

# ROMA: A map-making algorithm for polarised CMB data sets

G. de Gasperis<sup>1</sup>, A. Balbi<sup>1,2</sup>, P. Cabella<sup>1</sup>, P. Natoli<sup>1,2</sup>, and N. Vittorio<sup>1,2</sup>

<sup>1</sup> Dipartimento di Fisica, Università di Roma “Tor Vergata”, via della Ricerca Scientifica 1, 00133, Roma, Italy  
e-mail: giancarlo.degasperis@roma2.infn.it

<sup>2</sup> Sezione INFN Roma 2, via della Ricerca Scientifica 1, 00133, Roma, Italy

Received 9 December 2004 / Accepted 3 February 2005

**Abstract.** We present ROMA, a parallel code to produce joint optimal temperature and polarisation maps out of multidetector CMB observations. ROMA is a fast, accurate, and robust implementation of the iterative generalised least-squares approach to map-making. We benchmark ROMA on realistic simulated data from the last polarisation-sensitive flight of BOOMERanG.

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

## 1. Introduction

A direct consequence of the presence of anisotropies in the Cosmic Microwave Background (CMB) radiation is a certain degree of polarisation (Rees 1968; Kaiser 1983). Constraining CMB polarisation provides valuable cosmological information complementary to that encoded in temperature anisotropy, and it significantly tightens the constraints on cosmological parameters (Seljak & Zaldarriaga 1996).

Unfortunately, the polarised component of the CMB is expected to be small compared to total intensity, making its measurement an experimental challenge. For this reason and until very recently, the experimental effort has not focused on the polarised component. However, the situation is quickly changing. A first detection of CMB polarisation, in agreement with theoretical predictions, has been announced by the interferometric experiment DASI (Kovac et al. 2002). The WMAP satellite (Bennett et al. 2003) has measured the predicted correlation spectrum between CMB temperature anisotropy and polarisation (Kogut et al. 2003). Many other CMB polarisation experiments (DASI, CBI, CAPMAP, BICEP, QUEST, BOOMERanG 2K, MAXIPOL, SPORT, BarSPORT, PLANCK, etc.<sup>1</sup>) have either already taken data or are expected to do so in the very near future.

It is well known that in CMB science the data reduction process requires almost as much care as data gathering. Analysis of temperature data has now reached full maturity and has proved successful on several datasets. The same is not true for polarisation. Although the overall analysis scheme can be borrowed from the temperature case, and many of the algorithms involved admit a more or less straightforward

generalisation, polarised data sets present peculiar aspects that call for specific treatments. One basic problem is that the tensor-like nature of polarisation poses some constraints on the measurement process (see Cabella & Kamionkowski 2004; or Lin & Wandelt 2004, for recent reviews). As a consequence, not all possible instrumental configurations are equally advantageous from the point of view of data reduction (Couchot et al. 1999). The so called “optimal” configurations comprise several different polarisation sensitive detectors (polarimeters) placed at a convenient reciprocal orientation. Usually data streams from different polarimeters are jointly analysed to seek a faint signal. The difficulties involved in this process and their relation to the design of robust and efficient data mining algorithms are seldom considered in the literature (see Revenu et al. 2000, for a notable exception).

Two major steps in the data analysis process are (1) production of sky maps from a set of Time Ordered Data (TOD) and (2) extraction of the angular power spectrum from such maps. Here we focus on the first step and will leave the discussion of the power spectrum estimation to a forthcoming paper. In a previous paper (Natoli et al. 2001) we described implementation of an optimal map-making algorithm for the case of temperature-only TOD collected by a one horned CMB experiment, and showed the feasibility of the method by applying it to PLANCK and BOOMERanG simulated data. Here we discuss the generalisation of this algorithm to the case of TOD produced by an arbitrary number of polarisation sensitive detectors. The software implementation of this method, which we call ROMA (Roma Optimal Map-making Algorithm), allows fast and efficient production of optimal multidetector maps of CMB total intensity and polarisation. The code is currently being used to analyse the polarised data set from the second Antarctic flight of the BOOMERanG experiment

<sup>1</sup> See e.g. the site [lambda.gsfc.nasa.gov](http://lambda.gsfc.nasa.gov) for up to date references to ongoing and forthcoming Polarisation experiments.

(hereafter B2K; Montroy et al. 2003), which took place in January 2003. ROMA is part of a full end-to-end data analysis pipeline for B2K, completely developed in Rome, and its details will be described elsewhere.

The plan of this paper is as follows. In Sect. 2 we derive the least-squares equations for multidetector map-making; then in Sect. 3 we describe the details of the implementation, while in Sect. 4 we report application of ROMA to highly realistic B2K simulated data. Finally, in Sect. 5, we draw our main conclusions.

