Determining the optimum scan map strategy for *Herschel*-SPIRE using the SPIRE photometer simulator

T. J. Waskett,[1]* B. Sibthorpe,[1] M. J. Griffin[1] and P. F. Chanial[2]

[1]*School of Physics and Astronomy, Cardiff University, Queens Buildings, The Parade, Cardiff CF24 3AA*
[2]*Physics Department, Imperial College London, South Kensington Campus, London SW7 2AZ*

ABSTRACT

The forthcoming *Herschel* space mission will provide an unprecedented view of the far-infrared/submillimetre Universe, with the SPIRE instrument covering the 200–670 μm wavelength range. To obtain the best quality of astronomical data from such an expensive mission the observing modes must be optimized as far as possible. This paper presents the possible scanning strategies that can be utilized by the SPIRE photometer, within the limitations imposed by the *Herschel* spacecraft. Each strategy is investigated for effectiveness by performing simulated observations, using the SPIRE photometer simulator. By quantifying the data quality using a simple metric, we have been able to select the optimum scanning strategy for SPIRE when it begins taking science data within the next couple of years.

Additionally, this work has led to the development of a specific SPIRE mapmaking algorithm, based on the CMB code MADmap, to be provided as part of the SPIRE data pipeline processing suite. This will allow every SPIRE user to take full advantage of the optimized scan map strategy, which requires the use of maximum likelihood mapmakers such as MADmap.

Key words: instrumentation: photometers – methods: observational – submillimetre.

1 INTRODUCTION

This paper explores the capabilities of the *Herschel*-SPIRE instrument, through the application of the SPIRE photometer simulator (henceforth referred to simply as 'the simulator', Sibthorpe et al. 2004). This paper concentrates on the practicalities of observing with SPIRE using scan map mode, and how to optimize the scan map strategy.

SPIRE (Griffin et al. 2006) is one of three science instruments to fly on ESA's forthcoming *Herschel Space Observatory* (hereafter *Herschel*, Pilbratt 2005) and covers the long-wavelength range of the far-infrared (far-IR)/submillimetre waveband, 200–670 μm, with both a photometer and imaging spectrometer. Before *Herschel* flies, many of the telescope and instrument characteristics will be well known and understood but some will only be fully revealed in flight, making preparation, in the form of simulations, an important undertaking.

The *Herschel* spacecraft offers only a limited number of telescope movements for performing observations. These movements are restricted by the orbit geometry, the necessity of keeping the spacecraft sun-shield facing towards the Sun, and the fact that a space-based observatory does not have the large inertia of the Earth against which movements can be made.

When optimizing instrument observing modes it is crucial to take the spacecraft limitations into account so that data quality is not compromised. Data quality can mean many things but here we are primarily concerned with the effects that $1/f$ noise has on the ability of SPIRE to reconstruct submillimetre flux from the sky.

The aim of this paper is to demonstrate the possible scan map observations that can be performed by *Herschel*-SPIRE and to select the strategy that provides the best possible quality of SPIRE data, under the widest range of potential observations.

We shall first describe the limitations within which SPIRE can perform observations before describing $1/f$ noise in more detail. These two factors form the core of this work, so they deserve a thorough description. The SPIRE photometer simulator is also briefly described. The rest of this paper details the simulated observations that we performed in order to test the various observing possibilities, and the pros and cons of each. The data quality is quantified for each observing strategy and, combined with practical considerations, a selection of the optimum strategy is made. This selection is now implemented as the default strategy for SPIRE scan map mode, so when *Herschel* flies all scan map observations will benefit from this work.

1.1 Observatory restrictions

SPIRE observations are bound by the restrictions imposed by the *Herschel* spacecraft. There are two main types of restrictions: those
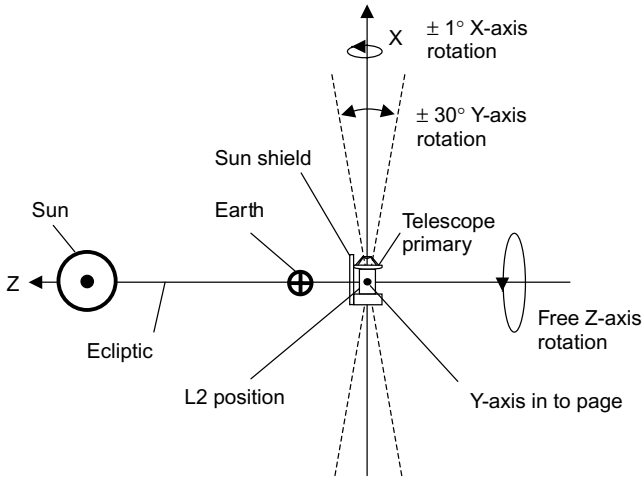
---

*E-mail: tim.waskett@astro.cf.ac.uk

Figure 1. *Herschel* orbits around the unstable Lagrangian point L2, co-orbiting the Sun along with the Earth. Three axes are defined with respect to the *Herschel* spacecraft: *X* is the boresight, *Z* is the *Herschel*–Sun direction and *Y* is the perpendicular to these other two. The necessity of keeping the main mirror in the shade restricts the direction in which spacecraft can point on the sky, as indicated by the dashed lines and arrows.

produced by the orbital geometry of *Herschel*; and those caused by the limited mobility of the spacecraft itself. The SPIRE instrument also dictates the way in which observations are carried out, to a certain extent, with only a limited number of possible options allowed within the limits imposed by the observatory restrictions.

