Reconstruction of electron precipitation characteristics from a set of multiwavelength digital all-sky auroral images

P. Janhunen

Finnish Meteorological Institute, Helsinki, Finland

Abstract

We present a method for reconstructing the precipitating electron flux from a set of multiwavelength digital all-sky auroral images. The method involves solving a large linear inversion problem, and it works for any number of all-sky imager stations and wavelengths. The idea in our method is to invert the energy-dependent precipitating electron flux directly and not the three-dimensional volume emission rate as an intermediate step. In this way, the inversion is automatically restricted to physical emission rate profiles. We show the effectiveness of the method in the all-sky camera case by simulated examples. The method is also applied to real data. The method is capable of correctly reconstructing not only auroral arc positions but also the precipitating electron energy and number fluxes, although not necessarily the full differential flux. A natural limitation is that arcs that are significantly off-zenith and are viewed sideways by only one all-sky camera station are smeared out. However, even for these arcs the average arc position and energy characteristics come out with reasonable accuracy.

1. Introduction

Ground-based digital all-sky camera (ASC) images contain a lot of information on the physics of auroral electron precipitation. For example, the four ASCs (soon to be expanded to six) belonging to the portion of the Magnetometer, Ionospheric Radar and All-sky Camera Large Experiment (MIRACLE) ground-based network that is located on the Scandinavian mainland produce 512×512 red (630.0 nm) and green (557.7 nm) images every 60 and 20 s, respectively, with 8-bit quantization and usually 1-s exposure time [*Syrjäsuo et al.,* 1998]. The cameras are running continuously and are now being properly calibrated so that accurate photon fluxes can be measured. However, thus far the MIRACLE all-sky images have been utilized in scientific studies only in a qualitative form. The only processing has been that the raw images have sometimes been mapped to the 110-km altitude level and then viewed in a geographic latitude-longitude grid.

In this paper we develop a data-processing method that goes much deeper: our aim is to invert the precipitating electron flux $F(\theta,\varphi,E)$ as a function of geographic colatitude θ, longitude φ, and energy E from a set of raw all-sky images. We formulate the problem in a flexible way so that any number of stations and wavelengths can be used. In the simplest case one has only a single ASC station that has taken a picture at one wavelength, but usually we have at least two stations, each of which has taken 557.7-nm and 630.0-nm images.

Methods for reconstructing the volume emission rates have been developed by using tomographic techniques [*Frey et al.,* 1996, 1998; *Nygrén et al.,* 1996; *Gustavsson,* 1998]. In all cases some a priori information is needed in order to regularize the inversion problem, which is otherwise ill-posed. Our approach differs from the previous ones in that we do not try to invert the three-dimensional (3-D) photon volume emission rates, but instead combine the geometric and physical problems to invert directly the differential electron fluxes. We think that this approach is better suited for processing all-sky camera images, which provide less accurate data than more sophisticated imaging systems such as Auroral Large Imaging System (ALIS) [*Steen,* 1989; *Steen and Brändström,* 1993]. The strength of our method is that by not treating the volume emission rate as an arbitrary function of three variables but deriving it from a differential electron flux $F(\theta,\varphi,E)$, we effectively restrict the space of possible solutions in a physically correct way. For example, volume emission rate height profiles containing a gap in some altitude range cannot occur in our method. A somewhat similar idea, but still remaining completely in the tomographic context, has been presented earlier [*Doe et al.,* 1997; *Semeter et al.,* 1999].

While $F(\theta,\varphi,E)$ is the quantity to be solved for in our method, it is not necessarily realistic to expect that the full $F(\theta,\varphi,E)$ is reconstructed with quantitative correctness in cases where the precipitating flux is strongly non-Maxwellian. For this reason, in this paper we concentrate on verifying the correctness of the first two energy moments of F (the number and energy flux).

2. Model

2.1. Inversion Method

The general methods for doing linear inversion are well known [*Press et al.,* 1992]. First, one must be able to solve the direct problem, which we write as $m = Au$, where u is the unknown vector (denoted also by F in this paper), m is the measurement vector, and A is the matrix (usually not square) defining the linear relationship. Then one incorporates the a priori regularization information by defining the quadratic form to be minimized,

$$f(u) = \frac{1}{2}\left[C^{-1/2}(Au - m)\right]^2 + \frac{1}{2}\lambda u^T H u, \quad (1)$$

where H is a regularization operator, λ is a parameter controlling the relative importance of regularization versus data, and C is the covariance matrix of Gaussian measurement errors. C is diagonal if the errors for each data point are independent, as we shall assume, and the diagonal entries are then equal to the squared measurement errors. The minimization of $f(u)$ corresponds to assuming a measurement error η probability distribution $\sim \exp(-(1/2)\eta^T C^{-1}\eta)$ and an a priori distribution of u of the form $\sim \exp(-(1/2)\lambda u^T H u)$.

We can write $f(u)$ in the form

$$f(u) = \frac{1}{2}u^T Q u - c^T u + \mathrm{const}, \qquad (2)$$

where $Q \equiv A^T C^{-1} A + \lambda H$ and $c \equiv A^T C^{-1} m$. The additive constant is not relevant for the minimization process. We assume that the matrices H and C are symmetric and positive definite. It then follows easily that also Q is symmetric and positive definite, for any matrix A and any $\lambda > 0$. Thus the minimization problem is a convex one, and it can be solved by, e.g., the conjugate gradient (CG) algorithm. The minimal vector u satisfies $\nabla_u f(u) = 0$, being equivalent with the linear equation $Qu = c$. For definiteness, we always use the initial guess $u = 0$.

In cases we have studied, negative solution vector components have not occurred, so separate positivity constraints do not seem to be needed. If needed, such constraints could be added by projections onto convex sets (POCS) methods [*Press et al.,* 1992].

The application of the CG method requires that we provide subprograms to compute Au for a given u as well as $A^T m$ for a given m. We can do Au as soon as we know how to solve the direct problem. Doing $A^T m$ is also not difficult and will be discussed shortly. The matrix A is applied to the particle flux vector, and it produces the corresponding all-sky images. Each energy component of the differential flux affects all altitude levels and thus contributes to all pixels that touch a one-dimensional line. Thus, although A is sparse, each column of A still contains at least several tens of nonzero components. To reduce the memory consumption, we decompose A as

$$A = GP, \qquad (3)$$

where the physics matrix P maps particle fluxes to volume emission rates, and the geometry matrix G maps the volume emission rates to all-sky images. Since the physics is the same independent of latitude and longitude (we do not consider horizontal variations in the neutral atmosphere), the P matrix is block diagonal with identical blocks P_1. One block, P_1, maps a single electron flux vector at fixed latitude and longitude into a volume emission rate altitude profile. We normally use 12 energy and 20 altitude levels, both logarithmically spaced, so that P_1 is a 20 by 12 matrix, except that in Tables 1 and 2 below we will also report 24×40 resolution results. The energy range is 0.1–8 keV, and the altitude range is 90-300 km. P_1 can be precomputed and stored as an ordinary dense matrix. The subprogram that applies P to a given vector simply does a P_1 matrix multiplication at each latitude and longitude grid point.

The geometry matrix G is stored explicitly using a columnwise sparse matrix storage scheme [*Duff et al.,* 1986; *Press et al.,* 1992]. Since it describes the contribution of each voxel (volume element containing photon emission) to each ASC pixel, the average number of nonzero elements per column is smaller than with the A matrix and thus fits in the available memory more easily.

Since $A^T = P^T G^T$, coding the computation of $A^T m$ for any m, which is required in the CG method, is now easy using the explicit sparse and block diagonal representations of G and P.

