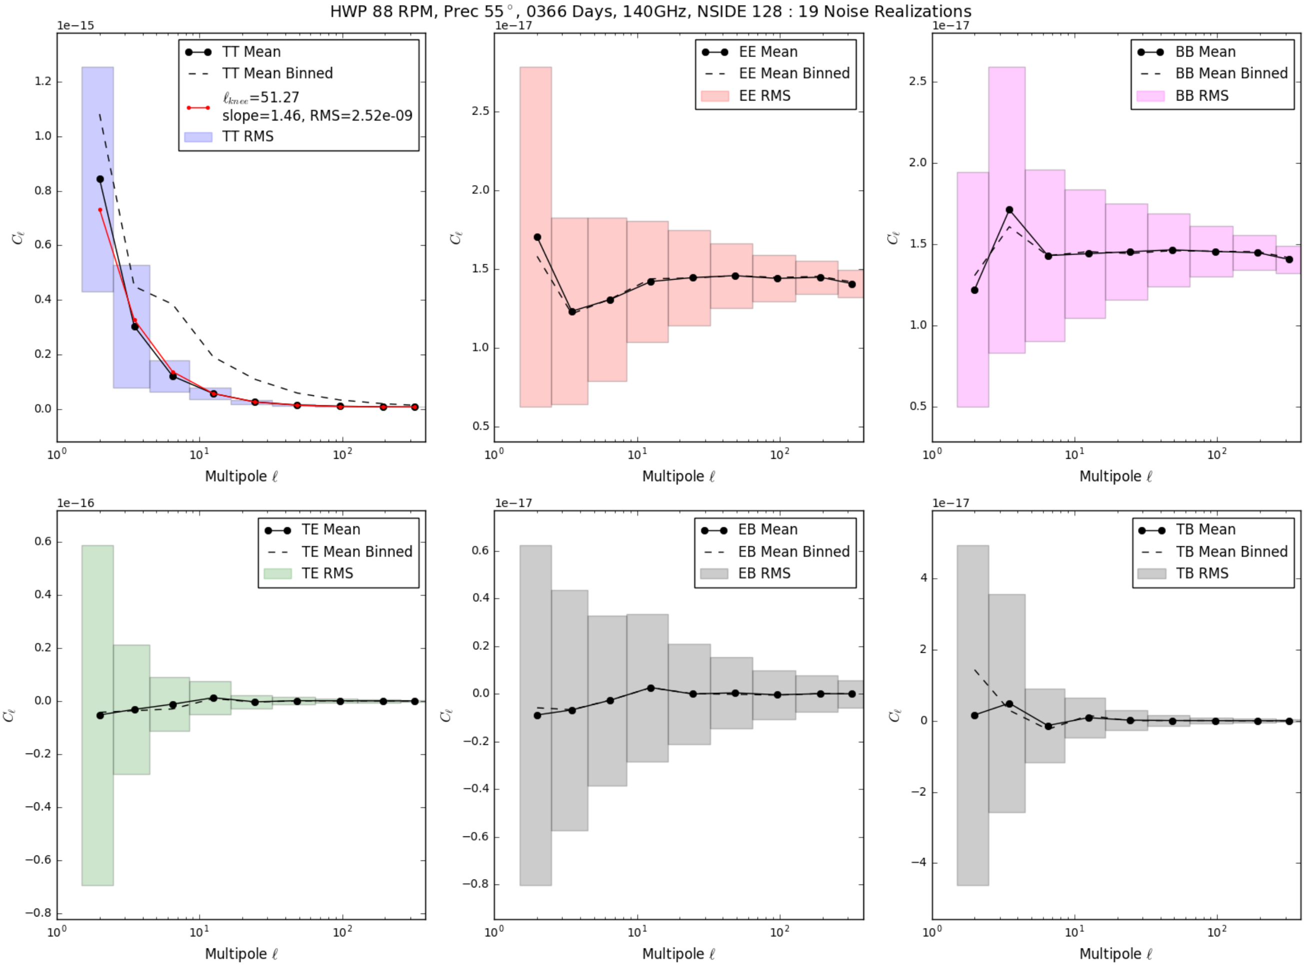

Fig. 14. Angular power spectra from an ensemble of noise timestream realizations for the case of no HWP and a 35° precession opening angle.

