Doc No:	SC2/ANA/S210/001
Vers:	1.7
Category:	Musings
Doc Type:	LaTeX
State:	Issued
Author:	Douglas Scott
Date:	2003/12/08

Map-making in different noise regimes

1 Introduction

In astronomical imaging, we are usually trying to obtain an accurate estimate of a map of a chunk of the sky using a detector with real noise.

Let us start by thinking of a part of the sky being split into a pixelised map, m_p, labelled by $p = 1, \ldots, N_{\rm pix}$ in some simple scheme. This could be a rectangular grid of pixels, with $p = 1$ being the top left pixel and the numbering going left to right along each row, for example.

Now consider a collection of detectors (bolometers for our purpose) labelled by $b = 1, \ldots, N_{\rm bol}$, each taking simultaneous measurements. We can ideally think of these measurements as being instantaneous, a reasonable approximation provided that the spatial binning is small compared with the beam-size – otherwise we need to deal with the extra complication of integrating the data over a finite time (probably limited by the bolometer time constant).

We shall not deal with that issue here, but focus on the sources of noise which enter into observations, and which limit the estimation of the underlying map.

If the detectors are in a filled array and are 'fully sampled', then there might be a simple relationship between the pixel array and the detector array. But the array will usually be moving around relative to the sky, so this is only ever true at one instant. It is also not true in general for real detector arrays, since there will be curvature in the focal plane and other effects that we don't want in our map array.

The problem is to estimate the underlying map m_p given a set of measurements taken at many time intervals (labelled by $t = 1, \ldots, N_{\rm tim}$) *in the presence of noise*.

Without noise the solution is trivial – one simply throws the flux from a detector into whichever pixel it is looking at, and averages. In the presence of purely white noise one just takes the weighted average of all the hits of a pixel. It is in more realistic noise scenarios that things are more complicated.

2 The general case

In general the data for bolometer b at time t can be written:

$$d_t^b = g_t^b \times \left(\mathcal{A}_{tp}^b m_p\right) + n_t^b, \tag{1}$$

where the term inside the brackets is a matrix multiplication. Here $\mathcal{A}_{tp}^b$ is the pointing matrix for bolometer b at time t, assigning each time sample to a pixel or set of pixels to which the bolometer is pointing; there is one of these matrices for each bolometer (but for regular array structures they will be closely related). These will typically be very sparse matrices, with only a few entries per row. The other symbols are g_t^b, which is the gain for bolometer b at time t, and n_t^b, which is the noise term.

The solution for m_p given d_t is clearly an inversion problem. The most general problem, where g and n have arbitrarily complex variations in time and are unrelated between detectors, is insoluble. However, things are never quite this hopeless in real systems (otherwise we would never see astronomical images!). How the inversion problem can be solved in practice depends on what noise regime one is in, as discussed below.

But first a few remarks about the pointing matrix. $\mathcal{A}$ could in principle contain the beam function, i.e. at time t a bolometer would be measuring signal from a group of sky pixels with weights given by the beam shape. However, we will typically *not* be attempting beam deconvolution here (a whole other topic, which is only valuable in the high signal-to-noise regime). For symmetric beams the linearity of the algebra means that we can pull the beam convolution through, and assume that what we are trying to estimate is actually $m_p = m'_p * \text{Beam}$, where m' is the underlying (not convolved with the beam) sky, which we don't imagine we can actually measure. If the beam is very asymmetric, then things are more complicated because what you measure depends on the orientation of the beam at the time. It is desirable (and often the case) that beams are fairly symmetric and Gaussian. So let us restrict ourselves to the case of symmetric beams, where what we are trying to estimate is the beam-convolved sky image.

The next point about the pointing matrix is related to chopping. In a 'total power' detector, the rows of $\mathcal{A}$ (describing the pointing at each time) will contain a '1' corresponding to the entry for the pixel that a bolometer is pointing at, and '0' everywhere else. In a 'single difference' (or '2-beam') experiment, there will be a '$+1$' and '-1' entry in each row. In a 'double difference' (or '3-beam') experiment, there will be a '-0.5', '$+1$', '-0.5' pattern, etc.

3 Single Detector

The typical example from Cosmic Microwave Background experiments (e.g. see[9]) has been to assume that $g = 1$ (or that it is slowly varying, and fit with a piecewise constant function) and if there is more than a single detector to treat them one at a time.