2.2. Physics Matrix

The physics matrix P_1 translates the electron differential number flux F (m^{-2} s^{-1} keV^{-1}) into photon volume emission rate e (m^{-3} s^{-1}) under the assumption of an electron energy spectrum that can be written as a superposition of (possibly many) Maxwellians. There is a different physics matrix for each wavelength. In this work we consider the 427.8-nm blue, 557.7-nm green, and 630.0-nm red spectral lines, which are generally the most prominent auroral emissions [*Rees and Luckey,* 1974]. Currently, the MIRACLE ASCs take pictures with green and red filters, and blue images will be acquired starting from fall 2000. It turns out that, approximately, the blue emission rate is proportional to the energy deposition rate, and the green/blue and red/blue ratios are, again approximately, functions of altitude only [*Rees and Luckey,* 1974; *Strickland et al.,* 1989]. From Figure 7 of *Rees and Luckey* [1974] or from Figure 6 of Strickland *et al.* [1989], we infer that the yield of 427-nm photons is ∼160 R/(erg cm^{-2} s^{-1}). This value is actually valid only for about 1-2 keV characteristic energy [*Rees and Luckey,* 1974], but we use it approximately for all energies (the maximum approximation error is ∼25% and occurs for high energies). Recall that one obtains the emission in Rayleighs by expressing the height-integrated volume emission rate in m^{-2} s^{-1} and multiplying by 10^{-10} [*Hunten et al.,* 1956]. From these numbers we infer that 1 keV of deposited energy gives 0.256 427.8-nm photons on the average. We assume that this number does not depend on altitude. The value adopted here affects all of our results in a multiplicative fashion. For the red/blue and green/blue ratio profiles, we use Figure 2 of *Rees and Luckey* [1974] and model the intensity ratio profiles given there by the following expressions:

$$\begin{aligned} \frac{I630}{I427} &= 10^{1.3-8\times10^{-5}(260-h)^2-4\times10^{-7}(260-h)^3}, \\ \frac{I557}{I427} &= 10^{1.1-10^{-4}(200-h)^2}, \end{aligned} \tag{4}$$

where h is height expressed in kilometers. Above 260 km the ratios are set to the values at 260 km. The profiles of (4) are shown in Figure 1. Figure 1

Thus we are ready with the physics matrix computation, once we specify how to calculate the height-dependent energy deposition rate for a given electron flux. Our method is based on formulas originally given by *Rees* [1963]. They are also given by *Rees* [1989]. We first need to specify the neutral density model. We use the simple expression

$$\begin{aligned} n_n = (&10^{-4.68-0.085(h/\mathrm{km}-80)} + \\ &10^{-7.53-0.018(h/\mathrm{km}-80)})\ \mathrm{kg\ m^{-3}}. \end{aligned} \tag{5}$$

The electron range is then computed from

$$R = \left[4.30 \times 10^{-6} + 5.36 \times 10^{-5}(E/\mathrm{keV})^{1.67}\right] \mathrm{kg\ m^{-2}}, \tag{6}$$

where R is the range and E is the particle energy. The atmospheric depth D is defined as $D(h) = \int_h^\infty dz n_n(z)$, and the height-dependent energy deposition rate ε is then given by

$$\varepsilon = \frac{\Lambda(D/R) n_n E}{R} \Phi, \tag{7}$$

where $\Lambda(x)$ is the "cosine" dependence function (corresponding to an isotropic source distribution) plotted by *Rees* [1963] and Φ is the precipitating number flux (dimensionality $\mathrm{m}^{-2}\ \mathrm{s}^{-1}$). For $\Lambda(x)$ we use the polynomial fit

$$\Lambda(x) = \max(0, -11.639x^6 + 32.1133x^5 - 30.8543x^4 + 14.6063x^3 - 6.3375x^2 + 0.6138x + 1.4946). \quad (8)$$

The energy deposition rate ε is expressed in keV $\mathrm{m}^{-3}\ \mathrm{s}^{-1}$ and multiplied, as discussed above, by 0.256 to obtain the blue volume emission rate in $\mathrm{m}^{-3}\ \mathrm{s}^{-1}$. The red and green rates are finally obtained by using (4).

2.3. Geometry Matrix

The geometry matrix G operates on a volume emission rate profile vector $e(\theta, \varphi, h)$ and produces a set of synthetic ASC images, one image for each station and wavelength. Besides pure geometry, G also models the ASC intensity calibration and atmospheric attenuation. Each column of G contains the coefficients telling how much a unit volume emission at the given voxel (θ, φ, h) contributes to the ASC pixels. If the projected voxels are chosen to be approximately the same size as pixels, each column of G has only a few (four, say) nonzero entries.

There are two possible principally different approaches to construct the geometry matrix. Either one can process each pixel in turn and integrate the emission numerically along the line of sight or one can process each voxel in turn and "render" it on the pixel plane of each camera. Both approaches should lead to the same result if implemented accurately. We selected the rendering approach. One of the benefits of the rendering approach is that it can be made exactly conservative in the sense that no emitted light is lost or spuriously gained owing to numerical inaccuracies: the discretization error can affect only the way in which light is distributed among neighboring pixels. In the line integration approach, on the contrary, the numerical sampling along the line of sight would have to be dense enough so that also the thinnest representable arcs would not be undersampled, which could be computationally a rather severe requirement, especially in cases where multiple cameras are present and the viewing direction is far from the zenith of the camera under consideration. The rendering algorithm is explained in Appendix A.

Atmospheric attenuation is a nonnegligible effect, at least for the blue emission. We use a very simple model that includes ozone absorption and Rayleigh scattering as well as some estimate for aerosol scattering and absorption and took the numbers from Table 1 of *Oikarinen* [2001]. This table was computed for a 23-km surface visible range. Typically, viewing conditions in northern Scandinavia are more favorable, so we reduced the aerosol effects by a factor of 3, corresponding to ~80 km visible range. The following vertical column optical depths τ for different wavelengths were obtained: $\tau_{427} = 0.42, \tau_{557} = 0.22$, and $\tau_{630} = 0.17$. The incident photon flux must be multiplied by $\exp(-\tau/\cos z)$, where z is the zenith angle, to get the attenuated flux. Scattering is treated in the same way as absorption; that is, scattered light is assumed to be lost. Taking scattered light self-consistently into account in the inversion would, of course, be quite a complicated task.

2.4. Regularization

The purpose of the regularization operator H is to incorporate the a priori probability distribution of u in the form $\exp[-(1/2)\lambda u^T H u]$. After experimentation we found that a straightforward smoothing operator aimed at minimizing the L^2 norm of the second derivative,

$$\int d^3\mathbf{x}(\nabla\nabla u)^2 = \int d^3\mathbf{x} u \nabla^4 u, \qquad (9)$$

works well. Here, ∇ operates in the (θ, φ, E) space. In other words, the smoothing operator is given by $H = \nabla^4$. We also tried versions of H that minimize the first or third derivatives but obtained slightly inferior results. The coefficient λ is selected so that $[C^{-1/2}(Au - m)]^2$ becomes comparable to the number of elements in u, which means physically that the model is made to fit the data with about one standard deviation at each data point. The results are not sensitive to the exact value of λ.

2.5. Spatial Grid

To reduce computing time, we use a regular but stretched latitude-longitude grid. The grid spacing varies in such a way that it is finest at the zenith of the stations, and the grid spacing is scaled in the same way as the horizontal extents of pixels projections on the 110-km altitude plane. In the case of multiple cameras, the minimum spacing with respect to all cameras is taken, so that the grid is dense at all camera zeniths. Figure 2 shows the spatial grid.

Figure 2

3. Inversion With Simulated Data

To test the inversion algorithm, we create an artificial electron differential flux $F(\theta, \varphi, E)$ that involves three separate arcs containing "zigzag" folds at irregular intervals. The electron distribution is assumed to be an accelerated Maxwellian [*Evans,* 1974] with temperature equal to one half of the accelerating voltage V. V is assumed to vary linearly from 2 keV at geographic longitude 20 to 4 keV at geographic longitude 30. For longitudes smaller and larger than 20 and 30, V is set to 2 and 4 keV, respectively. The precipitating energy flux is set uniformly to 35 mW m^{-2} at the center of the arcs, and the half width of the arcs is 0.1°.

