Earth Observing System (EOS)

Atmospheric Infrared Sounder (AIRS)

AIRS Level 1C Algorithm Theoretical Basis

Hartmut H. Aumann, Steve Broberg, Evan M. Manning, Tom Pagano and Brian Sutin

Caltech/JPL

L. Larrabee Strow
University of Maryland Baltimore County

Version 6.7
15 January 2020

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California 91109-8099

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Release Record

Version	Date released	Comments
1.0	March 1, 2011	Initial Release
2.0	March 10, 2012	Update
3.0	January 2015	Release limited to synthesizes radiances.
6.7	January 15, 2020	Full release

TABLE OF CONTENTS

1. Introduction

1.1 AIRS Overview

The Atmospheric Infrared Sounder (AIRS) [Aumann et al., 2003], launched aboard NASA's EOS Aqua spacecraft on May 4, 2002, is a grating array spectrometer having 2378 channels sensitive in the range 3.7 to 15.4 microns. The spectral resolution ($\lambda/\Delta\lambda$) is ~1200. A combination of a design philosophy having radiometric accuracy as a foremost goal, cooled and temperature-controlled spectrometer hardware (including most of the optics), and a thorough pre-flight calibration have made the AIRS a superb instrument that produces high quality radiance data. AIRS completed seventeen years of routine operations at the end of August 2019. The instrument remains healthy.

The 2378 AIRS channels are optically and electrically grouped in 17 detector modules, which share the scan mirror, the On-Board Calibrator (OBC) and Space Views (SV). Table 1 lists the wavenumber ranges covered by each of the 17 AIRS detector modules.

Table 1: The wavenumber ranges covered by each of the 17 AIRS detector modules.

Module Name	Start ν (cm^{-1})	End ν (cm^{-1})
m1a	2552	2677
m2a	2432	2555
m1b	2309	2434
m2b	2169	2312
m4a	1540	1614
m4b	1460	1527
m3	1337	1443
m4c	1283	1339
m4d	1216	1273
m5	1055	1136
m6	973	1046
m7	910	974
m8	851	904
m9	788	852
m10	727	782
m11	687	729
m12	649	682

The detectors, electronics and optical filters are module unique, with different sensitivity to radiation damage and surface contamination. Each of the 2378 channels is separately calibrated, and has the potential for noise behavior that is different from the other channels. The detector modules are spread across the focal plane. Small gaps between detector modules create small gaps in the spectral coverage. These gaps, as well as some overlaps in spectral coverage between separate modules, complicate the tasks of spectral calibration and integrating the spectrum over the passband of interference filter radiometers.

All of the AIRS infrared detectors are made of HgCdTe, but each detector module is made with a different doping for different spectral sensitivity, and each module has a unique Read-Out Integrated Circuit (ROIC). For redundancy, the image of the entrance slit for 15 of the 17 modules falls onto two detectors, referred to as A-side and B-side, which are read out independently. The weighted sum of the signals from the two sides is transmitted. The weights are stored in a 2378 element weight

table. Noise from each detector is characterized by it Noise Equivalent Delta Temperature (NEDT) when viewing a 250 K scene. The NEDT for the 2378 AIRS spectral channels is in the 0.07-0.6 K range, the median is 0.2 K.

The definition of what constitutes a "good" channel is subjective. Of the 2378 channels, 2225 channels were classified at launch as "good", i.e., with NEDT less than 1 K. As of 9/2019, 2223 channels were classified as "good". Of the 155 "bad channels, 63 were already dead during pre-launch testing. Some channels exhibit non-Gaussian noise, including "pops" (temporary changes in output level) and cold-scene noise (scene-dependent noise larger at lower signal levels) [Weiler *et al*., 2005]. The AIRS detectors and their ROIC's have different susceptibilities to radiation hits and to the slow build-up of total radiation dosage throughout the mission. A detector that exhibited very low noise for years can suddenly undergo a temporary or permanent noise increase.

The 17 detector arrays are located at the dispersed image of the spectrometer entrance slit. The position of each array is y_o in focal plane coordinates. The position of each detector in the array is defined relative to y_o. Each detector measures the radiance at a unique spectral frequency and with a unique Spectral Response Function (SRF). Changes in the spectral location of the centroid of the SRF can occur when changes in temperature gradients within the spectrometer optical train shift the image of the entrance slit. The temperature of the AIRS spectrometer is tightly controlled at one location, the base of the grating, with a choke heater, but changes in internal gradients can occur, which result in small shifts of the frequency with time, notably with orbital and seasonal periodicities.

AIRS instrument data products are categorized as Level 1A (L1A), Level 1B (L1B), and Level 1C (L1C). The L1A products contain raw detector counts. The L1B product consists of radiometrically calibrated radiances [Gaiser *et al*., 2003] at the nominal spectral frequencies for each of the 2378 channels. The L1A and L1B products have been produced since the start of routine instrument operations on August 31, 2002.

1.2 L1C Functional Requirements

The L1B product has a number of characteristics related to the design of the instrument, such as gaps in the spectral coverage, spectrally overlapping channels, and the presence of dead or noisy channels. The L1C product is an augmentation of the L1B product, which addresses these complications. The L1C Functional Requirements are as follows:

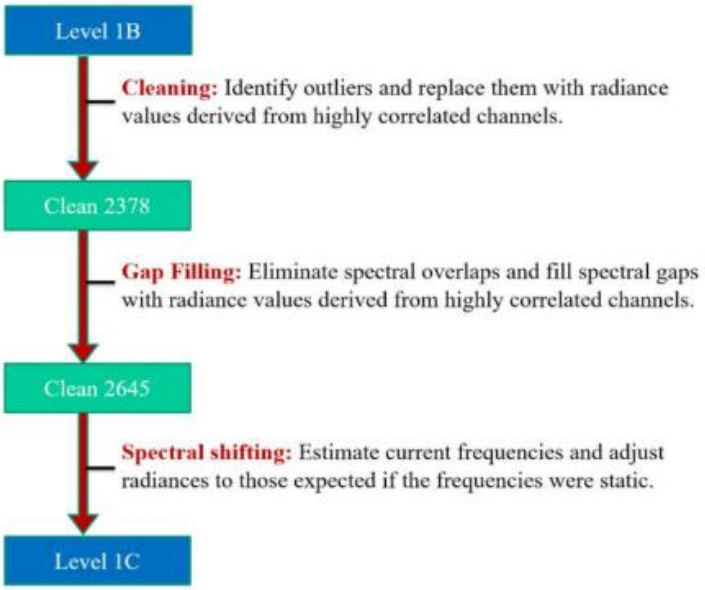

1) Identify and replace dead or noisy channels
2) Mitigate the radiometric effects of spatial inhomogeneity.
3) Define the Spectral Response Function (SRF) centroid positions as a function of time.
4) Present the AIRS spectrum as a monotonic sequence of 2645 channels and fill the gaps and overlapped regions of the spectrum as listed in Table 2.
5) Shift the radiances of the 2645 channels to a fixed frequency grid.

Figure 1. L1B to L1C flowchart

Figure 1 shows a processing flowchart from the 2378 channels in L1B, to Clean2378, the cleaned 2378 channel version of L1B, to Clean2645, the gap-filled version of Clean2378, to L1C, where the radiances are shifted to a fixed grid. The Clean2645 version has been un-officially available from the GSFC DISC as L1C v6.1. While the L1C trades a slight decrease in SI-traceability for a much more user-friendly product, it should be the first AIRS radiance product for new users.

Figure 2 shows a typical L1B spectrum before and after gap-filling. The panel on the left shows an overlay of 2223 of the 2378 "good" channel (blue), overlaid on the 2645 channel L1C spectrum (red), overlaid on the 2378 channel PCR (principal component reconstruction) using the first 100 EV (eigenvectors). The red-only regions of the spectrum are gap filled. The right panel zooms in on the 780-980 spectral region. The PCR for any one channel is less noisy than the observed noise, because correlated information from all other channels is included. The PCR values are used only for the replacement of "bad" channels (see section 2).

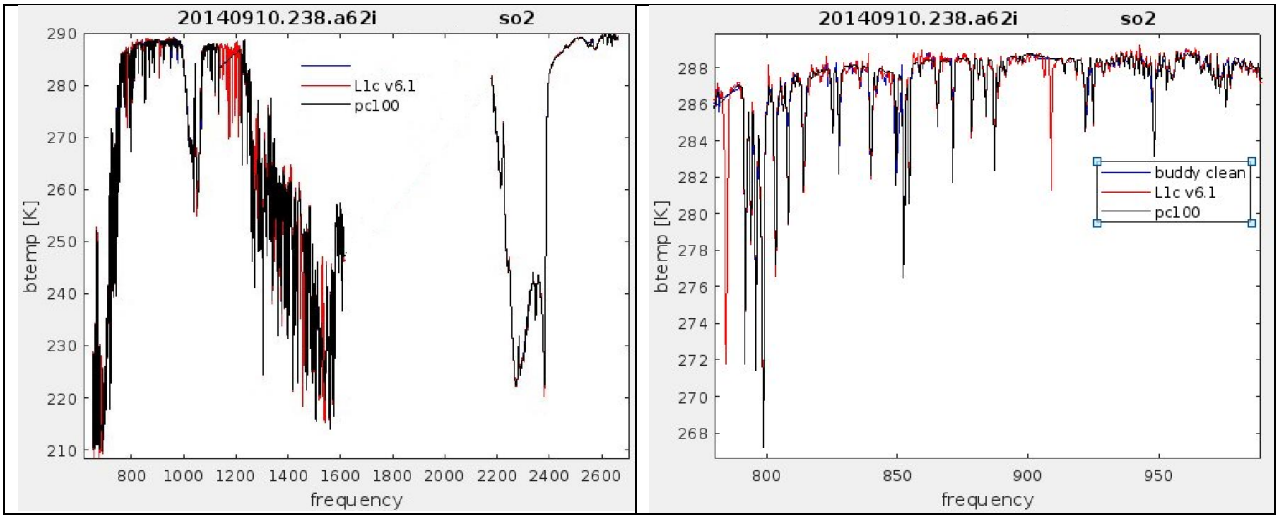

Figure 2. AIRS nighttime spectrum from granule 238 on September 10, 2014.

The AIRS L1C ATBD (this document) describes the algorithms used for converting L1B to L1C. The L1C acceptance test report will separately show that the implementation of the L1C as described in the ATBD meets the L1C functional requirements.

