

Making sky maps from Planck data

M. A. J. Ashdown^{1,2}, C. Baccigalupi^{3,4}, A. Balbi⁵, J. G. Bartlett⁶, J. Borrill^{7,8}, C. Cantalupo^{8,7}, G. de Gasperis⁵, K. M. Górski^{9,10,11}, E. Hivon^{10,14}, E. Keihänen^{12,13}, H. Kurki-Suonio¹², C. R. Lawrence⁹, P. Natoli⁵, T. Poutanen^{12,13}, S. Prunet¹⁴, M. Reinecke¹⁵, R. Stompor^{7,8,6}, and B. Wandelt^{16,17}
(The Planck CTP Working Group)

¹ Astrophysics Group, Cavendish Laboratory, J. J. Thomson Avenue, Cambridge CB3 0HE, UK

e-mail: maja1@mrao.cam.ac.uk

² Institute of Astronomy, Madingley Road, Cambridge CB3 0HA, UK

³ Institut für Theoretische Astrophysik, Universität Heidelberg, Albert-Überle-Str. 2, 69120 Heidelberg, Germany

⁴ SISSA/ISAS, Via Beirut 4, 34014 Trieste, and INFN, Sezione di Trieste, Via Valerio 2, 34127, Italy

⁵ Dipartimento di Fisica, Università di Roma “Tor Vergata”, via della Ricerca Scientifica 1, 00133 Roma, Italy

⁶ Laboratoire Astroparticule & Cosmologie, 11 place Marcelin Berthelot, 75231 Paris Cedex 05, France
(UMR 7164 CNRS, Université Paris 7, CEA, Observatoire de Paris)

⁷ Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley CA 94720, USA

⁸ Space Sciences Laboratory, University of California Berkeley, Berkeley CA 94720, USA

⁹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena CA 91109, USA

¹⁰ California Institute of Technology, Pasadena CA 91125, USA

¹¹ Warsaw University Observatory, Aleje Ujazdowskie 4, 00478 Warszawa, Poland

¹² University of Helsinki, Department of Physical Sciences, PO Box 64, 00014 Helsinki, Finland

¹³ Helsinki Institute of Physics, PO Box 64, 00014 Helsinki, Finland

¹⁴ Institut d’Astrophysique de Paris, 98 bis boulevard Arago, 75014 Paris, France

¹⁵ Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

¹⁶ Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana IL 61801, USA

¹⁷ Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 West Green Street, Urbana IL 61801, USA

Received 14 June 2006 / Accepted 9 February 2007

ABSTRACT

Aims. We compare the performance of multiple codes written by different groups for making polarized maps from Planck-sized, all-sky cosmic microwave background (CMB) data. Three of the codes are based on a destriping algorithm; the other three are implementations of an optimal maximum-likelihood algorithm.

Methods. Time-ordered data (TOD) were simulated using the Planck Level-S simulation pipeline. Several cases of temperature-only data were run to test that the codes could handle large datasets, and to explore effects such as the precision of the pointing data. Based on these preliminary results, TOD were generated for a set of four 217 GHz detectors (the minimum number required to produce I , Q , and U maps) under two different scanning strategies, with and without noise.

Results. Following correction of various problems revealed by the early simulation, all codes were able to handle the large data volume that Planck will produce. Differences in maps produced are small but noticeable; differences in computing resources are large.

Key words. cosmology: cosmic microwave background – methods: data analysis

1. Introduction

Cosmic microwave background (CMB) observations have driven a remarkable advance in cosmology over the past decade (Smoot et al. 1992; de Bernardis et al. 2000; Hanany et al. 2000; Benoît et al. 2003; Bennett et al. 2003 and references therein), and will continue to furnish invaluable data in the years to come. As the data volume and precision demanded of these observations increases (for example, Bond et al. 1999), the complexity of the analysis methods required to deal with the data increases also. Map-making – the process of turning time-ordered scan data into an image of the sky – is an example of a crucial step whose technical complexity has grown significantly. This is particularly so in the case of total power measurements, where one removes signal drifts due to $1/f$ -spectrum noise in the map-making step. If left unchecked, these drifts leave stripes in the final map with amplitudes greater than the cosmic signal,

potentially compromising the scientific goals of a precision instrument such as Planck. For example, in the simulations presented below, the magnitude of the striping signal (estimated from the rms difference between a simple coadded map and the output map from one of our codes) was $336 \mu\text{K}$, more than three times the output map’s residual rms error of $\sim 100 \mu\text{K}$ due to white detector noise (see Table 5). (These numbers are for $1/7$ pixels and four polarized detectors. For $5'$ pixels, corresponding to the resolution of the detectors, and for the full set of twelve detectors, the residual error is $\sim 15 \mu\text{K}$ for the temperature map.) In other words, map-making effectively removed striping with three times the target sensitivity. Proper map-making is thus crucial to mission objectives.

Planck¹, to be launched in 2008, will be the third-generation satellite dedicated to observations of CMB anisotropies.

¹ <http://www.esa.int/science/planck/>

Its primary objective is to measure the temperature anisotropies to the cosmic variance limit out to multipoles $l > 2000$; other scientific goals include detailed measurements of the polarized power spectrum, the extraction of catalogs of galaxy clusters and extragalactic sources, searches for non-Gaussianity, and in-depth studies of the Galaxy. To achieve these goals, Planck will image the sky in nine frequency bands, with resolution and sensitivity in the CMB-dominated bands of $5\text{--}15'$ and $5\text{--}10\ \mu\text{K}$, respectively. Crucial to the success of the mission is the production of sky maps approaching the instrumental white-noise limit; drifts and artifacts must be removed and the noise properties well-understood. The formidable challenge of doing this for maps containing millions of pixels lies at the heart of the effort of one of the Planck working groups, the CTP Working Group.

It is important to develop and test well before launch efficient algorithms for inclusion in the data analysis pipeline. Besides preparing the pipeline, this helps to identify potential sources of systematic error and inform mission operations (e.g., scanning strategy). It also allows us to better quantify the mission's expected scientific output.

In this paper, we evaluate a suite of map-making techniques using simulations of several channels of the Planck High Frequency Instrument (HFI) 217 GHz time-ordered data (TOD). The simulations model non-white noise and primary CMB temperature and polarization anisotropies. The suite of methods includes both destriping and optimal map-making algorithms. We gauge the quality of our recovered temperature and polarization maps by looking at rms pixel residuals, and the residual power spectrum. This gives us an evaluation through second order statistical measures of noise and artifact residuals. The complexity of the problem requires (at this stage) that we impose a number of simplifying assumptions; these are clearly spelled out in the text and are the focus of on-going work. We emphasize that the ability to produce the maps shown here is a notable achievement requiring intensive computation.

1.1. Planck scanning strategy

Planck will make its observations from the 2nd Earth-Sun Lagrange point, approximately 1.5×10^6 km from the Earth. The satellite is spin-stabilized, and during science observations it will rotate on its axis once per minute. The telescope points at angle of 85 degrees to the spin axis, so the detectors follow small circles on the sky. The satellite will perform a repointing manoeuvre once per hour to keep the spin axis close to the anti-solar direction. Thus during the one hour periods between manoeuvres the detectors will make repeated observations of the same "rings" on the sky. Some of the algorithms presented in this paper can take advantage of these repeated observations to reduce the computational burden.

Within the constraints imposed by Planck's design, there is freedom to choose the precise pointings of the spin axis to optimize the scientific returns of the mission. The choice of spin axis pointings – the scanning strategy – is one of the factors we examine in this paper.

1.2. Planck science goals for sky maps

Planck is designed to image the sky at nine frequencies from 30 to 857 GHz, with angular resolution from 33 to 5 arcmin. The raw sensitivity is sufficient that inferences about the underlying distribution of fluctuations on the sky should be limited not by noise, but rather by cosmic variance. To achieve this state,

systematic errors and processing artifacts must be controlled to microkelvin levels. It is the latter challenge that we are addressing in this paper and its predecessor (Poutanen et al. 2006).

2. Map-making algorithms

Maps of the Cosmic Microwave Background (CMB) sky signal are derived from long time series of data. These data are generally collected by telescopes fitted with detectors based on either HEMT amplifiers or bolometers. Both detection systems have noise characterized by a power spectrum that rises at low temporal frequencies (often referred to generically as "1/f noise", even though this is only approximate at best). These systems are better-suited to differential measurements, where a short period of time lapses between two measurements being differenced, than to absolute ones. To achieve this, such CMB telescopes are scanned rapidly across the sky, with the scanning pattern crossing over itself multiple times to enable determination of the low frequency noise in the time stream.

2.1. The problem

We assume that the data \mathbf{d} recorded from a detector can be written as the sum of two contributions, one from the sky signal \mathbf{s} and the other from the noise \mathbf{n} in the detection chain

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{n}, \quad (1)$$

where the pointing matrix \mathbf{A} describes the path of the detector across the sky. If we have data from more than one detector, we can treat all of the data together by concatenating their data vectors and pointing matrices. The sky signal is represented by a vector with discrete entries which are the pixels in a map. This implies that the signal has a constant value across each pixel, so the pixels of the map must be smaller than instrumental beam by a sufficient amount to satisfy this assumption. If the detectors only measure the intensity of the sky emission, then the signal vector contains one entry in each pixel p , giving the value of the I Stokes parameter, s_{pI} . If, however, the detectors also measure linear polarization, then the signal vector must contain three entries for the I , Q , and U Stokes parameters in pixel p , s_{pI} , s_{pQ} , and s_{pU} .