## 2. Polarised data map-making

In this section we derive the Generalised Least Squares (GLS) equations for the polarised map-making problem, given an arbitrary number of polarimeters. With “polarimeter” we mean, here and in the following, a generic detector measuring total intensity plus a linear polarisation component. Other types of detectors do exist and rely on different strategies to measure polarisation. We prefer to focus on the linear polarimeter case because of its widespread adoption. In any case it should be straightforward to generalise this scheme once the data model is properly modified.

The sky signal seen by a polarimeter can be expressed as (Chandrasekhar 1960)

$$\mathcal{D} = \frac{1}{2} (I + Q \cos 2\xi + U \sin 2\xi), \quad (1)$$

where  $I$ ,  $Q$ , and  $U$  are the Stokes parameters for total intensity and linear polarisation, and where the angle  $\xi$  identifies the polarimeter orientation with respect to the chosen celestial frame. Note that we do not include the contribution of circular polarisation, associated to the Stokes parameter  $V$ . As a consequence, a  $V$  signal would be seen by our polarimeter as a contribution to  $I$ . This fact, however, is not a problem in CMB science, because circular polarisation is not generated by Thomson scattering of unpolarised radiation.

All three relevant Stokes parameters can be extracted by either combining the output of many detectors with different mutual orientation or by allowing enough focal plane rotation. The data stream output of a given detector is a combination of sky signal and instrumental correlated noise

$$\mathcal{D}_t = \frac{1}{2} A_{tp} (I_p + Q_p \cos 2\phi_t + U_p \sin 2\phi_t) + n_t. \quad (2)$$

Here index  $t$  labels time, while  $p$  runs over the pixelized images of  $I$ ,  $Q$ , and  $U$ . Pointing matrix  $A_{tp}$  couples the time and pixel domains. Information about beam smearing can, in general, be included in this matrix. We prefer, however, to assume that  $A$  is pointwise (i.e. a single nonzero entry per row that occurs when a pixel falls into the line of sight) and to hide the beam in the  $I$ ,  $Q$ , and  $U$  maps (i.e. we solve for a pixelized beam-smear image of the sky). This is only meaningful if the beam is, to good approximation, symmetric with respect to boresight and common to  $I$ ,  $Q$ , and  $U$ ; see Armitage & Wandelt (2004), for an algorithm that takes the effect of an asymmetric beam into account. In Eq. (2)  $n_t$  represents a vector of correlated noise.

By inserting the trigonometric functions within the pointing matrix and considering a set of  $k$  polarimeters, one can recast Eq. (2) into a more compact formalism:

$$\mathcal{D}_t = A_{tp} \mathcal{S}_p + \mathbf{n}, \quad (3)$$

where the datastreams of each polarimeter are combined as

$$\mathcal{D}_t \equiv \begin{pmatrix} \mathcal{D}_t^1 \\ \vdots \\ \mathcal{D}_t^k \end{pmatrix},$$

and the generalised pointing matrix becomes:

$$A_{tp} \equiv \frac{1}{2} \begin{pmatrix} A_{tp}^1 \cos 2\phi_t A_{tp}^1 & \sin 2\phi_t A_{tp}^1 \\ \vdots & \vdots \\ A_{tp}^k \cos 2\phi_t A_{tp}^k & \sin 2\phi_t A_{tp}^k \end{pmatrix}.$$

Similarly, the sky signal can be expressed as

$$\mathcal{S}_p \equiv \begin{pmatrix} I_p \\ Q_p \\ U_p \end{pmatrix},$$

while the noise stream is:

$$\mathbf{n}_t \equiv \begin{pmatrix} n_t^1 \\ \vdots \\ n_t^k \end{pmatrix}.$$

Equation (3) defines a generic linear algebra problem whose unknown parameters (the map pixel values) can be constrained by means of the standard GLS solution (e.g. Lupton 1993)

$$\widetilde{\mathcal{S}}_p = (A^t N^{-1} A)^{-1} A^t N^{-1} \mathcal{D}, \quad (4)$$

with

$$N \equiv \langle \mathbf{n}_t \mathbf{n}_t^t \rangle = \begin{pmatrix} \langle n_t^1 n_t^1 \rangle & \cdots & \langle n_t^1 n_t^k \rangle \\ \vdots & \ddots & \vdots \\ \langle n_t^k n_t^1 \rangle & \cdots & \langle n_t^k n_t^k \rangle \end{pmatrix}, \quad (5)$$

where  $\langle \cdot \rangle$  denotes the expectation value. The  $N$  matrix becomes block diagonal when assuming that the noise in different polarimeters is uncorrelated:

$$\langle n_t^j n_t^i \rangle = \langle n_t^j n_t^j \rangle = 0 \quad (i \neq j). \quad (6)$$

We defer the actual implementation of Eq. (4) to the next section.