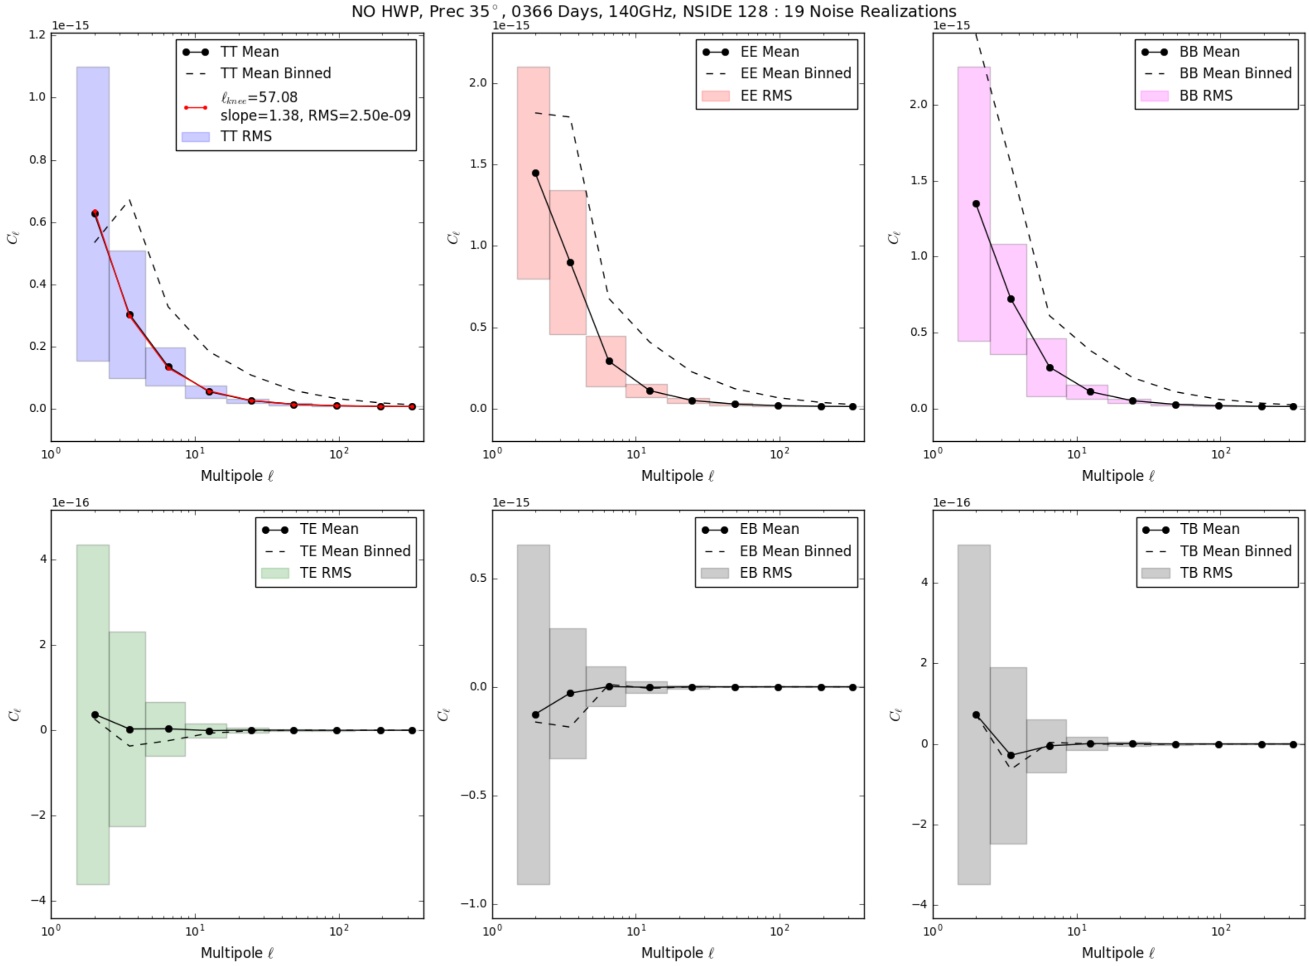

Fig. 15. Angular power spectra from an ensemble of noise timestream realizations for the case of a HWP rotating at 44 RPM and a 35° precession opening angle.