Table 2. AIRS Spectrum Gaps and Overlapped Regions

Type	Frequency Range (cm^{-1})	Comments
Gap	681.99-687.60	Filled by L1C
	781.88-789.26	Filled by L1C
	903.78-911.23	Filled by L1C
	1046.20-1056.07	Filled by L1C
	1136.63-1216.97	Filled by L1C
	1272.59-1284.35	Filled by L1C
	1443.07-1460.27	Filled by L1C
	1527.00-1541.10	Filled by L1C
	1613.86-2181.49	Large gap (not filled by L1C)
	2557.41-2558.53	Filled by L1C
Overlap	727.83-728.36	M-11/M-12
	851.20-852.72	M-08/M-09
	973.48-974.63	M-06/M-07
	1337.64-1339.18	M-03/M-04c
	2301.72-2320.79	M-01b/M-02b
	2545.19-2565.21	M-01a/M-02a

2. Identification and Replacement of Bad Channels

This section describes the identification of bad channels and the bad channel replacement algorithm.

The major steps are:

- Preselection of bad channels.
- First-order replacement of these bad channels using the buddy system.
- Principle Component Reconstruction (PCR) of the spectrum using the first 100 eigenvectors (EV).
- Replacement of the first order identified bad channels by their PCR values.
- Identification of scene-specific impact from scene inhomogeneity (Cij) and PCR replacement of Cij impacted channels.
- Final scan of the spectrum for outliers and their replacement by their PCR values.

2.1. Identification of bad channels

As discussed earlier, 2225 of the 2378 channels were defined at launch as “good”, with 2223 “good” channels as of 9/2019. Of the 155 “bad” channels, 63 were already dead during pre-launch testing. Typically an additional 95 channels exhibit non-Gaussian noise characteristics (referred to as “popping”). For each spectrum typically 155 of 2378 channels with NEDT>1 are preselected for replacement, and about another 95 are marked as “suspect”. This process relies on noise levels, error flags, and radiances from the Level-1B product, so in principle different channels could be identified as bad or suspect in each spectrum. Details on the identification of bad channels are found in Appendix A. The number of bad channels varies slightly over time depending on changes in the A-side/B-side channel weights from periodic table uploads to improve channel performance. The final dynamical detection of "bad" or "noisy" channels with the help of PCR (Principal Component Reconstruction) will be discussed later.

2.2. The Buddy Replacement

The PCR fails when applied directly to the AIRS L1B because of the presence of “bad” channels. In order for the PCR to work, the bad channels have to be replaced in what is referred to as “buddy replacement”. The 2378 channel AIRS spectra contain highly correlated information. We make use of the radiometric correlation between channels in the Buddy Replacement. The Buddy Replacement requires the conversion of the L1B spectra into the Brightness Temperature (BT) domain. The buddies are identified using the same training set as used for the calculation of the EV used in the PCR. Details are included in Appendix B. Each channel is replaced with buddies from the same spectrum (defined as the most correlated channels using minimum standard deviation). When the channel is identified as “bad”, its BT is replaced by a BT calculated from the BTs of most correlated channels in the list of potential buddies.

The correlated channel replacement list is calculated based on minimization of

$$dT(k,j) = \sqrt{\frac{1}{n}\sum_{i}^{n}(T_{ij} - T_{ik})^2} \qquad (1)$$

where:

n is number of spectra in the training set
i is spectrum index, range from 1 to n
j is channel number range from 1 to 2378
k is channel to be filled range from 1 to 2378

The $\delta T(k,j)$ represents the averaged deviation of each individual channel *j* from the filled channel *k*. For each channel *k*, $\delta T(k,j)$ values are sorted in the ascending order and the first 100 *j*s (with the least deviation from channel *k*) are selected and will be used to replace or fill the bad channels. This process is repeated for ten 15 K wide scene BT ranges, 220-235. 235-250, etc. to 355-370 K. The resulting 100 *j*'s (integer array size 2378x100) and the associated deviations (double array size 2378x100) and biases are calculated offline and stored in an ancillary file, read at runtime.

The brightness temperature of the channel to be filled or replaced is the average of the four best-correlated of the buddy channels weighted by the deviations from the filled or replaced channel channels

$$T_k = \frac{\sum_{j}(T_j + fB_r(k,j))\frac{1}{dT_r(k,j)}}{\sum_{j}\frac{1}{dT_r(k,j)}} \qquad (2)$$

Where:

j is the channel number range from 1 to 2378
k is a channel to be filled range from 1 to 2378
r is the brightness temperature range from 1 to 10
B is the brightness temperature bias
f is a bias scale factor.

The use of a bias allows us to find much better matches than otherwise, but there is a catch. The degree of bias between neighboring channels is highly scene dependent: the more clouds, the less spectral contrast. Therefore, for each channel to be replaced and in each spectrum, we first determine a bias scale factor *f*. Then *f* is selected from among nine possible values [0.00, 0.25, 0.50, 0.75, 1.00, 1.25, 1.50, 1.75, 2.00] by finding the value that minimizes the penalized standard deviation of the fill value candidate $T_j + f * Br(k,j)$. The penalty function requires that the standard deviation be 4 times smaller for the extreme bias scale factors of 0.00 and 2.00 than for the nominal value of 1.00, which makes the process favor using the nominal value. The penalty factors for each *f* are [4.00, 3.25, 2.50, 1.75, 1.0, 1.75, 2.50, 3.25, 4.00].

2.3. PCR Replacement

The transmission characteristics of the atmospheric gases creates a high degree of correlation between the channels measuring the upwelling spectra. These correlations can be identified using

eigenvectors (EV). The EV were derived from Radiative Transfer Model (RTM) calculated spectra for a representative set of atmospheric profiles, referred to as "the PC training set". The training set is decomposed in a 2378 EV, each of length 2378. The state of the atmosphere on the day of the training set are frozen in the EV. The EV are static, i.e. the same EV have been used since launch. Details on the creation of the training set are presented in Appendix A.

Each AIRS spectrum can be exactly reproduced by the dot product of 2378 unique Principle Components (PC) with the pre-calculated (static) EV. In practice the spectrum can be reconstructed within the noise with the first 100 PC using the first 100 EV. This is referred to as the pc100 reconstruction in Figure 1. All channels preselected for replacement are replaced by their pc100 values and are marked in an intermediate "L1bc" product in a field named "L1cSynthReason".

2.4 Dynamical Replacement

2.4.1 Dynamical Replacement for scene inhomogeneity effects (Cij)

Each AIRS channel has a slightly different spatial response. This can lead to cases where adjacent channels produce radiances which are inconsistent with the known spectral correlation. The "Inhomo850" indicator flags scenes where spatial inhomogeneity is an issue. In these cases, extra channels may be replaced by their pc100 values. Details of this are contained in Appendix C.

2.4.2 General Dynamic Replacement

Rare cosmic ray or solar energetic particles hits may create outliers (spikes) in detectors that are not preselected as "bad". For all channels and scenes there is a dBT= abs(pc100-L1bc) check to eliminate such outliers. In the majority of spectra, no channels are replaced in this step. An ancillary table provides thresholds for each channel and each 10 K BT range, designed to be 1.25 times the level at which 1 in 1,000 observations from a good channel would be flagged as an outlier of the Gaussian noise distribution. The minimum dBT threshold is set to 2.0 K and several additional adjustments are made:

- Thresholds are raised by an additional factor of 1.5 for the photoconductive arrays M-12 and M-11 (650-728.4 cm-1) because these detectors are of high quality.
- For M-09, M-08, and M-07 (789-974 cm-1) the threshold is set to 2.0 because there are many channels with non-Gaussian noise characteristics.
- In the ozone band (1040-1058 cm-1), the threshold is set to 4.0.

For channels marked suspect in section 2.1 the threshold is lowered 20%.

2.4.3. Broad spectral feature exclusion

When the absolute value of dBT exceeds a threshold, it does not automatically trigger a pc100 replacement. There are cases where the L1B observed radiances may be more correct than the pc100. These are cases where the actual scene has geophysical properties not represented in the training set. This is not an issue for CH_4, CO and SO_2, where climatology abundances were included in the training sets. Known issues include $CaCO_3$ and other desert surface emissivity features and variable minor trace gases which were not included in the training set. The algorithm calculates a metric of how many neighboring channels in the same spectrum have disagreement between observed and reconstructed spectra. A small percentage of cases with high "neighborliness" metric are excluded from dynamic replacement. In testing this is sufficient to preserve signals from $CaCO_3$ and others. There is still some danger that some gases have very narrow spectral lines and may be effectively removed by the dynamic replacement process. The

neighborliness metric looks at dBT for the 20 nearest channels and assigns points for close channels that also deviate, and a bonus for those with the same sign. It is normalized so that 100% corresponds to all nearby channels deviating significantly with the same sign. Any channel which is a candidate for dynamic replacement but which also has a neighborliness metric over 10% is preserved. All other channels for which |dBT| exceeds the threshold, the reconstructed value is substituted and the reason for the substitution is marked in L1cSynthReason.

3. Spectrum Gap Fill and Overlap Elimination

The L1bc is an internal product with 2378 channels per spectrum, with bad channels replaced, but with spectral gaps and spectral overlap regions. The L1C uses L1bc to create spectra without the gaps (except between 1613.86-2181.49 cm^{-1}) in a monotonically increasing sequence of 2645 channels. 2314 of the 2645 channels are identical with the frequencies from L1B and 331 new channels are added to fill in the gaps. 64 channels in spectral overlap areas were eliminated. The gap between 1613.86 and 2181.49 cm^{-1} was excluded. The SRF centroid frequencies of the channels in the gaps increases from nu1 in steps of (nu1+nu2)/4800 to nu2, where nu1 and nu2 are the start and the end frequency of the gap. The frequencies of the channels in the gaps are listed in Table 3. Their SRF widths were defined to be nu/1200.

The same redundancy seen in the 2378 channel spectra which allows bad channels to be replaced, also allows the calculation of the BT in the small gaps between detector modules. The atmospheric states in the ECMWF training set and kCARTA were used to calculate a training set. A buddy-system algorithm simpler than the one discussed in section 2.2 was used to get the best-correlated L1B channels for the gap channels. It is a simple 4-channel substitution. For each of the 331 gap channels we have 4 channel numbers [ch1, ch2, ch3, ch4]. There are 3 coefficients a1, a2, a3 giving weights for the first three channels. The weight for the 4^{th} channel is calculated as a4 = 1.0 – a1 – a2 – a3 so that the sum of the four coefficients is guaranteed to be 1.0. The BT for the gap channel then is simply:

$$BT = a1*BT(ch1) + a2*BT(ch2) + a3*BT(ch3) + a4 * BT(ch4)$$

Figure 3 below shows a typical L1bc spectrum (black) with the gaps filled (red).

