

Analysis of the accuracy of a destriping method for future cosmic microwave background mapping with the PLANCK SURVEYOR satellite

J. Delabrouille^{1,2}

¹ Institut d'Astrophysique Spatiale, CNRS & Université Paris XI, Bât. 121, 91405 Orsay Cedex, France

² Enrico Fermi Institute, University of Chicago, 5460 South Ellis Avenue, Chicago, IL 60510, U.S.A.

Received March 20; accepted May 9, 1997

Abstract. A major problem in cosmic microwave background radiation (CMBR) anisotropy measurements is the presence of low-frequency noise in the data streams. This noise arises from thermal instabilities of optical elements or of the thermal bath, gain instabilities and $1/f$ noise in the electronics, and other poorly understood processes. If improperly monitored or processed, this excess low-frequency noise might lead to striping in the maps, compromising the success of the experiment. In this paper, we show that a simple destriping method will clean the maps obtained with the High Frequency Instrument of the PLANCK SURVEYOR mission of any significant additional noise from low-frequency drifts, provided that the knee frequency of the low frequency noise is less than the spinning frequency of the satellite, i.e. $f_{\text{knee}} \leq 0.017$ Hz. For the High Frequency Instrument of PLANCK, the nominal knee frequency of the noise is $f_{\text{knee}} \simeq 0.01$ Hz or less, and thus no significant striping nor increase of the noise rms is expected due to low-frequency drifts. In addition, we show that even if the knee frequency of the low frequency noise were somewhat higher than the spinning frequency of the satellite one could estimate and remove the striping with a excellent accuracy.

Key words: methods: data analysis — cosmology: cosmic microwave background — space vehicles

1. Introduction

After the encouraging results of the DMR experiment on the COBE satellite (Smoot 1992), there has been a burst of renewed interest in the anisotropies of the Cosmic Microwave Background Radiation (CMBR), both on the

experimental side and on the theoretical side. While theorists refined calculations to evaluate how individual parameters of the theories affect the expected properties of the tiny fluctuations of the CMBR, experimenters, in answer to announcements of opportunities by several space agencies, proposed sophisticated new generation satellites to map the anisotropies of the CMBR with a sensitivity and angular resolution an order of magnitude better than those of COBE. Of these, the Microwave Anisotropy Probe (MAP) experiment has been selected by NASA as one of the next medium-size explorer, or midex, and the PLANCK satellite (formerly COBRAS/SAMBA) has been selected by ESA as the next medium-size mission M3.

The accuracy with which the useful cosmological information can be deduced from the data of such a mission depends on the global characteristics of the instrument and on the observing strategy: sensitivity of the detectors, spectral coverage, resolution, susceptibility to systematics, scanning strategy... the optimal solution is often a trade-off between several marginally compatible constraints, and very different strategies can be adopted.

Because a large telescope is necessary in order to achieve the high angular resolution that is mandatory to distinguish between cosmological models, the option of differential measurements implies complicated optics. Fortunately, the availability of space-qualified cryogenic devices (Benoit et al. 1994) and the development of new readout electronics (Gaertner et al. 1997) now permits to use in space bolometers cooled to 0.1 K. These are much more sensitive than available radiometers at the frequencies most interesting for cosmology, and so much more stable that for the PLANCK bolometer instrument (HFI, for High Frequency Instrument), the conservative differential approach has been abandoned in favour of total-power measurements.

In the case of the nominal PLANCK mission, the scanning of the sky is performed very simply, by rotating the

satellite at 1 rpm around a spin axis which position on the sky is roughly anti-solar (to first order the spin axis remains in the ecliptic plane, and its position is shifted by $\sim 5'$ every 2 hours). There is some flexibility as to the direction of the spin axis, however, so that the scanning strategy can be adapted for optimal sky coverage or rejection of systematics within technical constraints (the thermal stability of the payload puts a limit on the solar aspect angle of 15° , and the telemetry rate a limit on the earth aspect angle of 15° during the dumping of the data).

The beam axis makes an angle of 70° with the spin axis, and thus scans 140° diameter circles on the sky. The PLANCK orbit is a Lissajous orbit around the sun-earth Lagrange point L2. More details on PLANCK can be found in the COBRAS/SAMBA report on the phase A study (1996).

For this observing strategy, there are several important characteristic time-scales. One scan corresponding to one complete rotation of the satellite around itself is performed in a time $T_{\text{spin}} = 1$ minute. Data circles are obtained by averaging 120 such scans, and correspond to a period of two hours. Each of these data circles crosses in two points all other data circles obtained less than about 20 weeks before or after. Thus, they share a common area on the sky of at least two pixels (and more for circles that are tangent or nearly tangent). Finally, data circles corresponding to measurements separated by a 1 year period coincide on the sky. Figure 1 shows four PLANCK SURVEYOR circles on a sinusoidal projection of the sky in ecliptic coordinates.

SINUSOIDAL PROJECTION OF FOUR "PLANCK SURVEYOR" SCAN CIRCLES




Fig. 1. Sinusoidal projection of four scans for the nominal PLANCK SURVEYOR scanning strategy. The second, third and fourth scans from the left (long-dashed, short-dashed, and dotted line) are obtained 2 weeks, 2 months, and 5 months respectively after the first (plain line). Because these circles have 140° diameters, intersections are distributed everywhere along circles, not only at ecliptic poles as for great circles. Not all pairs of circles intersect: here for instance the first and last circles, separated by a period of 5 months, have no pixel in common

All these redundancies at very different time-scales make it possible to minimise low-frequency noise effects by comparing the values the signal takes at times where

the useful astrophysical signal is supposed to be the same (because the antenna is pointed at the same place on the sky) and thus estimating and correcting for low-frequency drifts.

It has recently been suggested that the scanning strategy and destriping method adopted for PLANCK might lead to striping on the maps due to excess low-frequency noise even if there were no intrinsic low frequency noise in the measurements (Wright 1996).

