Spatial Calibration and Imaging Performance Assessment of the Advanced Land Imager

David R. Hearn

In 1998, the imaging properties of the Advanced Land Imager on the Earth Observing 1 (EO-1) satellite were calibrated at Lincoln Laboratory under simulated flight conditions. The spatial calibrations estimated the system transfer functions for the various spectral bands, and the line of sight of each detector, relative to a reference cube on the instrument. The Advanced Land Imager is a push-broom scanner. For each spectral band, a two-dimensional image is acquired one line at a time as the spacecraft moves over the scene. To reconstruct a normal image, we take into account the staggered placement of the detectors on the focal plane, the distortions in the optical system, and the speed and orientation of the spacecraft. Applying the line-of-sight calibration data, we have produced system-corrected images from cues in the image data alone. We create browse images to show the entire scanned scene in any three of the available multispectral bands or in the panchromatic band. The panchromatic data have also been combined with the red, green, and blue multispectral data to produce natural-color images with ten-meter ground sampling distance. Following the launch of EO-1 in November 2000, the imaging performance was assessed from the on-orbit data. This was primarily done through analysis of images of bridges. In addition, the moon, stars, and planets were scanned. To reconstruct celestial images requires reference to the position and attitude data from the telemetry stream. The moon in particular provided valuable data to assess weak stray-light and ghost-image artifacts. In general, the images returned from orbit confirm the pre-launch calibrations, and fully satisfy the requirements developed for them at the start of the program.

During the design phase of the Advanced Land Imager (ALI), we estimated its spatial resolution analytically. After the instrument was built, we measured its performance, both in the laboratory and on-orbit. In addition to spatial resolution, the lines of sight of all of the detectors had to be measured precisely, to enable the reconstruction of multispectral images with a minimum of geometric errors. We refer to the measurements made during the pre-launch phase, mostly performed in the laboratory, as spatial calibration. Once the instrument was launched and on-orbit, we refer to the measurements as imaging performance assessment.

Our goal was to do a thorough characterization during the preflight phase in order to minimize the questions and uncertainties over whether the instrument was performing correctly on-orbit. A major concern was to make sure that the instrument was properly in focus after the integration of the focal plane. There was no provision for adjusting the focus other than by changing the focus shim between the telescope and the focal-plane system. Much of the on-orbit performance assessment would center on verifying that the instrument was still in focus after the rigors of the launch. In the end, the ALI spatial resolution performance on-orbit is essentially identical to the performance predicted by analysis, and agrees with the pre-launch spatial calibrations.

The following sections review the spatial calibration requirements we established, the equipment we assembled to perform the calibrations, and the procedures we followed. We describe the process of reconstructing system-corrected images from the radiometrically calibrated data, using the results of the detector line-of-sight calibration. We describe on-orbit performance assessments, using images of bridges, stars, planets, and the moon, and we describe a process to produce color images with 10 m resolution by combining 30 m multispectral and 10 m panchromatic data. Finally, we summarize the imaging performance results.

Spatial Calibration Requirements

The ALI was designed to demonstrate imaging performance equal to that of the Landsat multispectral imaging satellites [1]. In particular, the Enhanced Thematic Mapper Plus (ETM+) of Landsat 7 has a basic ground sampling distance (GSD) of 30 m, and an imaging swath width of 185 km. In addition, Landsat 7 has a panchromatic band with 15 m GSD. For the ALI, the panchromatic band has 10 m GSD. This corresponds to an angular instantaneous field of view of 14.2 μrad, or 2.93 arc sec. At the altitude of Landsat 7 (and EO-1), the 185 km swath subtends 15°. The ALI was therefore designed to have a 15° field of view, but to make this technology demonstration system more economical only 3° of the telescope's field of view was furnished with detectors. The detectors were placed at one end of the elongated field of view, where the optical aberrations and distortion are largest.

Our overall calibration plan was to make use of measurements at different levels of integration, from components to subsystems to the full instrument system. The system-level calibrations were done with the instrument operating as much as possible under the thermal and vacuum conditions it would experience on-orbit.

Modulation Transfer Function

A generally accepted way to characterize spatial resolution of an imaging system is by its modulation transfer function (MTF) [2]. This is the magnitude of the Fourier transform of the system point-spread function. It specifies the response of the system to objects as a function of spatial (or angular) frequency. We adopted the Landsat 7 specification for minimum MTF at the Nyquist frequency (half of the sampling frequency) as the spatial resolution specification for the ALI. The size of the optical aperture and the maximum wavefront error of the telescope were derived from this specification. In addition, the Landsat 7 specification included minimum values for the MTF at one-half and one-third of the Nyquist frequency.

Our challenge was to measure the MTF of the completed instrument while it was operating in its thermal vacuum environment. We desired a thorough and accurate MTF measurement to facilitate the image-sharpening process called MTF compensation. We anticipated that after launch, MTF compensation would be applied to the images, and the quality of the results would depend on the accuracy of the MTF calibrations we performed.

Lines of Sight

As illustrated in Figure 1, the optical detectors on the focal plane of the ALI are not arranged in a simple matrix of rows and columns. Instead, there are four sensor chip assemblies (SCA), each with ten separated rows of detectors for the various wavelength bands. Within each row, the odd and even detectors are staggered in the scan direction by two to six detector pitches. The four SCAs are also staggered to provide some overlap in the cross-scan direction. All of these detector placements must be taken into account in order to produce a clear image from the raw image data. In addition, the optical system has distortions that must be allowed for. Our spatial calibrations thus had to include measurements to determine the apparent line of sight (LOS) of each of the detectors on the focal plane.

The LOS calibration consists of two parts. One is the absolute pointing direction of the telescope, in the reference frame of the spacecraft attitude control system. The other part is the relative LOS of each detector

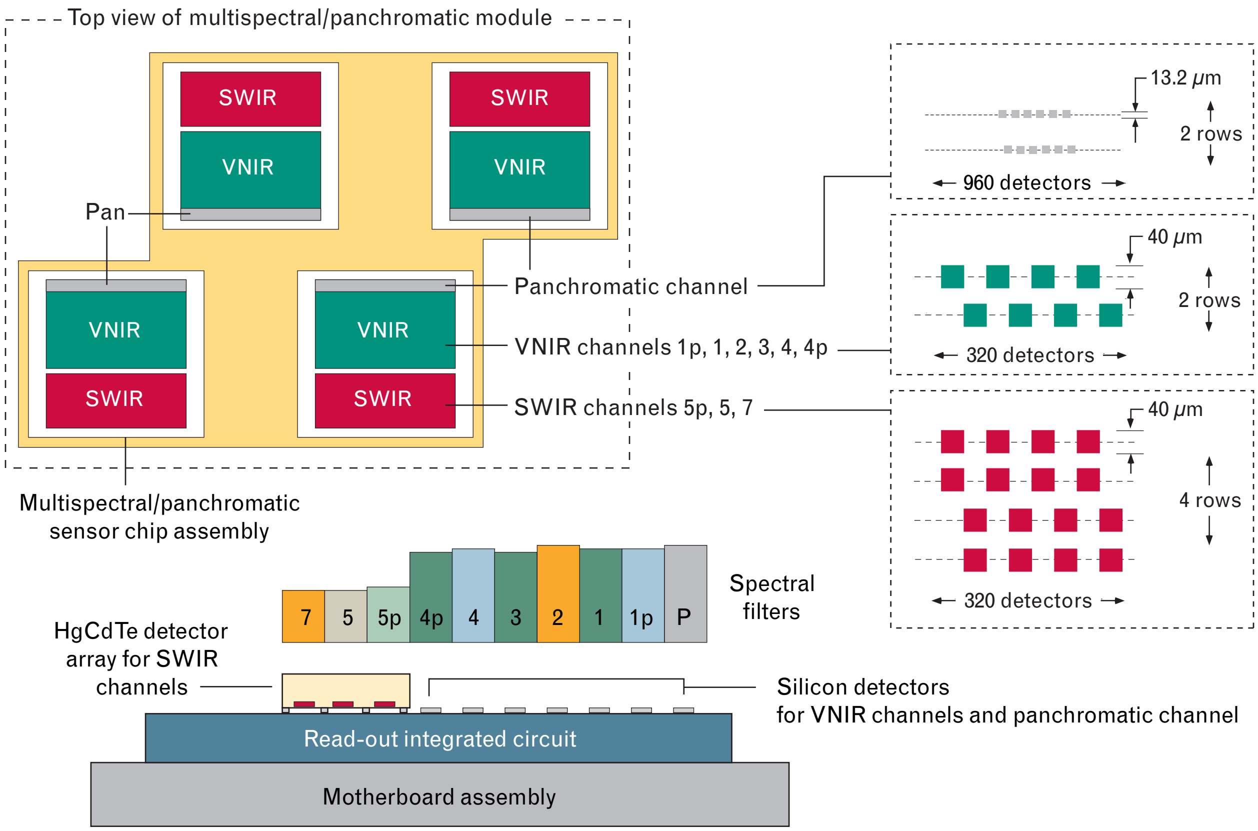

FIGURE 1. Arrangement of detectors within a sensor chip assembly (SCA). Four SCAs are mounted on the focal plane, with a small overlap in the cross-scan direction. The visible to near infrared (VNIR) detectors are PIN photodiodes formed in the silicon of the read-out integrated circuit (ROIC). The shortwave infrared (SWIR) detectors are HgCdTe photodiodes in a separate chip, indium bump-bonded to the ROIC. For the SWIR bands, a double row of detectors was provided for redundancy. In practice, only the outer rows are used.

in relation to the telescope. Since the LOS calibrations had to be done on the ALI before it was integrated with the spacecraft, we measured the telescope optical axes relative to a reflective reference cube mounted for that purpose on the base pallet of the instrument. After integration with the spacecraft, the relationship of that reference cube to the attitude control system was to be measured.

The relative LOS values of the focal plane are described by a set of parameters from which the relative LOS of any individual detector can easily be calculated. First, the layout of the detectors on each SCA follows a design produced by the focal-plane vendor. The microlithographic process for producing the sensor chip is accurate to a fraction of a micron, so we considered it to be free of positioning errors within the chip. Other focal-plane array parameters include the placement positions of the SCAs on the focal-plane module (two translations and one rotation each). The optical distortions of the telescope are described by general cubic polynomials in the in-scan (x) and cross-scan (y) dimensions. Sixteen polynomial coefficients are used to specify each of the x and y components of the distortion vectors. Finally, the telescope focal length is a key parameter. A data file produced from all these parameters lists the focal length and the apparent position on the focal plane of every detector. The angular LOS values can then be calculated easily by treating the telescope as a simple distortion-free lens.