If at time-index t an unpolarized detector is pointing at pixel p , then the datum recorded from the detector is

$$d_t = s_{pI} + n_t, \quad (2)$$

where n_t is the noise contribution to the datum. If the detector is sensitive to polarization, the datum recorded is

$$d_t = s_{pI} + s_{pQ} \cos(2\chi_t) + s_{pU} \sin(2\chi_t) + n_t, \quad (3)$$

where χ_t is the angle of the detector's polarization direction with respect to the polarization basis in that pixel. It can be seen from (2) and (3) that the pointing matrix \mathbf{A} is very sparse. Each row contains only one or three non-zero entries, depending on whether the detector is polarized.

The assumptions above imply that there is no attempt to deconvolve the instrumental beam during the map-making process. If the beam is symmetrical, then the resulting map will be smoothed with the same beam. If instead the beam is asymmetrical, then the effective smoothing of the map will vary with position. The smoothing at a particular position will depend on the orientations in which the detector has passed over that point in the sky. If the detector visits all points on the sky in all orientations, then the effective smoothing of the map will be given

by a symmetrized version of the beam (Wu et al. 2001). This is not the case for the Planck scanning strategy which is highly inhomogeneous.

It is assumed that the statistical properties of the noise \mathbf{n} are known a priori. It is possible to perform a joint estimation of the signal and the noise properties, but this is not addressed in this paper. Here we assume that the noise is Gaussian with zero mean and covariance \mathbf{N} ,

$$\begin{aligned}\langle \mathbf{n} \rangle &= 0, \\ \langle \mathbf{n}\mathbf{n}^T \rangle &= \mathbf{N},\end{aligned}\quad (4)$$

where $\langle \cdot \rangle$ denotes the average over noise realisations. It is usually assumed that the noise properties are stationary or piecewise stationary in the time domain. Because of its stationarity the noise can be described by a correlation function in the time domain or equivalently by a power spectrum in the (temporal) frequency domain. Thus each stationary block of the noise covariance matrix is a symmetric Toeplitz matrix (Golub & van Loan 1996).

2.2. Optimal solution

The optimal solution for the map is given by the maximum-likelihood estimate $\hat{\mathbf{s}}$ which is obtained by solving the generalized least-squares (GLS) equation

$$\left(\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A}\right) \hat{\mathbf{s}} = \mathbf{A}^T \mathbf{N}^{-1} \mathbf{d}.\quad (6)$$

The term in brackets in the left-hand side is an $N_{\text{pix}} \times N_{\text{pix}}$ matrix, where N_{pix} is the number of pixels in the map. It is the inverse of the pixel-pixel covariance matrix of the resulting map,

$$\mathbf{S} = \left(\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A}\right)^{-1}.\quad (7)$$

For Planck-resolution maps, it is impossible in practice to calculate \mathbf{S}^{-1} , let alone invert it, so (6) is solved using iterative methods. The implementations of optimal map-making tested in this paper all use the preconditioned conjugate gradient (PCG) method (Golub & van Loan 1996) to find the solution. In each iteration of the PCG method, it is necessary to apply $\mathbf{S}^{-1} = \mathbf{A}^T \mathbf{N}^{-1} \mathbf{A}$ to a vector. This can be achieved by using its factors in the following procedure:

1. Use the pointing matrix, \mathbf{A} , to project the map-domain vector into the time domain.
2. Apply the inverse noise covariance matrix \mathbf{N}^{-1} to the time-domain vector.
3. Use the transpose of the pointing matrix \mathbf{A}^T to map the time-domain vector back into the map domain.

Since the pointing matrix is sparse, steps (1) and (3) can be achieved very efficiently in $O(N_{\text{tod}})$ operations, where N_{tod} is the number of TOD samples. Step (2) can be applied by taking advantage of the piecewise-stationary properties of the noise. The stationary blocks of the noise covariance matrix \mathbf{N} are symmetric Toeplitz matrices, and it is assumed that they are large enough so that their inverses are well-approximated as Toeplitz matrices too (Stompor et al. 2002). This being so, the blocks can be applied using a convolution that is evaluated using a fast Fourier transform (FFT). The computational complexity of this step is $O(N_{\text{tod}} \log N_{\text{tod}})$. In practice, the convolution kernel is truncated to a length N_τ where the long-range correlations are dying away, so the computational complexity is reduced to $O(N_{\text{tod}} \log N_\tau)$. The convolution is applied using the overlap-add or overlap-save

methods (Press et al. 1992). Since the FFT is the most time-consuming of these operations, all of the optimal map-making codes in this paper use the highly-performant and portable FFTw library (Frigo & Johnson 1998).

The preconditioner used in the PCG solution is obtained by approximating the noise covariance matrix \mathbf{N} as diagonal,

$$\mathbf{N} \sim \mathbf{N}_u = \text{diag}(\dots, \sigma_t^2, \dots)\quad (8)$$

where σ_t is the noise standard deviation at time index t . Thus the preconditioner matrix

$$\mathbf{M} = \mathbf{A}^T \mathbf{N}_u^{-1} \mathbf{A}\quad (9)$$

is diagonal in the unpolarized case and block diagonal in the polarized case.

2.3. Approximate solution: the destriping approach

In calculating the optimal solution to the map-making problem, many computationally-expensive operations are required on data in the time domain. If these could be reduced or eliminated by making approximations of the noise, then the map-making process could be made faster.

In the destriping approach (Burigana et al. 1997; Delabrouille 1998; Maino et al. 1999, 2002; Revenu et al. 2000; Keihänen et al. 2004), the noise is divided into a low-frequency component represented by a number of basis functions multiplied by coefficients and a high-frequency part which is uncorrelated,

$$\mathbf{n} = \mathbf{B}\mathbf{x} + \mathbf{n}_u.\quad (10)$$

The coefficients \mathbf{x} are multiplied the matrix \mathbf{B} containing the basis functions which are usually assumed to be a series of offsets. The noise covariance matrix (8) is diagonal, though its diagonal elements need not be equal. The data are now given by

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{B}\mathbf{x} + \mathbf{n}_u.\quad (11)$$

It is thus possible to write the map-making problem in a maximum-likelihood form (Keihänen et al. 2004), where the parameters to be solved for are \mathbf{s} and \mathbf{x} . The maximum-likelihood coefficients $\hat{\mathbf{x}}$ can be found from the TOD by solving

$$\left(\mathbf{B}^T \mathbf{N}_u^{-1} \mathbf{Z} \mathbf{B}\right) \hat{\mathbf{x}} = \mathbf{B}^T \mathbf{N}_u^{-1} \mathbf{Z} \mathbf{d},\quad (12)$$

where

$$\mathbf{Z} = \mathbf{I} - \mathbf{A} \left(\mathbf{A}^T \mathbf{N}_u^{-1} \mathbf{A}\right)^{-1} \mathbf{A}^T \mathbf{N}_u^{-1}.\quad (13)$$

The GLS equation (12) is solved using the conjugate gradient method with or without a preconditioner. The offsets are then subtracted from the TOD, which has the effect of “whitening” the noise, so the map can be obtained by binning the resulting TOD

$$\hat{\mathbf{s}} = \left(\mathbf{A}^T \mathbf{N}_u^{-1} \mathbf{A}\right)^{-1} \mathbf{A}^T \mathbf{N}_u^{-1} (\mathbf{d} - \mathbf{B}\hat{\mathbf{x}}).\quad (14)$$

It is possible to incorporate prior information about the offsets in the destriping process (Keihänen et al. 2005). The drifts in the TOD are caused by the low-frequency part of the noise spectrum. The prior distribution of the noise is assumed to be Gaussian with zero mean (4), so the prior on the offsets must be Gaussian with a zero mean too. Thus the prior distribution of the offsets can be summarized by a covariance matrix \mathbf{P} which is related to the prior noise covariance matrix (5). Incorporating this into the solution for the offsets, (12) becomes

$$\left(\mathbf{B}^T \mathbf{N}_u^{-1} \mathbf{Z} \mathbf{B} + \mathbf{P}^{-1}\right) \hat{\mathbf{x}} = \mathbf{B}^T \mathbf{N}_u^{-1} \mathbf{Z} \mathbf{d}.\quad (15)$$

The prior covariance of the offsets can be derived from the properties of the noise spectrum or it can be estimated from the data.

2.4. Degenerate pixels

If the map being made is polarized, then it is possible that some pixels have not been observed by detectors in a sufficient number of orientations to well constrain all of the I , Q , and U Stokes parameters. If this is the case, then the sky signal estimate of all three components may not be well-defined, and the corresponding map-making equations are ill-conditioned. A procedure for eliminating the badly-observed pixels is thus required in order to solve the map-making equations in a general case.

Such a procedure can be based on the preconditioner matrix, \mathbf{M} (9). The matrix plays a crucial role in both the optimal and destriper approaches, and in the latter case, directly determines the numerical stability of the sky map computations (14). In the optimal case, the (numerically stable) invertibility of the matrix \mathbf{M} is a necessary (and for many scanning strategies, nearly sufficient²) condition for successful calculation of the map.

