CIPS RAA Level 2A Data
Last updated February 2023

1. Introduction

CIPS Rayleigh Albedo Anomaly (RAA) files are provided for levels 2A, 2B, 2C, and 3A. At the current time, the retrieval version is being updated from v1.10r06 to v1.10r07. This document describes the level 2A data product for v1.10r07. The AIM Calibration and Measurement Algorithm Document (CMAD) provides a detailed description of the RAA retrieval algorithm, so users should consult the CMAD for retrieval details (see https://lasp.colorado.edu/aim/documentation). Users should also read the CIPS instrument and data overviews (also available on the documentation web page).

Briefly, CIPS measures the single-scatter albedo at 265 nm, which in the absence of polar mesospheric clouds (PMCs), mesospheric clouds at mid-latitudes, and possibly space debris, is due entirely to atmospheric Rayleigh scattering. The source function of the observed Rayleigh scattering peaks near an altitude of 50–55 km (*Bailey et al*., 2009). The Rayleigh scattering albedo at 265 nm is controlled by the atmospheric neutral density and is modulated strongly by ozone absorption (*McPeters*, 1980). Coherent perturbations to the observed Rayleigh scattering signal on scales of tens to hundreds of kilometers are generally indicative of variations in the neutral density and/or ozone near 50–55 km induced by gravity waves (GWs). To quantify these variations, the RAA is defined as the residual difference between the observed Rayleigh scattering albedo and a baseline albedo that would be observed in the absence of any small-scale atmospheric variations, expressed in % (*Randall et al*., 2017). The RAA signal therefore provides a direct measure of dynamical perturbations near the stratopause or in the lowermost mesosphere. In particular, CIPS detects GWs with typical horizontal wavelengths from about 15-600 km and vertical wavelengths longer than 15 km.

The RAA processing stream (Figure 1) begins with the same level 1A data used for PMC data processing (*Lumpe et al.,* 2013) and produces level 2 RAA files for individual scenes (level 2A), full orbit strips (level 2B), global, daily maps that combine all level 2B orbit strips (level 2C), and global, 1-day and 5-day gridded maps (level 3A). Validation of the RAA data is ongoing, but *Randall et al*. (2017) presented several case studies that compared CIPS GW retrievals to near-coincident observations by the Atmospheric Infrared Sounder (AIRS) instrument. These comparisons showed good agreement in the location, orientation and horizontal wavelengths of GW features in both the CIPS and AIRS imagery. The brightness temperature weighting function for the 4.3 μm AIRS data peaks around 35 km, whereas the CIPS weighting function peaks around 50-55 km, enabling investigation of upward propagating GWs.

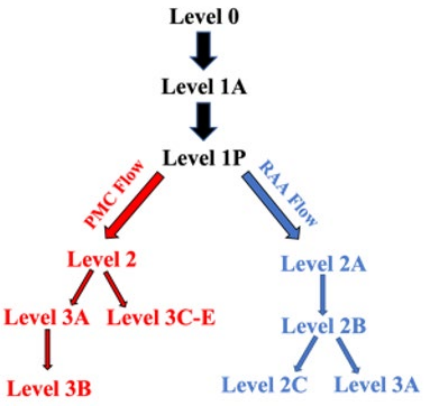

Figure 1. CIPS data processing flow for PMC and RAA processing.

2. CIPS Observing Modes and Data Availability

The RAA data products are derived using all available dayside CIPS images throughout the year.

CIPS instrument observing modes, which dictate the frequency and location of the CIPS images, have changed several times throughout the AIM mission. The CIPS data product overview document (https://lasp.colorado.edu/aim/cips/data/repository/docs/cips_data_overview.pdf) includes a table that summarizes the different observing modes and data availability. Here we summarize the observing modes and effects on CIPS scene locations.

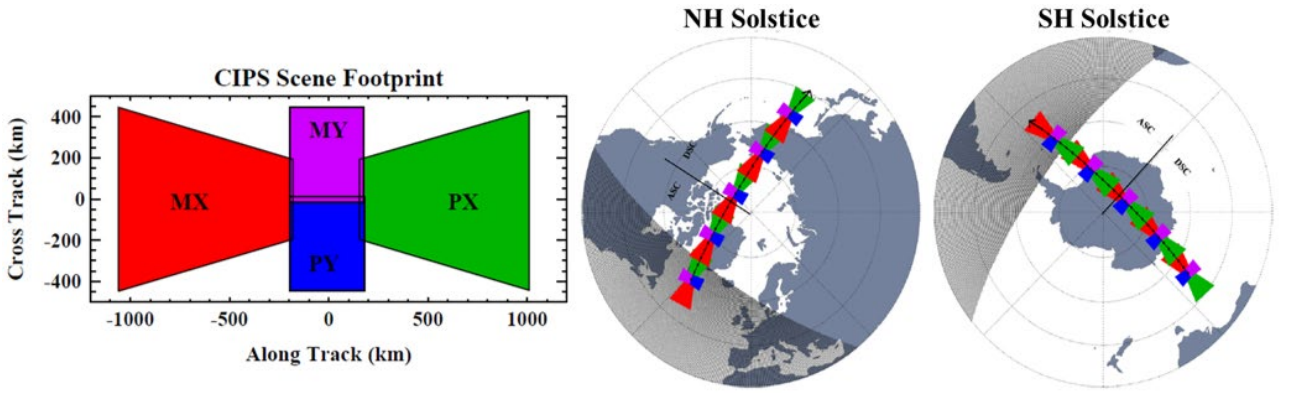

Figure 2. This diagram illustrates the relative alignment of the four CIPS cameras and "summer pole" sampling for an example orbit. The two "X" cameras point in the forward and aft direction, while the "Y" cameras have a primarily nadir orientation. The distance scales indicate the geometric footprint of a CIPS scene at the nominal PMC cloud deck altitude of 83 km. The PX camera always points in the sunward direction, which corresponds to the satellite ram (anti-ram) for Northern (Southern) Hemisphere measurements. The sampling along the orbit pertained to satellite operations from launch until 11 February 2016. The same camera colors are used in later figures. Adapted from *Lumpe et al*. (2013).

From 24 May 2007 through 11 February 2016 CIPS operated in the "summer pole" imaging mode. In this mode, overlapping, four-camera images ("scenes"; see Figure 2) were obtained every 43 seconds at northern hemisphere (NH) latitudes of about 40°N-85°N from March equinox to September equinox, and at southern hemisphere (SH) latitudes of about 40°S-85°S from September equinox to March equinox. At the current time, summer pole imaging data are only available for v1.10r05 (not r06 or r07), and only during the months of March, April, September, and October (outside of PMC seasons). We expect year-round RAA data during summer pole imaging mode to be available as part of v1.10r07, but as of February 2023, v1.10r07 retrievals have not yet been processed for data acquired while operating in the summer pole imaging mode.

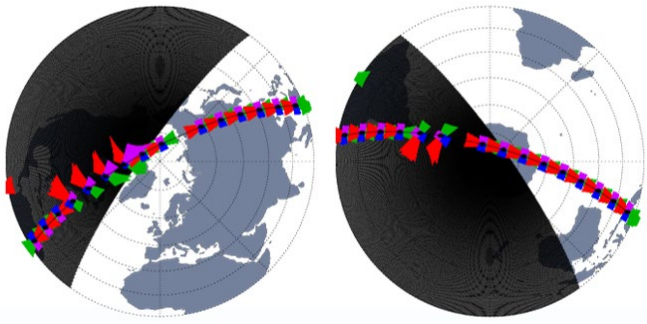

Figure 3. Example of CIPS sampling for full-orbit Continuous Imaging including eclipses. Pole-centered views are shown for the NH (left) and SH (right); black shading indicates the night side. For illustration only; the exact geometry varies from orbit-to-orbit and day-to-day.

