Two New Approaches to Spatial Interpolation with Inherent Sidelobe Suppression for Imaging Riometers

Martin Grill, Andrew Senior, Farideh Honary

All at Department of Communication Systems, Lancaster University, UK

Abstract

Absorption images as obtained by imaging riometers such as IRIS are usually created by interpolating between absorption values for individual beams. For IRIS, the locations of the beam centres serve as grid points for subsequent linear interpolation. Although generally producing good results, the fact that the actual shape of the imaging beams is not considered, potentially introduces errors and can lead to misinterpretations.

In this paper, two alternative interpolation methods are introduced. Method A is based on measuring the similarity between simulated reception of individual point sources and actually received data. Method B uses a mathematical model of the sky brightness distribution parametrised by the received data.

All interpolation methods are applied to power data, as opposed to absorption data, in order to avoid any errors that might be introduced by intermediate processing steps, especially QDC (quiet-day curve) generation.

We apply all methods to synthetically generated test data as well as to three exemplary real datasets which are also compared to a calculated sky brightness distribution obtained from a skymap.

1   Introduction

Riometers (Relative Ionospheric Opacity Meters) have been around for decades [9]. They measure absorption of cosmic background noise, caused by energetic particles ionising layers of the Earth's ionosphere. The first riometers were widebeam riometers with a large field of view in the order of $60°$ around zenith. The search for smaller scale structures led to the development of imaging riometers which divide a similar field of view into smaller areas by forming multiple beams, usually employing phased array antennas. Current imaging riometers have between 49 beams [2] and 256 beams [13], an even higher resolution imaging riometer based on a Mills Cross beamforming technique [12] currently being under development [14].

Along with the imaging capabilities of such riometers came the necessity to spatially interpolate between the measurements in order to form a real image of the observed area. Time series of these images can then be used for further studies like analysing the motion of absorption patches [10] and serve as a source for more advanced diagrams such as keograms, movies and virtual beams. These images can also be directly compared to images from other imaging instruments, for example optical cameras [5].

Currently, riometer images are usually created by interpolating between absorption values for individual beams. For IRIS, the Imaging Riometer for Ionospheric Studies [2], the locations of the beam centres serve as grid points for subsequent linear interpolation. This technique generally produces good results. However, the fact that the actual shape of the imaging beams is not considered, potentially introduces errors and can lead to misinterpretations. In particular, any given imaging beam receives signals not from one direction but from a range of directions around the beam centre, depending on the beamwidth, which itself is inversely proportional to the aperture of the receiving antenna. Also, a not always negligible fraction of signal is received from sidelobes that point in a significantly different direction from the main beam.

This paper suggests different interpolation methods that take the shape of the receiving beams into account. In doing so, they have the potential to compensate for the effects of sidelobes, to overcome the spatial constraints of linear interpolation (thus extending the field of view) and to uncover features that may not show up in traditionally interpolated images.

2   Prerequisites

In this section we introduce some general facts that are used for the observations in the following sections. We introduce the dataset that we will be using throughout this paper, highlight the role that obliquity factors play for the observations and introduce the coordinate system and the FLATM projection method that we use in this paper.

2.1   Power data, role of obliquity factors

The interpolation algorithms as discussed in this paper all deal with interpolation of power data on the positive hemisphere seen by the receiving instrument. The data is not interpreted in any way prior to processing. In particular, no assumptions are implied as to the media that the incoming signals traversed prior to reception, i.e. no correction factors (in this case known as obliquity factors) for the varying observed thickness of the absorbing layer etc. are applied to the data. In order to compare actual signals to theoretical signals based on convolution of beampatterns and skymap (see section 6), obliquity factors need to be taken into account as soon as there is an absorbing layer of electrons present. This layer appears thicker with decreasing elevation angles. See figure 1: The apparent thickness d' of the absorption layer decreases with
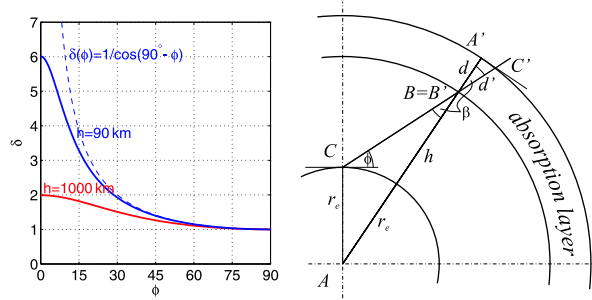

Figure 1: Obliquity factor δ correcting for apparent thickness of absorption layer (drawing not to scale)

observation elevation angle ϕ. The law of sines for $\triangle ABC$ allows us to derive angle β:

$$sin\beta(\phi) = \frac{sin(\phi + 90°)}{r_e + h} \cdot r_e \tag{1}$$

from which we can derive the obliquity factor $\delta(\phi)$ simply by looking at $\triangle B'C'A'$ (as long as $d \ll (r_e + h)$):

$$\delta(\phi) = \frac{d'(\phi)}{d} = \frac{1}{cos\beta(\phi)} \tag{2}$$

The left panel in figure 1 shows a plot of $\delta(\phi)$ for elevation angles $0° \leq \phi \leq 90°$ and $h = 90$km (blue line). For comparison reasons, $\delta(\phi)$ is also shown for an (unrealistic) height $h = 1000$km (red line). This is essentially a 'lower amplitude' version of $\delta(\phi)$ for $h = 90$km. The dashed line is $\delta(\phi) = 1/cos(90° - \phi)$ as used in [7]. This is an approximation to equation 2 that works well for elevation angles $\phi > 30°$.