In that paper, the author argues correctly that relying on no more than a pixel or two per scan (namely, ecliptic poles) to readjust relative offsets might cause striping in the maps. However, when suggesting that for this reason the PLANCK maps will be striped, he seems to disregard completely two essential characteristics of the PLANCK SURVEYOR scanning strategy, which are the 15° freedom of motion of the spin axis and the 70° off-axis spin angle, any of which characteristics modifies completely the way circles on the sky intersect each other. Thus, his suggested conclusions should be regarded with extreme caution.

In the following, we investigate how well a simple destriping technique can remove the striping in the maps in the context of the PLANCK SURVEYOR mission.

2. Analysis of the “noise processing”

It is in general a reasonably good assumption that the total noise (excluding systematics that might be scan-synchronous or correlated to the signal), be a Gaussian process that can be described by a Power Spectral Density (PSD, in Volts² of electrical signal per Hz) typically of the form:

$$\mathcal{S}_n(f) = a \left(1 + \left(\frac{f_{\text{knee}}}{f} \right)^\alpha \right). \quad (1)$$

Here f_{knee} is a “knee frequency” at which low-frequency noise and white noise contributions to the power spectral density are statistically equal, and α is a spectral index, typically between 1.0 and 2.5, depending on the dominant physical process which generates low frequency noise. It is a good approximation to assume that $\mathcal{S}_n(f)$ vanishes for $f < f_{\text{min}}$ and $f > f_{\text{max}}$, with $f_{\text{min}} \sim 1/T_{\text{total}}$ and $f_{\text{max}} \sim 1/2T_{\text{sampling}}$.

In an experiment for which the scanning strategy consists in scanning repeatedly the same circle of the sky, the sky signal for any given pixel is calculated by averaging the “samples” corresponding to this pixel. When the white noise contribution dominates the total noise, i.e. when f_{knee} and α are such that

$$f_{\text{max}} - f_{\text{min}} \gg \int_{f_{\text{min}}}^{f_{\text{max}}} \left(\frac{f_{\text{knee}}}{f} \right)^\alpha df, \quad (2)$$

which can be simplified to $f_{\text{max}}/f_{\text{knee}} \gg \ln(f_{\text{max}}/f_{\text{min}})$ if $\alpha = 1$, and $f_{\text{min}} \ll f_{\text{max}}$, or to $f_{\text{min}}/f_{\text{max}} \ll (\alpha - 1) \times (f_{\text{min}}/f_{\text{knee}})^\alpha$ if $\alpha > 1$ and $f_{\text{min}} \ll f_{\text{max}}$, then the rms

of the final noise on one scan circle obtained by averaging N consecutive scans is just the rms of the original noise divided by \sqrt{N} . This is not the case when the noise is significantly coloured (i.e. not white), since averaging consecutive scans actually reduces the bandwidth of the signal in such a way that most of the low-frequency noise contribution at frequencies $f < 1/T_{\text{spin}}$ is filtered out. In Fourier space, the spinning and averaging process keeps all components of the signal that are harmonics of the spinning frequency, and cuts other components (this is only an approximation, but good enough for this discussion). In our case, the low-frequency contribution to the standard deviation will be much larger on a 2-hour data stream than on the corresponding circle of data obtained by averaging.




Fig. 2. Example of 8 consecutive hours of simulated low-frequency noise, with $\alpha = 2.0$ and $f_{\text{knee}} = 0.10$ Hz. These values are typical for ground-based experiments which suffer from thermal fluctuations and atmospheric noise (for the PLANCK HFI, due to the extremely favourable observing conditions, we expect the low-frequency drifts to be dominated by electronics noise, i.e. $\alpha \simeq 1.0$ and $f_{\text{knee}} \leq 0.01$ Hz instead). The top panel shows drifts due to the low-frequency part of the noise only, with a spectrum of the form $S(f) = a \times (f_{\text{knee}}/f)^2$. The effect of adding the white noise contribution (which has a standard deviation of 1.00) would be a strong broadening of the line. The bottom panel shows the absolute value of the FFT of the total noise. The knee frequency of 0.10 Hz is clearly visible

The top plot of Fig. 2 is a plot of 8 hours of the low-frequency part of a noise with a spectrum of the form of Eq. (1), with $\alpha = 2.0$ and $f_{\text{knee}} = 0.10$ Hz (which is extremely pessimistic for the PLANCK HFI, but illustrates

our point better than more realistic noise: drifts on the scale of 8 hours are almost imperceptible for $\alpha = 1.0$ and $f_{\text{knee}} = 0.01$ Hz unless the noise is significantly smoothed to filter high frequency components). Such figures for the noise are quite typical for ground-based CMB experiments. Only the low-frequency component is represented in the top panel of Fig. 2 (which corresponds thus to a spectrum or PSD of $S_n(f) = (f_{\text{knee}}/f)^2$). Units for noise generation have been normalised so that the rms of the white part of the noise per sample is 1.00 (and the rms per resolution element on a data circle, obtained from averaging the 120 consecutive scans of one 2-hour period for one detector, would be $1/\sqrt{120} = 0.091$ if there were no low frequency noise). In this simulation, the sampling frequency is 2034 samples per minute, which for PLANCK corresponds to one sample per 10 arcminute pixel. The bottom plot of Fig. 2 is the absolute value of the Fast Fourier Transform of the noise realization with white noise included as well as the low-frequency part. The knee frequency of about 0.10 Hz is clearly visible.




