# Retrieval algorithms for the EOS Microwave Limb Sounder (MLS)

Nathaniel J. Livesey, W. Van Snyder, William G. Read and Paul A. Wagner

*Abstract*— The 'retrieval' algorithms for the EOS Microwave Limb Sounder (MLS) on the Aura spacecraft, launched on 15 July 2004, are described. These algorithms are used to produce estimates of geophysical parameters such as vertical profiles of atmospheric temperature and composition, ('Level 2' data) from the calibrated MLS observations of microwave limb radiance ('Level 1' data). The MLS algorithms are based on the standard optimal estimation approach. New aspects include the adaptation to a 'two dimensional' system, and an approach to the issues of retrieval 'phasing' and error propagation that differs from that taken for previous similar instruments. Important new aspects of the software that implements these algorithms are also described, along with the algorithm configuration for the 'version 1.5' dataset. Some examples are shown from MLS in-orbit observations.

## I. Introduction

**T**HE EOS Microwave Limb Sounder (MLS) instrument [1], [2] is one of four instruments on the Aura spacecraft that was launched on 15 July 2004. EOS MLS observes thermal microwave emission from the Earth's limb in order to measure the composition and temperature of the atmosphere in the region from ∼8 km to ∼90 km.

This paper describes the 'retrieval' algorithms used in the ground data processing for MLS. The task of these algorithms is to convert calibrated measurements of microwave radiances (known as Level 1B data) into estimates of atmospheric temperature and composition. The approach chosen is the standard 'optimal estimation' method [3], [4]. The retrieval process divides into two main parts: the forward and inverse models. The forward model computes estimates of radiances that would be observed by MLS, were the atmosphere in a given state. The task of the inverse model, as its name implies, is to 'invert' this calculation and deduce an atmospheric state from a given set of MLS radiance observations. Details of the forward models used in the MLS retrieval algorithms are given in other papers in this issue [5], [6]. This paper reviews the general retrieval approach and details the implementation of the inverse model. More detailed information on these calculations are included in the 'MLS Retrieval processes Algorithm Theoretical Basis Document' (ATBD) [7].

Most of the MLS data retrieved by these algorithms describe vertical profiles of geophysical parameters along the measurement track of the instrument. Such products are known as Level 2 data. The task of producing these is called Level 2 processing, and is the main subject of this paper. Some MLS data products have a sufficiently poor signal to noise ratio

All authors are at the Jet Propulsion Laboratory, California Institute of Technology.

that individual vertical profiles are not useful for scientific study due to their poor precision. For these 'noisy' products, some form of averaging is required to produce data with better precision, such as daily or weekly zonal means, or monthly global maps. Products such as these, being on regular latitude and/or longitude grids, are known as 'Level 3' data. While these products can be obtained by averaging together large volumes of Level 2 data, an alternative approach for retrieving them directly is described in section IV.

## II. Fundamentals of retrieval theory

### A. The retrieval problem

The task of the retrieval algorithms is to determine the state of the atmosphere that best matches the observed MLS radiances. This state is represented by the 'state vector' $\mathbf{x}$ of length $n$, which in the MLS case, as is typical, represents vertical profiles of atmospheric temperature and composition, along with selected other parameters described later. A 'measurement vector' $\mathbf{y}$ of length $m$ is constructed describing the radiance observations. A *forward model* is formulated that describes the radiances that MLS would expect to see, were the atmosphere in a state represented by the value of $\mathbf{x}$ according to

$$\hat{\mathbf{y}} = \mathbf{f}(\mathbf{x}). \tag{1}$$

These predicted radiances $\hat{\mathbf{y}}$ are compared with the observed MLS radiances $\mathbf{y}$ and the minimum is sought of a $\chi^2$ statistic defined as

$$\chi^2 = [\mathbf{y} - \mathbf{f}(\mathbf{x})]^{\mathrm{T}} \mathbf{S_y}^{-1} [\mathbf{y} - \mathbf{f}(\mathbf{x})], \tag{2}$$

where $\mathbf{S_y}$ is the matrix describing the error covariance of the measurements. The MLS algorithms invoke the standard Gauss-Newton approach to the minimization which iteratively arrives at a value of $\mathbf{x}$ that minimizes $\chi^2$ by invoking

$$\mathbf{x}^{(q+1)} = \mathbf{x}^{(q)} + \left[\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\mathbf{K}\right]^{-1} \mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1} \left[\mathbf{y} - \mathbf{f}(\mathbf{x}^{(q)})\right], \tag{3}$$

where $q$ is the iteration counter and

$$\mathbf{K} = \left. \frac{\partial \mathbf{f}(\mathbf{x})}{\partial \mathbf{x}} \right|_{\mathbf{x}=\mathbf{x}^{(q)}} \tag{4}$$

is known as the matrix of 'weighting functions' or 'Jacobians'. In most cases the $\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\mathbf{K}$ matrix to be inverted in (3) is singular, indicating that there are aspects of the state vector about which the measurements have yielded no information. This is remedied by adding additional vectors describing 'virtual measurements'. Often (and in the case of the MLS retrieval system), these take the form of an *a priori* estimate

of the state vector $\mathbf{a}$ with covariance $\mathbf{S_a}$. The forward model for these virtual measurements is an identity operation, leading to a weighting function matrix that is simply the $n \times n$ identity matrix. The iteration then becomes

$$\mathbf{x}^{(q+1)} = \mathbf{x}^{(q)} + \left[\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\mathbf{K} + \mathbf{S_a}^{-1}\right]^{-1}$$
$$\left[\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\left(\mathbf{y} - \mathbf{f}(\mathbf{x}^{(q)})\right) + \mathbf{S_a}^{-1}\left(\mathbf{a} - \mathbf{x}^{(q)}\right)\right]. \quad (5)$$

This is similar to the actual iteration used in the MLS case, the next section will detail additional constraints on the smoothness of the retrieved profiles, and the use of a Levenberg-Marquardt parameter to aid convergence.

This algorithm also gives an estimate of the uncertainty in the state vector according to

$$\mathbf{S_x} = \left[\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\mathbf{K} + \mathbf{S_a}^{-1}\right]^{-1}. \quad (6)$$

### B. Retrieval 'phasing' and constrained quantities

Most retrieval algorithms are implemented as a series of phases. Typically an initial retrieval of temperature and pressure information is performed using observations of radiance emitted by a molecule whose abundance is well known (usually $CO_2$ in the infrared and $O_2$ in the microwave). Later phases of the retrieval process use this temperature and pressure information in retrievals of species abundances such as ozone and water vapor. Many previous retrieval approaches, such as the UARS MLS version 5 algorithms [8] fix the temperature and pressure at the previously retrieved values in these later phases. When doing retrievals in this manner, it is often important to budget for the uncertainty in these 'constrained' quantities by inflating the errors on the radiances used to retrieve the species abundances according to

$$\mathbf{S_y} \rightarrow \mathbf{S_y} + \mathbf{K_c}\mathbf{S_c}\mathbf{K_c}^{\mathrm{T}}, \quad (7)$$

where $\mathbf{K_c}$ represents the derivative of the species radiances with respect to the constrained quantities $\mathbf{c}$ retrieved earlier (e.g., temperature and pressure) and $\mathbf{S_c}$ is the uncertainty in these quantities as reported by (6) in the earlier phase.