Obliquity factors will also have to be taken into account when it comes to deriving absorption from the input data, but in this paper we are solely dealing with (spatial) interpolation of the underlying raw received data. Existing algorithms can then be used to derive quiet-day curves (QDCs) and absorption data for arbitrary directions ('virtual beams') within the usable working area of the given interpolation algorithm. Note, however, that equation 2 only works well if there is actually a well-defined absorption layer present. For quiet-day absorption, equation 2 generally exaggerates the expected results. This is because quiet-day absorption happens over a large range of heights, not within one narrow layer. In this paper, we only ever use obliquity factors for qualitatively comparing data derived from a skymap to actual measurements. To qualitatively compensate for quiet-day absorbtion we use the red curve in figure 1 when

we plot the sky brightness distribution as derived from a skymap (see section 6). This curve was empirically found to produce good results.

Note that this use of obliquity factors is different from the traditional use where obliquity factors are applied to received power values from each beam, either ignoring the beam shape or including the effects of the beampattern in the 'effective' obliquity factor [7]. The obliquity factor in equation 2 does not relate to beams but simply to a certain viewing direction.

2.2 Time period

The time period chosen for the observations in this paper is 23rd September 1999. This day (as observed by the imaging riometer IRIS [2]) has quiet times, absorption events and 'negative absorption events' i.e. peaks in power due to interference (solar radio emissions). See figure 2 (obtained from the Multi


Figure 2: IRIS widebeam absorption during 1999-09-23

Instrument Analysis Toolkit MIA [11]). It is therefore an excellent dataset for comparing the performance of interpolation algorithms in various circumstances.

2.3 FLATM projection

Antenna radiation patterns and sky brightness distributions are most easily described in some sort of spherical coordinate system. In this paper we shall use a spherical coordinate system made up from azimuth angle θ (in the XY plane, counterclockwise, starting from the positive x-axis) and elevation angle ϕ (between XY plane and direction of interest, positive ϕ denotes positive z). For visualisation purposes and further studies, this spherical dataset often needs to be mapped onto a flat two-dimensional surface. For IRIS absorption images, we commonly use a projection that we will henceforth call the 'FLATM' (for 'flat metres') projection. This is a two-stage projection that first calculates the intersection between a ray in the given direction and the ionosphere (at a given height, 90km by default) and then maps the result onto a flat surface. The whole process is shown in figure 3. This projection offers relatively little distortion around the zenith, with distortion increasing with lower elevation angles. As the surface area of the ionosphere is much greater than the area usually covered by any given instrument, the FLATM projection gives a good (i.e. relatively undistorted) representation of the signal distribution at the height of the ionosphere (or any other sphere centred on the Earth's centre for that respect).

Note that from the interpolation algorithms presented in this paper, only the original IRIS interpolation algorithm inherently uses the FLATM projection. The other algorithms interpolate in spherical coordinates, only the results shown in this paper are mapped back to FLATM projections for presentation and comparison purposes.
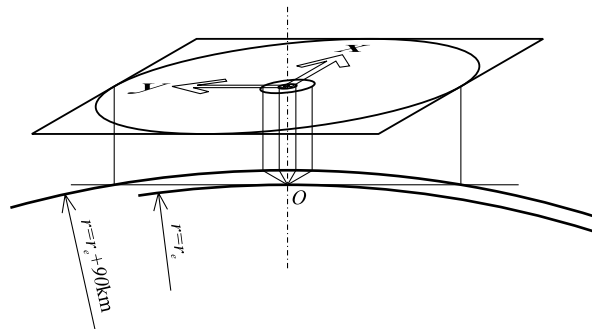
Figure 3: FLATM projection as used for IRIS image data. Drawn to scale. Instrument is situated at O.

3 Original IRIS interpolation algorithm

IRIS power/absorption images as created by the Multi-Instrument Analysis Toolkit (MIA) [11] use the FLATM projection to map the location of the 49 beam centres onto a flat two-dimensional grid[1]. MIA assumes that the recorded power values originate from the respective beam centre (the direction of maximum gain, also referred to as beam axis or boresight) and then uses linear interpolation to fill the space between the beam centres. An example of this can be seen in figure 4 (left panel), together with the triangles that
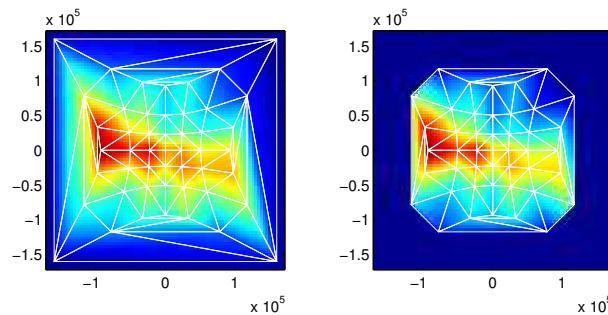

Figure 4: 'Traditional' IRIS Image Interpolation, distances are in m. Left panel shows Delaunay triangulation for all 49 beams, right panel shows Delaunay triangulation for only the 'good' beams.

are used internally by the two-dimensional linear interpolation algorithm. The vertices of the triangles represent the 49 beam centres. As the four corner beams (1, 7, 43, 49) have significant sidelobes [1] and the assumption of all power being concentrated at the beam centre does not even aproximately hold, these beams are usually ignored leading to an interpolated image as depicted in figure 4 (right panel). Therefore the useable working area for this algorithm is defined by the convex hull of the beam centres projected onto a flat grid by the FLATM projection, and we will overlay a white bounding box in subsequent images generated by other algorithms (which are not inherently limited to this working area) to simplify visual comparison of the results.

---

[1]Recent versions of MIA now use a default grid based on geographic latitude and longitude for interpolation, as geographic coordinates are more universally useful and accepted. This does not change the basic interpolation algorithm, however. We will stick with the FLATM projection in this paper, because it does not introduce any asymmetric distortion effects and is independent of instrument location.