The model energy flux is shown in Figure 3a. It contains four arcs, which we number 1–4 so that 1 is the southernmost arc. Because the arc is sampled in a discrete grid, there are slight variations in the energy flux even though it is, in principle, set exactly to 35 mW m^{-2}. The computed physics and geometry matrices are used to obtain simulated all-sky images at green and red wavelengths for two ASC stations: Kilpisjärvi (KIL; latitutde 69.02, longitude 20.79) and Kevo (KEV; latitude 69.76, longitude 27.01). The simulated green images for the two stations are shown in Figure 4. The images were synthesized in 200×200 size. A realistic amount of brightness-dependent Gaussian noise was added to the images. The arc positions were selected so that arc 1 does not overlap with the others but is viewed significantly off-zenith, arc 2 has the best viewing geometry, and arcs 3 and 4 are both off-zenith and close enough to each other to overlap visually.

Figure 3

Figure 4

The inverted energy flux is shown in Figure 3b. Arc 1 total energy flux has been reconstructed almost correctly, but the arc is smeared out owing to oblique viewing from both stations. Arc 2 is very well reconstructed. Arc 3 is partly well reproduced, but at around 20° geographic longitude the reconstruction suffers from overlap with arc 4. Arc 4 (the northernmost arc) is rather poorly reconstructed because of interference with arc 3, but its average position can still be determined. The Figure caption supplies the energy fluxes integrated over the whole region. These values are also in rather good agreement, but the reconstruction tends to overestimate the total energy flux slightly (in this case, ∼15%). This will be discussed more below.

Figure 5

Figure 5 shows the model and reconstructed electron number fluxes (in units of current density). Since the model electron energy is smaller on the westward side, the number flux is correspondingly larger there, the energy flux being constant. The inversion algorithm recognizes this trend correctly, and for arc 2 the agreement is quantitatively good. The integrated values are in close agreement (5% error).

Figure 6

Figure 7

To investigate the wavelength effects, we show in Figure 6 the energy flux reconstructed from the green images only. The result is surprisingly good, almost indistinguishable from the version where red images were also used (Figure 5b). For comparison, in Figure 7 we show the reconstructed number flux using only red images. Now the result is clearly inferior: the arcs are more smeared-out, and the peak number flux has been overestimated. However, the total current (293 kA) carried by the precipitating electrons is still in good agreement with the true value (298 kA).

Table 1

More quantitatively, in Table 1 we show the relative reconstruction errors in the L^2 norm for all seven possible wavelength combinations. The first column, L^2(eflux), is the global average energy flux error calculated using doubled energy and altitude resolution (24 energy and 40 altitude levels). All grid points closer than 250 km to one of the cameras are included in this norm. The best wavelength combination is RB (20.6 % error), RGB and GB coming rather close. The second column shows the same result computed for the standard resolution (12×20). The errors are a couple of percent larger than for the doubled resolution (for the less reliable single-wavelength reconstructions, the errors are about the same or even smaller, which may be due to random effects), from which we can conclude that the standard resolution is already good enough for practical work. The third column, L^2_{loc}(eflux) is the "local" error norm of the energy flux computed over the smaller rectangular domain shown in Figure 7. Not surprisingly, the errors are clearly smaller, but again, RB is best, with RGB and GB closely following. Finally, the fourth column shows the corresponding quantity for the number flux. Here the results are similar to the energy flux case, except that the B case is now clearly inferior with 20.3% error (in the energy flux, the B case comes close to RGB and GB). Both local norms are computed for the doubled resolution.

Table 2

Table 2 lists the errors of the total integrated energy and number fluxes in doubled resolution, in both global and local versions. The results are similar to Table 1 in that RB again emerges as the winner (exception: the local R norms are also quite small, but probably just by chance). All of the three combinations RGB, RG, and GB are about equally good, except that RG has more problems with the energy flux than the others.

Generally, the error norms reported in Tables 1 and 2 show that the results are rather good because errors that are less than about 30% can already be considered satisfactory because other factors such as atmospheric attenuation and errors in emission physics modeling are probably larger. Looking into the performance of different wavelength combinations, one would expect the error to decrease when more wavelengths are added, and one would also expect the blue emissions to be superior to the other ones, especially in the energy flux reconstruction, because the blue emission rate is directly proportional to the energy deposition rate with no altitude corrections. This is what Tables 1 and 2 show, except for two slight anomalies: RGB is inferior to RB, and the reconstruction more often than not overestimates the total fluxes (i.e., the numbers in Table 2 are mostly positive). These anomalies are not due to insufficient energy or altitude resolution, because they are present in both standard and doubled resolution (Table 1). We also found that the anomalies do not disappear if one increases the upper height limit from 300 to 500 km.

One possible explanation for the RGB/RB anomaly is that when as many as three wavelengths are available, precipitation of different energy may get assigned to slightly different portions of the auroral arcs because the tomographic data in the case of only two cameras may not fully resolve such ambiguities. This can slightly broaden the arc and cause a spurious increase in the fluxes. This may also explain why the method usually overestimates rather than underestimates the fluxes, especially the energy flux, which is more sensitive to particle energy than the number flux.

In practice, the performances of the combinations RGB, RB, RG, and GB are similar enough that the choice of combination is probably more dictated by data availability and quality, uncertainties in atmospheric attenuation (which affects especially the blue images), uncertainties in emission physics (which plague the green images), and temporal constraints (which is a problem of red emissions having 110-s lifetimes). It is interesting from Tables 1 and 2 that the GB combination performs surprisingly well, and thus it could be useful in cases where high time resolution green and blue images are available and one wants to track rapidly moving auroras. In other cases, performing reconstructions with several combinations (RGB, RG, and RB, for instance) and comparing the results might be useful to gain understanding of the relative errors of the method when applied to real data. With currently available MIRACLE data, we are limited to red and green images, however.

4. Inversion With Measured Data

In section 3, we showed that the method works well in a simulated case. This, however, was a test of the inversion algorithm only, not a test of the correctness of the emission physics modeling. When processing real ASC data, we must also consider possible nonideal weather conditions, moonlight interference, camera-pointing errors, badly known intensity calibration, etc. At the moment of this writing, an effort is going on to properly calibrate the MIRACLE all-sky cameras, but in the data available thus far, the intensity calibration is known only roughly: 1 digital unit corresponds to $\sim$100 R.

To have some idea of how much all of these uncertainties might affect the inversion, we select an event where the conditions are far from ideal but which has conjugate FAST satellite [*Carlson et al.,* 1998] particle data from which we get an independent estimate of the arc precipitation characteristics. The event is November 3, 1998, 1736 UT, and in an earlier analysis the FAST energy and number flux at the arc were estimated to be about 10 mW m^{-2} and 2.5 μA m^{-2}, respectively [*Janhunen et al.,* 2000]. The raw Kevo green image is shown in Figure 8. The full Moon is to the southeast. It was manually removed, together with the surrounding sky area, whose luminosity is mostly scattered moonlight. Pixels marked for removal are disregarded by the inversion procedure. Manual moonlight removal is not essential for the working of the

Figure 8

inversion, but it makes plotting the results more convenient.