Spatial Calibration Equipment

The system-level calibrations of the ALI were determined in a class-1000 clean room used for earlier projects at Lincoln Laboratory. Inside the clean room is a thermal vacuum chamber large enough to mount the ALI. The front door of the chamber has a window twelve inches in diameter. We found it possible to mount the ALI inside the chamber so that the instrument would have a clear view out through the window. The existing chamber window, of fused silica, was measured to have excellent optical transmission over the required wavelength range, and acceptably small transmitted wavefront errors. This enabled us to place all the calibration optical equipment in the more spacious area outside the chamber, where the equipment would not have to be compatible with the thermal vacuum environment.

Imaging Collimator

Spatial calibrations of the ALI required a collimated beam of light at least 12.5 cm in diameter to fill the entrance pupil of the instrument. The collimator optics had to be diffraction limited, in order to calibrate the instrument, since it is close to diffraction limited. We also wanted the collimator to provide a field of view up to 3° in diameter in order to test the ALI with a full, simulated scene such as it would scan on-orbit. This requirement posed a challenging optical design problem, since such a large field of view is difficult to achieve while maintaining diffraction-limited imaging over the field. The collimator also had to perform over a very wide wavelength range, from 400 to 2500 nm. Calls to vendors of stock and custom collimators failed to find any commercial solutions.

Bert Willard, an optical designer at Lincoln Laboratory, came up with a neat solution, shown in Figure 2. He called it a Schmidt-sphere imaging collimator [3]. It is based on a spherical mirror with a 3 m radius of curvature and 30 cm diameter. The entrance pupil of the ALI is placed at the center of curvature of this mirror. The success of this mirror relies on the Schmidt principle that a collimated beam through the entrance pupil would come to a focus on a spherical surface of half the radius of the mirror. By symmetry, the quality of the spot does not depend on the field angle. A true Schmidt system uses a corrector plate to correct spherical aberration, but in this instance, the f number is 12, which makes the spherical aberration quite small. Unlike a parabolic-mirror collimator, there is no coma to limit the usable size of the field.

To place the imaging surface of the Schmidt-sphere collimator outside the collimated beam, Bert inserted a thin beamsplitter. A convex lens is provided close to the focal surface to flatten the field. An imaging target or reticle is placed at the flattened focal surface. A quartz-tungsten halogen lamp, a six-inch integrating sphere, and a condensing lens behind the image plane provide illumination. The condensing lens is designed so that the output port of the integrating sphere is reimaged to fill the input pupil of the ALI, which is at the output pupil of the collimator. In this way, the ALI input pupil is fully and uniformly illuminated.

The reticle at the focal plane of the collimator is

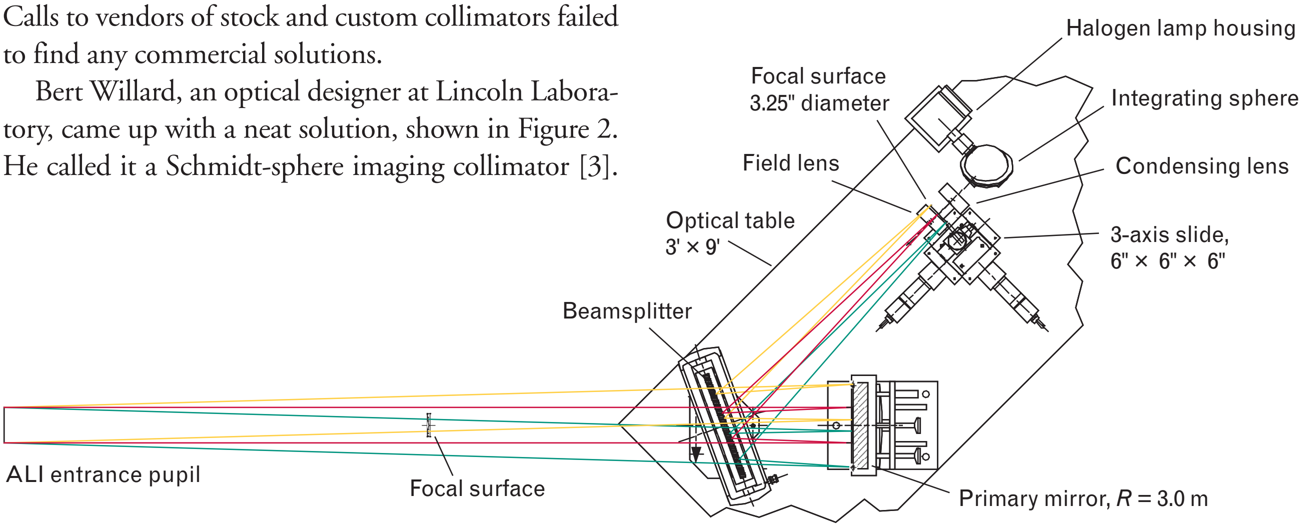

FIGURE 2. Optical layout of the Schmidt-sphere imaging collimator. The combination of the primary mirror, the field lens, and the condensing lens forms an image of the output port of the integrating sphere at the entrance pupil of the ALI, three meters from the primary mirror.

FIGURE 3. The setup of the Zygo interferometer, showing the flat mirror in place to measure the optical power of the vacuum chamber window. The fringes showing on the monitor at the left were formed by interference between the reflections from the front and back surfaces of the window. Their curvature indicates optical power, which is analyzed and displayed by the Zygo system.

supported by a three-axis translation stage. Since the ALI was placed in the chamber with its *x*-axis vertical, the vertical translation stage was used to simulate the motion of the scene as it would appear from orbit. This was essential in order to obtain two-dimensional images from the linear detector arrays of the ALI. The *y*-axis stage could be used to move targets in the cross-scan direction, and the *z*-axis was for focus adjustment. In order to scan knife edges slowly and smoothly across the field of view, two small linear stages were placed on the main stages for fine *x* and *y* motions. All of the stages were controlled by a computer programmed in LabVIEW.

Thermal Vacuum Window Problem

Our tests of the fused-silica chamber window, when it was removed from the chamber, showed it to be of excellent quality for transmitting the test images from the collimator. This result was confirmed after instrument integration by a series of test images made by the ALI from the imaging collimator. Images were formed first with no intervening window, then through the window but not under vacuum. Next, the vacuum was applied. All of these images were excellent. When the thermal chamber wall was cooled to about 100 K to produce the full thermal vacuum conditions, however, the instrument appeared to go out of focus! This was not supposed to happen, as the ALI optical system had been carefully designed to be athermal, with an Invar metering truss. We discovered that when the chamber wall was cold, the window developed a temperature gradient of several degrees Celsius from the mounting rim to its center. The roughly parabolic temperature distribution, when combined with the temperature coefficient of the index of refraction of the fused silica, was enough to turn the window into a weak lens [4].

We set up a Zygo interferometer, as shown in Figure 3, to measure the optical figure of the window while it was on the chamber. Fringes were formed by interference between the two reflections from the front and rear surfaces of the window. Curvature of the fringes indicated the optical power. We did an experiment to apply a heater outside the window to counteract the cooling from the chamber wall. We showed that the window optical power could be eliminated in this way, but since the heater had to be removed to conduct calibrations, the window would be cooling during the calibrations, and changing its power. This effect happened too quickly to provide a sufficient period for the calibrations.

A much simpler solution to the powered-window dilemma was to measure the power of the window with the Zygo interferometer, then offset the focus position of the target in the collimator. The amount of offset was computed to bring the collimator/window optical system to true collimation. Each time a series of imaging collimator measurements was to be made, we inserted a 12-inch-diameter flat mirror between the window and the collimator, to direct the Zygo beam to the window. The Zygo system computed the power of the window, and we applied the correct focus shift to the collimator. The 12-inch-diameter flat mirror was then removed. Repeated measurements showed that the window power was very stable over hours and days, since all temperatures were quite stable.

Positioning Fixtures

During the process of integration in the clean room, the ALI was supported by a Flotron fixture under a class-100 clean hood. The Flotron fixture is designed so that the instrument could be rotated to various tilt angles. Jackscrews were added to lift the fixture off its casters when it was time to perform optical measurements with the imaging collimator or a theodolite.

In the vacuum chamber, the ALI was mounted to a special fixture to permit it to rotate 15° about a vertical axis, to present any part of the focal plane to the view of the collimator. The axis of rotation was placed to pass through the entrance pupil of the ALI, to maintain the alignment with the imaging collimator as the ALI was rotated. The rotation position was sensed by an inductosyn, with arc-second precision. Control of the rotation stage was managed through a computer running a LabVIEW program.

Data Acquisition and Control Computers

A number of computers were essential to the calibration operations in the laboratory [5]. A custom bus interface box was built to emulate the spacecraft processor and communicate with the ALI control electronics (ALICE) through a 1773 fiber-optic bus link. ALICE received commands from the ALI calibration control node (ACCN), a Windows PC running LabVIEW. The ACCN also controlled the positioning stages in the imaging collimator.

Image data were received through high-speed data acquisition cards added to a Sun Microsystems workstation called the Electronic Ground Support Equipment 1 (EGSE1). A program on the EGSE1 descrambled the image data from the native telemetry format of the ALI. The image data were transferred from the EGSE1 to a Silicon Graphics workstation called the performance assessment machine (PAM). This machine was equipped with a digital linear tape drive for permanent archiving of the data. The PAM was also used during the laboratory data acquisitions to perform quick-look analyses. The results were displayed on a monitor in the clean room, next to the ACCN, so that we could verify that satisfactory data had been recorded.

Theodolite

For measurements of the absolute alignment of the telescope optics with respect to the optical reference cube, we used an electronic theodolite, Zeiss model Eth-2. A sturdy elevating tripod was also essential for this task. The theodolite automatically referred itself to the vertical direction of gravity. We had a 12-inch-diameter flat mirror in a fixed location to provide an azimuthal reference by autocollimation. Finally, we used a pair of bright flashlights to illuminate the focal plane while we sighted it through the ALI optics.