Thus the preconditioner matrix can provide a method common to the optimal and destripping approaches for eliminating badly-observed pixels. The preconditioner matrix is block-diagonal with 3×3 blocks that correspond to the approximate covariance of the I , Q , and U Stokes parameters of individual pixels. Using (3) and (8), it can be seen that the block of the matrix corresponding to pixel p is given by

$$\mathbf{M}_p = \sum_{d \in \text{dets}} \sum_{t \in p} \frac{1}{\sigma_{d,t}^2} \begin{bmatrix} 1 & c_t & s_t \\ c_t & c_t^2 & c_t s_t \\ s_t & c_t s_t & s_t^2 \end{bmatrix}, \quad (16)$$

where $c_t = \cos(2\chi_t)$ and $s_t = \sin(2\chi_t)$. The inner sum is over time-samples t of the detector d falling in the pixel p , and the outer one is over the detectors included in the data set. If the matrix is not invertible, then it is impossible to constrain the Stokes parameters for this pixel. We test for the matrix invertibility using its condition number, κ_p , (Golub & van Loan 1996), defined as,

$$\kappa_p \equiv \kappa(\mathbf{M}_p) = \frac{\max_{\rho_i \in \rho(\mathbf{M}_p)} (\rho_i)}{\min_{\rho_i \in \rho(\mathbf{M}_p)} (\rho_i)}, \quad (17)$$

where $\rho(\mathbf{M}_p)$ denotes the eigenvalues $\{\rho_i\}$ of the matrix \mathbf{M}_p . The condition number is thus calculated using an eigenvalue decomposition routine (Golub & van Loan 1996; Press et al. 1992). If it is larger than a chosen threshold, then the Stokes parameters are considered to be ill-constrained and the corresponding pixel is eliminated from the map. (Note that the lower bound on the value of the condition number is in our case equal to 2, corresponding to the perfectly isotropic observations of a given pixel.)

Some of the codes presented in this paper implement a post-processing phase in which further pixels are eliminated after the map-making. This step uses either a lower value of the block condition number threshold or an additional criterion using the

² The cases when this is not a sufficient condition are, for example, the ones when all the observations falling into a given pixel are nearly 100% correlated. Such cases can be readily avoided if each pixel on the sky is revisited multiple times on different time scales. The additional pixel excision criterion could then just read: $\Delta t > 1/f_{\text{knee}}$, with Δt denoting the time interval between the first and the last pixel observation. Note that for the scanning strategies and the characteristics of the instrument considered here, such degenerate cases are unlikely to happen, as whenever the pixel is observed from multiple, sufficiently different directions, that implies multiple visits of the pixel on a time scales comparable to or longer than the noise correlations.

absolute values of the block eigenvalues or diagonal elements. The pixels eliminated in this way are not so poorly-conditioned as to make the solution impossible, but they tend to produce large noise residuals and thus are undesirable in the final map.

Note also that for the optimal approach removing pixels prior to map-making needs to be done with care in order not to affect the continuity of the time ordered data (Stompor et al. 2002). In the optimal runs described here, we excised a pixel by setting the sky signal at the pixel to zero, and introducing a fictitious extra pixel with zero sky signal observed instead. In this way, we preserved the noise stationarity across the entire segments of the data, without a need for involved time stream processing. None of these complications is relevant for the destriper algorithm. In the destripping runs presented here, samples corresponding to the pixel to be removed are simply dropped from the map-making.

3. Map-making implementations

Six implementations of the two basic algorithms described in Sect. 2 were tested on simulated Planck data. Three are based on optimal map-making (MADmap, MapCUMBA, and ROMA), and three are based on destripping (Springtide, Polar, and MADAM). Common features of the codes are described in Sect. 2. In this section, we highlight implementation details and differences.

3.1. ROMA

ROMA (ROMA Optimal Map-Making) is a Fortran 95 parallel (MPI) implementation of an iterative optimal GLS map-making algorithm. The first implementation of ROMA (Natoli et al. 2001) showed that it is feasible to solve the optimal map-making Eq. (6) for a Planck-sized dataset using an FFT-based, PCG iterative solver. Over time, ROMA has evolved into a multichannel, polarization-capable code, which includes an iterative noise estimation stage. ROMA has been integrated at the Planck LFI Data Processing Center; it has also been used to analyze data from the last (2003) Antarctic flight of BOOMERanG. A detailed discussion can be found in (de Gasperis et al. 2005).

The noise was considered stationary over the whole timeline, and its power spectral density was assumed known *a priori*. The instrumental beam is assumed to be axisymmetric and common to I , Q , and U . This means that the code solves for a beam-smearred (I , Q , U) set of maps. The condition number of each block of the preconditioner matrix (9) is checked and, if deemed satisfactory, the block is inverted and the result used for preconditioning. The convolution is performed using the overlap save method (Press et al. 1992), exploiting the fact that the noise correlation function is truncated after a lag (equal to 16385 samples. Double precision is used throughout.

3.2. MADmap

MADmap, part of the MADCAP3 package, is a massively parallel implementation of the map-making equations (6); see Borrill et al. (2007) for a detailed description. Communications between processors are handled by MPI, and data are distributed over the memory associated with each processor.

The major data objects involved in the calculation are distributed over the memory associated with each processor in a balanced way. All time domain data, including the time stream vector, the pointing matrix, and the correlation functions that define the inverse time-time noise matrix, are distributed with an

even partition of the data over the processors so that every processor stores the data associated with an equal number of contiguous time samples. A single processor may store time stream data that belong to more than one data set (a data set is the data recorded by a single detector in the simulations described in this paper) if the processor stores the last samples from one data set and the first samples from the next.

In calculating the pixel-pixel noise correlation matrix with the PCG algorithm, five pixel-domain vectors must be stored. These vectors are distributed over the processors so that if a pixel is observed during the time interval associated with a given processor, then the vector elements associated with that pixel are stored in the processor's memory. This does not represent an even partition of the pixel domain, but it does limit the size of the pixel vectors stored on a given processor to be less than the size of the pointing matrix stored on the processor (with most experimental scanning strategies it will be significantly smaller). For the Planck scanning strategy in particular, the memory footprint required for the pixel vectors will be approximately 60 times smaller than the pointing matrix memory footprint.

The pointing matrix is stored as a sparse matrix in zero-compressed row major order; it can take on a variety of forms to reflect different models of the observation. If there are N_z non-zeros per row, then \mathbf{A} takes $O(N_z N_t)$ storage space and $O(N_z N_t)$ calculations to operate with, where N_t is the number of time samples. MADmap allows for an arbitrary pointing matrix, and this flexibility allows for very complicated time stream models. By using an arbitrary pointing matrix we can remove parasitic sky-asynchronous signals, calibrate from a known signal, pixelate a systematic effect and remove it, and so forth.

The inverse time-time noise correlation matrix operator has a complex implementation (Borrill et al. 2007). The computational cost of the \mathbf{N}^{-1} operator is $O(N_t \ln(N_t))$ and the storage space required is $O(N_s N_t)$, where N_t is the number of time samples, N_τ is the length of the correlation function, and N_s is the number of stationary intervals. It is important to note that \mathbf{N}^{-1} can be arbitrary within the structural definition outlined above. This includes the possibility for each stationary interval to have a different noise correlation function associated with it, and these noise correlation functions can take any form.

MADmap has the functionality to read an arbitrary sparse preconditioner matrix and use this precomputed data product as the preconditioner. If no preconditioner is provided, then MADmap uses a diagonal preconditioner composed of the inverse of the diagonal elements of $\mathbf{A}^T \mathbf{A}$.

MADmap uses the M3 library to access CMB data. Applications that use the M3 library for reading data can be daisy chained, so that the output of one application can be used as input for a subsequent application. (A suite of massively parallel applications using the M3 library is detailed in Borrill et al. 2007). This functionality is afforded by the combination of a versatile XML-based data model description and the ability of the M3 library to read data stored a variety of file formats while translating the data to conform to a single application programmer interface.

Use of the M3 library also gives MADmap a comprehensive and versatile user interface designed specifically to enable the analysis of simulated data. For each detector and for each six day period there are three files, containing the pointing information, the CMB signal time stream, and the detector noise time stream. These 732 files were organized in the M3 XML document describing the simulation, and the M3 library is able to add the signal and the noise time streams in a weighted fashion as specified by the user.

3.3. MapCUMBA

MapCUMBA is a set of codes dealing with the analysis and map-making of polarized CMB data. The current version (2.1) has been modified from previous versions (Doré et al. 2001) to solve the GLS map-making Eq. (6) by the preconditioned conjugate gradient (PCG) method. It can now process polarized data from multiple detectors, filling the gaps in the data stream due to cosmic ray hits or data acquisition drop-out with a constrained noise realization. The noise properties, assumed piece-wise stationary, can be estimated from the TOD through an iterative process of map-making and noise extraction. The codes are written in Fortran 90 and are parallelized for distributed memory architecture using the message passing interface (MPI). They have been applied to the analysis of the 2003 flight of Boomerang.

In the current analysis, the instrumental noise is assumed to have a known power spectrum and to be statistically stationary over the whole flight. We assume its time-time correlations to be negligible beyond 5.4 min, which corresponds to $N_\tau = 65536$ samples. This latter number has been chosen to be a power of two to speed up the FFT operations described below, but could be set to any other value larger than the noise correlation length.

The convolution is applied using an overlap-add method (Press et al. 1992). This approach allows for a simple parallelization of the code, reducing the computation and memory load of each CPU. Since the number of pieces created is generally larger than the number of CPUs, and they have a predictable length (they are all equal here), load balancing between CPUs is fairly simple.

In order to apply the PCG techniques, the matrix to be inverted $\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A}$ is approximated by $\mathbf{M} = \sum_c \mathbf{A}_c^T \sigma_c^{-2} \mathbf{A}_c$, where c is a data stream of a given detector for which the instrumental noise is assumed stationary with rms σ_c . \mathbf{A}_c is the pointing matrix of the detector during that stretch of time, and the sum is done over all such pieces for all detectors. For polarization, if the condition number (ratio of the largest to the smallest eigenvalues) of a given block of \mathbf{M} is small enough, the block is inverted to provide a local preconditioner. Otherwise the pixel is degenerate and flagged as unobserved.