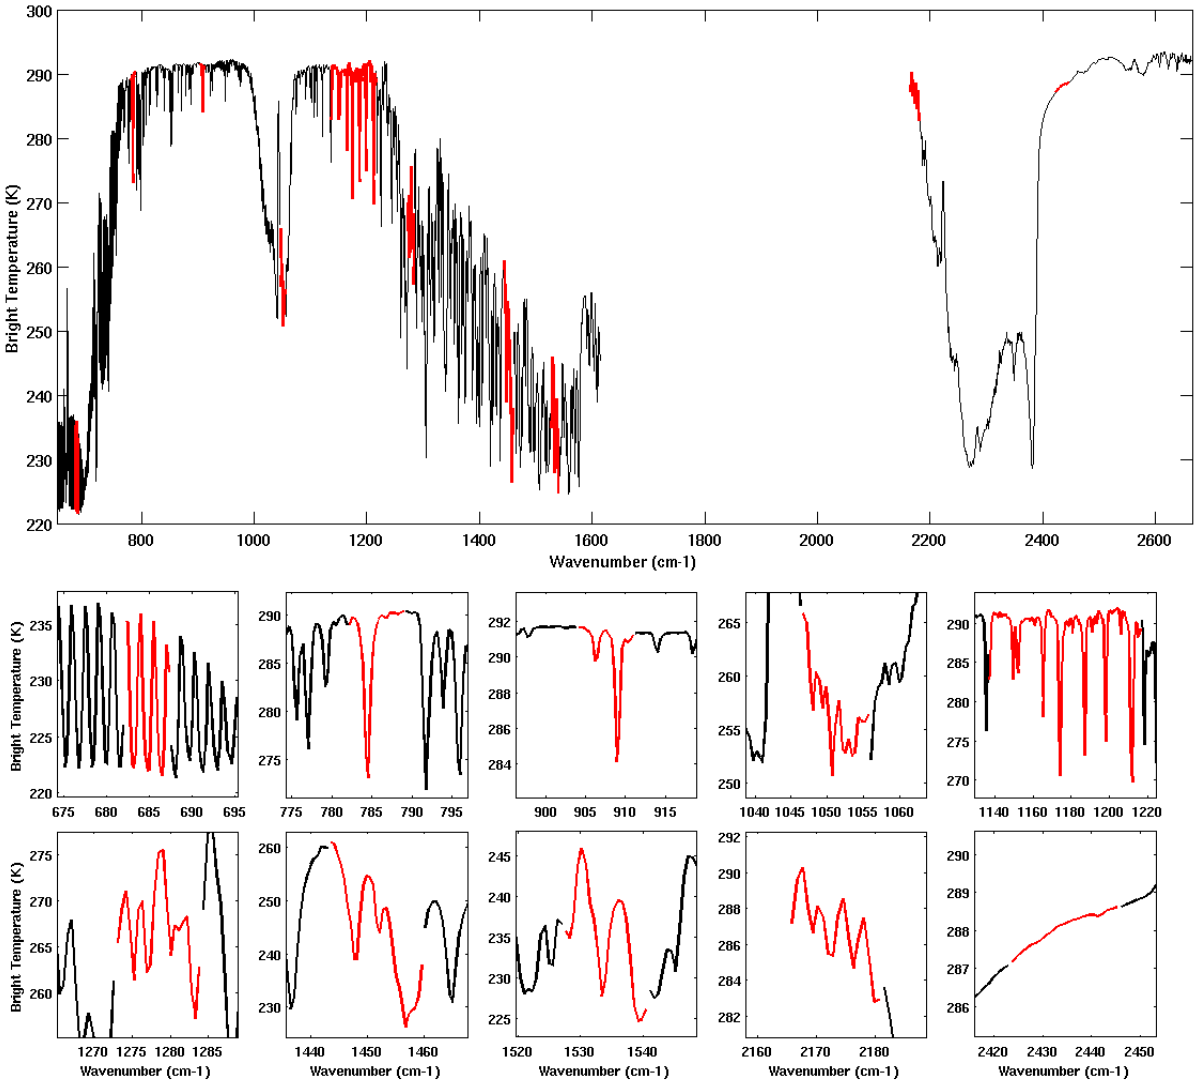

Figure 3. Sample spectrum from the full spectrum training set. The black lines represent the existing channels, and the red lines are the gap channels.

The validation that the cleaning and gap-filling did not produce radiometric artifacts is presented in the L1C acceptance test report.

Table 3. AIRS L1C Synthetic Channel Frequencies (cm^{-1})

Gap #1	Gap #2	Gap #3	Gap #4			Gap #5				Gap #6	Gap #7	Gap #8	Gap #9
682.25	782.22	904.12	1046.65			1137.16	1159.59	1182.03	1204.98	1273.13	1443.65	1527.65	2423.83
682.51	782.56	904.46	1047.10			1137.68	1160.11	1182.55	1205.50	1273.66	1444.23	1528.29	2424.80
682.76	782.89	904.80	1047.55			1138.20	1160.64	1183.07	1206.02	1274.20	1444.80	1528.93	2425.78
683.02	783.23	905.14	1048.00			1138.73	1161.16	1183.59	1206.55	1274.73	1445.37	1529.57	2426.75
683.27	783.56	905.48	1048.45			1139.25	1161.68	1184.11	1207.07	1275.27	1445.95	1530.21	2427.72
683.53	783.90	905.82	1048.90			1139.77	1162.20	1184.63	1207.59	1275.80	1446.52	1530.85	2428.70
683.78	784.23	906.15	1049.34			1140.29	1162.72	1185.16	1208.11	1276.34	1447.09	1531.49	2429.67
684.04	784.57	906.49	1049.79			1140.81	1163.25	1185.68	1208.63	1276.87	1447.66	1532.14	2430.64
684.29	784.90	906.83	1050.24			1141.33	1163.77	1186.20	1209.15	1277.40	1448.24	1532.78	2431.61
684.54	785.24	907.17	1050.69			1141.86	1164.29	1186.72	1209.68	1277.94	1448.81	1533.42	2432.59
684.80	785.58	907.51	1051.14			1142.38	1164.81	1187.24	1210.20	1278.47	1449.38	1534.06	2433.56
685.05	785.91	907.85	1051.59			1142.90	1165.33	1187.76	1210.72	1279.01	1449.96	1534.70	2434.53
685.31	786.25	908.19	1052.04			1143.42	1165.85	1188.29	1211.24	1279.54	1450.53	1535.34	2435.50
685.56	786.58	908.53	1052.49			1143.94	1166.38	1188.81	1211.76	1280.08	1451.10	1535.98	2436.48
685.82	786.92	908.87	1052.94			1144.46	1166.90	1189.33	1212.28	1280.61	1451.68	1536.62	2437.45
686.07	787.25	909.21	1053.39			1144.99	1167.42	1189.85	1212.81	1281.15	1452.25	1537.26	2438.42
686.33	787.59	909.54	1053.83			1145.51	1167.94	1190.37	1213.33	1281.68	1452.82	1537.90	2439.40
686.58	787.92	909.88	1054.28			1146.03	1168.46	1190.89	1213.85	1282.22	1453.39	1538.54	2440.37
686.84	788.26	910.22	1054.73			1146.55	1168.98	1191.42	1214.37	1282.75	1453.97	1539.18	2441.34
687.09	788.60	910.56	1055.18			1147.07	1169.51	1191.94	1214.89	1283.28	1454.54	1539.82	2442.31
687.35	788.93	910.90	1055.63			1147.59	1170.03	1192.46	1215.41	1283.82	1455.11	1540.46	2443.29
						1148.12	1170.55	1192.98	1215.94		1455.69		2444.26
						1148.64	1171.07	1193.50	1216.46		1456.26		2445.23
						1149.16	1171.59	1194.02	1273.13		1456.83		
						1149.68	1172.11	1194.55	1273.66		1457.41		
						1150.20	1172.64	1195.07	1274.20		1457.98		
						1150.72	1173.16	1195.59	1274.73		1458.55		
						1151.25	1173.68	1196.11	1275.27		1459.12		
						1151.77	1174.20	1196.63	1275.80		1459.70		
						1152.29	1174.72	1197.15	1276.34				
						1152.81	1175.24	1197.68	1276.87				
						1153.33	1175.77	1198.20	1277.40				
						1153.85	1176.29	1198.72	1277.94				
						1154.38	1176.81	1199.24	1278.47				
						1154.90	1177.33	1199.76	1279.01				
						1155.42	1177.85	1200.28	1279.54				
						1155.94	1178.37	1200.81	1280.08				
						1156.46	1178.90	1201.33	1280.61				
						1156.98	1179.42	1201.85	1281.15				
						1157.51	1179.94	1202.37	1281.68				
						1158.03	1180.46	1202.89	1282.22				
						1158.55	1180.98	1203.41	1282.75				
						1159.07	1181.50	1203.94	1283.28				
								1204.46	1283.82				

4. Spectral calibration

The L1B products present the radiances on a 2378 channel grid, but the spectral locations of the Spectral Response Function (SRF) centroid frequencies can shift due to second order thermal relaxation effects. These may cause 1) the angle of incidence on the diffraction grating to change, or 2) the projection of the entrance slit on the focal plane to shift, or 3) the focal plane to shift. Shifts on the scale of 1% of the SRF width have been observed since launch. In addition, the rotational motion of the Earth relative to the EOS-AQUA satellite also introduces frequency shifts in the AIRS radiances due to the Doppler Effect. Drifts and seasonal shifts in the AIRS frequencies were observed relatively early in the AIRS mission (Strow et al. 2006).

The final task in creating the L1C product, after cleaning and gap-filling, is to present the radiances on a fixed frequency grid. The resulting changes in brightness temperatures are small, but important at the climate level. First, it is important to understand that most of the L1B radiometric changes due to a shifting frequency scale are only easily measured on the sides of AIRS absorption/emission lines. Most retrieval and data assimilation avoid using this class of channels, since their weighting functions are quite broad. However, for climate trend studies the highest level of stability is needed for all channels. Figure 4 shows the maximum differences seen in the AIRS brightness temperature over the first 17 years of the mission in blue, for a clear-ocean tropical scene. These are mostly dominated by the secular drifts in frequency from 2002 through 2007. As shown here there are also seasonal frequency shifts and shifts during every orbit. The maximum orbital variations (which are always present) are shown in red, and our estimated maximum uncertainties in the AIRS brightness temperatures in the L1C product after shifting the of-date frequencies to a fixed frequency grid is shown in yellow. This last curve is plotted separately in Fig. 9 below showing that the maximum brightness temperature uncertainties are ~0.06K.