Fig. 3. Residual total noise on 4 consecutive circles, each obtained by averaging 120 consecutive 1 minute scans of the 8-hour sample of the simulated low-frequency noise plotted in Fig. 2. The data for each circle looks flat at this scale, but each circle has a very different offset

Low frequency drifts with an amplitude of a few (whereas the rms of the white part of the noise is 1.00) are obvious in Fig. 2. The standard deviation of the total 8-hour noise (white + low-frequency) is about 2.49, that of the low-frequency part is 2.28, and that of the white part is 1.00. However, what actually happens on

circles after 120 consecutive scans have been averaged is shown in Fig. 3, where the data corresponding to the four circles obtained from this 8-hour signal are plotted with the same units on the vertical axis. Here, the white noise component has been included. The variance on each individual circle is much smaller than that of the original signal, because each circle is obtained by averaging 120 scans as explained above. Also, it is impressive to notice that almost all the power of low frequency noise now appears in the form of a different “offset” A_i for each circle i . The average level of all the circles is not the same, but low frequency drifts on each circle are much smaller than on the original time sequence.




Fig. 4. Here we show the centred signal for each of the circles of Fig. 3. Note the scale compared to that of Fig. 2 and Fig. 3. Even for the very pessimistic assumed values for α and f_{knee} , low frequency drifts are barely visible in the white noise

It is important to check whether the low frequency noise will have any effect other than to add a different offset to the data corresponding to each circle. Figure 4 shows the centred noise for each of the above 4 circles. Even with the very pessimistic values of α and f_{knee} used here, low-frequency drifts are almost too low to be distinguished in the dominating white noise. If however we remove the white part of the noise and look at the four circles (Fig. 5), we see clearly that some low frequency drifts are still there, which would appear if a drastic smoothing of the data were performed.

A crude estimate of the increase of the standard deviation of the noise on a circle (compared to the standard deviation of white noise only) can be obtained by inte-

grating the PSD between $f_{\text{min}} \sim f_{\text{spinning}} \simeq 1/60$ Hz and $f_{\text{max}} \simeq f_{\text{sampling}}/2$. This method underestimates slightly the rms increase because drifts at frequencies lower than the spinning frequency are not totally cut out by the spin-chopping. The steeper the noise spectrum, the less accurate this method is. A rigorous calculation can be found in (Janssen et al. 1996) for the special case of $\alpha = 1$.

Although this additional noise power is quite low, it is at low frequency only, and the effect on the accuracy of the measurement of interesting cosmological quantities (e.g., the power spectrum of the fluctuations) should be investigated. For instance, it should be kept in mind that when the resolution of maps is degraded by smoothing, the standard deviation of pure white noise scales as the inverse of the size of the pixel, whereas the standard deviation of parallel stripes scales roughly as the square root of the size of the pixel. Quantifying the effect of striping on sophisticated statistical tests or pattern recognition methods is even harder, and may require the help of numerical simulations of two-dimensional noise maps.




Fig. 5. Residual of low frequency component of the noise only for each of the circles of Fig. 4. These structures, which can barely be distinguished in Fig. 4, would appear more and more clearly if the data were smoothed

3. From scan circles to sky maps

What is the effect of reprojecting data circles as those of Fig. 3 on sky maps? The answer depends on the statistical properties of the noise along the “data circles”, on the scanning strategy, and on the way that the relative

“offsets” of the data circles are readjusted with respect to one another.

3.1. Readjusting the offsets of the circles

It is clear from looking at Fig. 3 and from the above discussion that we should first find a way to adjust the relative values of the offsets A_i of the different data circles before reprojecting on the sky. Even for very low values of the knee frequency, there is no hope to keep the signal from diverging on time-scales of a few months. Just setting the average level of each circle to zero is not good enough, since there is no reason the average level of the useful signal on each circle should be the same, and the only measurement we have access to is signal+noise.

For PLANCK, the observing strategy ensures that, in addition to the short term redundancy provided by “spin-chopping”, there is also redundancy at long periods. In particular, as each data circle (obtained from 120 superimposed scans) crosses 3360 other such circles in two spots for a 1 year mission, it is quite natural, as a first order method for suppressing low-frequency noise effects, to try to readjust the average relative offset of the data circles by imposing that all the differences between signal over the same sky pixels but observed with different rotation axis of the satellite (at different times) are minimised. One key requirement on the scanning strategy for this method to work properly is that any given circle do not cut all the other circles in just a small number of pixels, in which case the adjusted value for the offset of that circle would depend drastically on the realization of the (white part of the) noise on the pixels used for the adjustment. For PLANCK, the scanning strategy is such that all the pixels of a circle are seen by at least one other circle, and thus all of them are used for the readjustment.

This method allows one to estimate the drifts due to low-frequency noise independently of the useful signal from the sky, since only differences of the total (signal+noise) at times where the beam is pointing on the same direction of the sky are used.

What should the accuracy of this method be? For a simple scanning strategy where the spin axis of the satellite is kept in the ecliptic, all circles play the same role (total symmetry), and thus $\sigma_{\text{stripes}}(i)$, the rms of the error on readjusting the offset for circle i , should be the same for all circles. If we denote by σ_n the rms of the noise per pixel on a single circle obtained by averaging 120 scans, the accuracy of the determination of the offsets will be of the order of $\sigma_n/\sqrt{N_j}$, where N_j is the number of independent pixels per circle.

It would be possible to optimise the scanning so that there be no preferred direction for this residual striping (which is anyway small, as for the PLANCK HFI N_j is about 2000 to 4500, depending on the channel), but it does not seem worth compromising the monitoring of other systematics and noise in the process.

Note that if we could build an ideal experiment with no low-frequency noise, the value for σ_{stripes} that we would estimate *just from sample variance on a circle of N_j pixels* would be $\sigma_{\text{stripes}}^2 = \sigma_n^2/N_j$, and thus no significant additional striping should be added by this inversion.

In order to check the above assertions about the accuracy of offset readjusting, and to evaluate the effect of low frequency noise on PLANCK maps, a low-resolution version of the expected PLANCK data has been simulated. For the moment, the simulation of the complete data stream for one detector at the actual PLANCK resolution is out of reach of our computers. However, the conclusions obtained with a degraded resolution can be scaled to the actual resolution.

