

Radiometric and Spectral Performance and Calibration of the GHz Bands of EOS MLS

Robert F. Jarnot, Vincent S. Perun and Michael J. Schwartz

Abstract—This paper describes radiometric performance and pre-launch radiometric and spectral calibrations of the GHz component of the Microwave Limb Sounder (MLS) experiment on NASA's Aura spacecraft. Estimated systematic scaling uncertainties (3σ) on limb radiances are $\sim 0.5\%$ from radiometric calibration, with an additional 0.55% (118 GHz bands) to 0.8% (640 GHz bands) equivalent scaling error from spectral calibration uncertainty. Operational noise performance is consistent with pre-launch expectations, and in-orbit measurements to date indicate no changes in instrument noise performance, and no observable calibration drifts. Spectral baseline has remained stable to 20 mK since launch. Refinements to calibrations based on in-flight data are discussed. Level 1 radiometric calibration algorithms are also described.

Index Terms— Calibration, limb sounding, microwave, submillimeter wave, GHz.

I. INTRODUCTION

EOS MLS on Aura is a follow-on to the successful MLS instrument on the Upper Atmosphere Research Satellite (UARS) [1]. It measures thermal emission from Earth's atmospheric limb to infer vertical profiles of minor constituents of interest to improving understanding atmospheric chemistry and dynamics in regions spanning the upper troposphere to mesosphere. A more detailed description of the instrument, its operation, mission and measurements, is given in a companion article in this journal [2]. In this paper we describe in-orbit radiometric performance, the radiometric and spectral calibration of the GHz portion of Aura MLS, and the algorithms used to convert raw data into calibrated radiances. Companion papers in this journal describe the field-of-view (FOV) calibrations of the GHz radiometers [3], and the THz portion of MLS together with all aspects of its calibration [4].

Compared to its predecessor, the current instrument implements twice as many radiometers, each with substantially greater measurement bandwidth and spanning a considerably larger overall frequency range. The number of measurement channels, a useful indicator of the magnitude of the calibration effort, is increased by an order of magnitude, and the instrument data rate by two orders of magnitude. The calibration techniques described below evolved from those employed on UARS MLS [5], but with significant enhancements in data quality, and measurement efficiency to support the greater than an order of magnitude increase in the quantity of calibration data collected.

Manuscript received March 2005. This work was done at the Jet Propulsion Laboratory, California Institute of Technology, under contract to the U.S. National Aeronautics and Space Administration.

The authors are with the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109-8099. R. Jarnot may be contacted at 818-354-5204 or Robert.F.Jarnot@jpl.nasa.gov.

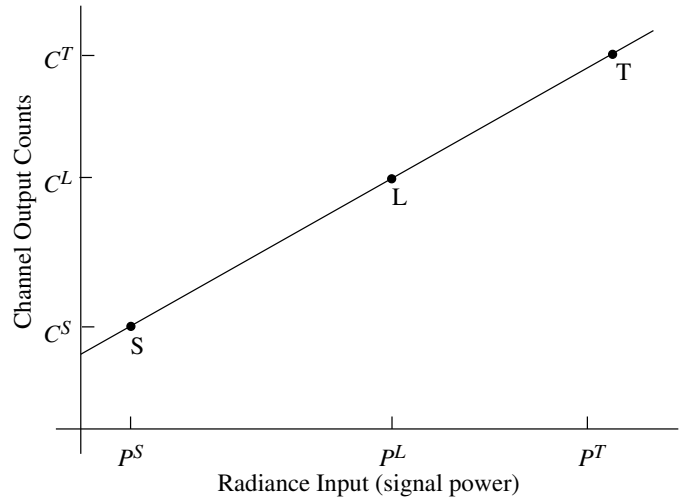


Fig. 1. This figure indicates the linear relationship between measured radiance P and channel output C . S and T indicate views to cold space and an ambient calibration target, and L indicates a view to Earth's limb via the main antenna.

A. Measurement Sequence and Radiance Calibration Process

Aura MLS radiance observations consist of a continuous, regular sequence of data integrations, each of 161 ms nominal duration, with 5 ms dead time between them. Each 166 ms integration/dead time period is referred to as a minor frame (MIF). MIFs are grouped into major frames (MAF), each of which contains 148 MIFs (24.7 s), the first 120 of which are used to view the atmospheric limb. The remaining MIFs are used for viewing cold space (12 MIFs) and one of two radiometric calibration targets (6 MIFs). A switching mirror common to all GHz radiometers is used to direct the receiver FOVs to one of 4 ports (limb, space, ambient or cooled calibration target). 10 MIFs of each MAF are used to step and settle the switching mirror, and measurements are made in a total-power (non-chopped) mode. The ambient target is near room temperature, and the other target is cooled passively by $\sim 40^\circ\text{C}$. Only one target is viewed routinely as part of the radiometric calibration process, but both targets were viewed by special switching sequences implemented during the instrument activation period in order to verify linearity. Details of atmospheric scanning are provided in [2].

The Aura MLS radiance calibration process and algorithms are very similar to the UARS MLS implementation. Since all GHz radiometers share a common switching mirror, all channels are radiometrically calibrated simultaneously. We can represent the output of a single measurement channel by the linear transfer function shown in Figure 1.

II. RADIOMETRIC CALIBRATION ALGORITHMS

The software which converts raw data into calibrated radiances is called Level 1 data processing. Radiometric calibration algorithms are described in detail in [6]. The switching mirror FOV's to the space, target and limb ports are partially restricted by the apertures of the corresponding ports of the cavity in which it is located. Representing the transmission through port X by η^{MX} , the output of a filter channel, C^X , is given by:

$$\begin{aligned} C_i^L &= g_i \left(\eta_r^{ML} \dot{P}_i^A + (1 - \eta_r^{ML}) \dot{P}_r^{BL} \right) + C_i^O, \\ C_i^T &= g_i \left(\eta_r^{MT} \dot{P}_r^T + (1 - \eta_r^{MT}) \dot{P}_r^{BT} \right) + C_i^O, \\ C_i^S &= g_i \left(\eta_r^{MS} \dot{P}_r^S + (1 - \eta_r^{MS}) \dot{P}_r^{BS} \right) + C_i^O. \end{aligned} \quad (1)$$

where C_i^X are the outputs of channel i for the three switching mirror positions X , and g_i is the radiometric gain of channel i (expressed in counts per Kelvin of signal brightness). C_i^O represents the offset counts from the digitizer of channel i , set by a combination of the electronics offsets in the measurement system, and the noise contribution of the receiver electronics (system temperature). η^{MX} is determined for each radiometer as part of the FOV calibration activity, and \dot{P}^{BT} is determined from engineering telemetry. Subscript r indicates the radiometer dependence of calibration parameters and radiances.

The first step in the radiance calibration process is to estimate radiometric gain, \hat{g}_i , at the time of each limb-viewing MIF:

$$\hat{g}_i = \frac{(\hat{C}^T - \hat{C}^S)}{\left(\eta^{MT} \epsilon \dot{P}^T - \eta^{MS} \dot{P}^S - (1 - \eta^{MT}) \dot{P}^{BT} + (1 - \eta^{MS}) \dot{P}^{BS} \right)} \quad (2)$$

where subscripts i and r on the right hand side of the equation have been dropped for clarity. \dot{P}^{BT} and \dot{P}^{BS} are radiation offsets from the baffles obscuring the edges of the FOVs from the switching mirror to the target and space views. ϵ is the emissivity of the calibration target T , discussed later. Radiometric calibrations are performed on MAF-sized chunks of limb data. The estimates of target and space view signals at the time of each limb view in the MAF being calibrated, \hat{C}^T and \hat{C}^S , are obtained by quadratic interpolation of the space and target view data from 6 groups of calibration views, 3 on either side of the group of limb being processed. Note that the *difference* between interpolated target and space counts is used in this expression, removing the need to estimate C_i^O in Equation 1, but requiring that any drifts in this offset be well fit by the quadratic interpolators over timescales of ~ 3 minutes. Adequate stability is obtained through appropriate choice of electronic components, careful design of electronic subsystems, and good thermal stability. Prior experience with UARS MLS led us to design for temperature stability better than 0.01 K per 100 s in the signal processing electronics. This performance is achieved in orbit, and drift characteristics have been acceptably small in both magnitude and derivative.

Once radiometric gain has been determined, the atmospheric radiance entering the limb port from the external telescope,

\dot{P}_i^L , is determined from

$$\begin{aligned} \dot{P}_i^L &= \frac{1}{\eta_r^{ML}} \left(\frac{C_i^L - \hat{C}_i^S(L)}{\hat{g}_i(L)} + \eta_r^{MS} \dot{P}_r^S - (1 - \eta_r^{ML}) \dot{P}_r^{BL} \right. \\ &\quad \left. + (1 - \eta_r^{MS}) \dot{P}_r^{BS} \right) \end{aligned} \quad (3)$$

where $\hat{C}_i^S(L)$ is the interpolated value of the space counts at the time of the limb view¹. In addition to the atmospheric radiance, there are significant additional radiance components reaching the switching mirror limb port, in particular emission from the antenna surfaces, and scattered/diffracted radiation from the spacecraft and Earth. These components are corrected in the limb port radiances estimated by the Level 1 software, and discussed further in [3].

1) *Antenna Baseline Correction:* Views to space via the space and limb ports (with the antenna viewing well above the Earth's limb) differ by up to ~ 10 K in offset, and ~ 0.5 K spectral structure. The large spectrally-flat component arises from thermal emission, scattering and diffraction at antenna reflectors. The smaller spectrally-varying components are a natural consequence of the diffraction-limited optics and frequency dependence of the illumination and diffraction sidelobes. Level 1 software makes a first order correction for these radiance differences (see accompanying FOV paper), but significant residual errors remain. Two approaches are used in Level 1 to address these remaining residuals:

- 1) Spectrally-averaged (DC) offsets are determined on a MAF-by-MAF basis independently for each band by averaging all limb observations above a tangent height of 85 km in channels which have negligible predicted atmospheric radiances.
- 2) Spectrally-varying (AC) offsets are periodically measured by viewing Earth's limb at tangent points well above those at which atmospheric signals are present.

A DC offset estimate is provided in the Level 1 radiance files for each measurement band, and for each of the 12 individual ~ 500 MHz wide filter channels in the 118 and 240 GHz radiometers. To accommodate drifts, the reported value for each MAF is the mean of the measured offset in the current and both adjacent MAFs.

Special measurement sequences run during instrument activation confirmed the stability and lack of scan dependence of the AC baseline, down to the few mK level. The AC baseline is reported in the Level 1 radiance file as a static quantity for all filter channels based on the data from these extended high tangent point scans.

2) *Calibration Targets:* The emissivities of the calibration targets were measured prior to launch using a set of swept RF sources and broadband receivers. The return signal from the targets was compared to that from a reference reflector over the full range of angles and orientations seen by the radiometers. This provided verification that no harmful diffractive effects were created by the grooved surface structure of the microwave

¹By convention, radiances are computed in units of temperature so that the measure converges to the absolute temperature in the long wave (Rayleigh-Jeans) limit.

absorbers. Return loss was better than 30 dB over the full frequency range spanned by the GHz receivers. During routine operation only one of the two GHz calibration targets is used as a radiometric reference every MAF, and we currently use the one attached to the radiative cooler. This is because of a pre-launch prediction that the surface of the other target would see a brief indirect reflection of sunlight from the antenna structure as the instrument came into solar illumination. Data from an instrument activation test during which the switching mirror stared at each target for just over 2 orbits indicate that any impact of this illumination transient on target radiance is too small to be observed. Use of the cooled target gives rise to a small difference between the temperature of the target reported in engineering telemetry, and the surface of the target (which ‘sees’ an environment $\sim 40\text{K}$ warmer). Tests run during instrument activation in which the switching mirror dwelled at the cold space and both target reference ports indicated that the radiative temperature of the cooled target was between 0.5 and 0.6 K higher than that reported in telemetry. The range arises from the varying skin depth of the target surface absorbing material, the largest discrepancy arising in the highest frequency radiometer. This temperature ‘error’ is corrected in the Level 1 processing algorithms.

The finite reflection from the surface of the calibration targets give rise to small standing waves, of order a few $\frac{1}{10}$ ’s of a Kelvin in each band. We have no way of measuring these separately for each target in flight, leading to an uncorrectable radiometric gain error component of up to $\sim 0.2\%$.