1.1.1 Herschel orbit geometry

Herschel is a far-IR/submillimetre observatory and as such requires low-temperature operation. The primary source of background radiation for the science instruments is thermal emission from the primary and secondary telescope mirrors, so to keep the telescope cool (~80 K) the whole spacecraft is shaded from the Sun by a large sun-shield. Additionally, *Herschel* will be placed into orbit around the second Lagrangian point (L2) of the Sun–Earth system, ~1.5×10^6 km anti-Sunward from the Earth. The spacecraft will be orientated such that the sun-shield always faces towards the Sun.

Herschel can rotate freely about the line joining itself to the Sun (the *Z*-axis). It can also tilt by up to 30° either towards or away from the Sun (the *Y*-axis) producing a visibility annulus 60° wide. However, the size of the sun-shield dictates that rotations about the boresight (the *X*-axis) are restricted to only ±1°, else the telescope and spacecraft cryostat would be exposed to sunlight. This final restriction effectively fixes the orientation of the *Herschel* field of view (FOV) on the sky for any given date. Fig. 1 shows the *Herschel* orbit geometry in a schematic way, summarizing the above points (restrictions imposed by the Earth and Moon are ignored here, for simplicity).

As the spacecraft travels around its orbit the FOV rotates with respect to any particular point on the sky, which can provide certain advantages, described in Section 2.3. High latitudes are visible to *Herschel* throughout the year and so can be observed with any FOV orientation; however, fields on the ecliptic plane are only visible to *Herschel* twice a year, roughly 6 months apart. Between these two times the *Herschel* FOV rotates by exactly 180° on the sky. Areas of the sky between the ecliptic plane and the poles will have varying

degrees of visibility, with corresponding ranges of possible FOV orientations, depending on the date the observation is carried out.

1.1.2 Herschel pointing modes

The *Herschel* spacecraft can perform various pointing operations that can be combined together, along with instrument operations, to make up an observation. In their simplest forms, they involve either pointing at a fixed position, accelerating up to speed or slewing the telescope in a straight line (a great circle on the sky) at a constant speed.

For SPIRE, an observation of a large area requires scanning mode, which combines a series of parallel slews together. All the slews (scan legs) must be the same length and the telescope must come to a stop after each scan leg before traversing to the starting point for the next one.

A single scan map observation is carried out as follows. Starting from rest, some way off one corner of the target area, the telescope accelerates up to full scanning speed by which point the leading edge of the SPIRE FOV enters the target map area. The array then passes over the target area, taking data continuously along the scan leg, until the trailing edge leaves the far side. The telescope then decelerates to a stop, traverses orthogonally to the scan direction and accelerates back the other way to perform another scan leg. Scan legs are added until the whole target area has been observed. This type of scan pattern is called a variety of names, including raster mapping, serpentine scanning and boustrophedon (turning as an ox in ploughing) scanning. We shall refer to it as simply scan mapping, or scan map mode. Fig. 2 shows a schematic of a typical SPIRE scan map observation. The SPIRE arrays are rectangular with a FOV of ~4 × 8 arcmin ^{2} and are rotated with respect to the scan direction
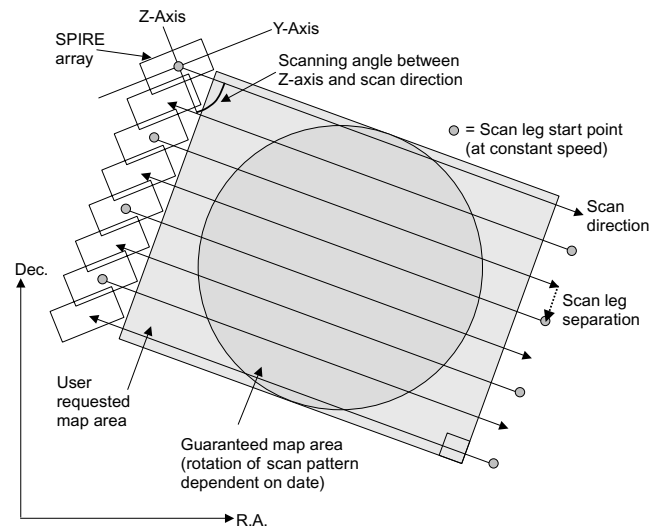

Figure 2. A typical SPIRE scan map is built up by a series of parallel scan legs as the telescope slews back and forth over the required area. Between each scan leg is the transit period, involving deceleration; traversing to a point in advance of the next scan leg starting point; and then acceleration, so that the telescope is travelling at full speed when the SPIRE FOV reaches the edge of the required map area. Note that the guaranteed map area is a circle even though the total map area is a rectangle. This is due to the rotation of the *Herschel* FOV over time, as explained in the text. If the same observations were carried out at some later date, the entire scan pattern may be rotated with respect to the first. The circular area is always observed, however, and so is the only part of the map that can be guaranteed for every possible observing date.

because the arrays do not instantaneously fully sample the sky. A particular scan angle is chosen so that the final map is fully sampled, as explained in Section 1.1.3.

More complicated scan patterns, such as the Lissajous scan employed by SHARC-II on the Caltech Submillimeter Observatory (Dowell et al. 2001), or the 'Pong' scan pattern to be employed by SCUBA-2 on the James Clerk Maxwell Telescope (Holland et al. 2006), are impractical for *Herschel*: Lissajous requires curved scan lines, which *Herschel* simply cannot perform; and the difficulty of moving a large spacecraft quickly in different directions would cause Pong to require a large run-off area outside of the requested area, to allow space for the turnarounds, resulting in a low observing efficiency.

A second mapping mode is available to SPIRE for observing small fields. Jiggle map mode involves *Herschel* pointing at a fixed part of the sky and using the SPIRE internal beam steering mirror to ensure the map is fully sampled, as well as chopping between an on- and off-source position. A single jiggle pointing will produce a map of $\sim 4 \times 4$ arcmin ^{2}. Any SPIRE observation larger than this will be performed in scan map mode. Sibthorpe, Waskett & Griffin (2006) describes these and other SPIRE observing modes in more detail.