In this simulation, a vector of spin-axis positions on the sky corresponding to a scanning strategy of PLANCK is generated. Each position of the spin-axis is distant from the previous one by a step equal to the resolution at which the simulation is performed, and for each such position, a vector of beam positions on a circle of radius 70° on the sky is generated. For each beam position, a single one degree by one degree pixel on a simulated map of the sky is selected. A two-dimensional set of data corresponding to these beam positions is generated. Here we assume that the knee frequency of the low frequency noise is small enough that there be no significant low-frequency contribution to the noise along individual data circles. We simulate the effect of low frequency drifts by adding to each data circle some offset A_i .

The actual data taking process is such that the zero level of the measurements is frequently readjusted, in order to avoid the saturation of the detectors due to slow drifts in the signal. In the end, the average value of the signal on each circle is not known precisely, but certainly does not drift by more than a couple orders of magnitude (extremely conservative assumption!).

The actual inversion is performed by the following least mean square method. The linear system of equations on the constants A_i to be determined is obtained by requiring a function $S(A_1, \dots, A_n)$ to be minimal. Here we take:

$$S(A_i) = \sum_{p \in \text{sky}} \sum_{i_1=2}^{n(p)} \sum_{i_2=1}^{i_1-1} w_{i_1, i_2} \times (\delta(i_1, i_2) - A_{i_1} + A_{i_2})^2 \quad (3)$$

where $\delta(i_1, i_2)$ is the estimated difference of offsets obtained by comparing the measurements of circle i_2 and circle i_1 on pixel p , $n(p)$ is the number of times pixel p is seen by the experiment, and w_{i_1, i_2} is a weight attributed to measurement $\delta(i_1, i_2)$. Since pixel p contributes $n(p) \times (n(p) - 1)/2$ terms in the sum, we set $w_{i_1, i_2} = 2/(n(p) - 1)$ (the total weight of the contribution of the $n(p)$ measurements at pixel p is thus proportional to $n(p)$).

We get the linear system to be inverted by writing that all partial derivatives of S with respect to all constants A_i should vanish. This system is degenerated, since if a set of constants A_1, \dots, A_n is solution, the same set with

a constant K added to all the A_i is also solution. This degeneracy is lifted by adding the requirement that the mean of the constants A_i vanish.

We first use the nominal scanning strategy of PLANCK, with a spin axis kept rigorously anti-solar, and check that the accuracy of the inversion does not depend on the input offsets for the circles, but does improve with the resolution. For each simulation, the data of the mission are stored as a table of N_i scan circles of N_j samples each. For a simulation with a 1 degree resolution, $N_i = 421$ and $N_j = 339$ for a 14 month mission and a 140° scan diameter. We evaluate the amount of striping by calculating the mean value of the signal on each scan circle, and computing the square root of the variance of the collection of N_i numbers so obtained. This is denoted as σ_{stripes} .

For a 1° resolution simulation, we generate a 421×339 element table of Gaussian random numbers with rms 1. This simulates the white noise of the mission. To this noise we add offsets which simulate the low frequency drifts. Various types of offsets were considered: a) no offset, in order to test that the method does not add striping where there is none; b) a linear drift of $100i/N_i$, where i is the index of the circle and N_i the number of circles; c) a sinusoidal drift of $100 \sin(2\pi i/N_i)$; d) a sequence of random offsets with an rms of 100 (hundred times the rms of the noise); e) a random walk of rms 10 per step; f) in order to check the intrinsic precision of the inversion, we try to recover the offsets of the random walk without adding any white noise.

In order to be able to compare the performance of the method in all cases above, let us consider one specific realization of the white part of the noise here, and add to it the offsets described above to get five different “noise signals”. In all cases a) through e), the value of σ_{stripes} after the inversion is 0.0412, which proves that the accuracy does not depend on the offsets to be corrected for to any significant level. It is interesting to realize also that if there is no offset, the expected value of σ_{stripes} before any inversion *just from sample variance* is 0.0543 (and on the particular noise realization we used here it happened to be 0.0522), so that the method does not only remove the striping due to drifts in the offsets due to low-frequency noise if any, but does even readjust to some extent the variations in the average level of the circles due to the sample variance of the average of the N_j points on a circle. This is anecdotic, of course, but does show that no additional noise or bias is generated in the readjusting process - in fact, for a 14 month mission the method partly suppresses low frequency components of the white noise itself. It is worth stressing that the success of the method is due for a large part to the fact that all the points of any circle contribute to the evaluation of the value of its offset, and not only two points at the North and South ecliptic poles. The precision of the inversion (performed in double-precision) is evaluated from the results of case f) above, for which we get a totally negligible residual striping of $3.83 \cdot 10^{-7}$.

In order to check that the accuracy does indeed depend on the resolution, we repeat the simulation of a) through e) for resolutions of 2° and 5° . In these cases, the value of σ_{stripes} after the inversion is 0.0549 and 0.0800 respectively, independent of the original stripes again, and below the value we get from sample variance on the original data. Note that the weight given to each term in the least-square sum above is critical to reach this accuracy. For instance, if we put the same weight to each term, too much importance is given to points on the circles that are close to the north and south ecliptic poles (because they are “seen” by many more circles than the ones in the ecliptic), and some residual striping of the order of magnitude of the white noise per pixel remains, because the value of the offset recovered depends most on the realization of the white noise at points near the poles.

We tried this inversion scheme with all kinds of offsets, with various kinds of scanning strategy and different mission durations (from 3 to 14 months). In all cases the inversion works extremely well. To be a little more illustrative, we show images of a noise generated according to method e) above, for a 1 year mission, at the resolution of 1° , reprojected on the sky, with no destriping treatment (Fig. 6) and after destriping (Fig. 7). For this simulation we used a scanning strategy for which the spin axis has been made to oscillate sinusoidally out of the ecliptic with an amplitude of 15° and a frequency of 8 oscillations per year, in order to maximise the sky coverage. Note the different scales for the amplitude of the structures that can be seen. No striping whatsoever remains on the map after the treatment.