Figure 9

The reconstructed energy flux is shown in Figure 9. The main arc that was studied by *Janhunen et al.* [2000] (called arc 2 there) resides at $\sim$70° latitude. FAST passed through the arc at $\sim$18° longitude at 1736:20 UT. For reconstruction, we select the epoch 1736:00 UT, however, since a red image was not taken at 1736:20; the difference is immaterial since the green images taken at 1736:00 and 1736:20 are very similar. The reconstructed arc intensity is $\sim$12 mW m^{-2}, which is actually in good agreement with FAST ($\sim$10 mW m^{-2}). The same is true for the number flux (Figure 10): both FAST and the reconstruction give $\sim$2.5 μA m^{-2} current density at the flyover point. About 100 km north from arc 2 there is arc 1, which was judged as "not a real arc but a precipitation blob" in the earlier analysis. The reconstruction suggests that arc 1 consists of at least four "blobs", the westernmost of which (at $\sim$17°-18° longitude) is the one that FAST passed through. Probably arc 1 is an active one containing a periodic curl structure, which, when viewed from the side, looks like a series of blobs. The energy and number flux seem to be similar to those of arc 2, which was also the FAST result. The parabolic structures in the southeastern parts of Figures 9 and 10 are due to the moonlight removal. In this event there is considerable amount of haze, which scatters moonlight. Some of the luminosity seen by Kevo in the northeast may be due to a cloud of fog reflecting moonlight. Not being told otherwise, the inversion program considers it auroral precipitation anyway and generates a rather large energy flux region at about 72.5° latitude, 32° longitude. For this reason the integrated energy and number flux are probably overestimations for this event.

Figure 10

5. Discussion and Conclusions

We demonstrated that electron energy and number fluxes can be successfully reconstructed from MIRACLE ground-based auroral images with reasonable accuracy, even under non-ideal viewing conditions. The most important uncertainties involved are within the emission photochemistry, which was modeled rather crudely in this paper. For red and blue emissions the modeling could be somewhat improved by using the best currently available models, but then one would also have to consider questions like the variable chemical composition of the neutral atmosphere. For green emissions, which are the brightest, the photochemistry is in any case known with less certainty [*Rees and Luckey,* 1974].

Other uncertainties are involved with weather conditions and the amount of atmospheric attenuation and refraction. For good weather conditions, however, these effects should be almost negligible. Further improvement in reconstruction accuracy may be obtained starting from fall 2000 data, when the fully calibrated cameras start operating.

The spatial reconstruction accuracy is good for arcs that are close to the zenith of one of the stations and especially between the stations where tomographic information is obtained. Arcs that are viewed sideways from more than 100 km horizontal distance are smeared out. However, this seems unavoidable in any reconstruction method, except by using some new type of a priori information with strong preference to narrow arcs. Reconstruction in energy succeeds best when both red and green (or blue) emissions are present because then the red/green (or red/blue) ratios give good information on the electron energy. Very low energy precipitation gives only red emission, while very high energy precipitation gives no red emission. Thus we have restricted the energy range to 100 eV – 8 keV. Precipitation more energetic than 8 keV will likely give a roughly correct energy flux but an overestimated number flux.

The computer requirements of the program are, by today's standards, rather modest. The solver phase takes ~3 min for 200×200 images on a 400-MHz Pentium-2 processor, the memory requirement being ~100 MB. The geometry matrix computation takes a little bit longer, but it can be precomputed and stored in a file if the same set of ground stations is used many times.

The reconstructed differential flux $F(\theta, \varphi, E)$ contains a lot of information on many interesting physical quantities. For example, the 3-D ionospheric local conductivity tensor components, as well as the height-integrated Hall and Pedersen conductances, could be calculated from F [*Frey et al.,* 1998]. In this paper we looked only at the first two energy moments of F, the number flux and the energy flux, and demonstrated that they are reproduced well. The full differential flux $F(\theta, \varphi, E)$ is not necessarily reconstructed well if the true F is sufficiently far from a Maxwellian.

The number flux that one derives from F contains only the part of the total current that is carried by energetic magnetospheric downward moving electrons. Within auroral arcs, this is likely to be approximately equal to the total current density because at least at Freja satellite altitude (~1700 km), the current densities estimated from the satellite magnetometer and from the electron detector instrument generally match [*Olsson et al.,* 1998]. The number of grid points in height and energy (20 and 12, respectively) as well as their ranges have been selected to suit the particular type of data that MIRACLE currently produces (8-bit 512x512 images) as well as the level of computing power that is at this moment (2000) easily available. Having better cameras and faster computers could warrant using a larger number of grid points and could result in achieving more accurate results.

Appendix A: The Voxel Rendering Algorithm

The rendering task needed to construct the geometry matrix is to "render" each voxel on the pixel plane of each camera. The voxels cover the photo-emitting ionosphere, and they are slightly nonrectangular cells aligned with latitude, longitude, and the local magnetic field direction. Upon assuming that the volume emission rate is constant within the voxel, this amounts to finding the volume of intersection between the voxel and the narrow quadrilateral viewing cone of each pixel on the pixel plane. In two-dimensional language this means finding the average luminosity within each pixel cell (Figure 11), where the luminosity is proportional to the distance that the line-of-sight ray travels inside the voxel. Figure 11

The problem is made much easier by first decomposing the voxel into tetrahedra (for quasi-rectangular voxels, five tetrahedra are enough) and rendering each of them separately. Since the voxels are slightly nonrectangular, the decomposition into tetrahedra is not completely exact (there may be a slight overlap/hole between adjacent voxels after decomposition), but the error introduced by this is quite small. The 2-D projection of a tetrahedron is composed of either three or four nonoverlapping triangles, depending on the viewing angle (Figure 12). If one neglects the perspective correction, which is a good approximation since the voxel dimension is small compared to its distance from the cam- Figure 12 era, one finds that the distance of intersection between the line-of-sight ray and the tetrahedron changes linearly within each projected triangle. A linear function of two coordinates inside each triangle is uniquely fixed as soon as we define its value at the three triangle corner points. For each triangle being part of the 2-D projected image of a tetrahedron, two of the three corner points lie along the outer edge of the tetrahedron, where the distance of intersection (the luminosity) is, of course, zero. The third triangle corner point is common to all triangles forming the tetrahedron: it is the central point where the luminosity reaches its peak value. Thus there is only one remaining free parameter in the linear functions within the 3 or 4 triangles. This parameter is uniquely fixed by the requirement that the total luminosity of the tetrahedron has the correct value.

The remaining problem is to calculate the 2-D integral of a linear function $u(x, y) = a + bx + cy$ (a, b, and c are known) over a set S that is an intersection of a pixel cell (i.e., a square) and a triangle defined by its three corner points $(X_1, Y_1), (X_2, Y_2)$, and (X_3, Y_3) (Figure 13). Being an intersection of two convex polyhedra, the set S is a convex polyhedron, and we must find its corner points $\{(x_i, y_i)\}_{i=1}^{n}$. The corner points consist of (1) all corners of the triangle that fall within the pixel cell, (2) all corners of the pixel square that are inside the triangle, and (3) all edge intersections between the edges of the triangle and the edges of the pixel square. We can then sort the corner point list so that they form a loop in the positive winding direction (counterclockwise). This can be done by computing the phase angle with respect to the mean point $(\bar{x}, \bar{y}) = \sum_{i=1}^{n}(x_i, y_i)/n$ and sorting the points in ascending order by the phase angle. The computation of the phase angle involves the arcus tangent operation $\mathrm{atan}(y/x)$, which is computationally slower than basic arithmetic. The arcus tangent call can, however, be avoided by finding another function $f(x, y)$ that is fast to compute and that is of the form $f(x, y) = g(\mathrm{atan}(y/x))$, where g is a monotonically increasing function of one variable, because such a function used in place of $\mathrm{atan}(y/x)$ gives an identical ordering of the corner points. A convenient choice for such f is given by

$$f(x, y) = \begin{cases} y/(x + |y|) & \text{if } x \geq 0 \\ y/(x - |y|) + 2.1 & \text{otherwise.} \end{cases} \tag{A1}$$