3) *Precision*: The noise on a time series of data from an individual channel, ΔT , is commonly represented by:

$$\Delta T = T_{total} \sqrt{\frac{1}{B\tau} + \left(\frac{\Delta G}{G}\right)^2} \quad (4)$$

Where T_{total} is the combination of receiver system temperature and scene radiance, B is the channel pre-detection noise bandwidth, and τ is the post-detection integration time of each measurement. The $\frac{1}{B\tau}$ component is commonly referred to as radiometer noise [7], and is uncorrelated between channels. The gain variation term indicated by $\frac{\Delta G}{G}$ has been measured to be essentially 100% correlated between all channels of a given radiometer in Aura MLS. A requirement on the Level 1 data processing software is that in addition to determining calibrated radiance at the limb port, estimates of both spectrally-averaged and spectrally-varying noise of calibrated radiances shall be provided. The spectrally-varying noise component of uncalibrated data is provided by routinely determining T_{sys} , and applying the radiometer equation (Equation 4 with the gain variation component omitted). T_{sys} is a stable quantity, with orbital variation of $< 1\%$, and is readily estimated in Level 1 processing from the expression:

$$T_{sys} = \frac{(C^S - C^Z)}{g} - \dot{P}(2.7) \quad (5)$$

where C^S and C^Z are space and ‘zero’ counts, and $\dot{P}(2.7)$ is the radiance of cold space. C^S is measured 12 times per MAF, and C^Z , the spectrometer output with no RF signal input, is known from pre-launch calibrations. UARS MLS experience is

that C^Z remained stable even after several years of continuous on-orbit operation. We have verified typical drifts in C^Z in EOS MLS filter channels of ~ 10 counts from pre-launch values after more than 8 months in orbit, which is negligible in comparison to the typical operating levels of $\sim 30,000$ counts.

The quadratic interpolator used in Level 1 processing is implemented so that it provides an estimate of the spectrally-varying uncertainty with each interpolate. The uncertainties on interpolated space and target views are used to determine the additional precision errors introduced in the calibration process. The Level 1 estimate of spectrally-varying noise on each measurement, $\Delta \dot{P}^L$, is determined from the equation:

$$\Delta \dot{P}^L = \sqrt{\frac{(T_{sys} + T_{limb})^2}{B\tau} + (\Delta R)^2 + \left(T_{limb} \times \frac{\Delta g}{g}\right)^2} \quad (6)$$

ΔR and Δg are the noise on the interpolated space references and channel gains respectively. Δg is determined from the precisions of the interpolated space and target signals at the time of each limb measurement.

For signals close to balance (i.e., scene radiances close to those of cold space) the total-power radiometric calibration process increases the noise on an individual measurement by $\sim 4\%$ from the radiometer equation noise. For out-of-balance signals there is an additional component of uncertainty arising from the finite precision of the gain estimate, \hat{g} .

4) *Spectrally-averaged noise*: The spectrally-averaged component of noise in Equation 4 can be readily determined for each measurement channel if ΔT and T_{total} are known. As a routine performance diagnostic from Level 1, the observed ratio of $(\Delta T)^2$ to that predicted by the radiometer equation is computed for all space view sequences. This diagnostic, called space view χ^2 , proved valuable on UARS MLS where its time dependence enabled us to rapidly diagnose a periodic noise problem arising from worn bearings in the switching mirror when spacecraft battery voltage was near its orbital minimum. We were able to implement changes to switching mirror operation which circumvented this problem. For EOS MLS we report χ^2 in the radiance file, allowing Level 2 (retrieval) software to trivially determine the spectrally-averaged noise component on each calibrated radiance. In-orbit χ^2 and T_{sys} data are presented later in this paper.

χ^2 has a typical value close to unity for signals whose noise contribution is predominantly spectrally-varying, the case for the narrower ($< 24\text{MHz}$) channels at the center of each spectrometer. For the outermost (96 MHz bandwidth) channels at the edges of each 25-channel filterbank χ^2 increase to ~ 1.2 , and as high as ~ 2 in the 500 MHz Wide Filter channels. The observed χ^2 in R4 (640 GHz) range between 2.5 and 5 in the 96 MHz filter channels, indicating a much higher than expected level of spectrally-averaged noise in this radiometer. An unusual characteristic of the spectrally-averaged noise in R4 is that its magnitude *decreases* with increasing scene radiance, contrary to expectations, and to the behavior of all other radiometers in this instrument. The nature of the pre-launch characterization environment was such that sensitivity was measured when viewing ambient temperature scenes, under which condition the observed noise

TABLE I

IN-ORBIT SPECTRALLY-VARYING NOISE IN EOS MLS GHZ BANDS FOR $\frac{1}{6}$ S INTEGRATION TIME. ΔI_{\min} DATA ARE FROM LONG (MULTI-ORBIT) CHARACTERIZATION TESTS.

Band	T_{sys}	ΔI_6 MHz	ΔI_{96} MHz	ΔI_{\min}
R1A	1200 K	.2 K	0.35 K	0.02 K
R1B	1400 K	1.3 K	0.40 K	0.02 K
R1BWF	1200 K	$\Delta I_{500} = 0.23$ K		0.02 K
R2	1000 K	3.2 K	0.9 K	0.02 K
R3	1400 K	4 K	1 K	0.02 K
R3WF	1400 K	$\Delta I_{500} = 0.4$ K		0.02 K
R4	4200 K	10 K	3 K	0.1 K

in R4 met expectations. This led to the situation that the excess spectrally-averaged noise was not recognized before launch. We do not yet know its source.

A. Radiometric Performance and Accuracy

Table I indicates the measured in-orbit spectrally-varying noise for all GHZ bands expressed as T_{sys} (double sideband except for the R1 radiometers) and ΔI (time-series noise) in 6, 96 and 500 MHz (WF) filter channels for single MIFs of data. ΔI_{\min} is the level down to which signals are required to integrate, and it was determined that this was met by differencing sets of spectra measured when viewing space via the main antenna for an extended period. We expect the long data integrations such as are required for measuring BrO to indicate better ΔI_{\min} performance than indicated here. All spectrally-varying noise requirements are met.

Spectrally-averaged noise requirements on individual data integrations are that it shall be less than $4 \times 10^{-4} \times T_{\text{sys}}$ in all GHZ channels. For R1 through R3 (118 to 240 GHz radiometers) the measured spectrally-averaged noise is approximately $2 \times 10^{-4} \times T_{\text{sys}}$, meeting requirements. Even with the anomalous spectrally-averaged noise behavior in R4, the observed noise contribution is $\sim 3.5 \times 10^{-4} \times T_{\text{sys}}$, meeting requirements.

Requirements on radiometric calibration accuracy are:

- (1) The systematic uncertainty (i.e., not including the estimated contributions of random noise) in the *absolute value* of the atmospheric/Earth radiances measured through each spectral channel shall be less than 3 K (at the 90% confidence level).
- (2) The systematic uncertainty (i.e., not including the estimated contributions of random noise) in the *spectrally-varying component* of the atmospheric/Earth radiances, measured from one channel or filter to another throughout a given radiometer, shall be less than 1% or $\Delta I_{\min}/3$ where ΔI_{\min} for each spectral region is given in Table IV of [2].

ΔI_{\min} requirements for the GHZ bands range between 0.02 K and 0.1 K, and the measured performance shown in Table I indicates that requirements are met. The primary sources of systematic radiance errors are summarized in Table II. These have been expressed as a percentage of the limb port radiance reported by Level 1 software. Errors for signals close to

TABLE II

SYSTEMATIC ERROR BUDGET FOR LIMB PORT RADIANCES, EXPRESSED AS A PERCENTAGE OF REPORTED LIMB RADIANCE.

Calibration Target temperature uncertainty	0.1%
Standing Waves	0.1%
Switching Mirror Baffle uncertainty	0.15%
End-to-end linearity	0.1%
Total	0.45%

radiometric balance have negligible standing wave error contribution because of the baseline corrections at Level 1. Linearity errors are a maximum for signals mid-way between cold space and calibration target in radiance, and negligible for signals close to either reference. The estimates in Table II should be considered ‘worst-case,’ and it is clear that requirements are met.

B. Level 1 Daily File Sizes

Level 1 software processes data in daily (GMT) increments. The input data consists of 240 individual files of spacecraft and raw instrument data totaling ~ 1.2 GB. Daily output data files are listed in Table III. The radiance files include spectrally-varying and spectrally-averaged noise estimates for each calibrated radiance, as well as AC and DC baseline estimates. Radiance data are flagged as ‘bad’ when the cold space or limb views are contaminated by the moon in their FOV, or when instrument configuration is changing in a way that precludes calibration of the limb data (e.g., when gain levels are being tuned). The contents of the engineering and diagnostics files are plotted and inspected daily to monitor instrument health and safety, and also trended for determination of any long-term drifts in instrument characteristics.

III. RADIOMETRIC CALIBRATION

Pre-launch radiometric calibrations consisted of the determination of calibration target emissivities and switching mirror baffle transmissions, and verification of linearity and noise performance. Calibration target measurements have been discussed in Section II, and switching mirror baffle properties in [3]. Figure 2 illustrates the setup used to verify linearity of the GHZ radiometers. Focussing mirrors M1 and M2 direct the view from the GHZ radiometer space port to external

TABLE III

LEVEL 1 DAILY OUTPUT FILE SIZES. EXCEPT WHERE NOTED, FILES ARE HDF5 FORMATTED.

GHZ filter channel radiances	1.4 GB
GHZ autocorrelator radiances	1.8 GB
THz radiances	434 MB
Orbit/Attitude data	292 MB
GHZ diagnostics	20 MB
THz diagnostics	1 MB
Engineering data (binary)	95 MB
Log files (text)	11.5 MB
Metadata (text, 4 files)	35 kB
Daily Total	4 GB

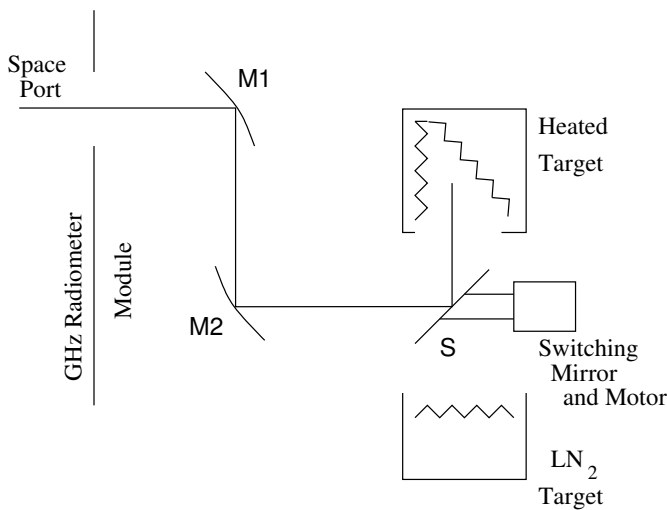


Fig. 2. This schematic illustrates the configuration to verify end-to-end linearity of MLS signal chains.

switching Mirror S which in turn directs the receiver FOVs towards either a liquid nitrogen (LN_2) cooled or a heated calibration target. The GHZ module switching mirror additionally provides a view to either of the calibration targets in the radiometer assembly. The external heated target can be varied in temperature between ambient and 100°C , and its controller provides good stability. This target is a spare UARS MLS calibration load very similar to the internal Aura MLS ones mounted in an insulated enclosure, with strip heaters on the aluminum backing plates of the absorbing surfaces, and PRTs to both control and monitor target temperature. The LN_2 target consisted of a circular panel of standard microwave absorbing material at the bottom of a stainless steel dewar. Standing waves in the views to the external targets were typically of order a few $\frac{1}{10}$'s of a Kelvin.

Figure 3 is a picture of the linearity measurement setup. The instrument GHZ module is mounted with the space port pointing down. Focussing mirror M1 is mounted below the space port, and directs the receiver FOVs towards M2. External switching mirror S is shown directing the FOV into the LN_2 dewar. The UARS heated target aperture, A, is clearly visible to the left of mirror S. The additional elements labeled in the figure are part of the high-resolution spectral sweep setup. Linearity measurements consisted of the 3-point observing sequence ambient/heated/ LN_2 targets. This sequence was synchronized to the instrument MAF, with the GHZ switching mirror providing the view to the internal ambient target, and the external one providing the views to the heated and LN_2 targets. By synchronizing the measurement sequence is this way it was possible to use standard Level 1 software to process the heated target views into calibrated radiances, using the ambient target view as the primary reference, and the LN_2 view as the radiometric gain reference. This type of flexibility was designed into the Level 1 processing software from the outset, and allowed early testing of the software with instrument data. In addition, this method of processing the data automatically takes care of instrumental gain and offset drifts. An example of the raw data from 1 MAF during a




Fig. 3. Radiometric linearity and high resolution spectral sweep calibration setup. See text for additional details.

linearity test run is shown in Figure 4. The three views to the different radiance targets are clearly visible. Figure 5 shows the same data for entire test, and the effects of gain and offset drifts are clearly visible. This measurement cycle was repeated with the heated target elevated in temperature in $\sim 15^\circ\text{C}$ steps, with ~ 45 minutes of stable data taken at each temperature plateau. The same data are shown in Figure 6 after processing to Level 1, clearly illustrating the effective drift compensation of this software.