The treatment above is clearly idealistic; however, it is possible for the formalism to incorporate correction for a nominal level of cross polarisation, one of the most common systematic effects occurring in CMB polarimetry. Cross-polarisation occurs because of cross-talk between the two orthogonal polarisation modes; that is, a polarimeter may be sensitive, with efficiency  $\chi_{\text{pol}}$ , to radiation linearly polarised  $90^\circ$  off its natural polarisation mode. If we assume that the cross-polarisation contamination is measurable (by on-ground and/or in-flight tests) and is constant across the mission lifetime, the formalism expressed above can be generalised to take this effect into

account. If we introduce a calibration constant  $C$  and a cross-polarisation factor  $\chi_{\text{pol}}$ , the data model Eq. (2) can be generalised as

$$\mathcal{D}_t^i = \frac{C}{2} A_{ip}^i \left[ (I_p + Q_p \cos 2\phi_t^i + U_p \sin 2\phi_t^i) + \chi_{\text{pol}} (I_p - Q_p \cos 2\phi_t^i - U_p \sin 2\phi_t^i) \right] + n_t^i$$

that is

$$\mathcal{D}_t^i = \frac{C_\chi}{2} A_{ip}^i \left[ I_p + \frac{1 - \chi_{\text{pol}}}{1 + \chi_{\text{pol}}} (Q_p \cos 2\phi_t^i + U_p \sin 2\phi_t^i) \right] + n_t^i, \quad (7)$$

where we have embedded the cross-polarisation factor into the new calibration constant  $C_\chi = C(1 + \chi_{\text{pol}})$ . The new data model expressed in Eq. (7) is then solved by Eq. (4), provided the generalised pointing matrix  $\mathbf{A}$  is rewritten accordingly. Note also that if  $\chi_{\text{pol}}$  is close to unity, the map-making problem expressed by Eq. (4) becomes ill-conditioned since the matrix  $\mathbf{A}'\mathbf{N}^{-1}\mathbf{A}$  is singular if  $\chi_{\text{pol}} = 1$ . For real-life experiments we expect a cross-polarisation level of <10% (e.g. Montroy et al. 2003, for B2K), not enough to hinder the search for a solution to Eq. (4). It is important to realise that the scheme outlined above can correct for a nominal and known cross polarisation level. Any uncertainty in this value, as well as in the overall calibration, is a potential source of systematics.

### 3. Implementation

Straightforward implementation of Eq. (4) for a modern CMB experiment would require storing and inverting a huge matrix (the map noise correlation matrix), a task often beyond the reach of today's supercomputers. In contrast, iterative methods only require matrix to vector products and appear well suited to tackling the problem. One such scheme proposed by Natoli et al. (2001) has been shown to be particularly convenient in the case of temperature map-making even for a high resolution full sky survey such as PLANCK. The idea (Wright 1996) is to decompose the product  $\mathbf{A}'\mathbf{N}^{-1}\mathbf{A}$  ("unroll, convolve and rebin") to avoid computing and storing the map correlation matrix and to make use of a preconditioned conjugate gradient (PCG) solver. The key assumptions are: (1) assume that the beam is axisymmetric, so as to keep the structure of the pointing matrix  $\mathbf{A}$  as simple as possible; and (2) assume that the noise is (piecewise, at least) stationary and that its correlation function decays after a time lag much shorter than the duration of the timeline piece being processed. Under the last assumption, the  $\mathbf{N}$  matrix is approximately circulant, and as such it is diagonalized by a Fourier transform. This approach achieves linear scaling (with the number of time samples) per PCG iteration; if the system is preconditioned by the pixel hits counter (the number of times each pixel is seen), an accurate solution can be obtained in a few tens of iterations (see Natoli et al. 2001, for a discussion of algorithmic details; other implementations of this approach include Doré et al. 2001; Cantalupo et al.; and Jones et al. 2005).

#### 3.1. The IQU approach

Our approach is to extend the above mentioned method to polarisation. Since the timestream includes contributions from

both total intensity and linear polarisation (see Eq. (2)) it is desirable to solve jointly for the temperature and polarisation maps in the spirit of Eq. (4). This approach, which we call IQU, is computationally more intensive than solving for  $I$  and  $(Q, U)$  separately, but preferable for at least two reasons: first, it minimises the  $Q$  and  $U$  contaminations on the  $I$  map; second it is more general, because it does not require any special constraint on the noise properties of the polarisation sensitive detectors, as would instead be the case when the TOD are explicitly differenced to remove the intensity component. It is easy to realise that in the last case the matrix  $\mathbf{N}$  would not be circulant except in the special case of having identical noise properties in different channels.