We are now ready to compute the integral

$$I = \int_S dx dy (a + bx + cy). \tag{A2}$$

Applying Stokes' theorem, we write it as a loop integral,

$$I = \oint_{\partial S} d\mathbf{l} \cdot \mathbf{A}, \tag{A3}$$

where $\mathbf{A} = (A_x, A_y, 0)$ is a "vector potential" satisfying $(\nabla \times \mathbf{A})_z = \partial A_y/\partial x - \partial A_x/\partial y = a + bx + cy$ and the line integral is performed along the piecewise linear path formed by the above constructed sorted corner point list. A solution for $\mathbf{A}$ is $A_x = -cy^2/2, A_y = ax + bx^2/2$. The loop integral (A3) can then be explicitly computed along each of the linear contour pieces:

$$\begin{aligned} &\int_{(x_1,y_1)}^{(x_2,y_2)} d\mathbf{l} \cdot \mathbf{A} = \\ &\frac{1}{6}\Big[(x_1 - x_2)c(y_1^2 + y_1 y_2 + y_2^2) + \\ &(y_2 - y_1)(3a(x_1 + x_2) + \\ &b(x_1^2 + x_1 x_2 + x_2^2))\Big]. \end{aligned} \tag{A4}$$

For computational efficiency, our algorithm contains some further checks to detect empty intersection sets S as early as possible. The rendering algorithm is called only once before the start of the CG iterations, and the results of the algorithm are stored in the sparse geometry matrix G. The G matrix can also be precomputed and stored in a disk file to skip the rendering algorithm entirely, e.g., in cases where a sequence of images is analyzed with similar geometry.

Acknowledgments. The authors thanks Johanna Tamminen for discussions about inverse problems, Liisa Oikarinen for discussions about atmospheric attenuation, and Kirsti Kauristie, Sanna Mäkinen, and Mikko Syrjäsuo for maintaining and calibrating the MIRACLE all-sky cameras.

Janet G. Luhmann thanks Joshua Semeter and other referees for their assistance in evaluating this paper.

References

Carlson, C.W., R.F. Pfaff, and J.G. Watzin, The Fast Auroral SnapshoT (FAST) mission, *Geophys. Res. Lett., 25,* 2013–2016, 1998.

Doe, R.A., J.D. Kelly, J.L. Semeter, and D.P. Steele, Tomographic reconstruction of 630.0 nm emission structure for a polar cap arc, *Geophys. Res. Lett., 24,* 1119–1122, 1997.

Duff, I.S., A.M. Erisman, and J.K. Reid, *Direct Methods for Sparse Matrices,* Clarendon, Oxford, England, 1986.

Evans, D.S., Precipitating electron fluxes formed by a magnetic field aligned-potential difference, *J. Geophys. Res., 79,* 2853–2858, 1974.

Frey, S., H.U. Frey, D.J. Carr, O.H. Bauer, and G. Haerendel, Auroral emission profiles from three-dimensionally reconstructed arcs, *J. Geophys. Res., 101,* 21,731–21,741, 1996.

Frey, H.U., S. Frey, B.S. Lanchester, and M. Kosch, Optical tomography of the aurora and EISCAT, *Ann. Geophys., 16,* 1332–1342, 1998.

Gustavsson, B., Tomographic inversion for ALIS noise and resolution, *J. Geophys. Res., 103,* 26,621–26,632, 1998.

Hunten, D.M., F.E. Roach, and J.W. Chamberlain, A photometric unit for the airglow and aurora, *J. Atmos. Terr. Phys., 8,* 345–346, 1956.

Janhunen, P., A. Olsson, O. Amm, and K. Kauristie, Characteristics of a stable arc based on FAST and MIRACLE observations, *Ann. Geophys., 18,* 152–160, 2000.

Nygrén, T., M. Markkanen, M. Lehtinen, and K. Kaila, Application of stochastic inversion in auroral tomography, *Ann. Geophys., 14,* 1124–1133, 1996.

Oikarinen, L., Polarization of light in UV-visible limb radiance measurements, *J. Geophys. Res., 106,* 1533–1544, 2001.

Olsson, A., L. Andersson, A.I. Eriksson, J. Clemmons, R.E. Erlandsson, G. Reeves, T. Hughes, and J.S. Murphree, Freja studies of the current-voltage relation in substorm related events, *J. Geophys. Res., 103,* 4285–4301, 1998.

Press, W.H., S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing,* 2nd ed., Cambridge Univ. Press, New York, 1992.

Rees, M.H., Auroral ionization and excitation by incident energetic electrons, *Planet. Space Sci., 11,* 1209–1218, 1963.

Rees, M.H., *Physics and Chemistry of the Upper Atmosphere,* Cambridge Univ. Press, New York, 1989.

Rees, M.H., and D. Luckey, Auroral electron energy derived from ratio of spectroscopic emissions, 1, Model computations, *J. Geophys. Res., 79,* 5181–5186, 1974.

Semeter, J., M. Mendillo, and J. Baumgardner, Multispectral tomographic imaging of the midlatitude aurora, *J. Geophys. Res., 104,* 24,565–24,585, 1999.

Steen, Å, ALIS - an auroral large imaging system in northern Scandinavia, *Eur. Space Agency Spec. Publ., ESA SP-291,* 299–303, 1989.

Steen, Å, and U. Brändström, A multi-station ground-based imaging system at high latitudes, *STEP Int. Newsl., 3,* 11–14, 1993.

Strickland, D.J., R.R. Mejer, J.H. Hecht, and A.B. Christensen, Deducing composition and incident electron spectra from ground-based auroral optical measurements: Theory and model results, *J. Geophys. Res., 94,* 13,527–13,539, 1989.

Syrjäsuo, M.T., et al., Observations of substorm electrodynamics by using the MIRACLE network, in International Conference on Substorms-4, Terra Sci., Tokyo, 1998.

P. Janhunen, Geophysical Research, Finnish Meteorological Institute, P.O.Box. 503, Vuorikatu 15A, FIN-00101, Helsinki, Finland. (Pekka.Janhunen@fmi.fi)

Received June 30, 2000; revised October 20, 2000; accepted October 20, 2000.

This preprint was prepared with AGU's LATEX macros v5.01. File ascinv formatted August 27, 2001.

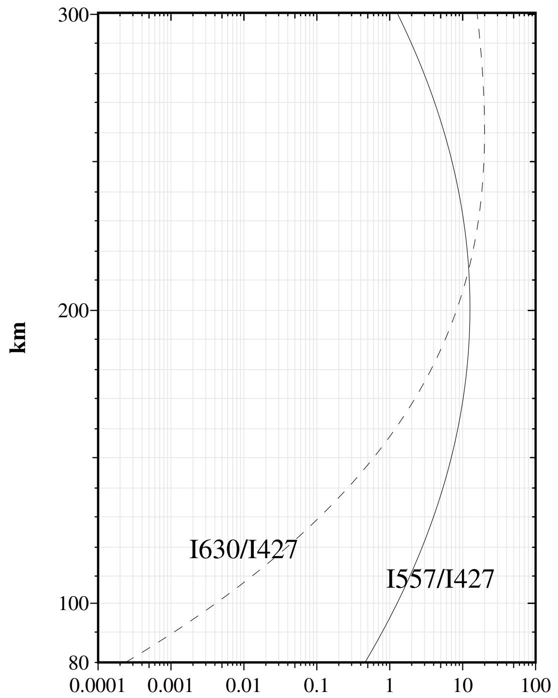

Figure 1. Equations (4) for the intensity ratio profiles based on *Rees and Luckey* [1974].