While for many instruments (including EOS MLS) the original $\mathbf{S_y}$ matrix is diagonal (or can be assumed to be so to a reasonable level of accuracy), the addition of errors on constrained quantities typically results in a fully populated matrix. As this matrix is required to be inverted in (5) this represents a significant increase in computational effort. Indeed, in the case where $m \gg n$ (as with EOS MLS) this computation becomes by far the most intensive aspect of the retrieval computation. Clearly it is desirable to retain a diagonal form for $\mathbf{S_y}$ if possible.

The solution to this issue adopted by the MLS algorithms is to continue to retrieve the previously estimated quantities (temperature, pressure etc.) in the later retrieval phases, including the same radiance information as used to retrieve them in the earlier phases. Continuing to include these quantities in the state vector retains the efficient diagonal form for $\mathbf{S_y}$. This begs the question as to what purpose is served by the earlier phases — if the same parameters are to be retrieved in the later phases, with arguably better quality, as more radiance

information is available. The later phases involve large state and measurement vectors, and therefore significant forward model effort. Reducing the number of iterations required for such expensive phases is clearly an important goal. As the parameters targeted by the early phases (such as temperature and tangent pressure) are typically the most non-linear, performing several iterations of the simpler earlier phases should leave the state vector close to the 'correct' solution for these parameters in the later phases, reducing the number of iterations likely to be required. The task of the early phases can be regarded as that of obtaining a good 'starting point' for the later phases. More details on the specific implementation of this approach are given in section VI.

### C. Characterizing retrieval results

*1) Estimated precision:* Care should always be taken when interpreting results obtained from retrieval algorithms such as those used in the MLS Level 2 processing, and attention should be paid to several important diagnostics. Firstly, the estimated precision of the retrieved products should be considered. In the EOS MLS case (as is typical), this is reported as the square root of the diagonal elements of the solution covariance matrix from (6). Careful comparison should be made between this and the uncertainty quoted for the *a priori* information by the $\mathbf{S_a}$ matrix. If little difference is observed between elements of the estimated precision and the corresponding *a priori* uncertainties, it implies that the instrument has contributed little additional information on these elements of the state vector, and that they should probably not be used in scientific study.

*2) $\chi^2$ statistics:* In addition to the estimated precision, the value of $\chi^2$ obtained at the solution should also be considered. Solutions where $\chi^2 \gg m - n$ should be avoided as a poor fit to the radiances has been achieved. In the MLS case, $\chi^2$ quantities are reported in a normalized manner, being divided by the number of measurements considered. Under these circumstances $\chi^2$ should be about unity at the solution, and results corresponding to significantly larger values are flagged as suspicious.

*3) Averaging kernels:* Another important quantity in diagnosing retrieval performance is the 'Averaging Kernel' matrix, given by

$$\mathbf{A} = \frac{\partial \hat{\mathbf{x}}}{\partial \mathbf{x}} = \left[\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\mathbf{K} + \mathbf{S_a}^{-1}\right]^{-1}\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\mathbf{K}. \quad (8)$$

For the purposes of this equation $\mathbf{x}$ is the unknown true state of the atmosphere (so far as it can be represented in state vector form) and $\hat{\mathbf{x}}$ is the retrieved estimate from the iteration in (5). Each row of this matrix describes how the corresponding element of the retrieved state vector has been influenced by all the elements of the true vector. Each column of the matrix represents the influence of a delta function perturbation of the corresponding element of the true state vector on the retrieved vector.

In the case where the state vector represents vertical profiles of atmospheric temperature and composition, measures such as the full width at half maximum of the averaging kernels

are often used as a measure of the vertical resolution of the measurement system.

*4) Profile representation issues:* In the particular case of the MLS retrieval system, the state vector represents vertical profiles of atmospheric species in a 'piecewise linear' manner. The only exception is water vapor where the representation is piecewise linear in $\log(\text{vmr})$. This contrasts with the common alternative approach which considers the state vector to represent 'layer means'. Special consideration of the MLS interpretation of the state vector needs to be made when comparing retrieved MLS data to other datasets.

Consider a simple MLS state vector $\mathbf{x}_{\text{coarse}}$ representing a single temperature profile on fixed pressure surfaces, and another vector $\mathbf{z}_{\text{fine}}$ describing a temperature measurement on a finer vertical grid (such as from a radiosonde, or GPS). It transpires [5] that to compare the two measurements, in addition to applying the averaging kernels, one needs to transfer the higher resolution profile into the MLS grid by applying the transformation

$$\mathbf{z}_{\text{coarse}} = \left[\eta^{\text{T}}\eta\right]^{-1}\eta^{\text{T}}\mathbf{z}_{\text{fine}}, \tag{9}$$

where the $\eta$ matrix describes a linear interpolation converting the low resolution MLS representation to that of the radiosonde

$$\mathbf{x}_{\text{fine}} = \eta\mathbf{x}_{\text{coarse}}. \tag{10}$$

## III. APPLICATION TO MLS RETRIEVALS

### A. The 'two dimensional' approach to the retrieval problem

Figure 1 shows the viewing geometry for the EOS MLS instrument. It can be seen that the limb observations from consecutive scans cover significantly overlapping regions of the atmosphere. The MLS Level 2 software retrieves individual vertical profiles at approximately the same horizontal spacing as the individual limb scans (known as major frames). The overlap in the limb observations is such that the radiances for each limb scan are influenced by several consecutive retrieved profiles. If the retrieval algorithm is to be accurate, it must take this influence into account.

The approach taken in the EOS MLS Level 2 software is to divide the data into 'chunks', typically consisting of a 15° span of great circle angle's worth of observations (about 10 vertical scans). Retrievals are performed for each of these chunks independently. The state vector consists of $N$ sets of vertical profiles ($N$ temperature profiles, $N$ ozone profiles etc.), with retrievals being performed using measurements from $M$ radiance scans. The results from the individual chunks are joined together to produce a complete set of output for the day. The chunks overlap slightly to reduce 'edge effects' (horizontal homogeneity is assumed in the forward model beyond the first and last profiles, introducing inaccuracies). The spacing of the retrieved profiles is typically chosen to match that of the scans, so that $N \simeq M$. Occasional differences between $N$ and $M$ arise as the relationship between the scans (which are approximately evenly spaced in time) and the retrieved profiles (evenly spaced in great circle angle) wanders due to orbit eccentricity and earth oblateness.

The state vector is broken down into $N$ subvectors $\mathbf{x}_i$ each of length $n$ consisting of a vertical profile of temperature, composition etc., along with possible additional terms described by $\mathbf{x}^\diamond$ which are constant throughout the chunk (such as instrument calibration or spectroscopy terms).

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}^\diamond \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_N \end{bmatrix} \tag{11}$$

The measurement vector can be similarly broken down into $M$ sets of measurements $\mathbf{y}_i$ (each of length $m$) corresponding to the individual scans.

The weighting function matrix $\mathbf{K}$ describes the sensitivity of the $M$ radiance scans to each of the $N$ profile sets. This matrix will have a significant amount of sparsity as, for example, the state of the atmosphere for profile number 10 will have no influence on the radiances observed in scan number 1. This can be described by a sparsity parameter $p$, indicating the number of profiles influencing a single vertical scan. The weighting function matrix $\mathbf{K}$ will be of a 'singly bordered block band diagonal' form as illustrated for the case where $N = M = 6$, $p = 3$.

$$\mathbf{K} = \begin{bmatrix} \times & \times & \times & 0 & 0 & 0 & 0 \\ \times & \times & \times & \times & 0 & 0 & 0 \\ \times & 0 & \times & \times & \times & 0 & 0 \\ \times & 0 & 0 & \times & \times & \times & 0 \\ \times & 0 & 0 & 0 & \times & \times & \times \\ \times & 0 & 0 & 0 & 0 & \times & \times \end{bmatrix}, \tag{12}$$

where 0 denotes a block in the matrix that has zero for all elements, and $\times$ denotes a block with one or more non-zeros. The block rows of $\mathbf{K}$ correspond to the $M$ individual scans, the first block column indicates the sensitivity of those scans to the $\mathbf{x}^\diamond$ term while the other columns indicate the sensitivity to the $\mathbf{x}_i$ subvectors. The value of $p$ is determined by the geometry of the MLS measurements. However, for computational efficiency it can be useful to limit $p$ to a smaller value (such as 5), leading to sparser forms for $\mathbf{K}$. This is achieved by assuming horizontal homogeneity beyond the region $p/2 + 1$ profiles away from each scan.