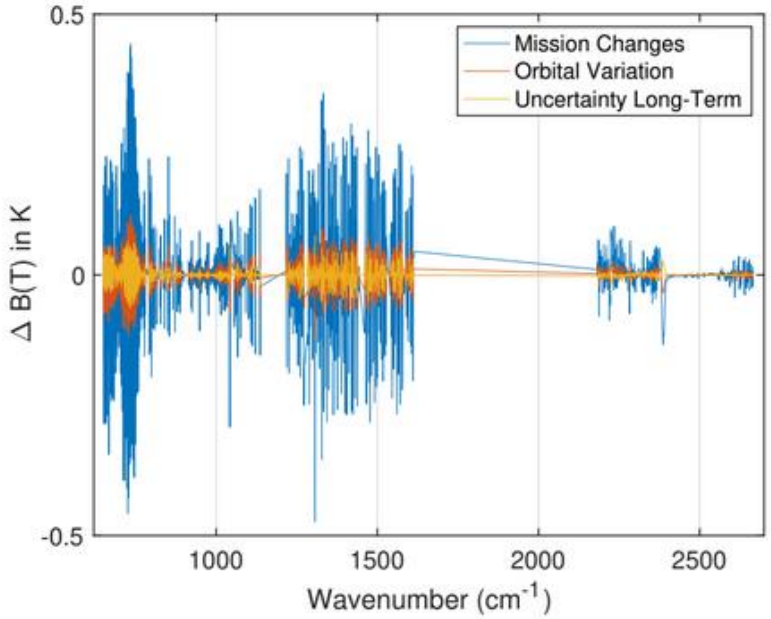

Figure 4. Maximum differences seen in the AIRS brightness temperature between 2002 and 2019.

In the following we outline the approach for determining the of-date spectral frequencies and for presenting then shifted to a fixed frequency grid as L1C radiances.

4.1 The AIRS Spectral Frequency Model

AIRS is a grating array spectrometer. The dispersed image of the spectrometer entrance slit is projected by afocal relays onto the AIRS detectors on 17 detector arrays. The width of the image of the spectrometer entrance slit is 100 µm. The position of each array is defined by its location, y_o, in focal plane coordinates and units of micrometers [µm]. The position of each detector in the array is defined relative to y_o in steps of 50 µm. The conversion for focal plane coordinates to channel frequencies is based on the standard grating equation,

$$\nu_{oi} = m/(d * (\sin(\alpha_{1,2}) + \sin(\beta_i))) \qquad \text{Eq. 3}$$

where m is the grating order, ν_{oi} the center wavenumber of channel i, d = 77.560 μm is the groove spacing, α the angle of incidence and β the angle of diffraction. AIRS has two incident angles, which we denote α_1 (= 0.55278) and α_2 (= 0.56423). The diffraction angle can be expressed in terms of position on the focal plane using

$$\beta_i^k = \tan^{-1}\left(\frac{y_i^k}{F}\right) \qquad \text{Eq. 4}$$

where y_i is the detector position in the focal plan, F is the focal length of the focusing mirror, and k is an index assigned to each detector module. We use the fact that each AIRS detector is 50 μm wide to write y_i^k for each array k as

$$y_i^k = y_o^k + i * 50. \qquad \text{Eq. 5}$$

y_o^k is the position in the dispersed direction of the short wavelength side of each array. Note that the grating order m can vary with k, as does α. Pre-flight testing of AIRS indicated that a very small correction was needed to model the measured channel centroid frequencies given by

$$\nu_i = \nu_{oi} + a_k \times (\nu_{oi} - \nu_k)^2 \qquad \text{Eq. 6}$$

where ν^k is fixed for each array to be the nominal center wavenumber of the array, and the parameters a_k are experimentally determined (L L Strow et al. 2003). These corrections are extremely small, at most 2% of the spectral response function (SRF) width, equivalent to a ~15 ppm fractional change in the centroid frequency. In this document we will often refer to y_o^k as the module Y-offset. The values of y_o^k were measured during pre-flight testing.

4.2. Measurement of Channel Frequencies

Gaiser et al. (2003) found that the early in-orbit AIRS frequencies were offset from pre-flight testing value by the equivalent of a Y-offset shift of approximately -13.4 μm, likely due to a combination of launch vibrations and gravity release. Subsequent examination of the first 29 months of AIRS radiances revealed time-dependent shifts of the AIRS frequencies (Strow et al. 2006). Figure 5 shows the time dependence of the Y-offset values for many of the AIRS detector modules between 2002 and 2010. Static offsets between arrays have been removed. The m4a and m4c time dependencies are plotted separately and offset to highlight their different secular drift compared to the other arrays.

One purpose of the AIRS L1C product is to remove these trends and provide an AIRS Level 1 product with a fixed frequency grid. In this section we review the methodology used to detect AIRS frequency trends and how they are parameterized for removal from the L1B radiances.

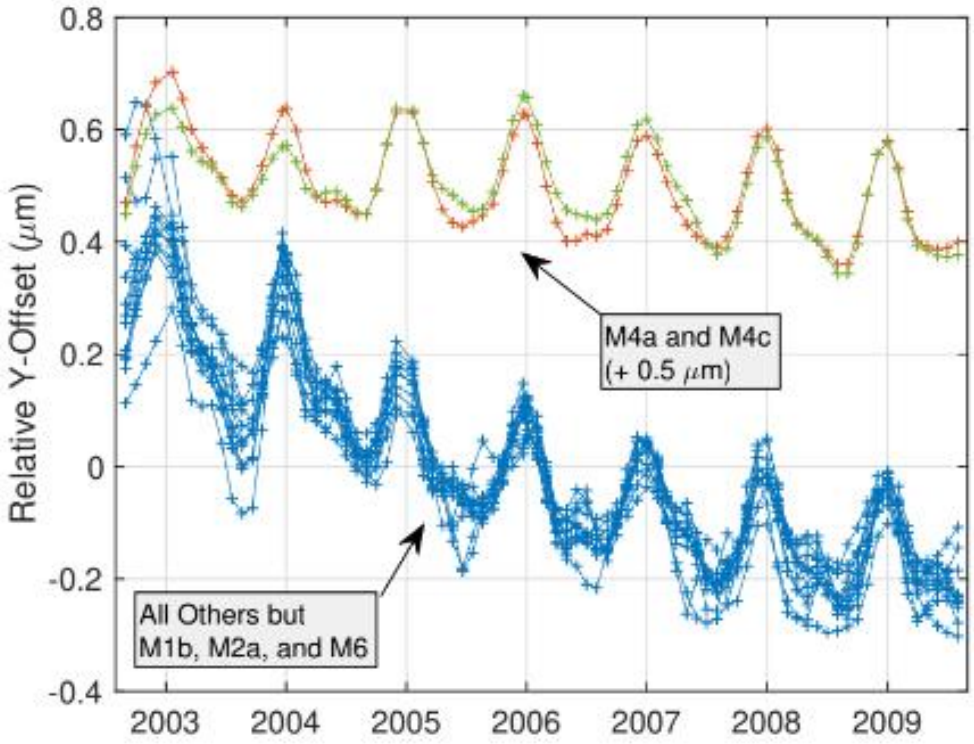

Figure 5. Time dependence of the Y-offset values for many of the AIRS detector modules. Static offsets between arrays have been removed. The m4a and m4c time dependencies are plotted separately and offset to highlight their different secular drift compared to the other arrays.

The in-orbit AIRS frequencies are measured on a per module basis by determining the optimum values of y_o^k. The detector spacing is set at 50 μm during the manufacturing of the arrays. Consequently the channel center frequencies are determined by the values of y_o^k for each of the 17 detector modules. Although the angle of incidence could also change, we subsume any incident angle drifts into the diffracted angle drifts, via the observed y_o^k values. The module names and their nominal start and end wavenumbers are listed in Table1. Under ideal circumstances empirical measurements over time of the y_o^k values for each module would allow computation of the AIRS channel center frequencies via Eq. 6.

The module averaged frequency shifts are measured by cross-correlating simulated spectra with the observed AIRS spectra for each module (Strow et al. 2006). The simulated spectra were computed using the SARTA Radiative Transfer Model (RTM) with time and space collocated atmospheric model fields from either the ECMWF three-hour forecast/analysis or from the ERA-Interim re-analysis fields. The simulated BT spectra were spline interpolated by the equivalent of Y-offset = 0.01 □m until a peak in the cross-correlation between the observed and simulated spectra is detected. The frequency of the simulated spectrum that gives the maximum correlation defines the optimum Y-offset for the module. In practice high-quality measurements are not possible for several modules since they contain insufficient spectral contrast. However, in these cases high-accuracy spectral drift measurements are not required since the spectral contrast is small. Spectral calibration is especially difficult in the m1b modules where there is little contrast, but there is a sharp radiance shoulder from 2380 to 2400 cm^{-1} that can confuse the cross-correlation measurement.

Two types of data sets were used to determine and validate the measured frequency shifts.

- AIRS Level-2 cloud-cleared radiances from mission start (Sept. 2002) through mid-2009 averaged every 2 degrees in latitude in 16-day bins.
- AIRS Level 1C (without frequency calibration) clear-ocean scenes separated into 40 zonally averaged equal-area latitude bins also time-averaged over 16-days.

The model fields were co-located with the AIRS observations for each individual spectral radiance. The simulated radiances were then binned identically to the observed radiances.

The observed AIRS frequency shifts from these data exhibit four types of time-dependent behaviors:

(1) A slow secular drift that largely stopped in ~2009,
(2) A seasonal variation that appears correlated with the AQUA spacecraft solar beta angle,
(3) Periodic drifts that follow the orbit phase, and
(4) Sharp jumps that correlate with spacecraft and/or AIRS events such as the protective shutdown of the AIRS instrument in November 2003 due to concerns about a large solar storm.

For any given 16-day average, the Y-offsets versus orbit phase for all modules had nearly identical periodic variations with an amplitude of roughly 0.3 μm.

Examination of the derived Y-offsets for each module indicated very similar orbit-phase and longer-term time variations for all modules except m4a and m4c. We found that the m3 module, which covers a range of strong water-vapor absorption, gave the most stable set of Y-offset measurements, especially as a function of orbit phase. For example, the peak cross-correlation of this module with simulated spectra was quite high, but more importantly, was nearly constant over all latitudes for both ascending and descending orbits. Consequently, we used the m3 Y-offset data to parameterize the orbit-phase and longer-term time dependence of the AIRS detectors. As discussed below, the long-term time dependence of the m4a and m4c modules was slightly different, so they were parameterized separately.

A per-module mean Y-offset was determined using the 2002 to 2009 data. These offsets were quite small, varying by ~1 μm in focal plane coordinates (equivalent to 1% of the SRF width) and centered roughly between -13 and -14 μm compared to the pre-launch values. Over the full mission we found Y-offsets for all modules to vary at most from -12.4 to -14.3 μm. The AIRS spectrometer grating temperature is servo-controlled, but the grating and other spectrometer optics are not perfectly isothermal. In consequence, the seasonal variation in the spectral calibration has a pattern similar to the AQUA solar beta angle, and the orbital changes are presumably caused by the periodic variation in solar heating during each orbit, resulting in temperature gradients in the spectrometer.

Figure 5 shows the time dependence of the Y-offsets for most modules up to mid-2009. The mean Y-offsets for each module have been removed to illustrate how similar the modules vary in time. Modules m4a and m4c are plotted separately since their secular shift varies more slowly in time (more linear then exponential). Although not clearly evident in this plot, the seasonal variations have shoulders that are similar to the AQUA solar beta angle; later plots will show this more clearly.