1.1.3 SPIRE restrictions

So for large maps we are limited to the simple back-and-forth scan map strategy. The key parameters for SPIRE when using this mode are the scan speed, the angle between the array *Z*-axis and the scan direction (the scan angle) and the separation between the adjacent scan legs. The scan speed has already been optimized by the requirement to compromise between mapping speed and keeping the turnaround overheads as low as possible. In operation there will be a choice of either 30 or 60 arcsec s^{-1}, the faster speed being used only for the largest maps. The scan angle has only a limited number of allowed values, while the scan leg separation is set by the scan angle and the requirement for the final map to be covered uniformly by the SPIRE arrays (Sibthorpe et al. 2006).

Because the SPIRE arrays are hexagonally packed, feedhorn-coupled bolometers the detector–detector spacing is twice the diameter of the beam full width at half-maximum (FWHM). Therefore, SPIRE does not instantaneously fully sample the sky. To fill in the gaps between the detectors the arrays must be scanned at a particular angle with respect to one of the three symmetry axes, to ensure sampling on the sky at least every half a beam FWHM. There are thus 12 possible directions in which SPIRE can scan. For practical purposes, however, there are three unique scan angles, all others being either reflections or rotations of these basic three. One direction lies close to the short axis of the array, one lies close to the long axis and one lies in between, roughly on the diagonal. As such we shall refer to these possibilities as the 'short', 'long' and 'diagonal' scan directions. Fig. 3 shows these three possibilities in schematic form.

Table 1 summarizes the three basic angles along which the *Herschel* spacecraft can scan with respect to the *Z*-axis (short axis) of the SPIRE arrays. These angles all provide data that are sampled at double the Nyquist frequency (i.e. 1/4 of a beam FWHM), something that is possible because of the large number of detectors in the SPIRE arrays (Sibthorpe et al. 2006). To ensure uniform coverage of the sky the scan leg separations are different for the three different scan directions because the SPIRE FOV is rectangular. This also results in the different scanning angles having very different effec-
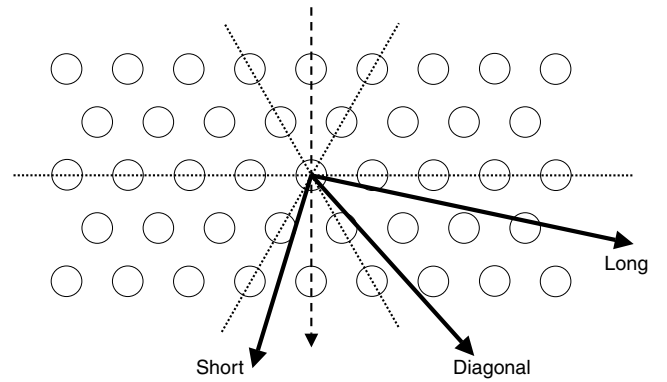

Figure 3. Schematic showing the SPIRE 500-μm array with circles representing the feedhorns. The dotted lines are the three symmetry axes of the hexagonal close packing arrangement, the dashed line is the *Z*-axis and the thick solid lines are the three possible scan directions mentioned in the text.

Table 1. Basic parameters for the different scan angles. 'Scan angle' refers to the angle between the short axis of the SPIRE array (*Z*-axis) and the direction of the scan leg, in degrees. 'Step size' indicates how far apart adjacent scan legs are on the sky, in arcseconds. 'Relative τ' indicates the relative effective integration time, per map repeat, with 'long direction scanning' being unity.

Direction	Scan angle	Step size (arcsec)	Relative τ
Long	77°.6	235	1.000
Diagonal	$\pm 42°.4$	348	0.665
Short	$-17°.6$	455	0.515

tive integration times. Effective integration time is defined here as the length of time that any point within the observed map area is actually observed by the SPIRE arrays as they pass over it during the course of an observation. Long scanning has nearly double the effective integration time of short scanning because the arrays are roughly twice as long as they are wide. Diagonal scanning falls in between these two extremes. One map repeat is a single coverage of the sky by the scan pattern in question. To increase the effective integration time of an observation the scan pattern can be performed again to add an additional map repeat. Therefore, for every one long map repeat, two short map repeats would be required to achieve approximately the same sensitivity.

The signs in front of the angles in Table 1 will become relevant for Section 2.3.

1.2 1/f noise

Noise is present in all instrumental systems and takes many forms depending on the type of detector and wavelength of light being detected, amongst other factors. The detectors used in SPIRE are semiconductor bolometers sensitive to far-IR/submillimetre radiation. Far-IR bolometers are essentially very sensitive thermometers – a change in the absorbed radiant power results in a change in the resistance of the bolometer, and hence the voltage across it, which is measured by the readout electronics. Gaussian noise is imprinted on the voltage time-line by both the bolometer itself (Johnson, or thermal noise) and the electronics, and there is also a contribution from Poisson photon noise. Inevitably the noise voltage spectral density of any bolometer shows extra power at low frequencies, above the white noise spectral density level. This additional noise contribution is called 1/f noise because the voltage spectral density tends to be

inversely proportional to frequency. The final voltage noise spectral density can then be well represented by a white noise component, with constant spectral density, plus a $1/f$ component. The $1/f$ component is defined by a parameter called the 'knee frequency' which is the frequency at which the noise voltage spectral density rises by a factor of $\sqrt{2}$ above the white noise level.