Full coverage of the sunlit hemisphere became possible starting in March 2016

when the CIPS instrument began operating in “continuous imaging” (CI) mode. In CI mode from March of 2016 until early November of 2018, images were taken on a 3-minute cadence over the entire orbit. Figure 3 shows an example of this "full-orbit" CI mode sampling, with pole-centered views for the NH (left) and SH (right). Since the images were acquired throughout the orbit, they included the eclipsed (night side) of the earth; however, images taken on the night side are not valid for scientific analysis.

During the full-orbit CI time period in 2016-2018 the beta angle (i.e., the angle between the orbital plane of the satellite and the vector to the sun) changed significantly. This resulted in several prolonged periods of data interruption. From 24 February 2017 until 23 November 2017 the satellite was in “full-sun” conditions, without eclipses. Unfortunately, satellite roll control issues during this period of time prevented the acquisition of reliable calibration data, so no CIPS data are available. In addition, from 7-25 Feb 2018 CIPS was off in preparation for entering the second full-sun period. For this second full-sun period, the sampling during which is illustrated in Figure 4, the satellite remained nadir-pointed, avoiding the roll control problems encountered during the 2017 full-sun period. CIPS data are thus available from 26 February until 20 September 2018. CIPS was turned off to prepare to exit full-sun conditions from 21-28 September 2018, at which point the sampling returned to a geometry similar to that depicted in Figure 3.

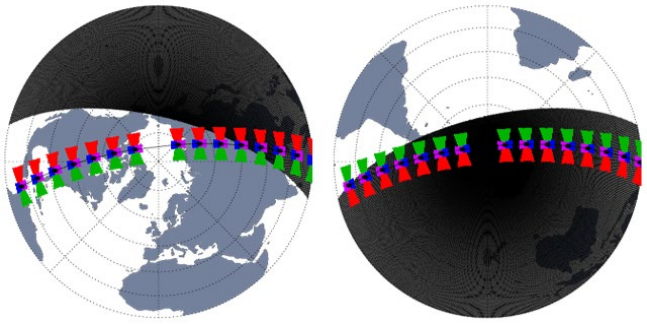

Figure 4. As in Figure 3, but for “full-sun” conditions, when CIPS images were near the terminator.

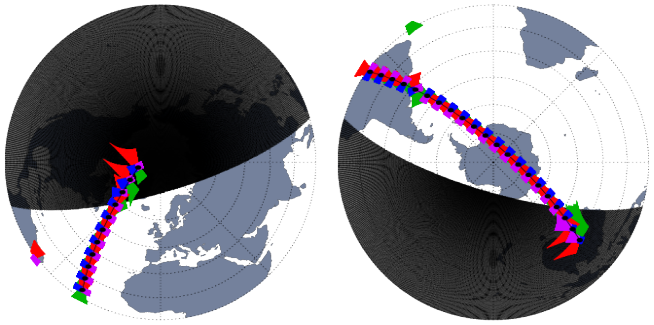

Figure 5. As in Figure 3, but for CI mode confined to the sunlit side of the orbit.

On 3 November 2018 CIPS was commanded to take all of its images over a period of ~70 minutes in order to confine imaging to sunlit latitudes. An example of this sampling is shown in Figure 5. The image acquisition time was modified to 60 minutes in October of 2019 when the eclipse duration lengthened.

3. RAA Level 2A Data Product

Level 2A files contain CIPS RAA data for individual scenes. A CIPS scene contains simultaneous images from the four CIPS cameras (see Figure 2), with a footprint of approximately 2000 km by 900 km (*Lumpe et al.,* 2013). The orientation of the scenes with respect to the orbit track has varied over the course of the mission, as shown in Figures 2-5. Although the level 2A files contain data for individual scenes, the files are provided on an orbit-by-orbit basis. For RAA v01.10r06 and v01.10r07, three level 2A NetCDF data files for each orbit are provided, with each file containing data for all scenes individually in a given orbit (items 1-3 below). In addition, one png

file for each scene is provided separately.

(1) Geolocation: Includes variables such as date, time, latitude, longitude, solar zenith angle, etc. The file name is *cat.nc ("cat" is short for "catalog").

(2) Rayleigh Albedo Anomaly (RAA): Includes Rayleigh Albedo, Rayleigh Albedo Anomaly (RAA), RAA uncertainty, FFT-filtered RAA, FFT-filtered RAA variance, and FFT-filtered RAA variance uncertainty. These files also include the 2D FFT wave amplitudes vs. wavenumber (1/wavelength) and wave orientation. The file name is *alb.nc.

(3) Measurement geometry: Includes satellite view angles and scattering angles for each scene. The file name is *ang.nc.

(4) Plots of RAA, FFT filtered RAA, and RAA variance for each scene (one png file per scene). Regions of significant wave detections are indicated. *As of February 2023 the level 2A plotting code is being updated to correct formatting issues that appear in some seasons in v1.10r06*.

Variables contained in the NetCDF files are described in tables at the end of this document. There are ~15 orbits per day. In summer pole imaging mode, each orbit contains 30 PX camera images and 27 images from each of the other three cameras (i.e., 27 four-camera scenes plus 3 extra PX images). The number of scenes per orbit varies from about 20-30 for CI mode, depending on image cadence and eclipse length. Level 2A (and 2B) RAA and associated geolocation variables are provided with a horizontal resolution of ~56.25 km^2 over the entire orbit track. More precisely, an oblique equirectangular projection is used for the retrieval, with resolution elements that are 7.5 km × 7.5 km in the nadir. With this projection the cross-track dimension remains constant, but the along-track dimension decreases as 7.5 km x cos(ϕ), where ϕ is the along-track angle from nadir. However, the decrease never exceeds 0.4%, so for most purposes an assumption of 7.5 km x 7.5 km is sufficient.

Data arrays in the level 2 files are reported for all level 2 pixels, with array dimensions corresponding to the number of elements in the along-track and cross-track directions. For convenience in data handling, the level 2A arrays span the entire bounding box defined by a single CIPS scene. However roughly half of these elements correspond to locations where no measurements are made and therefore have fill (NaN) values. The compressed level 2A geolocation, albedo and measurement geometry NetCDF files are ~11, 6, and 18 MB in size, respectively. Uncompressed file sizes are much larger due to the significant fraction of fill (NaN) values in these files. IDL software tools to read the Level 2 RAA NetCDF files are available for download from the CU-LASP and NASA SPDF web sites. NetCDF readers for other software packages are available elsewhere (e.g., http://www.unidata.ucar.edu/software/netcdf/software.html).

Users interested in conducting a quantitative analysis of GW wavelengths and amplitudes should use the level 2A data product. Users interested in a more qualitative view of the data, but not in the wave properties themselves, are encouraged to use the level 2B data product. The advantage of using level 2B is that it provides the context of an entire orbit, and waves are often visible in multiple adjacent scenes. The disadvantage is that, for most of the AIM data set, the scenes overlap one another, so wave properties in the overlap regions can be obscured in the level 2B data. The

level 2B data files thus do not contain information about the wave properties.