Figure 6 brings together the three time dependences of the Y-offsets using module m3 as an example. The secular and seasonal variations are shown versus time on the x-axis, while the short-term periodic variations over the 99 minute orbit are expressed in units of the orbit phase. The small Y-offset shift in March 2014, for example, can be seen in this figure.

AIRS has had several events during the mission, these include the AIRS instrument shutdown in Nov. 2003, and two AIRS anomaly events in January 2010 (28V power supply) and March 2014 (Cooler A SEU) that apparently induced small offsets in the frequency shifts. These were measured using the m3 module, but were also found in other modules, including m10 and m12. These offsets are included in the parameterization of the Y-offsets versus time.

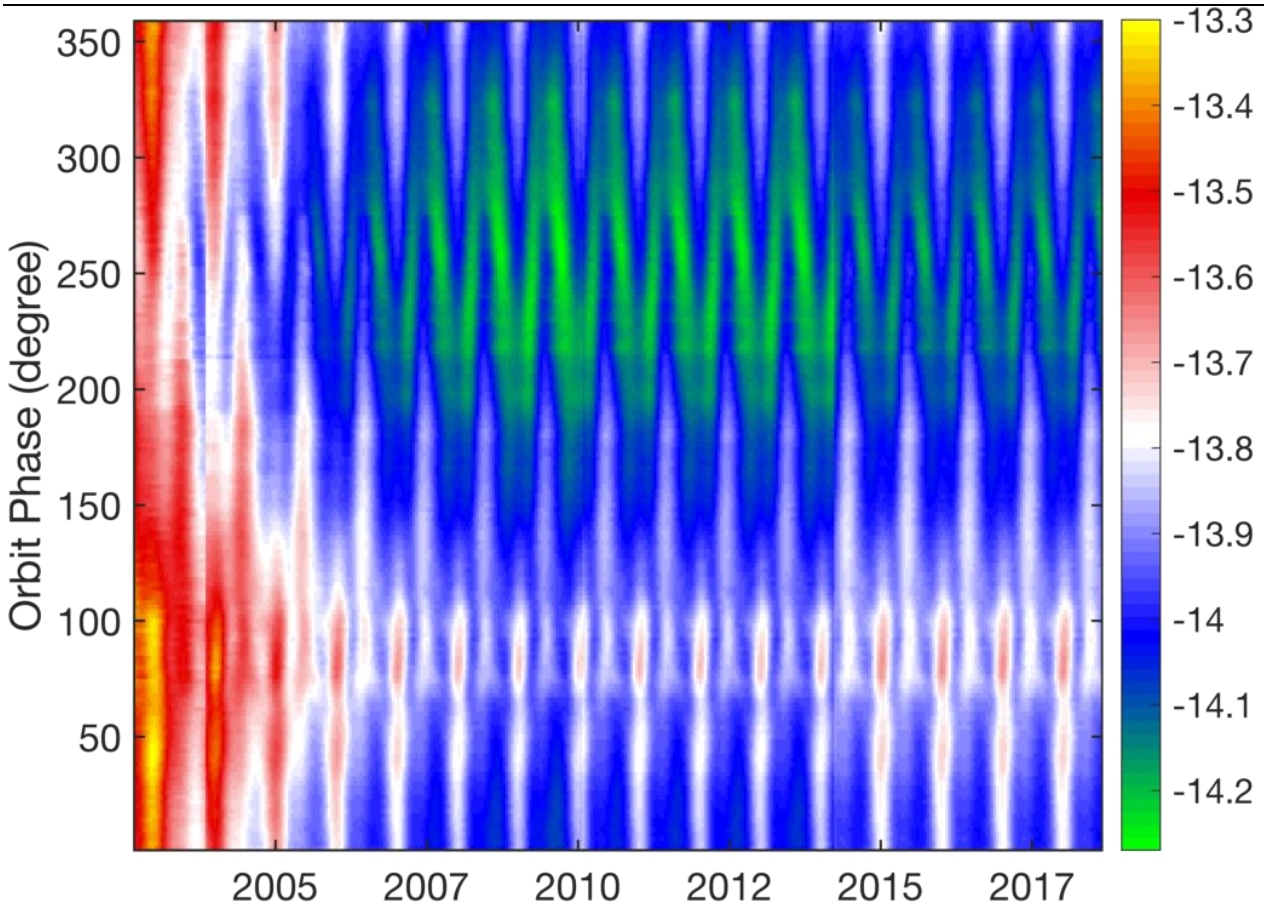

Figure 6. Time dependence of the m3 Y-offsets showing both longer-term time variations and the short-term orbit phase changes. The color scale is the Y-offset in μm.

4.3. Parameterization of the Y-offset Shifts

The 2002 to 2009 date was used to parameterize the shifts of each module versus time and orbit phase as follows. Modules m4a and m4c appear to be changing more linearly in time than the others, and were parameterized as shown below.

$$y_o^k = y_{\text{pre-flight}}^k + a_1^k - a_2 \Delta t + \sum_{i=1}^{2} c_i \sin(2i\pi\Delta t + d_i) \qquad \text{Eq.7}$$

where we set $y_{pre-flight}^k$ to -14 μm. The other modules were also parameterized with Eq. 7 for the pre-Nov. 2003 data only. The time span of the pre-Nov. 2003 data was insufficient for detecting/modelling with an exponential, which is used for data after Nov. 2003.

The a_1^k are the time-independent offsets of each module relative to the pre-flight value. All the a_1^k have the same relative dependence on orbit phase, so they only differ by a constant. a_2 is the amplitude of the module Y-offset drift where Δt is either the time since launch, or the time since the Nov. 2003 AIRS shutdown. Therefore there are two distinct sets of coefficients, mission start to Nov. 2003, Nov. 2003 to the present. The sin terms, $c_{1,2}$, and $d_{1,2}$, are identical for all arrays and only vary slightly with orbit phase. For the other modules after Nov. 2003 we switch to Greek notation for the parameterization constants, and now use an exponential term for the secular drift:

$$y_o^k = y_{\text{pre-flight}}^k + \alpha_1^k - \alpha_2 \exp^{-\Delta t/\alpha_3} + \sum_{i=1}^{3} \beta_i \sin(2i\pi\Delta t + \gamma_i) \qquad \text{Eq. 8}$$

As in Eq. 7 the α_1^k parameters are time-independent offsets of each module but with identical dependence on orbit phase. The α_2 parameters vary slightly with orbit phase but are identical for all arrays (except m4a and m4c which use Eq. 7). The decay rate of the exponential, α_3, is the same for all orbit phases and modules. The orbit phase dependence of the higher order terms ($\beta_{1\text{-}3}$ and $\gamma_{1\text{-}3}$) are identical for all modules, and again vary only slightly with orbit phase. The orbit phase dependence of the sin terms are identical for all arrays and mostly only modify the ascending part of the orbit (when solar radiation is hitting the spacecraft).

In summary, with regard to orbit phase dependence, all terms except the α_2 and α_2 parameters depend on orbit phase, identically for all modules. However, all but about 20% of the orbit phase dependence are provided by the α_1^k and α_1^k parameters.

Figure 7 shows an example of the time dependence of y_o^{m3} near the equator for both descending and ascending passes. Most evident is the steep secular negative increase in early years that levels out in the 2009 time range. Also evident are the seasonal cycles and the inversion of the amplitudes between ascending and descending. There are small "shoulders" on the nearly sinusoidal oscillation in the April-June time that one sees in a plot of the AQUA solar beta angle, suggesting that solar heating of the spacecraft is dominating the seasonal variability. Although difficult to see in this plot, small shifts due to AIRS shutdowns occur in Nov. 2003, Jan. 2010, and March 2014.

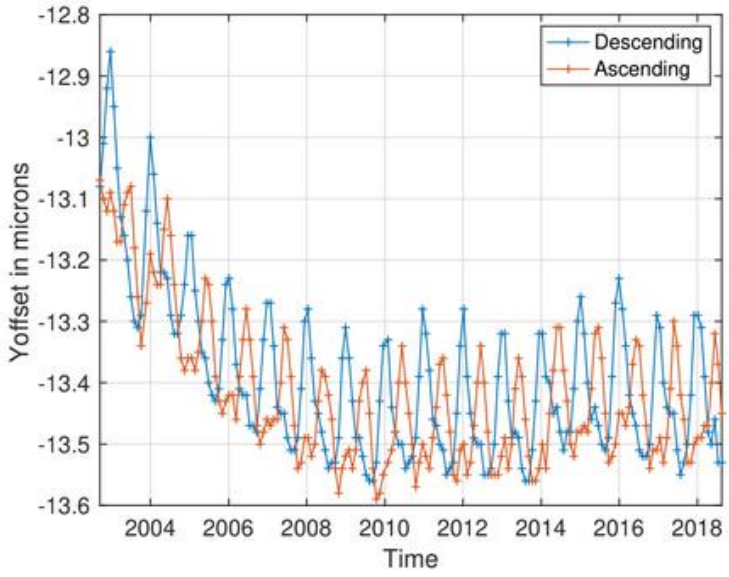

Figure 7. m3 Y-offsets near the equator versus time for both descending and ascending nodes.

4.4. Conversion of AIRS radiances to a fixed grid

The fixed frequency grid for the AIRS L1c product was chosen to be our measurement of the AIRS observed frequency grid on January 20, 2010 at the first descending equator crossing. This is a somewhat arbitrary choice. Factors for this choice include:

(1) By this date the secular shift of AIRS had pretty much stopped, and
(2) Just before this date AIRS had a shutdown, which caused a small 0.04 μm Y-offset shift, and thus we avoid interpolations for this shift after that date.

Given an accurate model for the frequency shifts of each channel in AIRS, we can then interpolate the measured spectra with a time-varying frequency grid to a fixed frequency grid. This is an essential step for accurate studies of radiance trends, and allows the use of a single frequency grid radiative transfer model for AIRS retrievals.

The frequency shifts required to put the AIRS radiances on a fixed grid are quite small, with maximum shifts of ~10 ppm in frequency, 1.2 μm in focal plane coordinates. We noted earlier that the

width of the image of the spectrometer entrance slit is 100 μm wide when projected on the array of 50 μm wide detectors. This means that the SRF is nearly Nyquist sampled and can be shifted by interpolation. This could be done by interpolating the radiances on the observed frequency scale (from the above discussed model) to a fixed frequency scale using a cubic spline. The spline interpolation to the fixed frequency grid is improved by modifying the spline derivative using small regression coefficients as follows,

$$BT_{l1c} = BT_{obs} + \left[a\left(\frac{\partial BT}{\partial \nu}\right) + b\right]\delta\nu. \qquad \text{Eq. 9}$$

where a and b are linear regression coefficients, for each channel, that are derived from simulations. BT_{l1c} is the correct brightness temperature for the fixed frequency Level 1c grid, BT_{obs} is the observed brightness temperature, and $\delta\nu$ is the frequency difference between the observed and fixed grids. The BT derivative is computed from

$$\frac{\partial BT}{\partial \nu} \equiv \frac{BT_{obs} - BT_s}{\delta\nu} \qquad \text{Eq. 10}$$

where BT_s is a spline fit to BT_{obs} evaluated at the ν_{l1c} freqencies.