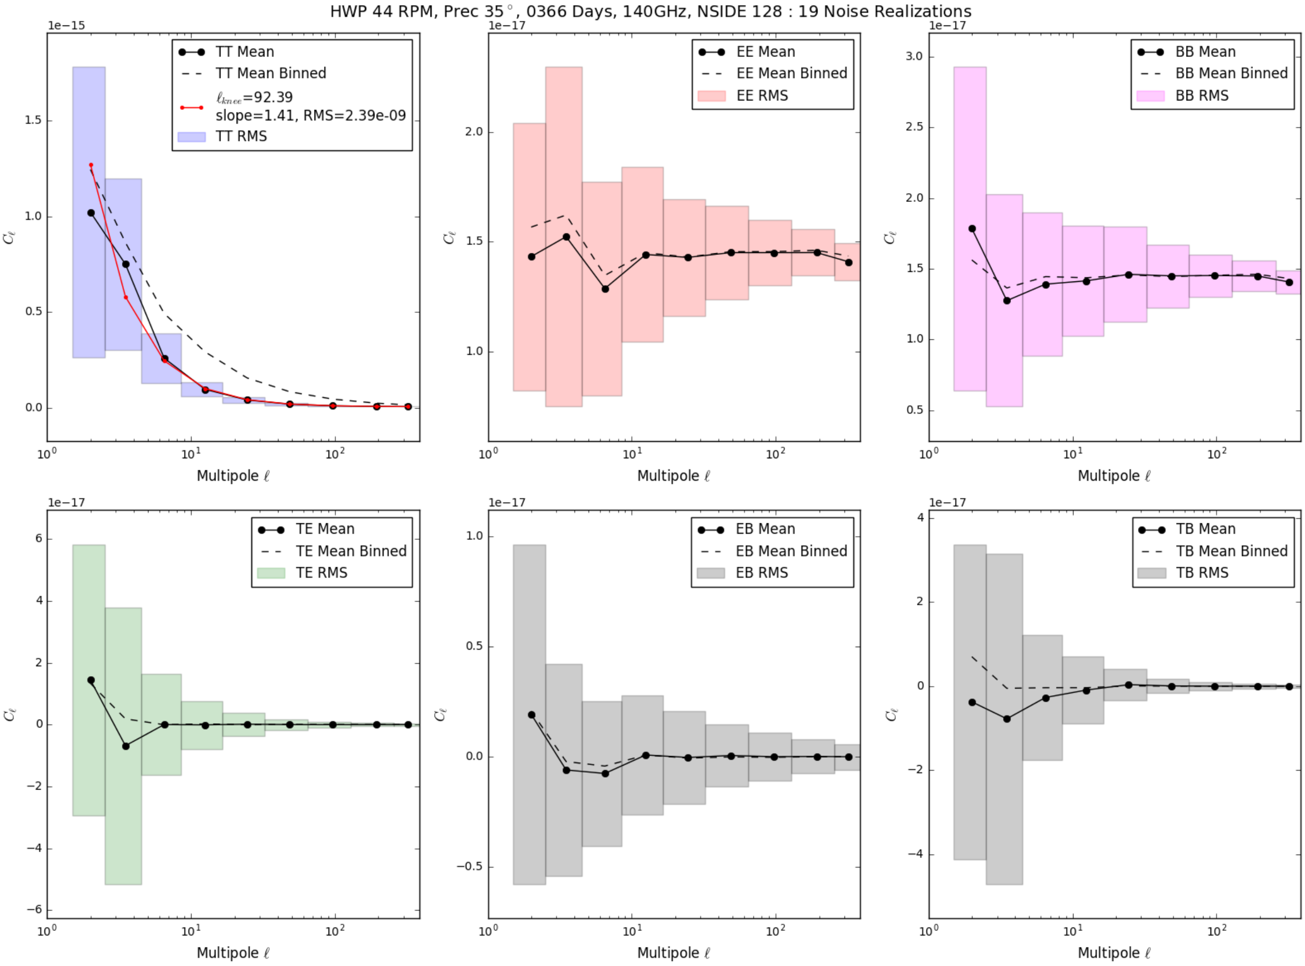

Fig. 16. Angular power spectra from an ensemble of noise timestream realizations for the case of a HWP rotating at 88 RPM and a 35° precession opening angle.