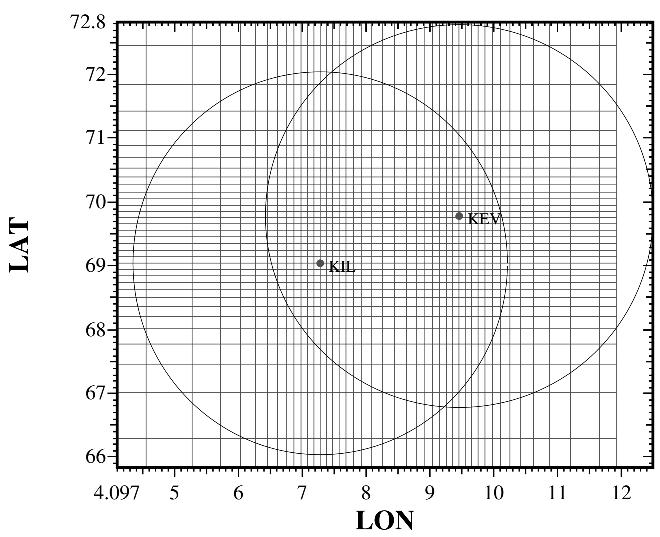

Figure 2. The stretched spatial grid where the reconstruction is carried out in the cases presented in this paper. Only every fourth grid line is shown. Nominal 300-km radius viewing areas of Kilpisjärvi (KIL) and Kevo (KEV) stations are shown by circles.

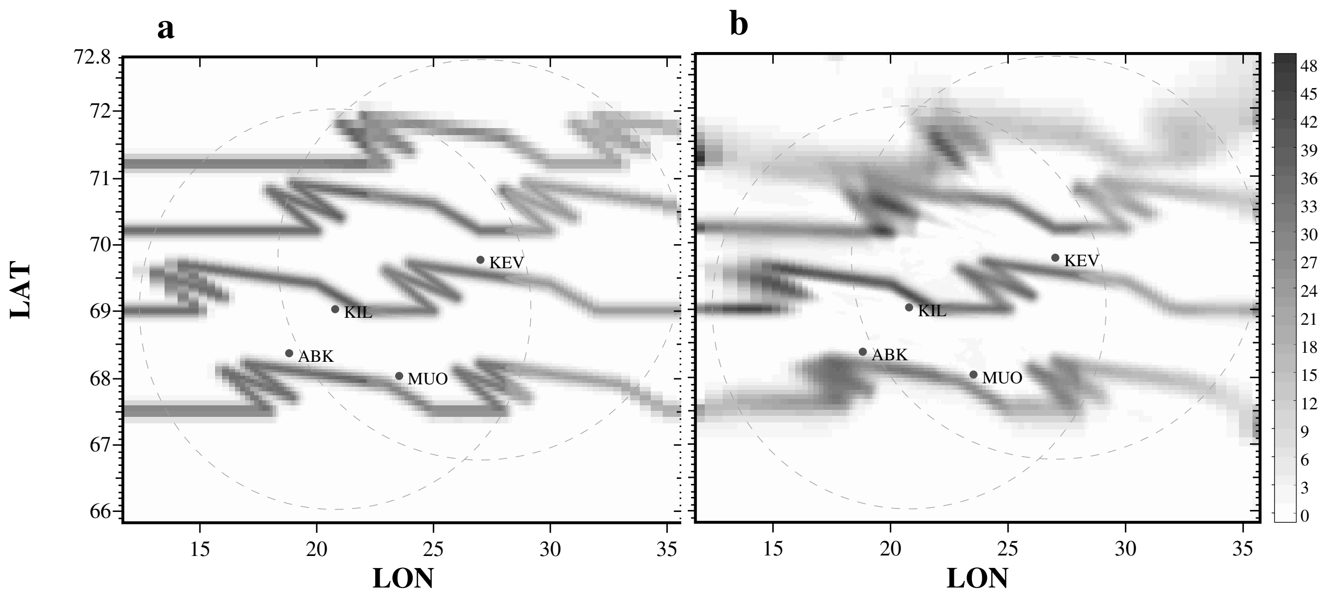

Figure 3. Energy flux in mW m^{-2} of the (a) model and (b) reconstructed auroral arc distribution. Nominal 300-km radius viewing areas of Kilpisjärvi (KIL) and Kevo (KEV) stations are shown by dashed circles. In (a), the total energy flux is 1.23 GW, in (b) it is 1.41 GW.

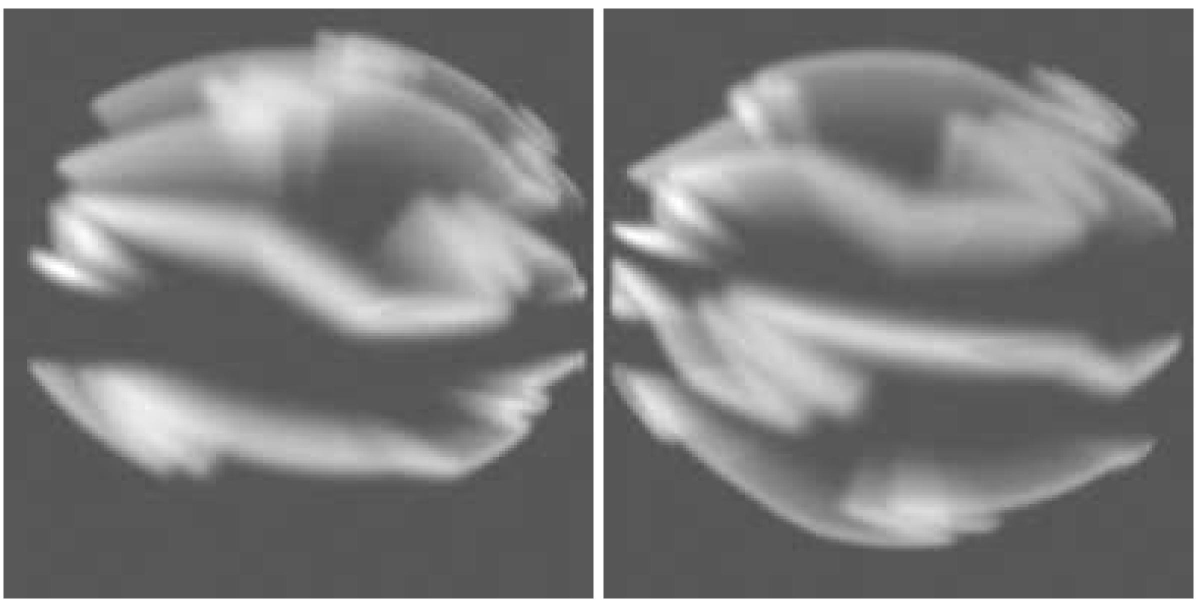

Figure 4. Simulated green (557.7 nm) images for Kilpisjärvi (left) and Kevo (right) all-sky camera stations. North is upward and west is to the left. Gamma-correction value of 3 has been used to amplify faint parts of the arcs. The visible area for each station is approximately a circle with 300 km radius. The stations are about 200 km apart.