Assuming that the noise covariance function is known (or equivalently the power spectrum of the noise) there is a minimum variance solution for the map

$$m_p = \left(\mathcal{A}^{\mathrm{T}}\mathcal{N}^{-1}\mathcal{A}\right)^{-1}\mathcal{A}^{\mathrm{T}}\mathcal{N}^{-1}d_t, \tag{2}$$

and the pixel-pixel correlation matrix is

$$\mathcal{N}^{\mathrm{pix}} = \left(\mathcal{A}^{\mathrm{T}}\mathcal{N}^{-1}\mathcal{A}\right)^{-1}. \tag{3}$$

If the time-domain noise covariance function is not known, then it has to be estimated from the data in some way, e.g. using previous data, or perhaps simultaneously along with the estimate of the map.

The formal calculation of equation (2) is usually impractical for large data-sets. In particular $\mathcal{N}^{-1}$ requires inversion of a $N_{\rm tim} \times N_{\rm tim}$ matrix, which can be extremely large. However, CMB experimental teams have developed approximate solutions, using Fourier techniques, iterative approaches, etc.

When there are several detectors, the CMB approach has usually been to simply make a map for each detector and then to co-add them.

4 Difference measurements

With a set of difference measurements one can still uncover the underlying sky, but there are some caveats. Consider the case of a single difference experiment for simplicity. Then the measurements are not sensitive to the overall DC (monopole) component of the map, which means that this DC level cannot be determined by the matrix inversion. In other words the inversion is singular for this mode. The method still works provided that this fact is taken into account. Similarly, if there is only a single chop throw and direction, then there are a set of modes (stripes perpendicular to the direction of the chop and with separation equal to the chop throw) to which the measurements are insensitive. Again this makes the matrix inversion singular in a set of known modes, which can be dealt with explicitly.

In the presence of noise there will be many more modes for which the inversion is in principle stable, but in practice numerically unstable. The way to avoid this is to have a lot of redundancy, i.e. to have differences between each pixel and many other pixels. Another way to say the same thing is that in the timestream you want to be able to distinguish spatial structure (i.e. real stuff in the map) from temporal structure (which is just drifts in the noise). A stable inversion can therefore be achieved by rapidly scanning across an image with cross-linking scans. Differencing over large angles (the *COBE* and *WMAP* approach) can make things more stable still. But differencing is not always possible, and there are some simplifications with using 'total power' devices, such as is the case for SCUBA-2.

In fact the *only* way to robustly measure the longest wavelength modes in an image is to move the detectors across the image in a time which is not long compared with the timescale of drifts in the gain and noise.

5 Stable gain

Now let us look at the gain term. This can be split into 2 contributions: one is the gain of the detector itself; and the other is the atmospheric transmission. Both are varying with time, and both are in principle different for each detector.

For SCUBA[4] the relative gains of the bolometers are extremely stable, and the values are read by the observer from a file which is updated only intermittently. However, it is important to realise that this is the gain for a *chopped* measurement, and is therefore going to be much more stable than the gain for a 'total power' measurement.

If we can assume that the gain from both detectors and atmosphere is relatively stable, and that the noise is relatively white, then we are in a simple regime. A reasonable estimate of the map can be made by using the known gain values and just co-adding the data into the pixels.

6 Common-mode gain

The detector gain may also be split into 2 parts: a part which is fluctuating independently for each detector; and a 'common-mode' part (due to temperature variations, for example) which is the same for each detector. This latter contribution has essentially the same effect as atmospheric gain (i.e. transmission) variations, which will affect all of the detectors at once.

For SCUBA there is little indication that the atmosphere varies across the ~ 2 arcmin size of the array.[1] This may not be true for larger arrays of course, but for SCUBA-2 the sampling rate is much higher and so this may still be a good approximation. The atmospheric transmission variations can be corrected for using an independent monitor of the atmosphere. For SCUBA this comes from the CSO $\tau-$meter or sky-dips performed over timescales of many minutes, or from the water vapour radiometer on a timescale of seconds. These approaches will not generally correct for the instrumental common-mode drifts, however. But it is also possible to correct for the gain drifts using the temporal variations of a bright source in the image, which would remove both atmospheric and detector common-mode gain variations.

In any case, what is needed is an estimate of the relevant drift timescales, or even better, the gain autocorrelation function. Are these timescales long or short compared with the rate at which atmospheric opacity is monitored or at which data are taken? And do the atmospheric fluctuations have better or worse long term drifts than the detector fluctuations?

