AVIRIS ground data-processing system

John H. Reimer, Jan R. Heyada, Steve C. Carpenter, William T. S. Deich, and Meemong Lee

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, California 91109

ABSTRACT

The Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) has been under development at the Jet Propulsion Laboratory (JPL) for the past four years. During this same time period, a dedicated ground data-processing system has been designed and implemented to archive and process the large amounts of data expected from this instrument. This paper reviews the objectives of this ground data-processing system and presents the hardware implementation. An outline of the data flow through the system is given, and the software and incorporated algorithms developed specifically for the systematic processing of AVIRIS data are described.

1. INTRODUCTION

The Airborne Visible/Infrared Imaging Spectrometer (AVIRIS), the second in a series of imaging spectrometer instruments developed at JPL for Earth remote sensing,[1,2] has recently become operational. This instrument, flown aboard a U-2 aircraft, uses a scanner and four line arrays of detectors to image a 614-pixel swath simultaneously in 224 contiguous spectral bands (0.4 to 2.4 μm) with a ground instantaneous field of view (GIFOV) of 20 m. Ten-bit data are recorded onto 14-track high-density digital tape (HDDT) at a rate of 17 Mbit/sec, which is equivalent to 12 image lines per second. Data are written to the HDDT as a series of major frames. Each major frame contains one full-spectrum image line with associated navigation, engineering, and dark current data, as well as imbedded sync words, frame IDs, and frame counts. Record time is limited to 43 minutes by the HDDT's capacity of 5.5 Gbyte.

The need to process and archive the large volume of data generated by the AVIRIS instrument required the development of a dedicated ground data-processing system. This system, based on a VAX 11/780 minicomputer with the VMS operating system and running under the TAE/VICAR2 image processing executive, was developed and is currently operated within JPL's Image Processing Laboratory.

2. OBJECTIVES

The basic objectives of the AVIRIS ground data-processing system are to provide the following capabilities:

(1) Decommutate and archive AVIRIS data.

(2) Provide retrieval processing of archived AVIRIS data as requested by science investigators and apply appropriate radiometric and geometric rectification.

These two main processing steps, archival processing and retrieval processing, have further objectives that are more specific.

2.1 Archival processing

Archival processing is performed on all data recorded by the AVIRIS instrument. Specific archival processing requirements are as follows:

(1) Decommutate the HDDT data with playback on an Ampex HBR-3000 HDDT recorder and transfer the data directly to the host computer system.

(2) Segment the data stream into flight lines.

(3) Separate engineering, platform, calibration, and image data.

(4) Provide a scrolling video display of the image data from any one selected spectral band during the archival process for data quality assessment.

(5) Provide printouts and/or plots of engineering data to assess the health of the instrument.

(6) Generate an automated data base entry for record-keeping and data retrieval purposes.

(7) Generate black-and-white 8" x 10" photographic prints of the image data from four spectral bands (one band from each spectrometer). One print set with negatives is to be maintained as part of the archival records, while another set of prints is to be sent to the investigator(s) who requested specific flight lines.

(8) Generate archival CCT labels.

(9) Block the data as necessary to achieve an archival storage capacity of 512 full-spectrum image lines and associated auxiliary data per 6250-bpi 9-track CCT.

(10) Require no more than two weeks to archive one AVIRIS HDDT.

(11) Maintain printed records of all archived AVIRIS data.

2.2 Retrieval processing

Retrieval processing occurs when a request for data is received from a science investigator. Specific retrieval processing requirements are as follows:

(1) Provide the ability to locate data within the data base by geographic site name, latitude/longitude, or flight ID and run ID.

(2) Possess the capability to subset data in the band and/or line dimension.

(3) Conduct radiometric rectification to provide full-spectrum reconstruction accommodating spectral overlap between spectrometers, with dark current subtraction and resampling to correct for spectral readout time delays.

(4) Conduct geometric rectification to perform geometric warping to achieve correct pixel sizes and correct the relative spatial position of pixels.

(5) Block retrieved data and write them to 6250-bpi tape(s) for delivery to the investigator.

(6) Accomplish complete retrieval processing of one 10.5 km by 17.5 km full-spectrum scene (224 bands by 1024 lines by 614 samples) within two working days.

3. HARDWARE DESCRIPTION

A diagram of the AVIRIS ground data-processing system is shown in Figure 1. The system is based on a Digital Equipment Corporation (DEC) VAX 11/780 with 10 Mbytes of memory. The VAX is configured with three UNIBUS channels and one MASSBUS channel. Standard DEC peripherals include a console terminal, a line printer/plotter, two TU78 9-track 6250/1600-bpi tape drives, and six RA81 456-Mbyte disk drives. Additionally, a DEC LA100 printer is configured to print tape labels.

An Ampex HBR-3000 high-density tape drive is used to play back data recorded by the AVIRIS instrument onto HDDT. The HBR-3000 is interfaced to the VAX via a JPL-designed and -built controller interface, which in turn is connected to a DEC DRE11 alternate-buffered direct memory access (DMA) interface on one of the UNIBUS channels. The interface passes tape transport commands from the VAX to the HBR-3000 and passes data from the HBR-3000 to the VAX. Data leave the HBR-3000 as a serial bit stream. The JPL controller interface searches for and detects a specific 40-bit sync pattern imbedded in the data stream at the start of each major frame. It places each 10-bit data value following the sync pattern into the 10 least significant bits of a 16-bit word and passes the word to the DRE11, which then does a DMA transfer to the VAX. The DRE11 alternately accesses two separate 64K word areas within main memory. This allows one area to be processed while the other area is being filled, resulting in a virtually continuous data transfer.