1) *Integral and Differential Linearity*: Figure 7 shows the individual channel linearity data for R1A (the primary 118 GHz radiometer). In the upper panel the Level 1 radiance at each target plateau is shown. The lower panel shows the same data, but with the mean radiance (weighted by channel bandwidth) subtracted. Standing waves are clearly visible, and present at similar levels in the data from most bands. Of most importance is the *difference* between the plots at different scene temperatures, since this a direct indication of differential linearity, or channel-to-channel breakup for out-of-balance

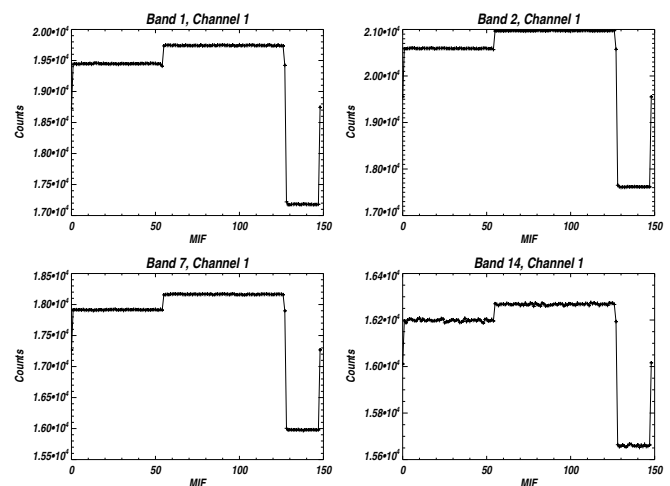


Fig. 4. Raw data for 1 MAF taken during a linearity measurement test. The four panels are data from individual 96 MHz filter channels in each of the GHZ radiometers.

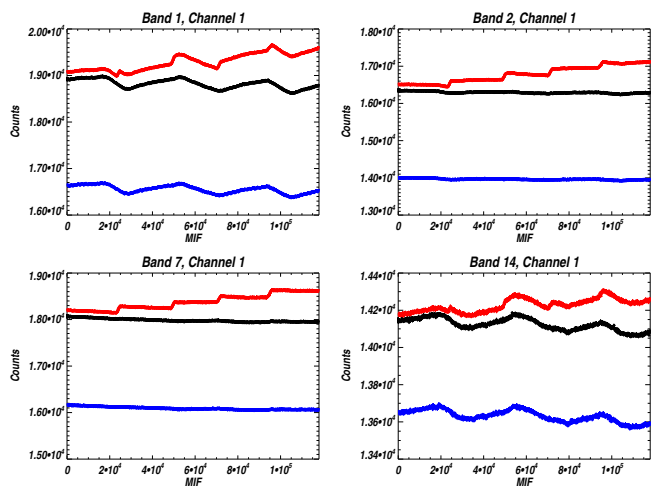


Fig. 5. Data as Figure 4 for an entire linearity measurement run. The blue/black/red data are for the LN₂/ambient/heated target views respectively, each plotted point being a MAF-average. Note the cyclic drifts in these data.

scenes. The baseline correction at Level 1 effectively removes the effect of standing waves for measurements near balance. With UARS MLS we have measured clean H₂O₂ spectra with channel-to-channel differences of 0.001 K, a factor 10⁶ below T_{sys}, as shown on Figure 12 of [9]. Standing waves similar in magnitude to those in the lower panel of Figure 7 exist in the space-target differences for EOS MLS, and are not corrected in the current software, leading to the 0.1% standing wave systematic error contribution in Table II.

Figure 8 plots the band-averaged radiance for R1A versus scene temperature. Departures from a linear transfer function indicate imperfect integral linearity. All GHz bands indicate similar performance to that shown in the figure.

We place an upper limit of 0.1% on linearity error in this instrument based upon pre-launch linearity measurements discussed above, and from in-flight data obtained from special switching sequences viewing both GHz calibration targets. The in-flight measurements extend the dynamic range of the

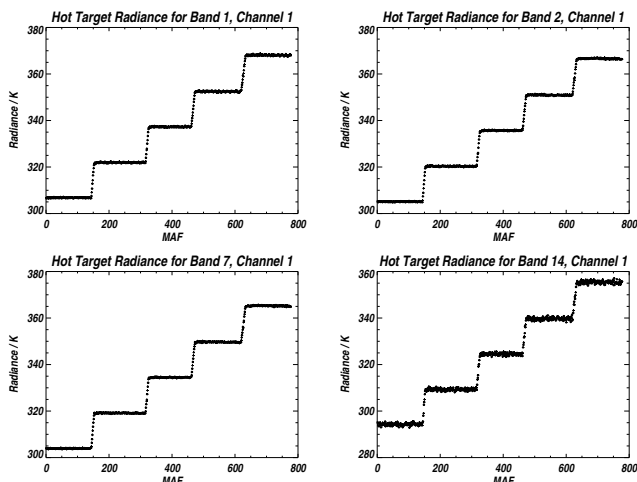


Fig. 6. The data of Figure 5 processed by Level 1 software into heated target radiances is shown here. Note the removal of all instrumental drift-related artifacts.

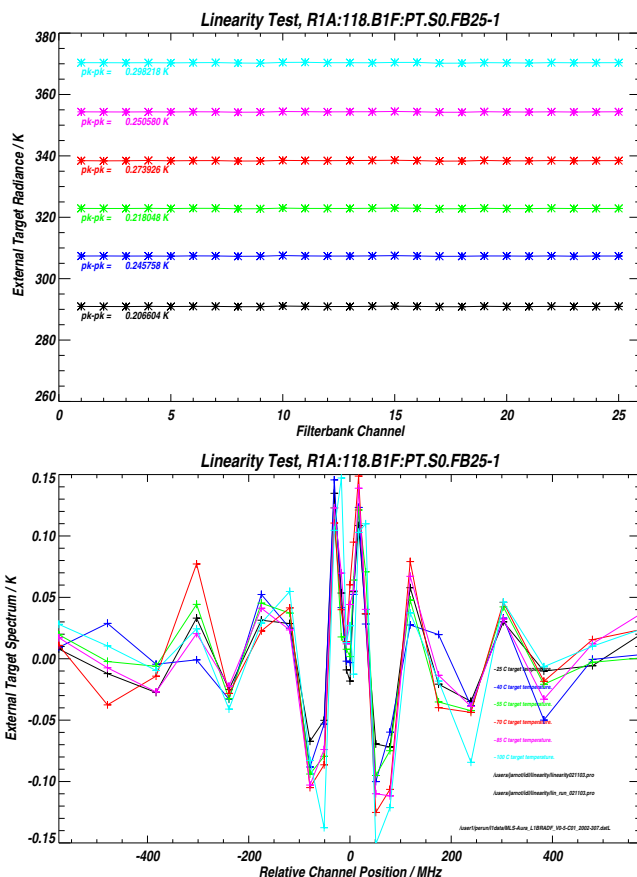


Fig. 7. Hot Target Radiances for all 25 channels of Band 1 (118GHz radiometer) for 6 different target temperatures. The lower panel shows the data from the upper panel with the weighted mean radiance subtracted from each measurement group.

linearity measurements to cover approximately 1/3rd of full dynamic range of atmospheric scene temperatures. Since the linearity measurements encompass the portion of the signal chain operating range that can be expected to be the least

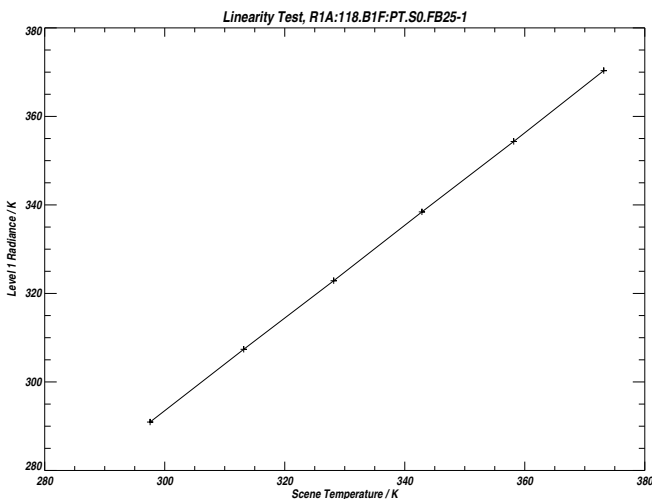


Fig. 8. Integral linearity plot for R1A. The offset between scene (target) temperature and Level 1 radiance is a consequence of the expression of radiance in units of Planck temperature, not a performance problem with the instrument.

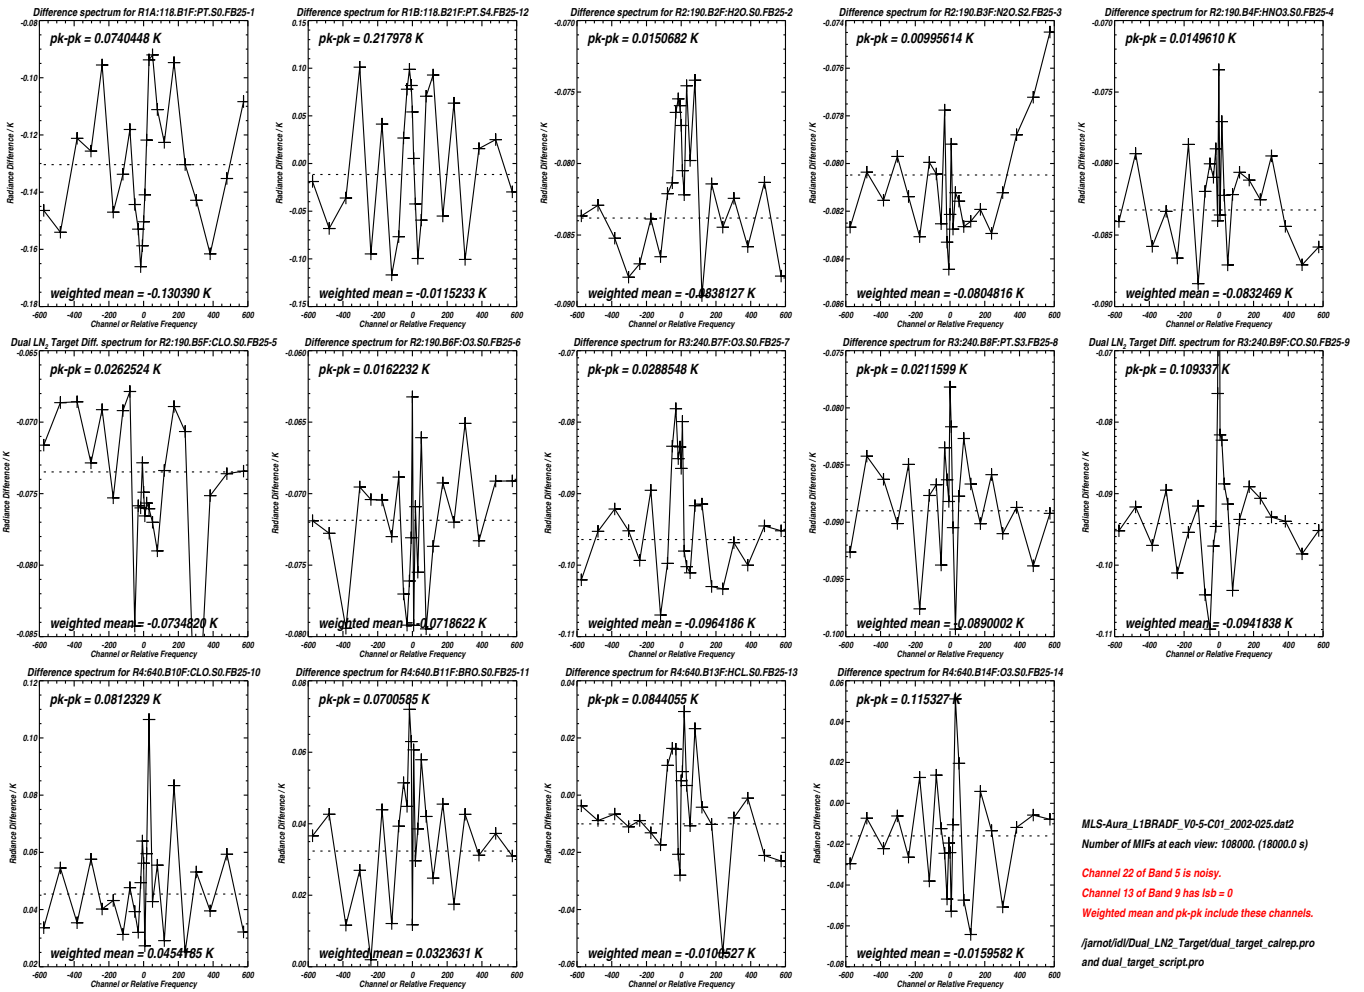


Fig. 9. Differences between limb and space port views to similar LN₂-cooled targets for the GHz 25-channel filterbanks. This data is used to verify the ability of MLS to integrate out-of-balance signals with the required precision

linear, we are confident in extrapolating these data to encompass the full working dynamic range. This outstanding linearity performance was obtained by running all RF amplifiers at least 12 dB below their 1 dB compression points, and operating all filterbank detector diodes with input power levels of -29 dBm or less. The ability to integrate signals down to levels well below the intrinsic resolution of the raw telemetry data is obtained by appropriate design of post-detector digitizers. For both UARS and EOS MLS the digitizers are implemented as voltage-to-frequency converters (VFC) operating over a small portion of their dynamic range.