Simulated data were created using the kCARTA pseudo line-by-line radiative transfer algorithm (Strow et al. 1998) for 49 highly variable profiles that largely span global profile variability. (These profiles are used to create the AIRS SARTA forward model (Strow et al. 2003)). The monochromatic simulations are convolved with the AIRS spectral response functions for -13 and -14 μm frequency offsets. Then Eq. is used with the simulated data to regress for a and b for each channel.

These regression coefficients reduce the application of Eq. 9 with simulated data from errors of up to 0.05K in BT to errors well below 0.01K. The channel averages of the "a" coefficients are close to unity with a standard deviation of 0.6. The channel average of the "b" coefficients is near zero with a standard deviation of 0.9. Note that the b coefficients are multiplied by the frequency offset between the observed and fixed frequency grid, generally very small numbers in the in the range of 0.005 to 0.025 cm^{-1}. On average the a and b coefficients modify the BT interpolation by ~14% and ~7% respectively. Given that the AIRS detector noise values are nominally around 0.2K, the frequency shift corrections are only seen in averaged data, especially as a function of time.

The left panel of Figure 8 shows the observed m3 Level 1B Y-offsets over time for a range of orbit phases, the node type and latitude locations are shown in the label. The right panel shows the same as on the left, but for the Level 1C data after adjustment for frequency drifts. While the L1B frequencies shifted from -12.9 to -13.5 μm with a 0.4 μm p-p seasonal modulation, the effective Y-offset is stable at -13.3 μm with about 0.1 μm residual seasonal modulation.

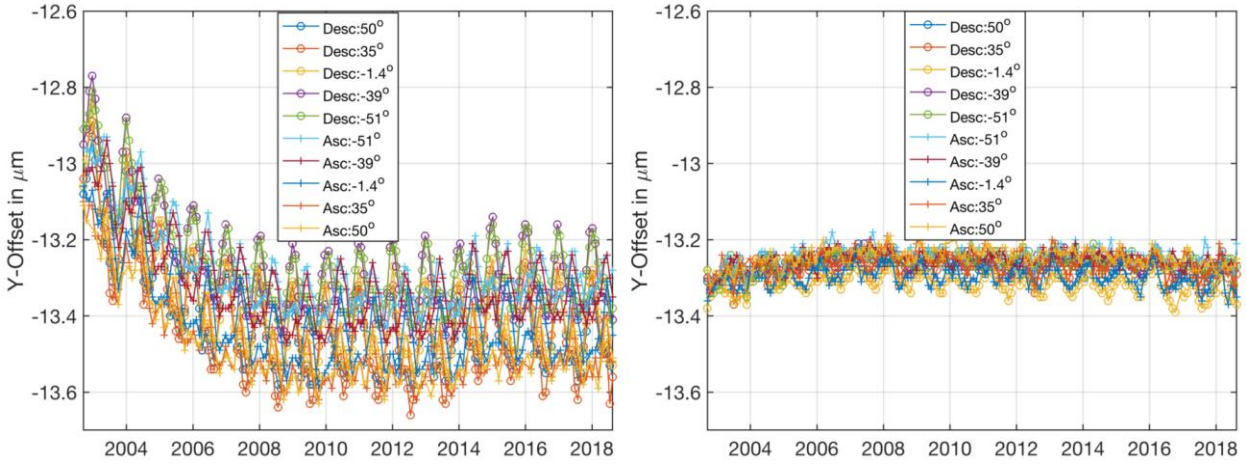

Figure 8. Observed (left) and frequency stabilized (right) Y-offsets over time for a range of orbit phases, the node type and latitude locations are shown in the label.

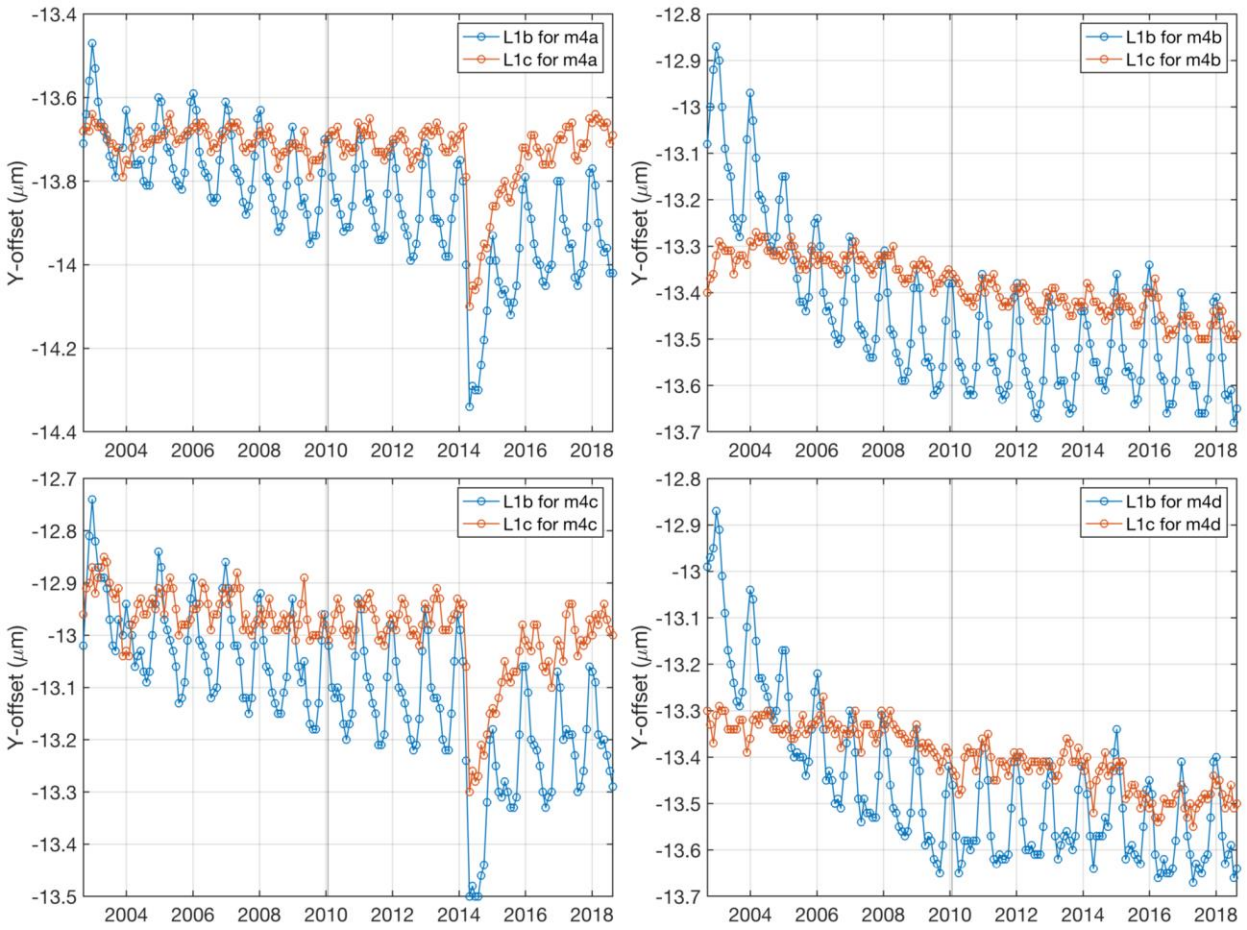

Figure 9. Observed Y-offset shifts for the m4 arrays before and after Level 1c frequency calibration, descending node data near the equator. Note the un-corrected shifts in the Level 1C product in modules m4a and m4c caused by the March 2014 AIRS focal plane warm-up.

Detector modules M4a and M4c show an anomalous behavior. Figure 9, left hand side, shows both the uncalibrated L1B and calibrated L1C frequency shifts. However, there is a striking offset in these two modules that starts directly after the March 2014 AIRS focal plane warm-up. Apparently this shift recovers after about two years. At this time, we have not implemented any fix for these frequency shifts in L1C. There is the possibility that these apparent frequency shifts are not true frequency shifts but radiometric offsets that corrupt the cross-correlation frequency shift measurement algorithm. More work is needed to fully understand this issue and correct any offsets.

Figure 9 also shows the M4b and M4d modules (right hand side). Although the frequency calibration of these two modules greatly improves the frequency stability, it is clear that a small drift down in Y-offset is present, with a total drift of about 0.2 □m over 17 years. This is small enough to be ignored for most applications, but we will re-visit this issue in future updates to the L1C frequency calibration.

The combination of uncertainties in the frequency determination and the interpolation algorithm used to shift the frequencies to a fixed grid are shown in Figure 10 (same as the yellow points in Fig. 4). It shows that the maximum brightness temperature uncertainties are ~0.06K. These uncertainties are well below the estimated 100 mK absolute radiometric uncertainties (Pagano et al. 2018).

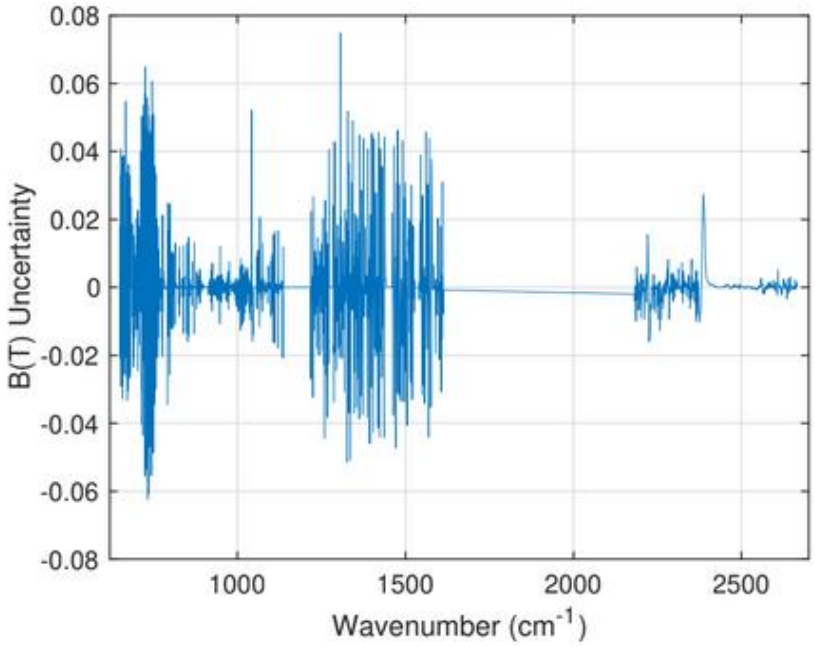

Figure 10 details of the yellow points in Fig. 4, for the final estimated uncertainty in the AIRS brightness temperature for insufficient correction for AIRS instrument frequency drifts.