In the case where $\mathbf{S_y}$ is diagonal, the $\mathbf{K}^{\text{T}}\mathbf{S}_y^{-1}\mathbf{K}$ matrix of (5) has the form

$$\mathbf{K}^{\text{T}}\mathbf{S_y}^{-1}\mathbf{K} = \begin{bmatrix} \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & 0 & 0 & 0 \\ \times & \times & \times & \times & \times & 0 & 0 \\ \times & \times & \times & \times & \times & \times & 0 \\ \times & 0 & \times & \times & \times & \times & \times \\ \times & 0 & 0 & \times & \times & \times & \times \\ \times & 0 & 0 & 0 & \times & \times & \times \end{bmatrix}. \tag{13}$$

This is a 'doubly bordered block band diagonal matrix' having a block-bandwidth of $2p - 1$. The formation of this matrix, which, because $m \gg n$ is the most time consuming aspect of the inverse model, scales as $Np^2n^2m$ (neglecting the $\mathbf{x}^\diamond$ terms). The key point to note is that it scales linearly in $N$,

Fig. 1. The top plot shows the viewing geometry of EOS MLS, which observes limb radiances in the forward direction. The lower plot is an expansion of the boxed region in the upper plot. Here, 12 of the 120 limb ray paths for five scans are shown by the nearly horizontal lines. The loci of the geometrical limb ray tangent points are shown by the thin, angled lines. The kinks in these lines are due to a change of vertical scan rate (the instrument spends more time observing the troposphere and lower stratosphere than the upper regions of the atmosphere in order to improve the information yield from the lower regions). The thicker 'curved' lines show the loci of the refracted (i.e. true) tangent points. The 'vertical' lines represent the location of the retrieved atmospheric profiles.

so, the same amount of CPU time is required to retrieve 10 chunks of 10 profiles each as to retrieve a single 100 profile chunk, excluding the consideration of overlap regions.

Significant efficiency of storage can be gained by recognizing that

$$\mathbf{K}^{\mathrm{T}}\mathbf{S}_{\mathbf{y}}^{-1}\mathbf{K} = \sum_{i=1}^{M} \mathbf{K}_i^{\mathrm{T}}\mathbf{S}_{\mathbf{y}_i}^{-1}\mathbf{K}_i, \qquad (14)$$

where $\mathbf{K}_i$ represents the $i$th of the $M$ block rows of $\mathbf{K}$ and $\mathbf{S}_{\mathbf{y}_i}$ similarly for $\mathbf{S}_{\mathbf{y}}$. Similar simplifications are possible with the $\mathbf{K}^{\mathrm{T}}\mathbf{S}_{\mathbf{y}}^{-1}\left[\mathbf{y} - \hat{\mathbf{f}}(\mathbf{x}^{(q)})\right]$ terms. This means that each iteration can be accomplished needing storage for only one block row of $\mathbf{K}$ rather than the entire matrix, a significant saving. The $\mathbf{K}^{\mathrm{T}}\mathbf{S}_{\mathbf{y}}^{-1}\mathbf{K}$ matrix still needs to be stored in its entirety. However this matrix is significantly smaller than $\mathbf{K}$ as $m \gg n$, and, being sparse, its storage can be efficient.

### B. Major components of the MLS state vector

The 'standard products' for the MLS data processing are vertical profiles of temperature and species abundances for selected molecules on fixed pressure levels. For most of the MLS products, there are six surfaces per decade change in pressure, starting from 1000 hPa, with the grid spacing coarsening to three surfaces per decade for pressures less than 0.1 hPa. Of course the true vertical resolution of the information obtained by the retrieval algorithms is often poorer than the spacing of this 'reporting grid'. In addition to these abundances and temperatures, the geopotential height of a fixed pressure surface (typically 100 hPa) is included. A full vertical profile for geopotential height is not required (or indeed desired). This is because, through consideration of hydrostatic balance, the geopotential height profile can be computed from the temperature profile and the single

geopotential height. Including a full geopotential height profile would lead to redundancy in the state vector and consequent instability in the retrieval system.

Further parameters are required in the state vector in order for the forward model to accurately predict the radiances that would be observed by MLS. The most critical of these are the tangent pressures for the mid-point in time of each individual limb integration period (known as a *minor frame*) for both the GHz and THz vertical scans. These are defined as the atmospheric pressure at the tangent point of the limb ray (taking into account refraction effects). Figure 2 shows an example of how the radiances change as a function of pressure. As most of the MLS radiances signals are strongly determined by pressure broadening, the tangent pressure is clearly an important coordinate by which to define the observations (as opposed to for example tangent height).

### C. Continuum emission and 'baseline'

Most of the MLS products are derived from observations of spectral contrast. The frequency resolution of the instrument is such that it can resolve individual spectral lines of the target molecules over most of the vertical retrieval range. It is the MLS observations of the shape of these lines, as captured in the radiance differences from channel to channel, that yield the retrieved estimate of the geophysical parameters, rather than the absolute radiance values themselves.

Our knowledge and understanding of absolute radiance is generally poorer than that of their spectral shape, due both to instrumental effects (such as thermal emission from the MLS antenna and other reflectors or spectrally broad variations in instrument gain), and forward model limitations (such as uncertainties in continuum spectroscopy). Although the spectral signatures of the target species are largely orthogonal

Fig. 2. Sample MLS radiances as a function of tangent pressure. The different lines represent different MLS spectral channel observations of emission from 118 GHz $O_2$ line. High altitude observations typically correspond to measurements of an optically thin atmosphere, leading to small radiance signatures. As the instrument looks lower down, radiances increase as the atmosphere thickens. Toward the bottom of the scan the radiances 'saturate' (the 'knee' on the radiance curves). Here the atmosphere is sufficiently optically thick that the instrument does not see all the way to the tangent point, instead the bulk of the signal emanates from a region higher up, closer to the instrument. Channel 1 is furthest from the line and so in a region where the atmosphere is optically thin down to around 10 hPa. Channel 13 is centered on the line and so the atmosphere is optically thick in this spectral region even when looking in the lower mesosphere. The dependence of radiance on tangent point pressure is clearly non-linear in nature. These data are taken from the average of the first 100 scans taken by MLS on 30th August 2004.

to these spectrally flat effects, some residual impact can be seen on the MLS products if the spectrally flat terms are not considered in the retrieval algorithm.

The MLS state vector includes terms designed to account for such spectrally flat phenomena, and these are retrieved along with the target species. These terms fall into two categories known as 'extinction' and 'baseline'. Typically only one of these is retrieved, as the signatures of the two are highly correlated and thus hard for the retrieval to appropriately distinguish. Baseline terms are spectrally flat radiances that can be uniformly added to each radiance observation in a radiometer. The software allows these to be described either as a set of vertical profiles as a function of tangent pressure, or as an independent value for each minor frame of radiance observation. This is useful for capturing most spectrally flat artifacts that have an instrumental origin.