4. RAA Level 2A Example

As noted in Section 1 above, users should refer to the CMAD for details of the RAA retrieval algorithm. A significant improvement in retrieval version 1.10r06 and later is the addition of RAA variances, which are derived from the level 2A RAA data. This data product is meant to facilitate automated analyses of the RAA data in order to derive statistical information such as might be used for climatological analyses. As an example of the level 2A data, Figure 6 shows plots that users will find in the level 2A png files. The plot of RAA data (Figure 6a) depicts fluctuations in the observed Rayleigh scattering signal that from visual inspection are clearly indicative of GWs. Note that the RAA data included in the level 2A files have been corrected for systematic errors as described in the CMAD.

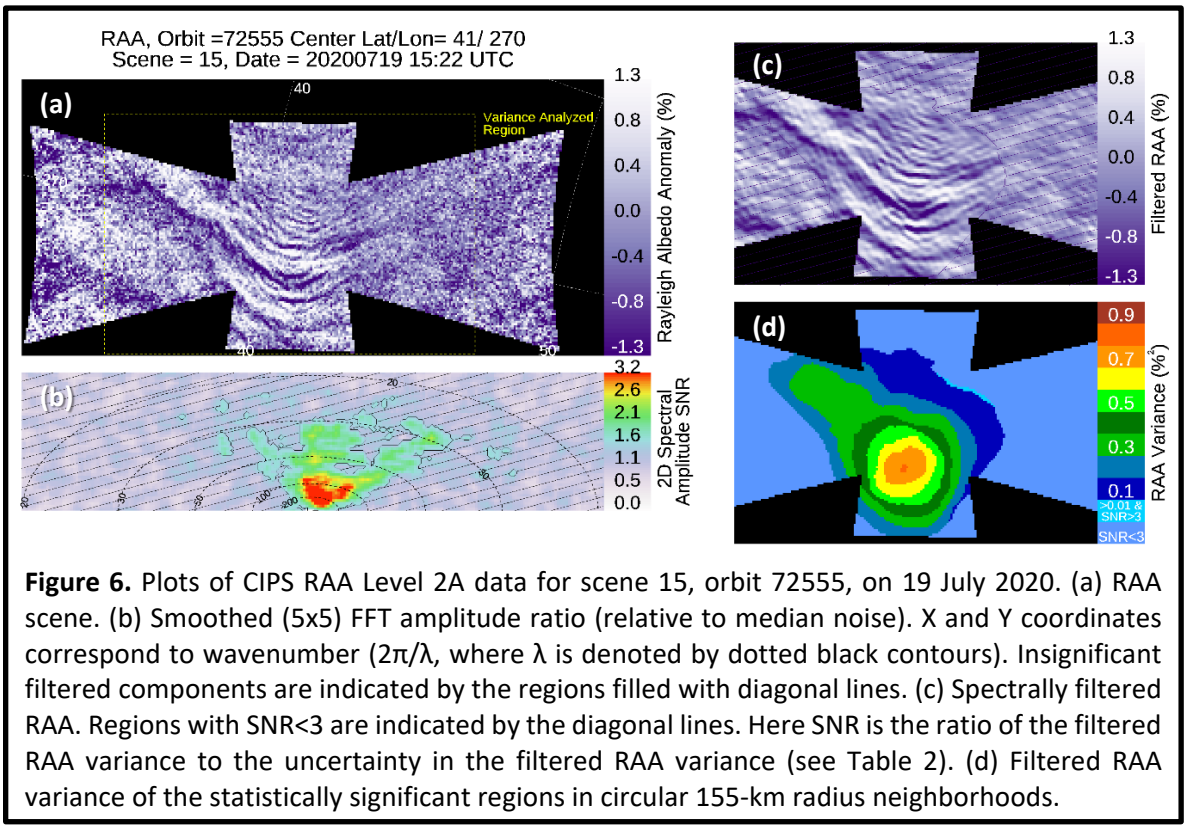

Figure 6. Plots of CIPS RAA Level 2A data for scene 15, orbit 72555, on 19 July 2020. (a) RAA scene. (b) Smoothed (5x5) FFT amplitude ratio (relative to median noise). X and Y coordinates correspond to wavenumber ($2\pi/\lambda$, where λ is denoted by dotted black contours). Insignificant filtered components are indicated by the regions filled with diagonal lines. (c) Spectrally filtered RAA. Regions with SNR<3 are indicated by the diagonal lines. Here SNR is the ratio of the filtered RAA variance to the uncertainty in the filtered RAA variance (see Table 2). (d) Filtered RAA variance of the statistically significant regions in circular 155-km radius neighborhoods.

Quantifying the properties of waves observed by CIPS is accomplished by applying a Fast Fourier Transform (FFT) analysis to convert the RAA image to wavenumber (spectral) space. The FFT yields the amplitudes, horizontal wavelengths, and orientations of the waves identified in the CIPS observations. A ninth-order Butterworth filter is used to remove structure with horizontal wavelength less than 20 km or greater than 400 km. Figure 6b shows the FFT parameters as a plot of the signal-to-noise ratio (SNR) of the spectral amplitude as a function of wavelength, denoted by the dotted black contours. We define the spectral amplitude SNR as the smoothed (5x5 pixels) wave amplitude divided by the median noise, which was determined from a statistical set of wave-free scenes. The spectral amplitude SNR threshold that defines the region of statistical significance (i.e.,

area that is not filled with diagonal lines) is 1.7, which conservatively limits the probability that noise would be detected as a wave to no worse than 1-in-1000. The X and Y coordinates correspond to the wavenumber grids found in the FFT_WAVENUMBER_X and FFT_WAVENUMBER_Y arrays in the level 2A NetCDF *alb files (see Table 2 below). The phase speed direction is given by the arc tangent of the ratio FFT_WAVENUMBER_Y / FFT_WAVENUMBER_X. Thus Figure 6b can be interpreted as a polar coordinate plot where the wavelength is represented by the radial distance and the phase speed direction is represented by the radial angle.

The FFT data are inverted to provide a wave-filtered RAA image in the original spatial domain, with pixels matching the RAA level 2A scenes. The filtered RAA resulting from this transformation is shown in Figure 6c. In this spatial domain, the statistically insignificant data (indicated by the diagonal lines) are those values for which the filtered RAA corresponds to SNR < 3. SNR is defined here as the filtered RAA variance divided by the uncertainty in the filtered RAA variance. Figure 6d shows the filtered RAA variance ($\%^2$) for the statistically significant regions (i.e., the regions without diagonal lines in Figure 6c).

5. Cautions to Users

Users are of course encouraged to consider the RAA uncertainties, which are provided with both the level 2A and 2B data, when analyzing the data. This section describes systematic anomalies or artifacts in the RAA data of which users should be aware, which by their nature are not captured in the uncertainty analysis. Some of these are simply results of the observing geometry, but others are caused by retrieval artifacts or PMC contamination. If images are found to exhibit suspicious behavior that is not described here, please inform the CIPS team by sending an email to aimsds@lasp.colorado.edu.