A. Additional Radiometric Performance Tests

In addition to the linearity measurement discussed above, several other important tests were run to investigate radiometric performance. In order to verify the ability to integrate down data from scenes out of balance with the primary reference we took an extended (12 hour) data set with the GHz module switching between limb and space ports, and with the FOVs of both limb and space ports directed by plane mirrors into LN₂ loads like the one shown in Figure 3. This test was run when the GHz radiometer module was not integrated with the main

antenna. In addition to the limb and space port views, a short view to the internal GHz ambient target was performed during each MAF. This allowed the data to be processed by Level 1 software, using the target view as the primary reference, the LN₂ view via the space port as the gain reference, and the limb port view as the calibrated scene. The key information from this test is the limb-space port difference, shown plotted in Figure 9 for all GHz FB25 channels. Of most significance are the peak-to-peak spectral differences noted at the top left of each panel, which are at levels indicating that the ΔI requirements on sensitivity are met with adequate margin.

An end-to-end test (the ‘Blue Sky’ test) was run in which the instrument was located at the air lock to the assembly building, oriented in such a manner that two large plane aluminum faced panels could be used to direct the FOVs of both antenna and space port near zenith (Figure 10). The signal levels at the two ports were balanced by adjusting the angles of the lower panel, and multi-hour data sets taken while switching between the two internal calibration targets and the space and limb ports. For this test one of the internal targets was heated (by ~25°C) using its built-in heater elements, and the internal target views enabled Level 1 software to continuously determine channel




Fig. 10. ‘Blue Sky’ measurement setup. See text for details.

gain.

The exceptionally low atmospheric humidity brought about by Santa Ana conditions at the time of this test resulted in high tropospheric transmission at the lower measurement frequencies, and atmospheric spectra can be clearly seen in several of these bands. Another feature of these data is the significant offsets in the mid-band sections of the R4 25-channel filterbanks. Given the low thermal contrast between the internal targets and the atmospheric views, such an offset should not exist. This was traced down to inadequate voltage regulation in the spectrometer module, and a fix was implemented prior to instrument delivery. The relatively large standing waves visible in many of the spectral plots are attributable to standing waves in the internal target views similar to those shown in Figure 7. The small temperature difference between the two targets in conjunction with the radiometric calibration process at Level 1 leads to these enhanced standing waves, a situation that does not arise in orbit. Data sets were recorded with the antenna scan actuator at its nominal and extreme settings. These data were used to verify the absence of any significant scan dependence in the limb port signals.

IV. SPECTRAL CALIBRATION – HIGH RESOLUTION

Details of measurement channel widths and placements are given in [2]. All GHz measurement channels were swept using synthesized fundamental sources coupled into the space port of the GHz radiometer module. Each channel was swept with approximately 100 frequencies uniformly spread across each nominal filter passband, with the sweeps extending two filter widths beyond the nominal outer channel filterbank bandpasses to capture any extraneous responses. Wide filters were swept over a range covering three times their nominal width. All filter channels were swept in both sidebands (except for the single-sideband 118 GHz receivers), and with all combinations of switch network settings. The measurement setup is shown in Figure 3. The GHz module (M) space port is directed vertically down towards focusing mirror M1, seen edge on in the figure. Mirror M2 directs the instrument FOV towards switching mirror S which is shown pointing towards the LN₂-cooled target in dewar LN₂. During active sweeps S directs

the instrument FOV through the small hole visible in absorbing baffle B, through lens L, and finally onto the small horn visible at the left hand end of fundamental source F. A schematic of the sweep setup is shown in Figure 12.

A. Alignment

Optimal data quality requires precise alignment of the instrument FOV through the calibration path for all calibration configurations and setups. Focusing mirrors M1 and M2 are positioned correctly with respect to the GHz module by means of metering rods and tooling balls located on the periphery of the mirrors and the GHz module space port. Four tooling balls are clearly visible on M2. All external mirrors are oversized relative to the beams to ensure that the signals are fully captured. The next step in the alignment process consisted of positioning and orienting switching mirror S so that the beam from M2 illuminated S centrally, and also passed through the center of the small aperture visible in absorbing baffle B.

A mirror/LN₂ target combination similar to the one shown in Figure 3 was located in the place of the RF source F, and a custom-built low noise power detector/amplifier used to monitor the total IF power from a band in the radiometer to be swept. The power meter had sufficient S/N to allow the receiver FOV to be traced with a hand-held absorbing iris similar to the one in Figure 3, starting at S. This mirror assembly could be easily translated and rotated, then clamped in place once the receiver FOV struck the mirror centrally, and the reflected beam was directed through the small aperture in B. The remainder of the alignment process was similar, the entire alignment process taking only ~10 minutes to complete after some practice.

The absorbing baffle B serves two purposes: it absorbs spillover from the source F through focusing lens L illuminating the aperture on the baffle; also by replacing the baffle with one having an aperture size smaller than the beam, the source power could easily be attenuated when necessary. A small selection of absorbers with different aperture sizes was kept on hand to facilitate rapid setting of signal levels.

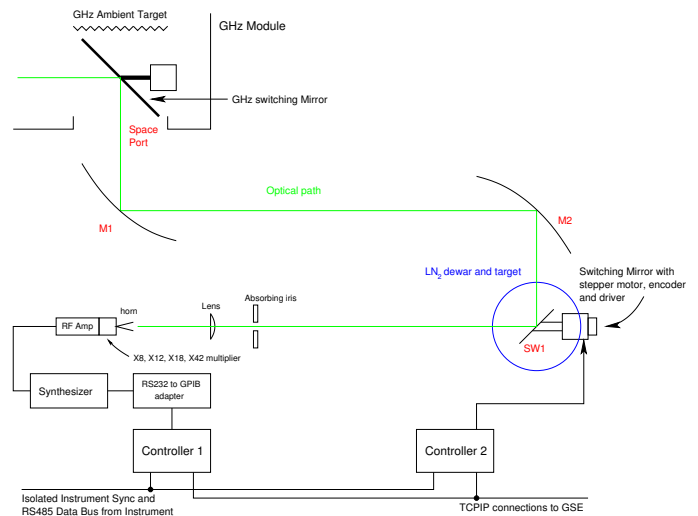
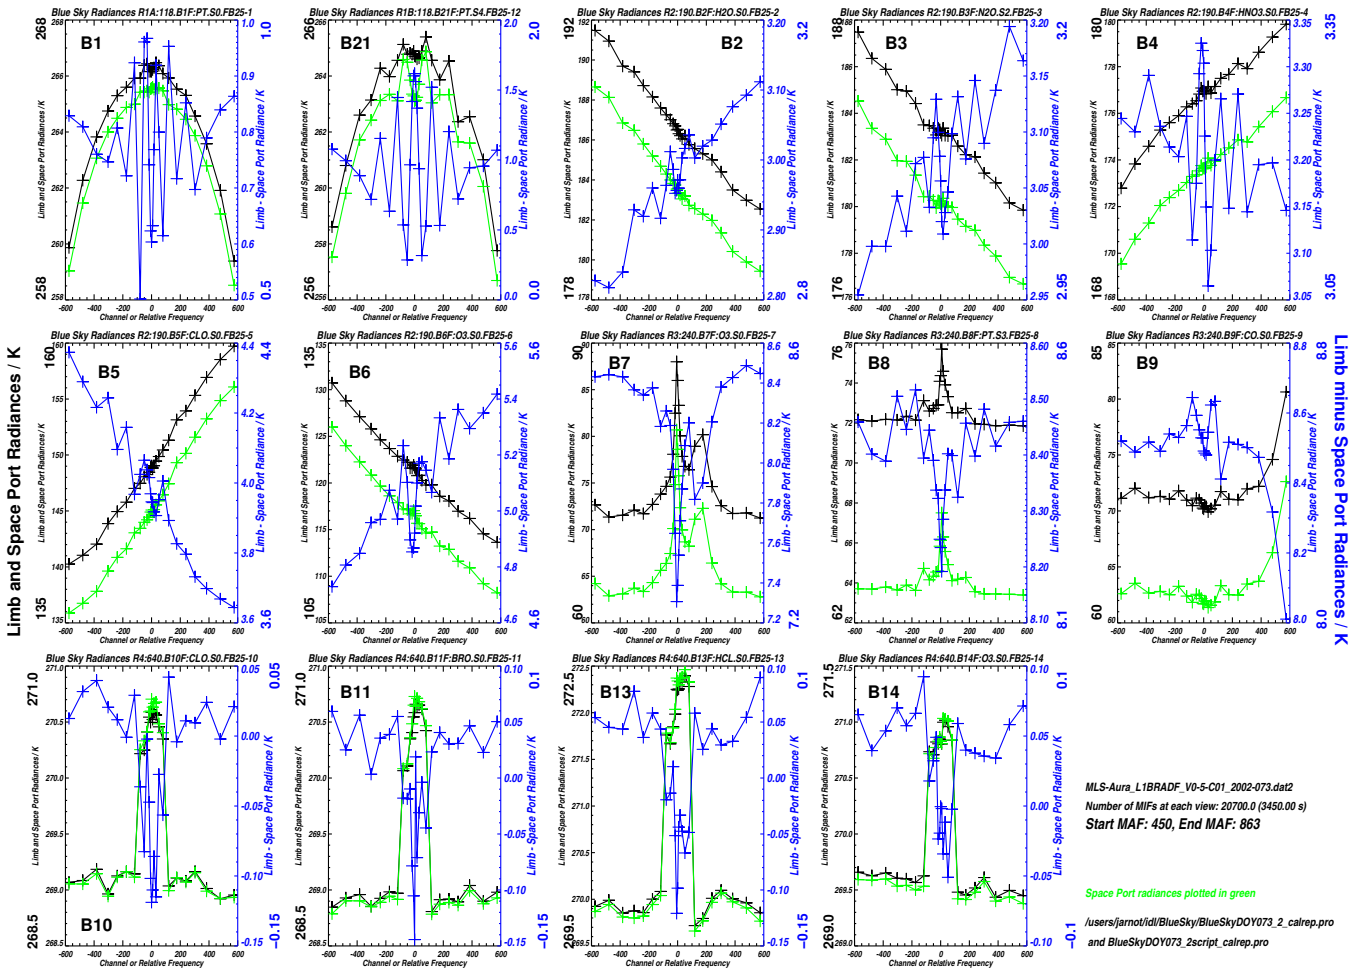


Fig. 12. Schematic of high resolution front-end sweep setup.



MLS-Aura_L1BRADF_V0-5-C01_2002-073.dat2
 Number of MIFs at each view: 20700.0 (3450.00 s)
 Start MAF: 450, End MAF: 863

Space Port radiances plotted in green
 /users/jarnot/dl/BlueSky/BlueSkyDOY073_2_calrep.pro
 and BlueSkyDOY073_2script_calrep.pro

Fig. 11. Limb and space (green) port data from the Blue Sky test processed to show sky radiances measured through the both ports. The blue line indicates the difference between Limb and Space port radiances.

As shown in Figure 3, the signal source F is wrapped in copper foil to prevent leakage from the relatively strong signal source directly into the closely-situated GHz space port. A simple way of verifying that all unwanted paths from the source to the receiver had been dealt with was to center the source in a filter channel and look for any changes in observed signal level as a hand-held piece of absorber was moved around the measurement setup. For repeatable results it was essential to remove all unwanted signal paths between source and space port.

When calibration measurements were underway, additional absorbing panels were placed between B and the instrument space port to eliminate potential problems arising from diffraction/reflection of the source signal from both baffle B and lens L. Additional large panels were placed around the entire calibration setup to prevent any interactions from people moving around the instrument. With these precautions in place we observed no interference or interactions of any kind from concurrent activities taking place around the instrument.

B. Signal Source

A family of signal sources, F, were used to span the 114 to 662GHz range covered by the GHz receivers. A single

source covered all of the bands in a given receiver, and all sources were fed by a programmable synthesizer operating in the 13.5 – 17.1GHz range. The sources multiplied the synthesizer output frequency by factors of 8, 12, 18 or 42 depending upon the band/sideband being swept. The small fractional bandwidth of each channel ensured that the power from the source remained constant over the channel width. Minor problems were encountered sweeping a handful of channels when unwanted harmonics from the multiplier passed through the channel response in the opposite sideband to the one being targeted. There were also some instances where the mode of an unused harmonic in the source jumped, causing a small change in RF power. Both of these conditions were readily identifiable in the data, and fixes were built into the analysis software to correct these potential error sources [10].