Alternatively, the 'extinction' parameter is better suited to capturing spectrally flat features due to problems with modeling of atmospheric transmittance. While these have a spectrally flat impact on a 'local' scale in the atmosphere,

their impact on the MLS radiance signals will not necessarily be spectrally flat. For example, the impact of the extinction at 100 hPa on radiance observations having 100 hPa tangent pressures will be significant in 'transparent' channels, and negligible for channels where the stratosphere is optically thick. Unfortunately, this parameter has a highly non-linear impact on the measurement system, most notably when it is invoked to reduce radiances (i.e., negative mixing ratios of the extinction 'molecule' are required). This can lead to significant instability in the retrieval, particularly in cases where cloud scattering can lead to radiance suppression. Therefore, in the v1.5 algorithms (producing the first publicly released dataset, as described later), only the baseline term is considered in the retrievals.

### D. Minor components of the 'state vector'

There is a distinction drawn in retrieval algorithms between parameters passed to the forward model for which solutions are sought (such as the quantities described above), and other quantities required by the forward model for which *a priori* knowledge is sufficient. The latter quantities include such things as the spacecraft velocity, used for determining Doppler shift effects, and knowledge of the microwave background space radiance. Although not retrieved, such quantities are often loosely referred to as being 'in the state vector' in the MLS algorithms. In addition, the forward model also requires detailed spectroscopic and instrument calibration information [5].

### E. The use of tangent height information

In addition to the MLS radiance observations, the tangent point altitude information obtained from the MLS antenna position encoder and Aura attitude determination system can considered to be indirect measurements of the state vector. Given the state vector description of the atmospheric temperature profile, the estimated pressures at the tangent points, and the geopotential height of a reference pressure surface, a forward model estimate of the Level 1 tangent heights can be constructed based on considerations of refraction and hydrostatic balance [5].

Typically the MLS radiance observations convey information on tangent pressure over a somewhat limited vertical range of limb tangents. Near the top and bottom of the scans the spectral contrast in the MLS radiances is not strongly dependent on tangent point pressure, as shown in figure 2. The inclusion of the tangent point height information in the retrieval system can extend knowledge of tangent point pressure into these regions.

The inclusion of this height information has made the need for an *a priori* constraint for tangent pressure unnecessary, and indeed undesirable. Any such estimate would have to be taken from the same tangent point altitude information given in the measurement vector. Thus, including an *a priori* term for tangent pressure would amount to using the same information twice. This is implemented in the retrieval by setting the appropriate rows and columns of the $\mathbf{S_a}^{-1}$ matrix to zero. Although the absence of these virtual measurements

can in principal lead to the matrix inversions in (5) being singular, it can be shown that the inclusion of the tangent height information guarantees successful inversion.

### F. The Tikhonov smoothing constraint

In retrievals from remote sounding instruments, there is often some trade off to be made between the precision and resolution (typically vertical, but horizontal is also relevant) of the retrieved product. This trade off can be controlled by adding constraints on the 'smoothness' of the results in the retrievals. Often, such as in the UARS MLS version 5 dataset [8], these constraints are implemented by adding off diagonal terms to the estimated error covariance matrix for the *a priori* terms ($\mathbf{S_a}$).

For the EOS MLS retrievals however, an alternative approach is implemented using a second order Tikhonov constraint. This amounts to placing a 'soft' constraint (i.e., additional term in $\chi^2$) on the magnitude of the second derivative (vertically and horizontally) in the retrieved profiles, by introducing virtual measurements with weighting function matrices based on binomial coefficients similar to the $(n-2) \times n$ matrix

$$
\mathbf{K} \simeq \begin{bmatrix}
-\frac{1}{4} & \frac{1}{2} & -\frac{1}{4} & 0 & \dots & 0 & 0 & 0 & 0 \\
0 & -\frac{1}{4} & \frac{1}{2} & -\frac{1}{4} & \dots & 0 & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & 0 & \dots & -\frac{1}{4} & \frac{1}{2} & -\frac{1}{4} & 0 \\
0 & 0 & 0 & 0 & \dots & 0 & -\frac{1}{4} & \frac{1}{2} & -\frac{1}{4}
\end{bmatrix}
\tag{15}
$$

Constraints on higher order derivatives can be achieved by using higher order binomial coefficients. The actual form of the $\mathbf{K}$ matrices used is more complicated, partly because we introduce one set of virtual measurements to describe horizontal smoothing and another for vertical, and partly because we wish to provide a height dependent weighting for the smoothing terms [7].

The horizontal and vertical smoothing terms, scaled as described above are combined in a matrix $\mathbf{R}$, giving a modified Gauss-Newton iteration according to:

$$
\mathbf{x}^{(q+1)} = \mathbf{x}^{(q)} + \left[ \mathbf{S_a}^{-1} + \mathbf{R}^{\mathrm{T}}\mathbf{R} + \sum_i \mathbf{K}_i^{\mathrm{T}}\mathbf{S}_i^{-1}\mathbf{K}_i \right]^{-1}
$$
$$
\left\{ \mathbf{S_a}^{-1}\left[\mathbf{a} - \mathbf{x}^{(q)}\right] + \mathbf{R}^{\mathrm{T}}\mathbf{R}\left[\mathbf{a} - \mathbf{x}^{(q)}\right] + \right.
$$
$$
\left. \sum_i \mathbf{K}_i^{\mathrm{T}}\mathbf{S}_i^{-1}\left[\mathbf{y}_i - \mathbf{f}_i(\mathbf{x}^{(q)})\right] \right\} \tag{16}
$$

The choice of weighting for the smoothing parameter is largely made on scientific grounds according to the desired precision and/or resolution for each product. Typically the smoothing constraint is turned off at the higher reaches of the retrieved profiles, where the inherent resolution of the measurement (i.e., that it would have were no smoothing constraint applied) is already poorer than about 6 km in the vertical.

### G. Numerical stability concerns and scaling

The physical disparity of the MLS state vector (temperatures have typical values of 150–300 K, while BrO mixing ratios have values of order $10^{-11}$ vmr) gives rise to numerical stability concerns for the algorithm. As it transpires, however, the measurement vector does not present a concern in the MLS case, as it contains either radiances in the 0–300 K range or scan residuals of order $\pm 300$ m. In addition the multiplication by $\mathbf{S_y}^{-1}$ effectively scales the measurements into a dimensionless space in any case.

This large dynamic range within the state vector leads to large variations from column to column in the values of $\mathbf{K}$. Despite this, the formation of the $\mathbf{K}^{\mathrm{T}}\mathbf{S_y}^{-1}\mathbf{K}$ matrix is actually numerically stable in our case where $\mathbf{S_y}$ is diagonal. The matrix product is essentially the results of a set of dot products of every column of $\mathbf{K}$ with every other column of $\mathbf{K}$. The terms summed together to form one of these dot products are all related to the same pair of state vector elements, so they all have the same physical units. Thus, no special care need be taken of 'small' numbers that might get lost in the summation, small terms are by their nature insignificant.

However, the matrix inversion in (16) is a concern. The matrix to be inverted contains a wild range of numbers (being related to the square of the state vector), which need to be combined in a series of multiply / addition operations. To alleviate this concern, an element-by-element scaling is applied to the state vector before the inversion such that the diagonal elements of the matrix all become unity. Once the inversion and the rest of the computation is completed, the resulting state vector and covariance matrices are returned to their original units.