3.2. Improving the method

3.2.1. Modelization of nominal PLANCK HFI noise

Whereas it has been made clear that our method is satisfactory to remove the striping due to low-frequency noise off maps obtained by PLANCK in the case where the knee frequency of the noise is low enough that all significant low-frequency noise appears in the form of different offsets for different circles on the sky, it is worth mentioning that if there is significant low-frequency noise power at frequencies higher than the spinning frequency then the method above may not be sufficient anymore: After relative offsets are readjusted, the next effect of low-frequency noise is to add very low level drifts along individual circles, as shown in Fig. 4 and Fig. 5, and one should investigate whether or not these fluctuations may cause a problem.

It is not easy to generate degraded resolution maps that preserve exactly both the visual aspect (i.e. do we see striping or not) and the statistical properties of the noise for real $1/f$ noise on all scales. For instance, the rms amplitude of low-frequency noise relative to white noise depends on the sampling frequency (the variance of white noise is proportional to the sampling frequency, and the




Fig. 6. Map of reprojected noise and offsets if no processing of the signal is done. The map is heavily striped. This should not be interpreted as a noise map for PLANCK, as at least some offset readjustment can be made before reprojecting. For this data set, the average level of circles has been left to drift some hundred sigma away!

variance of $1/f$ noise diverges proportionally to the logarithm of the sampling frequency). For these simulations, we decide to scale both the white noise and low frequency noise rms to their rms values for 10 arcminute pixels, independently of the size of the pixel of the simulation (here 1 square degree pixels). The next problem is that for 1 degree resolution simulations, instead of one circle every two hours, we generate one circle every day. However, we wish to preserve the possible correlations between consecutive circles of the simulation. Thus, we have to face the problem that either the correlations between consecutive circles is right, or the relative value of the offset between circles taken at a 1 day time interval is right. As the paragraph above demonstrates that relative levels of the circles do not play a role in the accuracy of the readjustment of the offsets, and as offsets will not drift too far away because they will be readjusted by the instrument anyway, I decided, in this next part of the simulations, to preserve the correlation between consecutive circles rather than their relative offsets.

Thus, a time sequence for the noise is generated by the following method: for a 1 year mission at a 1 degree resolution, we need 361 circles (so as to re-observe the first circle at the end of the mission) of 339 points each. Each of these circles is obtained from averaging 120 scans, and thus we need to generate a noise of almost 15 million data points. We wish that these data points preserve the relative rms amplitudes of low-frequency noise to white noise of the full 10 arcminute resolution maps, i.e. we do not smooth the noise down to the lower resolution, but rather affect to each 1 degree by 1 degree pixel the noise of its “central” 10 arcminute by 10 arcminute pixel, which corresponds to reprojecting on a 1 degree resolution map a subset of the 10 arcminute resolution pixels. This method preserves best the relative rms of the low-frequency residual noise and the white noise part of true full resolution maps.

Generating a 15 million sample dataset with a “true” $1/f$ spectrum is out of the reach of the computer I used. In order to generate the $1/f$ noise, an under-sampled (by a factor 32) white noise with a unit variance was




Fig. 7. Map of reprojected noise and offsets after signal-preserving destriping (see text). Compare with Fig. 6 (note the different colour scale!)

generated, converted to Fourier space (by FFT) where it was multiplied by $(f_{\text{knee}}/f)^{\alpha/2}$ (with the appropriate rescaling to take into account the effect of under-sampling). Then we compute the inverse FFT and interpolate between the samples, add consecutive circles by packs of 120 to get a 361×339 array of data points (corresponding to 361 circles), and add to it a randomly generated array of 361×339 data points with a Gaussian statistic and a rms of $1/\sqrt{120}$. We then re-scale everything conveniently by multiplying by $\sqrt{120}$. We check that the relative rms values and the visual aspect of the circles obtained in this way are correct by comparing with circles obtained from fully fast-sampled simulated low frequency noise of a few circles only.

First, let us generate such a noise for nominal low-frequency noise parameters for the PLANCK SURVEYOR High Frequency Instrument, i.e. $f_{\text{knee}} = 0.01$ Hz and $\alpha = 1$, and readjust the offsets by the method of the previous sub-section. Again we use the scanning strategy for which the spin axis has been made to oscillate sinusoidally out of the ecliptic with an amplitude of 15° and a frequency of 8 oscillations per year. The resulting output noise map is shown in Fig. 8. No striping

whatsoever is apparent on the map, and the increase of the rms of the noise is about 0.47%. In order to check the effect of the out-of-ecliptic motion, we do the simulation with the same realization of the noise but with a nominal anti-solar spin axis. In this case the increase in the rms of the noise is 0.49% (no significant difference). Repeated simulations with different noise realizations show that the small difference is always in favour of sinusoidal out-of-ecliptic motion.

In their 1996 paper, Janssen et al. computed the maximum noise increase along a scan circle. Adapted to our notations, their formula can be written as

$$F_{\text{max}} \simeq \left(1 + \frac{T_{\text{spin}} f_{\text{knee}}}{m} (2 \ln m + 0.743) \right)^{1/2} \quad (4)$$

where m is the number of independent data points along one circle. For PLANCK, $m = 2034$ for a 10 arcminute resolution, $T_{\text{spin}} = 60$ sec., $f_{\text{knee}} = 0.01$ Hz, and we get $F_{\text{max}} \simeq 1.01$. The average additional noise that we find by simulations is about half the maximum noise increase along one circle they predict (for diametrically opposed pixels), so our result is fully consistent with theirs. This




Fig. 8. Map of reprojected simulated $1/f$ noise, $f_{\text{knee}} = 0.01$ Hz, $\alpha = 1.0$, after correction by simple offset readjustment. The very-low level residual striping can not be distinguished by eye. This is a 1-degree resolution projection of what the processed PLANCK all-sky noise could look like. The relative standard deviations of reprojected white noise and reprojected residual striping are those of a 10 arcminute resolution mission (see text)

is consistent with most of the additional noise being due to small drifts along individual circles.