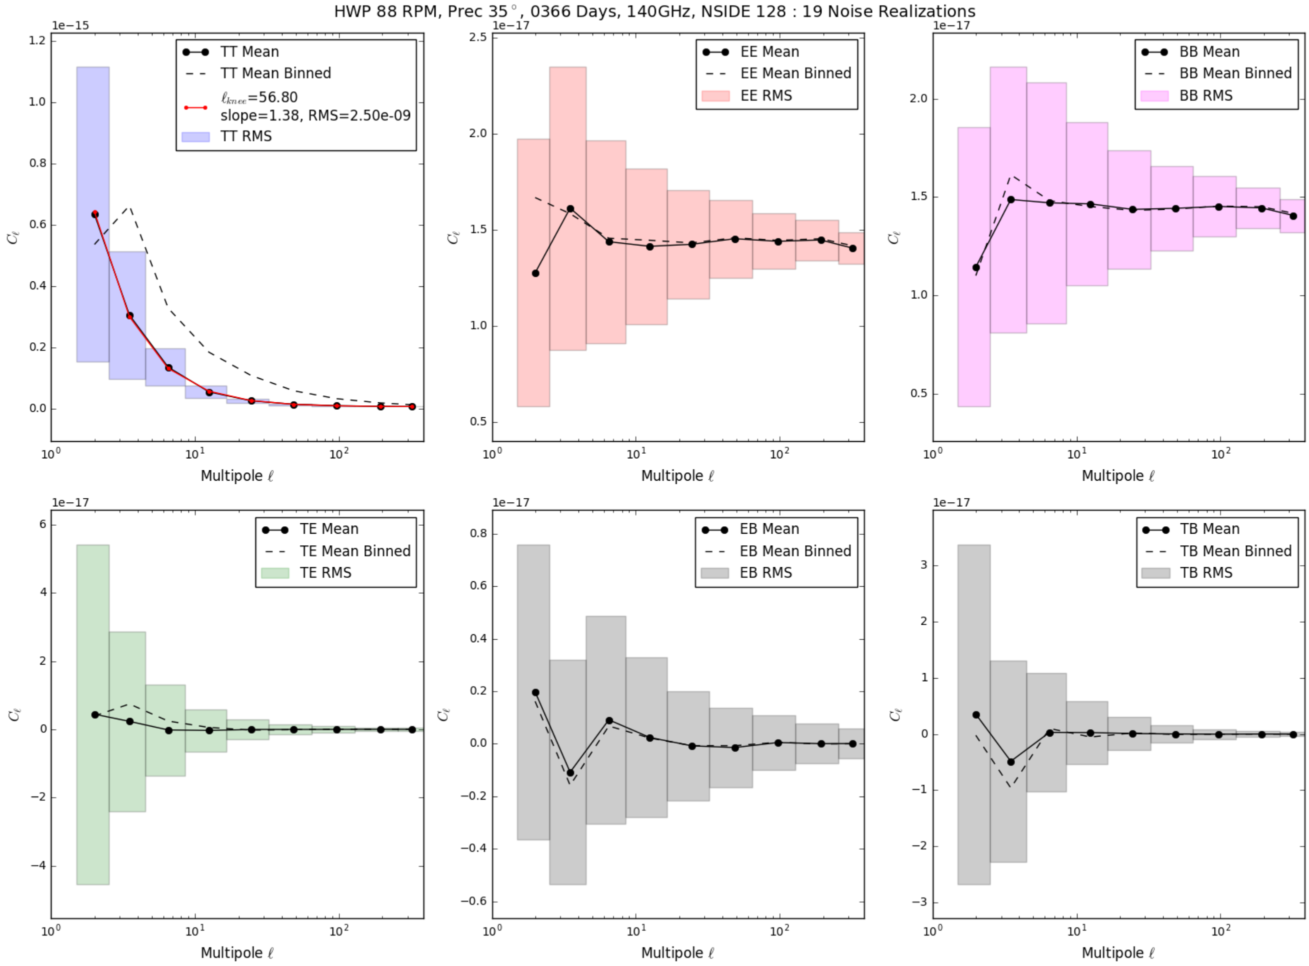

Fig. 17. Comparison of the mean angular power spectra for the three HWP cases (0, 44, and 88 RPM) and two precession angles with some fiducial total $\langle BB\rangle$ signal spectra for three different values of r.