Making maps of the Planck HFI data with an optimal method and a PCG algorithm poses a double challenge in terms of data volume. The full data stream must be kept in memory or re-read regularly from disk to allow for the filtering operations described earlier, and the PCG algorithm requires five intermediate maps and the preconditioner \mathbf{M}^{-1} , each covering the same sky area as the final map with the same resolution and number of fields, to be kept in memory at the same time. For the full-sky HFI applications considered here, each unpolarized map with 0'8 pixels ($N_{\text{side}} = 4096$) takes 1.6 GB in double precision, while a polarized map with 1'7 pixels ($N_{\text{side}} = 2048$) takes 1.2 GB. The maps used in the PCG solver are therefore broken into N_{cpu} non-overlapping pieces of almost equal number of pixels, and each piece is assigned to one CPU. This scheme simplifies the operations such as scalar products of maps, and application of the preconditioner performed during the PCG iterations, and ensures a perfect load balancing independently of the scanning strategy.

3.4. Polar

Polar (Keihänen et al. 2006) is based on the destriping approach. We assume no a priori knowledge of the baseline amplitudes, assigning uniform prior probability to \mathbf{x} . The baseline amplitude vector \mathbf{x} is solved from the TOD using (12), without preconditioning.

In the present study we use a baseline length of 1 min, the nominal spin period of the Planck spacecraft. The number of baselines is then $N_b = N_{\text{det}} \times 366 \times 24 \times 60$, where N_{det} is the number of detectors used. We have found that 1-min baselines produce lower noise in the maps than longer baselines (Keihänen et al. 2006). Note that no coaddition of the data was performed, unlike in Keihänen et al. (2004) and Poutanen et al. (2006), where the use of idealized pointing combined with the Planck scanning strategy made it easy to compress the TOD into smaller size by coadding repeated scans over the same sky circles.

Solving for the baselines and the final map-making step involve multiplication by the inverse of the matrix $\mathbf{A}^T \mathbf{N}_u^{-1} \mathbf{A}$ (9). This matrix has been discussed in more detail in Section 2.4. As noted there, that in order to ensure its invertibility, we need to remove the degenerate pixels. The Polar code has two options for the criteria to determine those pixels. One is a determinant criterion, defined as $\det \equiv \sigma^6 \det \mathbf{M}_p$, where the same white noise level $\sigma_i = \sigma$ is assumed for each detector and sample and \mathbf{M}_p denotes the p -th diagonal block of the matrix $\mathbf{A}^T \mathbf{N}_u^{-1} \mathbf{A}$. The second criterion is the condition number requirement introduced in Section 2.4. The pixels, whose determinant or reciprocal condition number is less than some predefined value, will be rejected.

The code is written in Fortran 90, and parallelized using MPI. The full data stream is kept in memory, leading to a large memory requirement. Each processor handles a section of the TOD. In this study the sections were six days long. The total number of sections and processors required, therefore, was $N_{\text{det}} \times 61$. Because of the memory requirement, however, 512 (1536) processors were reserved for $N_{\text{det}} = 4$ (12). The map is divided into submaps (in this case 1024 pixels each). Each processor handles a number of submaps. Processors handling TOD and processors handling maps (the same processors may do both tasks) communicate by sending these submaps to each other.

3.5. MADAM

The MADAM map-making method is described in Keihänen et al. (2005). The method is based on the destriping technique, but also uses information on the noise spectrum.

The first implementation of MADAM was serial, and included no polarization. It allowed fitting several basis functions, such as Fourier components, and also allowed for coadding of the TOD. For the purpose of this study we wrote a new parallel implementation of the method, capable of handling larger data sets, and specially designed for uncoadded data. It allows fitting of uniform baselines only, which simplifies the code somewhat, but is able to handle polarization data.

As in Polar, the low-frequency part of the noise is represented by a sequence of constant baselines, and the remaining part of the noise is treated as white noise. The baseline amplitude vector \mathbf{x} is solved from Eq. (15). The covariance matrix \mathbf{P} of the baseline amplitudes (15) is computed from the noise spectrum, assumed to be known. The final output map is constructed from the cleaned TOD as in Eq. (14).

The code has some parts in common with Polar. In particular, the binning of TOD into the map and handling of degenerate pixels are done in the same way. Also, the parallelization scheme and distribution of data among processors are similar. The main difference between the codes is the term $\mathbf{P}^{-1} \mathbf{x}$ in (15), which allows MADAM to fit accurately much shorter baselines than Polar and thus model the correlated noise better. In the present study we used a baseline length of 4 s. The chosen baseline

length represents a trade-off between accuracy and computation time.

MADAM solves the linear system (15) by conjugate gradient iteration, with a preconditioner matrix

$$\mathbf{M} = (\mathbf{B}^T \mathbf{N}_u^{-1} \mathbf{B} + \mathbf{P}^{-1}). \quad (18)$$

The preconditioner is circulant and its multiplication with a baseline vector can be handled by a Fourier technique, like the term $\mathbf{P}^{-1} \mathbf{x}$.

The term $\mathbf{P}^{-1} \mathbf{x}$ is computed using the overlap-save method (Press et al. 1992). It is calculated independently for each detector. The sequence of baselines is split into sections, whose length in the present study was 6 or 3 days. Each section is padded at both ends with baselines from the two neighbouring sections, wrapping around at the ends. Each section is then convolved with \mathbf{P}^{-1} using a FFT technique and cut back to the original length. The convolution takes into account noise correlations over the length of one data section.

Formulae for computing the prior offset covariance \mathbf{P} were given in Keihänen et al. (2005). The code assumes stationary noise over the whole TOD sequence.

3.6. Springtide

Springtide (Ashdown 2007) is an implementation of destriping that is optimized to enable medium-sized supercomputers or clusters to analyse Planck-sized data sets. The map-making implementations described above require all of the data and pointing to be stored in memory at once, at large cost in memory. For example, to process the data from all 12 detectors of the HFI 217 GHz channel requires several terabytes of RAM. Springtide reduces the memory requirements by dividing the map-making into two stages.

The first stage takes advantage of the repeated observations of the same ring during one pointing period to compress the data. The baselines that will be subtracted from the time-ordered data are taken to be equal to the pointing period, during which the noise is assumed to be white. The data can therefore be compressed by making a map for each pointing period for each detector – a “ring-map” – by simple binning. The low-frequency part of the noise, modelled by an offset for each pointing period, is unaffected by the binning operation. The high-frequency noise in the resulting ring-maps is uncorrelated, with an amplitude modulated by the hit-count in each of the ring-pixels. Making ring-maps in this way compresses the TOD by a factor of 20–30, depending on the amplitude of the nutation of the satellite. The ring-maps are made using the same pixelation that will be used to make the final map to avoid any re-pixelation errors.

The second stage of the map-making is to combine the ring-maps into a global map by destriping. In each ring-map, the low-frequency noise is modelled as an offset. The maximum-likelihood solution for the offsets is found using the conjugate gradient algorithm. The maximum-likelihood offsets are subtracted from the ring-maps, which are then binned to make the global map. An iterative procedure can be used to estimate the low-frequency properties of the noise.

Springtide is written in Fortran 95 and is parallelized using MPI.

4. Simulations

4.1. Detectors

In this study, our goal was to use data from detectors in a CMB-dominated Planck channel with the highest angular resolution ($\sim 5'$). The HFI 217 GHz channel meets these requirements, so we chose to use it in this study.

For temperature-only simulations we used data from the unpolarized HFI bolometer 217-4. For polarization maps, a minimum of four polarisation-sensitive detectors is needed, so we used 217-5a, 217-5b, 217-7a, and 217-7b. A pair of detectors (a and b) share a horn in the focalplane and are oriented perpendicularly to each other. The two horns we used scan the same path on the sky. The polarization axes of 217-7a and 217-7b are aligned parallel and perpendicular to the scan direction, whereas the polarization axes of 217-5a and 217-5b are oriented at 45° to the scan direction.

4.2. Assumptions

We assumed the following:

- Noise is known a priori. (In an experiment it will have to be determined from the data themselves. Here we unrealistically avoid the effects of both the error and the statistical uncertainty that would exist in dealing with real data.)
- Noise consists only of white noise and a low frequency component, with no peaks or other features at high frequencies, and no gaps in the noise spectrum (as would result from, e.g., microphonics of slightly varying frequency).
- Noise is pseudo-stationary and Gaussian. The data are divided into 6-day pieces; noise is uncorrelated between pieces, but noise properties are identical in all pieces.
- The white noise level is $\sigma = 1.280$ mK for unpolarized and $\sigma = 1.811$ mK for polarized timelines.
- The knee frequency $f_{\text{knee}} = 0.03$ Hz, minimum frequency $f_{\text{min}} = 1.15 \times 10^{-5}$ Hz, and spectral slope $\alpha = -2$.
- Noise is uncorrelated between detectors.
- In this paper the map statistics are in thermodynamic microkelvins and the angular power spectra in these units squared.
- Beams are symmetric and Gaussian.
- The *FWHMs* of the beams are 4:7247, 4:7269, 4:7370, 4:7123, and 4:7116 for 217-4, 217-5a, 217-5b, 217-7a, and 217-7b, respectively. In practice the beam-widths will not be known to this level of precision; however, we give additional significant figures to show the level of variation in the beam-widths, and to reflect what was actually used in the simulations.
- Every bolometer has the same ideal top-hat spectral response and the shape of the beam is constant over the passband.
- There are no missing data in the time-stream
- There are no foreground signals.