C. Standing waves

With the 118GHz radiometers it was observed that the relatively large iris aperture in B needed for these sweeps, combined with the precise optical alignment between test equipment and receivers, created noticeable standing waves. These were readily observed by performing a sweep of a 96MHz channel twice, with the source moved by $\frac{\lambda}{4}$ between

sweeps. It became standard practice to check for standing waves, and perform a pair of $\frac{\lambda}{4}$ separated sweeps for afflicted bands. The analysis software then combined these data so as to remove the effects of the standing waves. To make precision movements of the source easy and repeatable, it was mounted on a micrometer-driven translation stage. The source was mounted on a rotation stage located on the micrometer-controlled translation stage so that polarization of source and receiver could be lined up with ease.

D. Measurement sequence and timing

To maximize measurement efficiency the operation of the internal and external switching mirrors, and of the synthesizer/sweeper, were precisely synchronized. As with the radiometric calibration measurements, the measurement cycle was similar in timing to the standard in-orbit operating mode, allowing the use of Level 1 processing to remove the effects of gain and offset drifts. The first 20 s of each MAF was used to step the external source through 20 frequencies, and the remaining 4.7 s was used for a reference measurement in which the source was moved out of band, and a view to an LN₂ target to provide a periodic gain reference. The source was moved to an out-of-band frequency rather than the synthesizer output being turned off in order to avoid any potential adverse thermal effects as the multiplier amplifier chain settled between low and high signal states. As an aid to the operators, the start of the sweep was signalled by a MAF in which the source frequency was moved between the center of the first channel being swept and the out-of-band frequency in a regular and easily recognizable pattern. The code in the microcontroller operating the synthesizer had sweep tables for all GHz bands and wide filter channels in both sidebands, and automatically generated the start-up sequence when a sweep was begun. With this highly automated system in place we were able to sweep a 25-channel filterbank (2,900 discrete frequencies) in just under 1 hour, and perform 8 complete and verified sweeps in a standard working day.

E. Results

Figure 13 shows the results of an upper sideband sweep of Band 7 of R3, covering a frequency range of ~ 242.7 to 244.4 GHz. The plot was made with the ‘quick look’ tools used to verify data quality in near real time. Figures 14 and 15 show the detailed responses of the lowest and highest frequency FB25 channels in the GHz module. The legends cf and width are the measured center frequency (in final IF space) and 3 dB width of the channel. The numbers in parentheses are the nominal values of these parameters.

V. SPECTRAL CALIBRATION – LOW RESOLUTION

In a double-sideband measurement system like EOS MLS it is necessary to determine the *relative* response of each channel in the two sidebands. We measured relative sideband response by observing the difference between ambient and LN₂ loads viewed through a scanning Fabry-Perot interferometer. A schematic of this setup is shown in Figure 16, and a




Fig. 13. Front-end sweep of the upper sideband of Band 7. The x-axis is the spectrometer input frequency. All channels have been independently normalized to unity peak response.

photograph in Figure 17. The sideband sweep setup uses two external switching mirrors. The first mirror (SW1) directs the receiver FOVs to either an LN₂ load or the Fabry-Perot. The second mirror (SW2) directs the FOV through the Fabry-Perot to either ambient or LN₂ loads. The internal GHz switching mirror directs the FOVs to either an internal ambient load or the external calibration setup.

A family of calculated Fabry-Perot transmission curves is shown in Figure 18 for the 240 GHz radiometer. The widths of the transmission orders are seen to be considerably wider than the individual bands being characterized. The orders of the Fabry-Perot are seen to be well separated in all bands in this example. The same is true for the sweeps of the 640 GHz bands, which were performed with finer Fabry-Perot grids. Sweeps with both grid sets were performed over large enough ranges of grid separations that the sideband responses of all




Fig. 14. The end-to-end frequency responses of the first 4 channels of Band 1 of R1. The responses are plotted on a linear scale in blue, logarithmically in green.




Fig. 15. The end-to-end frequency responses of the first 4 channels in the upper sideband of Band 14 of R4. The responses are plotted on a linear scale in blue, logarithmically in green.

bands were well resolved over several Fabry-Perot orders. In this figure the ‘brightness temperature’ difference of the views to the ambient and LN₂ loads has been plotted, assuming unity peak transmission.

A. Alignment

The alignment process for the Fabry-Perot sweep system was the same as that described for the high resolution sweeps. The grids were removed from the Fabry-Perot, and the receiver beams traced along their optical path using a small iris in a hand-held piece of absorber. At the end of the alignment, the beams were reflecting off the centers of each mirror, and passed centrally through the Fabry-Perot grid holders. Once this stage had been reached, the grids were placed into their holders on the translation stage, and aligned by maximizing transmission (i.e., minimizing channel counts when viewing the LN₂ load through the Fabry-Perot) by iterative adjustments

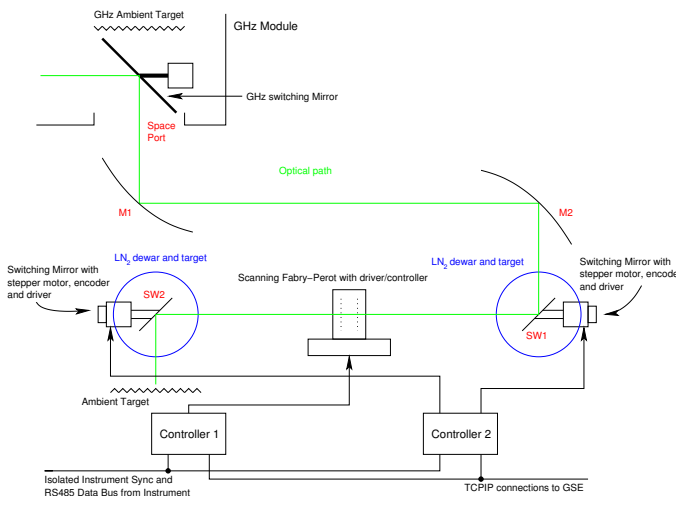


Fig. 16. Schematic of the Fabry-Perot sweep system used to measure relative sideband response in the GHZ radiometers.




Fig. 17. Photograph of the Fabry-Perot sweep system depicted in Figure 16.

to grid spacing and angle. A complete alignment of this measurement setup took about 20 minutes.

For these measurements no attempt was made to orient the Fabry-Perot slightly off axis, a common practice to reduce or eliminate standing waves in the cavity formed by the front grid and the receiver. Instead we chose to allow the presence of these small (but not insignificant) standing waves. Since these artifacts were present when viewing both the ambient and LN₂ loads through the Fabry-Perot, their effect was eliminated in the difference of these two measurements, the quantity used in the data analyses.

B. Measurement sequence and timing

As with the high-resolution sweeps, the measurement sequence closely matched the in-flight one to allow use of Level 1 software to remove the effects of gain and offset drifts, and to convert the measured signals into radiances. The first 20s of each MAF is spent viewing through the Fabry-Perot

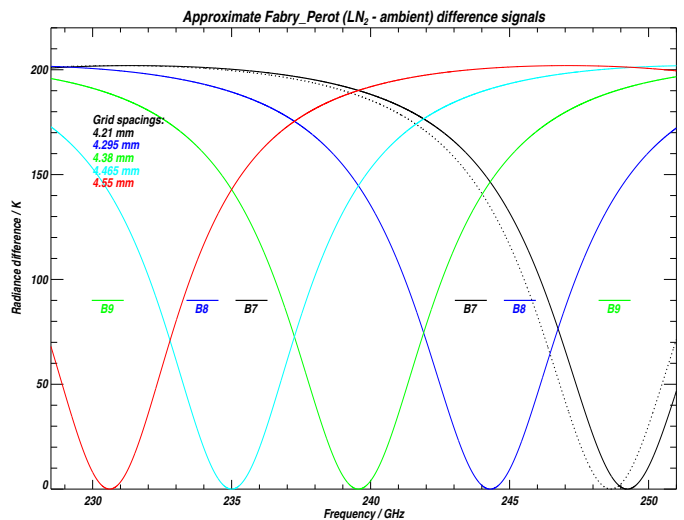


Fig. 18. Computed radiance signatures generated by the Fabry-Perot when sweeping the 240 GHz radiometer bands. The dashed line shows the spectral shift corresponding to a 0.01 mm increase in grid spacing from the initial spacing. Band positions in both sidebands are indicated for reference.

with the separation of the grids stepped every 5 s. For the R4 sideband sweeps, the switching sequence was kept the same, but the grid spacing was increased every 10 s to compensate for the poorer signal-to-noise in this radiometer. Half of the time at each grid separation was spent viewing the LN₂ load, the other half viewing the ambient load via SW2. The remaining 4.7 s of each MAF was used for viewing the external LN₂ via SW1 and the internal GHz ambient load via the internal switching mirror. Raw data from 1 MAF for Channel 1 of Band 2 is plotted in Figure 19. Clearly visible in these data are the alternating views through the Fabry-Perot to the LN₂ and ambient targets via SW2. The increasing transmission of the Fabry-Perot is a result of the 10 μm steps in grid spacing between measurements. The views to the LN₂ load via SW1 and to the internal GHz ambient load can be seen at the end of the MAF. The grid spacing for these measurements was 90 lines per inch (lpi), and by stepping the grid spacing from ~3.74 to ~9.93 mm it was possible to measure the sideband response all bands in R2 and R3 from a single data set. Figure 20 shows the sweep data for channel 1 of bands 2 and 9, the lowest and highest frequency double sideband channels swept with the 90 lpi grids.

For sweeping R4, the 90 lpi grids were exchanged for 200 lpi ones, the Fabry-Perot step size was reduced to 5 μm, and the scan range was from ~3 to ~8 mm. Figure 21 shows the data for channel 1 of bands 10 and 14, the lowest and highest frequency channels swept with the 200 lpi grids. The data plotted in Figures 20 and 21 are the output from Level 1 expressed as the radiance difference between the ambient and LN₂ views through the Fabry-Perot.

C. Data Analysis

The first step in the analysis of Fabry-Perot sweep data was the conversion from raw counts into calibrated radiances using Level 1 software. We evaluate the transmission, τ , of the Fabry-Perot using the simple expression:

$$\tau(s, \nu, r) = \left(\frac{1}{1 + 4 \sin^2(s * 4.19169 * 10^{-2} * \nu / 2) * r / (1 - r)^2} \right) \quad (7)$$

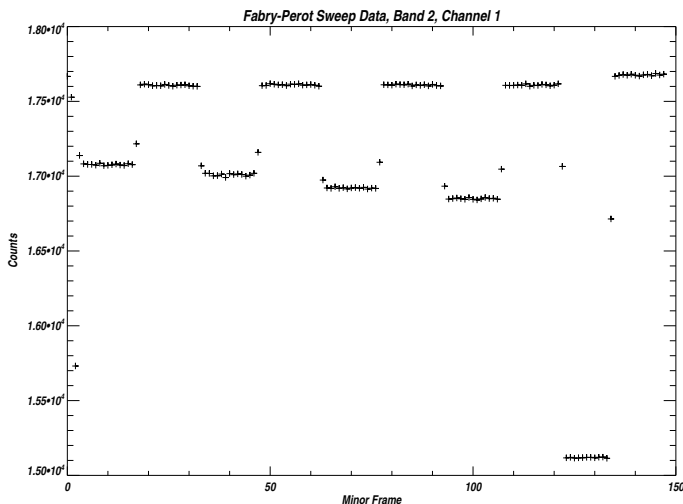


Fig. 19. Relative sideband sweep raw data for a single MAF. These data are for Channel 1 of Band 2.




Fig. 20. Level 1 radiances from channel 1 of band 2 (upper) and band 9 (lower) as the Fabry-Perot grid spacing stepped from ~3.74 to ~9.93 mm. These data were taken concurrently.

where:

s = grid spacing, mm
 ν = signal frequency, GHz
 r = grid reflectivity, and

$$4.19169 * 10^{-2} = 4 * \pi / c \text{ in appropriate units.}$$

The estimated double-sideband difference radiance measured in the target channel, T , is then given by:

$$T = 220 * (R_l * \tau(s, \nu_l, r_l) + R_u * \tau(s + \delta s, \nu_u, r_u) * P) / 2 \quad (8)$$

where:

220 = the approximate temperature difference between the 2 external targets, Kelvin,
 ν_x = the lower ($x = l$) and upper ($x = u$) sideband frequencies, GHz
 s = the grid spacing, mm
 δs = a correction applied to the upper sideband grid spacing to allow for the change in phase of grid reflections as a function of signal frequency, mm
 R_l = sideband response (lower sideband)
 R_u = sideband response (upper sideband)
 P = a Planck correction – see text below, and
 r_l, r_u = lower and upper sideband grid reflectivities.