### H. Non-linearity and convergence issues

While the Gauss-Newton minimization technique is excellent for linear and moderately non-linear systems, its underlying assumption of linearity can lead it to take inappropriate steps in more seriously non-linear situations. For these situations, the Levenberg-Marquardt [9], [10] stabilization is a common solution. This involves simply adding a matrix $\lambda\mathbf{I}$ to the matrix to be inverted in (16), where $\lambda$ is a scalar chosen each iteration, and $\mathbf{I}$ is the $n \times n$ identity matrix. This term is added to the matrix after the scaling described in the previous subsection is applied. When small values of $\lambda$ are chosen, the step taken is clearly close to that which would be taken by a regular Gauss-Newton iteration. Larger values of $\lambda$ result in a smaller step, closer to the more conservative 'steepest descent' iteration. The value of $\lambda$ chosen for the first iteration is a 'user input', chosen based on the degree of non-linearity anticipated. In later iterations the value of $\lambda$ is chosen based on the progress made to that point.

Before each iterative step is taken in the minimization of $\chi^2$ in (16) with the additional $\lambda\mathbf{I}$ term above, it is possible to compute the value of $\chi^2$ we would find at the destination, were the system truly linear. This is done by substituting the value of $\mathbf{x}$ from (16) into (2) and recognizing that in the linear regime

$$
\mathbf{f}\left(\mathbf{x}^{(q+1)}\right) = \mathbf{f}\left(\mathbf{x}^{(q)}\right) + \mathbf{K}\left[\mathbf{x}^{(q+1)} - \mathbf{x}^{(q)}\right]. \tag{17}
$$

In addition, it is similarly possible to compute the value we would expect $\chi^2$ to take at the minimum, again assuming linearity. For the case of Gauss-Newton iteration (where the Levenberg-Marquardt parameter $\lambda$ is zero) these are of course the same location.

These two pieces of information can be useful in determining iteration strategy. Comparison of the value of $\chi^2$ found at the destination with that which was expected yields information on the amount of non-linearity in the system. This can be considered when choosing the value of $\lambda$ to use in the next iteration. Alternatively, in the case where particularly poor reduction (or even increase) in $\chi^2$ is observed, the algorithm may choose to retreat and select a different $\lambda$ for the current step. More details of this decision process are given in Appendix C of [7].

The value of $\chi^2$ estimated to be at the true minimum is useful in deciding when satisfactory convergence has been obtained. For example, for the first post launch version of the MLS Level 2 software, iterations cease when $\chi^2$ is within 2% of the value that is predicted to exist at the minimum.

## IV. AN ALTERNATIVE APPROACH FOR 'NOISY' PRODUCTS

Some of the molecules EOS MLS is designed to observe have particularly small mixing ratios, and weak emission lines. The corresponding radiance observations thus have poor signal to noise ratios, leading to noisy retrievals. For these products, more useful scientific information can be obtained by considering averaged products, such as daily zonal means, or monthly maps. There are several ways in which to compute such quantities.

One approach is simply to retrieve the products in the same manner as all the others, and then use whatever averaging technique is appropriate afterwards. The disadvantage of this method is that, unless special care is taken, the *a priori* information can significantly bias the results, as it is included in each separate retrieval. This is the approach taken for the version 1.5 of the MLS data processing software (using appropriately large values for the *a priori* uncertainty for the species of interest).

A second approach is to average the radiances from the relevant bands in whatever manner is appropriate, and then to perform retrievals on the averaged radiances. This method has a problem however when the lines of interest are contaminated by strong emission from other, highly variable molecules. This is the case for example with some of the MLS BrO radiance observations which are close to a strong $O_3$ line.

The best approach to this problem is to retrieve the averaged products as a separate task, after the main processing has occured. Rather than using averaged radiances as above, however, the full radiance data set for the relevant band is considered. Consider the iterative retrieval expression given in (16). In the linear (i.e. single iteration, with initial guess

$x = a$) case, this reduces to

$$\mathbf{x} = \mathbf{a} + \left[ \mathbf{S_a}^{-1} + \mathbf{R}^T\mathbf{R} + \sum_i \mathbf{K}_i^T \mathbf{S}_i^{-1} \mathbf{K}_i \right]^{-1}$$
$$\sum_i \mathbf{K}_i^T \mathbf{S}_i^{-1} \left[ \mathbf{y}_i - \mathbf{f}_i\left(\mathbf{a}\right) \right]. \quad (18)$$

Now for the case of the noisy products, take $\mathbf{x}$ to be a specific component of an averaged dataset (e.g., a single profile corresponding to one latitude in a monthly zonal mean retrieval). Consider the measurement vectors $\mathbf{y}_i$ to represent each individual scan in the relevant spectral band that contributes to this component (e.g., all the scans in the latitude range under consideration that month.) The forward models for each scan use the previously retrieved values for the other molecules and parameters that affect the radiance measurements ($O_3$, temperature, tangent pressure etc.) as constrained quantities.

It is possible to take this method further by defining $\mathbf{x}_0$ as the value of the product retrieved by the standard Level 2 processing. Let the vector $\mathbf{b}$ contain all the other aspects of the state retrieved by the Level 2 algorithms (ozone, temperature etc.).

It is clear therefore that

$$\mathbf{y}_i - \mathbf{f}_i\left(\mathbf{a}, \mathbf{b}\right) = \mathbf{y}_i - \left[ \mathbf{f}_i\left(\mathbf{x}_0, \mathbf{b}\right) + \mathbf{K}_i\left(\mathbf{a} - \mathbf{x_0}\right) \right] \quad (19)$$

The retrieval calculation (18) then reduces to:

$$\mathbf{x} = \mathbf{a} + \left[ \mathbf{S_a}^{-1} + \mathbf{R}^T\mathbf{R} + \sum_i \mathbf{K}_i^T \mathbf{S}_i^{-1} \mathbf{K}_i \right]^{-1}$$
$$\left[ \sum_i \mathbf{K}_i^T \mathbf{S}_i^{-1} \left[ \mathbf{y}_i - \mathbf{f}_i\left(\mathbf{x}_0, \mathbf{b}\right) \right] - \mathbf{K}_i^T \mathbf{S}_i^{-1} \mathbf{K}_i \left[ \mathbf{a} - \mathbf{x}_0 \right] \right]. \quad (20)$$

Accordingly, by having the Level 2 software gather appropriate sums of the $\mathbf{K}_i^T \mathbf{S}_i^{-1} \mathbf{K}_i$ matrices and the $\mathbf{K}_i^T \mathbf{S}_i^{-1} \left[ \mathbf{y} - \mathbf{f}_i\left(\mathbf{x}_0\right) \right]$ vectors, the 'noisy products' algorithm need not invoke any forward model calculations. All that is required is that values of the above matrix and vector are collated together appropriately, including the correction term $\mathbf{K}_i^T \mathbf{S}_i^{-1} \mathbf{K}_i \left[ \mathbf{a} - \mathbf{x}_0 \right]$ where $\mathbf{x}_0$ is taken from the standard Level 2 product, and the final state computed as the result of (20).

Extending this to allow it to follow from a full two dimensional Level 2 calculation is achieved by defining $\mathbf{x}$ to be the mean of all or several profiles in the chunk and collapsing together the appropriate block columns of the $\mathbf{K}_i$ matrices. The issue of errors on constrained quantities has not been considered for this problem. It is possible that these should be considered, and a non-diagonal form for $\mathbf{S}_i$ used. This would make it harder for the 'noisy products' algorithm to avoid invoking forward model calculations. The issue will be investigated when these algorithms are developed.