4   Method A: Correlation method

Given the radiation patterns of all beams, we can (per definition) calculate the response of each beam to a noise source in any given direction. When calculating theoretical QDCs, this is done for all directions simultaneously, yielding the total response of each beam to the given sky brightness distribution as defined by a sky map. In the correlation method, a different approach is suggested:

First, we agree on a grid of 'directions of interest.' This can be any arbitrary list of I direction vectors $\vec{d_i} = [\theta_i, \phi_i]$, they do not have to form a regularly spaced grid.

We then treat each direction separately: For each direction $\vec{d_i}$ we calculate the response of all N beams to a power source of a fixed intensity in that direction. This gives us N received power values p_n, one for each beam.

We then collectively compare this theoretical response of all beams $\{p_1...p_N\}$ to the actual recording from these beams $\{b_1...b_N\}$. The correlation as represented by the normalised cross correlation coefficient

$$c_i = \frac{1}{\sum_{n=1}^{N} p_n} \sum_{n=1}^{N} p_n b_n \tag{3}$$

between the two sets of values turns out to be a good estimate of how much the recording was influenced by a noise source in the particular direction in question ($\vec{d_i}$).

We repeat these steps for each direction $\vec{d_i}$ and in doing so build up an image of the sky, represented by the values c_i.

It is interesting to note that equation 3 resembles the equation used for finite impulse response (FIR) filters, suggesting that this algorithm can be looked upon as some sort of matched directional filtering, the p_n being the expected responses.

4.1   Implementation note

Provided we know all potential directions of interest $\vec{d_i}$ from the outset, the theoretical beam responses p_n can be calculated in advance. For any given set of input values b_n we then only need to calculate the cross-correlation coefficients c_i according to equation 3.

5   Method B: Parametrised model

We will first describe the general approach, this is not specific to using spherical harmonics, but is indeed valid for any kind of model that can be parametrised with a set number of parameters less than or equaling the number of concurrently available datapoints. We will then proceed to describe two approaches based on spherical harmonics.

5.1   (General) Curve Fitting Method

Rather than starting with an unknown power distribution (which is of course what we will be doing later on), let us assume for a moment that we know the spatial power distribution P across the visible hemisphere, i.e.

$$P(\theta, \phi) = \text{known} \tag{4}$$

We also know the radiation pattern for each of the given instrument's N beams for all possible directions (θ, ϕ):

$$B_{1...N}(\theta, \phi) = \text{known} \tag{5}$$

Given the power distribution and the radiation patterns, we can now calculate the power response (the power received) $p_{n,simul}$ for each beam n. This is the same method that can also be used to derive theoretical quiet-day curves (see for example [16]):

$$p_{n,simul} = k \cdot \int_{\theta,\phi} P(\theta,\phi) B_n(\theta,\phi) cos\phi d\theta d\phi \tag{6}$$

k is a constant that can be used for calibration purposes. The N values $p_{n,simul}$ directly correspond to the received power as measured by the receivers.

If we now find a way of representing the power distribution P in equation 4 by means of $M \leq N$ parameters instead of an infinite number of discrete values, we can work our way backwards from the actual received power values and derive (an approximation of) the original power distribution P, denoted P_{model}.

Let us assume that P_{model} is a linear combination of M functions of (θ,ϕ), weighted by γ_m, i.e. a function of the direction as specified by (θ,ϕ) and of M parameters $\gamma_1...\gamma_M$ as follows:

$$P_{model}(\theta,\phi,\gamma_1...\gamma_M) = \gamma_1 \cdot f_1(\theta,\phi) + \gamma_2 \cdot f_2(\theta,\phi) + ... + \gamma_M \cdot f_M(\theta,\phi) \tag{7}$$

In order to determine the parameters $\gamma_1...\gamma_M$, we make use of equation 6, replacing the simulated results $p_{n,simul}$ with the actual measurement results p_n $(n = 1...N)$ from the N beams and the 'known' sky brightness distribution P (equation 4) with the modelled brightness distribution P_{model} (equation 7):

$$p_n = k \cdot \int_{\theta,\phi} B_n(\theta,\phi) P_{model}(\theta,\phi,\gamma_1...\gamma_M) cos\phi d\theta d\phi \tag{8}$$

We expand equation 8 using the definition of our model in equation 7:

$$p_n = k \cdot \int_{\theta,\phi} B_n(\theta,\phi) \cdot [\gamma_1 \cdot f_1(\theta,\phi) + \gamma_2 \cdot f_2(\theta,\phi) + ... + \gamma_M \cdot f_M(\theta,\phi)] cos\phi d\theta d\phi \tag{9}$$

Rearranging:

$$
\begin{aligned}
p_n \quad = \quad & \gamma_1 \cdot k \cdot \int_{\theta,\phi} B_n(\theta,\phi) f_1(\theta,\phi) cos\phi d\theta d\phi \\
+ \quad & \gamma_2 \cdot k \cdot \int_{\theta,\phi} B_n(\theta,\phi) f_2(\theta,\phi) cos\phi d\theta d\phi \\
+ \quad & ... \\
+ \quad & \gamma_M \cdot k \cdot \int_{\theta,\phi} B_n(\theta,\phi) f_M(\theta,\phi) cos\phi d\theta d\phi
\end{aligned}
\tag{10}
$$

Combining all constant terms into constants c:

$$p_n = c_{n,1}\gamma_1 + c_{n,2}\gamma_2 + ... + c_{n,M}\gamma_M \tag{11}$$

which can be written as one matrix equation for all N beams:

$$P = C \times \Gamma. \tag{12}$$

In other words, we get one linear equation with M unknowns $\{\gamma_1...\gamma_M\}$ for each of the N beampatterns. This set of equations can be solved (possibly in a least-squares sense for $N \neq M$) and therefore the unknown model parameters $\{\gamma_1...\gamma_M\}$ can be determined. Once these parameters are known, the (model) sky brightness can be calculated in any arbitrary direction (θ,ϕ) by using equation 7.