Focal-Plane Integration

A major milestone was reached when the focal-plane assembly was mated with the telescope of the ALI. For the first time, the ALI would be capable of producing image data. At this point, as planned, we conducted an end-to-end imaging test. We placed a standard 1951 U.S. Air Force resolution target in the collimator and set the collimator for infinite focus (collimation). The target was translated vertically at the correct speed while image data were being recorded. The raw data were descrambled, and simple frame-number shifting was done on the odd and even detectors to reconstruct the image. We were gratified to see the Air Force reso-

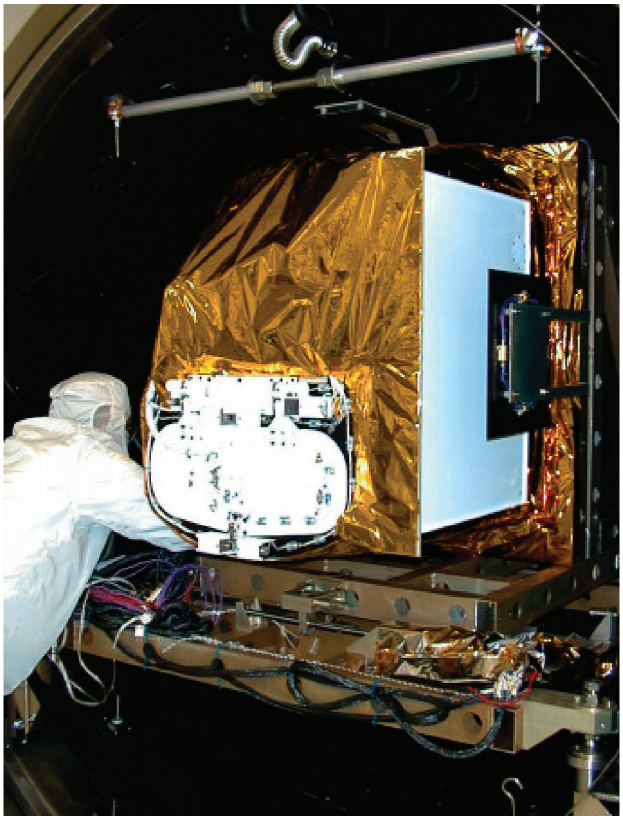

FIGURE 4. Installation of the ALI in the thermal vacuum chamber. The platform immediately supporting the instrument is pivoted at the rear, and driven by a lead screw visible at the front.

lution target appear as it should in the reconstructed image. This result verified that everything from the electronic read-out circuits on the focal plane to the reconstruction software was working correctly.

The telescope and the focal-plane array were separated with a shim, a few millimeters thick, which was trimmed so that on assembly the focal-plane array would be at the true focus of the telescope. It was critical to get this adjustment right, because there was no focus adjustment mechanism available after this integration. We tested the focus of the ALI by shifting the focus of the imaging collimator by precisely known amounts. At each setting, a knife edge was scanned across a portion of the field to record an edge-spread function for the detectors there. The steepness of the edge was a measure of the focus. When the steepness was plotted against the collimator focus offsets, the point of maximum steepness showed the best focus. The first such set of edge scans showed that the focal-plane array was almost a millimeter away from the true focus position. After the focus shim was machined to the new indicated dimensions, the focal-plane array was reintegrated and the procedure was repeated. This time no significant focus error was indicated. The process could have been repeated as many times as necessary, but one iteration was enough.

Spatial Calibration Procedures

For the spatial calibrations, the ALI was mounted on the positioning fixture in the thermal vacuum chamber, as shown in Figure 4. The chamber was evacuated and its cold wall cooled to about 100 K. This produced the operating environment that the instrument would experience on-orbit. The collimator was positioned in front of the chamber window so that its exit pupil coincided with the entrance pupil of the ALI. For each data collection period, the collimator was first tested to set the targets at the true focus, using a laser unequal-path interferometer and a shear plate [3]. Preliminary tests revealed the problem (described above) of the optical power of the vacuum window.

During the spatial calibrations, normal scans were acquired with a starburst pattern target, as well as with the Air Force target as described for the end-to-end tests. The starburst pattern has alternating black and white lines at all angles radiating from a common center. This pattern can provide a quick estimate of the lower-frequency MTFs at a variety of angles.

MTF Procedures

For an optical system, the MTF is the magnitude of its optical transfer function. A full description of the imaging properties must include the phase transfer function. We chose to treat the MTF of the ALI system (optics plus detectors) as a complex quantity, incorporating both magnitude and phase. To avoid confusion, we call this the system transfer function, or STF, though colloquially we may still call it the MTF.

The STF is the Fourier transform of the point-spread function (PSF). A one-dimensional slice through the center of the STF is the Fourier transform of a line-spread function of a line source orthogonal to the direction of the slice. Since it is difficult to present either a true point source or line source to the instrument in the laboratory, we chose to analyze measured edge-spread functions (ESF). A line-spread function is the spatial derivative of an ESF. The measurement data for an ESF can be acquired readily by projecting the image of a sharp edge onto the detector plane.

Knife-Edge Scans

We required measurements of the STF for spatial frequencies well beyond the Nyquist frequency. This required many data samples per detector pitch. For a two-dimensional imaging detector such as a charge-coupled device, this measurement is usually accomplished by placing the sharp-edge image at a small angle to the columns (rows) of detector pixels. Successive rows (columns) then sample the edge image at finely spaced fractional pixel phases. For the ALI focal plane, this placement was not an option, at least for an ESF in the cross-track direction. The detectors are arranged in single rows, with odd and even detectors staggered in the in-track direction.

To acquire the ESF data, we mounted the knife-edge target on small precision-translation slides at the focus of the collimator. The slide was then moved slowly while we collected focal-plane data, so that approximately seventy samples were recorded as the edge traversed a single detector pitch. For a knife-edge target, we used the clear opening, approximately 9 mm square, in the 1951 Air Force resolution target (the rest of the target was covered with an opaque mask). In that way, the same target served for both horizontal (cross-track) and vertical (in-track) scans. Good data were acquired from approximately forty detectors per band in a single scan. Several bands were also scanned at once. Scans were made to cover all of the spectral bands of each of the SCAs.

Data Analysis

The knife-edge scan data were first radiometrically corrected to remove variations of offset and gain among the detectors. We selected a subset of the scan data that included only the detectors that were fully scanned by the knife edge. Next, each individual detector's ESF data were fitted to a hyperbolic tangent function, merely to ascertain the edge-crossing sample for that detector. All the good detectors' data from a single band were then resampled to co-align the crossing points, as plotted in Figure 5. Fourier transforms were used to differentiate the ESF data to obtain line-spread functions.

To avoid difficulties caused by the difference between the values at the two ends of the scan, we appended a copy of the ESF, in reverse order, to each ESF scan array. The now-symmetrical ESF array was then transformed to the frequency domain with the fast-Fourier-transform procedure in the IDL language. High spatial frequencies were rolled off with a filter function, based on the diffraction-limit frequency. Next, the array was multiplied by the corresponding spatial frequencies, resulting in the one-dimensional STF.

The inverse transform of the STF is the line-spread function. The ESF, STF, and line-spread function arrays were computed for each detector individually, and the means and standard deviations were computed for each spectral band and SCA combination. Figures 6 and 7 show typical plots of these data. The imaginary part of the one-dimensional STFs was generally smaller than a few percent in magnitude. The in-track scans produced peak imaginary components as large as ±0.06, but with no clear pattern.

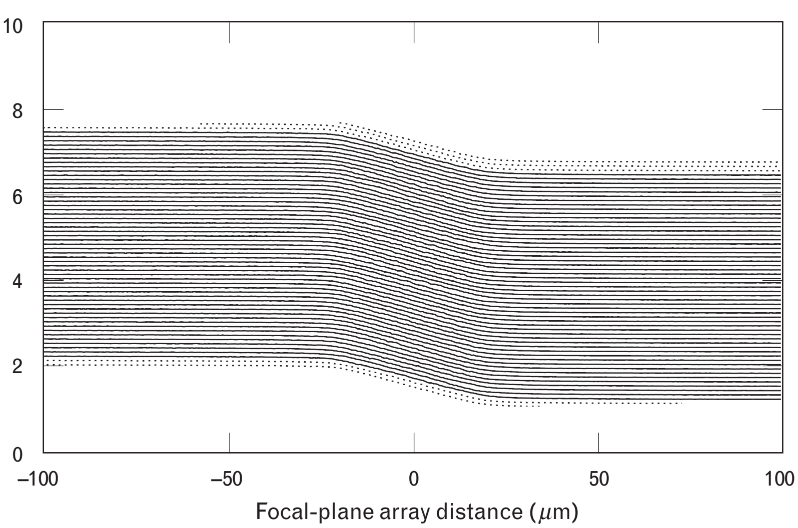

FIGURE 5. Edge-spread functions of the band 4 detectors of SCA 2, scanned in the cross-track direction. Pixels not used in further analysis are shown with dotted lines

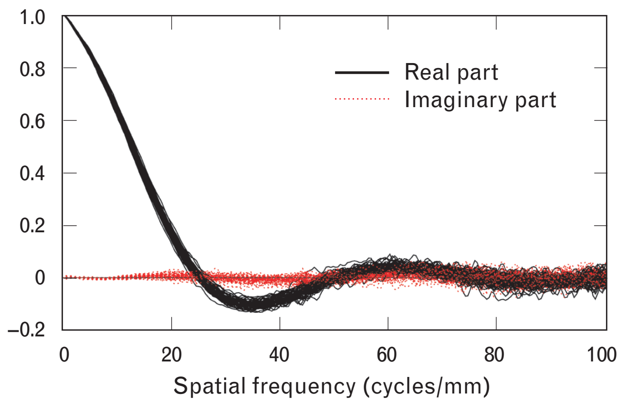

FIGURE 6. Normalized system transfer functions derived from the data shown in Figure 5. This graph shows both real and imaginary parts, though the imaginary parts are consistent with a zero value.

We suspect that this result may be an artifact caused by the vertical slide, which we later learned was moving somewhat erratically.

Full Two-Dimensional STF Model

The estimated STFs derived from the knife-edge measurements are one-dimensional. Furthermore, they are a large collection of individual data arrays, representing the various spectral bands and SCAs. We wished to have a compact and comprehensive description of the full, two-dimensional STF of the ALI. This description would be needed for the assessment of on-orbit imaging performance. We also anticipated that the land images would be sharpened up by MTF compensation, a technique frequently applied to Landsat images. In MTF compensation, the higher spatial frequencies (up to the Nyquist limit) are carefully amplified to compensate for the decreasing modulation with frequency, described by the MTF of the instrument. To do this amplification precisely, we needed a good estimate of the instrument's MTF.