## V. IMPLEMENTATION IN SOFTWARE

The algorithms described in section III have all been implemented in a single 'MLS Level 2 software program' in the Fortran 95 programing language, chosen because of its great suitability to handling the complex matrix and vector entities involved. The program was designed to be as flexible

as possible, and is controlled by a 'Level 2 Configuration File' (L2CF) that is effectively a high level programming language.

The L2CF controls all aspects of the software, defining the contents of state and measurement vectors, defining the configurations of the various forward models available, reading appropriate *a priori*, spectroscopic and calibration data, performing retrievals, doing forward model runs for simulations or off-line analysis, post-processing results, computing diagnostics and outputting results in appropriate files. In 'production' mode, the software operates in a parallel form, with one instance of the program acting as a 'master', coordinating the work of multiple 'slave' instances on a cluster of computers, each computing the results for individual 'chunks' of data as described in section III-A. The software can also be run 'online' for single chunks or simple one dimensional retrievals of individual profiles. The on-line mode is enhanced by the ability of the software to communicate with a separate program (written in the IDL language from Research Systems Inc.) that presents a graphical interface into the algorithm, allowing the user to see the current state and measurement vectors and many diagnostics, and to monitor the progress of retrievals (effectively acting as a graphical 'debugger').

In addition to the ability to do conventional retrieval calculations, producing Level 2 Geophysical Product (L2GP) output files from input L1B radiance data, the software can also produce L1B files of simulated radiances based on a state vector formed from a set of L2GP files taken as input. This capability was used extensively prior to launch for generating radiance fields corresponding to known atmospheric states. Developing a different program to perform that essential task would entail the duplication of all the relevant code for initializing the forward model and constructing the state vector. Combining both the retrieval and simulation tasks in a single piece of software makes it far easier to ensure that identical forward model algorithms and parameters are used for both tasks, and dramatically reduces the complexity of the code maintenance effort.

In addition to these tasks, the Level 2 software has proved flexible enough to have been used for a large variety of other tasks, from those as mundane as translating MLS Level 2 data files from an older version of their format (based on HDF-EOS version 4) to a newer one (HDF-EOS version 5), to as complex as doing a retrieval using monthly zonal mean radiances, or pre-computing tables to be used in the linear forward model [5].

## VI. RETRIEVAL APPROACH FOR VERSION 1.5

To this point, this paper has described the EOS MLS retrieval algorithms in a fairly general sense. In this section we will describe the particular configuration of the software used to generate the version 1.5 (v1.5 hereafter) EOS MLS Level 2 data, the first publicly released MLS dataset. This section should be regarded as a 'snapshot' of the current configuration, as future versions of the data processing algorithms are planned.

### A. 'Standard' products, and 'Diagnostic' products

The geophysical products from the MLS retrieval algorithms can be divided into two categories. Each of the 'standard' MLS geophysical products are output in separate files with a daily granularity. The 'standard' products are the science team's 'best' estimate of that product from the MLS observations. Typically they are taken from the MLS observation of that species in a particular frequency region. For example, in the version 1.5 processing, the standard product for ozone is the ozone as retrieved from the 240 GHz radiance information.

In later versions, it is intended that the standard products for many species will be formed from some optimal combination of the information obtained from all the relevant MLS radiances (ideally from one 'comprehensive' retrieval phase). However, changes in instrument configuration, such as the temporary power down of one radiometer, or the changing of the MLS switch network [1] will impact these products, as they change the whole MLS measurement system. Such retrieval schemes have not yet been implemented, pending improvements in computer resources, and in understanding of any systematic differences seen between the estimates obtained in the different MLS radiance signals.

For analyses such as trend studies, it may be more appropriate to consider the 'diagnostic' MLS products. These are simply the products retrieved from each retrieval phase independently. Being based on only a single radiometer (plus the 118 GHz signal used in the retrieval temperature and tangent pressure in all phases), these may be less sensitive to any changes in instrument configuration. As an example, consider nitrous oxide, for which there are two diagnostic products: N2O-190, and N2O-640, corresponding to retrievals using 190 and 640 GHz radiance observations respectively. Of these N2O-640 is generally considered superior, and so the standard product for $N_2O$ is currently simply a copy of the N2O-640 product. Table I details how each of the standard products is derived from the diagnostic products in v1.5.

### B. The Core, Core+Rn approach

The phasing approach described in section II-B has been implemented in what is known as the 'Core, Core+Rn' approach in the v1.5 algorithms. In the 'Core' phase of the retrievals (actually three separate phases), retrieved estimates are obtained for the tangent pressure, temperature, and upper tropospheric humidity aspects of the state vector. These are obtained from the R1A 118 GHz observations of emission from $O_2$ (mainly for temperature and pressure) and selected channels from the R2 190 GHz observations (mainly for upper tropospheric water vapor). This is followed by phases such as 'Core+R2A' and 'Core+R2B' where, in addition to temperature and pressure, other species such as water vapor, ozone and nitric acid are retrieved.

Section II-B described how the 'cumulative' approach to retrieval phasing was preferable to performing constrained quantity error propagation. However, in pre-launch testing, it was decided that this approach was not universally appropriate. In particular, it was found that retrievals including tropospheric water, a species whose impact on the MLS radiances is very

TABLE I

THE ORIGIN OF EACH OF THE 'STANDARD PRODUCTS' FROM V1.5

| Product | Origin |
|---|---|
| BrO | Core+R4 (640 GHz) |
| $CH_3CN$ | Core+R2B (190 GHz) |
| ClO | Core+R4 (640 GHz) |
| CO | Core+R3 (240 GHz) |
| $H_2O$ | Core+R2A (190 GHz) |
| HCl | Core+R4 (640 GHz) |
| HCN | Core+R2B (190 GHz) |
| $HNO_3$ | Core+R3 (240 GHz) for 10 hPa or greater, Core+R2 (190 GHz) for lesser pressures |
| $HO_2$ | Core+R4 (640 GHz) |
| HOCl | Core+R4 (640 GHz) |
| $N_2O$ | Core+R4 (640 GHz) |
| $O_3$ | Core+R3 (240 GHz) |
| OH | Core+R5 (2.5 THz) |
| Temperature | Core for pressures of 1 hPa or greater, Core+R2 for lesser pressures. |

nonlinear, were prone to instability. Accordingly, once an appropriate estimate for this is obtained (in the Core+R2A phase), it is constrained in later phases. No propagation of the errors associated with constraining water vapor is performed in the later phases. Investigation showed that the error involved in neglecting this propagation is insignificant, mainly because the impact of upper tropospheric water vapor on most of the MLS radiances is fairly spectrally flat and largely orthogonal to the signatures of other species.

Other minor deviations from the strict implementation of the planned scheme have been chosen. In particular, in the 'Core+R4B' phase, targeting $N_2O$ from the 640 GHz radiometer, the ozone abundance, rather than being retrieved is constrained to previously retrieved values with no error propagation. Knowledge of the ozone signature in the $N_2O$ spectral region is somewhat uncertain, and this approach was found to produce generally preferable results for $N_2O$.

### C. Flagging the presence of clouds in v1.5

While microwave signals are far less sensitive to the presence of clouds in the atmosphere than shorter wavelength signals, very thick clouds can have an impact large enough to affect MLS measurements of gas phase composition. The Level 2 software needs to be able to identify such radiances and deal with them appropriately. In addition, the cloud signatures can be used to measure cloud properties [11].