PX Camera Edge Artifacts in 2007-2008

As noted above, from 2007 until February 2016 CIPS was operating in the summer pole imaging mode. Also, at the time of this writing, only v1.10r05 (not r06 or r07) data are available for summer pole imaging, and only during the months of March, April, September, and October. Much of the RAA v1.10r05 data in 2007-2008 exhibit positive RAA artifacts on the edges of the PX camera. At the time of this writing we believe that the improved calibration in v1.10r07 will avoid such artifacts, but data for summer pole imaging operations

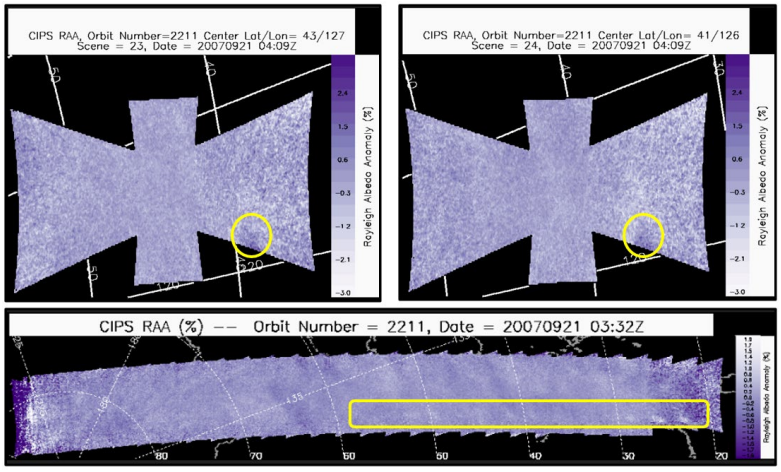

Figure 7. Plots of v1.10r05 Level 2A (top) and 2B (bottom) RAA data for orbit 2211 on 21 September 2007. The level 2A scenes shown are scene 23 (left) and 24 (right). Artifacts mentioned in the text, as they appear in the 2A and 2B data, are highlighted in yellow. Note that the level 2A images are plotted with a color scale that goes from light at the low end to dark at the high end, the opposite of the level 2B color scale.

have not yet been retrieved with v1.10r07. Figure 7 shows an example of this for orbit 2211, on 21 September 2007. Early indications are that these signals are small enough that they will not lead to significant values of the wave variance, but at the time of this writing, caution is warranted when interpreting such data. More generally, if repeating patterns occur in multiple adjacent scenes, caution is warranted.

Apparent Camera-to-Camera Discontinuities

Users will notice that occasionally one camera in a scene shows, on average, systematically higher or lower RAA values than the other cameras in that scene. Two examples are shown in Figure 8. The top left four panels show the level 2A plots for scene 3 of orbit 72288 on 1 July 2020. The top right four panels show the level 2A plots for scene 13 of the same orbit. For context, the level 2B orbit strips of RAA and RAA variance are shown below the level 2A data. Note that scene 1 in this orbit was taken as CIPS was emerging from satellite eclipse, at SZAs that were too high for the CIPS retrievals. Thus the level 2B orbit strip begins on the left with scene 2, and even this scene is partially truncated because of the high SZA values. Scene 13 was acquired just before the 180-degree satellite yaw, which occurs twice per orbit in order to keep the PX camera, which is equipped with a sun shade, solar-pointed.

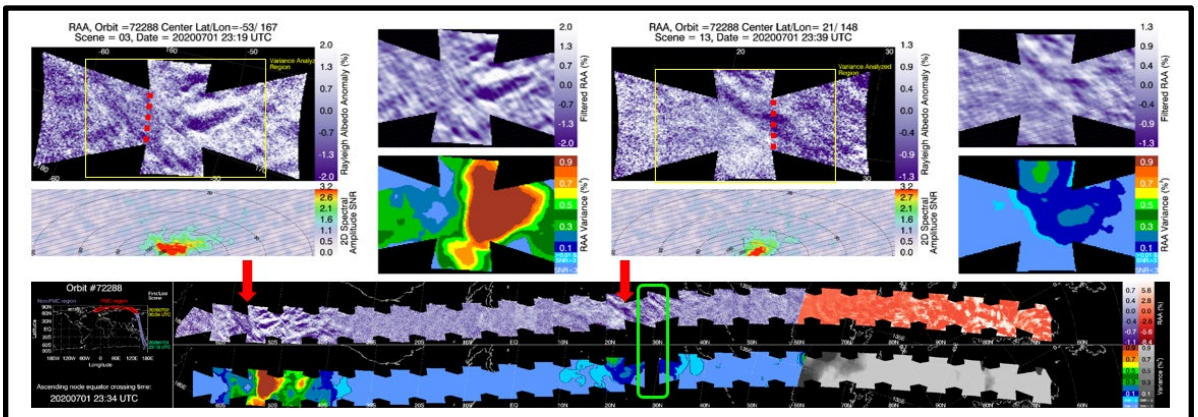

Figure 8. Examples of apparent camera-to-camera discontinuities, for orbit 72288 on 1 July 2020. Level 2A plots for scene 3 (four panels at top left) and scene 13 (four panels at top right) are shown along with the level 2B orbit strip (bottom). Scene 3 was taken shortly after exiting eclipse, at high SZA, and scene 13 was taken shortly before a 180-degree satellite yaw. Dashed red lines mark the intersection between the nadir cameras and MX (left-most) camera in scene 3, and between the nadir cameras and PX (right-most) camera in scene 13. In scene 3 the MX camera has systematically lower RAA values, on average, than the PX or nadir cameras. In scene 13 the PX (right-most) camera has systematically lower RAA values, on average, than the MX or nadir cameras. Note that because the Butterworth filter screens out any waves with $\lambda > 400$ km, these apparent discontinuities are no longer evident in the FFT-filtered scenes and variances. Red arrows in the level 2B orbit strip denote the same locations as the red dashed lines on the level 2A plots, and are just meant to guide the eye. The green box is centered on the location of the scene that is missing due to the 180-degree satellite yaw. The gap in variances inside the green box, which is filled in the RAA plot itself, is caused by the fact that the variance calculation includes only part of each scene, as indicated by the yellow box in the top left panels of the level 2A plots. The red/orange (gray) shading in the orbit strips denotes the nominal PMC region poleward of 60°, in which signals from PMCs can obscure gravity wave signals.

Figure 8 shows that the MX (left-hand, southernmost) camera on scene 3 is uniformly darker purple (lower RAA) than the rest of the scene. It is possible that this indicates the presence of large-scale atmospheric waves. However, it is also possible that the apparent large-scale structure is an artifact of the Rayleigh subtraction at high SZA due to an error that is not yet understood. In any case, the Butterworth filter applied in the variance calculation screens out any waves with $\lambda > 400$ km. Therefore, the camera-to-camera differences shown in the RAA itself are no longer evident in the reconstructed FFT-filtered RAA (top right of the 4-panel plots) or variances (bottom right of the 4-panel plots; the high variances here are due to the smaller-scale waves). More generally, however, users should exercise caution when interpreting RAA and RAA variances for high SZA conditions, since the Rayleigh scattering signal is so low at high SZA.

Figure 8 also shows that the PX (right-hand) camera in scene 13 is uniformly darker purple (lower RAA) than the rest of the scene. As for scene 3, it is possible that this is caused by the presence of large-scale atmospheric waves. However, this behavior is often seen in scenes immediately before or after a yaw, which suggests that this is a retrieval artifact. Briefly, since images are smeared if the satellite is yawing during image acquisition, yawed images are omitted from the retrievals; this explains the missing scene in the level 2B plot of orbit 72288 near 30°N latitude, highlighted by the green box. Because of this omission, the region viewed by the PX camera immediately adjacent to a yaw is viewed only once – there are no overlapping views by the PX camera from overlapping scenes. When a region of the earth is viewed more than once by the same camera (i.e., on successive scenes), the RAA retrieval averages the results. Thus regions that are viewed only once by the PX camera can appear to have systematically higher or lower RAA than regions for which multiple views are averaged. Again, because the variance calculation filters waves with $\lambda > 400$ km, the camera-to-camera differences are not evident in the reconstructed FFT-filtered RAA or variances. Users should exercise caution, however, when interpreting the unfiltered RAA values for scenes adjacent to a yaw.

Finally, significant camera-to-camera discontinuities can very occasionally occur in the nadir cameras (not shown). This most likely results from errors in stitching together RAA data from the four cameras into a single scene, but the specific cause has not yet been determined. Since this happens so rarely (a few scenes per year), the CIPS team removes these scenes from the data base.