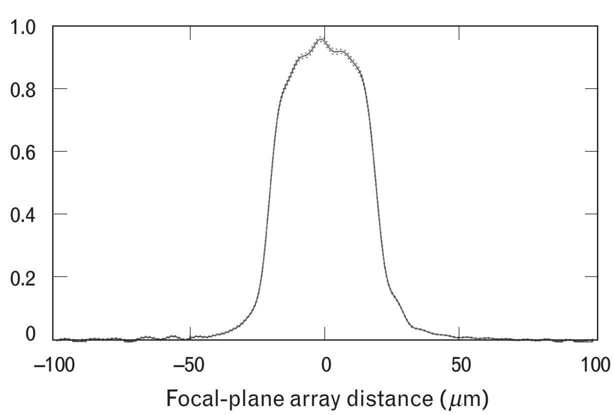

FIGURE 7. An average line-spread function derived from the data shown in Figure 5. The dotted lines are one standard deviation above and below the mean.

Our approach was to construct an analytical model of the full two-dimensional STF, incorporating the effects of the optics, the detectors, and the motion of the image during the frame integration time. The parameters of the model were adjusted to provide the best match to the one-dimensional STF estimates from the knife-edge scans. Fortunately, the optical telescope had been measured by the builder, SSG Inc., with a laser unequal-path interferometer. These measurements consisted of a series of interferograms, representing the optical wavefront error across the pupil, as seen from a series of points distributed over the focal plane. SSG analyzed the interferograms and fit each one with a set of Zernike polynomials. The Zernike polynomials are a set of orthogonal functions defined over a circular pupil (see the sidebar entitled "What Is a Zernike Polynomial?"). With the Zernike coefficients we wrote procedures in IDL to describe the wavefront error across the pupil (see Figure 8), the optical transfer function (see Figure 9), and the PSF (see Figure 10) of the telescope for any location on the focal plane. One of the coefficients, proportional to focus error, would be adjusted to match the measured results. The rest were left fixed.

The system transfer function for static images is the product of the optical transfer function and the MTF of an individual detector. The detectors were initially modeled as uniform rectangles, with a sharp cutoff in response at the edges. The MTF of such a detector is simply a product of sinc functions running in the x and y directions.

$$MTF_{\text{pixel}} = \operatorname{sinc}(wf_x)\operatorname{sinc}(hf_y)\,,$$

where w and h are the width and height of the detecting area, and f_x and f_y are the spatial frequencies in the x and y directions, respectively.

We found that the data were much better fitted if we applied an additional factor in the detector MTF

WHAT IS A ZERNIKE POLYNOMIAL?

To have a compact description of the aberrations of an optical system, it is useful to expand the wavefront error over the pupil in a complete series of orthogonal functions. One such set defined on the unit circle is the set of Zernike polynomials, also called *circle polynomials* [1]. They can be written as

$$V_n^l(r\sin\theta, r\cos\theta) = R_n^l(r)e^{il\theta},$$

where (r, θ) are the polar coordinates within the circle $(r \le 1)$, n and l are integers, with $n \ge |l| \ge 0$, and $n - |l|$ is even. To stay with real functions, we can use

$$U_n^m = R_n^m(r)\cos m\theta$$
$$U_n^{-m} = R_n^m(r)\sin m\theta,$$

where $m = |l|$.

One formula for the radial functions is

$$R_n^{\pm m}(r) = \sum_{s=0}^{\frac{(n-m)}{2}} (-1)^s \frac{(n-s)!}{s!\left(\frac{n+m}{2}-s\right)!\left(\frac{n-m}{2}-s\right)!} r^{n-2s}.$$

Other indices are often used in listing the set of U_n^m, but the functions are the same [2]. Figure A shows plots of a few of the lower-order Zernike polynomials. The first polynomial, with (n, m) = (0, 0), would simply represent a constant phase shift, and has no effect on an optical image. The next two polynomials, (1, 1) and (1, –1), simply represent tilts of the wavefront, causing a lateral shift of the image. The polynomial (1, 0) represents a focus error, or axial shift of the image. The (2, 2) and (2, –2) polynomials correspond to astigmatism. Spherical aberration can be described in terms of the $n > 0$, $m = 0$ polynomials.

The wavefront error of good optical systems typically has very small amplitude at high spatial frequencies. Therefore, a relatively small set of Zernike amplitude coefficients is sufficient to characterize this error. The coefficients are usually in units of a test wavelength, such as 632.8 nm.

References
1. M. Born and E. Wolf, *Principles of Optics,* 6 ed. (Oxford Pergamon, New York, 1980), section 9.2.1.
2. D.R. Hearn, "EO-1 Advanced Land Imager Modulation Transfer Functions," Technical Report 1061, Lincoln Laboratory (22 Mar. 2000), DTIC #ADB-252266.

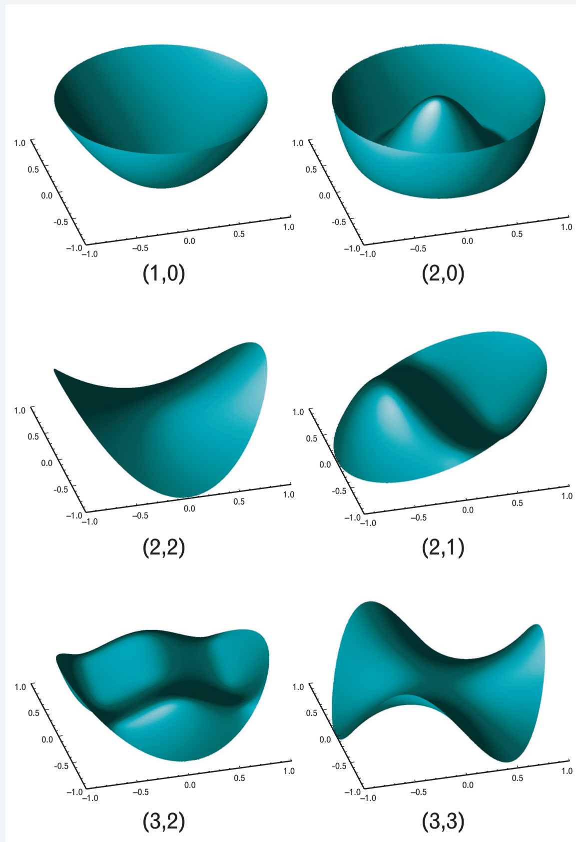

FIGURE A. Six of the low-order Zernike polynomials. Each one is labeled with its (n, m) indices. For $m \neq 0$, the (n, –m) polynomial looks the same as the corresponding (n, m) polynomial, but rotated by 90°/m.

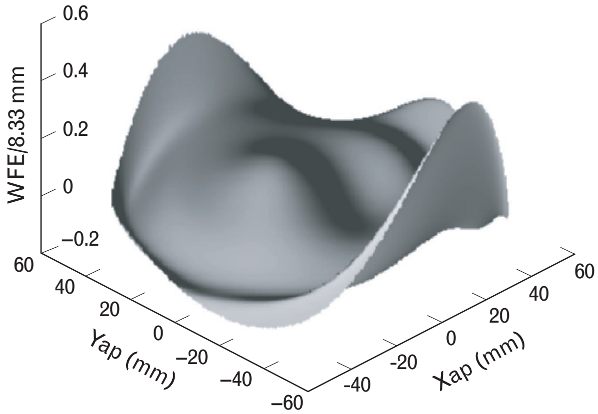

FIGURE 8. Surface plot of the wavefront error (WFE) near the middle of the multispectral/panchromatic focal-plane array. The focus term is not included here.

to reduce the highest frequency components. We attributed this problem to diffusion of charge carriers within the detectors. The carrier diffusion factor we applied was

$$MTF_{\text{diffusion}} = \exp\left(-\left|\frac{f}{f_0}\right|^g\right)$$

$$f^2 = f_x^{\,2} + f_y^{\,2} \,.$$

The detector MTF is then

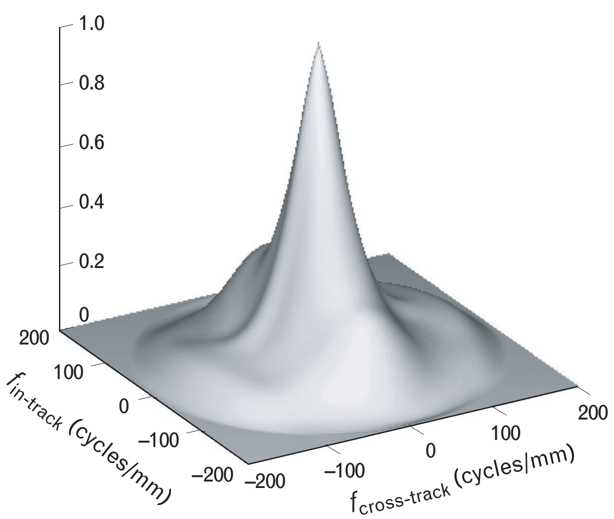

FIGURE 9. Optical transfer function of the panchromatic band (real part), computed from the wavefront error, including the focus term.

$$MTF_{\text{detector}} = MTF_{\text{pixel}} MTF_{\text{diffusion}} \,.$$

For the visible to near infrared (VNIR) detectors we found the values $g = 1$ and $f_0 = 200$ cycles/mm agreed with the knife-edge data, assuming the nominal design values for the detecting area (39.6 by 40 μm). The HgCdTe shortwave infrared (SWIR) detectors were fitted best with $g = 1.5$ and $f_0 = 35$ cycles/mm, with the geometric height and width reduced to 36.8 μm. Using these values, we computed the detector point-spread functions for the panchromatic band, the VNIR multispectral bands, and the SWIR multispectral bands, as shown in Figure 11.