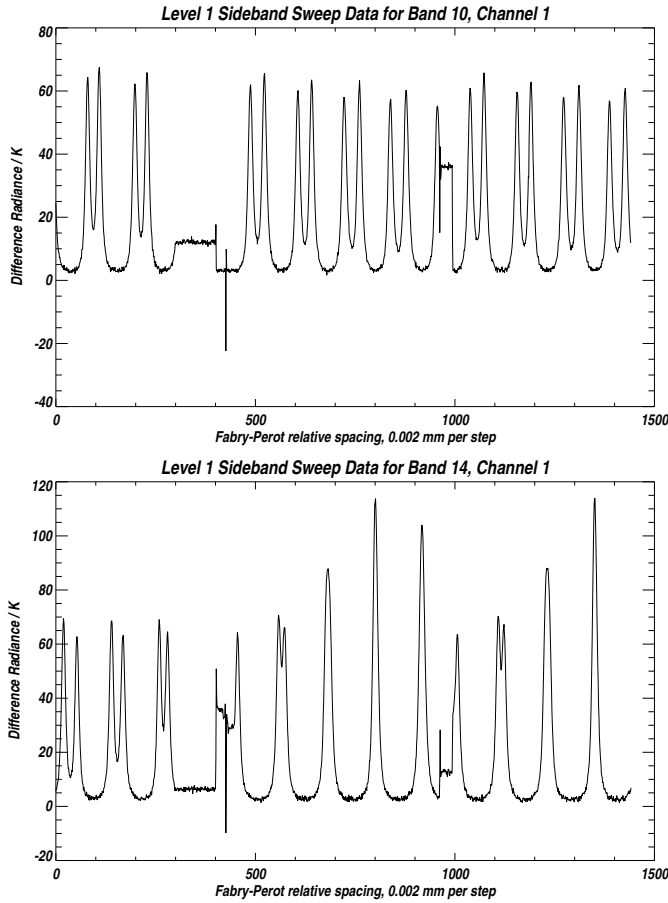


Fig. 21. Level 1 radiances from channel 1 of band 10 (upper) and band 14 (lower) as the Fabry-Perot grid spacing stepped from ~ 3 to ~ 8 mm. These data were taken concurrently. The ‘gaps’ in the data indicate where the Fabry-Perot spacing was increased manually to allow all channels in R4 to be swept with well-resolved sidebands in a reasonable time.

In Equation 8 above we assume that the ambient and LN₂-cooled targets observed through the Fabry-Pérot are stable and at temperatures of 300 K and 80 K for the duration of the measurements used in the data analysis. We use approximately 60 measurements (Fabry-Perot spacings) in each fit, taking approximately 12 minutes of data (about double this time for R4), sufficiently short timescales to maintain adequate stability in the climate-controlled environment in which these tests were performed. The factor 220 in the equation is just the temperature difference between the ambient and LN₂ loads seen through the Fabry-Perot. This is clearly an approximation that does not model the actual temperatures of the targets exactly. Both sideband responses will be in error by the same (small) multiplicative factor due to departures of scene temperatures from those assumed above, and this error is eliminated when the responses in each sideband are divided (to obtain the relative sideband response of each channel).

Although this approach takes care of the minor issue of not knowing the precise target temperatures, it does not account for the difference in Planck function between the two sidebands. This is accomplished by means of parameter P in

Equation 8, where:

$$P = \frac{B(\nu_u, 300) - B(\nu_u, 80)}{B(\nu_l, 300) - B(\nu_l, 80)} \quad (9)$$

and $B(\nu, T)$ is the Planck function at frequency ν and temperature T . This correction is extremely close to unity in all bands, and could have been ignored with no significant loss in accuracy.

The parameter δs in Equation 8 allows for the change in phase of the grid reflections between lower and upper sidebands, but does not account for the actual phase shift in the lower sideband. This is accounted for by fitting both parameters s and δs during sideband retrievals, the actual values of the phase shifts being unimportant to the end result.

The goal of the sideband retrieval is to determine the values of s , δs , w , r_l , r_u and R which provide the best match (in the least-squares sense) between measurements and model. We chose to perform the sideband retrievals on a band by band basis (25 channel filterbank), and to allow R_l and R_u to vary smoothly across each sideband, expressed as separate cubic polynomials for each sideband:

$$R_l(f) = a_r^l + b_r^l \times f + c_r^l \times f^2 + d_r^l \times f^3 \quad (10)$$

with a corresponding equation for the upper sideband response. Parameter f is the nominal channel center frequency at the input to each spectrometer.

The final measurement model thus had 12 parameters to be retrieved:

s	initial grid spacing, mm
δs	upper sideband phase correction, mm
r_l	lower sideband grid reflectivity
r_u	upper sideband grid reflectivity
a_r^l	constant coefficient,
b_r^l	linear coefficient,
c_r^l	quadratic coefficient, and
d_r^l	cubic coefficient of lower sideband polynomial
a_r^u	constant coefficient,
b_r^u	linear coefficient,
c_r^u	quadratic coefficient, and
d_r^u	cubic coefficient of upper sideband polynomial

Only the initial grid spacing of the measurement set, s , needs to be retrieved, since all other grid spacings of the measurement set are accurately known by adding the appropriate grid relative movement to this initial spacing. Note also that the phase difference between lower and upper sideband grid reflections is modelled as a small virtual grid separation offset, δs , in the upper sideband. This was purely for coding convenience, and is mathematically equivalent to modelling the phase shift difference as an angle.

The retrieval method was to perform a conventional least squares fit of the measurements to the model, allowing a sufficient number of iterations for adequate convergence. Typical retrievals took about 5 iterations and completed in a few seconds on a desktop PC. Partial derivatives were determined from small perturbations of the model on each iteration. The measured and fitted data for band 8 for a single Fabry-Perot

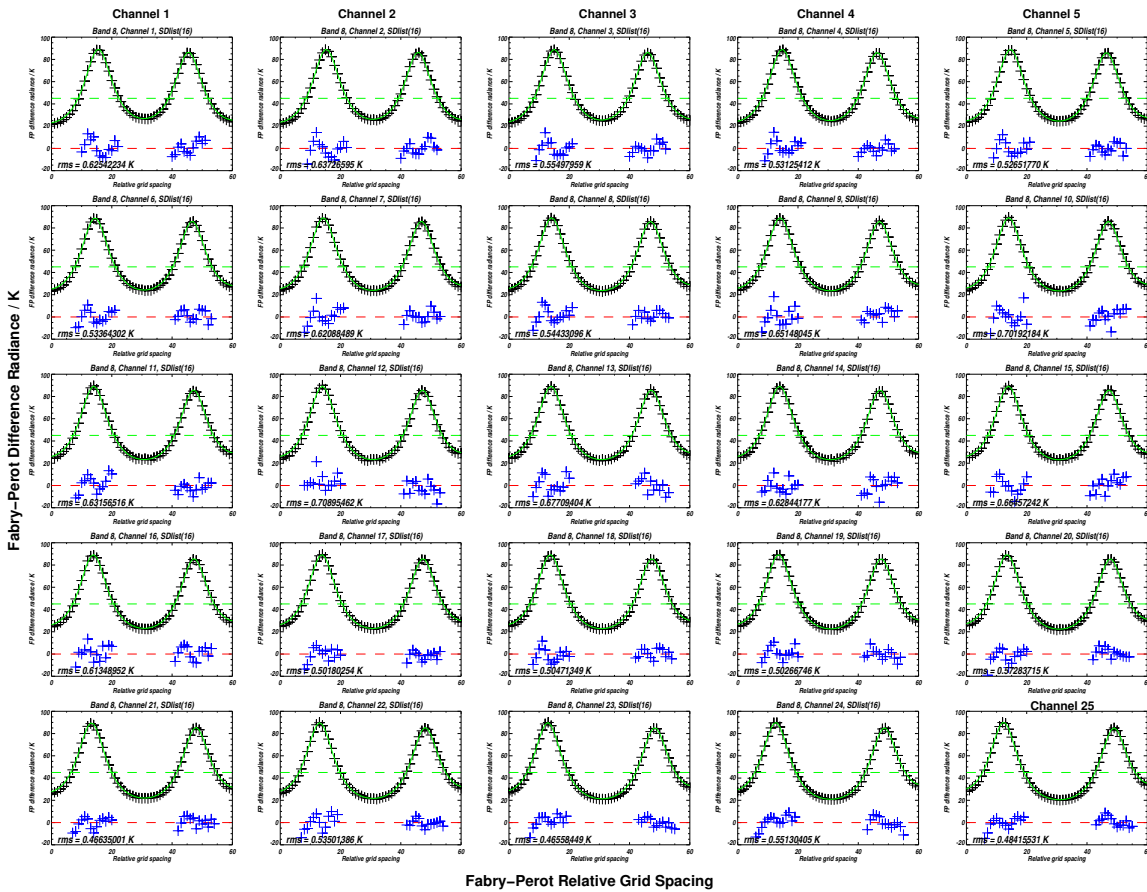


Fig. 22. Fabry-Perot data for Band 8 (black crosses) and theoretical prediction for the grid spacings chosen for the fit. Vertical range for each panel is -20 K to 100 K, and channel order is 1 to 25, left to right, top to bottom. The blue crosses are the residual to the fit (difference between measurements and model) multiplied by a factor of 10.

order are shown in Figure 22. The dashed horizontal green line indicates the radiance cutoff below which data are not included in the fit. The blue crosses at the bottom of each panel are the measurement residuals (difference between model and measurement) magnified by a factor of 10. The cutoff radiance was selected on a band by band basis to avoid systematics in the residuals which appeared in the data when the Fabry-Perot transmission was near its minimum. The retrieved data for the same band are plotted in Figure 23. The upper set of crosses is the retrieved relative sideband response (upper/lower) for each channel. The lower sets of crosses are the individual upper and lower sideband responses which do not add up to unity because of the finite throughput of the Fabry-Perot.

The analysis software was capable of fitting up to four orders simultaneously, limited only by available memory on the analysis computer. No significant differences were obtained when fitting to one or multiple orders.

The sideband measurements for all double-sideband GHz bands are plotted in Figure 24. Fits to individual 25-channel filterbank bands (FB25) are indicated by the black curve segments, where channel positions are indicated by the small plus signs. The small red and blue plus signs are the results from performing relative sideband retrievals on the data from individual channels, shown for information. Since there is overlap in several bands, and relative sideband response is

defined by the front-end receiver, we chose to perform a global fit to all of the FB25 sideband data in a given radiometer. This defines sideband response as a function of IF frequency, and




Fig. 23. Results of analysis of data in Figure 22 plotted by channel. The lower and upper relative sideband responses are shown as sideband fractions (blue and green crosses), and the ratio of the response in each sideband is indicated by the black crosses (upper/lower sideband response).