In SPIRE there are two main sources of $1/f$ type noise: a contribution from thermal drifts in the telescope/instrument (following a $1/f^2$ spectrum), and a component that is inherent in the individual bolometers and their readout electronics. The former manifests itself as a correlated signal across all the detectors in the arrays, more or less simultaneously. The same thing happens in ground-based far-IR/submillimetre observations and is due principally to variations in the Earth's atmosphere. The latter is uncorrelated from detector to detector, so each detector time-line will have its own low-frequency variation unrelated to any of the others.

Clearly $1/f$ noise is a major problem and if astronomers wish to extract the most meaningful and reliable data as possible from an observation then an understanding of $1/f$ noise and how to deal with it is essential.

There is much that can be done to alleviate the effects of uncorrelated $1/f$ noise in practical observations, the principal method being to employ 'cross-linking' together with a sophisticated mapmaking algorithm. Cross-linking is simply two (or more) co-added observations of the same field performed with different scan directions. If an observation is performed with a particular scan direction, by performing a second observation with a different scan direction then the spatial and temporal information for a specific point on the sky are mixed up. This makes it possible to distinguish between structure in a map due to sources from that due to $1/f$ noise.

In this paper we test different scan map strategies, designed to deal with $1/f$ noise, to find the most effective one compatible with the constraints provided by the operation of the *Herschel* spacecraft.

1.3 SPIRE photometer simulator

The simulator is an IDL (Interactive Data Language) coded virtual model of the photometer half of the *Herschel*-SPIRE system. It incorporates the main physical instrumental and telescope characteristics in a computationally practical and user-friendly program. Full details of the individual modules and their interaction with each other are given in Sibthorpe et al. (2004).

Briefly, the simulator acts on a suitable input sky – based either on observations taken by another instrument and scaled to the appropriate units, or generated by some other numerical simulation – for each of the three SPIRE bands. These are fed into the simulator where they are convolved with a representative beam profile and then 'observed' with the bolometer arrays. The observing mode and associated parameters for the observation are pre-defined by the user in the same way that a real observation would be planned. The astronomical power from the sky and the background radiation from the telescope and internal instrument components are all passed into a module containing a model of the individual detectors, which calculates their response to the incident radiation. This bolometer model also calculates and superimposes realistic noise on the output detector time-lines. The detector time-lines are then filtered and sampled to produce output voltage time-lines identical in form to those produced by the real instrument electronics. Additionally, a pointing time-line is generated based on the observation parameters.

The resulting data file represents a stage in the data processing pipeline after the spacecraft telemetry packets have been pre-processed and unpacked but prior to any further processing, such as mapmaking.

2 SIMULATED OBSERVATIONS

The available parameter space that can be investigated by the simulator is huge. Any characteristic of the instrument can be modified and tested using a series of simulations, and a full list is beyond the scope of this work. Here we have chosen to test the key characteristic that is likely to have a significant effect on the quality of the data obtained by the real instrument, $1/f$ noise.

2.1 Assumptions

For the purposes of this investigation we must make a number of assumptions and simplifications. Of the two sources of $1/f$ noise present in the real system only the uncorrelated component is included in these simulations. The correlated $1/f$ component is assumed to be taken care of in the pipeline processing, since thermal drifts within the telescope/instrument will be measurable and accounted for. The presence of 'dark detectors' and thermometers within each array as well as thermometers attached to various other parts of the telescope/instrument will enable this. And of course the arrays themselves will provide some measure of the correlated noise by taking a median of the detector ensemble. Additionally, thermal control of the photometer detector arrays will be implemented to some degree, which should minimize this effect in the first place.

The knee frequency of the uncorrelated $1/f$ noise imposed on the SPIRE detector time-lines is set to 100 mHz. This is a pessimistic assumption based on the instrument requirements. In reality, and based on instrument level tests, the knee frequency should be lower than this with typical values in the range 30–100 mHz.

The uncorrelated $1/f$ noise is assumed to have the same knee frequency for all detectors. Related to this, the detectors within each array are assumed to be of perfectly uniform behaviour and the bolometer yield is assumed to be 100 per cent – that is, there are no 'dead' detectors. Again, this is a simplifying assumption but one that should not significantly affect the conclusions of this work, as the variation in sensitivity from detector to detector will not be great and any difference will be flat-fielded out of the data by the pipeline processing.

It is also assumed that the reconstruction of the telescope pointing is perfect, so that the commanded observation is identical to what is actually observed. In reality there will be some error in the actual versus commanded pointing of the telescope, and some error in reconstructing the actual pointing by virtue of imperfect star tracker information. Both of these errors can be simulated but they are expected to be small enough in the real system that they can be ignored in this work.

Finally, no glitches (cosmic ray hits, electronics errors) are present in the simulated data, so no additional data processing is required between the simulator output data and creating the maps for analysis.

All simulations were performed with a scanning speed of $30\ \mathrm{arcsec\ s^{-1}}$, which will be the default for most scan map observations.

2.2 Input maps

Any input map can be used for an investigation of this type but we choose to perform our simulated observations on something resembling a realistic piece of sky. For this paper we use an input

Table 2. Approximate confusion limits for the GALICS model in the three SPIRE wavebands.

Band (μm)	Confusion limit (mJy)
250	21.7
350	16.7
500	12.2

Table 3. Table showing how the different map repeats were combined to produce the final observations. These combinations ensure that the total effective integration times for the different strategies are almost identical (4.00, 4.06, 3.99, respectively, using the same units as in Table 1).

Type	Combination	Array rotated?
Delayed	2×long-1 + 2×long-2	Yes
Concatenated	2×long + 4×short	No
Concatenated	3×diagonal-1 + 3×diagonal-2	No

map representative of a deep, extragalactic field, containing only point-like sources.