While the intensity (temperature) field can theoretically be reconstructed from a single one-horned detector, up to the same order of magnitude as the polarisation contribution to the timeline, the same thing is not true for polarisation, unless special care is taken to ensure that the detector observes all pixels of the sky under very different orientations. The latter strategy has indeed been chosen by a few experiments (e.g. MAXIPOL that uses a rotating grid, Johnson et al. 2003), but the majority (e.g. B2K, PLANCK) plan to extract the polarisation maps by comparing channels with different mutual orientation of the plane sensitive to linear polarisation. If such a strategy is used, it is necessary to reduce several channels simultaneously in order to obtain a polarisation map; i.e. one is forced to perform multi-detector map-making. Of course, multi-detector map-making can also be performed in the case of polarisation insensitive detectors, when an optimal temperature map has to be produced out of different channels at a given frequency, not necessarily taken by the same experimental system. The IQU approach is thus very general, as it also allows us to merge naturally data taken by different channels, yet for temperature and polarisation at the same time.

Due to optical constraints, different detectors do not observe exactly the same region of the sky during the mission lifetime, with the notable exception of full sky surveys. A polarisation map can only be extracted for pixels observed with sufficient variety in focal plane orientation, in order not to incur an ill-conditioned  $\mathbf{A}'\mathbf{N}^{-1}\mathbf{A}$  matrix. Hence, it is a very natural choice to restrict the polarisation map to intersection of the sky regions surveyed by different detectors. We impose this constraint by flagging the time samples associated to pixels outside of the common region. A more refined discrimination would entail explicitly selecting those pixels observed with enough spread in the polarisation angle. For the sake of simplicity we choose to restrict the  $(Q, U)$  maps to the intersection region of detector coverage. The same constraint is neither necessary nor advantageous for the  $I$  map. It is not necessary, because the temperature is normally well defined outside of the commonly observed region, although in general with a larger noise contribution; it is not advantageous, because of the way the solver works. In fact, scans would often go continuously in and out of the common region; when unrolling a tentative solution map  $\mathbf{m}$  over the timestream (i.e. when performing the operation  $\mathbf{A}\mathbf{m}$ ), the latter would display time gaps if the solution is restricted to the common region. One way to tackle this problem is to fill the gaps with a constrained noise

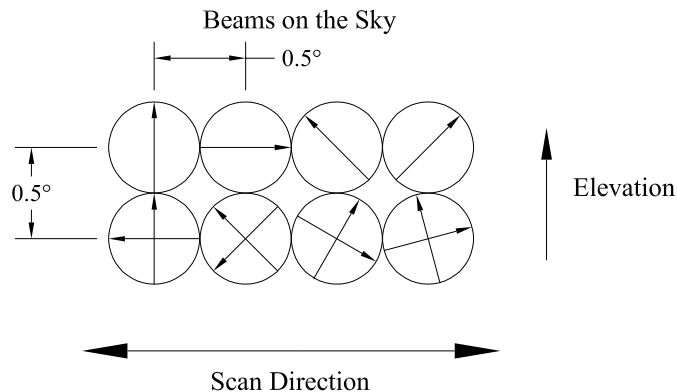
realisation. This approach has the disadvantage of increasing the computational cost of the algorithm, because standard methods to produce constrained noise realisations scale as the third power of the noise correlation length (Hoffman & Ribak 1991). We instead prefer to avoid introducing extra gaps in the unrolled timestream, to solve for  $I$  in the whole “union” region, and to obtain a continuous timestream using the  $I$  tentative solution map.

### 3.2. Noise estimation

The GLS algorithm described above requires knowledge of the noise correlation properties of each detector system. These must be estimated from the data alone, which consist of a combination of noise and (often non negligible) signal; hence the problem of noise estimation. One possible solution is to start from a non optimal, yet unbiased, estimate of the signal (e.g. obtained by naive coadding), subtract it from the data, and then use the residual to produce an estimate of the underlying noise power spectrum. The latter can be used to compute a GLS map, employed in turn to produce a new noise estimate, and the scheme can then be iterated. This is basically the approach described in Prunet et al. (2000) and Doré et al. (2001), and is de facto identical to a scheme proposed earlier by Ferreira & Jaffe (2000), although this latter algorithm is cast in a Bayesian formalism. An alternative approach proposed by Natoli et al. (2002) is to stop the noise iteration at the first step (thus saving computational time) and to increase statistical efficiency by performing a parametric fit on the estimated noise power spectral density. Although computationally more efficient, this scheme depends on the parametric assumption for the functional form of the noise power spectral density. For the purpose of this paper we do not want to lose generality, so we opt for estimating the noise by means of the iterative scheme. However, the parametric approach is a viable option for those cases where the noise spectral density can be conveniently modelled.

In the case of multi-detector, polarisation-sensitive map-making, it is a logical choice to use the global IQU maps as an estimate of the underlying signal. We will show below that this strategy achieves a substantial reduction in the estimated noise residual bias, by a factor depending on the number of detectors considered, as one would naively guess by applying a heuristic argument.