It is probably a good approximation to assume that the atmospheric transmission fluctuations are the same for all the detectors. Of course this approximation could be relaxed to assuming that there is a low order polynomial across the array, or perhaps a checker-board pattern of constant values. Similarly one can imagine detector gain fluctuations which are only common to groups of detectors (rows across the array, or bolometers on the same electronics card, for example).

If the detector gain and atmospheric transmission variations can be corrected for, then we are back to the problem of solving $d = \mathcal{A}m + n$.

7 Slowly varying gains and offsets

It is now worth describing a little about the regime in which astronomers observing in the near-IR often find themselves. One immediate difference with sub-mm observing is that the detectors are 'read-noise' limited if they are read very quickly. And hence there is no realistic option for taking data rapidly. Instead integrations are allowed to build up with enough counts that the read noise is sub-dominant. The observer then ends up with a stack of images to combine.

The technique of 'dithering' (moving the frame centres around in a pseudo-random or predetermined pattern) helps to move defective detectors around on the sky, as well as to provide better sampling in the reduced image. 'Dark frames' are taken to determine the offsets of the detectors (a contribution to the noise n_t^b), and 'flat fields' to determine the relative gains. This method can only work provided the atmospheric noise does not dominate over the integration time and that the gains and the offsets are both slowly varying relative to the timescales on which the frames are taken. If atmospheric emission (or transmission) was varying rapidly (as it does in the sub-mm) then some method would have to be found of removing that variation in real time.

However, in practice the gains and offsets are fairly constant, and so can be determined using darks and flats taken as close as possible to the actual observations. Fixsen et al. (2000,[3] also Arendt et al. 2000[2]) describe an approach whereby the data themselves are used to solve for a single gain and offset for each detector element at the same time as solving for the image. Their method could obviously be extended to cover cases where the gains and offsets are fit by some slowly varying function or piecewise constant. They develop a rule of thumb for assessing whether the dither pattern has coupled the pixels well enough to be able to separate these detector parameters from the pixel estimates. It basically says that you want good sampling in the 2D Fourier plane. This rule of thumb may also be applicable to other regimes. However, for time-varying drifts and offsets (or atmosphere) it is clear that there is also a tight requirement on moving detectors across the sky rapidly compared with the timescale of drifts.

8 Common-mode noise fluctuations

For any individual bolometer

$$\mathcal{N}^b = \langle n_t^b . n_{t'}^b \rangle \tag{4}$$

is the noise correlation function. This is usually assumed to be a function of $|t - t'|$ only (stationarity).

The form of the power spectrum of the noise is often

$$P(f) = P_0 \left(1 + f_{\rm knee}/f\right), \tag{5}$$

where P_0 is the white noise power and $f_{\rm knee}$ is the 'knee-frequency', below which the $1/f$ noise term dominates. A spectrum which increases as $1/f$ is very common in real physical systems, although in practice systems may have even stronger long-term drifts, with power spectra varying approximately as $1/f^{1.5}$ or $1/f^2$.

These drifts cause 'baseline' variations in the timestream. They are removed by collecting data faster than about $f_{\rm knee}$. However, if this isn't possible, then the drifts can be at least partly removed by cross-linking scans, or otherwise rapidly connecting different pixels in a way which allows one to distinguish spatial from temporal structure.

The noise in the timestream, n_t^b, comes partly from atmospheric emission and partly from the detector offsets. The atmospheric part is expected to be largely common to all detectors and hence can be removed by removing the average of the array at every timestep t. This will of course have the effect of correlating all of the detector values, but for a large enough number of detectors this effect is negligible. Subtracting this average will also remove any common-mode offsets among the detectors.

Once this average-removal has been carried out, the remaining noise fluctuations may be close to white. This seems to be true for SCUBA, for example. For SCUBA-2, the data will be taken at the much faster rate of 200 Hz, and so one might expect the $1/f$-type effects to be very much reduced. However, one would like to know whether atmospheric or detector noise dominates in the raw data, and then which one dominates in the average-removed data.

For SCUBA-2, if it proves to be the case that almost all the noise drifts are removed by removing the array average at 200 Hz, then it is possible that the simple 'STARE' mode[5] will work. It is also feasible that a slightly more complex procedure (for example removing a common-mode noise among some groups of detectors, while also removing a 2D 2nd-order polynomial across the array) will work fairly well instead.

9 Residual noise drifts

However, one must be concerned about the possibility of all of the atmospheric and instrumental noise drifts not being fully removed with such a simple procedure.