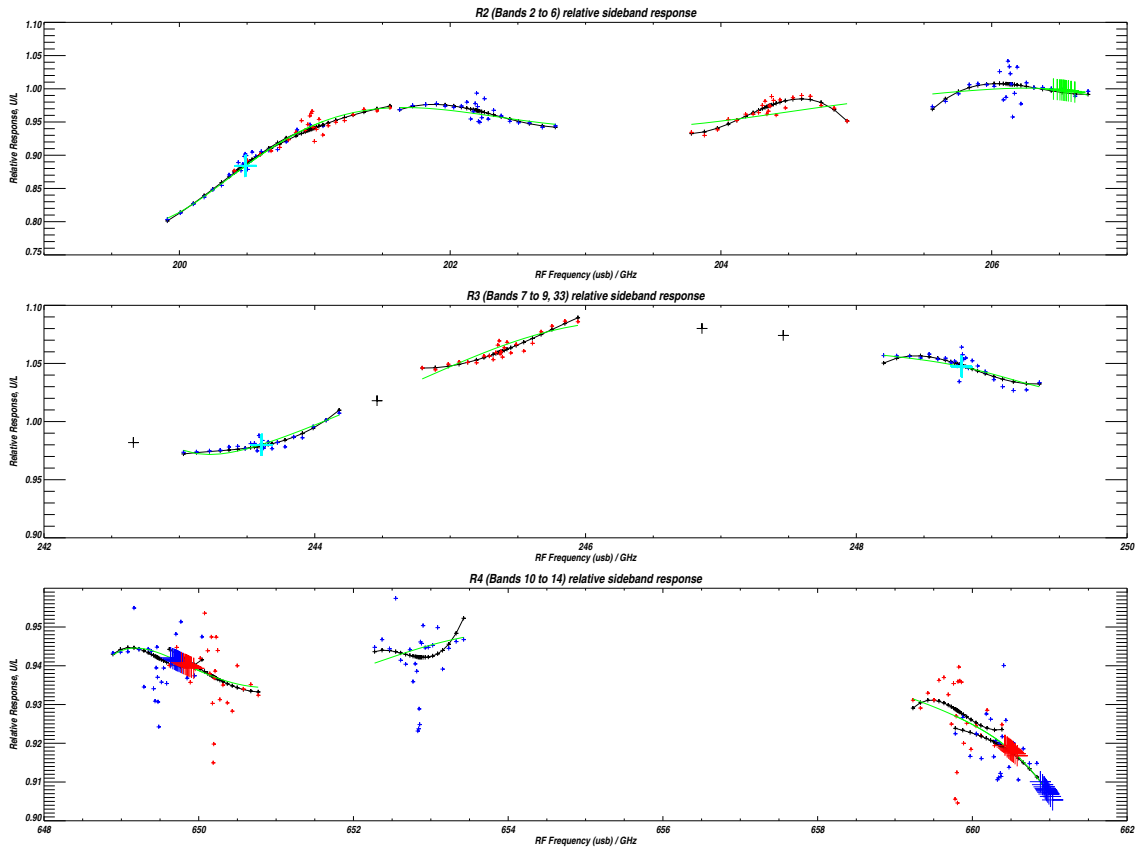


Fig. 24. Relative sideband response for all double sideband channels of the GHz radiometers. The large plus signs in the center panel are sideband data for the wide filter channels in R3. The large green, red and blue plus signs in the other panels are the sideband responses in the mid-band filterbanks (bands 27 to 31). Black line segments are results of sideband retrievals on individual bands, and the green segments are a global fit to these data for all bands in a given radiometer. Individual small red and blue plus signs are from sideband retrievals performed on individual channels.

the results of these fits are shown by the green lines, which are the data used by the forward model [11] in the Level 2 software [12]. The large green, red and blue plus signs indicate the inferred sideband responses in the 11-channel filterbank bands (FB11), and the large plus signs in the center panel are sideband responses for the wide filter channels in R3, which were retrieved on a channel-by-channel basis.

D. R4 Sideband Response

After delivery of the MLS instrument for integration on the Aura spacecraft, the whisker-contacted tripler in the 1st LO source for R4 was replaced by a more robust planar technology one. The R4 receiver front-end was removed and returned to JPL for this rework. One consequence of this change to the receiver was a small change in LO drive level and matching, resulting in a small change in relative sideband response. The sideband response of the receiver was measured with a simplified Fabry-Perot setup, before and after the tripler changeout. Without IF and filterbank subsystems it was only possible to measure sideband response at 3 IF frequencies, corresponding to the approximate centers of the 3 groups of channels in the lower panel of Figure 24. The effect of the tripler changeout was a decrease in the R4 relative sideband responses shown in Figure 24 by $\sim 4\%$.

VI. MASTER OSCILLATOR

All critical frequencies sources in the MLS instrument are referenced to a 5 MHz Master Oscillator (MO). Trend data taken over a ~ 480 day interval prior to instrument launch is shown in Figure 25 for both the primary and backup units. The requirement is for drift of no more than 1 part in 7×10^{-8} of the starting frequency at the end of the mission (6 years after launch). Extrapolation of these data indicate that requirements are met. In-orbit spectra observed by the digital autocorrelators (DACs) provide a very sensitive measure of first LO placements in R1 and R3, and provide a valuable indicator that no unexpected drift has taken place in the MO during the first 6 months of in-orbit operation.

VII. SPECTRAL CALIBRATION ACCURACY

Requirements on spectral calibration accuracy are similar to those for radiometric calibrations, and may be summarized as: the systematic error contribution in calibrated radiances shall be no more than 1% of the calibrated scene radiance. This allocation is separately applied to high-resolution and relative sideband data.

For the high resolution sweep data the most sensitive parameter is channel position, and since all downconversions are locked to the ultra-stable MO, we anticipate no frequency drifts in the signal chains that precede the spectrometers. The DACs process their input signals digitally, and are similarly

immune to drifts. The only likely potential source of significant channel drift arises in the LC filters used to define the bandpasses of the filter spectrometers. The inductors in these filters are hand-wound, and annealed twice during filter tuning. The spectrometers were extensively temperature-cycled before instrument delivery, and no evidence of channel drift was detected. The on-board synthesizer, which uses the MO as its reference, can determine filter characteristics in flight should the need arise. Level 2 residuals to date give no indication of drift in the placement of any channel.

The end-to-end sweeps of channel position/shape were of very high precision as evidenced by the sample data shown in Figures 14 and 15, making shape errors also a negligible source of error. The remaining potential error source in these measurements arises from ‘out of band’ responses which could be at a low level compared to the peak response of any channel, but could also span a much larger frequency range than the nominal filter width. All filterbanks were swept over a frequency range wider than the nominal spectrometer passbands by twice the width of the broadest filter in the spectrometer, and the data carefully examined for spurious responses. No such responses were observed, but to be conservative we have chosen to allow the error budget for filter channel position and shape to be 0.3% of the measured radiance.

The error introduced by uncertainty in relative sideband response is difficult to estimate. For the 118 GHz radiometers their single sideband response was adequately verified by front-end sweeps in the image sideband using a fundamental source, and by examination of the Fabry-Perot sweep data used to determine band 2 to band 9 (i.e., R2 and R3) sideband responses (the R1 radiometers were active during sideband sweeps). For relative sideband responses close to unity, the case for all of the GHz double-sideband MLS radiometers, an fractional error of δ translates into a corresponding radiance error of $\frac{\delta}{2}$ in each sideband.

For R2 the largest deviations between the individual band and IF-wide fit data are ~ 0.02 in relative sideband response,

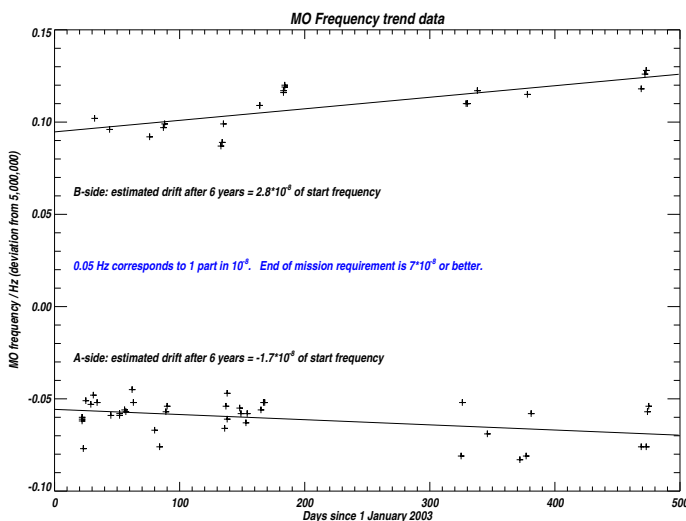


Fig. 25. Trend plot showing drift of the primary (lower) and backup master oscillator frequencies.

corresponding to 1% error in single sideband radiance. For R3 the errors are approximately half as large, and for R4 half as large again, indicating peak relative sideband uncertainties of ~ 0.01 and ~ 0.005 , corresponding to radiometric errors of 0.5% and 0.25% in these radiometers. The presence of unidentified systematics in these measurements may result in these uncertainty estimates being somewhat optimistic, but even so it is reasonable to assume that the combination of errors in high-resolution and relative sideband calibration data comfortably meet the intent of the requirements.

We expect significant feedback from Level 2 data in the future in regard to relative sideband response, and such information will be documented in [10].

A. The Common Calibration Setup

An important factor in the efficiency of all of the instrument calibrations and characterizations discussed so far was the use of a common setup which allowed simultaneous setup of GHz radiometric, spectral and FOV calibrations. For most of the pre-launch calibrations the GHz radiometer module was positioned with the space port viewing downwards, and the antenna system viewing the FOV calibration setup with the long axis of the GHz primary reflector horizontal. Radiometric and spectral calibration setups were located below the GHz radiometer module, coupled to the space port via a pair of focusing mirrors. Radiometric and spectral calibration equipment could be placed and aligned in the working area below the radiometer module, and any of the calibration setups chosen by appropriate positioning of the switching mirrors, a task consisting of just issuing commands to the instrument and calibration controllers. This flexibility allowed spectral and radiometric calibrations to be performed during the day, and the longer FOV measurements to be performed at night (when measurement conditions were best), with reconfiguration taking just minutes.

Another key contributor to the high efficiency of radiometric and spectral calibrations was the use of controllers which ‘snooped’ the internal instrument data and timing buses via optically isolated interfaces. This allowed the operation of external calibration equipment to be precisely synchronized to the $\frac{1}{6}$ s instrument measurement frames. Quick-look software allowed near-instantaneous feedback on data quality, and was also of great help during alignment of calibration setups to the instrument FOV from the GHz module space port.

A useful benefit of the highly-efficient calibration setups was that the reduced measurement times (compared to UARS MLS) resulted in improved thermal stability of the signal chains. This was a very useful benefit given the total-power (unchopped) implementation of the MLS instruments. In addition, by making all calibration sequences mimic the timing of the instrument in its nominal operation mode, Level 1 software could be used to convert the data into calibration ‘radiance,’ taking full advantage of the gain and offset drift compensation in the Level 1 algorithms.

VIII. IN-ORBIT PERFORMANCE

In-orbit performance of the GHz radiometers is basically the same as seen during ground testing, but with improved $\frac{1}{f}$




Fig. 26. Plots of T_{sys} for all GHZ 25-channel bands active during nominal operation. The noise performance shown here has been stable for the duration of the mission to date.

characteristics due to the more benign thermal environment and lack of convection-induced artifacts. Channel noise is characterized by T_{sys} and χ^2 , shown in Figures 26 and 27 for the GHZ FB25 bands. These Figures are samples taken from the daily diagnostic plots used to monitor instrument health and safety, and are representative of instrument performance from the beginning of the mission.

IX. IN-ORBIT CALIBRATION

The most significant in-orbit calibrations are refinement of pointing knowledge through moon scans, and updates to antenna scattering, emission and ohmic loss parameters. These changes were expected based on prior UARS MLS experience, and are described in [3].

An important in-orbit calibration has been the implementation of static and dynamic baseline corrections in Level 1 software. These updates were anticipated from UARS MLS experience, and Level 1 software was designed to accommodate the expected changes after launch. Static baseline is the spectral difference between the signals received when viewing space by the GHZ switching mirror space port and main antenna. This baseline signature was measured with long data integrations during the instrument activation phase, and again approximately 8 months later, to determine its stability.

These measurements indicate typical baseline signatures of a few $\frac{1}{10}$'s of a K peak-to-peak in most bands, with long-term stability of $\sim 0.02\text{K}$ (limited by the duration of the calibration data sets) in all bands. For R4 we see no spectral baseline signature down to the 20 mK level, possibly due to the more highly apodized illumination of the main optics for this radiometer. We plan on verifying the stability of this important data periodically.

Dynamic baseline refers to the offset in mean radiance between the space and limb port views to space. This component of the baseline signature is radiometer-dependent, ranging from $\sim 3.5\text{K}$ in R2 to $\sim 8\text{K}$ in R3 and R4, with orbital variations of just over 1K in the bands with the greatest offsets. This offset is measured on a MAF-by-MAF basis in Level 1 processing using radiances from limb views with tangent heights greater than 85 km for channels with negligible atmospheric contribution. Both static and dynamic baselines, together with precision estimates, are reported in the Level 1 radiance files for the GHZ data. Since the THz system uses the same scanning mirror to view both limb and references, it has no corresponding baseline signatures.