A significant fraction of the data gathered from a real world experiment are badly contaminated and cannot be used for science extraction. Corruption of these data normally occurs because of annoying events, either unpredictable (e.g. cosmic ray hits) or necessary (e.g. the ignition of a calibrator). Data loss can also affect pointing information; in this case for instance, the corresponding sky signal must be discarded even if usable per se. In any case, excerpting the affected samples destroys the timeline continuity and noise correlation. Both things must be reintroduced somehow, as is normally done by inserting a noise realisation, constrained to reproduce the correct statistical properties, while continuous (in a stochastic sense) at the gap boundaries (see Hoffman & Ribak 1991; Stompor et al. 2002). At the same time, the map making code must be warned that



**Fig. 1.** The B2K focal plane (Masi et al. 2003). The 145 GHz PSB’s are located in the lower row. The upper detector row hosts unpolarised (spider web) bolometers coupled to a polarimeter.

such samples do not contain any useful signal, a requirement trivially accomplished by associating a timeline “flag” to all samples. Flagged samples do not project over the sky but are assigned to a virtual pixel whose value is estimated by the solver. Note that this scheme is equivalent to an ideally modified experiment that observes the sky and, for flagged samples, a null signal (virtual) pixel. The latter is eventually assigned a noise variance (and covariance to real pixels) compatible with instrumental properties. As shown in the next section, we find that this approach allows for precise signal recovery within the noise free limit, while correctly minimising the noise contribution in the weak signal regime.

## 4. Application to a realistic dataset

To test ROMA we simulated and reduced a full mission scan (about 10 days) of the 145 GHz channels of the B2K experiment, using the nominal pointing solution, which includes a polarisation angle for each detector horn (de Bernardis, private communication).

B2K (Montroy et al. 2003) is basically the follow up of BOOMERanG (Crill et al. 2003) featuring polarisation sensitive bolometers (PSB) and a scanning strategy optimised to measure polarisation of the CMB and of Galactic foregrounds. The portion of the sky observed by B2K overlaps almost completely with the target of the 1998 Antarctic campaign of BOOMERanG (de Bernardis et al. 2000). However, the scanning strategy of B2K differs from the one implemented for the previous flight. The survey area of B2K is split between a  $\sim 300$  square degrees region close to the Galactic plane and a low foreground emission area of  $\sim 1300$  square degrees for CMB science. This area is in turn divided between a “shallow” and a “deep” region which differ by more than an order of magnitude in integration time per pixel. The deep region is totally embedded in the shallow one and covers about 120 square degrees.

In Fig. 1 we show the experimental setup of the B2K focal plane, with the 145 GHz PSB’s in the lower row.

In simulating the data the following scheme is employed. First we generate a high resolution map ( $1.7$  or NSIDE = 2048

in HEALPix<sup>2</sup> language) of a standard CMB sky with  $\Lambda$ CDM cosmology, and then assume a Gaussian beam of 10' FWHM for all the PSBs. The signal time streams are then obtained using the nominal B2K scan. The noise time streams are simulated by assuming the following form for the noise spectral density

$$P(f) = A[1 + (|f|/f_k)^\alpha], \quad (8)$$


where  $f_k$  is the knee frequency. We choose  $f_k = 0.1$  Hz,  $\alpha = -2$ , and amplitude  $A$  corresponding to the expected white noise level of each of the 145 GHz PSB receivers; variation of  $A$  between different receivers is within a factor of two (de Bernardis, private communication). The minimum frequency in the TOD is set by the inverse of the mission lifetime, while the nominal sampling frequency of the experiment (60 Hz) determines the Nyquist frequency. For the sake of simplicity we assume that no cross-polarisation contamination is present in the simulations described below. However, we performed several tests by varying  $\chi_{\text{pol}}$  up to a 20% level; they all concordantly show that the correct solution has been found with a small increase in the iteration count and a slightly higher residual noise level in the final maps. We fix the output map resolution to 3:7 (NSIDE = 1024) to properly sample the 145 GHz beams.

As a first step we perform noise estimation. We follow the iterative scheme explained above and compare the quality of the noise spectrum recovered when using the single PSB (temperature only) maps as the signal baseline, as opposed to using the full (global) temperature map. Not surprisingly, in the latter case, the residual noise bias is substantially minimised, especially at intermediate frequencies (see Fig. 2) when compared to the by definition unbiased estimate of the power spectrum obtained from the noise-only TOD used in the S+N simulation.