We chose an input catalogue of galaxies extracted from the GALICS numerical simulations (Hatton et al. 2003). Initial tests of the GALICS model show it to be quite successful in reproducing many of the observed galaxy population properties in the optical and IR (Hatton et al. 2003). However, its success at reproducing far-IR/submillimetre properties is yet to be tested (something that *Herschel* will undoubtedly do). Despite not necessarily being a perfectly realistic representation of the far-IR Universe, GALICS has the advantage it extends to flux levels far lower than the *Herschel*/SPIRE confusion limit (simplistically assumed here to be 1 source per 40 beam areas) and naturally includes realistic clustering – something that will affect the detection of faint sources in particular.

For this work a catalogue of some 58 590 sources was extracted from the GALICS data base, covering an area of 1 deg^2 and with a flux limit of 0.2 mJy in the SPIRE 250-μm band. For comparison Table 2 gives the approximate confusion limit of the GALICS model in the three SPIRE wavebands.

2.3 Cross-linking

The key to dealing with $1/f$ noise is to perform cross-linked observations, as mentioned in Section 1.2. In Section 1.1 we explained how the orientation of the *Herschel* FOV will change with time, so this could be used to implement cross-linked observations, simply by delaying half of the map repeats until a later date. However, this is not an option for fields close to the ecliptic plane, where only 180° rotation is possible. There are two possibilities for tackling this issue: first, the very limited degree of flexibility in the roll angle of *Herschel* along its boresight could be utilized to provide a very small cross-linking angle between two map repeats; secondly, rather than using the same scan strategy for every map repeat and using the rotation of the array to provide the cross-linking angle, instead the array can be scanned at a complementary angle, e.g. one long direction scan and one (or two) short direction scan(s), providing a nearly orthogonal cross-linked observation. This second strategy has several advantages over the first; principally it allows cross-linking to be performed in one observation, without the need to wait until a suitable window of time opens up at some future date. It also produces nearly square maps in a natural way, which cannot be done easily with, say, a 45° cross-linked observation.

For this work, three types of cross-linked observation are performed: a long direction scanning strategy, with cross-linking provided by delaying half of the map repeats to allow a rotation of the scan pattern ('delayed'); and two versions of the naturally cross-linked observation, with either long plus short direction scanning, or two complementary diagonal direction scans ('concatenated'). For these latter strategies, the SPIRE array remains at the same angle on the sky at all times, since the second set of map repeats is performed immediately after the first set; it is the scan direction that is changed in between.

Table 3 summarizes these various scan map strategies. To ensure that roughly half of the effective integration time for any strategy is spent during each of the complementary sets of map repeats, two short direction scans are performed for every one long direction scan. Likewise, three diagonal scans are performed for every two long scans, or four short scans. This way all the strategies end up with very similar total effective integration times, and so can be compared easily. The remaining small differences between the total effective integration times can be calibrated out by scaling the results appropriately.

For the delayed-type observations, various cross-linking angles were investigated to determine how this parameter affects the final data quality. The angles investigated were 0°, 5°, 10°, 20°, 40° and 90°. For the concatenated cross-linked observations the cross-linking angle is fixed at 84°.8 (or 95°.2, depending on how it is defined), which is a consequence of keeping the array orientation fixed on the sky and using the complimentary scan angles to provide the cross-linked data; this can be easily seen by taking the difference between the two complementary scan angles in Table 1.

Two sets of simulations were performed, one with white noise only and the other with $1/f$ noise switched on with a knee frequency of 100 mHz, as described in Section 2.1. All other observing parameters were identical for the two sets of simulations.

In Fig. 4 we show examples of integration time maps, or coverage maps, for the two concatenated-type observations of 30 arcmin^2. The 250-μm SPIRE array, shown in these examples, always lies horizontally, leading to different orientations of the final map coverage. These figures are only illustrative and so they do not represent the orientation or coverage of the actual simulations used for the rest of this paper, which were designed to cover the full 1 deg^2 input maps as far as possible. Data taken while the telescope was accelerating or decelerating are not included in these figures so only data taken with the array travelling at full scanning speed are shown.

3 MAPMAKING

The simulated observations come in the form of detector voltage time-lines and a pointing time-line that must be calibrated and re-gridded to form a map representing the sky flux density. The simplest form of map is created by simply averaging the data points falling within a given map pixel. This is called naive mapmaking. However, naive mapmaking makes no attempt to deal with $1/f$ noise and the result is the appearance of stripes in the map and the swamping of faint sources by excess noise (see Fig. 5a).