A small change made to radiometric calibrations since launch has been the determination of the thermal gradients in the passively-cooled GHZ target discussed earlier, and the

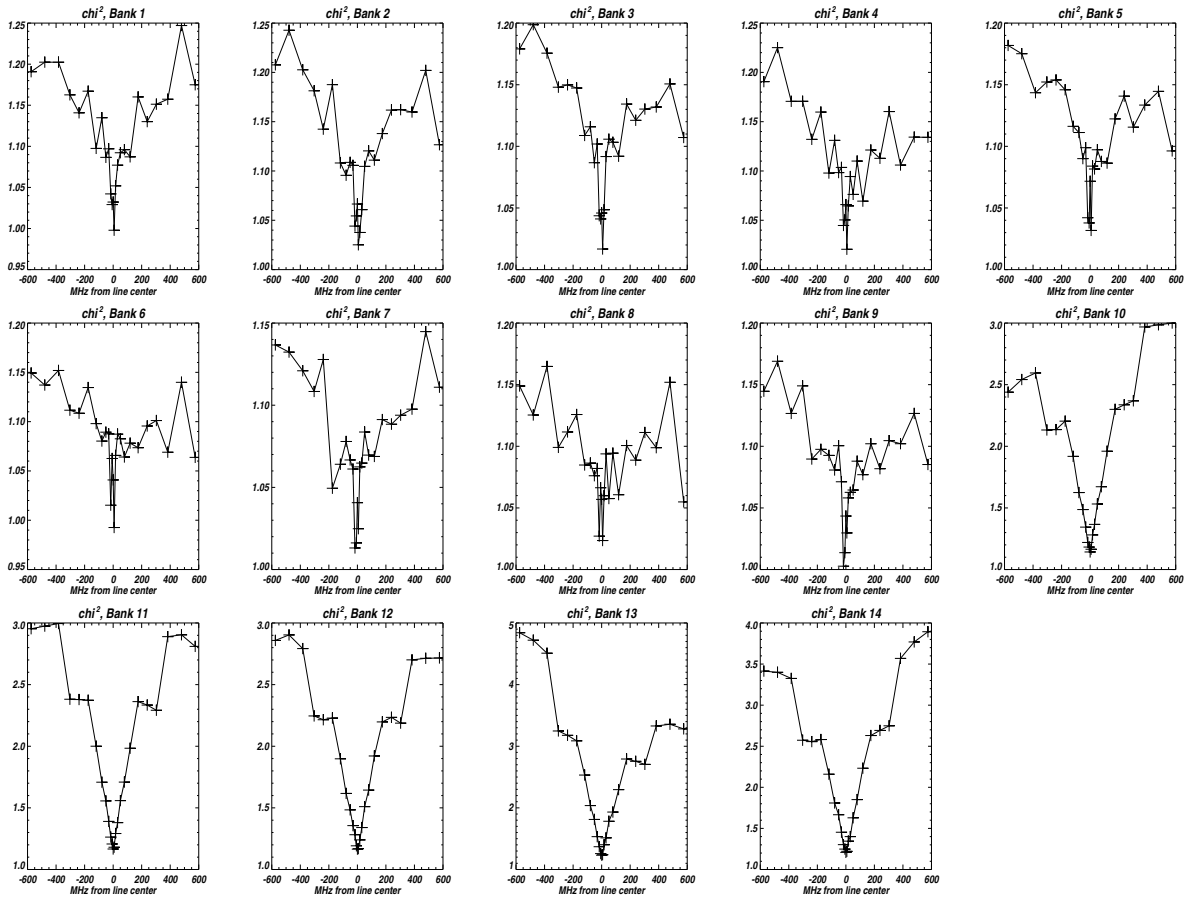


Fig. 27. Plots of χ^2 for all GHz 25-channel bands active during nominal operation. The noise performance shown here has been stable for the duration of the mission to date.

corresponding changes to Level 1.

Possible future in-orbit calibrations include the use of saturated atmospheric radiances to validate/refine relative sideband ratio knowledge, and also to develop methods for reducing the small impacts of standing waves in the views to the calibration loads.

A. R3 Spur Compensation

An unexpected behavior observed in flight was the presence of varying spurs in the R3 first LO, which manifested as clear artifacts in the DACS data radiance residuals. The presence of these spurs was known before launch, but they were expected to be stable. These spurs are offset ± 714 kHz from the LO fundamental frequency. There is a second, much smaller, pair of spurs at $\pm 1,428$ kHz. During the first six months of MLS operation the fraction of LO power in the spurs has ranged from less than 15% to more than 36%, with large changes associated with relatively small changes in operating conditions that effect the temperature of RF components.

As a result, Level 1 software has been modified to estimate the magnitude of spurs, and the spectra are corrected for their effects. Estimates are made separately for B24 and B25 (DACS bands 24 and 25) to simplify handling of missing data, even though the underlying source is the R3 LO in both cases.

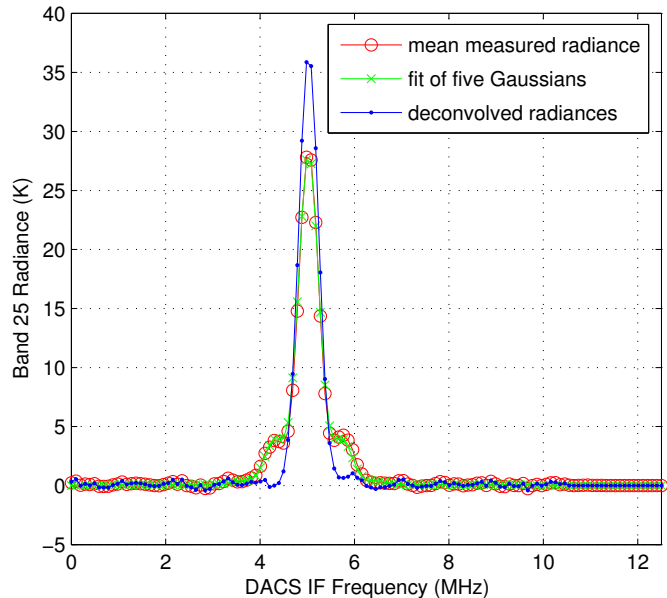


Fig. 28. CO spectrum from B25, before and after correction for 1st. LO spurs.

The spur amplitudes are retrieved from daily-mean, high-altitude (tangent point above 78 km) spectra from B24 and B25, which contain multiple images of the Doppler-broadened 235-GHz O₃ line and 230-GHz CO line, respectively. Since these Doppler-broadened emission lines are not saturated, the observed radiances have a Doppler shape. The high-altitude O₃ line is modeled as a Doppler shape centered at B24 channel 47.384 with σ of 1.67 channels. The high-altitude CO line is similarly modeled as being centered on channel 50.44 of B25 with σ of 2.2 channels. The DACS channel spacing is 97.66 kHz. A linear least squares fit of the mean spectrum with the modeled fundamental image and two pairs of images as well as a constant baseline offset is performed. All spectra are divided by the DCT of the retrieved LO shape in the Fourier domain to effectively deconvolve the LO shape from the spectra.

We must account for the effect of this deconvolution on the channel precisions. Receiver noise (which dominates measurement noise) is uncorrelated between channels in the raw measurements, but becomes correlated when the LO shape is deconvolved from the spectra. The noise in an individual channel increases, although a properly-constructed average over the band, using the full noise covariance matrix, beats down to the same noise level in both cases.

MLS level 2 retrieval algorithms do not currently account for off-diagonal noise covariance matrix elements, and the only means of accounting for the effects of correlated noise is to report the inflated diagonal elements. If we consider deconvolution to be multiplication by W in the time domain, where the elements of W are the inverses of the elements of the DCT of the LO shape, then the diagonal elements of the noise covariance matrix are inflated by the *rms* of W . Using estimated precisions inflated by this factor, Level 2 will have a correct value for the noise on an individual channel but will overestimate the noise on the average of a block of channels. This effect is small, with noise underestimated by a factor of ~ 2 on a block of data 4 s in duration. A similar situation exists for the filter channels, since Level 1 does not report the temporal noise correlations that are introduced by the use of interpolated space and target references.

Convolution in Level 2 of the forward model with the LO shape, rather than deconvolution of the LO shape in Level 1, permits Level 2 to work in a space with a nearly-diagonal noise covariance matrix. This approach is used for constant LO widths and spurs of other radiometers. At the time of writing such convolution has not been implemented in the Level 2 forward model, but it is intended to be implemented for later versions of the algorithms.

X. CONCLUSION

This paper has described the radiometric performance, and radiometric and spectral calibrations of the GHz component of Aura MLS. All performance and calibration requirements have been met, and, as with its UARS predecessor, we expect instrument calibration to be maintained, and possibly even enhanced, during the duration of the mission.

ACKNOWLEDGMENT

Numerous people contributed to the success and efficiency of the pre-launch characterization and calibration activities, including Rick Cofield for providing the external optics for the calibration setups, and the means to rapidly align them; Mario Loo, José Garcia and Gerson Melgar for operating the instrument during calibrations; Don Benson and Pim Vosse for fabricating the many brackets and supports that made up the calibration setups; Karen Lee and Kumar Chandra for their support of the spectral sweeps; Tim Crawford for providing the grid holders used in sideband sweeps; Joe Waters for having faith in the aggressive schedule for this activity. Many other colleagues contributed to this effort, too numerous to mention here, and we wish to thank them all.

XI. GLOSSARY

ΔI	noise on a calibrated measurement
ϵ	calibration target emissivity
η_r^X	transmission through switching mirror port X in band r
τ	data integration time
B	noise bandwidth
C_i^X	raw data ‘counts’ in channel i for view to scene X
C_i^O	raw data counts for channel i with no signal at the spectrometer input
DACS	128 channel digital autocorrelator spectrometer
FB11	11 channel filterbank spectrometer
FB25	25 channel filterbank spectrometer
g_i	radiometric gain (Counts/K) of channel i
GB	gigabyte
GHz	gigahertz, 10^9 Hz
kB	kilobyte
Level 1:	the software that converts raw instrument data into calibrated form
Level 2:	the software that transforms Level 1 output into retrieved atmospheric quantities/profiles
MAF	Major Frame, nominally 148 MIFs, comprising a complete limb scan/calibration cycle
MB	megabyte
MIF	Minor Frame, nominally $\frac{1}{6}$ s, during which science data are integrated
\dot{P}_i^X	radiance in channel i from port X
T_{sys}	System Temperature, the noise of the measurement system in temperature units

REFERENCES

- [1] F. T. Barath et. al., *The Upper Atmosphere Research Satellite Microwave Limb Sounder Instrument*, J. Geophys Res., vol 98, pp. 10,751 – 10,762, 1993.
- [2] J. W. Waters, et. al., *IEEE Trans. Geosci. Remote Sensing*, this issue.
- [3] R. E. Cofield, and P. C. Stek, *IEEE Trans. Geosci. Remote Sensing*, this issue.
- [4] H. M. Pickett, *IEEE Trans. Geosci. Remote Sensing*, this issue.
- [5] R. F. Jarnot, R. E. Cofield, J. W. Waters and D. E. Flower, *Calibration of the Microwave Limb Sounder on the Upper Atmosphere Research Satellite*, J. Geophys Res., vol 101, April 1996, No. D6, 9,957 – 9,982.
- [6] R. F. Jarnot, H. M. Pickett and M.J. Schwartz, *EOS MLS Level 1 Data Processing Algorithm Theoretical Basis*, JPL D-15210, available on request.

- [7] F. N. H. Robinson, *Noise and Fluctuations in Electronic Devices and Circuits*, Clarendon Press, Oxford, 1974, Chapter 19.
- [8] J. W. Waters and R. F. Jarnot, *Science Requirements on the EOS MLS Instrument and Data Processing Software*, JPL D-14421, Rev. 3, 26 February 2002.
- [9] J. W. Waters, *Observations for Chemistry (Remote Sensing): Microwave*, pages 1516-1528 in *Encyclopedia of Atmospheric Sciences*, edited by J. Holton, J. Curry and J. Pyle, Academic Press, 2003.
- [10] R. F. Jarnot, R. E. Cofield, H. M. Pickett and P. C. Stek, *EOS MLS Instrument Calibration Report*, JPL D-26280, available by request from jarnot@mls.jpl.nasa.gov
- [11] William G. Read, et. al., *The Clear-Sky Unpolarized Forward Model for the EOS Aura Microwave Limb Sounder (MLS)*, this issue.
- [12] Nathaniel J. Livesey, et. al., *Retrieval algorithms for the EOS Microwave Limb Sounder (MLS) instrument*, this issue.



Dr. Robert F. Jarnot Dr. Jarnot obtained his D. Phil in Atmospheric Physics in 1977 from Oxford University during which time he made the first stratospheric NO measurements using a balloon-based infra-red pressure modulator radiometer. Since then he has been at JPL working on various incarnations of MLS instruments ranging from balloon to satellite borne versions. His roles have included C&DH design, on-board and ground based support and analysis software design and implementation, responsibility for Level 1 algorithms, instrument calibration and general troubleshooting.



Mr. Vincent S. Perun Vincent Perun studied physics and electrical engineering at Case Western Reserve University. In 1983 he joined the Jet Propulsion Laboratory, Pasadena, California, to write software to support development and operation of the various MLS instruments. Responsibilities include Level 1 programming and data analysis tools.



Dr. Michael J. Schwartz began his undergraduate education at Deep Springs College, received his B.A. in Physics from Carleton College in 1985 and his Ph. D. in Physics from Massachusetts Institute of Technology in 1998. His thesis, supervised by Prof. D.H. Staelin, involved atmospheric remote sensing with millimeter-wave oxygen lines. He was on the research staff of the MIT Research Laboratory of Electronics in 1999, where he contributed to the development of an aircraft-based atmospheric sounding package. In 2000, he joined the Microwave Limb Sounding group at the Jet Propulsion Laboratory. His responsibilities have included the development of digital autocorrelator spectrometer hardware and processing algorithms for the EOS-MLS instrument, oversight of the MLS temperature and geopotential height data products, and development of the polarized forward model for Zeeman-split oxygen lines.