For each detector, measurements are taken at a rate of 200 samples/s. The detectors scan the sky along $\sim 85^\circ$ -opening-angle circles, at 1 rpm. The spin axis is changed once per hour by $2.5'$. Two scanning strategies were simulated (Dupac & Tauber 2005):

- nominal scanning, in which the spin axis follows the ecliptic plane in the anti-Sun direction during the whole flight;
- cycloidal scanning, in which the spin axis follows a cycloidal path 7° from the anti-Sun direction with 6-month period.

For both scanning strategies, a small nutation of the spin axis was included, along with a small change in the spin rate. The nutation amplitude and spin rate were changed every hour when the spin axis was repointed, and then kept constant for an hour. Their values were selected randomly from truncated Gaussian distributions with $0.5'$ rms and $2'$ maximum for nutation and 0.12 s^{-1} rms and 0.3 s^{-1} max for spin rate deviation from the nominal 1 rpm.

4.3. The CMB template

The CMB reference sky was constructed to reproduce the large scale anisotropy pattern observed by WMAP. Specifically, we chose a threshold in the angular domain, l_{WMAP} , and used the observed harmonic coefficients for $l \leq l_{\text{WMAP}}$ from the internal linear combination (ILC) CMB map³ (Bennett et al. 2003) as described below. To reduce the residual effect of foregrounds and WMAP instrumental noise and systematics, we chose the threshold to be much larger than the actual WMAP angular resolution, i.e., $l_{\text{WMAP}} = 70$. The CMB reference consists of total intensity and *E* mode polarization, but zero power on *B*, and has been obtained as described below. Extensive use was made of HEALPix⁴ routines. The template and the codes for producing it are publicly available as part of the Planck Component Separation Working Group effort for building a background and foreground reference sky⁵.

4.3.1. $l \leq l_{\text{WMAP}}$

For total intensity, we take directly the $a_{lm}^{T,\text{WMAP}}$ from the outputs of the HEALPix anafast routine running on the WMAP ILC map. The coefficients for the *E* polarization mode, a_{lm}^E , were obtained as follows. For all m , the average amplitude of the component of a_{lm}^E which is correlated with *T* is given by C_l^{TE}/C_l^T ; the remaining component has average amplitude $\sqrt{C_l^E - (C_l^{TE}/C_l^T)C_l^{TE}}$. Here by C_l we mean the angular power spectrum taken from the best fit theoretical power spectrum to the WMAP, ACBAR, and CBI data, described below. We define the angular power spectrum C_l in the usual way

$$C_l = \frac{1}{2l+1} \sum_{m=-l}^l |a_{lm}|^2 \quad (19)$$

and include no $l(l+1)/2\pi$ normalization. The relative weights of the real and imaginary part of the a_{lm}^E component which is correlated with *TE* are set by $a_{lm}^{T,\text{WMAP}}$. Those of the remaining part are chosen randomly. The result is

$$a_{lm}^E = a_{lm}^{T,\text{WMAP}} \cdot \frac{C_l^{TE}}{C_l^T} + \left(\frac{x_{lm}^E + iy_{lm}^E}{\sqrt{2}} \right) \sqrt{C_l^E - \frac{C_l^{TE}}{C_l^T} C_l^{TE}}, \quad (20)$$

where x_{lm}^E and y_{lm}^E are Gaussian distributed numbers with zero mean and unit variance; the $\sqrt{2}$ as well as the imaginary part are absent for $m = 0$.

4.3.2. $l > l_{\text{WMAP}}$

The harmonic coefficients on these scales were obtained by running the synfast HEALPix routine to generate a template for *T*,

³ http://lambda.gsfc.nasa.gov/product/map/m_products.cfm

⁴ <http://healpix.jpl.nasa.gov/>

⁵ see <http://www.planck.fr/heading79.html> for details.

Table 1. Inputs and outputs of the four temperature simulations without noise, designed to test our ability to reproduce the input map exactly.

Case	Bolometer time constant	Pointing precision	Output resolution / N_{side}
1	No	Double	4096
2	No	Single	4096
3	Yes	Double	4096
4	No	Double	2048

Q , and U out of the WMAP, ACBAR, and CBI best-fit theoretical CMB power spectrum with no running spectral index⁶, specified by $h = 0.71992$, $\omega_b = 0.02238$, $\omega_{CDM} = 0.11061$, $\tau = 0.11027$, $n_s = 0.95820$, and $r = 0$, where h is the Hubble constant in units of $100 \text{ km s}^{-1} \text{ Mpc}^{-1}$, ω_x/h^2 is the ratio between the density of the species x and the critical one, τ is the optical depth at reionization, n_s is the scalar perturbation spectral index, and r is the ratio between the primordial tensor and scalar perturbation amplitude.

4.4. Temperature

An initial set of simulations for temperature only, summarized in Table 1, was run with no noise, in order to test our ability to reproduce exactly the input map. These simulations comprised one year of CMB temperature data from HFI 217-4, using the cycloidal scanning strategy over an input temperature map generated at HEALPix $N_{\text{side}} = 4096$ resolution. The pointing data were generated in double precision and stored in either double or single precision (see Table 1); other data were stored in single precision.

Three main issues were identified which impacted the ability of the map-making codes to reproduce the noiseless signal in the input maps:

- Pointing precision – Information about the satellite pointing is specified in terms of θ and ϕ , the co-polar and azimuthal angles in ecliptic coordinates. Initially, the Level-S pipeline calculated the pointing angles in double precision but stored them in single precision. This truncation led a small fraction of the pointings being assigned to the wrong pixel. This mis-assignment could only affect pointings that lay within the truncation error limit of a pixel boundary, but the number of such pointings increases as the pixel size is reduced.
- Bolometric time constants – Each time sample corresponds to the integration of the signal along the scan for the duration of the sampling period. The Level-S pipeline simulates this by performing a convolution of the input signal along the scan. The time stamp associated with each sample is the time when the sample is read by the onboard electronics, at the end of the integration period. In addition to this integration, bolometers have a finite time impulse response. This leads to an additional smearing which is modeled in the simulations by a convolution with a causal filter (implemented through averaging over ‘fast samples’). The smearing of the signal due to both of these effects will have to be taken into account in the analysis. Ideally the filter response will be inferred from the data themselves. We conjecture that the ability to do so will depend on the availability of bright point sources and the ratio of the bolometric time constant to the sampling period. Except for the noiseless case 3, where this feature was included in the simulation, to see its effect when not

Table 2. RMS pixel residual for the four temperature simulations without noise. For the optimal map-making codes * indicates that noise correlations were ignored and † that they were taken into account; see the text for further details.

Code	rms residual / μK			
	Case 1	Case 2	Case 3	Case 4
Polar	0.00	10.9	0.0753	1.85
MADAM	0.00	–	–	1.85
Springtide	0.00	–	–	1.85
MADmap*	2.08×10^{-18}	10.9	0.0682	1.85
MADmap†	4.50×10^{-4}	–	–	–
MapCUMBA*	2.06×10^{-5}	10.9	0.0682	1.85
ROMA*	0.00	–	–	1.85
ROMA†	0.00	–	–	–

corrected for, we decided to switch it off from the simulation pipeline for several reasons:

- In practice, these effects will have to be included in the time-stream pre-processing and noise estimation, not the map-making, and this stage was not being considered here.
- The effect of the bolometric time constants is coupled to the effects of asymmetric beams and correlated noise. Our current exercise does not consider these effects, and it does not make sense to consider the effect of bolometric time constants in isolation.
- The bolometric time constants used in Level-S were not based on actual hardware performance, which was not yet available, and therefore were not likely to be representative of the actual Planck instrument.
- Output map resolution – If the output map has lower resolution than the input map there is a residual due to the additional pixel smoothing.

Once these issues were sorted out, a fifth temperature-only simulation was run, the same as Case 1 but with noise added.

4.5. Polarization

We then ran simulations with polarization. Based on the lessons learned in the temperature cases, all used double precision pointing, no bolometer time constant, and output maps at the same resolution as the input data. Data were generated for four polarization-sensitive bolometers, 217-5a, 217-5b, 217-7a & 217-7b. Input and output I , Q , and U maps were made at $N_{\text{side}} = 2048$ to reduce memory requirements during the simulation process. Four cases were run: with and without noise for both nominal and cycloidal scanning strategies.

5. Results

The output map was subtracted from a reference map at the same resolution, and the mean, minimum, maximum, and root mean square residual pixel values were calculated, together with the angular power spectrum of the residual map. The full results are available on the Helsinki results website⁷. Tables 2 (no noise), 3 (with noise), and Figure 1 (with noise) summarize the results for the single detector temperature-only runs for the four cases listed in Table 1.

In the noiseless cases, the optimal map-making implementations can be run with two different assumptions about the noise

⁶ <http://lambda.gsfc.nasa.gov/product/map/lcdm.cfm>

⁷ <http://www.mrao.cam.ac.uk/~maja1/ctp/helsinki/>

Table 3. RMS pixel residual for the temperature simulation with noise.


Code	rms residual / μK
Polar	379.12
MADAM	378.96
Springtide	401.19
MADmap	378.96
MapCUMBA	378.96
ROMA	378.96

covariance matrix. The first is to assume that the prior information about the noise is correct, even though no noise is included in the simulations. The second is to ignore the prior information about the noise and to assume it is white. For the noiseless case the latter is closer to the truth, but ignores the correlations that will be introduced when noise is added to the simulations. Tables 2 and 4 indicate which assumption was made.