It is the aim of the 'DREAM' mode of SCUBA-2[10,7] to provide a way of removing these variations by performing a rapid, closed dither pattern over a small region, and effectively removing a gradient in the timestream around the loop. This method had some partial success on SCUBA,[8] but it cannot be claimed that it has been demonstrated to work effectively.

If there are residual drifts, then in a 'SCAN' mode[5,6] these would result in 'striping' across the maps. Such stripes can be reduced by removing low order polynomials across each scan, although that can also remove signal. It is better to simultaneously distinguish the drifts in the timestream while estimating the underlying map. This can be done provided that the scanning pattern is complex enough to inter-link pixels. A large array of detectors has a distinct advantage, because each pixel in the final map is scanned over by multiple detectors at different times.

Van Engelen, Borys & Scott (in preparation) tried this technique for SCUBA data taken along 3 scan directions. This proved feasible provided an assumption was made that each bolometer has the same shape of power spectrum, only allowing a scaling with the r.m.s. of each bolometer. An iterative technique simultaneously estimating the noise power spectrum and map seemed to be successful. Presumably one can develop a model for the detector noise as data are acquired, which can improve upon such methods.

With a large number of bolometers and the ability to move rapidly on the sky, one can utilise the ability to directly remove common-mode terms, and to cross-link distant pixels on relatively short timescales. Then it is only a matter of figuring out the optimal strategy for scanning, removing low-order effects among bolometers, and iteratively producing the best final map from the timestream of data.

10 Conclusions

It is important to understand the various sources of noise and other instrumental and environmental effects which are part of the timestream of data. The various terms need to be quantified for each particular experimental set-up in order to determine which regime the data-set falls into. Different regimes call for different analysis approaches, or combinations of approaches. For SCUBA-2 we should develop a clear idea of the regimes we are in so that we can choose the most appropriate data analysis techniques, and carry out simulations in order to optimize these techniques.

References

[1] E. N. Archibald, T. Jenness, W. S. Holland, I. M. Coulson, N. E. Jessop, J. A. Stevens, E. I. Robson, R. P. J. Tilanus, W. D. Duncan, and J. F. Lightfoot. On the atmospheric limitations of ground-based submillimetre astronomy using array receivers. *Monthly Notices of the Royal Astronomical Society*, 336:1–13, October 2002.

[2] R. G. Arendt, D. J. Fixsen, and S. H. Moseley. Dithering Strategies for Efficient Self-Calibration of Imaging Arrays. *Astrophysical Journal*, 536:500–512, June 2000.

[3] D. J. Fixsen, S. H. Moseley, and R. G. Arendt. Calibrating Array Detectors. *Astrophysical Journal Supplement*, 128:651–658, June 2000.

[4] W. S. Holland, E. I. Robson, W. K. Gear, C. R. Cunningham, J. F. Lightfoot, T. Jenness, R. J. Ivison, J. A. Stevens, P. A. R. Ade, M. J. Griffin, W. D. Duncan, J. A. Murphy, and D. A. Naylor. SCUBA: a common-user submillimetre camera operating on the James Clerk Maxwell Telescope. *Monthly Notices of the Royal Astronomical Society*, 303:659–672, March 1999.

[5] B. D. Kelly. SCUBA-2 modes and data processing. SCUBA-2 Project SC2/ANA/S100/028, 2001.

[6] B. D. Kelly. SCUBA-2 scanmap simulation. SCUBA-2 Project SC2/ANA/S100/038, 2002.

[7] B. D. Kelly. SCUBA-2 noise-free DREAM simulations. SCUBA-2 Project SC2/SOF/S200/012, 2003.

[8] R. S. Le Poole and H. W. van Someren Greve. DREAM: Dutch real-time acquisition mode for SCUBA. In *Proc. SPIE Vol. 3357, p. 638-643, Advanced Technology MMW, Radio, and Terahertz Telescopes, Thomas G. Phillips; Ed.*, pages 638–643, July 1998.

[9] R. Stompor, A. Balbi, J. D. Borrill, P. G. Ferreira, S. Hanany, A. H. Jaffe, A. T. Lee, S. Oh, B. Rabii, P. L. Richards, G. F. Smoot, C. D. Winant, and J. P. Wu. Making maps of the cosmic microwave background: The MAXIMA example. *Physical Review D*, 65:22003–+, January 2002.

[10] H. W. van Someren Greve. DREAM algorithms for SCUBA-2. SCUBA-2 Project SC2/ANA/S100/046, 2002.