5.2   Implementation employing Spherical Harmonics (B1)

Spherical harmonics [8] can be used to describe intensity distributions around a sphere. See, for example, [17] for a geophysical application. Spherical harmonics are useful because they form an orthogonal basis for functions on the sphere in a manner analogous to sines and cosines on the interval $[0, 2\pi]$. We are, of course, only using a finite number of spherical harmonics to approximate these functions, as is often done for Fourier sine/cosine series.

Spherical harmonics are solutions to Laplace's equation in spherical coordinates [19]. Without going into any further mathematical details, we stick to the practical description of the properties of spherical harmonics as found in [15]:

"The spherical harmonic $Y_{lm(\theta,\phi)}$, $-l \leq m \leq l$, is a function of the two coordinates θ, ϕ on the surface of a sphere. The spherical harmonics are orthogonal for different l and m, and they are normalised so that their integrated square over the sphere is unity."

The degree of the spherical harmonic in question is depicted by l, the order by m. The spherical harmonics are related to *associated Legendre polynomials* P_l^m by the following equation ([15]):

$$Y_{lm}(\theta, \phi) = K_l^m P_l^m(cos\theta)e^{im\phi} \tag{13}$$

Since we are only dealing with a real numbers, we will use real spherical harmonics defined as follows:[2]

$$Y_{lm}(\theta, \phi) = \begin{cases} \sqrt{2}K_l^m cos(m\phi)P_l^m(cos\theta), & m > 0 \\ \sqrt{2}K_l^m sin(-m\phi)P_l^{-m}(cos\theta), & m < 0 \\ K_l^0 P_l^0(cos\theta), & m = 0 \end{cases} \tag{14}$$

In both equations 13 and 14, the scaling factor K is defined as:

$$K_l^m = \sqrt{\frac{(2l+1)}{4\pi}\frac{(l-|m|)!}{(l+|m|)!}} \tag{15}$$

Note that this scaling factor (equation 15) is constant for any given l, m. For our purposes we can therefore ignore this part, the weighting coefficients will automatically adjust themselves to accommodate for this missing factor.

For modelling the sky brightness distribution according to equation 7, we need a linear combination of a number of spherical harmonics equalling the number of available beam power readings N (49 in case of the IRIS observations used in this paper). We establish a unique order between the spherical harmonics starting with $f_1 = Y_{l=0;m=0}$ and then moving on to higher degrees, cycling through all possible orders $m = -l, ..., m = l$ for each degree, until we arrive at $f_{49} = Y_{l=6;m=6}$.

5.3   Implementation using adjusted spherical harmonics (B2)

Ordinary spherical harmonics as used in 5.2 have the disadvantage of modelling the brightness distribution over the surface of a complete sphere, resulting in relatively low-resolution at any given region of interest. Especially, in our case, a riometer will only ever 'see' at most a hemispherical subsection of the whole sky. It would be sensible to choose a basis which represents the power distribution over the hemisphere alone. One possibility is capped spherical harmonics [6]. However, capped spherical harmonics can be complex to implement and can result in a range of computational problems [18]. Therefore, in this paper, we use a simple approximation to spherical capped harmonics called 'adjusted spherical harmonics' as proposed by DeSantis [17]. It is based on a coordinate transformation that 'adjusts' (compresses) the elevation angles ϕ in the original definition of the spherical harmonic to those of interest (in our case the visible hemisphere). By employing these adjusted spherical harmonics, we effectively double the resolution in the hemisphere of interest. We still use the same logical order of spherical harmonics as described in section 5.2 above.

5.4   Implementation notes

As for interpolation method A (section 4), some of the values needed for evaluating equation 7 can be pre-calculated. In case of method B we do not need to know about the required directions $\vec{d_i}$ from the outset, we only need to define a grid for numerically doing the integration in equation 10, this grid does not directly relate to the possible output directions, it is only used for the integration (summation) process, therefore limiting the accuracy of the results if chosen too wide-meshed. Once we have defined this 'internal' grid, the N-by-N matrix C (made up of the values $c_{1...N,1...N}$ in equation 11 can be computed. For any given set

---

[2]Robin Green, "Spherical Harmonic Lighting: The Gritty Details", http://www.research.scea.com/gdc2003/spherical-harmonic-lighting.html

of input values b_n we then only need to solve the system of linear equations (equation 11 for all $n = 1...N$) and use equation 7 to calculate the interpolated power values for arbitrary directions (θ, ϕ). Note that this method has 'self-interpolating' properties, meaning that once we have calculated the coefficients c, we are not limited to certain predefined directions. Instead, equation 7 will directly give results for all desired directions (θ, ϕ).

6 Results

In this section, we will present the results from the two interpolation methods for three different times during 23rd September 1999 as well as for two test-cases using synthetically generated data. We will compare these results to the result obtained by the original MIA interpolation method. The results for the synthetically generated brightness distributions are compared to these very distributions. For the 'real' datasets, this is of course impossible since we do not know the actual sky brightness distribution in these cases. We will compare the 'real' datasets to the expected sky brightness distribution as derived from a skymap [4], noting that the real distribution at the time might vary greatly from this theoretical distribution, especially during active times.