The data system also has two Rastertech one/25 512 x 512 display devices and a Matrix QCR hardcopy film recorder with both Polaroid and 35mm capability. An OSI Laserdrive 1200 optical disk drive with Perceptics software has recently been installed, and an Interlan Ethernet controller connects the AVIRIS data system to several other systems on the Imaging Spectrometer Local Area Network (ISLAN) located at JPL.

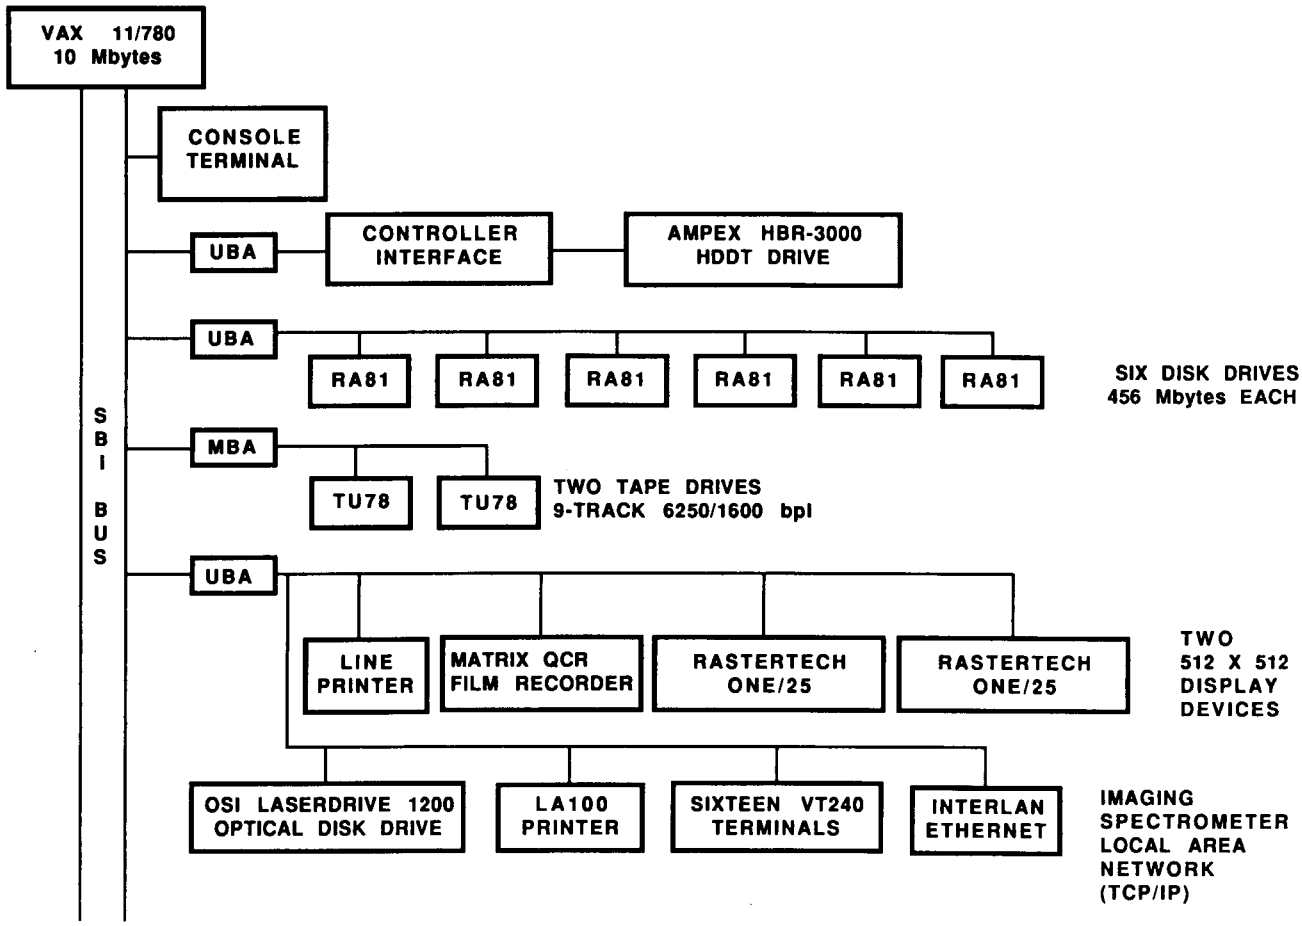

Figure 1. AVIRIS ground data-processing system.

4. OPERATION

Figure 2 shows the data flow through the AVIRIS ground data-processing system as it has been implemented. In this diagram, two subsidiary tasks, creation of an auxiliary input file and creation of the calibration file, have been shown in addition to the main archival and retrieval processing tasks.

4.1 Creation of the auxiliary file

The auxiliary file is a text file that provides the archival processing software with information not contained on the HDDT but necessary for a complete data base entry. The file is generated from the AVIRIS flight log and contains the following information for each flight line:

(1) Site name
(2) State (country for foreign flights)
(3) Principal investigator's name
(4) Principal investigator's affiliation
(5) Comments (cloud cover, visibility, etc.)

One auxiliary file is created using the editor for each HDDT and is used as a secondary input to the archival processing task.