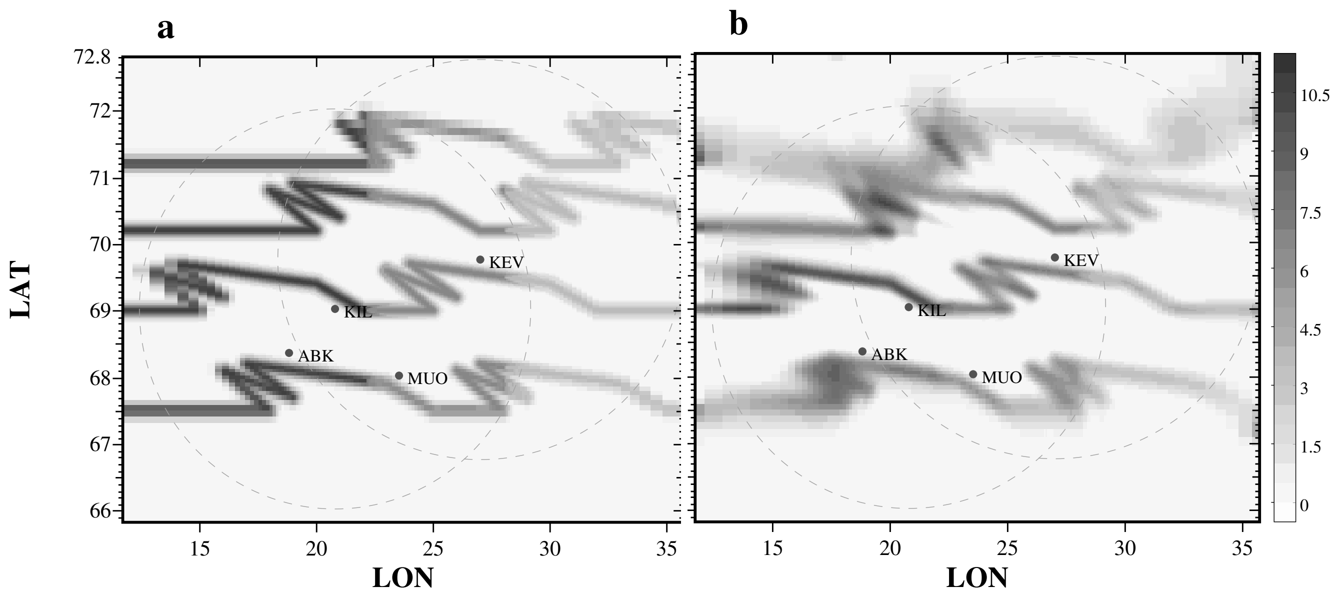

Figure 5. Number flux in μA m^{-2} of the (a) model and (b) reconstruction. The model number flux decreases eastward between 20 and 30 LON. In (a), the total flux in current units is 298 kA, in (b) it is 314 kA.

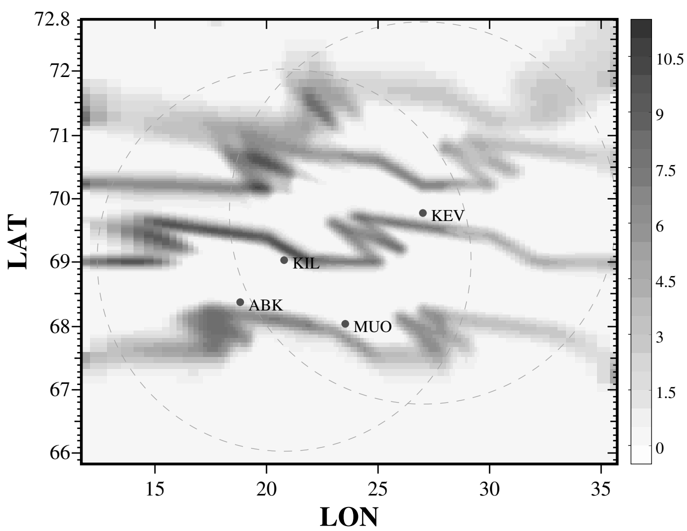

Figure 6. Number flux reconstructed from green images only. The result is almost indistinguishable from Fig. 5b. The total flux in current units is 316 kA.

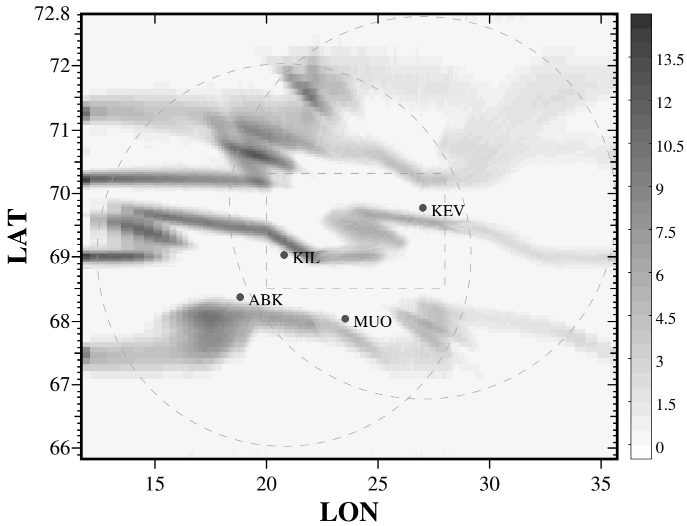

Figure 7. Number flux reconstructed from red images only. The result is not as good as in Fig. 5b

The local area where the 'local' norms shown in Tables 1 and 2 are calculated is shown by a dashed rectangle. The total flux in current units is 293 kA.

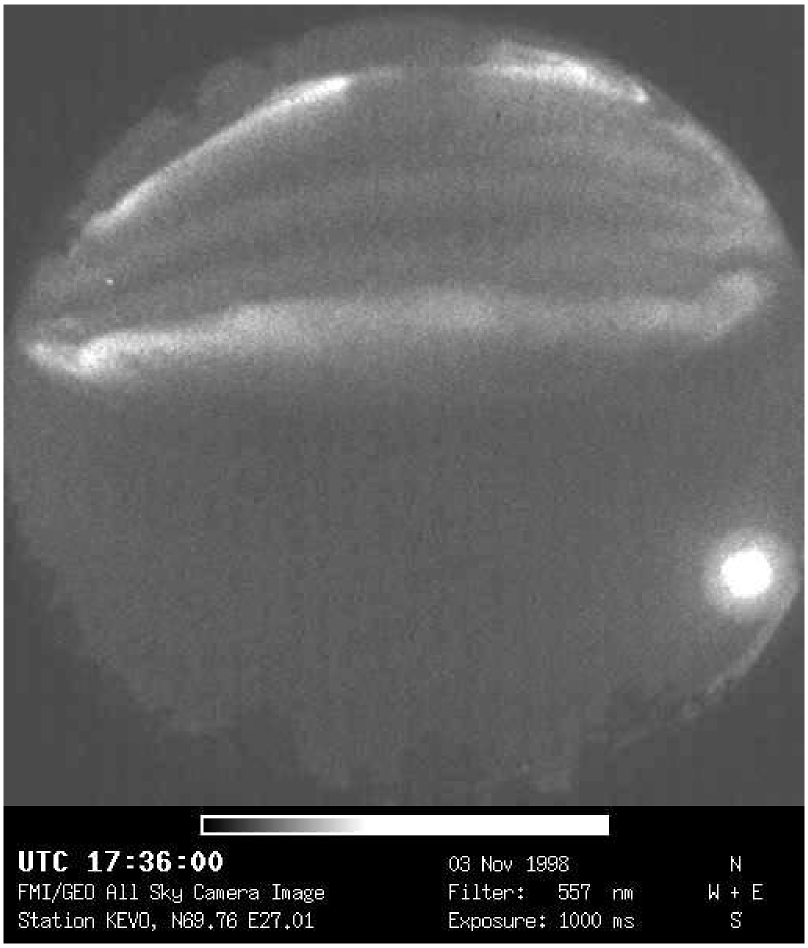

Figure 8. Raw green image at Kevo on Nov 3, 1998, 1736 UT. Intensity map has been edited, and saturated white corresponds to about 12 kR. Moon is in lower right-hand corner. Faint vertical stripes are due to 50 Hz interference from AC power source, an error source that we do not try to remove in this study as it causes no visible trace in reconstructed data

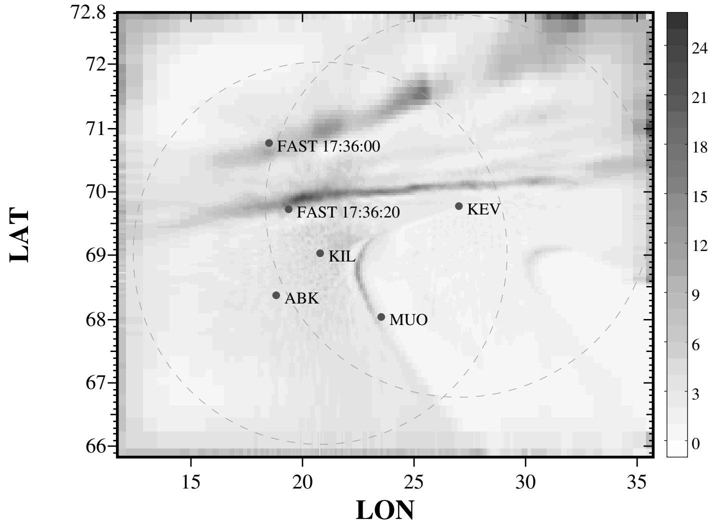

Figure 9. Reconstructed energy flux (mW m^{-2}) for Nov 3, 1998, 1736 UT. Parabolic shapes in lower right part are due to moon removal. FAST position at 1736:00 UT and 1736:20 UT is also shown.