Details of the validation of the frequency determination and the radiometric shift algorithms are presented in the L1C acceptance test report. It should be pointed out that the overall radiometric uncertainty of the AIRS radiances in the L1C product is the combination of the frequency shift uncertainty and the radiometric calibration uncertainty.

5. Summary

The AIRS L1C product corrects for effects that can make the L1B product difficult to use. The Level 1C product supports users who depend on the availability of outlier-free and gap-free spectra with monotonically increasing frequency for inter-instrument comparisons. An algorithm has been presented that corrects for the effect of static and dynamic (non-Gaussian) noise in AIRS L1B products. The algorithm also fills small gaps in these spectra and creates spectra monotonic in frequency with 2645 channels and only one (instrument design related) gap between 1650 and 2175 cm^{-1}. "Bad" L1B channels are filled using their PCR values. Gap channels are filled with a simpler buddy algorithm. The calibrated radiances are shifted to a fixed frequency grid to create the AIRS Level 1C product. This product will be made available to the public at the GES DISC. It is recommended that new users of the AIRS radiance product use the L1C product.

References

Aumann, H.H. , M.T. Chahine, C. Gautier, M. Goldberg, E. Kalnay, L. McMillin, H. Revercomb , P.W. Rosenkranz , W. L. Smith , D. H. Staelin, L. Strow and J. Susskind, (2003), "AIRS/AMSU/HSB on the Aqua Mission: Design, Science Objectives, Data Products and Processing Systems", IEEE Trans. Geosci. Remote Sens., Vol.41, 253-264.

Aumann, H. H. and 17 coauthors (2018) "Evaluation of Radiative Transfer Models with Clouds" JGR Atmospheres doi: 10.1029/2017JD028063

Chen, Yong, Yong Han, and Fuzhong Weng. 2013. "Detection of Earth-Rotation Doppler Shift from Suomi National Polar-Orbiting Partnership Cross-Track Infrared Sounder." *Appl. Opt.* 52 (25): 6250–7. https://doi.org/10.1364/AO.52.006250.

Clerbaux, C., Boynard, A., Clarisse, L., George, M., Hadji-Lazaro, J., Herbin, H., Hurtmans, D., Pommier, M., Razavi, A., Turquety, S., Wespes, C., and Coheur, P.-F.: Monitoring of atmospheric composition using the thermal infrared IASI/MetOp sounder, *Atmos. Chem. Phys.*, **9**, 6041–6054, doi:10.5194/acp-9-6041-2009 (2009).

Gaiser, S. L., Hartmut H. Aumann, L. Larrabee Strow, Scott E. Hannon, and Margaret Weiler, In-Flight Spectral Calibration of the Atmospheric Infrared Sounder, IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING, VOL. 41, NO. 2 (2003).

Jolliffe, I. T.,Principal Component Analysis. Springer-Verlag. pp. 487. doi:10.1007/b98835. ISBN 978-0-387-95442-4 (1986).

Motteler, Howard E., and L. Larrabee Strow. 2018. "AIRS Deconvolution and the Translation of AIRS-to-CrIS Radiances With Applications for the IR Climate Record." *IEEE Transactions on Geoscience and Remote Sensing* 57: 1793–1803. https://doi.org/10.1109/tgrs.2018.2869170.

Strow, L. Larrabee, Howard E. Motteler, Robert G. Benson, Scott E. Hannon, and Sergio De Souza-Machado. 1998. "Fast Computation of Monochromatic Infrared Atmospheric Transmittances Using Compressed Look-up Tables." *Journal of Quantitative Spectroscopy and Radiative Transfer* 59 (3-5): 481–93. https://doi.org/10.1016/S0022-4073(97)00169-6

Strow, L. L., Scott E. Hannon, Margaret Weiler, Kenneth Overoye, Steven L. Gaiser, and Hartmut H. Aumann, Prelaunch Spectral Calibration of the Atmospheric Infrared Sounder (AIRS), IEEE TRANSACTIONS ON GEOSCIENCE AND REMOTE SENSING, VOL. 41, NO. 2 (2003)

Strow, L. L., S. E. Hannon, S. De-Souza Machado, H. E. Motteler, and D. C. Tobin. 2006. "Validation of the Atmospheric Infrared Sounder Radiative Transfer Algorithm." *Journal of Geophysical Research: Atmospheres (1984–2012)* 111 (D9). https://doi.org/10.1029/2005JD006146.

Weiler , M. H., Kenneth R. Overoye, James A. Stobie, Paul B. O'Sullivan, Steven L. Gaiser, Steven E. Broberg, and Denis A. Elliott, Performance of the Atmospheric Infrared Sounder (AIRS) in the Radiation Environment of Low-Earth Orbit, SPIE International Symposium on Optical Science and Technology, Vol. 5882, San Diego (2005).

Appendix A: Principal Component Training Set

The training set used for principal component analysis (PCA) is a composite of two separate sets created by RTM from UMBC and NASA Langley.

A.1: UMBC PC training set

A training set was received from Larrabee Strow and Sergio DeSouza of UMBC and placed at JPL in /asl/data/rtprod_airs/2009/03/01/JPL2834/. This set represents an entire day of simulated data with 240 six-minute granule files in RTP format. The state of the atmosphere and clouds was taken from the collocated ECMWF data. The RTM converted the clear and cloudy data into AIRS spectra using SARTA (Aumann et al. 2018). In order to accommodate the use of this data for training a PC set which can be applied to AIRS data throughout the AIRS mission, spectra in the set have CO_2 levels randomly selected between 300 and 400 ppmv, and the SRF centroids are shifted randomly by up to 10% of the SRF width. For CH_4, CO and SO_2 the RTM used climatology abundances. These spectra have 2834 channels representing the 2378 AIRS instrument channels and 456 synthetic channels in spectral gaps. For the generation of this training set we used only the original 2378 channels.

The IDL program in /home/evan/L1C/global_spectra2.pro was used to extract from this day of data a PC training set of about 10,000 representative spectra. Details of this procedure are given in Aumann, H. H. et al. 2018. Note that this procedure creates a maximum variance data set and makes no explicit provision to ensure that the training set represents the full range of scan angles, CO_2 levels, or spectral shift. The resulting training set is train2.2009-03-01f.h5, with 7377 spectra in 100-spectrum objects.

Appendix A.2: Langley PC training set

A training set was received from Dan Zhou of NASA Langley Research Center and placed at JPL in /home/hha/DanZhou.AIRS.simulation/. It is organized in 84 files with names in the pattern "AIRS_RAD_xx_y1_y2.bin". xx is the satellite zenith angle for all spectra in that file and [y1, y2] is a range of Sun angles. The RTM used CRTM, which can handle clear and cloudy conditions (Aumann et al. 2018) for 2378 channels at the nominal fixed SRF centroids. Each file contains 35043 spectra, representing the same 35043 representative geophysical states. The first 26600 spectra are under clear conditions, and the rest are under cloudy conditions. An ancillary file named "data_info.asc" contains lat. (col #1), long. (col #2) and cloud optical depth (col #3).

The IDL program in /home/evan/L1C/zhou_training2.pro was used to extract a maximum variance PC training set of <~10,000 spectra using the same procedure as discussed above for SARTA. The resulting training set is zhoutrain2.h5, with 13021 spectra in 100-spectrum objects.

Appendix B: Static channel quality check

The flagging of channels statically (without comparison to a PC reconstructed spectrum) is based on a series of conditions, detailed below. Key inputs are noise levels, quality flags, and whether or not a channel has a fill value from L1B. A configuration file flags additional channels for exclusion when they would not be caught automatically. Some of these channels are marked permanently bad because they have been unusable since the launch of the AIRS.

The NEdT is used as the primary quality check. Level-1B determines noise from the jitter among the downlinked observations of space and of the blackbody for every data granule (6 minutes). A NEdT of -9999 is a flag value, indicating that there were no valid observations for one or both of these calibration sources. These are referred to as "dead" channels.

Channel replacement is required whenever:

1) Its NEdT at 250 K scene temperature exceeds a threshold (currently 0.85 K). Channels with more than 0.85 K of noise are likely to have >3 K outliers just from Gaussian distribution statistics.
2) Its NEdT at 250 K exceeds a threshold (currently 3.0) times the baseline NEdT. The baseline NEdTs are a fit to the best channels in a module (contained in a fixed ancillary file). The baseline is multiplied by an additional factor of sqrt(2) for channels in A-Only or B-Only modes. As with check #1, those channels with higher NEdT compared to a baseline are more likely to be non-Gaussian or biased.
3) NEdT is negative (indicating noise could not be characterized).
4) L1B provides no calibrated radiance value, i.e. radiance value in L1B is the flag value of -9999. In addition to the same cases caught by the negative NEdT test, this test will flag cases where the individual readings have saturated, caused either by cosmic ray hits or by bright glints.
5) L1B radiance is out of a configurable range of expected geophysical values. Currently this range is [170, 420] K, with an additional margin of 5 times the channel's noise level. This test is also designed to catch extreme values from cosmic ray hits.
6) Channels are on a list of permanently or temporarily bad detectors. These include five cross-wired channels and others that have been determined in testing to have undesirable characteristics but are not flagged according to the criteria above.

Channels not meeting the above conditions are instead marked as suspect if they meet any of the following conditions. Suspect channels are not automatically replaced but are excluded from use in replacing other channels and have a lower threshold for dynamic replacement:

1) NEdT at 250 K exceeds a threshold of 0.70 K.
2) NEdT at 250 K exceeds a threshold of 1.75 times the baseline NEdT for the module and channel A/B state.
3) Observed radiance is negative. Negative radiances are expected in the shortwave region for very cold scenes, but are still suspect.
4) L1B CalFlag indicates a problem with gain or offset calculations, bad telemetry, or a "pop" event. When a pop occurs the zero level from the detector has changed between the observation of space before the current scan and the one after. There is therefore an additional uncertainty in the calibration.
5) AB_State from the current channel properties file marks the channel lower quality with state > 2. These states have been judged to be low quality by the AIRS calibration team.
6) Cij in channel properties file is less than a specified value (currently 0.92) indicating the channel is not well aligned spatially with the nominal boresight.

Appendix C: Cij replacement.

All instantaneous spatial responses are basically circular and coaxial, but within this circle some channels are more sensitive on one side and others on the other. These differences are generally quantified in 2 ways, first with X and Y centroids, giving the offset of the center-of-mass for a given channel from the nominal boresight in each direction. The second way is with Cij, the coregistration of the *i*-th and *j*-th channels. The coregistration is reported relative to channel #2113.

The spatial misalignment matters only when scenes are spatially inhomogeneous. 'Cij' is used below as shorthand for the distortion of the spectra caused by this inhomogeneity.