4.2 Archival processing

Archival processing may be divided into three stages: decommutation, data reformatting, and data output. These three stages form a cycle that is repeated as data are processed in segments of 512 major frames at a time. This is the maximum amount of data stored on one archival CCT. Processing may be halted and restarted following the completion of a cycle. The time required to archive one segment is approximately 1.2

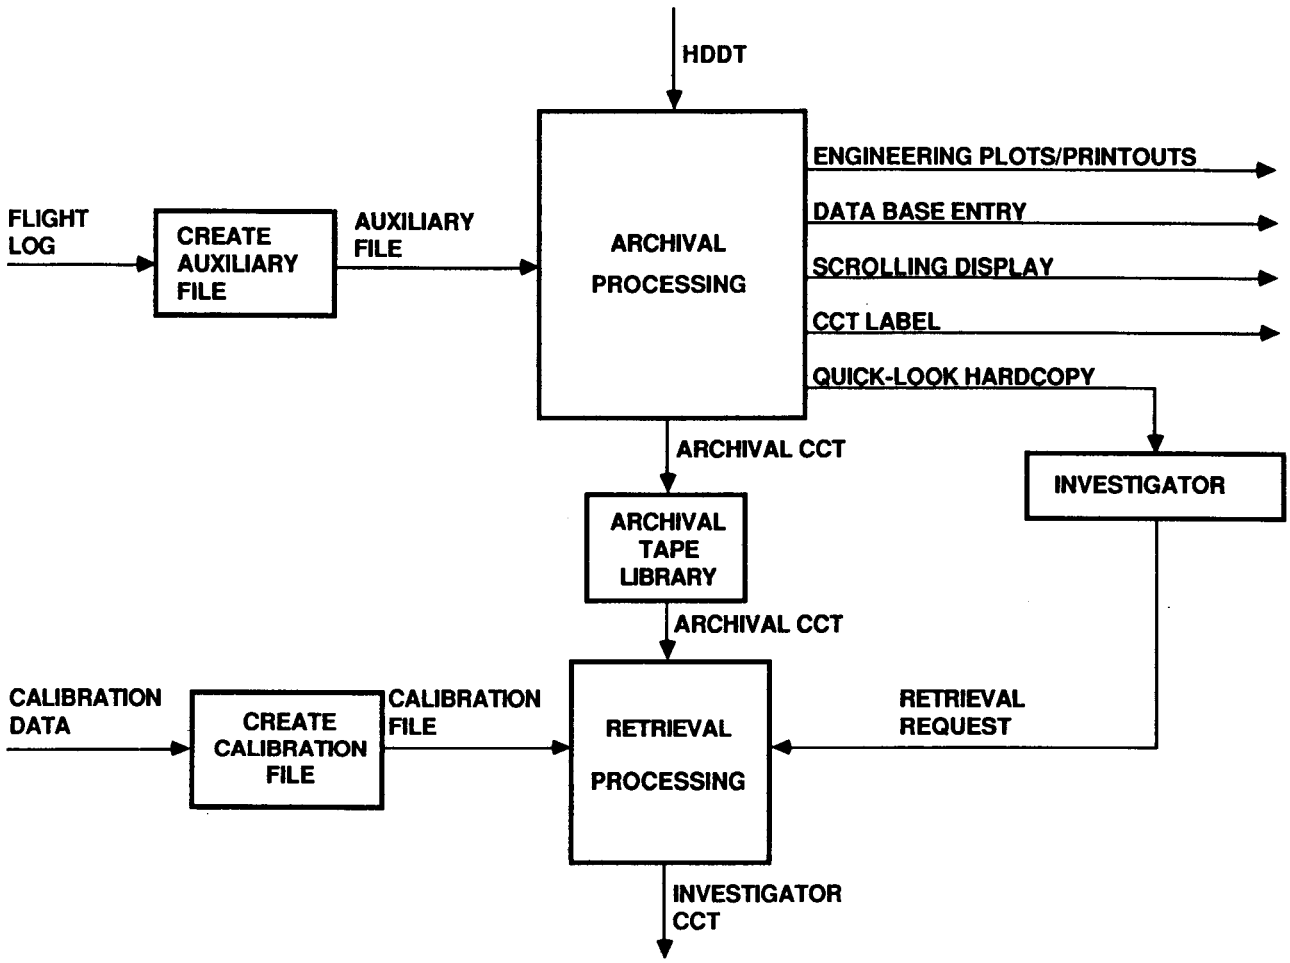

Figure 2. Data processing flow.

hours. Since a full HDDT contains 61 segments, the time required to archive a full HDDT is approximately 9 working days. Due to the real-time nature of the decommutation process, no other processing can be performed on the system during archival processing.

4.2.1 Decommutation. During decommutation, data are read into the VAX from the HBR-3000 and written directly onto magnetic disk, one major frame at a time. During this process the format of the major frame is verified and the major frame count is checked. Any corrupted or missing major frames are flagged with values of 4096. Playback of the HDDT on the HBR-3000 is done at 1/16 of the record speed, corresponding to a data rate of 1.06 Mbit/sec out of the HBR-3000 and, after the 10-bit to 16-bit conversion, 1.7 Mbit/sec into the VAX. Although the HBR-3000/VAX interface was run successfully at twice this speed, disk access/write times did not allow operation at this higher speed. Following decommutation of 512 major frames, the HBR-3000 tape transport is stopped and the HDDT is rewound slightly to position it for the next segment.

4.2.2 Data reformatting. Following decommutation, the data are read back from disk and reformatted. The data are separated into seven possible types: Pre-Cal, Navigation, Engineering, Dark Current, Offset, Image, and Post-Cal. Pre-Cal and Post-Cal data sets contain onboard calibration data taken immediately prior to and immediately following every flight line. Navigation and Engineering data are decoded from a bit-encoded format to INTEGER*4 and REAL*4 values. Image data are converted to a band interleaved by line organization and the pixels within each line are transposed to account for the scan mirror scanning from left to right. The center 512 samples of one spectral band are displayed on one of the Rastertech display devices in a scrolling manner as each line is processed. Histograms are calculated for each of the four spectral bands that are to be displayed in the quick-look hardcopy. Initial plans to simply use the eight high-order bits of each 10-bit pixel in producing the quick-look hardcopy did not result in images of acceptable quality.