Once we had settled on the set of model parameters that produced the best match to the knife-edge data,

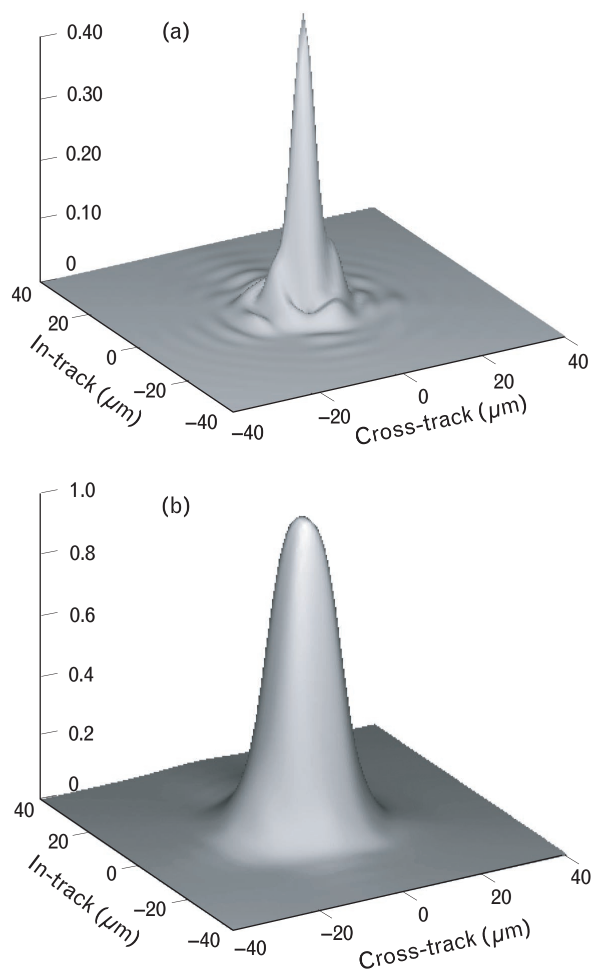

FIGURE 10. (a) Optical point-spread function of the panchromatic band, computed from the full wavefront error, including the focus term. (b) The system point spread function.

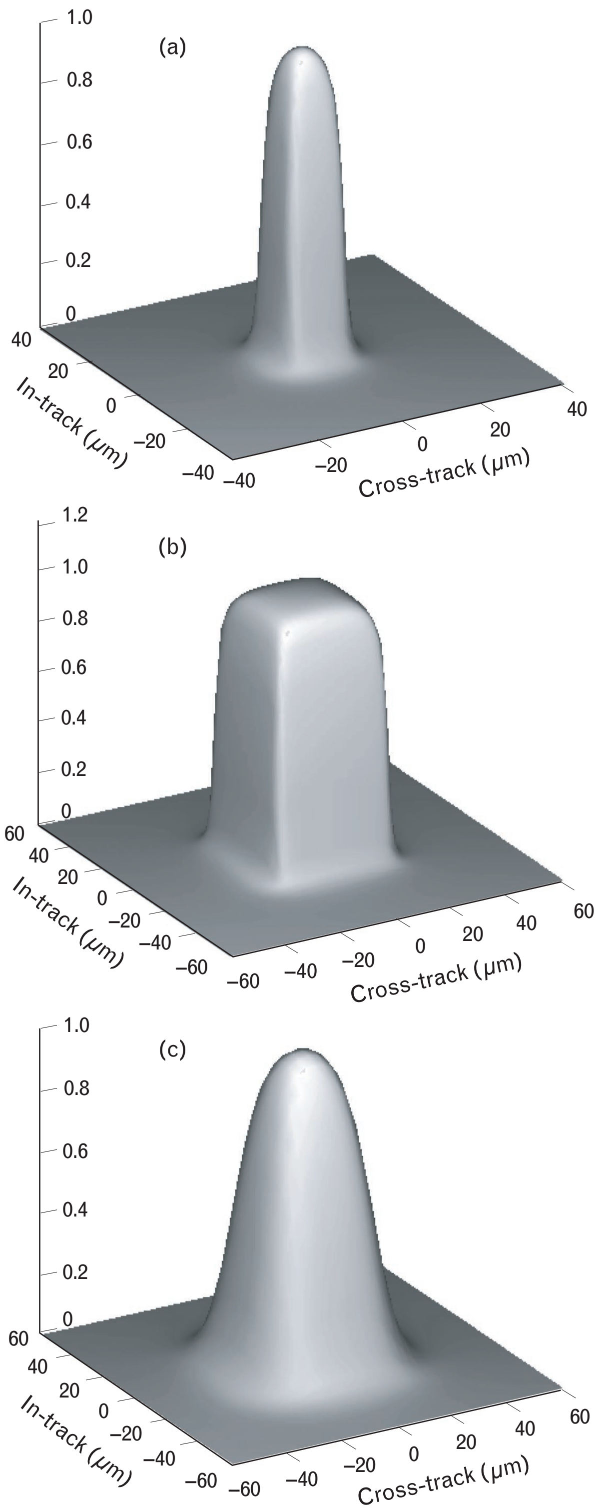

FIGURE 11. Detector point-spread functions for (a) the panchromatic band, (b) the 4p band, and (c) the 7 band. The plot for band 4p represents all VNIR multispectral detectors, and the plot for band 7 represents all multispectral detectors. Notice that the horizontal scale is larger for the panchromatic band plot.

we computed tables of values of STF over one quadrant of frequency space. The 17 by 17 tables extend to four times the Nyquist frequency for each band. The tables are contained in a previously published technical report [6]. A user can reconstruct the full STF from the one quadrant because the real part of the STF is symmetric about the origin, while the imaginary part is anti-symmetric.

We must apply an additional factor to the static STF described so far in order to obtain the complete on-orbit STF. Since the ALI scans the earth in a continuous motion, the image is slightly blurred by the motion during the frame integration time. The factor is

$$\mathrm{MTF}_{\text{integration}} = \operatorname{sinc}(\alpha \nu)\,,$$

where α is the angle of the apparent ground motion during the integration time, and ν is the angular frequency.

Line-of-Sight Procedures

The sparse layout of the detectors on the ALI focal plane also presented problems when it came to measuring the line of sight of the detectors. For a simple two-dimensional array such as a charge-coupled device (CCD), a precise reticle composed of a grid of fine lines or fiducial marks could have been placed in the collimator and imaged onto the focal plane. In the case of the ALI, many or most of the marks would fall between the detectors. To address this problem, we placed a Ronchi ruling in the collimator, with provision to rotate it about the optical axis. This type of ruling consists of an array of parallel lines at a uniform spacing, with equal clear and dark areas. When this ruling is projected on the focal plane, the relative position of each detector in the direction perpendicular to the lines can be estimated from the image data. The line-of-sight direction is derived from the position of the detector and the telescope's focal length.

Ronchi Images

Figure 12 shows a scanned image of the Ronchi ruling. We actually analyzed static images of the Ronchi ruling, made while the ruling was stationary in the collimator. Thus we eliminated any uncertainties related to the motion of the slide. Noise was reduced by averaging a hundred frames of data. In principle, two images taken with the ruling at different (preferably orthogonal) angles would be sufficient to determine the relative detector locations. Instead, we used three orientations, with the lines at 0°, 60°, and 120° from the vertical, in-track axis. This combination ensured that the detector positions would be over-determined. We corrected the raw detector readings by subtracting a dark frame and dividing by a light frame (taken with no ruling present). Some further overall adjustments were then applied to normalize the data range from 0 (dark) to 1 (light).

Data Analysis

To analyze the Ronchi image data, we performed a least-squares fit to the detector lines-of-sight model. Parameters of the lines-of-sight model included the positions of each SCA on the focal plane, a set of polynomial coefficients to describe the optical distortion, and the focal length of the telescope. The optical distortion field was modeled with cubic polynomials in x and z, resulting in sixteen coefficients each for the x and y components of the field. Initial values for the distortion coefficients were obtained by analysis of the optical distortion measurements of the telescope, performed by SSG with a theodolite.

In addition to parameters of the instrument, there were the setup parameters of ruling orientation and phasing at the origin. We assumed that the position of each detector within the SCA was described by the design layout of the SCA to sub-micron accuracy. To do the least-squares fit, we simulated the expected images from the line-of-sight model, and adjusted the parameters until the expected images best matched the data. Since the spatial-frequency content of the Ronchi ruling is well known, it was a simple matter to apply the system transfer functions already constructed in order to compute the simulated images.

The IDL procedure MPFIT was used in the least-squares fit. It is a more flexible version of the CurveFIT procedure supplied with IDL, which employs a gradient-expansion algorithm. In practice, the choice of a fine Ronchi ruling of 2.0 cycles/mm led to great difficulties in coaxing the fit to converge to the global minimum χ^2 value. This resulted in either 20 or 40 cycles of the ruling across each row of detectors. Unless the simulated signals had very nearly the same spatial frequency, the MPFIT procedure would settle into

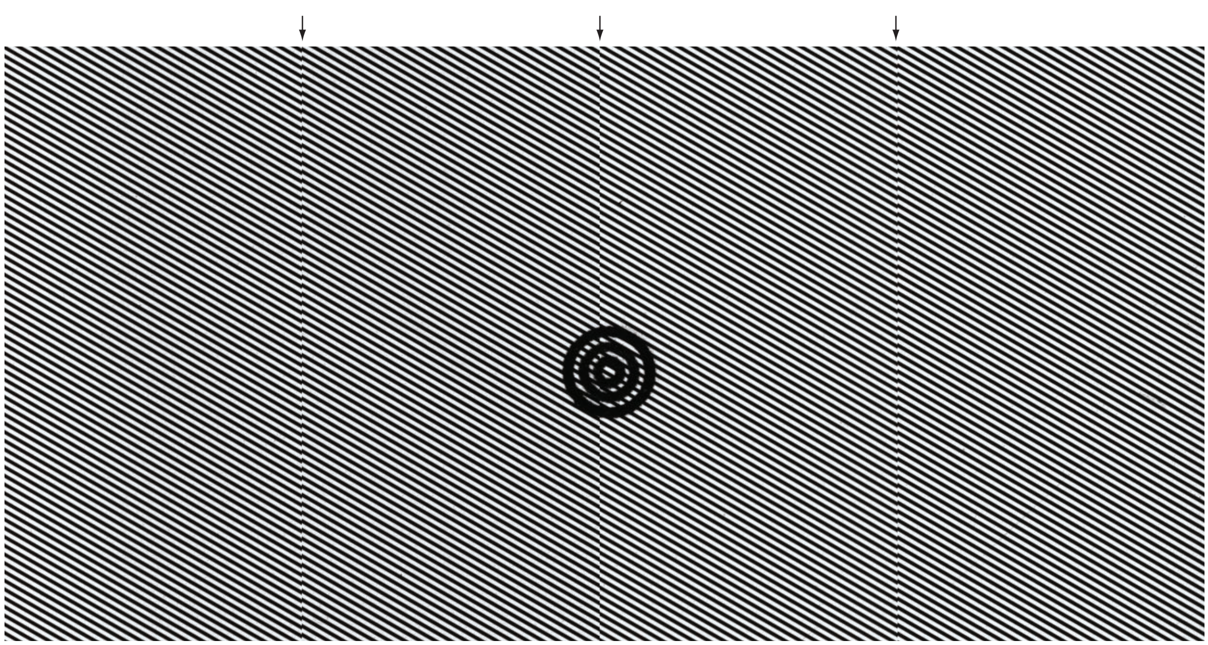

FIGURE 12. Scanned image of the Ronchi ruling. The bull's-eye indicates the axis of the collimator, so that the computed distortions from the collimator could be corrected. The parallel lines of the Ronchi pattern show offsets at the boundaries (indicated by arrows) between SCAs, because the four segments of the image were only approximately aligned.