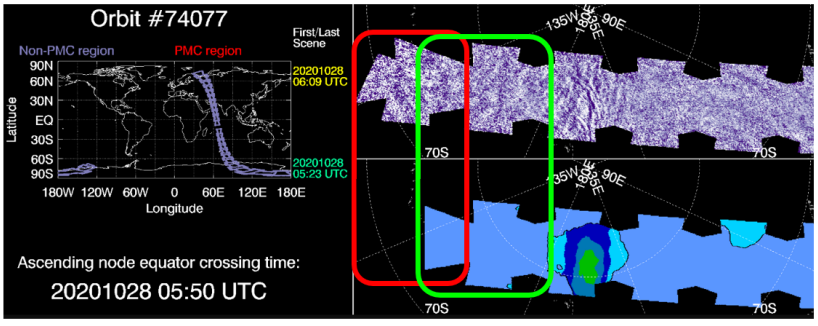

Figure 9. Partial plot from the Level 2B png file for orbit 74077 on 28 October 2020. Scenes 1 and 2 are highlighted in red and green, respectively. Because the high-SZA pixels in scene 1 have been screened, leaving only a partially complete scene, the variance analysis was not applied.

Partial Scenes

The v1.10r07 variance calculation applies a threshold for the minimum number of pixels in a scene for the variance analysis to be run (see

the yellow boxes in Figure 8). Scenes that do not exceed this threshold, which tend to be partial edge scenes that are truncated due to high SZA, are not included in the variance calculation. Figure 9 shows such an example, for orbit 74077 on 28 October 2020. The numerical RAA data for scene 1 in this orbit are available in the level 2A netcdf data files, but because the scene does not include enough data for the variance analysis, the variance arrays are set to NaN. In addition, there is no accompanying png file for scene 1. Likewise, the level 2A plots provided on the CIPS website for this orbit begin with scene 2.

Large-Scale Waves

As noted several times above, the RAA variance calculation filters out structures in the RAA field with wavelengths longer than 400 km. These structures might therefore appear in a visual inspection of the RAA data, but will not be evident in the RAA variance results. Thus, at least at the current time, statistical analyses that utilize the RAA variances cannot be used to investigate the occurrence of such large-scale waves.

Signatures of PMCs in the RAA Data

PMCs generally occur from mid-May through August in the north and mid-November through mid-February in the south. PMCs modify the amount of 265-nm radiation scattered by the atmosphere, so signatures of PMCs are apparent in the RAA data. These signatures can be seen, for example, in Figure 8 in the red/orange shaded region of the level 2B RAA orbit strip plot. Plots in the RAA png files use red/orange shading at summer hemisphere latitudes poleward of 60° during the PMC seasons to emphasize the need for caution to avoid misinterpreting PMC signatures in the RAA data. The default setting for the png file plots is to shade the PMC region red/orange from 15 May to 31 August in the NH and from 15 November to 28 February in the SH. For some (but not all) seasons that were reprocessed after the end of the PMC season, the first date of PMC region shading will correspond to the date the first cloud was detected.

Users should also be aware that, because of PMC contamination, the principal component analysis (PCA) used to correct for systematic errors in the RAA retrieval ignores latitudes poleward of 60° in the summer hemisphere between the dates of the first and last observations of PMCs each season. This often results in artifacts at high latitudes in the RAA data, an example of which is shown in

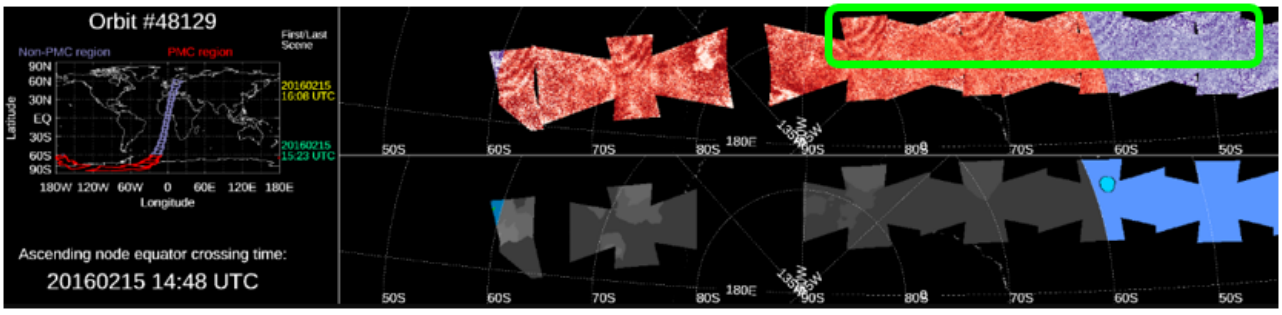

Figure 10. Level 2B RAA data at high southern latitudes for orbit 48129 on 15 February 2016. Repeating artifacts in the red/orange shaded PMC region (poleward of 60°S) extend to mid-latitudes (purple region), as indicated by the green highlight. The artifacts arise because the principal component analysis ignores the polar region during the PMC season.

Figure 10. Occasionally, these artifacts extend to latitudes equatorward of 60°. The variances associated with the artifacts outside the PMC region are small, but users should beware of camera-to-camera repeating patterns that extend from the PMC region into the mid-latitudes.

Although PMCs most often occur poleward of 60° latitude in each hemisphere, mesospheric clouds have been observed at latitudes as low as ~35°. Therefore, signatures of PMCs can sometimes be found in the RAA data equatorward of 60°. Figure 11 shows an example of mid-latitude PMC detections in the RAA data. PMC signatures in the RAA data generally appear to be much more randomly structured than GW signatures, and they are also often easy to identify because they appear as extensions of patterns that are apparent in the red/orange shaded region. In addition, PMCs would typically only appear equatorward of 60° in the CIPS data near the peak of the PMC season, in late June and July in the NH and in late December and January in the SH.