In the following figures, each panel is the FLATM projection (see section 2.3) of the respective results, covering a square area of 340×340km. For figures containing 'real' data, we also show in each panel the locations of the two major radio stars [3] in this area of the sky, Cassiopeia A (white circle) and Cygnus A (white X).
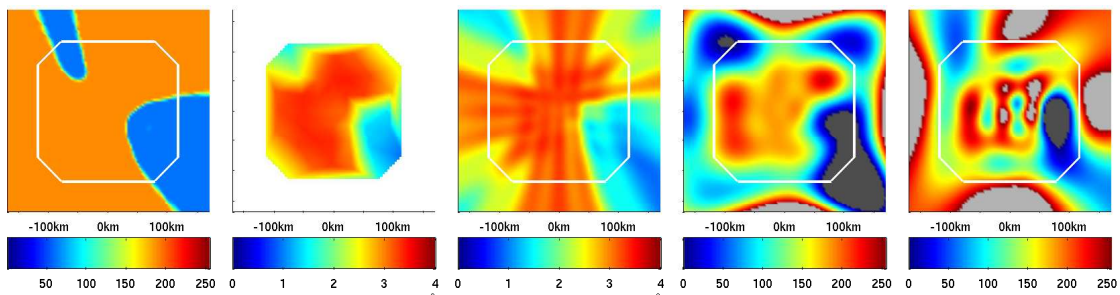
6.1 Synthetic test brightness distribution 1

Figure 5: Simple testimage used as input to the interpolation algorithms. From left to right: Synthetic skymap used to synthesise power readings for all 49 IRIS beams; original MIA interpolation algorithm applied to the synthetic power readings of all IRIS beams except the four corner beams ('good beams'); interpolation method A applied to simulated data from all 49 beams; interpolation method B1 applied to simulated data from all 49 beams; interpolation method B2 applied to simulated data from all 49 beams.

In a first step, the algorithms were compared using a synthetically generated dataset. The leftmost panel in figure 5 shows the image that was used as a skymap. Convolution of the 49 IRIS beampatterns with this skymap according to equation 6 results in 49 synthetically created 'power readings'. The four remaining panels show the results of applying these synthetically created power readings to the different interpolation algorithms.

The original MIA interpolation algorithm produces a good representation of the input distribution. The intensity values diminish towards the outside of the interpolated area. There are low-power artefacts due to the interpolation between distant beams, especially visible at the top and bottom of the image.

Algorithm A also produces a good representation of the input distribution. However, some effect of the beam patterns is visible, and the power values trail off towards the to right and bottom left of the image.

Algorithm B1 shows the general shape of the input distribution within the primary field of view. The distribution is not nearly uniform, however, and there are strong overshoots in both negative and positive

directions. Algorithm B2 clearly shows the doubled spatial resolution compared to B1, but also seems to emphasise the problems observed with B1. There are strong overshoots in positive and negative directions and a high level of variability even inside the primary field of view.
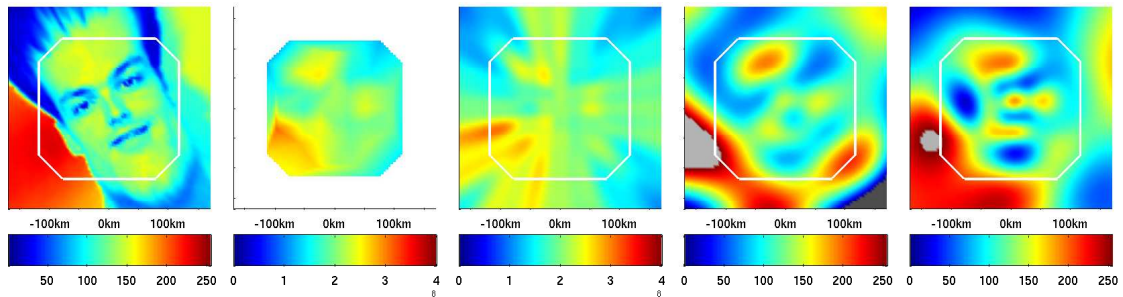
6.2 Synthetic test brightness distribution 2

Figure 6: More detailed test image used as input to the interpolation algorithms. From left to right: Synthetic skymap used to synthesise power readings for all 49 IRIS beams; original MIA interpolation algorithm applied to the synthetic power readings of all IRIS beams except the four corner beams ('good beams'); interpolation method A applied to simulated data from all 49 beams; interpolation method B1 applied to simulated data from all 49 beams; interpolation method B2 applied to simulated data from all 49 beams.

A more detailed input image was now used to investigate the reaction of the different algorithms to input data with higher spatial structure. See figure 6.

All algorithms seem to be able to resolve the main features (eyes, mouth). The influence of the beam-patterns is again clearly visible in the image created by algorithm A, while in the original MIA interpolated image the structure of the triangles used for the interpolation seems to show through. Both algorithms B1 and B2 seem to be doing rather well in this test case. They both emphasise the main structures rather better than the other algorithms. However, we can again observe a tendency to overshoot. Notice that both B1 and B2 tend to oscillate, especially visible for B2 as one moves from top to bottom.
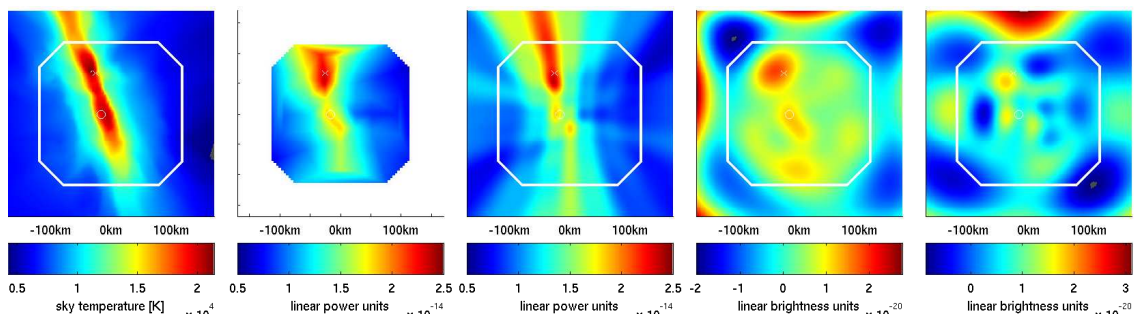
6.3 Quiet Dataset

Figure 7: Quiet Dataset (1999-09-23 21:47UT). From left to right: Skymap for this part of the sky at this particular moment in time; original MIA interpolation algorithm applied to actual data as received by IRIS at the time in question (only 'good beams' i.e. all beams except the four corner beams); interpolation method A applied to real data from all 49 beams; interpolation method B1 applied to real data from all 49 beams; interpolation method B2 applied to real data from all 49 beams.