4.2.3 Data output. At this point, samples of the data to be archived are output. The engineering data are printed and/or plotted. A linear auto-stretch with 1% saturation is performed separately on each of the four quick-look bands, converting them to byte data. These data are then sent to the QCR hardcopy device, exposing a 35mm negative. An example of a quick-look print is shown in Figure 3. A data base entry and a CCT label are generated using information from the data stream and the auxiliary file. Using this same information, VICAR annotation labels are generated which will be attached to the archived data sets. Finally, the data are written to the archival CCT as separate, blocked, and VICAR-labeled files.

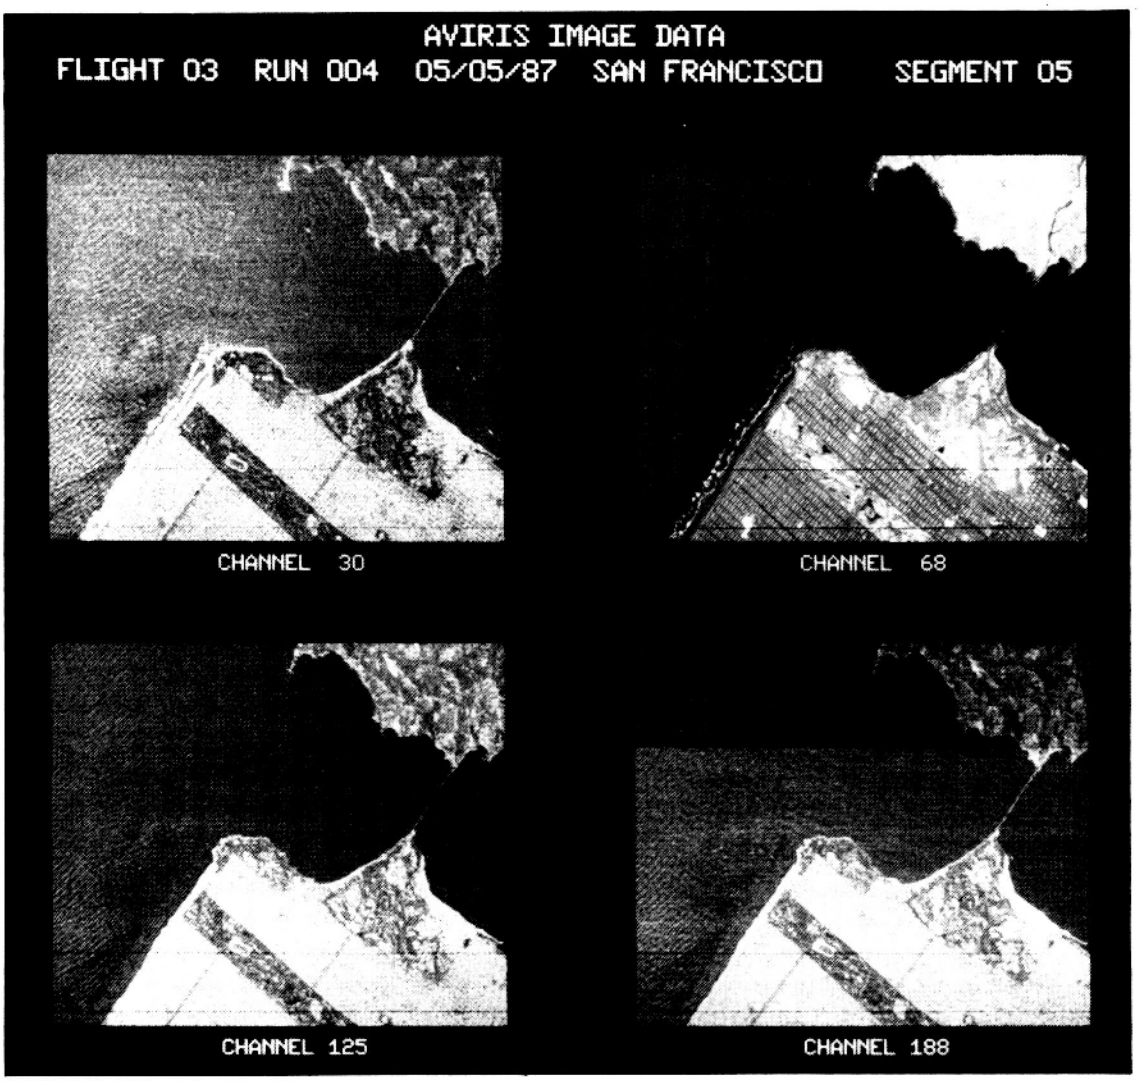

Figure 3. AVIRIS image data.

4.3 Creation of the calibration file

The calibration file, used in the radiometric rectification of AVIRIS data, is a set of REAL*4 multipliers used to convert AVIRIS DN values into radiance values in terms of $\mu W/cm^2$-nm-sr. The file is generated from laboratory calibration data taken prior to a flight season. These data consist of two components: radiance measurements of an integrating sphere light source taken by a spectroradiometer and data taken by the AVIRIS instrument while viewing the same integrating sphere. The AVIRIS laboratory calibration data are decommutated and archived in the same manner as flight data.