In order to give an estimate of the order of magnitude of striping and the statistical properties of the noise, we show, in Fig. 9, plots of cuts through the noise map of Fig. 8. On both panels of Fig. 9, total noise is represented with diamonds, and the component of noise due to residual low-frequency drifts (obtained by computing the difference between the output map of Fig. 8 and the map obtained by simple reprojection on the sky of the white part of the noise used for the simulation only) as a plain line. In these plots, although one point only is plotted per degree, the spread in represented points is typical of what we would get in a 10 arcminute resolution map, not a smoothed map with a 1 degree resolution. The top panel corresponds to a vertical cut in the middle of the map, and the bottom panel to an horizontal cut in the middle of the map (thus perpendicular to the expected striping, if any). In both cases the total contribution of striping to the total noise is very small. The structure of residual low-frequency noise is not similar in both directions, which can be understood

from the direction of scans. The small structure on the plot of residual low-frequency noise in the middle of the top panel is characteristic of very-low level striping. Finally, it is obvious from the distribution of the total noise that some regions of the sky are integrated more than others (as the region around ecliptic longitude $\lambda = 0^\circ$ and ecliptic latitude $\beta = -45^\circ$).

3.2.2. Monitoring unforeseen instabilities

Now we want to investigate what would happen if the low frequency noise were much worse than expected. The purpose of the following simulations, using parameters that are unrealistically pessimistic for the PLANCK HFI, is to show that the conclusions of the previous paragraphs do not depend drastically on noise assumptions, and that even unforeseen instabilities can be monitored quite well with PLANCK. Here we assume that some unwanted temperature fluctuations (for instance) generate low frequency noise with the parameters of $f_{\text{knee}} = 0.1$ Hz (ten times the nominal!) and $\alpha = 2.0$ (same parameters as were used

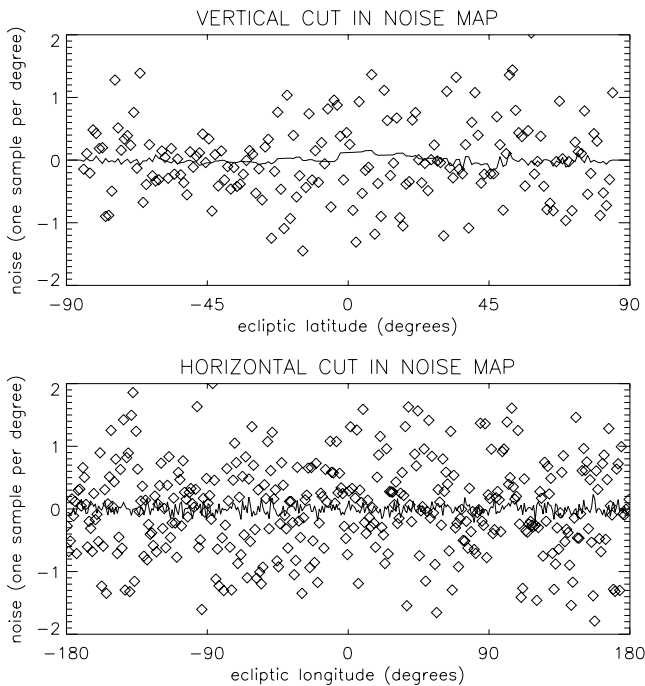


Fig. 9. Cuts through the noise map of Fig. 8. The top panel is a vertical cut at 0° ecliptic longitude, and the bottom panel an horizontal cut along the ecliptic plane. The plain line represents residual striping, obtained by computing the difference between the map of reprojected total noise after inversion, and the map of reprojected white part of the noise only. Diamonds correspond to the total reprojected noise. The structure due to striping is very small compared to the total noise, and the noise r.m.s. increase due to this residual striping is 0.47% on a 10 arcminute resolution map of the sky

to generate the plots of Figs. 2-5). The resulting increase in the noise standard deviation of the map (scaled to 10 arcminute pixel size as explained above) is 7.5%. This is not much, but it is significantly larger than the value of 0.5% obtained with nominal HFI noise. Because the excess power is not white, it could be distinguished on the maps, especially at degraded resolution. For instance, within the framework of standard assumptions of physical cosmology this could cause trouble for the estimation of the properties of primordial supra-horizon fluctuations, which are important for constraining the models of inflation. It could also be annoying in pattern recognition methods looking for discontinuities generated by cosmic strings. Finally, some optimal foreground separation methods rely on statistical methods which use prior knowledge of the spectrum (as a function of scale) of the various astrophysical components (i.e. cirrus clouds, primary CMB, free-free, synchrotron, etc...).

In order to get rid of this residual striping, we can try a more sophisticated treatment. The idea again is to find a method which does not depend on the real signal, and thus estimates low-frequency components by using signal

differences on common pixels. To do so, we can adapt the above method: instead of fitting just one constant for each data circle, we fit a function with more parameters, so that along a scan circle i we may write

$$n_i(j) \sim \sum_k A_{ik} f_k(j). \quad (5)$$

In the equation above, j indexes the samples along the data circle, and the functions $f_k(j)$ are vectors of a basis of functions on which to decompose the noise $n_i(j)$. Typically, the set of functions f_k can be sines and cosines (Fourier modes), or polynomials, or other well-chosen functions. The sum we want then to minimise is:

$$S(A_{ik}) = \sum_{p \in \text{sky}} \sum_{i_1=2}^{n(p)} \sum_{i_2=1}^{i_1-1} w_{i_1, i_2} \times (\delta(i_1, i_2) - F_1 + F_2)^2 \quad (6)$$

where

$$F_1(i_1, j_1) = \sum_k A_{i_1 k} f_k(j_1(p)) \quad (7)$$

and

$$F_2(i_2, j_2) = \sum_k A_{i_2 k} f_k(j_2(p)) \quad (8)$$

where $j_1(p)$ and $j_2(p)$ index the samples on circles i_1 and i_2 respectively for which pixel p on the sky is observed.