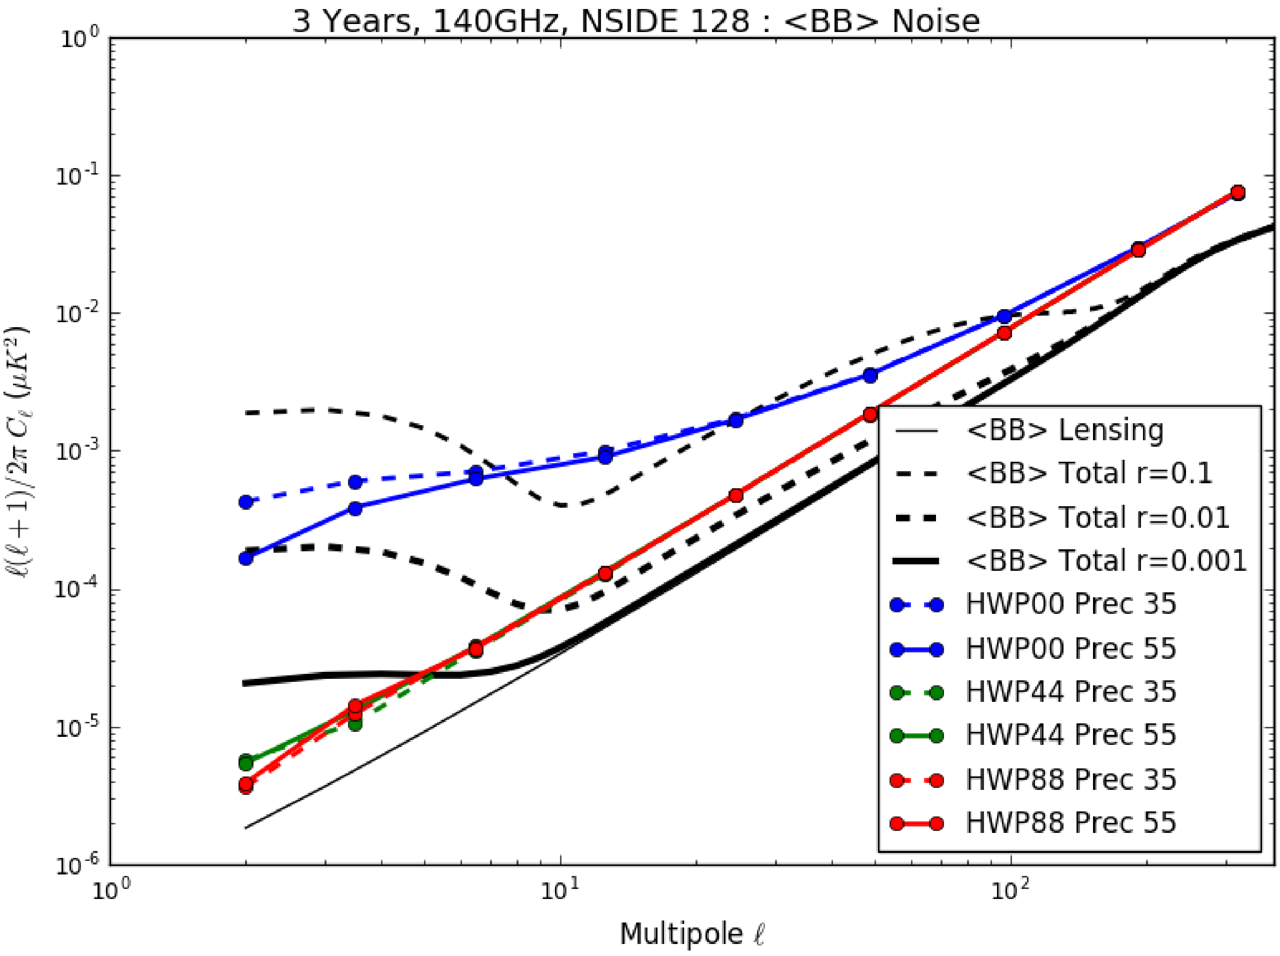

Based on this comparison, we find that the without the HWP there is significantly more low-ℓ noise power than the cases of a HWP at 44 or 88 RPM. Without a HWP, the $\langle BB\rangle$ noise power is lower in the case of a 55° precession opening angle. In the comparison figure 17, we see that without a HWP, the noise level is higher than the signal for all but the largest values of r.

4. Conclusions and Future Work

Based on two rotation rates of a (perfect) HWP and also the case of no HWP, we have identified the scan strategy parameters which maximize the hit map uniformity and ability to perform relative calibration on timescales close to that of the cooler cycle. Those simple pointing tests showed that precession opening angles of 35° and 55° seemed similarly good choices. Among these HWP cases, we compared the noise power spectra for the two scan strategies estimated using a small number of Monte Carlo realizations. We find that the cases without a HWP have more residual noise power at low-ℓ in the $\langle BB\rangle$ spectrum. The overall scan strategy produces similar coverage between frequencies, so we expect this qualitative result to hold for all frequencies.

Although this study shows the benefits of using a perfect HWP, it does not explore the impact of systematics created by the HWP and the removal of those systematics. That would be a logical next step for this work. We should also revisit the assumed noise levels in this study and also discuss whether the sample rate can be adjusted in any way.

References

Bond, J. R., Jaffe, A. H., & Knox, L. 1998, Physical Review D, 57, 2117