The calibration file is generated by first subtracting the dark current from the AVIRIS instrument data and then averaging these data over time as a function of spectral band (detector element) and cross-track sample number. A cross-track dependence is kept to account for nonuniform response with scan angle due to vignetting in the foreoptics of the

AVIRIS instrument. This results in a set of 224 x 614 mean DN values (224 spectral bands and 614 cross-track pixels). The 614 mean DN values of every spectral band are then associated with the corresponding radiance value for that spectral band, computed by resampling the spectroradiometer measurements to match the spectral spacing of the AVIRIS detector elements. Assuming detector linearity, the final set of multipliers is then computed as:

$$MULT(s,b) = RAD(b)/DN(s,b) \quad (1)$$

where s = the cross-track sample number and b = the spectral band number. This set of 224 x 614 multipliers is saved as the calibration file and used in the radiometric rectification of AVIRIS data taken during the flight season.

4.4 Retrieval processing

Retrieval processing is initiated by an investigator's request for AVIRIS data. The data are located, processed to the level desired by the investigator, and written to 6250-bpi CCT for delivery to the investigator. The time required to perform full retrieval processing on a full-spectrum 10.5 km by 17.5 km scene is approximately 2 working days. Since a full HDDT contains 30 scenes, full retrieval processing of an entire HDDT would require approximately 60 working days.

When a request for AVIRIS data is received, the first step in retrieval processing is to locate the desired data in the AVIRIS data base. This is a file-structured data base set up under DATATRIEVE. The AVIRIS data base contains information such as flight ID, run ID, flight date and time, site name, state or country of the site, site latitude and longitude, principal investigator's name and affiliation, name and location of the tape the data are stored on, and any comments noted during the recording of the data concerning instrument configuration and/or data quality. Software has been written which allows interactive interrogation of the data base. By specifying items such as flight ID, run ID, date, time, site, latitude/longitude, and/or investigator's name, information on any AVIRIS data an investigator might be interested in may be located, sorted, and printed.

After they are located, the desired data are subset to investigator specifications in the line and band dimensions, with the number of lines restricted to a maximum of 1024. The data are copied onto disk and, in the event that they span more than one archival CCT, the multiple data segments are concatenated together. Radiometric rectification and geometric rectification are then performed as requested and the data written to CCT for delivery to the appropriate investigator. Full retrieval processing of a 1024-line full-spectrum scene requires approximately 15 cpu hours.

4.4.1 Radiometric rectification. The radiometric rectification of AVIRIS data converts 10-bit DN values to REAL*4 radiance values in terms of $\mu W/cm^2$-nm-sr and linearly scales the values to 16-bit words. In the process, dark current subtraction, detector equalization, resampling to account for detector readout delays, and reconstruction of the full spectrum accounting for band overlap and varying spectral spacings between spectrometers are performed.

The first step in radiometric rectification processing is dark current subtraction. One dark current value for each detector element is recorded every major frame. Due to the presence of noise, dark current subtraction is done using a sliding mean of 101 dark current values as follows:

$$DN(s,b,l) = DN(s,b,l) - 1/101 \sum_{k=l-50}^{k=l+50} DC(b,k) \quad (2)$$

where s = the cross-track sample number, b = the spectral band number, l = the line number, and DC = the dark current.

The next step is detector equalization. Using the multipliers stored in the calibration file, DN values are converted to REAL*4 radiance values:

$$RAD(s,b,l) = DN(s,b,l)*MULT(s,b) \quad (3)$$

Resampling to correct for detector readout delays is then performed. These delays occur because the elements in the linear arrays are read sequentially. This results in the last element of an array (element #64) seeing a spot on the ground approximately one pixel further along the scan than was seen by the first element. Linear interpolation is performed between successive spatial pixels within a scan line as shown in Equation 4. As indicated by this formula and due to scan direction, time increases as the sample number decreases. The time between reading element #64 of an array and then reading element #1 of the same array (the next spatial pixel) is equivalent to reading two detector elements; thus the number 66 is used in Equation 4. Readout of the 32-element silicon array corresponds in time to readout of elements 1 through 32 within the three 64-element InSb arrays.

$$RAD'(s,b,l) = [(b'-1)/66]*RAD(s+1,b,l) + [(67-b')/66]*RAD(s,b,l) \quad (4)$$

where b' = the array element number of band b.

Following this, the full spectrum is reconstructed. Resampling is performed in the spectral direction to correct for band overlap between spectrometers and to create a spectrum with uniform spectral spacing. This step produces 210 bands from 0.4000 μm to 2.4482 μm with 9.8-nm spectral sampling. The center wavelengths of the AVIRIS detector elements are known from data taken in the spectral alignment procedure done in the laboratory. Ignoring detector element #1 from each linear array (unusable because of the readout architecture), each output band is computed by linearly interpolating between the two closest surrounding input bands:

$$RAD''(s,b,l) = \left(\frac{\lambda_b - \lambda_{b-}}{\lambda_{b+} - \lambda_{b-}} \right) [RAD'(s,b+,l) - RAD'(s,b-,l)] + RAD'(s,b-,l) \quad (5)$$

where b = the output band number, b- = the input band nearest to b with wavelength less than b, and b+ = the input band nearest to b with wavelength greater than b.

Finally, the REAL*4 radiance values are converted to 16-bit words by multiplying each radiance value by a factor of 100 and rounding to the nearest integer value. This is necessary for further processing by the geometric rectification and spectral analysis software and for the reduction of data storage requirements by a factor of two.