The mechanism whereby clouds affect the MLS radiances is mainly scattering of microwave radiation. The details of the impact vary both from channel to channel and as a function of limb ray tangent height. Optically thin observations in a channel (those that are less than about 50% of the radiance at saturation) can be affected by both scattering and emission from clouds which lead to an unexpected enhancement in radiance. Radiance observations lower down, where the radiances are close to or beyond saturation, can be affected by scattering from clouds, which leads to suppression of the radiance signal. While the cloud scattering and emission effects are spectrally broad in nature, the impact of these effects on the MLS limb observations are frequency dependent and become more severe as the radiances get closer to saturation. For example, a cloud at 100 hPa will affect channels where the atmosphere is optically thin enough to allow MLS see down to that level. However, channels that do not see down to 100 hPa will be unaffected.

The MLS Level 2 software therefore takes steps to avoid considering radiances that are thought to be strongly affected by cloud effects, and/or report an increased uncertainty on them. These impacted radiances are identified by comparing the MLS radiance observations in selected optically thin channels (most suited to cloud detection) in each radiometer with those predicted from forward model calculations. The gas phase retrievals are instructed to ignore or downplay radiances where large differences between observation and model are observed. This activity is performed at three distinct points during the v1.5 algorithms. The first two are during the Core group of phases where the 'current best' temperature and tangent pressure information from MLS are used in a forward model, in conjunction with a water vapor profile representing 110% relative humidity with respect to ice, to obtain a reasonable upper limit for clear sky radiances. Finally, after the Core+R2B phase, once the best information on water vapor has been obtained from MLS, a new forward model estimate is computed using the MLS retrieved water observations in order to re-flag cloud contaminated radiances for all radiometers to be used in later phases.

Appropriate thresholds for cloud contamination have been empirically determined for each radiometer, based on simulations. The thresholds have been chosen to maximize the use of radiances consistent with not severely affecting the retrieved species. Flagged radiances are either not used (in the case of the 240 GHz and some 118 GHz radiances), or weighted less in the retrievals (190 GHz radiances). To give one example, the software ignores MLS 240 GHz radiance observations in cases where the observed radiances are more than 5 K greater or 30 K smaller than those predicted by the forward model for a selected optically thin channel in the 240 GHz radiometer.

TABLE II

THE PHASES THAT FORM THE V1.5 RETRIEVAL ALGORITHMS.

| Phase | Target species[a] | Measurements | Comment |
|---|---|---|---|
| Init-pTan | T, pTan (GHz), GPH | R1A (118 GHz) | Very quick forward model |
| Update-pTan | T, pTan (GHz), GPH | R1A (118 GHz) | Slower more accurate model |
| Init-UTH | U.T. $H_2O$ | R2 (190 GHz) | |
| Core+R2A | T, pTan (GHz), GPH, $H_2O$, $N_2O$, $HNO_3$, $O_3$ | R1A (118 GHz), R2 (190 GHz) | $H_2O$ retrieved down to 316 hPa, other species to 215 or 100 hPa, uses costly full forward model. |
| Core+R2B | T, pTan (GHz), GPH, $H_2O$, $HNO_3$, ClO, $O_3$, HCN, $CH_3CN$ | R1A (118 GHz), R2 (190 GHz) | Main $H_2O$ radiances excluded, products retrieved down to between 316 and 100 hPa. Fast linear forward model used. |
| High-Cloud | Baseline terms as proxy for cloud contamination | R2 (190 GHz), R3 (240 GHz) | Used for flagging clouds in Core+R3 and later phases, in addition to forming basis for cloud water products. |
| Core+R3 | T, pTan (GHz), GPH, $O_3$, CO, $HNO_3$ | R1A (118 GHz), R3 (240 GHz) | Retrievals down to 316 hPa |
| Core+R4A | T, pTan (GHz), GPH, ClO, BrO, $HO_2$, HOCl, HCl, $O_3$, $HNO_3$, $CH_3CN$ | R1A (118 GHz), R4 (640 GHz) | Retrievals down to 147 hPa |
| Core+R4B | T, pTan (GHz), GPH, $N_2O$ | R1A (118 GHz), R4 (640 GHz) | Retrievals down to 147 hPa |
| Core+R5 | T, pTan (GHz, THz), GPH, OH, $O_3$ | R1A (118 GHz), R5H and R5V (2.5 THz) | Retrievals down to 68 hPa |

[a]Tangent pressure and Geopotential height have been abbreviated to pTan (GHz/THz) and GPH respectively. Minor state vector components such as 'baseline' have been omitted.

### D. A note on correlated noise

The MLS radiances measured by the 640 GHz radiometer show an unexpected signature of correlated noise in their observations of limb radiance. The noise signature is largely consistent from channel to channel across the whole of the 640 GHz band. This behavior is inconsistent with the pre-launch understanding of the behavior of the 640 GHz receiver formed from ground based calibration studies. The receiver has had this property since construction, but it was not recognized in the pre-launch calibration data because of differences between the calibration and limb observation regimes. The 2.5 THz receivers show a smaller manifestation of the same phenomenon.

The Level 1B algorithms report two estimates of the noise on individual radiances. The first noise is the spectrally varying component for each channel in each radiometer. The second component is that which is spectrally flat. The observed behavior in the 640 GHz receiver results in a larger than anticipated spectrally flat component to the noise. Prior to launch, it had been understood that the spectrally flat component would be sufficiently small that it could be essentially ignored in Level 2. The baseline terms, represented by a fairly coarse vertical profile on pressure surfaces, would be sufficient to account for unexplained radiances. However, the observed large amounts of spectrally flat noise which vary rapidly from minor frame to minor frame dictate a switch to having an independent baseline for each minor frame. It was decided for consistency to switch to this representation not just for the 640 GHz, but all the receivers.

### E. Some selected results from v1.5

A full discussion of the results from the v1.5 algorithms and the discussion of its performance are beyond the scope of this paper. A detailed discussion of data quality will be supplied to those wishing to use MLS data [12]. Figure 3 gives as an example a pair of retrieved stratospheric $N_2O$ profiles on 10th January 2005, one at the equator, one at high northern latitudes in the winter polar vortex. Other examples of retrieved MLS products are given in [13].

### F. Plans for future versions

The main goal for the next version of the Level 2 algorithm is to retrieve a water vapor product with a higher vertical resolution of twelve surfaces per decade change in pressure in the tropopause region, compared to the typical six per decade. Such a product, while having greater vertical resolution than the 'standard' water vapor product, will necessarily have a poorer precision.

Significant improvements in forward model efficiency (with a slight penalty in accuracy) are anticipated in future versions, through the optional use of a 'pre-frequency averaging' approximation. This will allow for the use of the 'full' non-linear forward model in cases where up to now its use has involved a prohibitive amount of computational effort (such as for most of the 640 GHz radiances).

Increase in knowledge of the spectroscopic parameters influencing the MLS radiance signals will undoubtedly lead to improvements in many MLS products. While most of the spectral lines 'targetted' by MLS are well characterized, some

Fig. 3.   An example of some results from the v1.5 algorithms for measurements made on 10th January 2005. The case shown is two retrievals of $N_2O$ abundance from the 640 GHz radiances. The left hand plots show retrieved $N_2O$ with the error bars indicating the estimated precision. The right hand plots show the measured (symbols) radiances in channels 4, 10, 11, 12 and 13 of the 640 GHz $N_2O$ band (red, green, blue, orange, magenta respectively). The widths of the symbols denote the reported noise (spectrally varying component) on the measurements. The solid lines show the fitted radiances estimated by the forward model, corresponding to the retrieved $N_2O$ profile. These are shown only for the radiances used in the retrieval, hence the absence of the solid lines below ∼17 km.

measurements are influenced by emission from other lines whose parameters are currently more poorly known. In addition, further (probably minor) improvements to our knowledge of the calibration of the MLS instrument are anticipated which will have an impact on the MLS data quality.