Borrill, J. 1999, in Proceedings of the 5th European SGI/Cray MPP Workshop, Bologna, Italy
Cantalupo, C. M., Borrill, J. D., Jaffe, A. H., Kisner, T. S., & Stompor, R. 2010, ApJS, 187, 212
Keihänen, E., Keskitalo, R., Kurki-Suonio, H., Poutanen, T., & Sirviö, A.-S. 2010, A&A, 510, A57
Keihänen, E., Kurki-Suonio, H., & Poutanen, T. 2005, MNRAS, 360, 390
Tegmark, M. & de Oliveira-Costa, A. 2001, Physical Review D, 64, 63001

Appendix A: CMB Map Making

Many areas of scientific measurement acquire some "data", which represents a projection of some underlying signal, plus noise from the apparatus performing the measurement:

$$d = \boldsymbol{A}\, S + n \tag{A.1}$$

where d is the measured data, n is the noise on that measured data, and S is the underlying signal. The matrix $\boldsymbol{A}$ is the *design matrix*, which describes how the underlying signal is projected into the space where the measurement is done. If n is assumed to be generated by a gaussian random process, then the maximum likelihood estimate for S is given by the solution to the generalized least squares problem:

$$\left(\boldsymbol{A}^T\boldsymbol{N}^{-1}\boldsymbol{A}\right)\hat{S} = \boldsymbol{A}^T\boldsymbol{N}^{-1}d \tag{A.2}$$

$$\hat{S} = \left(\boldsymbol{A}^T\boldsymbol{N}^{-1}\boldsymbol{A}\right)^{-1}\boldsymbol{A}^T\boldsymbol{N}^{-1}d \tag{A.3}$$

$$\hat{S} = \boldsymbol{C}\boldsymbol{A}^T\boldsymbol{N}^{-1}d \tag{A.4}$$

where the measurement-domain noise covariance has been defined as

$$\boldsymbol{N} = \langle n^T\, n\rangle \tag{A.5}$$

and the data covariance is given by

$$\boldsymbol{C} = \left(\boldsymbol{A}^T\boldsymbol{N}^{-1}\boldsymbol{A}\right)^{-1}. \tag{A.6}$$