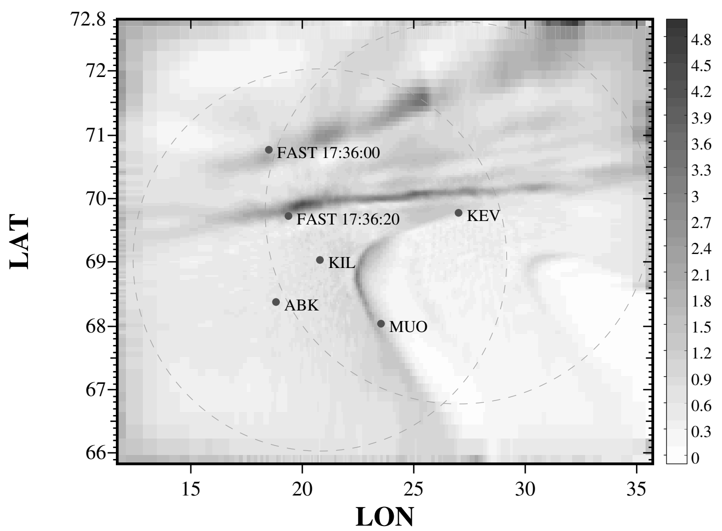

Figure 10. Reconstructed number flux (μA m^{-2}) for Nov 3, 1998, 1736 UT. Parabolic shapes in lower right part are due to moon removal. FAST position at 1736:00 UT and 1736:20 UT is also shown.

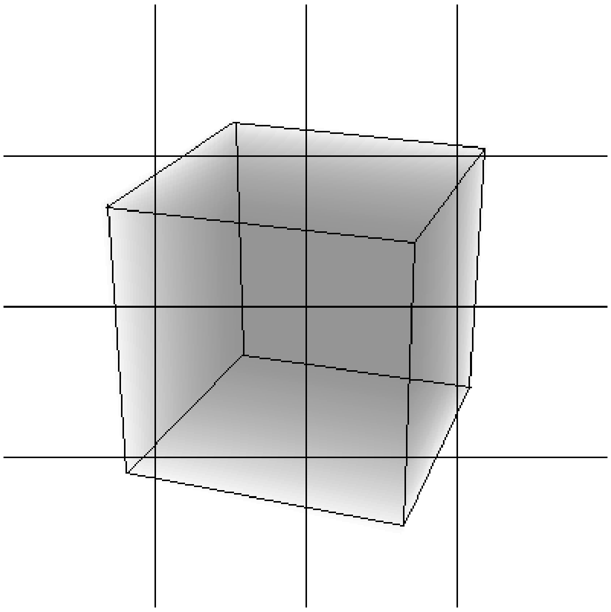

Figure 11. Rendering the voxel to the pixel plane amounts fo averaging the luminosity within each pixel cell. The luminosity (indicated by gray shading) is proportional to the distance that the line-of-sight ray travels inside the quasi-cubic voxel.

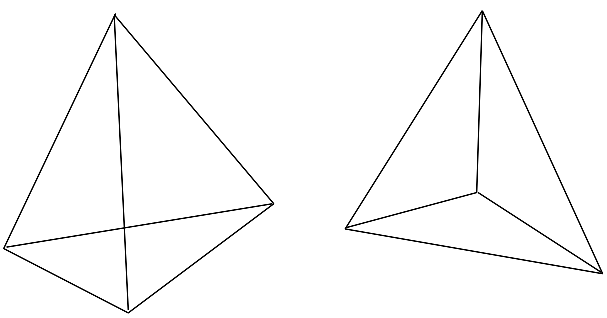

Figure 12. The 2-D projection of a tetrahedron is composed of either 3 or 4 non-overlapping triangles, depending on the viewing angle.

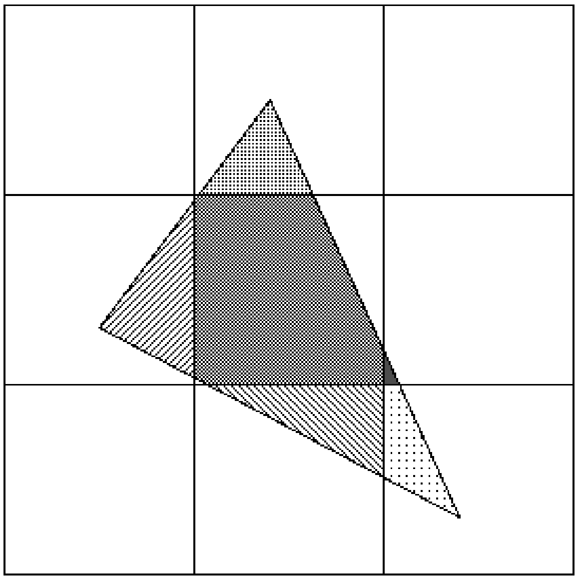

Figure 13. The linear function $u(x,y) = a + bx + cy$ must be integrated over a set which is an intersection of a square pixel and a triangle. In this example, the triangle intersects six pixels.

Table 1. Effect of Wavelength Combination to Relative L^2 Errors[a]

	L^2(eflux)	L^2_{std}(eflux)	L^2_{loc}(eflux)	L^2_{loc}(nflux)
RGB	22.5	25.1	12.1	13.0
RG	28.3	30.2	15.9	14.9
RB	20.6	22.2	10.3	11.5
GB	23.8	25.9	13.0	14.0
R	40.8	40.7	30.7	23.3
G	36.1	33.4	20.2	17.6
B	26.3	25.6	13.9	20.3

[a]Doubled resolution (24 energy, 40 altitude steps) global relative L^2 error norm of the energy flux (L^2(eflux)), standard resolution version of it (L^2_{std}(eflux)), local norm (L^2_{loc}(eflux)), and number flux local norm (L^2_{loc}(nflux)). In global norms, all points closer than 250 km to one of the cameras are included. Local norm is computed over the rectangle shown in Figure 7. Both local norms are also in doubled resolution. All values are percent. R, red; G,green; B,blue.

Table 2. Effect of Wavelength Combination on Integrated Quantity Error[a]

	$\int$ eflux	$\int$ nflux	$\int_{\mathrm{loc}}$ eflux	$\int_{\mathrm{loc}}$ nflux
RGB	8.3	3.7	6.9	6.0
RG	14.1	4.4	11.0	6.0
RB	4.1	-0.7	3.8	1.2
GB	9.1	4.7	7.2	7.0
R	4.1	-1.7	-3.4	0.7
G	23.1	3.3	15.8	4.9
B	5.0	-15.7	4.6	-7.2

[a]Integrated energy and number flux percent errors for different wavelength combinations. Values marked "loc" are integrated over the rectangular box shown in Figure 7, others are integrated over the whole grid. R,red; G,green; B,blue.