## VII. SUMMARY

This paper has reviewed the retrieval algorithms implemented for the EOS MLS instrument. In conjunction with its companion papers, this should serve as a useful reference for those wishing to better understand the EOS MLS measurement system, and as possible guidance for those implementing retrieval schemes for other instruments. The algorithms are

performing well on incoming MLS data, and yielding results whose quality is broadly in line with pre-launch expectations.

## APPENDIX
### THE CALCULATION OF COLUMN ABUNDANCES

In addition to retrieving profiles of atmospheric temperature and composition, the MLS Level 2 software also computes column abundances above the tropopause for most species. Tropopause pressure is derived from the MLS temperature profiles, according to the standard World Meteorological Organization definition, adapted appropriately for use with pressure rather than altitude coordinates.

Given a retrieved abundance profile $f_i$ for a linear representation basis (see section II-C.4) on a set of pressure surfaces

$P_i$, with $i = 1 \ldots n$, the column abundance (number per cm$^2$) above a pressure level $P^\star$ is given by

$$
\mathcal{C} = \frac{1}{mg} \left\{ f_n P_a \right.
$$
$$
+ \sum_{i=1}^{n-1} \frac{f_i}{\Delta \zeta_i} \left[ P_b \left( \zeta_{i+1} - \zeta_b \right) + \frac{P_{i+1} - P_b}{\ln 10} \right]
$$
$$
+ \sum_{i=2}^{n} \frac{f_i}{\Delta \zeta_{i-1}} \left[ P_c \left( \zeta_c - \zeta_{i-1} \right) + \frac{P_c - P_i}{\ln 10} - P_i \Delta \zeta_{i-1} \right]
$$
$$
\left. + f_1 \left( P_d - P_1 \right) \right\}, \quad (21)
$$

where $\zeta_x = -\log_{10}(P_x)$,

$$
\zeta_a = \max \left( \zeta_n, \zeta^\star \right),
$$
$$
\zeta_b = \min \left[ \max \left( \zeta_i, \zeta^\star \right), \zeta_{i+1} \right],
$$
$$
\zeta_c = \min \left[ \max \left( \zeta_{i-1}, \zeta^\star \right), \zeta_i \right],
$$
$$
\zeta_d = \min \left( \zeta_1, \zeta^\star \right),
$$

$\Delta \zeta_i = \zeta_{i+1} - \zeta_i$, $m$ is the molecular mass of dry air and $g$ is a nominal value of the earth's gravitational field. Using $\frac{1}{mg} = 0.789$ DU ppmv$^{-1}$ hPa$^{-1}$ gives the column in milli-atm-cm (Dobson units) for pressure in hPa and concentrations $f_i$ in ppmv.

## ACKNOWLEDGEMENTS

## REFERENCES

[1] J. W. Waters, L. Froidevaux, R. S. Harwood, R. F. Jarnot, H. M. Pickett, W. G. Read, P. H. Siegel, R. E. Cofield, M. J. Filipiak, D. A. Flower, J. R. Holden, G. K. Lau, N. J. Livesey, G. L. Manney, H. C. Pumphrey, M. L. Santee, D. L. Wu, D. T. Cuddy, R. R. Lay, M. S. Loo, V. S. Perun, M. J. Schwartz, P. C. Stek, R. P. Thurstans, K. M. Chandra, M. C. Chavez, G.-S. Chen, M. A. Boyles, B. V. Chudasama, R. Dodge, R. A. Fuller, M. A. Girard, J. H. Jiang, Y. Jiang, B. W. Knosp, R. C. LaBelle, J. C. Lam, K. A. Lee, D. Miller, J. E. Oswald, N. C. Patel, D. M. Pukala, O. Quintero, D. M. Scaff, W. V. Snyder, M. C. Tope, P. A. Wagner, and M. J. Walch, "The Earth Observing System Microwave Limb Sounder (EOS MLS) on the Aura satellite," *IEEE Trans. Geosci. Remote Sens.*, vol. This issue.

[2] J. W. Waters, W. G. Read, L. Froidevaux, R. F. Jarnot, R. E. Cofield, D. A. Flower, G. K. Lau, H. M. Pickett, M. L. Santee, D. L. Wu, M. A. Boyles, J. R. Burke, R. R. Lay, M. S. Loo, N. J. Livesey, T. A. Lungu, G. L. Manney, L. L. Nakamura, V. S. Perun, B. P. Ridenoure, Z. Shippony, P. H. Siegel, and R. P. Thurstans, "The UARS and EOS Microwave Limb Sounder (MLS) experiments," *J. Atmos. Sci.*, vol. 56, pp. 194–217, 1999.

[3] C. D. Rodgers, "Retrieval of atmospheric temperature and composition from remote measurements of thermal radiation," *Rev. Geophys.*, vol. 14, no. 4, pp. 609–624, 1976.

[4] ——, *Inverse methods for atmospheric science, theory and practice.* World Scientific, 2000.

[5] W. G. Read, Z. Shippony, M. Schwartz, and W. V. Synder, "MLS forward model for gas-phase retrievals in the troposphere and stratosphere," *IEEE Trans. Geosci. Remote Sens.*, vol. This issue.

[6] M. J. Schwartz, W. G. Read, and W. V. Snyder, "MLS forward model for polarized mesospheric signals," *IEEE Trans. Geosci. Remote Sens.*, vol. This issue.

[7] N. J. Livesey and W. V. Snyder, "EOS MLS retrieval processes algorithm theoretical basis," Jet Propulsion Laboratory, Tech. Rep., 2004, D-16159, available on the MLS web site http://mls.jpl.nasa.gov.

[8] N. J. Livesey, W. G. Read, L. Froidevaux, J. Waters, H. Pumphrey, D. Wu, M. Santee, Z. Shippony, and R. Jarnot, "The UARS Microwave Limb Sounder version 5 dataset: Theory, characterization and validation," *J. Geophys. Res.*, vol. 108, no. D13, p. 4378, 2003, doi:10.1029/2002JD002273.

[9] K. Levenberg, "A method for the solution of certain nonlinear problems in least squares," *Quart. Appl. Math.*, vol. 2, p. 164, 1944.

[10] D. W. Marquardt, "An algorithm for least-squares estimation of nonlinear parameters," *J. Soc. Ind. App. Math.*, vol. 11, pp. 431–441, 1963.

[11] D. L. Wu, J. H. Jiang, and C. P. Davis, "Aura MLS cloud ice measurements and cloudy-sky radiative transfer model," *IEEE Trans. Geosci. Remote Sens.*, vol. This issue.

[12] N. J. Livesey, "Data quality document for the EOS MLS version 1.5 level 2 dataset," Jet Propulsion Laboratory, Tech. Rep., 2005.

[13] L. Froidevaux *et al.*, "Initial validation of EOS MLS observations," *IEEE Trans. Geosci. Remote Sens.*, vol. This issue.

**Nathaniel Livesey** was born in 1970 in Cambridge England. He received his B.A. (hons) in Physics from the University of Oxford, England, in 1991, and his D.Phil in Atmospheric Physics, also from the University of Oxford, in 1995. From 1996 to the present he has been a member of the Microwave Atmospheric Science team at the Jet Propulsion Laboratory, and is a co-Investigator for the Earth Observing System Microwave Limb Sounder.