In this framework, for instance, we can take advantage of the fact that the most interesting property of low-frequency noise is that it does not have significant high frequency power. Because of that, low frequency noise itself can be estimated by sampling it at a much lower sampling frequency than the true signal. Thus, it seems to be a good idea to use as a basis of functions for noise estimation along one circle a set of a few ‘‘top-hat’’ functions, corresponding each to a constant on a fraction of a circle only.

Using the same realization of low-frequency and white noise, we inverted the data set by adjusting more than one constant for each circle. We do it for two constants (i.e. one for each half-circle), three constants (one for each third of a circle), and four constants (one for each fourth of a circle). We performed the inversion also using 1st through 3rd degree polynomials.

Table 1 gives the noise rms increase on maps in all the cases discussed above. Maps, too numerous to be shown here, can be provided by the author upon request. It is clear that the residual striping can be reduced substantially by this method (and the more so at high resolution, as more points are available to estimate individual parameters of the fitting functions f_k , whereas no more functions are needed at high resolution than at low resolution to estimate low-frequency structures in the noise). 4 constants per circle instead of one is the best that could be done because of computational limitations. For full-resolution

data sets, with a good computer, going to 10 constants or so should be possible and should improve the fit significantly. This method, used on very pessimistic noise here for illustrative purposes, could also be applied to destripe further full-resolution maps obtained with nominal PLANCK HFI noise, if one wished to do so.

Table 1. Increment of noise rms from residual striping for very pessimistic low-frequency drifts (assumed knee frequency $f_{\text{knee}} = 0.1$ Hz and spectral index $\alpha = 2.0$). Different sets of fitting functions have been used to remove striping effects (see text). Polynomials seem to perform better up to 3 constants per circle (second degree polynomial). This can be understood from looking at the shape of drifts along circles (Fig. 5). A polynomial of degree 3 is worse as far as reducing the rms of the noise is concerned for 1 degree resolution simulations. This is due to insufficient constraints and to the fact that the weighting of the terms in the least square sum is not adapted to such functions, but to step functions. In fact the 3rd degree polynomial induces a dipole on the sky, which can be understood as high degree polynomial fits have a known tendency to oscillate between the data points that carry the most weight. Most of the additional power comes from this dipole, although it is much lower in amplitude than the “real” physical (cosmological or Doppler-induced!) dipole. Maps can be made available by the author upon request

FIT	steps	polynomials
1 constant	+ 7.5%	+ 7.5%
2 constants	+ 3.3%	+ 2.9%
3 constants	+ 3.2%	+ 2.3%
4 constants	+ 2.9%	+ 6.3%

For each noise spectrum there must be an optimal set of functions f_k to use. For instance, a 1st degree polynomial is better than two constants per circle for very steep noise spectra, as most of the low-frequency noise contribution comes from very low frequencies. For each noise spectrum, there is also an optimal number of functions f_k to use, as the more functions one uses, the less constraints one gets on each of the functions. These optimal solutions are yet to be found.

4. Discussion and comments

4.1. Finding an optimal scanning strategy?

It is clear that for the methods discussed above to work there must be some optimal way of scanning the sky. The scanning strategy sets the position of the points of intersection of the different circles described by the main beam of PLANCK SURVEYOR.

For instance, scanning the sky with great circles while the spin axis is kept rigorously anti-solar insures that every circle crosses every other circle in exactly two points,

namely the North and South ecliptic poles. For this scanning strategy, pixels close to the ecliptic are seen by only one circle (obtained from adding together a great number of scans, 120 scans for PLANCK) every 6 months.

If instead of great circles the main beam scans 140 degree diameter circles while keeping the spin axis anti-solar, as for the nominal PLANCK mission, there is no common “reference” pixel seen by all circles. Circles cross each other on points that are spread out all over their length. Circles observed within a small time difference cross at high ecliptic latitudes. Circles observed at a time interval of about 140 days cross near the ecliptic plane, as shown in Fig. 1. For readjusting offsets, this scanning strategy should be superior, as all the pixels along one circle contribute to the evaluation of the offset of that circle.

This is not the case for great circles, and as emphasized by Wright (1996), the estimated offset for a great circle that crosses reference circles in only two points will depend on the realization of the noise at these two points. Of course, one could rely on more than two points, by using also points along the scan close to the pole which overlap substantially (and thus compare measurements in more extended “polar caps”). Some significant improvement can probably be obtained in this way, as after all consecutive great circles have a significant area in common, but one may have to worry about high spatial frequency signals on the sky (especially point sources) for fields of view that do not exactly coincide. Iterative correction of oversampled maps could be a solution, but this may lead to complications. Finally, this scanning strategy using great circles with anti-solar spin axis does not allow the natural improvement of the method where the low-frequency noise realization along each circle is fitted by sampling it at its own Nyquist frequency, as there are points of the circle (near the ecliptic plane) which are not seen by any other circle (here we talk about circles obtained by adding consecutive scans, not about individual scans).