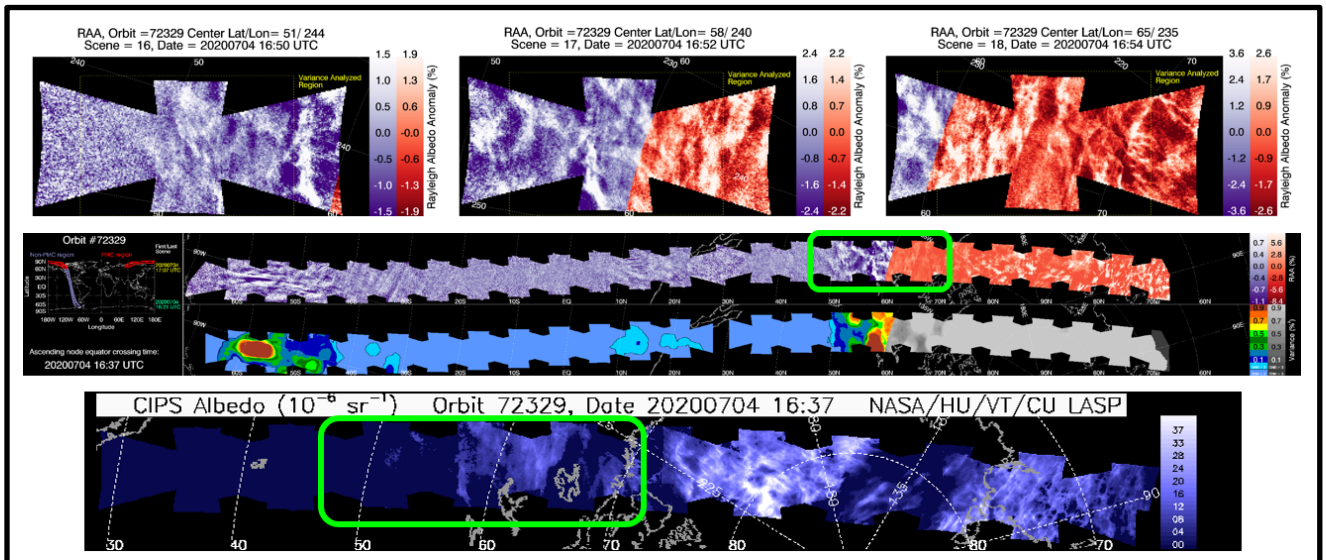

Figure 11. *Top*: Examples of PMC signatures in the level 2A RAA data from 4 July 2020, orbit 72329, scenes 16-18. Red/orange shading indicates the nominal PMC region poleward of 60° as a caution to users, but mesospheric clouds can occur equatorward of this region during peak PMC occurrence time periods. *Middle*: Corresponding level 2B RAA data; the green box highlights the region of scenes 16-18. *Bottom*: PMC albedo for the same orbit. The RAA data show evidence of many more clouds from 50-60°N, because no screen is applied to avoid false cloud detections. Although the RAA data cannot be used for quantitative analyses of PMCs, it can be used to identify dim PMCs that the operational PMC algorithm misses.

Using RAA Data for PMC Detection

The RAA retrieval is registered to an altitude near 53 km, so it cannot be used for quantitative analyses of PMCs. However, unlike the PMC retrievals, the RAA retrievals have no filtering to avoid false detections of PMCs. The PMC retrieval filtering effectively avoids false detections, but in so doing fails to identify some of the dimmest PMCs. Therefore, more PMC signatures are often evident in the RAA data than in the PMC data. This is why the PMC albedo shown in Figure 11 does not show the lowest latitude clouds that are evident in the RAA data in Figure 11 for the same orbit. Users interested in detecting the dimmest mesospheric clouds can thus take advantage of the RAA data, even though these clouds will not be identified as such in the operational PMC retrievals.

PMC vs. RAA Orbit Numbering

Users interested in relating PMC and RAA data should be aware of a change in the AIM orbit-numbering scheme for RAA data beginning on orbit 59351 on 27 February 2018. With the orbit numbering used for PMC retrievals and for RAA v1.10r05, orbits are defined to begin at the ascending node equator crossing. This follows the official NASA numbering convention. However, beginning in late February 2018 with the second full-sun observing period, this caused orbits to be "split" by a nightside traverse. That is, the ascending node equator crossing switched from the night side to the day side. This meant that the northernmost portion of the orbit was acquired before entering eclipse, and the southernmost portion after exiting eclipse. The result was that scenes from the same orbit were split into two sections that were separated by a significant gap in longitude.

The orbit split just described is not an issue for PMC retrievals, since these retrievals include only mid and high latitudes in the summer hemisphere. However, the split is inconvenient for users of the RAA data. Therefore, to avoid the split orbits the orbit numbering logic in RAA v1.10r06 and later combines the last half of "old" orbit n-1 (i.e., the orbit portion that extended from eclipse exit in the SH until just before the orbit n ascending node equator crossing) with the first half of "old" orbit n (i.e., the orbit portion that extended from the old orbit n ascending node equator crossing until eclipse entry in the NH). So beginning with orbit 59351 on 27 February 2018, the NH data in RAA v1.10r06 and later have the same orbit number as in the PMC data and RAA v1.10r05, but the SH data for orbit n in RAA v1.10r06 and later correspond to orbit n-1 in the PMC and RAA v1.10r05 data. Note that the official start time of orbit n is still the ascending node equator crossing time – since this occurs in the middle of the sunlit side of the orbit, some scenes in orbit n are acquired prior to the official start time.

Regarding PMC analyses, a major consequence of renumbering the orbits is that the PMC data for all SH seasons beginning with SH1819 will have different orbit numbers than the RAA data. For instance, PMC data in orbit 80596 on 1 Jan 2022 will correspond to SH RAA data in orbit 80597 on 2 Jan 2022.

Image Projection in Png Files

On rare occasions that usually correspond to orbits with missing data, the orientation of the orbit bounding box is skewed. In these situations the level 2A plotting code can produce results where the plots overfill the plot area. The numerical data in the NetCDF files is unaffected.

6. NetCDF Files

The NetCDF files listed in Section 3 enable users to quantitatively analyze the level 2A RAA data. The tables below summarize the contents of each data file and provide a brief description of all parameters and arrays. This section also includes additional information regarding the file contents.

File Name Convention: Each file has a day-of-year included in the file name. This is the UT day corresponding to the ascending node equator crossing time, which is the official orbit start time. When the equator crossing time is shortly prior to midnight UT, some or all of the data in the file

might have been acquired on the day after the day in the filename. Also, as noted above, when the ascending node equator crossing time is on the sunlit side of the orbit, some of the data in the file might have been acquired prior to this time. In addition, file names can include a "preliminary" tag, which indicates that the data have not yet been processed with both the final satellite ephemeris and final instrument calibration data. Preliminary RAA data products are typically released to the public within ~2 days of acquisition. Definitive data with final satellite ephemeris and instrument calibration data are released in 6-month increments, between the March and September equinoxes. At the time of this writing, data acquired during the period from the 2022 September equinox to the 2023 March equinox are still preliminary.

SZA Corrections: Because of systematic biases at high SZA that are not captured completely in the initial retrieval, SZA-dependent corrections are applied to the RAA data for all pixels that correspond to SZA > 75°. These corrections are given in the variables HIGH_SZA_SYSTEMATIC_CORRECTION_NORTH and HIGH_SZA_SYSTEMATIC_CORRECTION_SOUTH, at the SZA values given in the variable HIGH_SZA_SYSTEMATIC_CORRECTION_GRID. The corrections are based on the PCA and are given in %, and are simply values that are added to the initially retrieved RAA values. Most users will not need these corrections. However, if suspicious SZA-dependent anomalies become apparent in an analysis of the RAA data itself, these variables can be subtracted from the RAA to test whether removing the corrections has any significant effect on the apparent anomalies. Note that the exact correction for specific SZAs should be determined by interpolation from the high SZA systematic correction grid.

Orbit Track Grid: The 3-element arrays ORBIT_TRACK_X_AXIS, ORBIT_TRACK_Y_AXIS, and ORBIT_TRACK_Z_AXIS specify the axes on which the orbit track grid is defined. They are the same for the levels 2A and 2B files. These arrays can be used to recreate the CIPS Earth-Centered, Earth-Fixed (ECEF) level 2 grid, for instance to bin or sample another data set onto the CIPS grid without interpolating in latitude/longitude. This might be particularly useful when doing orbit and observation geometry calculations, such as when calculating coincidences with a limb viewing instrument. The level 2 grid is parameterized by the along-track angle (λ, behaves like longitude) and cross-track angle (ϕ, behaves like latitude). Analogous to longitude and latitude, the ECEF unit vector [x(*i,j*), y(*i,j*), z(*i,j*)] for each grid point (*i,j*) can be computed as follows:

$$[x(i,j), y(i,j), z(i,j)] = \cos[\phi(i,j)]*\cos[\lambda(i,j)]*\text{ORBIT\_TRACK\_X\_AXIS} + \cos[\phi(i,j)]*\sin[\lambda(i,j)]*\text{ORBIT\_TRACK\_Y\_AXIS} + \sin[\phi(i,j)]*\text{ORBIT\_TRACK\_Z\_AXIS}$$

The variables here are:

$\lambda = x\_grid\_index(i,j)*pixel\_length/R_e$ = along track angle in radians

$\phi = y\_grid\_index(i,j)*pixel\_length/R_e$ = cross track angle in radians

pixel_length = 7.5 km

R_e = Earth Radius = 6378.137 km

x_grid_index = BBOX(0,*k*) + *i*

y_grid_index = BBOX(1,*k*) + *j*

The 4-element bounding box array (BBOX) is given in the NetCDF file; *k* is the scene index for level 2A, which is not present in level 2B.