In Fig. 3 we display our results for a signal-only B2K timestream. Despite the highly realistic simulation of a complex experiment, the signal maps are recovered to high precision, better than 1% for most pixels in the  $I$  map, and better than 10% for the  $Q$  and  $U$  maps. Note that the second result is consistent, because the standard deviation of the CMB signal on the polarisation maps is roughly one order of magnitude lower than the one on the temperature map. Also note that in the deep region, the high number of hits per pixel allows recovery of the polarisation pattern with much greater precision.

The ‘‘signal plus noise’’ (S+N) maps are displayed in Fig. 4. We have accurately tested that the code preserves linearity for the GLS solution it implements: that is, running ROMA on a S+N timestream is equivalent to summing the output of two separate S and N runs. Strictly speaking this statement is true only if the code is allowed to iterate until the final solution is reached. We found in practice that the linearity is very well preserved after a reasonable number of iterations ( $\sim 100$ ) are completed.

The importance of having a joint  $IQU$  solver is underlined by Fig. 5. Here we present, as for the bottom row of Fig. 3 above, the difference (input minus output) maps for a signal-only case. However (see also the figure’s caption), here we



**Fig. 2.** The percentual ratio, as a function of frequency, between the noise power spectrum of a single detector estimated from the S+N data (by means of the iterative procedure described in the text) and the unbiased estimate of the same noise power spectrum obtained from the noise-only timestream. The black (*lower*) curve is for a single detector case, while the blue (*upper*) curve is obtained with a GLS map from all eight detectors. Note the strong suppression of residual bias in this case.

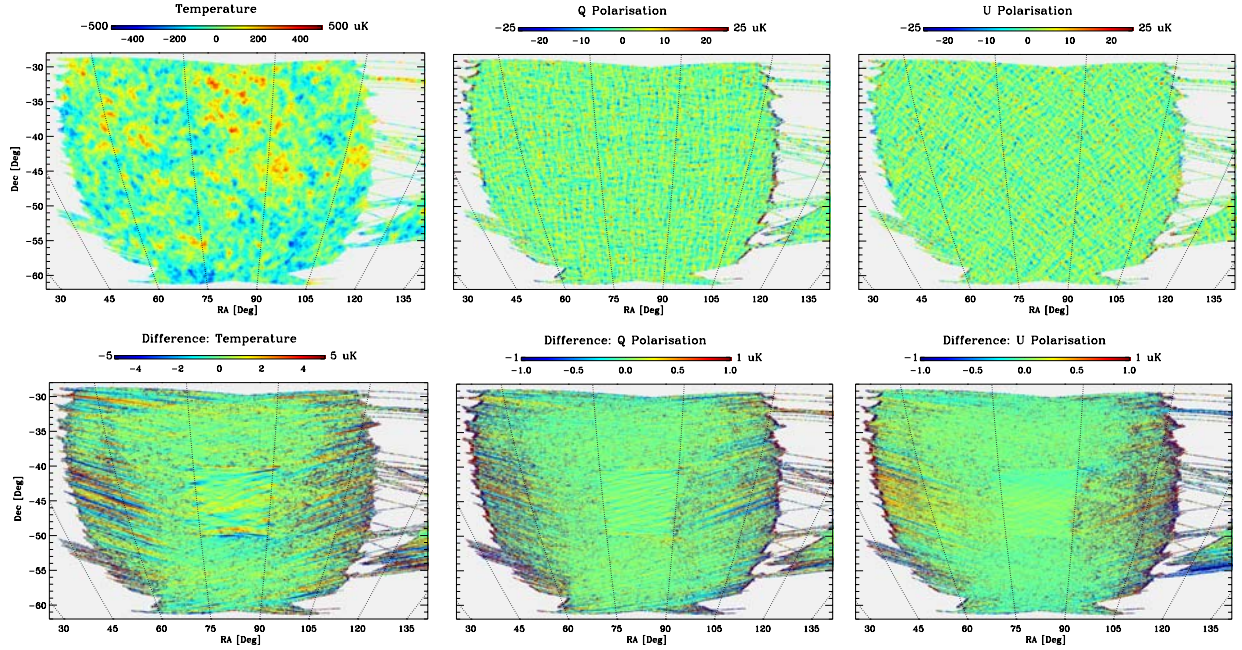
show the residuals obtained when processing a single channel map, for which no polarisation solution can be found. As one would expect, this residual is dominated by a polarisation-like pattern, which the map-making code is unable to distinguish from the temperature signal.