In the case of CMB data, the underlying signal is the pixelized map of the sky, and the matrix $\boldsymbol{C}$ is the full pixel-pixel noise covariance:

$$d_t = \boldsymbol{A}_{pt}\, S_p + n_t \tag{A.7}$$

$$\hat{S}_p = \boldsymbol{C}_{pp'}\boldsymbol{A}_{pt}{}^T\boldsymbol{N}_{tt'}{}^{-1}d_t \tag{A.8}$$

For small numbers of pixels, the maximum likelihood map and covariance can be solved exactly using Fourier techniques to apply the $\boldsymbol{N}_{tt'}{}^{-1}$ filter in a piecewise way and using distributed linear algebra to invert the pixel covariance Borrill (1999). This map and pixel covariance can then be used to compute the maximum likelihood angular power spectra **and** approximate the ℓ-space covariance between spectral bins (Bond et al. (1998), Tegmark & de Oliveira-Costa (2001)).

It is worth emphasizing that equation (A.8) is *linear*. If you had perfect knowledge of the signal and noise parts of d_t, then you could make separate maps of the noise and signal timestreams and their sum would equal the map of the combined timestream.

Appendix A.1: Iterative Techniques

As the number of pixels increases, it becomes impossible to compute and store the dense pixel covariance needed to solve for $\hat{S}_p$ in equation (A.8). Instead, we can solve (A.2) iteratively with a preconditioned conjugate gradient technique Cantalupo et al. (2010). This gives us the same result for the maximum likelihood map, but without the full pixel-pixel covariance. To estimate the errors on our results due to noise, we instead construct random realizations of the timestream noise which are consistent with the noise PSD that is either estimated (from data) or modelled (for pure simulations). These simulations are then processed with the identical map making procedure.

Appendix A.2: Destriping

The map making technique known as destriping has been used for more than a decade, and involves changing the description of the data from (A.7) to:

$$d_t = \boldsymbol{A}_{pt}\, S_p + \boldsymbol{B}_{at}\, n_a + w_t \tag{A.9}$$

Where the noise is now described by piecewise "steps" (baselines, n_a) and white noise (w_t). The full discussion of this technique can be found in Keihänen et al. (2005) and Keihänen et al. (2010). Here we simply summarize the results of those papers. The solution to (A.9) can be found by simultaneously maximizing the likelihood of S_p and n_a. In practice this is done by solving iteratively for the maximum likelihood baseline values given the dataset (equation 7 of Keihänen et al. (2010)), and then solving for the maximum likelihood map (equation 6 of Keihänen et al. (2010)). Note that this equation is just the usual "binned map" of the baseline-subtracted data.

The destriping formalism is completely independent of scan strategy. If the timestream noise is well described by baselines plus white noise, then the resulting map is identical to the solution given by (A.8). In the limit where the baseline length shrinks to one sample, then any timestream noise is well described by this formalism. We use this fact in the current study in order to remove all questions of baseline length from the discussion. We simply run madam with the true noise filter (which was used to simulate the noise timestream) and we set the baseline length to be well above the knee frequency of the noise.

Appendix B: Software Stack

The TOAST software stack is a collection of libraries and tools for manipulating astrophysical timestream data. It was originally written in C++ and included both timestream simulation and map-making functionality. Recently much of the user-facing code was rewritten in python, and many of the internals are in the process of being migrated from the old code base to the new. TOAST also now has an interface to some external libraries, including a library version of the original MADAM destriping map maker. The main code repository is at https://github.com/hpc4cmb/toast, and there are experiment-specific repositories for Planck, Core, and LiteBIRD. The LiteBIRD repository is available to collaborators at https://github.com/hpc4cmb/toast-litebird. TOAST currently makes heavy use of the library version of the MADAM destriping map maker (https://github.com/hpc4cmb/libmadam).