Figure 7 shows 1-minute integrated data for a quiet part of the day around 21:47UT. The skymap panel shows that we would expect to see a band of high power across the field of view. Both the traditional MIA interpolation algorithm and method A tend to break the band up into sections of varying intensity, the effects of individual imaging beams showing through. Note, however, that method A at least seems to show the continuity of this high power section to the top and to the bottom above and beyond the working area of the MIA algorithm. Opposed to this, the traditional MIA method shows areas of low power to the top and the bottom of the high power section, resulting from the linear interpolation between the outermost beam centres. This is clearly not an accurate representation of the real sky brightness. This can, of course, be eliminated by further limiting the working area of the algorithm, further reducing the field of view of the instrument.

Both methods B, whilst still showing the approximate direction of the high power band, produce several local maxima and minima within the visible area, that clearly do not resemble the actual expected brightness distribution. Notice that the two largest local maxima seem to coincide with the location of the two radio stars.
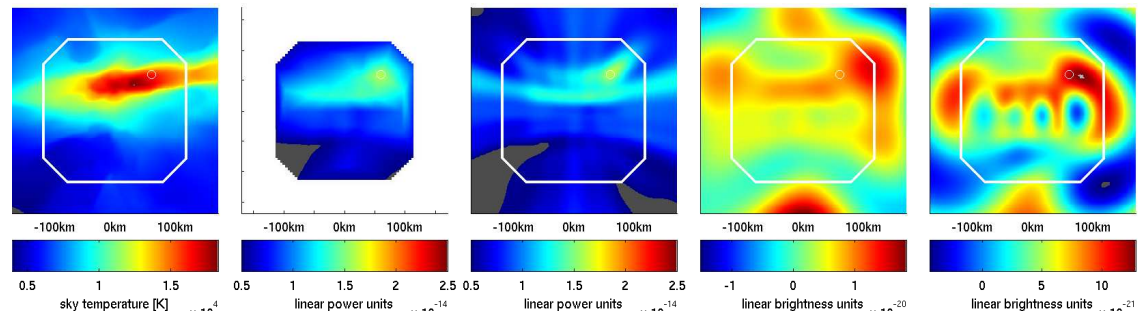
6.4 Absorption Event

Figure 8: Absorption Dataset (1999-09-23 05:23UT). From left to right: Skymap for this part of the sky at this particular moment in time; original MIA interpolation algorithm applied to actual data as received by IRIS at the time in question (only 'good beams' i.e. all beams except the four corner beams); interpolation method A applied to real data from all 49 beams; interpolation method B1 applied to real data from all 49 beams; interpolation method B2 applied to real data from all 49 beams.

Figure 8 is a 1-minute integrated dataset for a period of strong absorption around 05:23UT. Note the generally reduced power in all three images that are derived from real data. The image obtained by method A quite clearly shows individual beam shapes, even though those should in theory be completely removed by the algorithm. The image produced by the traditional MIA algorithm looks much smoother. It does, however, again show the sharp cutoffs, this time at the left and right, where we would expect the bright band to continue North and South. Also note the isolated brighter spot at the bottom. Looking at the skymap, we'd expect this spot to continue towards the bottom, as it indeed does for method A.

Method B, while clearly showing less overall power and the general direction of the high-power band, again seems to be the worst of the interpolation methods. It still manages to approximately show the location of the radio star. Especially method B2 shows strong oscillations in the horizontal direction.

6.5 Interference Event

Finally, figure 9 is an example of strong solar radio emissions as captured by the IRIS riometer at 03:39UT. The original MIA interpolation method clearly shows the high power received from the eastern side of the array. Similar to the previous figures, the image seems to contain errors in the extreme outer parts, especially visible at the right, where the interpolation algorithm produces a red band all along the right-hand side. In comparison, method A clearly separates the two power sources (galactic centre and solar
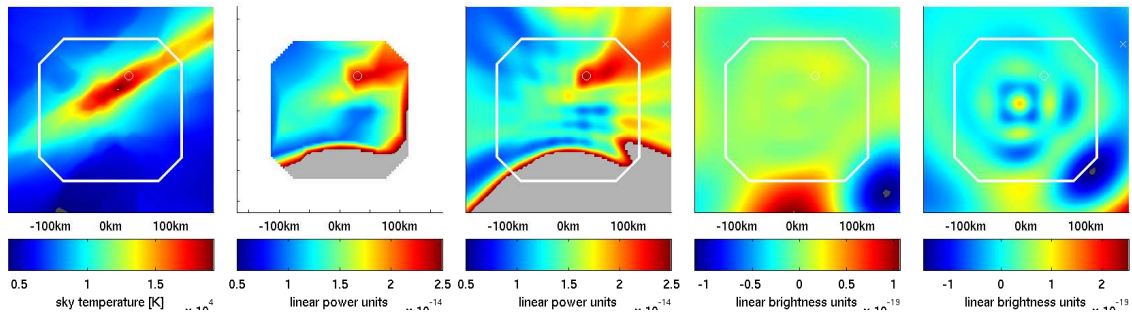
Figure 9: Interference Dataset (2004-09-23 03:39UT). From left to right: Skymap for this part of the sky at this particular moment in time; original MIA interpolation algorithm applied to actual data as received by IRIS at the time in question (only 'good beams' i.e. all beams except the four corner beams); interpolation method A applied to real data from all 49 beams; interpolation method B1 applied to real data from all 49 beams; interpolation method B2 applied to real data from all 49 beams.

radio emission). Overall, however, the image created by method A shows strong beam artifacts, wheras the MIA image looks smooth.