The spatial sensitivity of each channel is related to how the sensors are arranged on the focal plane of the AIRS instrument. Detectors are physically located on 17 separate detector modules. Afocal relays and mirrors project the dispersed image of the spectrometer entrance slit onto different detectors. In some cases spectrally adjacent channels are physically separated into different detector modules. Cij effects are most pronounced when we look across these boundaries.

The magnitude of Cij distortion of varies greatly from scene to scene, so it is necessary to have a metric of Cij in order to apply tight correction criteria only to Cij-impacted scenes. Cij effects are strongest for channels at the longwave end of the M-08 detector module, with frequencies near 850 cm^{-1}, so this region is used for the metric. We define metric Inhomo850 as follows:

1. Determine the ten "good" channels closest to the longwave end of M-08. Good channels are channels which are not bad according to the criteria in section 2.1 and also are not suspect, unless they are only suspect because of their Cij values.
2. Calculate the mean over these ten channels of (BT1B – BTpca). Call this dBTm8.
3. Calculate the mean over these ten channels of BTpca. Call this BTm8.
4. Determine the ten "good" channels closest to the shortwave end of M-09.
5. Calculate the mean over these ten channels of (BT1B – BTpca). Call this dBTm9.
6. Calculate the mean over these ten channels of BTpca. Call this BTm9.
7. Calculate mean BT near 850: BT850 = (BTm8 + BTm9) / 2
8. RawInhomo850 is dBTm9- dBTm8.
9. In order to prevent Inhomo850 from getting too large for cold scenes, calculate:
   CijFactor = dBdT(BT850, 850) / dBdT(250, 850)
   Where dBdT is the slope of the Planck function for the given BT and frequency.
   If BT850 is > 250 K, set CijFactor to 1.0
10. Inhomo850 = RawInhomo850 * CijFactor

Inhomo850 is a mean brightness temperature difference in K, and represents the mean magnitude of the Cij effect over the most impacted channels. It is output in the L1C product files. Any scene with |Inhomo850| > ~0.84 K is considered Cij-impacted and any scene with |Inhomo850| > ~1.69 K is strongly impacted.

The dynamic check for Cij related outliers is applied only to the approximately 5% of scenes where |Inhomo850| exceeds a threshold (currently 0.84 K). For scenes that are not extremely inhomogeneous (Currently 0.84 K < |Inhomo850| < 2.96), these tests apply only to channels known to be susceptible to Cij and that are currently marked with an "AB_State" that puts the channel in a vulnerable condition (Table E1). For scenes that are extremely inhomogeneous (currently 2.96 < |Inhomo850|), all channels are subject to these tests. In these cases the entire spectrum is so impacted that the resulting spectrum should be used only with caution.

Table C1 gives the Cij-sensitive frequency ranges. These are generally near the ends of detector modules where the optical path is furthest from the center. In some cases not all channels within a frequency range are vulnerable, only those with a particular set of 'AB_State's. For 745-785 cm^{-1}, there is little risk for channels in AB_State 0, indicating a good channel using both the A and B detectors, but channels that either use only one side or are of lower quality are vulnerable. In the remaining ranges all channels are sensitive no matter what AB_State is. Figure C1 shows the sensitive regions superimposed on a typical spectrum.

Table C1. Cij frequency ranges

Frequency range (cm^{-1})	AB_State conditions	Modules
745 - 845	AB_State not 0	Shortwave end of M-10, Longwave end of M-09
845 - 877		Shortwave end of M-09, Longwave end of M-08
877 - 895	AB_State not 0	Mid M-08
895 - 925		Shortwave end of M-08, Longwave end of M-07
925 - 970	AB_State not 0	Mid M-07
970 - 986		Shortwave end of M-07, Longwave end of M-06
986 - 1140	AB_State not 0	Shortwave end of M-06, All of M-05
1200 - 1225		Longwave end of M-04d
1338 - 1360		Longwave end of M-03

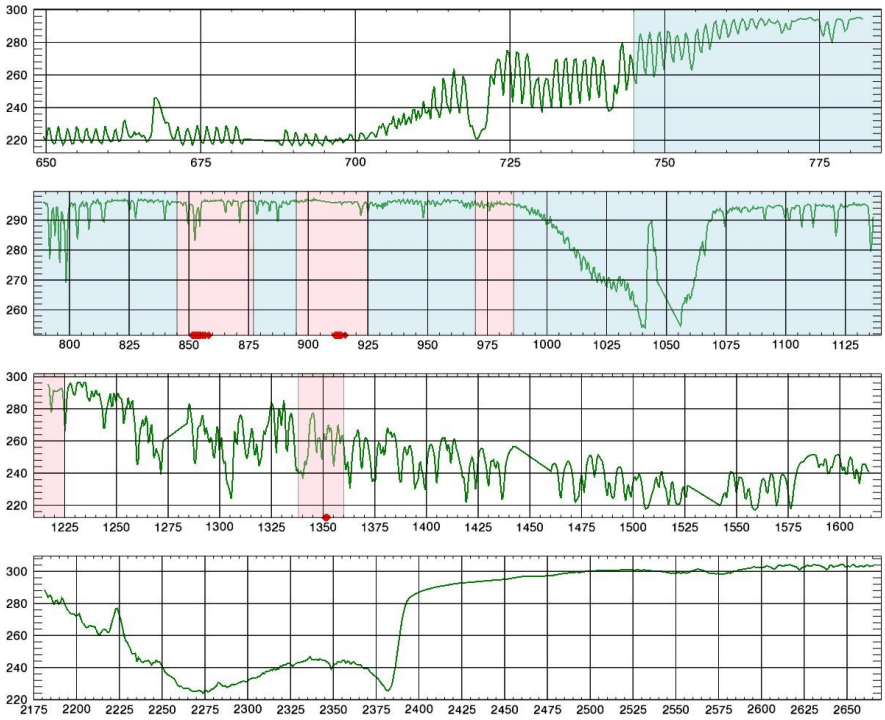

Figure C1. Cij-sensitive spectral ranges. Pink shading: all channels are sensitive. Blue shading: A-Only and B-Only channels are sensitive but not A+B. Red diamonds: the most sensitive channels.

Figure C2 shows the data on which Table C1 is based. One day of data was compared to PC reconstruction. The margin needed to ensure only 1 in 100,000 valid observations are flagged was calculated, separately for non-Cij scenes and for scenes with significant Cij. The difference is a measure of how important Cij is for each channel. Because channels can change between A-only,

B-only, and A+B states as gain tables are uploaded over the mission, Cij sensitivity in each frequency range is keyed to current A/B state instead of static.

Note that while the shortwave region with frequency > 2375 cm^{-1} shows some symptoms of Cij in Figure C2, it is not included in the frequency ranges listed in Table E1. This is because other testing suggests that the main problems here are either poor representation of solar effects and desert emissivity in the PCA training set or instrument calibration problems, not Cij.

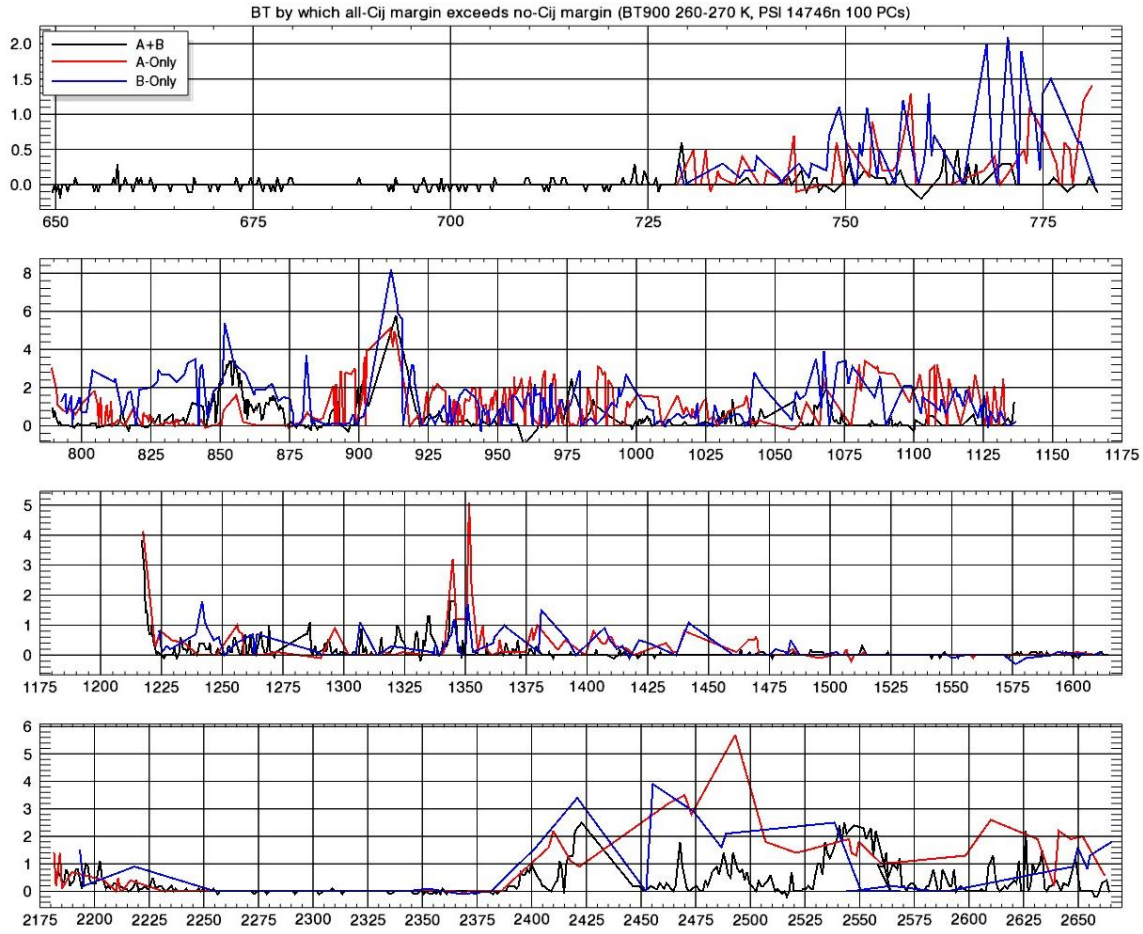

Figure C2. Estimated degradation from Cij effects depending on whether channels are A-Only, B-Only, or A+B

For scenes and channels that fit the criteria for Cij, we set the dBT threshold BT_thresh_Cij to 1.0 K if |Inhomo850| is in [0.84, 1.69) K or 0.7 K if |Inhomo850| is >1.69 K. Channels are flagged for replacement with reconstructed values if dBT > BT_thresh_Cij / CijFactor. For scenes with low spatial inhomogeneity (0.28 < |inhomo850| < 0.84), only the most Cij-sensitive channels (red diamonds in figure E1) are subject to this test.