Level 2A Example: Orbit 74077, 28 October 2020, Grid Point *i* = 130, *j* = 25, *k* = 3 (i.e., Element [130,25,3] from the LATITUDE and LONGITUDE arrays):
ORBIT_TRACK_X_AXIS = [0.34745139 , 0.88977518 , -0.29593522]
ORBIT_TRACK_Y_AXIS = [0.24064462 , 0.22041730 , 0.94525467]
ORBIT_TRACK_Z_AXIS = [0.90629339 , -0.39964526 , -0.13753531]
BBOX(0,3) = -1129 BBOX(1,3) = -73
x_grid_index = -1129 + 130 = -999
y_grid_index = -73 + 25 = -48
λ = -999*7.5/6378.137 = -1.17472 radians
ϕ= -48*7.5/6378.137 = -0.0564428 radians
cos[ϕ]*cos[λ] = 0.385191, cos[ϕ]*sin[λ] = -0.921111, sin[ϕ] = -0.0564128
[x , y , z] = [-0.13895197 , 0.16224945 , -0.97691732]

PCA Parameters: The PCA is run on a one- to two-week cadence to identify artificial, systematic patterns or gradients in the RAA that can be similar in magnitude to readily identifiable, geophysical wave signatures. Structures in the RAA images that are consistent with linear combinations of the PCA components are removed during the retrieval process. The RAA level 2A and 2B data files include the variable PCA_ORBIT_RANGE, which gives the first and last orbits included in each PCA. Unless otherwise noted in this documentation, the CIPS team is not aware of significant systematic errors that have not been caught by the PCA. However, if questionable patterns in the RAA data appear to correlate with the timing of a PCA, users should proceed with caution and are encouraged to contact the AIM CIPS team.

Tables of variables and references follow below.

Table 1. Variables in the CIPS RAA Level 2A geolocation (*cat.nc) file. Fill value is NaN.

Variable Name	Units	Type/Dimension	Description
AIM_ORBIT_NUMBER	N/A	Long / 1	NASA orbit number[a]
BBOX	pixels	Long / [4,Nscenes]	Bounding Box: Bottom-left indices (x0, y0) and size of box (xsize, ysize) within the larger orbit track grid [x0, y0, xsize, ysize]
CENTER_LON	degrees	Double / 1	Center longitude of orbit at maximum orbit track latitude (~82.5°) in the spring/summer hemisphere; divides ascending and descending node.
DATA_PRODUCT	N/A	String[b] / 1	Brief description of the data product
GAP_ADJACENT_FLAG	N/A	Byte / [Nscenes]	Data quality warning flag indicating the scene was adjacent to a data gap (such as a dark); 0=not adjacent, 1=adjacent
GPS_MICROS_TIME	micro-seconds	Double / [Nscenes]	GPS microsecond time for each scene
HIGH_SZA_SYSTEMATIC_CORRECTION_NORTH	%	Float / [201]	PCA-based, SZA-dependent correction added to northern hemisphere RAA with SZA>75°
HIGH_SZA_SYSTEMATIC_CORRECTION_SOUTH	%	Float / [201]	PCA-based, SZA-dependent correction added to southern hemisphere RAA with SZA>75°
HIGH_SZA_SYSTEMATIC_CORRECTION_SZA_GRID	degrees	Float / [201]	SZA grid for high SZA corrections; interpolate off this grid to recover corrections at other SZAs
JD_TIME	days	Double / [Nscenes]	JD time for each scene (days from noon UT on January 1, 4713 BC)
KM_PER_PIXEL	km	Float / 1	Linear dimension of square pixel occupying area of CIPS resolution element
LATITUDE	degrees	Float / [Xdim,Ydim,Nscenes]	Latitude of each element
LONGITUDE	degrees	Float / [Xdim,Ydim,Nscenes]	Longitude of each element; ranges from -180 to 180
NOTES	N/A	String[b] / 1	Any additional notes
NSCENES	N/A	Long / 1	Number of scenes in the orbit
ORBIT_END_TIME	micro-seconds	Double / 1	GPS end time of AIM_ORBIT_NUMBER (microseconds from 0000 UT on 6 Jan 1980)
ORBIT_START_TIME	micro-seconds	Double / 1	GPS start time of AIM_ORBIT_NUMBER (microseconds from 0000 UT on 6 Jan 1980)
ORBIT_START_TIME_UT	N/A	String[b] / 1	String for UT start time of orbit
ORBIT_TRACK_EPOCH	micro-seconds	Double / 1	GPS time at which the orbit track axis approximately matches the true orbit track[c]
ORBIT_TRACK_X_AXIS	N/A	Double / [3]	X axis on which the orbit track grid is defined, for levels 2A and 2B, in Earth-centered, Earth-fixed (ECEF) coordinates
ORBIT_TRACK_Y_AXIS	N/A	Double / [3]	Y axis on which the orbit track grid is defined, for levels 2A and 2B, in ECEF coordinates
ORBIT_TRACK_Z_AXIS	N/A	Double / [3]	Z axis on which the orbit track grid is defined, for levels 2A and 2B, in ECEF coordinates
PCA_ORBIT_RANGE	N/A	Long / [2]	Orbit range over which the PCA was run

PITCH_ANGLE	degrees	Float / [Nscenes]	Pitch of the spacecraft with respect to the velocity, cross-track, and zenith vectors in a right-hand coordinate system
PRODUCT_CREATION_TIME	N/A	String[(b)] / 1	String containing UT time at which data file was produced
REVISION	N/A	String[(b)] / 1	Data revision number
ROLL_ANGLE	degrees	Float / [Nscenes]	Roll of the spacecraft with respect to the velocity, cross-track, and zenith vectors in a right-hand coordinate system
UT_DATE	yyyymmdd	Long / [Nscenes]	UT date of each scene in the orbit
UT_DATE_ORBIT_START	yyyymmdd	String[(b)] / 1	UT date of AIM_ORBIT_NUMBER start time (ascending node equator crossing time); CIPS file names use this date
UT_TIME	hour of day	Float / [Nscenes]	UT time for each scene
VERSION	N/A	String[(b)] / 1	Data version number
XDIM	N/A	Long / 1	Number of along-orbit-track elements (pixels) in the RAA and RAA Variance data arrays
YAW_ADJACENT_FLAG	N/A	Byte / [Nscenes]	Data quality warning flag indicating the scene was adjacent to a yaw; 1 means adjacent
YAW_ANGLE	degrees	Float / [Nscenes]	Yaw of the spacecraft with respect to the velocity, cross-track, and zenith vectors in a right-hand coordinate system
YDIM	N/A	Long / 1	Number of cross-orbit-track elements (pixels) in the RAA and RAA Variance data arrays
ZENITH_ANGLE	degrees	Double / [Xdim,Ydim,Nscenes]	Solar zenith angle (SZA) of each element referenced to 50 km altitude

[(a)] After Sept 2017, orbits are spliced with second half of AIM_ORBIT_NUMBER-1 to produce continuous dayside data, in which case this is the orbit number for the northern half of the data.