Method B1 seems to show the location of the radio star, and both methods B show high intensity in the direction of the solar radio emission. As before, they fail to accurately map the overall brightness distribution, however.

7 Conclusion

We compared the traditional MIA riometer power image interpolation method to two other methods that try to take the actual antenna radiation patterns of the riometer into account. While method A seems to give good results for the 49 beam riometer IRIS, methods B1 and B2 fail to accurately reconstruct the sky brightness distribution in the general case. Some results seem to suggest that methods B1 and B2 may be useful for identifying strong but spatially small features (figure 6).

At this stage, method A seems to perform equally to the traditional MIA method, both methods leading to slightly different results due to the different underlying principles. Given that the shapes of the beampatterns still show through in the resulting images, there seems to be no clear advantage of this method over the original MIA method. Which method is more suitable for the given task at hand will need to be decided on a case-by-case basis.

In its current form, method B fails to accurately model the brightness distribution to a high enough level of accuracy. It also suffers from oscillations and overshoots.

Although failing to deliver the increase in quality that was originally envisaged, it is hoped that this paper will trigger an increased interest in the area of interpolation algorithms for improving images generated by riometers, eventually leading to higher quality absorption images as a basis for high quality riometry work. The following sections gives a few ideas as to possible improvements and ongoing work.

8 Suggestions for Improvements / Future work

In interpolation method A, the influence of the beam patterns seems to be visible, whereas ideally this method should completely eliminate any beam-related patterns from the resulting image. One option might be to weight different beams differently in equation 3, aiming to put less emphasis on beams whose beampatterns are less certain. With regard to both methods A and B, the theoretically derived beam patterns could be replaced by actual beam measurements (provided these are available), or the theoretical beam patterns could be replaced by simulated beam patterns that take imperfect ground conditions into account.

Research is ongoing with the aim of improving method B's results. Currently, the matrix problem is ill-conditioned due to the inability of the IRIS beam pattern to discriminate between a large number of configurations of the model power distribution, effectively resulting in fewer equations than unknowns. In section 5.1, it has already been hinted at the fact that the number of equations used to model the brightness distribution (M) does not have to equal the number available data points (N). In fact, having $M < N$ allows for least-square fits and also for error ranges to be included in the approximation. This has the potential to significantly increase the quality of method B's output, along with a built-in quality indicator. Along the same lines, it is possible to add additional spatial contraints to increase the number of available equations. This will lead to a better conditioning of the matrix problem.

It may also be possible to improve method B's results by choosing different interpolation functions for the curve fitting process.

Finally, it has to be noted that interpolation of power images is only the first step to deriving absorption data in arbitrary (virtual beam) directions. Based on the interpolation algorithms presented in this paper, quiet day curves (QDCs) can potentially be calculated for arbitrary directions, enabling the derivation of absorption in these directions and therefore the creation of absorption images on arbitrary grids and with arbitrary resolutions.

Appendix A: Image interpolation with Spherical Harmonics

To confirm the spatial resolution that image approximation using spherical harmonics is able to produce, three test images have been approximated by calculating the first 49 and 256 coefficients for both the standard spherical harmonics approach (section 5.2) and the adjusted spherical harmonics approach (see section 5.3), see the figures 10, 11 and 12. Clearly, spherical harmonics are able to reproduce the source image with sufficient details, but deriving the coefficients from real-world data is non-trivial.

References

[1] Noach Amitay, Victor Galindo, and Chen Pang Wu. *Theory and analysis of phased array antennas*. Wiley-Interscience, New York, 1972.

[2] S. Browne, J. K. Hargreaves, and B. Honary. An imaging riometer for ionospheric studies. *Electronics and Communication Engineering Journal*, 7(5):209–217, October 1995.

[3] Bernhard F. Burke and Francis Graham-Smith. *An Introduction to Radio Astronomy*. Cambridge University Press, 1st edition, 1997.

[4] H. V. Cane. A 30 MHz map of the whole sky. *Australian Journal of Physics*, 31:561–565, 1978.

[5] C. F. del Pozo, M. J. Kosch, and F. Honary. Estimation of the characteristic energy of electron precipitation. *Annales Geophysicae — Atmospheres, Hydrospheres and Space Sciences*, 20(9):1349–1359, 2002.

[6] G. V. Haines. Spherical cap harmonic analysis. *Journal of Geophysical Research*, 90:2583–2591, 1985.

[7] J. K. Hargreaves and D. L. Detrick. Application of polar cap absorption events to the calibration of riometer systems. *Radio Science*, 37(3):7/1–7/11, 2002.

[8] Ernest William Hobson. *The theory of spherical and ellipsoidal harmonics*. 1955.

[9] C. G. Little and H. Leinbach. The riometer - a device for continuous measurement of ionospheric absorption. *Proceedings of the Institution of Radio Engineers*, 47:315–320, 1959.