In both cases, the destripping method can be improved by allowing the spin axis to move away from the anti-solar. This is especially true if great circles are used to scan the sky, as then circles would not cross other circles in two points only. A small displacement of the spin axis of a circle along the ecliptic makes it cross other circles near the ecliptic poles, and a small displacement of the spin axis perpendicular to the ecliptic makes the circle cross nearby circles near the ecliptic plane. Thus, in order to insure that all points along a circle have other circles to be compared to, a reasonable solution *as far as destripping is concerned* could be to move the spin axis on circles around the anti-solar. This has the additional characteristic that the sun aspect angle is kept constant, which may help monitoring the thermal stability of the payload if the satellite is reasonably symmetric around its spin axis. Sinusoidal or tooth-saw motion out of the ecliptic while the spin axis turns around the ecliptic are also possible solutions.

It has been argued by some authors (Wright et al. 1996) that based on the inherent difficulty of stabilising the sensitivity of an instrument to enough precision, it is desirable for the raw data to be collected in differential form. However, we would like to stress that even differential measurements can potentially suffer from low-frequency noise, because any asymmetric source of fluctuations (as for instance thermal fluctuations of optical elements which are not commonly shared) generates low-frequency instabilities in the measurements. Differential measurements are not just desirable: in a way or another they are, on short time scales, a necessity for radiometers, less stable than bolometers by orders of magnitude for CMB applications. These authors proved that it is possible to invert megapixel differential data. This task may be much harder if in addition some unforeseen low-frequency drifts contribute to the data stream.

Finally, whereas the method works better if the spin-axis is not kept rigorously anti-solar, it should be kept in mind that moving the spin axis could generate increased thermal instabilities of the payload or increased sidelobe contamination. For PLANCK, it has been shown that thermal specifications could be fulfilled if this angle does not exceed 15 degrees, which is sufficient for destriping and for sky integration time optimisation if desired. The freedom to move the spin axis also allows one to readjust the scanning strategy after the preliminary analysis of the first few days or weeks of data if necessary. This flexibility is a powerful tool for monitoring systematic effects.

4.2. Comments on the method

Our method has the very nice property that the correction of low frequency drifts does not require the inversion of a huge matrix, contrarily to methods which try to make a least square fit of both the signal on all pixels of the sky and all additional parameters as those which allow for the correction of low-frequency drifts. A (relatively!) small matrix of a few thousands by a few thousands needs to be inverted, which is an easy task for modern computers.

However, the method relies heavily on the assumption that different measurements on the same pixel of the sky should generate the same “useful” signal, apart from different contributions of low-frequency drifts. If this is not the case, it is still possible to fit the noise in much the same way, but this will require the inversion of huge matrices of typically a few million by a few million entries. Iterative methods make this task possible if there are not too many non-vanishing elements. Potential sources of trouble include far sidelobe straylight, which depends on the orientation of the satellite and on the position of bright sources in the sky, polarisation (for detectors sensitive to polarisation, because again what is measured on a given pixel with a polarisation-sensitive detector depends on the orientation of the satellite, as shown in Fig. 10), and all other

potential systematic effects. The remark about the complications induced by the measurement of polarised light is also true when one wants to invert differential data (e.g. Wright et al. 1996). These authors’ method for inverting the differential signal expected from the MAP satellite also relies on the assumption that the useful signal from a given pixel is independent of satellite orientation. This assumption is not correct for polarisation sensitive measurements, obviously, and again this can make the inversion significantly harder than expected.

Bolometers are not sensitive to polarisation unless one places a polariser in front of them.




Fig. 10. Representation of the effect of polarisation-sensitive measurements on the comparison of relative measured signal on different scans. Small circles with arrows inside represent different fields of view of the instrument (main beam). On the common pixel in the middle of the figure, measurements on different scans do not see the same polarised light, as indicated by the arrows. This may complicate attempts to measure polarisation and/or the inversion of polarised differential data

5. Conclusion

In this paper, some aspects of the expected statistical properties of the noise for the future PLANCK satellite, both on data streams and on output maps, have been analysed. It has been shown that for the bolometer instrument of PLANCK, no significant striping on the maps is expected due to low-frequency noise only for the nominal knee frequency of the noise of 0.01 Hz, as a simple destriping method can remove low-frequency noise effects on maps. For the PLANCK HFI the increase in the noise rms on destriped maps is expected to be less than half a percent. Even if the low-frequency noise were much worse

than expected, correction for its effects are possible with the scanning strategy of PLANCK, contrarily to what has been suggested by Wright (Wright 1996).

However, we believe that testing the extraction of cosmological information using simulated noise maps is probably one of the best way to quantify the exact consequences of instrumental effects such as striping, and continuing efforts should be made towards simulating realistic noise maps in the future.

An optimised scanning strategy must be devised by including simulations of most foreseen systematics. We believe that this optimum can be found within the orbit specifications for PLANCK, which minimise most of the potentially harmful systematic effects.

Acknowledgements. Valuable comments from G. Smoot and helpful discussions with F.-X. Désert, R. Gispert, J.-M. Lamarre and J.-L. Puget are gratefully acknowledged. This work was supported by a PhD research grant from CNES, the French national space agency.

References

- Benoit A., et al., SAE Technical Paper Series 941276
COBRAS/SAMBA report on the phase A study: this report summarises the results of the scientific and technical study activities of the COBRAS/SAMBA phase A; copies can be obtained from J. Tauber or O. Pace, both at ESTEC, Keplerlaan 1, P.O. Box 299, 2200 AG Noordwijk, The Netherlands, or from S. Volonté, Programme Coordination and Planning Office, ESA HQ, D/SCI, 8-10, rue Mario Nikis, 75738 Paris Cedex 15, France
Gaertner S., et al., 1997 (accepted for publication in A&A)
Janssen M.A., et al., 1996 (submitted to ApJ) (astro-ph 9602009)
Smoot G.F., 1992, ApJ 396, L1
Wright E., paper presented 22 November 1996 at the IAS CMB Data Analysis Workshop (astro-ph 9612006)
Wright E.L., Hinshaw G., Bennett C.L., 1996, ApJ 458, L53