From the no-noise results we see that simulation errors exceed the numerical precision of the map-making codes and are recovered by all six methods. Residuals arising from the map-making itself are extremely small, as can be seen from the case 1 results in Table 2. The much higher rms of the codes in case 2 illustrates the point made in Sect. 4.2 about the impact of pointing precision on high resolution maps. case 3 shows that smoothing caused by the bolometer time constant adds appreciable residuals. (The marginally higher rms for Polar in case 3 is due to using a lower pixel resolution ($N_{\text{side}} = 2048$) in the ring offset fit. This was a temporary fix for a problem an early version of the code had with noiseless data. The TOD for cases 2 and 3 were later deleted to save disk space, so that they could not be redone with a consistent use of the higher resolution, after the problem was corrected.) Case 4 shows that the use of different resolutions for input and output maps contributes much larger residuals than the map-making itself. The most significant of these three sources of error is the pointing precision, followed by the output map resolution, and then the bolometer time constant. It is worth noting that the pointing precision error generates residuals with not only the largest amplitude, but also the most structured angular power spectrum.

In all cases the (pure) map-making residuals are much smaller than the residual noise level, as shown by the results given in Table 3 for the simulations including instrumental noise. As discussed in Sect. 5.1 below, the residual rms values in Table 3 are consistent with the expected white noise level. Figure 1 displays the power spectrum of the noise residuals. All the methods tend towards the white noise plateau at high l , and contain significant structure at low l . Residual noise from Springtide remains noticeably higher than in the other codes (especially at multipoles $l \lesssim 100$), because it works with ring baselines rather than the shorter time-stream baselines used by the other destripers. This greatly reduces the resources required for the code (see below), but it permits a larger amount of low-frequency noise to pass into the reconstructed map. Figure 1 also shows the spectrum of the CMB anisotropy signal. At low multipoles the residual noise is a tiny fraction of the CMB signal. At $l \gtrsim 900$ the noise exceeds the magnitude of the CMB. The differences of the residual noise of the map-making codes appear mainly in the large angular scales (at $l \lesssim 100$), where the CMB dominates. Therefore these results suggest that the differences in the temperature maps from our map-making codes are not likely to be significant for the science to be extracted from them.

Tables 4 (no noise) and 5 (with noise) and Figs. 2 (cy-cloidal scanning) and 3 (nominal scanning) show results for the

**Fig. 1.** Spectrum of residuals from the temperature simulation with noise. For comparison the spectrum of the CMB reference map is shown too (black curve). C_l is defined in Eq. (19).

polarization simulations. The beam sizes of the four detectors differ slightly; the reference map used to calculate the residuals was constructed by averaging the four input maps. Comparison of the two tables demonstrates that once again residuals caused by the map-making itself are negligible relative to the noise residuals. In Fig. 2 the low- l residuals of the ROMA code differ from the residuals of the other codes, whereas the ROMA residuals are well comparable to the other residuals in Figs. 1 and 3. For Fig. 2 we could have let ROMA to perform some more conjugate gradient iterations for better low- l recovery. We can see from Figs. 2 and 3 that the low- l ($l \lesssim 7$) EE power of the noise residuals is comparable to the power of the reionization bump of the CMB signal. At higher l the residual noise dominates. A closer inspection of Figs. 2 and 3 reveal that the low- l EE powers of the residual noise are nearly identical except for Springtide and Polar, which have slightly larger low- l power than the other map-making codes (at least in this realization of noise). For the detection of the weak polarization signal the larger residual noise is a disadvantage.

We also note that the U residuals in Table 4 are significantly higher than those on Q , an effect that we have traced to differences in the $FWHM$ of the beams (uncorrected in our map-making). In fact, the difference in beam $FWHM$ for the 217-5ab pair, aligned at 45 degrees to the scan direction, is larger than the difference for the 217-7ab pair, aligned parallel/perpendicular to the scan direction. Given these orientations and the scanning pattern on the sky, the 217-5ab pair primarily measures the U field (especially at the Ecliptic equator). The larger beam difference for this pair thus causes a larger residual in the U measurement. The effect is, however, much smaller than the noise residuals (Table 5). We discuss the import of these results on scanning strategy in the Conclusion section.

Polarization map-making introduces the additional complication of handling pixels in which the I , Q and U Stokes parameters are either mathematically or numerically degenerate. In the white noise case this degeneracy can be seen in the properties of the 3×3 (I, Q, U) pixel triplet covariance submatrices of the full block-diagonal inverse pixel-pixel noise correlation matrix $S^{-1} = A^T N^{-1} A$. In these runs we used the determinant criterion of Polar, so that Polar rejected all (I, Q, U) pixel triplets whose covariance matrix had a determinant below some threshold, while MapCUMBA and MADmap imposed a

Table 4. RMS pixel residuals for the polarization simulations without noise. For the optimal map-making codes * indicates that noise correlations were ignored and † that they were taken into account; see the text for further details.

Code	Cycloidal rms residual / μK			Nominal rms residual / μK		
	I	Q	U	I	Q	U
Polar	4.07×10^{-3}	3.42×10^{-3}	1.22×10^{-2}	2.30×10^{-3}	2.86×10^{-3}	1.27×10^{-2}
MADAM	2.09×10^{-3}	3.41×10^{-3}	1.22×10^{-2}	2.00×10^{-3}	2.86×10^{-3}	1.27×10^{-2}
Springtide	2.09×10^{-3}	3.41×10^{-3}	1.22×10^{-2}	2.00×10^{-3}	2.86×10^{-3}	1.27×10^{-2}
MADmap†	2.16×10^{-3}	3.71×10^{-3}	1.83×10^{-1}	2.02×10^{-3}	2.87×10^{-3}	1.46×10^{-2}
MapCUMBA*	2.09×10^{-3}	3.41×10^{-3}	1.22×10^{-2}	2.00×10^{-3}	2.86×10^{-3}	1.27×10^{-2}
MapCUMBA†	1.23×10^{-1}	1.70×10^{-2}	2.47×10^{-2}	5.39×10^{-3}	2.88×10^{-3}	1.27×10^{-2}
ROMA*	2.09×10^{-3}	3.41×10^{-3}	1.22×10^{-2}	2.00×10^{-3}	2.86×10^{-3}	1.27×10^{-2}
ROMA†	2.09×10^{-3}	3.41×10^{-3}	1.22×10^{-2}	2.00×10^{-3}	2.86×10^{-3}	1.27×10^{-2}

Table 5. RMS pixel residuals for the polarization simulations with noise.

Code	Cycloidal rms residual / μK			Nominal rms residual / μK		
	I	Q	U	I	Q	U
Polar	102.91	148.05	148.41	97.36	138.23	139.31
MADAM	102.86	147.97	148.33	97.32	138.16	139.25
Springtide	106.18	152.51	153.18	100.03	141.91	143.25
MADmap	102.86	147.97	148.34	97.32	138.16	139.24
MapCUMBA	102.86	147.97	148.33	97.32	138.16	139.25
ROMA	102.86	147.97	148.33	97.32	138.16	139.24

minimum condition number on this matrix for the pixel triplet to be accepted. The implications of these acceptance criteria, and the choice of threshold for each, is an ongoing research activity.

5.1. Expected white noise levels

The map residuals in Tables 3 and 5 are due mainly to remaining noise in the maps. They can be compared to the expected level of white noise, which we discuss in this subsection.

The white noise level in our simulated noise timelines was set to $\sigma = 1279.78 \mu\text{K}$ for the unpolarized detectors, and $\sigma = 1811.15 \mu\text{K}$ for the polarized detectors. A full year (366 days) TOD from one HFI detector contains $366 \times 86400 \times 200 = 6\,324\,480\,000$ samples. If the whole sky were sampled uniformly (same number of samples from each pixel), we would get $n = 31.4$ hits per pixel for an $N_{\text{side}} = 4096$ map, and $n = 125.7$ hits per pixel for an $N_{\text{side}} = 2048$ map, from one detector. This would lead to a white noise rms level of $228.34 \mu\text{K}$ ($114.17 \mu\text{K}$) per pixel for an $N_{\text{side}} = 4096$ ($N_{\text{side}} = 2048$) I map from a single unpolarized detector. Assuming also an optimal sampling of polarization directions, from the set of four polarized detectors, the white noise level would be $80.79 \mu\text{K}$ ($114.25 \mu\text{K}$) per pixel for the I (Q , or U) map with $N_{\text{side}} = 2048$. And from the one-year data from the full set of twelve HFI 217 GHz detectors (8 polarized, 4 unpolarized), the white noise level would be $40.38 \mu\text{K}$ ($80.79 \mu\text{K}$) per pixel for the I (Q or U) map with $N_{\text{side}} = 2048$.

A standard measure of performance is to give the white noise level for a pixel whose size corresponds to the resolution of the detector. From the above numbers the white noise level per a $5' \times 5'$ pixel would be $13.87 \mu\text{K}$ or $\Delta T/T = 5.1 \times 10^{-6}$ for the I map, and $27.75 \mu\text{K}$ or $\Delta T/T = 10.2 \times 10^{-6}$ for the Q and U maps. For fourteen months of data these numbers are smaller by a factor of $\sqrt{12/14}$ giving $\delta T/T = 4.7 \times 10^{-6}$ for I and 9.4×10^{-6} for Q and U , in accordance with the Planck instrument performance goals.

The actual average white noise level will be higher, since the sky will not be sampled uniformly, and also the polarization directions will not be sampled optimally for every pixel. On

the other hand, there will be regions of the sky where the noise will be much lower. We calculated the resulting theoretical white noise rms values for our maps using the actual pointing of the TOD used in this study.