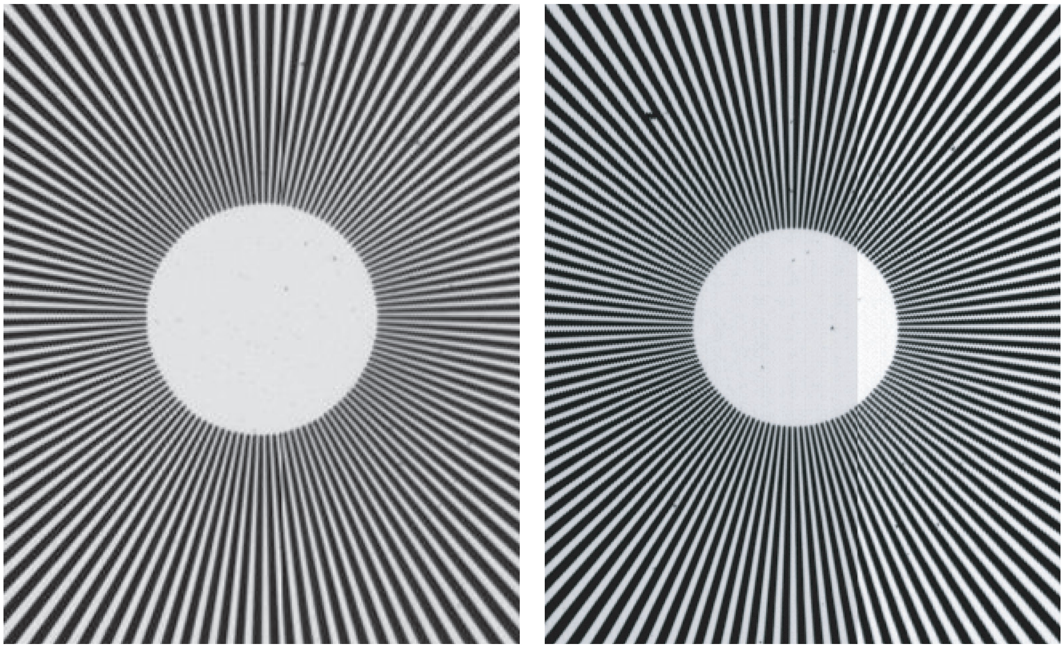

FIGURE 13. Close-up views of the starburst target before (left) and after (right) the environmental tests. The fact that the images are nearly identical shows that there was no loss of imaging performance as a result of environmental testing. The image on the right is not fully corrected radiometrically, which is responsible for the change in brightness from one SCA to the next.

a shallow local minimum in χ^2. Ultimately, by plotting the measured and simulated signals on screen, we were able to choose initial values very close to the true values for the important parameters that controlled the spatial frequency of the simulated signals. The fits then converged quickly to the global minimum.

We created a data file that gives the telescope focal length and the apparent position of each detector on the focal plane. From this, the line-of-sight direction is easily computed. All the parameters of the line-of-sight model were also published [7].

Reference Cube Sighting

The lines of sight derived from the Ronchi images are all relative to the optical axes of the ALI telescope. To align the telescope to the spacecraft attitude-control system, we fixed a mirrored optical reference cube to the pallet of the ALI. With an electronic theodolite in the clean room, we then measured the alignment of the reference cube relative to the telescope. We measured each visible cube face with an autocollimation sighting, in which a light beam from the theodolite is reflected directly back on itself, and the spot of light is aligned with a cross hair. Detector arrays on the focal plane were sighted directly with the theodolite looking into the entrance of the telescope. The theodolite has a vertical reference built in. For an azimuth reference, we placed a large flat mirror where it could be sighted in autocollimation from the several different positions of the theodolite. We converted all of the sighting directions in the laboratory frame defined by the local vertical and the azimuth reference mirror to the frame of the telescope axes. In this process, we relied on the relative line-of-sight calibration of the detectors.

Post-Environmental Testing

Following the integration of the ALI with the EO-1 spacecraft, a series of environmental tests were conducted at Goddard Space Flight Center (GSFC), including vibration tests. A post-environmental end-to-end imaging test was performed in a large clean room at GSFC to make sure that the ALI was still performing correctly. The EO-1 spacecraft was mounted on a heavy rotating table, and positioned with the ALI view axis horizontal. The imaging collimator was placed on extension legs to bring it to the necessary height, and placed in the proper position in front of the ALI. A normal scan was acquired with the starburst pattern target. After processing, this scan was compared with the corresponding starburst target scan acquired with the ALI alone at Lincoln Laboratory. As Figure 13 shows, no visible difference can be seen between the two images. We concluded that the ALI had withstood shipment, integration, and environmental testing without degradation of its optics and image-sensing systems, and that it could be launched with confidence in its spatial performance.

Reconstruction of Images from the Advanced Land Imager

The detectors on the focal plane of the ALI are sparsely arranged in staggered rows, as shown previously in Figure 1 [1]. As a consequence, the detector image samples must be rearranged to form a two-dimensional multiband image after the data stream is received. This rearrangement can be done to several levels. The image data are received as a series of frames, recorded at approximately 226 frames per second for the multispectral data. Each frame represents a simultaneous sampling of all the detectors. The simplest reconstruction is to shift detector samples by whole numbers of frames, to remove the offset of odd and even detectors within a band, and the offsets from band to band and from one SCA to another. The next level is to resample the data to take into account the true lines of sight of the detectors. This is called system correction. Finally, the data can be projected onto a map of the portion of the earth that was imaged. This is called geo-referencing. We have carried out the first two kinds of reconstruction; our techniques are described in the following section. We were not given the much larger task of making geo-referenced images. Nevertheless, we developed some of the necessary algorithms in the course of making maps of the moon and planets.

For scientific analysis of the ALI images, the Level 1R data files can be utilized. They contain radiometrically calibrated readings in every band [8]. If they have also been geometrically corrected, they are called Level 1G files. For the purpose of visual inspection, however, simpler files are appropriate. These contain “browse images.” We describe below how we make three-color browse images from the Level 1G files. The images reproduced in this issue are from our ALI browse image library [9]. We also describe below how we used the panchromatic-band data to produce color browse images with an apparent spatial resolution of 10 m ground sampling distance.

Whole-Frame Shifts

The basic pitch of the multispectral detectors on the focal plane is 40 μm in the in-track, or scanning, direction, and 39.6 μm in the cross-track direction. The panchromatic detectors are spaced on a 13.2 μm grid. We normally choose the frame rate so that the image moves across the focal plane at 40 μm/frame, which is the nominal image speed. The odd and even detectors within a single band are staggered in the scan direction by two or six detector pitches. In the image data files processed at GSFC, the odd detectors are shifted by that many frames, so that the image samples line up correctly in a simple raster image of one band. The focal-plane separation between multispectral bands is 20 pitch units. Therefore, a multi-band image can be brought into approximate registration by shifting the bands by multiples of 20 frames. The band 1p rows of adjacent SCAs are separated by 187 pitch units. This distance is approximately the additional shift required to match the image segments from the different SCAs.

The advantages of whole-frame shifts are that they are easy to do, and they can be reversed to restore the original information exactly. One disadvantage is that they do not account for optical distortion, the major factor in the detector line-of-sight correction. In addition, if the image speed is not quite right, the shifts can not produce exactly registered images.

System Correction

To correct for optical distortions, it is necessary to resample the image data. At the same time, we take into account the measured positions of the detectors. Both effects have been incorporated in the detector line-of-sight calibration file [7]. If the image of the scene moves across the focal plane with a known uniform speed, it is relatively straightforward to perform the interpolations. For normal scenes, we do a bilinear interpolation in two steps. First, each detector is resampled from the frame coordinate to the in-track distance coordinate in the image space. Then each row is resampled from detector number to the cross-track distance coordinate. We treat the frame and detector coordinates separately because the staggered placement of alternating detectors precludes a direct bilinear interpolation.

Our reconstructions of some early on-orbit images showed some mismatches between adjacent SCAs. As we learned when reconstructing scanned images simulated through our imaging collimator, it is necessary to make minor corrections to the assumed scan direction, as well as the speed. Typically, we found that the image moved across the focal plane at an apparent yaw angle on the order of one milliradian, and at a speed of about 0.97 detector pitches per frame. It was time consuming to adjust the speed and yaw parameters for every data collection event. That would not have been necessary if we had learned to read the telemetry files, and relate the image data to the attitude and position of the spacecraft.

An alternative approach was available to us, however. Since the SCAs were deliberately arranged to provide overlapping coverage between neighboring SCAs, we used the image data in the overlap regions to estimate the actual speed and yaw parameters. This was done with a procedure in IDL. We assumed that the speed and yaw were constant for the roughly twenty-five-second duration of the data collection. The procedure reads the Level 1R data and cross-correlates the 30-pixel overlap regions of the panchromatic data. It derives the speed from the in-track correlation, and the yaw from the cross-track correlation. Once the image velocity (speed and yaw) has been estimated, another procedure resamples the Level 1R data as described earlier, and writes out the reconstructed, system-corrected data, one file per SCA. Locally, we label these files as Level 1G data, though strictly speaking, that designation is usually given to geo-referenced images.