There are more sophisticated mapmaking algorithms, and to best exploit cross-linked observations we must turn to these methods. As part of the SPIRE pipeline development programme it was deemed necessary to provide a mapmaking algorithm within the pipeline processing suite. A selection procedure was initiated to find the most appropriate code, from which a SPIRE-specific version
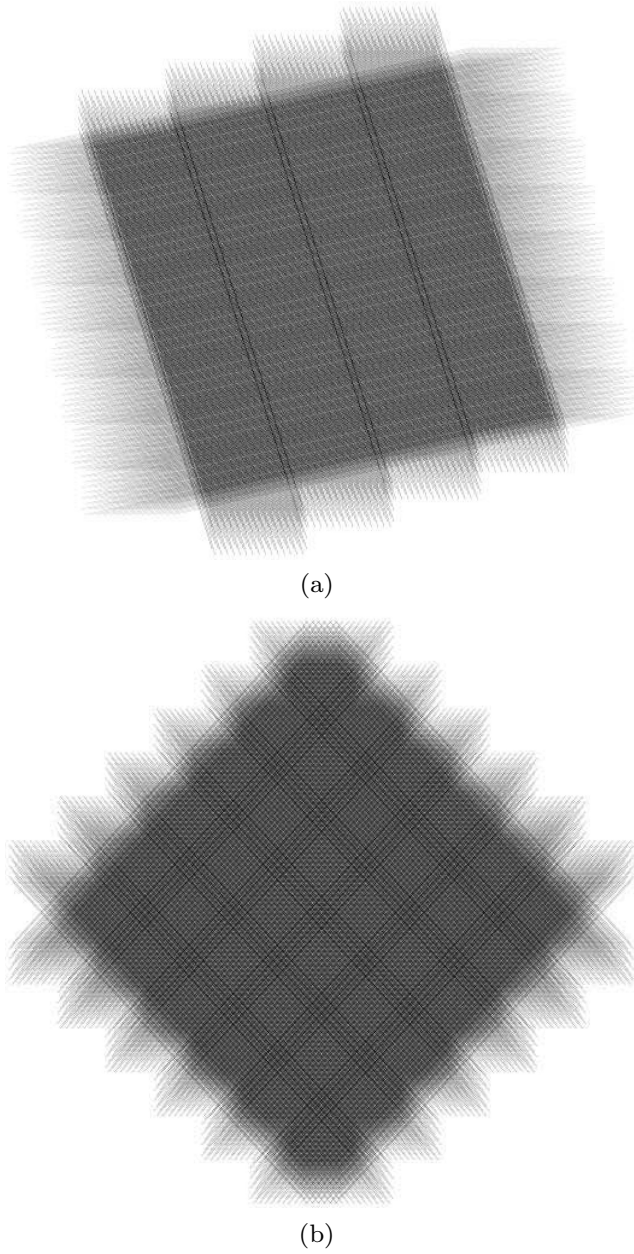
(a)

(b)

Figure 4. Examples of integration time maps: (a) one long plus two short observations; (b) two complimentary diagonal observations. The SPIRE array (250 μm in this case) is horizontal in all these example observations.

could be adapted. Having tested six different codes, including naive mapmaking, the maximum likelihood code MADmap (Cantalupo 2002) was selected.

Given a pixelized sky s_p, the time-ordered series d_t can be written as

$$d_t = \mathbf{A}_{tp}s_p + n_t, \qquad (t = 1, \ldots, n \text{ and } p = 1, \ldots, m), \qquad (1)$$

where $\mathbf{A}_{tp}$ is the pointing matrix and n_t the noise. Hence, a map-making algorithm can be seen as a solver for s in this set of linear equations. Lossless methods have been developed (Tegmark 1997). In particular, it can be shown that the maximum likelihood estimate of the map $\hat{s}$ is a solution of the equation

$$(\mathbf{A}^T\mathbf{N}^{-1}\mathbf{A})\hat{s} = \mathbf{A}^T\mathbf{N}^{-1}d, \qquad (2)$$

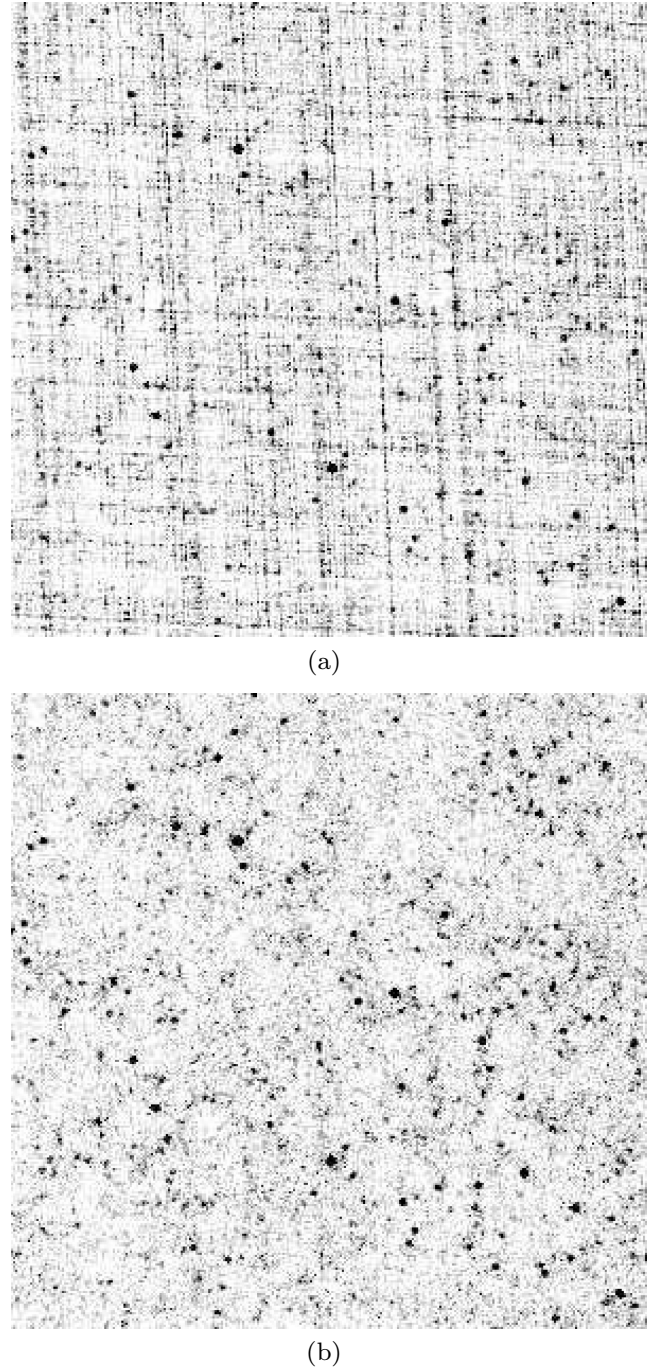where $\mathbf{N} = \langle nn^T \rangle$ is the time–time noise covariance matrix.