[10] R. A. Makarevitch, F. Honary, I. W. McCrea, and V. S. C. Howells. Imaging riometer observations of drifting absorption patches in the morning sector. *Annales Geophysicae — Atmospheres, Hydrospheres and Space Sciences*, 22(10):3461–3478, 3rd November 2004.
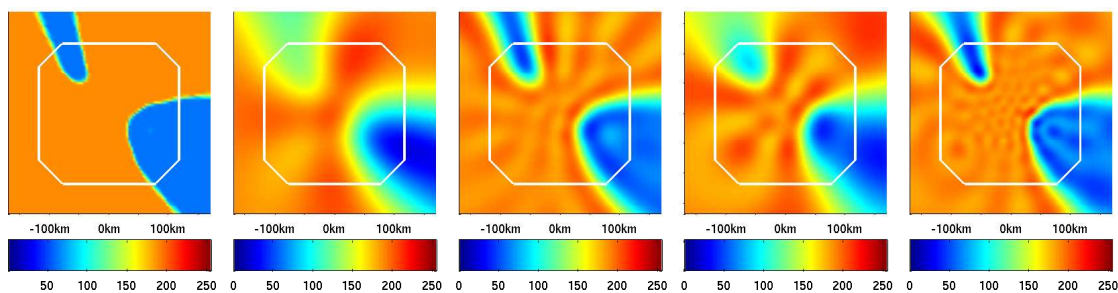
Figure 10: Interpolating an image using spherical harmonics. From left to right: original image; interpolated image using 49 coefficients and standard spherical harmonics; interpolated image using 256 coefficients and standard spherical harmonics; interpolated image using 49 coefficients and adjusted spherical harmonics; interpolated image using 256 coefficients and adjusted spherical harmonics.
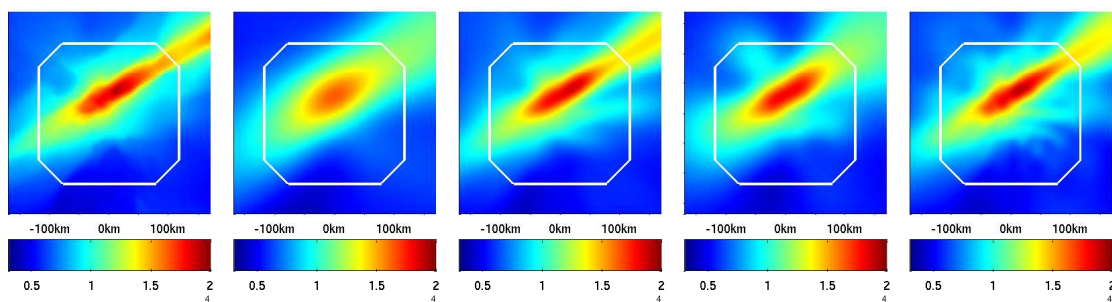

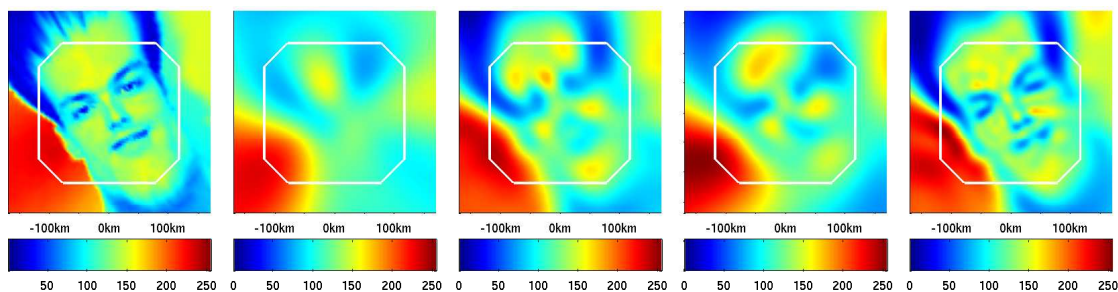Figure 11: Interpolating an image using spherical harmonics. Panels as in figure 10.

Figure 12: Interpolating an image using spherical harmonics. Panels as in figure 10.

[11] S. R. Marple and F. Honary. A multi-instrument data analysis toolbox. *Advances in Polar Upper Atmosphere Research*, 18:120–130, September 2004.

[12] B. Y. Mills and A. G. Little. A high-resolution aerial system of a new type. *Australian Journal of Physics*, 6:272–278, 1953.

[13] Yasuhiro Murayama, Hirotaka Mori, Shoji Kainuma, Mamoru Ishi, Ichizo Nishimuta, Kiyoshi Igarashi, Hisao Yamagishi, and Masanori Nishino. Development of a high-resolution imaging riometer for the middle and upper atmosphere observation program at Poker Flat, Alaska. *Journal of Atmospheric and Solar-Terrestrial Physics*, 59(8):925–937, 1997.

[14] E. Nielsen and T. Hagfors. Plans for a new rio-imager experiment in Northern Scandinavia. *Journal of Atmospheric and Solar-Terrestrial Physics*, 59(8):939–949, 1997.

[15] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vettering. *Numerical Recipes in C*. Cambridge University Press, 1st edition, 1988.

[16] T. J. Rosenberg, D. L. Detrick, D. Venkatesan, and G. van Bavel. A comparitive study of imaging and broad-beam riometer measurements: The effect of spatial strucure on the frequency dependence of auroral absorption. *Journal of Geophysical Research — Space Physics*, 96(A10):17793–17803, October 1991.

[17] A. De Santis. Conventional spherical harmonic analysis for regional modelling of the geomagnetic field. *Geophysical Research Letters*, 19:1065–1067, 1992.

[18] A. De Santis and J.M. Torta. Spherical cap harmonic analysis: a comment on its proper use for local gravity field representation. *Journal of Geodesy*, 71:526–532, 1997.

[19] Eric W. Weisstein. *CRC Concise Encyclopedia of Mathematics*. Chapman & Hall/CRC, 2nd edition, December 2002.