An $N_{\text{side}} = 2048$ pixel is typically just 1.7 across. The Planck spin axis is shifted by about 2.5 at every repointing. The simulations presented here involve either one or four detectors, rather than the full complement of twelve that Planck will have. As a result, the sampling of the sky in these simulations is quite nonuniform at small scale. This leaves “cracks” between the scanning rings of a single detector or a 4-detector, 2-horn polarized pair at low ecliptic latitudes in an $N_{\text{side}} = 2048$ map. The expected nutation of the satellite spin axis, included in our simulated data, reduced this effect somewhat by broadening the scanning rings, but many pixels were not hit at all, and many were hit only a few times. Pixels hit only a few times, moreover, were not necessarily hit by all the detectors, and the sampling of polarization directions may be poor.

For the one-detector case, 189 146 022 of the 201 326 592 pixels of the $N_{\text{side}} = 4096$ map were hit, with an average of $n = 33.4$ hits per hit pixel. If the hits were distributed uniformly over these pixels, the white noise level would be $\sigma_{\text{map}} = 1279.78 / \sqrt{n} = 221.32 \mu\text{K}$. Actually the hits are distributed quite nonuniformly, with $1/\langle 1/n_p \rangle = 11.4 \ll 33.4 = \langle n_p \rangle$, giving an expected rms white noise of $\sigma_{\text{map}} = 1279.78 \times \sqrt{\langle 1/n_p \rangle} = 378.95 \mu\text{K}$ on the map. Here n_p indicates the number of hits in a pixel, and $\langle \cdot \rangle$ an average taken over all hit pixels. We see that the map residuals in Table 3 are quite close to this expected level of white noise.

The corresponding results for the four-detector cases with polarization are presented in Table 6. Some of the pixels of the $N_{\text{side}} = 2048$ map are not hit at all. In the nominal scanning strategy there are also large holes at the ecliptic poles, where the scanning does not extend. Of the pixels that are hit, we have defined as *degenerate* those map pixels whose $\kappa_p^{-1} < 10^{-6}$, and *bad* those with $10^{-6} \leq \kappa_p^{-1} < 10^{-2}$, with the idea that degenerate pixels should be excluded from the map-making completely, whereas bad pixels can be calculated by the map-making codes,




Fig. 2. Spectra of residuals for the polarization simulations with noise in the cycloidal scanning strategy. The black curve shows the power spectrum of the CMB reference map. It has zero B mode power. C_l is defined in Eq. (19).

but are excluded from the calculations of the map rms residuals or the expected white noise levels (see Sect. 2.4). Inclusion of bad pixels would increase dramatically the Q and U residuals, because the low reciprocal condition number means that noise is amplified in the solution of these Stokes parameters.

We can see from Table 6 that the main reason the rms residuals in Table 5 are significantly larger than in the ideal case of uniform sky coverage and optimal polarization sampling is the nonuniform sky coverage. Nonoptimal sampling of the polarization directions gives an additional contribution. This latter contribution is sensitive to the chosen cut-off (here 10^{-2}) for “bad” pixels. Using $\kappa_p^{-1} < 10^{-1}$ as a criterion, we would halve the

difference between optimal and actual polarization sampling, at the cost of leaving out an additional 17196 (5540) pixels in the cycloidal (nominal) case.

Table 6 cannot be directly compared to Table 5, because the set of pixels included is not exactly the same. Different codes used slightly different criteria for excluding pixels, and Table 5 is calculated for the common set of pixels calculated by all codes, which is slightly smaller than the set of $\kappa_p^{-1} \geq 10^{-2}$ pixels. We have therefore added in Tables 6 and 7 a column giving the Polar rms residual for exactly the same set of pixels as the other entries in Tables 6 and 7. We can see that the expected white noise level for the actual scanning predicts well this Polar result, indicating




Fig. 3. Spectra of residuals for the polarization simulations with noise in the nominal scanning strategy. The black curve shows the power spectrum of the CMB reference map. It has zero B mode power. C_l is defined in Eq. (19).

that the rms residuals for the map-making codes come mainly from the white noise.

For the cycloidal scanning strategy we actually had available simulated TOD for the full set of 12 detectors. It can be seen from Table 7 how including 8 or 12 detectors leads to more uniform sky coverage. The small-scale nonuniformity due to the cracks between scanning rings goes away. The large-scale nonuniformity due to more visits per pixel near the ecliptic poles

remains. The degenerate and bad pixels, that were due to the cracks have practically disappeared. Thus the results are much closer to the ideal case.

5.2. Resource requirements

The codes described here operate on the full TOD, and their resource requirements are sizable. The one-year TOD

Table 6. The expected level of white noise in the maps for the 4-detector polarization simulations.

Pixels			Uniform hits	Actual hits	Actual hits	Polar with
Not hit	Discarded		optimal sampling	optimal sampling	actual sampling	$\kappa_p^{-1} \geq 0.01$
			$\sigma_{\text{map}}/\mu\text{K}$	$\sigma_{\text{map}}/\mu\text{K}$	$\sigma_{\text{map}}/\mu\text{K}$	$\sigma_{\text{map}}/\mu\text{K}$
Cycloidal						
40383	56140	<i>I</i>	80.71	102.82	102.97	103.06
Holes 0	Degenerate 54380	<i>Q</i>	114.14	145.41	149.96	150.02
Cracks 40383	Bad 1760	<i>U</i>	114.14	145.41	149.83	149.98
Nominal						
145626	10 346	<i>I</i>	80.66	97.26	97.29	
Holes 144021	Degenerate 7559	<i>Q</i>	114.07	137.55	138.31	
Cracks 1605	Bad 2787	<i>U</i>	114.07	137.55	139.43	

Table 7. The expected level of white noise in the maps for 8-detector and 12-detector simulations. One year of simulated TODs were generated for all 12 detectors. The scanning paths of the detectors were generated using a realistic focalplane. The beam-widths of 217-4, 217-5a, 217-5b, 217-7a, and 217-7b were as given in Section 4.2. The beam-widths of the rest of the detectors were $FWHM = 4'.7263, 4'.7055, 4'.7011, 4'.7192, 4'.7148, 4'.7438, \text{ and } 4'.7442$ for 217-1, 217-2, 217-3, 217-6a, 217-6b, 217-8a, and 217-8b, respectively.

Pixels			Uniform hits	Actual hits	Actual hits	Polar with
Not hit	Discarded		optimal sampling	optimal sampling	actual sampling	$\kappa_p^{-1} \geq 0.01$
			$\sigma_{\text{map}}/\mu\text{K}$	$\sigma_{\text{map}}/\mu\text{K}$	$\sigma_{\text{map}}/\mu\text{K}$	$\sigma_{\text{map}}/\mu\text{K}$
Cycloidal, 8 polarized detectors						
0	4	<i>I</i>	57.12	65.72	65.72	65.79
Holes 0	Degenerate 3	<i>Q</i>	80.79	92.94	93.16	93.23
Cracks 0	Bad 1	<i>U</i>	80.79	97.94	93.04	93.11
Cycloidal, all 12 detectors						
0	16	<i>I</i>	40.38	45.88	45.88	
Holes 0	Degenerate 3	<i>Q</i>	80.79	92.94	93.16	
Cracks 0	Bad 13	<i>U</i>	80.79	92.94	93.04	

from four HFI detectors contained $4 \times 366 \times 24 \times 3600 \times 200 = 25\,297\,920\,000$ samples. All codes except Springtide keep the whole dataset in memory. Tables 8 and 9 show the memory and CPU times required for both the one-detector (unpolarized) and four-detector (polarized) cases. The codes were run on the NERSC *Seaborg* supercomputer, an IBM SP RS/6000, with clock speed 375 MHz and peak performance 1.5 Gflops/processor. Note that for the same code these run times can vary depending on load and status of the supercomputer. Since not all timings were obtained in exactly the same manner, the comparison should be taken as indicative only.

All algorithms solve a linear system by iterative methods. In destriping (Polar, Springtide) and destriping-based (MADAM) methods, the quantities to be solved are a set of baselines for the TOD. As a result, the system is smaller than in optimal map-making (MADmap, MapCUMBA, ROMA). Thus the destriping codes are significantly faster than the optimal codes.

The processors in *Seaborg* are organized into nodes of 16 processors. Processors of the same node share memory. Most nodes have 16 GB of memory, but some have 32 GB or 64 GB. Because of the size of the TOD, the codes require lots of memory. Since the communication overhead between processors increases when the same task is performed with a larger number of processors, run time increases as memory per processor decreases. As a result, the 32 GB and 64 GB nodes are faster. However, queue times on the high-memory nodes are longer, so most of the codes were run on the 16 GB nodes. The ROMA code requires the 32 GB nodes, because at one point in the calculation it keeps a full sky map in the memory requested by a single processor. For the other codes, the memory requirement is dominated by the size of the TOD, and the variation between codes reflects differences in how the TOD are stored in memory

(single or double precision; how the pointing and orientation of the detector is represented). For the unpolarized case, MapCUMBA was run with an older version of the code, which required more than twice the memory amount of the current code and used a slower FFT library.

Springtide requires less memory than the other codes, because it keeps only ring-maps in memory simultaneously. A ring-map for a Planck one-hour pointing period contains many fewer pixels than the number of samples in the corresponding one-hour TOD. Springtide could achieve lower noise in the maps by using baselines shorter than one hour. This would reduce the compression factor from TOD to ring-maps; for 1-minute baselines, as used by Polar, there would be no compression. But this would also increase the size of the matrix equation to solve, increasing computation time.