4.4.2 Geometric rectification. The geometric rectification of AVIRIS data involves geometric warping to correct for oversampling and for distortions caused by aircraft attitude variations. AVIRIS data is oversampled in both the cross-track and along-track directions in that the center-to-center spacing of 20-m pixels is 17.1 m. The pixels of AVIRIS images do not correspond to a rectangular array of imaged points on the ground because the aircraft is traveling forward during the time required to image successive pixels within a scan line and the aircraft is undergoing possible rotations (pitch, roll, and yaw) and possible variations in altitude and/or ground speed during image aquisition.

The first step in the geometric rectification process is to use a lookpoint model to map the input image onto a "reference plane," a horizontal coordinate plane located at sea level. One would like to map the image onto coordinates of a map of the surface; however, given some amount of terrain relief, some error in pixel location will occur. This is shown in Figure 4, where a pixel's line of sight is projected onto a sea-level plane and placed at point X, while the true pixel location on a map of the surface would place the position at X'. The coordinate system of the reference plane is arranged so that at some initial time t_o, $\hat{x}$ is to starboard, $\hat{y}$ is in the direction of aircraft travel, and $\hat{z}$ is oriented upwards; the aircraft's position is (0,0,h). (See Figure 5.)

To map a given pixel into a position on the reference plane, we assume that the following is known:

(1) t_p = time of pixel exposure
(2) $v(t)$ = horizontal velocity, known from t_o to t_p
(3) h = aircraft altitude at t_p
(4) ψ = aircraft yaw at t_p
(5) θ = aircraft pitch at t_p
(5) ϕ = aircraft roll at t_p
(7) η = angle of observation from aircraft nadir

We wish to compute the position (x,y) where a line from the aircraft at an angle η and time t_p intercepts the reference plane. The method used is to compute rotation matrices for yaw, pitch, and roll; apply the matrices to a unit vector at an angle η from the nadir of the aircraft; compute the intercept of the unit vector with the reference plane; and finally, translate the unit vector's intercept point by the amount the aircraft moved between t_o and t_p.

The aircraft's "body" coordinate system is defined as a right-handed orthonormal system, O_{RPY}, in which the origin, O, is at the aircraft center of gravity; the roll axis, $\hat{R}$, points in the forward direction; the pitch axis, $\hat{P}$, points to starboard; and the yaw axis, $\hat{Y}$, points down. (See Figure 6.) We first convert from coordinates (R,P,Y) in the O_{RPY} system to a coordinate system $O_{R_oP_oY_o}$, in which $\hat{R}_o$ is parallel to the $\hat{y}$ axis of the reference plane, $\hat{P}_o$ is parallel to the $\hat{x}$ axis of the reference plane, and the $\hat{Y}_o$ axis points in the $-\hat{z}$ direction. Yaw, pitch, and roll are defined so that one may convert from O_{RPY} to $O_{R_oP_oY_o}$ using the standard order of the angular rotations,[3] namely:

(1) Roll through $-\phi$ about the R axis
(2) Pitch through $-\theta$ about the new P axis
(3) Yaw through $-\psi$ about the Y_o axis

The rotation matrices to take O_{RPY} into $O_{R_oP_oY_o}$ are:

$$\lambda_\phi = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & \sin\phi \\ 0 & -\sin\phi & \cos\phi \end{bmatrix}, \quad \text{(roll)} \tag{6a}$$

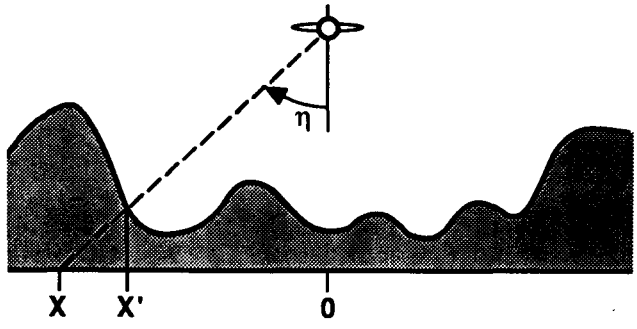

Figure 4. Pixel location error.

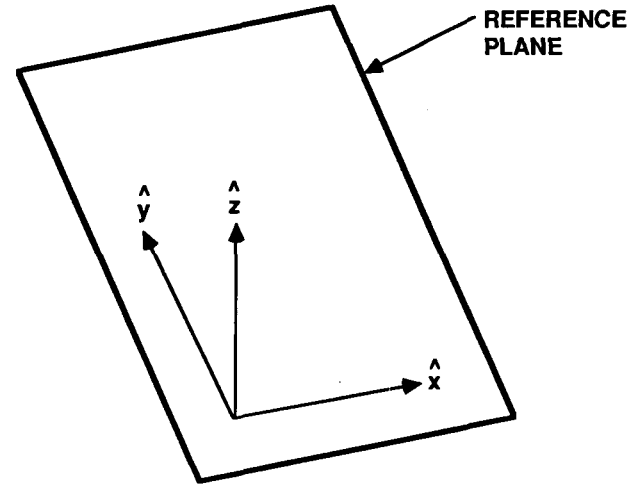

Figure 5. Reference plane coordinate system.