Browse Image Creation

In order to view the system-corrected images of the whole observed swath, we needed to create three-color browse images, which take into account the human visual response and some of the limitations of the computer display and printing devices. The browse image is created from selected bands of the Level 1G files for all four SCAs. We form a single array, abutting the adjacent SCA data by omitting half of the overlapping pixels of each neighbor. It is unnecessary to blend the two images to hide the joint, since the system-corrected sub-images generally match extremely well. For the multispectral images, the array dimensions are 1250 pixels across, 2000 to 6000 or more pixels long, and three colors deep. The panchromatic image dimensions are 3750 pixels across, three times as long as the multispectral array, and one color deep. Any three bands can be used in the browse images, which are then mapped into the red, green, and blue (RGB) color space of the computer display. A natural-color image results from using ALI bands 3, 2, and 1 for the red, green, and blue components, respectively. Another set of bands frequently chosen for browse images are 4, 3, and 2. The 4-3-2 images show healthy vegetation in bright red, so that it stands out better than in natural-color images.

To bring out the details of the images, we enhance the contrast of the darker areas and suppress the brightest areas by scaling the radiances logarithmically. Calculations were performed with MODTRAN to estimate the radiance that would be seen in each band as a function of the solar zenith angle, for surface reflectances of 0 and 100%. For the solar zenith angle of the current observation, we use those minimum and maximum expected radiances to scale the observed scene:

$$B_i = 256\left\lfloor\left(\frac{\ln\left(L_i\right)-\ln\left(L_{\min,i}\right)}{\ln\left(L_{\max,i}\right)-\ln\left(L_{\min,i}\right)}\right)\right\rfloor,$$

where L_i is the radiance in band i, and $L_{\min,i}$ and $L_{\max,i}$ are the minimum and maximum radiances. The result of the scaling is a byte value for each band, to form the 8-bit or 24-bit color array to display.

In practice, some adjustments were made to the predicted radiance ranges to allow for atmospheric variations. This scaling technique produces images with a good range of contrast, without saturating very dark or bright areas. It also suppresses path radiance from the atmosphere, which can cause the image to look blue and hazy. Once the image array is formed, the IDL procedure writes it to a file in JPEG format, which can be conveniently read by many other computer applications, especially Adobe Photoshop. In Photoshop, an image can be viewed, cropped, scaled, color balanced, and printed.

Map Projection of Images

The system-corrected images and browse images are created without reference to the geographic coordinates of the area that was imaged. They are like photographs taken by a well-corrected camera. Although we have not had the opportunity to produce geo-referenced images, for completeness we describe here the process of their creation.

The spatial position of the satellite relative to the center of the earth, and its attitude (orientation) are recorded for every second of a data collection event. With this information, and the geoid (the precise shape of the earth at mean sea level) and the rotational position of the earth, system-corrected images can be projected to geographic coordinates. A preferable procedure is to start with radiometrically corrected data, and apply the detector line-of-sight map directly in the process of producing the map projection [10]. This procedure avoids having to do a second resampling, and better preserves the resolution of the image. For full accuracy, it is necessary to apply a digital elevation model of the region, to correct for offsets caused by non-vertical lines of sight.

On-Orbit Image Assessment

Five days after EO-1 was launched, the ALI acquired its first earth image. Once we had received the data, and we were able to reconstruct the image, we were gratified to see that it was sharp, clear, and free of artifacts. Although we did not produce geo-referenced images, we applied the detector lines-of-sight calibration in the process of image reconstruction. Close examination of the system-corrected images shows that the band-to-band registration is well within one pixel, which validates the line-of-sight calibration.

Our detailed imaging performance assessment concentrated on the sharpness of the images. Instead of trying to derive an MTF from the on-orbit images, we applied the calibrated STFs to compare the actual sharpness of on-orbit image features with what we computed from the STFs. For this comparison, we needed scenes that contained sharp features that were reasonably simple to model. We chose bridges, the lunar limb, and stars for study. The performance assessments are summarized here, and are more fully described in another report [11].

Bridge Profiles

Cities next to bodies of water were placed on the priority scene list because the bridges they possess are well suited for our MTF evaluation. A bridge usually appears in the image as a bright strip on a fairly uniform dark background. If the bridge is reasonably long, then the pixel samples along its length will be at a variety of relative pixel phases. From the radiometrically calibrated—but not resampled—images, we extracted subimages containing just the bridge and the water on each side. Figure 14(a) shows an example of such an image. These pixel radiances we projected onto an axis perpendicular to the bridge. The result is an over-sampled profile of the apparent bridge radiance.

Next, we created a simple step-function model of the bridge radiance profile by clicking the mouse cursor at points on a plot of the observed profile. The parameters of the model are a set of radiances and the locations of the steps between them. We convolve this model with the line-spread function computed from the STF (including the integration-time factor) to produce the predicted instrument response. Our analysis program calls the CurveFIT procedure to adjust the model parameters to find the least-squares fit of the predicted response to the observed bridge profile. Figure 14(b) shows the predicted response for the example image.

This analysis has been applied to bridges scanned by each of the SCAs, and to bridges lying at a variety of angles with respect to the EO-1 ground track. For bridges with relatively simple structures, we find a good match between the predicted and observed bridge profiles, as shown in Figures 15 and 16.

Celestial Images

The ALI has made scans of the moon, stars, and planets. We have reconstructed images of these objects by projecting the image data to maps in celestial coordinates. This procedure is a simpler problem than making geo-referenced images. For star and planet scans, we read the attitude telemetry data and apply the de-

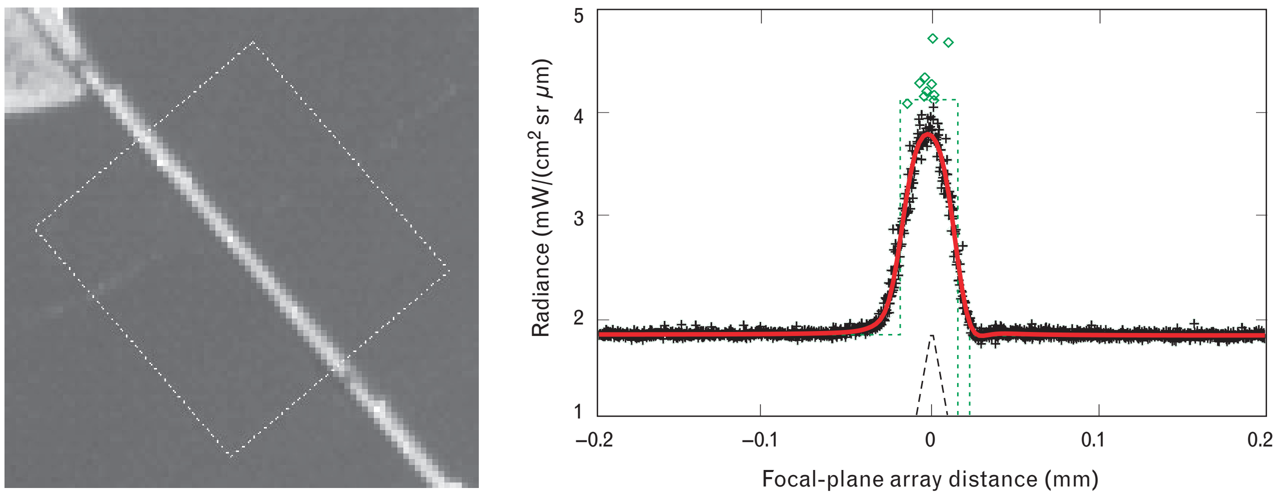

FIGURE 14. Rectangular region of interest (left) across the Bronx Whitestone bridge, New York City, in the panchromatic image acquired 20 March 2001. The plot (right) shows the fitted panchromatic radiance profile of the same bridge. Black crosses are the data points, the dotted green step function is the fitted radiance model (including shadow), and the red curve is the predicted response to the model. The green diamonds are data points omitted from the fit (probably vehicles on the bridge), and the dashed triangle shows the projected area of a single detector.

tector line-of-sight map to project the Level 1R data to celestial coordinates (right ascension and declination). Lunar scans also require us to use the changing positions of the spacecraft and the moon in order to project the data to a plane centered on the moon. For greatest accuracy, the velocity of the satellite relative to the speed of light should be taken into account, though this is a rather small effect. A larger problem is our uncertainty over the relative alignment of the ALI axes with the attitude control system of the spacecraft, and the exact time of the first frame of image data.

Lunar Scans

A regular part of the EO-1 operations is lunar calibration scans, conducted near the times of the full moon. The primary purpose of the lunar scans is radiometric calibration. These calibrations are valuable for cross-comparison of different instruments because the moon is an unchanging reference object, visible to all. We have also taken advantage of these scans to assess the spatial response of the ALI. The limb of the moon presents a sharp transition from the sunlit surface to the blackness of space. In a lunar calibration, the ALI scans slowly across the moon four times, so that the moon crosses each SCA once. The angular speed is chosen so that the speed of the moon's image on the focal plane is approximately one-eighth of the speed of the earth's image during a normal data collection.

Spatial analysis of the lunar scan begins once the image data are radiometrically calibrated. We read the telemetry files to find the position, velocity, and attitude of the spacecraft for every second of the lunar scan. The position of the moon comes from an IDL procedure MOONPOS.PRO, from the Astronomy User's Library at GSFC. We combine these data with the detector line-of-sight calibration data to determine where each detector was looking relative to the center of the moon during every frame.

Two things are then done. First, a radiance map of the moon, as shown in Figure 17, is created for each band, at a user-selected scale and resolution. Second, we write a file of data structures containing all of the detector pointing and radiance information for the part of the scan while the detector was within a narrow annulus containing the limb. We read the annulus data file to plot the observed radiance as a function of radius from the center of the moon, within a series of azimuth bins centered on the solar azimuth. In a manner similar to the bridge profile analysis, we fit a simple model of the limb radiance to the data. In this case, we model the radiance with a low-degree polynomial terminating at a sharp edge, with zero beyond.