6. Conclusions

The main goal of the simulations reported in this paper was to compare various map-making codes, and to demonstrate that they can deal with Planck-size data sets. The complexity of the simulated data was kept to a minimum in order to isolate software-induced systematic effects. No instrumental systematic effects other than noise correlations were included in the simulated data. Comparisons based on more realistic simulations will be made in future papers, which will assess the impact of strong gradients in the signal due to foregrounds, and include a more realistic treatment of the instrumental transfer function, e.g. through the inclusion of beam asymmetries. Nevertheless, some useful results can be identified.

Table 8. Memory use and run times for the different codes in the unpolarized case with one detector. See text for discussion.

Code	Number of processors	Memory required (reserved) / GB	Run time / min	Total CPU-hours
Polar	256	60 (256)	23	98
Springtide	128	45 (128)	21	45
MADAM	256	60 (256)	18	77
MADmap	512	200 (512)	90	768
MapCUMBA	768	300 (768)	52	666
ROMA	512	100 (1024)	129	1101

Table 9. Memory use and run times for the different codes in the polarized case with four detectors. See text for discussion.

Code	Number of processors	Memory required (reserved) / GB	Run time / min	Total CPU-hours
Polar	512	400 (512)	37	316
Springtide	512	70 (512)	36	307
MADAM	512	400 (512)	45	384
MADmap	2048	800 (2048)	90	3078
MapCUMBA	1024	800 (1024)	100	1707
ROMA	1024	800 (2048)	206	3515

6.1. Residual noise

The residual noise rms of the optimal and MADAM maps were nearly the same. The map noise rms of Polar and Springtide were larger, but the difference is small.

In our temperature maps the residual noise is smaller than the CMB anisotropy signal at large and intermediate angular scales. The differences in the residual noise of our map-making codes appear mainly at large angular scales where the temperature maps are CMB dominated. Therefore we expect, that these differences would not affect significantly the science extracted from these maps. On the other hand, the large scale noise power of our polarization maps is comparable to the power of the reionization bump of the CMB signal. In smaller angular scales the residual noise dominates over the CMB anisotropy signal in our polarization maps. Springtide and Polar seem to have slightly stronger large scale residuals (at least in this realization of noise) than our other map-making codes, which is a disadvantage for them when detecting the weak polarization signal.

6.2. Resource requirements

The map-making codes described in this paper require considerable resources for Planck-sized data. All codes except Springtide keep the entire data stream in memory. Memory requirements are therefore dominated by the size of the TOD. Springtide requires significantly less memory, since it first calculates ring-maps, which compresses the data by a factor of 20–30 (for a Planck-type scanning strategy).

The optimal codes (MADmap, MapCUMBA, ROMA) achieve slightly lower noise in the final maps than Polar and Springtide, but require an order-of-magnitude more CPU time. Polar achieved lower noise than Springtide, because it worked with shorter baselines (1 minute instead of 1 h). Using shorter baselines would increase the memory requirement of Springtide. MADAM, which combines optimal map-making ideas with destriping, achieves practically the same noise levels as the optimal codes, with similar memory requirements, but in a time comparable to the destriping codes.

6.3. Scanning strategy

The Planck design allows considerable flexibility in the choice of the scanning strategy, even in orbit. The refinement of the scanning strategy will therefore be an important aspect of the pre-launch simulation and analysis work. Our ability to assess the pros and cons of various candidate strategies will improve as our map-making algorithms include the functionality to deal with increasing levels of realism in the simulated data.

Keeping these qualifications in mind, we compare the map-making residuals for the nominal and cycloidal scanning strategies. The results are summarized in Table 5. There is a slight preference for the nominal strategy. This preference is a result of the smoother distribution of integration time on the sky for the case of the nominal scanning. Comparison of Figs. 2 and 3, on the other hand, shows that the power in the residuals is higher at the lowest multipoles for the nominal scanning, at least for this particular realization of the noise TOD. An accurate quantification of the effects of scanning strategy requires, however, ensemble averaging, e.g., using Monte Carlo simulations. This will be the subject of future work. The issue is important for proper understanding of the mission's ability to measure, for instance, the reionization bump at low multipole.

Furthermore, we anticipate that the nominal scanning will perform more poorly than the cycloidal scanning when we consider some other features of the scanning strategy (Dupac & Tauber 2005). Such features are, for example, the sky coverage (nominal scanning strategy leaves unobserved caps around the ecliptic poles, whereas the cycloidal scanning can provide full-sky coverage), the ability to revisit pixels on a range of timescales in order to be able to reject systematic effects (such as residuals from quasi-periodic signals such as cooler noise) and the ability to cross through pixels in several different directions in order to allow the accurate reconstruction of the polarization maps and the beam transfer functions.

We will return to the assessment of the relative benefits of scanning strategies in future publications.

6.4. Future improvements

The simulated data prepared for this work are the most advanced so far, but are still far from the reality of the Planck experiment. This holds for instrumental systematics as well as for accurate modeling of the sky signal. No foregrounds have been included, and several aspects of the CMB emission have been simplified. In particular, no gravitational lensing or tensor signals have been included. Both processes have their main impact on the B modes of polarization anisotropy. Lensing is a well-understood and inevitable effect in cosmology, distorting the CMB and producing B modes in a broad peak in the power spectrum centered at a $\ell \approx 1000$, with an amplitude much smaller than that of E modes. Tensor signals show up primarily as degree-scale B modes. Early reionization could make a tail of that component appear on very large angular scales, with an amplitude which might be detectable by Planck. The pattern of the total intensity anisotropies on scales of about three degrees or more has been taken directly by the WMAP data. This affects also a component of the E polarization mode pattern. The angular extension and reliability of the WMAP pattern in our simulation may certainly benefit from the future releases of WMAP data.

Acknowledgements. The authors would like to thank the University of Helsinki for its hospitality in June 2004 when the CTP Working Group met to undertake this work. MAJA is a member of the Cambridge Planck Analysis Centre, supported by PPARC grant PPA/G/R/1997/00837. Part of this work was supported by NASA. CB was partly supported by the NASA LTSA Grant NNG04CG90G. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098. We acknowledge the use of version 0.1 of the Planck reference sky model, prepared by the members of Planck Working Group 2 and available at <http://www.planck.fr/heading79.html/>. EK and TP were supported by the Academy of Finland grants no. 75065, 205800, 213984, and 214598. TP wishes to thank the Väisälä Foundation for financial support. Some of the results in this paper have been derived using the HEALPix (Górski et al. 2005) package.

References

- Ashdown, M. A. J. 2007, in preparation
 Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003, *ApJ*, 148, 1
 Benoît, A., Ade, P., Amblard, A., et al. 2003, *A&A*, 399, L19
 Bond, J. R., Crittenden, R. G., Jaffe, A. H. & Knox, L. 1999, *Comput. Sci. Eng.*, 1, 21
 Borrill, J. D., Cantalupo, C. M., Jaffe, A. H., & Stompor, R. 2007, in preparation
 Burigana, C., Malaspina, M., Mandolesi, N., et al. 1997, *Int. rep. TeSRE/CNR*, 198/1997, November, [arXiv:astro-ph/9906360]
 de Bernardis, P., Ade, P. A. R., Bock, J. J., et al., *Nature*, 404, 955
 de Gasperis, G., Balbi, A., Cabella, P., Natoli, P., & Vittorio, N. 2005, *A&A*, 436, 1159
 Delabrouille, J. 1998, *A&AS*, 127, 555
 Doré, O., Teyssier, R., Bouchet, F. R., Vibert, D. & Prunet, S. 2001, *A&A*, 374, 358; see also <http://ulyssse.iap.fr/cmbsoft/mapcumba/>
 Dupac, X., & Tauber, J. 2005, *A&A*, 430, 363
 Frigo, M., & Johnson, S. G. 1998 ICASSP Conference, 3, 1381, see also <http://www.fftw.org/>
 Golub, G. H., & van Loan, C. F., *Matrix Computations*, 3rd edition, 1996 (John Hopkins University Press)
 Górski, K. M., Hivon, E., Banday, A. J. B. D., et al. 2005, *ApJ*, 622, 759,
 Hanany, S., Ade, P., Balbi, A., et al. 2000 *ApJ*, 545, L5
 Keihänen, E., Kurki-Suonio, H., Poutanen, T., Maino, D., & Burigana, C. 2004, *A&A*, 428, 287
 Keihänen, E., Kurki-Suonio, H., & Poutanen, T. 2005a, *MNRAS*, 360, 390
 Keihänen, E., Kurki-Suonio, H., Poutanen, T., et al. 2006, in preparation
 Maino, D., Burigana, C., Maltoni, M., et al. 1999, *A&A*, 140, 383
 Maino, D., Burigana, C., Górski, K. M., Mandolesi, N., & Bersanelli, M. 2002, *A&A*, 387, 356
 Natoli, P., de Gasperis, G., Gheller, C., & Vittorio, N. 2001, *A&A*, 372, 346
 Poutanen, T., de Gasperis, G., Hivon, E., et al. 2006, *A&A*, 449, 1311
 Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. 1992, *Numerical Recipes in FORTRAN, The Art of Scientific Computing*, 2nd Edition (Cambridge: Cambridge University Press)
 Revenu, B., Kim, A., Ansari, R., et al. 2000, *A&AS*, 142, 499
 Smoot, G. F., Bennett, C. L., Kogut, A., et al. 1992, *ApJ*, 396, L1
 Stompor, R., Balbi, A., Borrill, J., et al. 2002, *Phys. Rev. D*, 65, 2003
 Wu, J. H. P., Balbi, A., Borrill, J., et al. 2001, *ApJS*, 132, 1