Most of the computational effort required by ROMA is claimed by the Fourier transforms needed to perform repeated convolutions with the  $N$  matrix. An efficient FFT library must therefore be used. Our choice fell on the publicly available FFTW3 library (Frigo & Johnson 1998), which claims, in our case, about 80% of the total computing time. Use of the FFT guarantees that the code scales linearly (per PCG iteration) with the number of timeline samples, i.e. with the dataset size. Strictly speaking, the FFT scales log-linearly with time samples. However, when performing convolutions, we only take the non-zero band of  $N$ , a tunable but constant factor (see Natoli et al. 2001). For the B2K test case under consideration, we find that retaining a noise bandwidth of size  $\sim 10^5$  is optimal and that about  $10^2$  PCG iterations are needed to reach convergence to better than  $10^{-5}$  precision in the cases under consideration. This can be achieved, for instance, in about 5 min on a 128 processor job of an IBM SP3, for the full (8 PSB,  $4 \times 10^8$  total samples) dataset (see also Fig. 6).


## 5. Summary and conclusions

We have presented our state of the art map-making code developed to jointly reduce multichannel CMB anisotropy and polarisation data. ROMA is a parallel (MPI) implementation of the GLS approach to map-making, solved by a PCG iterative method. The only assumptions are that the optical beam is purely scalar and axisymmetric and that the timeline noise is (at least piecewise) stationary and uncorrelated across different detectors. For the rest, great care was taken in tackling real

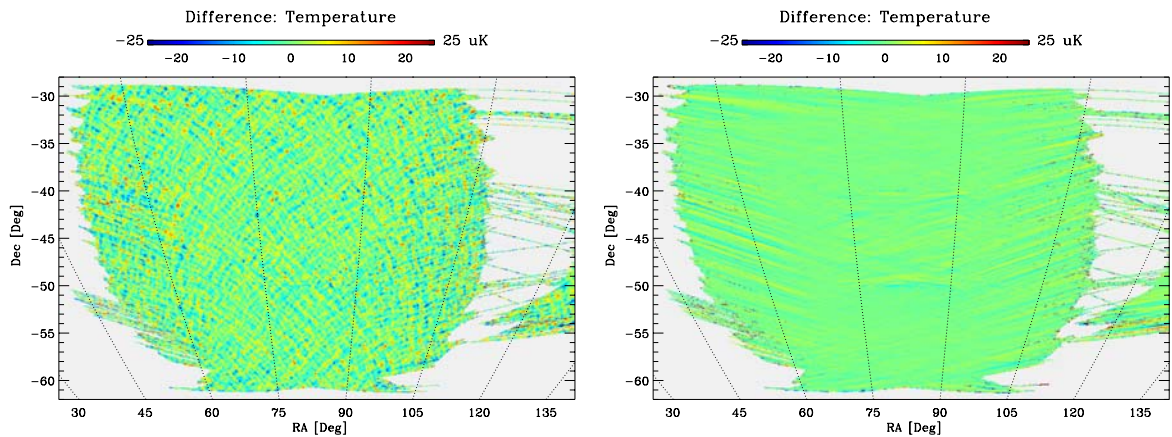
<sup>2</sup> <http://www.eso.org/science/healpix/>




**Fig. 3.** *Top row:*  $I$ ,  $Q$ , and  $U$  maps obtained by running ROMA on eight “signal-only” simulated timelines (one for each B2K bolometer). The display area is restricted to the “shallow” region. *Bottom row:* difference maps, computed by subtracting the above maps from the simulated sky that was input to the simulations. In doing so, the original maps have been degraded to NSIDE = 1024. Note the change in the colour scale.



**Fig. 4.** *From left to right,*  $I$ ,  $Q$  and  $U$  output ROMA maps on eight “signal plus noise” simulated timelines. Assumptions for the noise properties are explained in the text.



**Fig. 5.** *Left panel:* difference between the input, simulated sky, and the map output by ROMA, when processing a single-detector “signal-only” (i.e. noiseless) timeline. Both (input and output) are intensity maps; the input map was degraded to NSIDE = 1024 before subtraction. Not surprisingly, the difference is dominated by the polarisation pattern that is present in the timeline, but it cannot be properly reduced when solving for a single detector map (i.e. no polarisation map can be found). *Right panel:* same as left, but here the  $I$  map is estimated jointly from all eight detectors.  $Q$  and  $U$  maps are also computed but are not shown here. This figure is de facto the one displayed in the left panel, bottom row of Fig. 3, but on a different scale.



**Fig. 6.** Inverse time for a single PCG iteration as a function of the number of processor elements ( $N_{pe}$ ) used by ROMA, for the full B2K dataset (eight timelines, consisting of  $5 \times 10^7$  samples each). The machine used is an IBM Power3 RS/6000. ROMA scales optimally until  $N_{pe} \sim 80$ .

world issues, including cross-polarisation, multi-detector noise estimation, and the problem of missed data. As a test case we used ROMA to reduce eight PSB (145 GHz) timelines of highly realistic simulated B2K data. We show that the *IQU* maps can be recovered with great precision in the signal-only case, while attaining the usual GLS (“optimal”) noise suppression in the noisy case. To our knowledge, this is the first joint temperature and polarisation map-making code demonstrated to work on a realistic dataset. The code scales linearly with the dataset size, while its parallel behaviour is quasi-optimal, thus represents a viable option for reducing present and forthcoming large datasets, including PLANCK.