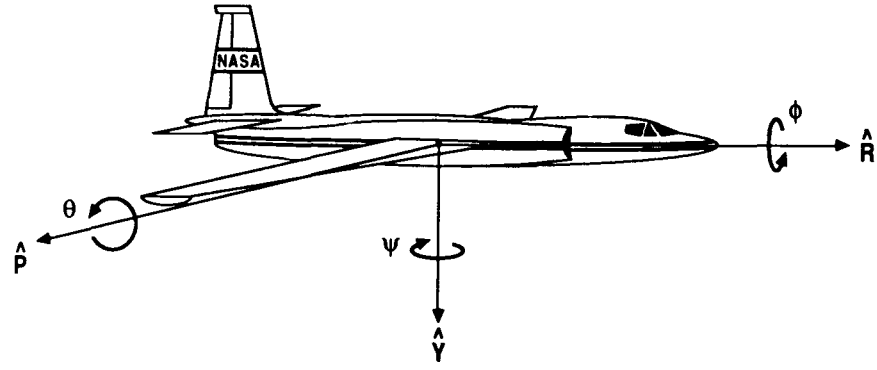

Figure 6. Aircraft coordinate system.

$$\lambda_\theta = \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix}, \quad \text{(pitch)} \tag{6b}$$

$$\lambda_\psi = \begin{bmatrix} \cos\psi & \sin\psi & 0 \\ -\sin\psi & \cos\psi & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \text{(yaw)} \tag{6c}$$

and the net rotation matrix to convert a vector in the O_{RPY} system into one in the $O_{R_oP_oY_o}$ system is:

$$\lambda_{b_o} = \lambda_\psi \lambda_\theta \lambda_\phi \tag{7}$$

In turn, a unit vector in the $O_{R_oP_oY_o}$ coordinate system may be expressed as a unit vector in the reference plane by first rotating 180 deg around the roll axis and then rotating 90 deg around the new yaw axis:

$$\lambda_r = \begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix} \tag{8}$$

The total rotation matrix is thus:

$$\lambda_T = \lambda_r \lambda_{b_o} \tag{9a}$$

$$= \lambda_r \lambda_\psi \lambda_\theta \lambda_\phi \tag{9b}$$

$$= \begin{bmatrix} -\sin\psi\cos\theta & (\cos\phi\cos\psi + \sin\theta\sin\phi\sin\psi) & (\sin\phi\cos\psi - \sin\theta\cos\phi\sin\psi) \\ \cos\psi\cos\theta & (\sin\psi\cos\phi - \sin\theta\sin\phi\cos\psi) & (\sin\theta\cos\phi\cos\psi + \sin\phi\sin\psi) \\ \sin\theta & \sin\phi\cos\theta & -\cos\theta\cos\phi \end{bmatrix} \tag{9c}$$

Now let d' be a unit vector in the body frame that points at a pixel viewed at an angle η from the yaw axis $\hat{Y}$ in the plane perpendicular to the roll axis $\hat{R}$:

$$d' = \begin{bmatrix} 0 \\ -\sin\eta \\ \cos\eta \end{bmatrix} \tag{10}$$

In the coordinate system of the reference plane, this unit vector has the coordinates:

$$d = \lambda_T\, d' \tag{11a}$$

$$= \begin{bmatrix} -(\cos\theta\cos\psi + \sin\theta\sin\phi\sin\psi)\sin\eta + (\sin\phi\cos\psi - \sin\theta\cos\phi\sin\psi)\cos\eta \\ -(\sin\psi\cos\phi - \sin\theta\sin\phi\cos\psi)\sin\eta + (\sin\theta\cos\phi\cos\psi + \sin\phi\sin\psi)\cos\eta \\ -\sin\phi\cos\theta\sin\eta - \cos\theta\cos\phi\cos\eta \end{bmatrix} \tag{11b}$$

The intersection of the unit vector d with the reference plane can be conveniently computed when d is expressed in cylindrical coordinates: Define α to be the azimuth of the observation (measured CCW from the $\hat{x}$ axis of the reference plane), and let ρ be the angle of d away from the $-\hat{z}$ axis; $\rho = 0$ means that d is normal to the reference plane, while $\rho = \pi/2$ means that d is parallel to the reference plane. We have:

$$\rho = \pi - \cos^{-1} d_z \tag{12a}$$

$$\alpha = \tan^{-1}(d_y/d_x) \tag{12b}$$

The aircraft position at time t_p is:

$$x_a(t_p) = \int_{t_o}^{t_p} v_x(t)dv_x \tag{13a}$$

$$y_a(t_p) = \int_{t_o}^{t_p} v_y(t)dv_y \tag{13b}$$

The intercept of the line of sight from the aircraft with the reference plane is then obtained from Equations 12 and 13:

$$x = x_a + h \tan\rho \cos\alpha \tag{14a}$$

$$y = y_a + h \tan\rho \sin\alpha \tag{14b}$$

Using this procedure and the navigation data from the U-2 aircraft, an array of control points is generated mapping the image onto the reference plane. These control points are spaced approximately every 100 lines and 100 samples in the input image.

The next step is to resample these control points to obtain a regularly spaced array of points in the reference plane. This is required by the geometric warping algorithm to be used. The algorithm employed in resampling the control points is called "triangularization".[4] It connects all the points into an optimal surface formulated by searching for the shortest line segment that does not cross existing line segments. The search is applied exhaustively (i.e., for all points) by computing the vector distances from a given point to all points and sorting the distance list. Once the connections are configured, each triangle mapping (geometric transformation) is computed as:

$$\begin{bmatrix} x_1 & x_2 & x_3 \\ y_1 & y_2 & y_3 \end{bmatrix} = A \begin{bmatrix} x'_1 & x'_2 & x'_3 \\ y'_1 & y'_2 & y'_3 \end{bmatrix} + C \tag{15}$$

where

$$A = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

$$C = \begin{bmatrix} c_1 & c_1 & c_1 \\ c_2 & c_2 & c_2 \end{bmatrix}$$

After the transformation coefficients are generated for every triangle, the program creates a regular grid. For each grid point, it determines which triangle the grid point is in and then applies the appropriate triangle transformation to compute the corresponding input grid point:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = A_i \begin{bmatrix} x \\ y \end{bmatrix} + C_i \qquad (16)$$

where i denotes the triangle within which grid point (x,y) is located.