(a)

(b)

Figure 5. Comparison images showing a $30 \times 30\,\text{arcmin}^2$ section of a map both before and after treatment using MADmap. This example comes from the concatenated diagonal direction scanning strategy. The contrast levels are the same for both images.

Briefly speaking, the MADmap algorithm (Cantalupo 2002) makes things computationally tractable by assuming that the $n \times n$ matrix $\mathbf{N}^{-1}$ is piecewise Toeplitz band diagonal, so that $\mathbf{N}^{-1}$ acts like a set of convolutions with band-limited kernels, which are straightforward operations in Fourier space. The inversion of the $m \times m$ matrix $\mathbf{A}^T\mathbf{N}^{-1}\mathbf{A}$ is done by the pre-conditioned conjugate gradient method.

Fig. 5 shows the dramatic improvement in the quality of the maps when MADmap is employed to remove the effects of $1/f$ noise.

Fig. 5(b) has none of the striping, characteristic of long time-scale drifts, as seen in Fig. 5(a) and fainter sources are visible.

4 RESULTS

Once the simulated observations have been processed by MADmap, we subtract the clean input map (having also been convolved with the SPIRE beam pattern) from the noisy output map to leave a map that contains what should be just noise residuals. We then select a 42×42arcmin^2 region inside the 1-deg^2 map area to ensure that we only consider what is covered by all the map repeats making up the observation (for some observations the corners of the map are missed by the rotated map repeats). To get an indication of the noise level in the map we measure the standard deviation of the pixel values within this region. This is the simplest possible metric, providing the most transparent analysis of the different scan map strategies.

The results of this analysis are presented in Fig. 6. Fig. 6(a) is the standard deviation of the residuals; Fig. 6(b) is the square of this quantity, to give an indication of how long it would take to integrate down to a particular sensitivity with each strategy; Fig. 6(b) is the inverse of Fig. 6(c), giving an indication of how quickly each strategy could map a given area to a given sensitivity, i.e. the mapping speed.

5 DISCUSSION

It is very encouraging to see that all three possible scan map strategies result in similar data quality, as long as the delayed-type observations are scheduled to have at least 20° of rotation between the two sets of map repeats. So there is no clear advantage of performing one strategy over another, at first sight.

However, data quality is not the only consideration here; the practicality of performing a given scan map strategy is almost as important as data quality. Therefore we shall now discuss the relative benefits of the three strategies by considering their practical implications to the operation of SPIRE.

5.1 Long direction scanning with delayed scheduling

Long direction scanning is an excellent way for SPIRE to map a given area because the greater length of the array provides more detectors along the scan direction. In the event of a few detectors failing there is greater redundancy along a scan and so less probability of gaps in the final map. If SPIRE were not required to perform cross-linked observations then long direction scanning would be the sensible choice.

However, the clear advantage of cross-linking to data quality makes this strategy less desirable. In order to achieve the necessary rotation of the scan pattern there is no choice but to schedule a delay between two sets of map repeats. As noted in Section 1.1, the orbital geometry of *Herschel* restricts the scheduling possibilities for this type of strategy, with the ecliptic plane proving to be particularly troublesome.

Additionally, if a rotation of much less than 90° were implemented then the final map coverage would be octagonal, rather than square, with the corners of each set of map repeats falling outside of the jointly covered area. To observe the desired field fully with both sets of map repeats a much larger area would need to be mapped in each case, requiring much greater observing times to achieve the same goal.
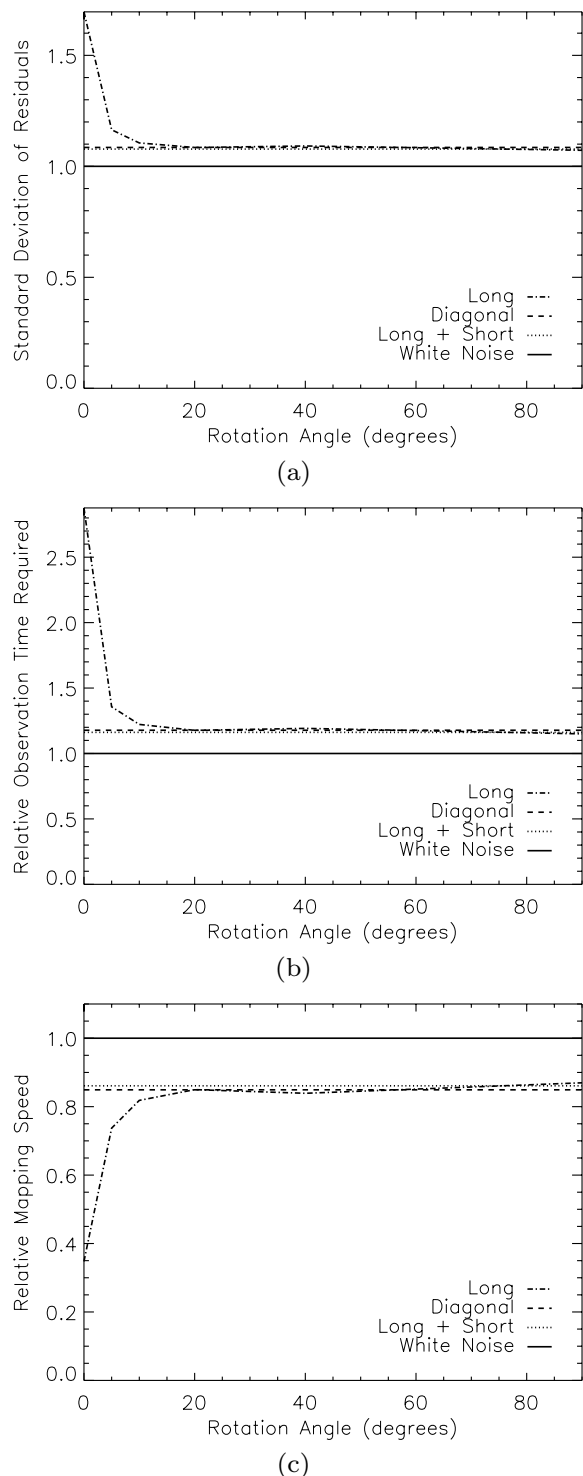