*Acknowledgements.* This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the US Department of Energy under Contract No. DE-AC03-76SF00098. We thank P. de Bernardis and the whole BOOMERanG collaboration for having provided us with the B2K scan and instrumental performances. We acknowledge use of the HEALPix package (Górski et al. 1999, 2004) and of the FFTW library (Frigo & Johnson 1998). We thank the CASPUR (Rome-ITALY) computational facilities for computing time and technical support. The authors wish to thank the PLANCK CTP working group and, in particular, C. M. Cantalupo and J. D. Borrill for stimulating discussions.

## References

Armitage, C., & Wandelt, B. D. 2004, *Phys. Rev. D*, 70, 123007  
 Bennett, C. L., Halpern, M., Hinshaw, G., et al. 2003, *ApJS*, 148, 1

- Cabella, P., & Kamionkowski, M. 2004  
 [arXiv:astro-ph/0403392]  
 MADmap is a publicly available GLS map-making code by Cantalupo, C. M. et al.  
 see: <http://crd.lbl.gov/~cmc/MADmap/doc/>  
 Chandrasekhar, S. 1960, *Radiative transfer* (New York: Dover)  
 Couchot, F., Delabrouille, J., Kaplan, J., & Revenu, B. 1999, *A&AS*, 135, 579  
 Crill, B. P., Ade, P. A. R., Artusa, D. R., et al. 2003, *ApJS*, 148, 527  
 de Bernardis, P., Ade, P. A. R., Bock, J. J., et al. 2000, *Nature*, 404, 955  
 Doré, O., Teyssier, R., Bouchet, F. R., Vibert, D., & Prunet, S. 2001, *A&A*, 374, 358; see also  
<http://ulyse.iap.fr/cmbsoft/mapcumba/>  
 Ferreira, P. G., & Jaffe, A. H. 2000, *MNRAS*, 312, 89  
 Frigo, M., & Johnson, S. G. 1998, *ICASSP Conf.*, 3, 1381, also see  
<http://www.fftw.org/>  
 Gorski, K. M., Hivon, E., & Wandelt, B. D. 1999, in *Proc. of the MPA/ESO Cosmology Conference Evolution of Large-Scale Structure*, ed. A. J. Banday, R. S. Sheth, & L. Da Costa, PrintPartners Ipskamp, NL, 37  
 [arXiv:astro-ph/9812350]  
 Gorski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759  
 Hoffman, Y., & Ribak, E. 1991, *ApJ*, 380, L5  
 Johnson, B. R., Abroe, M. E., Ade, P., et al. 2003, *New Astron. Rev.*, 47, 1067  
 Jones, B., Montroy, T., Contaldi, C., & Mactavish, C. 2005, *Priv. Comm.* JQU is a Jacobi based, iterative map-making code developed for the North American B2K pipeline  
 Kaiser, N. 1983, *MNRAS*, 202, 1169  
 Kogut, A., Spergel, D. N., Barnes, C., et al. 2003, *ApJS*, 148, 161  
 Kovac, J. M., Leitch, E. M., Pryke, C., et al. 2002, *Nature*, 420, 772  
 Lin, Y., & Wandelt, B. D. 2004  
 [arXiv:astro-ph/0409734]  
 Lupton, R. 1993, *Statistics in theory and practice* (Princeton: Princeton University Press)  
 Masi, S., Ade, P. A. R., Balbi, A., et al. 2003, *Mem. Soc. Astron. It. Suppl.*, 2, 54  
 Montroy, T., Ade, P. A. R., Balbi, A., et al. 2003, *New Astron. Rev.*, 47, 1057  
 Natoli, P., de Gasperis, G., Gheller, C., & Vittorio, N. 2001, *A&A*, 372, 346  
 Natoli, P., Marinucci, D., Cabella, P., de Gasperis, G., & Vittorio, N. 2002, *A&A*, 383, 1100  
 Prunet, S., Ade, P. A. R., Bock, J. J., et al. 2001, *Proc. of the MPA/ESO/MPA Conf. Mining the Sky*  
 [arXiv:astro-ph/0101073]  
 Rees, M. J. 1968, *ApJ*, 153, L1  
 Revenu, B., Kim, A., Ansari, R., et al. 2000, *A&AS*, 142, 499  
 Seljak, U., & Zaldarriaga, M. 1996, *ApJ*, 469, 437  
 Stompor, R., Balbi, A., Borrill, J. D., et al. 2002, *Phys. Rev. D*, 65, 022003  
 Wright, E. L. 1996, in *IAS CMB Data Analysis Workshop*  
 [arXiv:astro-ph/9612006]