Finally, the actual geometric warping of the image data is perfomed using the resampled control points. The software that performs this function is written to operate on data with three dimensions: lines, samples, and bands. It assumes that all bands are co-registered. The geometric warping process consists of pixel location mapping using the input control points (shown by Figure 7 and Equations 17a and 17b) and pixel intensity mapping (bilinear interpolation shown in Equation 18). Since the location mapping process is a function of the line and sample only, the location map can be shared over the third dimension, the band dimension. This sharing of the location map across the third dimension enhances the computation speed greatly.

$$s' = a1 + (s-x1)*DS1 + (l-y1)*DS2 + (l-y1)*(s-x1)*DSL \qquad (17a)$$

$$l' = b1 + (s-x1)*DL1 + (l-y1)*DL2 + (l-y1)*(s-x1)*DLS \qquad (17b)$$

where

$$DS1 = (a2-a1)/(x2-x1)$$
$$DL1 = (b2-b1)/(x2-x1)$$
$$DS2 = (a3-a1)/(y3-y1)$$
$$DL2 = (b3-b1)/(y3-y1)$$
$$DLS = [(b4-b3)-(b2-b1)]/[(x2-x1)(y3-y1)]$$
$$DSL = [(a4-a3)-(a2-a1)]/[(x2-x1)(y3-y1)]$$

$$\text{Pixel intensity } p(s,l) = P1 + (l'-il')*(P2-P1) \qquad (18)$$

where

$$P1 = p'(is',il') + (s'-is')*[p'(is'+1,il')-p'(is',il')]$$
$$P2 = p'(is',il'+1) + (s'-is')*[p'(is'+1,il'+1)-p'(is',il'+1)]$$
$$il' = INT(l')$$
$$is' = INT(s')$$

and the function INT() denotes real-to-integer conversion with truncation.

Geometric warping is performed by first reading in a grid row of the regularly spaced control points. The minimum and maximum input line numbers are computed and the corresponding portion of the input image is read into memory (see Figure 8). (This algorithm assumes that the geometric distortion does not have a large degree of rotation.) For each grid, geometric transform coefficients are computed using Equations 17a and 17b. The portion of the output image corresponding to this grid row is then constructed using the bilinear intensity interpolation (Equation 18) one line at a time for all bands. Following this, the next grid row is read and the process repeats until all grid rows have been processed. Pixels in the output image that map outside of the input image are set to zero.

5. FUTURE PLANS

Future plans for the AVIRIS ground data-processing system include the acquisition of a 1024 by 1024 International Imaging Systems IVAS display device. This will allow viewing of a full swath-width AVIRIS image. The possibility of archiving AVIRIS data on optical disk, rather than on CCT, is being investigated. The advantages to archiving the data on optical disk are a large reduction in physical storage space, increased media reliability and lifetime, and a reduced requirement for operator intervention. The disadvantages are increased media costs and slower read/write speeds. Finally, it is envisioned that sometime in the future the operation of both the AVIRIS instrument and the ground data-processing system will be shifted to Ames Research Center.

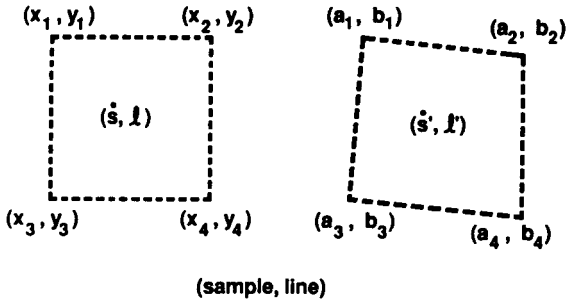

Figure 7. Location mapping.

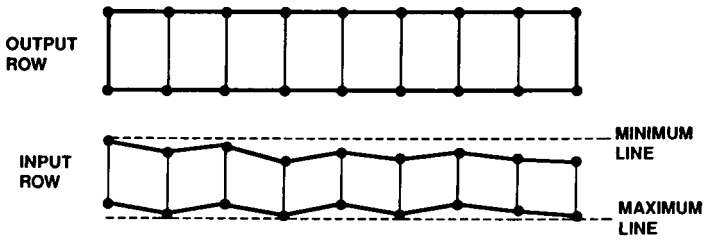

Figure 8. Grid row of control points.

6. REFERENCES

1. G. Vane, A. F. G. Goetz, and J. B. Wellman, "Airborne Imaging Spectrometer: A new tool for remote sensing," IEEE Transactions on Geoscience and Remote Sensing, Vol. GE-22, No. 6, pp. 546-549 (1984).

2. G. Vane, M. Chrisp, H. Enmark, S. Macenka, and J. Solomon, "Airborne Visible/ Infrared Imaging Spectrometer (AVIRIS): An advanced tool for earth remote sensing," Proceedings of the 1984 IEEE International Geoscience and Remote Sensing Symposium, Publication No. SP215, European Space Agency, Paris, pp. 751-757 (1984).

3. A. W. Babister, Aircraft Stability and Control, Pergamon Press, N.Y. (1961).

4. G. K. Manacher and A. L. Zobrist, "A fast, space-efficient average-case algorithm for the greedy triangulation of a point set," SIAM Journal on Computing (submitted).

7. ACKNOWLEDGMENT

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.