(a)

(b)

(c)

Figure 6. Three plots comparing the performance of the different scanning strategies. The rotation angle refers only to the delayed-type observations; the concatenated-type observations have a fixed rotation angle of 84°.8. All quantities have been normalized to the white noise level, to aid comparison. The slight differences in effective integration time have also been normalized out of these plots. Note how all the lines, apart from the white noise case, are almost indistinguishable. This shows that all three strategies are equally good at suppressing $1/f$ noise, as long as the rotation angle between the co-added observations is at least 20°. Although white noise performance is not reached with any strategy, the mapping speed is only ~15 per cent slower when cross-linking and MADmap are employed, compared to ~65 per cent slower if $1/f$ noise is left untreated.

5.2 Concatenated scanning strategy

Both sets of concatenated strategies neatly avoid all of the scheduling problems experienced by the delayed strategy. Because the SPIRE array can be scanned along complementary scan directions there is no need for a scheduling delay between the two sets of map repeats required for cross-linking. The obvious advantage of this type of strategy is that cross-linked observations can be performed in a single observing session, with the added advantage that the entire sky is accessible too. Although any particular point on the ecliptic plane will only be visible to *Herschel* for two short periods a year, it will be possible to perform cross-linked observations during these times. The issue of failing detectors is also alleviated somewhat by employing cross-linking because the complementary scan direction will fill in any gaps left in the map by the first scan direction.

The choice of long + short (one long + two short) or diagonal + diagonal makes little difference to the data quality or the length of time it takes to perform a cross-linked observation to a given sensitivity, so the distinction between these two options is less clear. However, in practical terms the diagonal + diagonal strategy just wins out over long + short because the two complementary map repeats required to produce a cross-linked observation have essentially the same characteristics. The effective integration time of a long direction scan map is roughly twice that of a short direction scan map, so combining them together in such a way as to equalize the sensitivity of each part of the cross-linked observation requires two short scans for every one long scan. The two complementary diagonal scan maps are, of course, just mirror images of each other and so naturally add together in pairs.

One final point in favour of the diagonal + diagonal strategy is the lower effective integration time of a single cross-linked observation, compared with long + short. A single diagonal + diagonal observation has only two-thirds of the effective integration time of a single long + short observation meaning that it would take less time to produce very large, shallow maps using this method. SPIRE is sufficiently sensitive that even a single long + short observation will be too deep for the very largest of the planned surveys, so diagonal + diagonal is the preferred choice since it will map large areas faster than long + short.

6 CONCLUSIONS

Concatenated-type observations have many clear advantages over delayed-type observations. Of the two possible concatenated types the Diagonal + diagonal method has the advantage of practicality and elegance and so is the scan strategy recommended in this paper. This strategy is now being implemented as the default for all

SPIRE scan map observations so that the archival data quality can be assured.

Of course, cross-linked observations are not in themselves sufficient to ensure good data quality. Maximum likelihood mapmaking algorithms must be employed to obtain the most out of the information encoded in the cross-linked data. Therefore, the SPIRE data processing pipeline suite will contain such an algorithm, tailored to SPIRE, so that every SPIRE photometer user can benefit from the optimized scan map observing mode.

This work also highlights the usefulness of developing instrument simulators, such as the SPIRE photometer simulator, when preparing for expensive missions like *Herschel*. The simulator is proving invaluable for helping to understand many aspects of the operation of SPIRE and will continue to be used to further optimize observing modes, and to help plan observations, as new information about the performance of SPIRE comes to light.

REFERENCES

Cantalupo C., 2002, MADmap: A Fast Parallel Maximum Likelihood CMB Map Making Code, http://crd.lbl.gov/cmc/MADmap/doc/
Dowell C. D. et al., 2001, BAAS, 33, 792
Griffin M. et al., 2006, in Mather J. C., MacEwen H. A., de Graauw M. W. M., eds, Proc. SPIE, Vol. 6265, Space Telescopes and Instrumentation I: Optical, Infrared, and Millimeter. SPIE, Bellingham, p. 62650A
Hatton S., Devriendt J. E. G., Ninin S., Bouchet F. R., Guiderdoni B., Vibert D., 2003, MNRAS, 343, 75
Holland W. et al., 2006, in Zmuidzinas J., Holland W. S., Withington S. D., William D., eds, Proc. SPIE, Vol. 6275, Millimeter and Submillimeter Detectors and Instrumentation for Astronomy III. SPIE, Bellingham, p. 62751E
Pilbratt G. L., 2005, in Wilson A., ed., The Dusty and Molecular Universe: A Prelude to *Herschel* and ALMA *Herschel* Mission: Status and Observing Opportunities. ESA, Noordwijk, p. 3
Sibthorpe B., Woodcraft A. L., Griffin M. J., Watkin S. L., 2004, in Mather J. C., ed., Proc. SPIE, Vol. 5487, Optical, Infrared, and Millimeter Space Telescopes. SPIE, Bellingham, p. 491
Sibthorpe B., Waskett T. J., Griffin M. J., 2006, in Silva D. R., Doxsey R. E., eds, Proc. SPIE, Vol. 6270, Observatory Operations: Strategies, Processes, and Systems. SPIE, Bellingham, p. 627019
Tegmark M., 1997, ApJ, 480, L87

This paper has been typeset from a TeX/LaTeX file prepared by the author.