[(b)] Depending on the read software, these variables might appear as Byte rather than string. If so, they should be converted to string.

[(c)] Due to the rotation of the earth the orbit track doesn't trace out a plane in earth-fixed coordinates.

Table 2. Variables in the CIPS RAA Level 2A Rayleigh Albedo Anomaly (*alb.nc) file. Fill value is NaN.

Variable Name	Units	Type/Dimension	Description
FFT_WAVENUMBER_X	1/km	Float / [XdimFFT]	Wavenumber grid for x-direction of FFT
FFT_WAVENUMBER_Y	1/km	Float / [YdimFFT]	Wavenumber grid for y-direction of FFT
FILTERED_RAA	%	Float / [Xdim,Ydim,Nscenes]	Spatial domain inversion of the spectrally filtered RAA
FILTERED_RAA_VARIANCE	$\%^2$	Float / [Xdim,Ydim,Nscenes]	Variance of spectrally filtered RAA
FILTERED_RAA_VARIANCE_UNC	$\%^2$	Float / [Xdim,Ydim,Nscenes]	Uncertainty in variance of spectrally filtered RAA (i.e., variance in the noise allowed through by the spectral filter)
OVERLAP_ERROR	%	Double / [5,Nscenes]	Residual median difference between the overlap regions after the overlap offset is applied (PX-PY, PX-MY, PY-MY, MX-PY, MX-MY)

OVERLAP_ERROR0	%	Double / [5,Nscenes]	Median difference between the overlap regions before applying overlap offset (PX-PY, PX-MY, PY-MY, MX-PY, MX-MY
OVERLAP_OFFSET	%	Double / [4,Nscenes]	Normalization applied to each image in a scene to minimize disagreement between overlapping camera fields of view. This number is added to the original RAA for each image (PX, MX, PY, MY)
RAA_FFT_AMPLITUDE	%	Float / [XdimFFT, YdimFFT,Nscenes]	2D-FFT amplitude (upper half plane)
RAA_FFT_MEDIAN_NOISE_AMPLITUDE	%	Float / [XdimFFT, YdimFFT,Nscenes]	Median noise amplitude in the spectral domain (noise in RAA_FFT_AMPLITUDE)
RAA_FFT_NOISE_RATIO	N/A	Float / [36, Nscenes]	FFT Signal-to-Noise Ratio (SNR) thresholds defined as the ratio RAA_FFT_AMPLITUDE / RAA_FFT_MEDIAN_NOISE_AMPLITUDE; elements of this array correspond to the percentile grid given by RAA_FFT_PERCENTILE_GRID
RAA_FFT_PERCENTILE_GRID	N/A	Double / [36]	Percentiles corresponding to the FFT SNR thresholds in RAA_FFT_NOISE_RATIO[(b)]
RAA_FFT_PHASE[(a)]	radians	Float / [XdimFFT, YdimFFT,Nscenes]	2D-FFT phase (upper half plane)
RAYLEIGH_ALBEDO	10^{-6} sr^{-1}	Float / [Xdim,Ydim,Nscenes]	Total observed albedo
RAYLEIGH_ALBEDO_ANOMALY	%	Float / [Xdim,Ydim,Nscenes]	100*(observed_albedo-baseline)/baseline baseline = basic_state_albedo
RAYLEIGH_ALBEDO_ANOMALY_UNC	%	Float / [Xdim,Ydim,Nscenes]	Uncertainty in RAA due to noise
XDIM	N/A	Long / 1	Number of along-orbit-track elements (pixels) in the RAA and RAA Variance data arrays
XDIMFFT	N/A	Long / 1	Number of elements in the x-direction of the FFT grid
YDIM	N/A	Long / 1	Number of cross-orbit-track elements (pixels) in the RAA and RAA Variance data arrays
YDIMFFT	N/A	Long / 1	Number of elements in the y-direction of the FFT grid = 1 + Ydim/2

[(a)] When initially provided, the RAA_FFT_PHASE arrays in the v1.10r06 Level 2A RAA data files were mistakenly filled with NaN values. Users who downloaded the RAA 2A data prior to March 2023 should re-download the data if they are interested in the phase information.

[(b)] FFT_PERCENTILE_GRID ranges from 0.5 to 0.99999392, which in theory would correspond to a range in probability that noise will be detected as a wave from 1-in-2 (element 0) to 1-in-164,000 (element 35). In practice, because of limited data in the statistical calculations, this probability is most reliable for the first 25 elements, from 1-in-2 to 1-in-1000. The probability that noise will be identified as a wave decreases monotonically from 1-in-1000 for elements 26-35, but not as steeply as the percentile grid elements would suggest.

Table 3. Variables in the CIPS RAA Level 2A measurement geometry (*ang.nc) file. Fill value is NaN.

Variable Name	Units	Type/Dimension	Description
ANGLE_ALONG_ORBIT_TRACK	degrees	Double / [Nscenes]	Angle of satellite along the orbit track measured from the midpoint (minimum SZA) on the day side
PIXELS_BINNED_FROM_1A	N/A	Double /	Number of Level 1A pixels binned into each grid

		[Xdim,Ydim,Nscenes]	point
SCATTERING_ANGLE	degrees	Double / [Xdim,Ydim,Nscenes]	Solar scattering angle
VIEW_ANGLE	degrees	Double / [Xdim,Ydim,Nscenes]	Angle between the line of sight and the earth surface normal at the binning height (50 km)
VIEW_ANGLE_DERIVATIVE	degrees / km	Double / [Xdim,Ydim,Nscenes]	Derivative of the view angle with altitude along the line of sight
ZENITH_ANGLE_DERIVATIVE	degrees / km	Double / [Xdim,Ydim,Nscenes]	Derivative of the solar zenith angle with altitude along the line of sight

7. References

Bailey, S. M., G. E. Thomas, D. W. Rusch, A. W. Merkel, C. D. Jeppesen, J. N. Carstens, C. E. Randall, W. W. McClintock, and J. M. Russell III (2009), Phase functions of polar mesospheric cloud ice as observed by the CIPS instrument on the AIM satellite, J. Atmos. Sol. Terr. Phys., 71, 373–380, doi:10.1016/j.jastp.2008.09.039.

Lumpe, J. D., S.M. Bailey, J.N. Carstens, C.E. Randall, D. Rusch, G.E. Thomas, K. Nielsen, C. Jeppesen, W.E. McClintock, A.W. Merkel, L. Riesberg, B. Templeman, G. Baumgarten, and J.M. Russell, III (2013), Retrieval of polar mesospheric cloud properties from CIPS: algorithm description, error analysis and cloud detection sensitivity, *J. Atmos. Solar-Terr. Phys.*, http://dx.doi.org/10.1016/j.jastp.2013.06.007.

McPeters, R. D. (1980), The behavior of ozone near the stratopause from two years of BUV observation, J. Geophys. Res., 85, 4545–4550, doi:10.1029/JC085iC08p04545.

Randall, C. E., J. Carstens, J. A. France, V. L. Harvey, L. Hoffmann, S. M. Bailey, M. J. Alexander, J. D. Lumpe, J. Yue, B. Thurairajah, D. E. Siskind, Y. Zhao, M. J. Taylor, and J. M. Russell, III (2017), New AIM/CIPS global observations of gravity waves near 50-55 km, *Geophys. Res. Lett*., 44, 7044–7052, http://dx.doi.org/10.1002/2017GL073943.

Created by Jerry Lumpe and Cora Randall, June 2017.