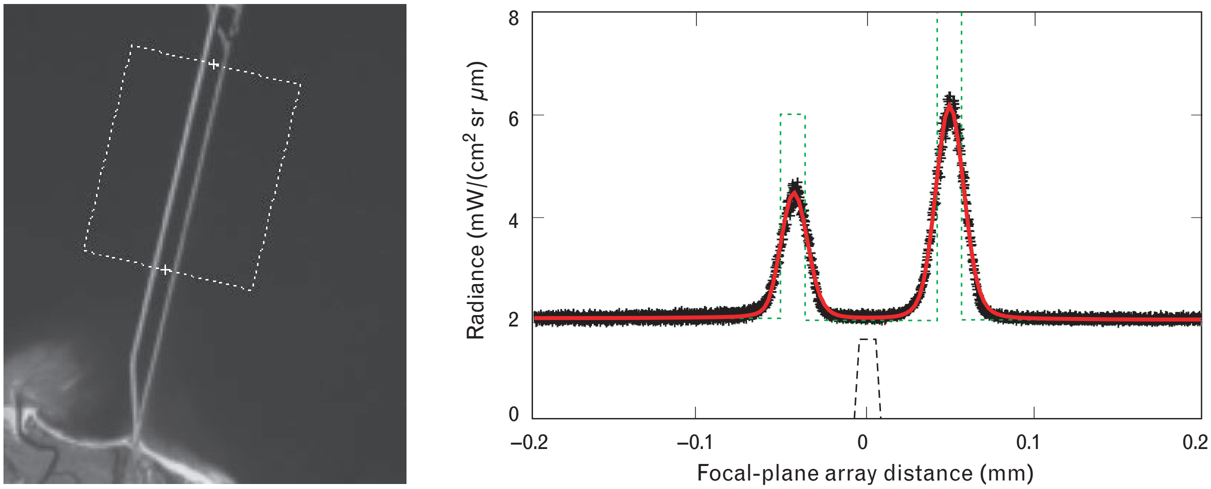

FIGURE 15. Panchromatic image and radiance profile of the Earle Naval Pier, Sandy Hook Bay, New Jersey, observed 20 March 2001.

We again compute the predicted ALI spatial response to this model from the calibrated STF.

The panchromatic-band edge response predicted from the calibrated STF corresponds very well with the observed response, within approximately 30 μrad of the edge, as shown in Figure 18. Farther out into dark space, however, we find an excess response, falling off much more gradually than the edge-response prediction. This response can be seen in the semi-log plots of radiance versus radius, as shown in Figure 19.

We knew that the mirror surfaces of the ALI scatter much more light than normal telescope mirrors. The calibrated STF model did not include this scattered-light contribution. In fact, the lunar scans provide the best opportunity to assess this effect at the system level. It would be difficult in the laboratory to produce an artificial scene with such a dark background next to a bright edge. The semi-log plots also reveal other interesting artifacts. Certain of the bands in various SCAs appear to show ghost images. There are differences between the stray radiances in the odd or even detectors of a single band. The ghosting is mostly confined to the SWIR bands. The worst case is 8% of the edge radiance in band 5p of SCA 1. Elsewhere, it is less than or equal to 4%. The leading hypothesis for the ghosts is that they result from scattering at the

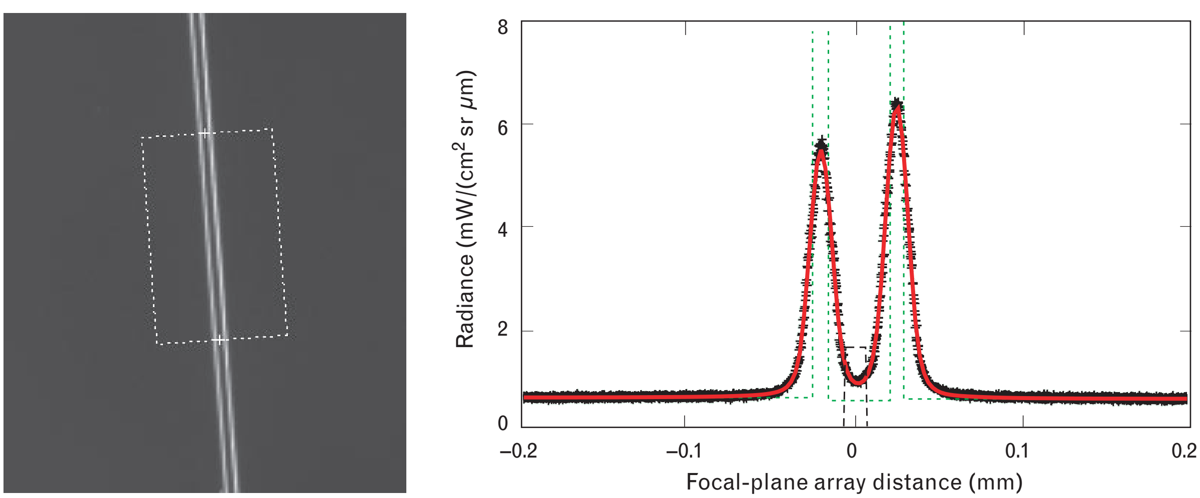

FIGURE 16. Panchromatic image and radiance profile of the Lake Pontchartrain causeway, near New Orleans.

FIGURE 17. Panchromatic image of the moon, from the lunar calibration of 26 February 2002, for SCA 2.

edge of the VNIR filters, which protrude above the SWIR filters. Except for the worst ghost artifacts, the stray-light effects are difficult to find in a normal earth scene because they are so small in comparison with scene radiance variations.

Stars and Planets

In March 2001, we obtained scans of Venus, Jupiter, Saturn, and the Pleiades. In May 2001, Venus was scanned again, as well as Vega. In all of these scans, the angular speed was approximately one-fourth of the apparent speed during a normal data collection. The images are oversampled in the scan direction, but not in the cross-scan direction. The stars appear as point sources to the ALI, so their images represent the PSF of the system. Vega appeared with a high signal-to-noise ratio. As for the lunar scans, we read the telemetry data to reconstruct the path of each detector across the sky. We plotted the apparent calibrated radiance versus scan angle, and compared that with the PSF computed from the STF, as shown in Figure 20. The agreement between the observed values and the expected values is excellent for the multispectral bands. The panchromatic data seem to indicate that the true PSF is a bit narrower than the calibrated STF suggests. Also, the small tail extending to the right

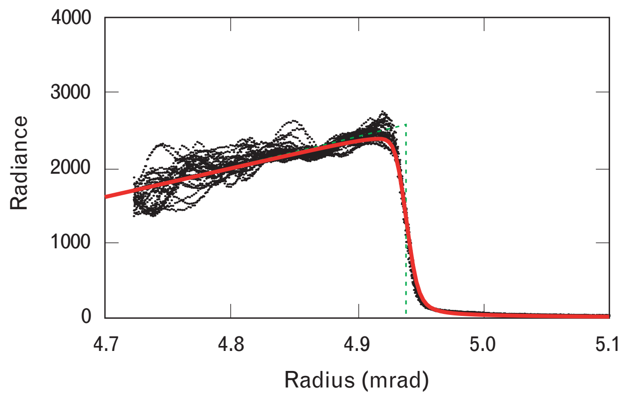

FIGURE 18. Lunar-limb radiance profile, from the scan of 26 February 2002, panchromatic band, for SCA 3, over 44° to 49° azimuth. The dotted green line is the fitted radiance model, and the red curve is the predicted response.

from the computed PSF is probably an artifact from the calibration process.

The Pleiades scan occurred shortly after the scans of Jupiter and Saturn. An early attempt to form an image of the star cluster was disappointing, since only a few pixels were not simply black. Then we analyzed the radiometrically corrected data detector by detector. We formed a matched filter, using the STF, and convolved it with the series of readings from each detector. We then correlated peaks significantly above background

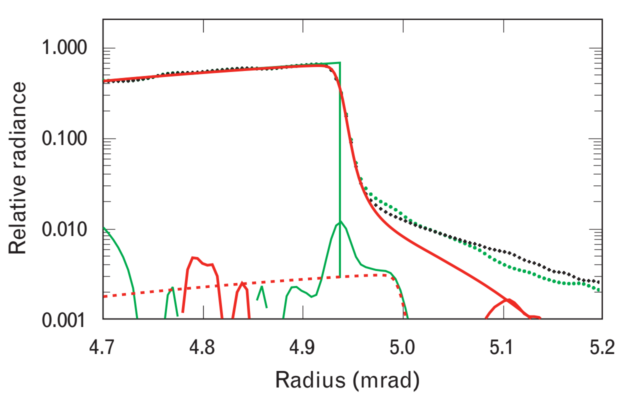

FIGURE 19. Lunar limb radiance profile in SCA 3, panchromatic band over 44° to 49° azimuth, as in Figure 18. The black dots are the average normalized radiances for the odd-numbered detectors, and the green dots are the radiances for the even-numbered detectors. The solid curves at the bottom show the even-minus-odd radiances (green) and the odd-minus-even radiances (red). The dashed red curve is a scaled and shifted copy of the predicted edge response, fitted to the excess signal of the even detectors.

with those from neighboring detectors to identify bright objects. We fitted the positions (right ascension and declination) and intensities of the bright objects in each band, using the PSF and the telemetry data. Ten bright objects were detected in this way, one appearing in both SCAs. From their spatial pattern and approximate positions, we identified the detected objects with specific stars. Next we compared the panchromatic intensities of the objects with the V magnitudes of the stars. We were able to correlate the two magnitudes rather well, as

$$m_{pan} = -2.5\log_{10}\left(\frac{I_{pan}}{2400}\right),$$

where I_{pan} is the fitted panchromatic-band intensity in scaled radiometric units, and m_{pan} is a pseudo-magnitude comparable to a stellar V magnitude. The stars detected in the Pleiades range from V = 6.17 to V = 2.90. The m_{pan} magnitudes generally differ from the V magnitudes by less than 0.1 magnitude. This observation was an interesting check on the radiometric sensitivity of the ALI. Looking again at the Vega scan, we find m_{pan} = 0.15, while the V magnitude of this A0 star is 0.03. This discrepancy is probably attributable to the difference between the panchromatic spectral band and the pass band defined for V magnitudes.

For the planet scans, we projected the radiance data onto maps in celestial coordinates centered on each planet, as we did for the moon. Figure 21 shows the resulting images for the panchromatic band. While the detail in these images is no better than what can be seen through a small telescope, they are at least clear and free of artifacts. The optical resolution is all that it was intended to be for imaging the Earth.

Panchromatic-Sharpened Color Images

The panchromatic images show much more detail than the multispectral images, since they are sampled at one-third of the multispectral GSD. They are not as visually appealing, however, since they are monochromatic. We succeeded in combining the spatial resolution of the panchromatic images with the color information of the multispectral images to produce panchromatic-sharpened color images of selected regions, as follows. We read the Level 1R data for bands 3, 2, and 1, which span the spectral response of the

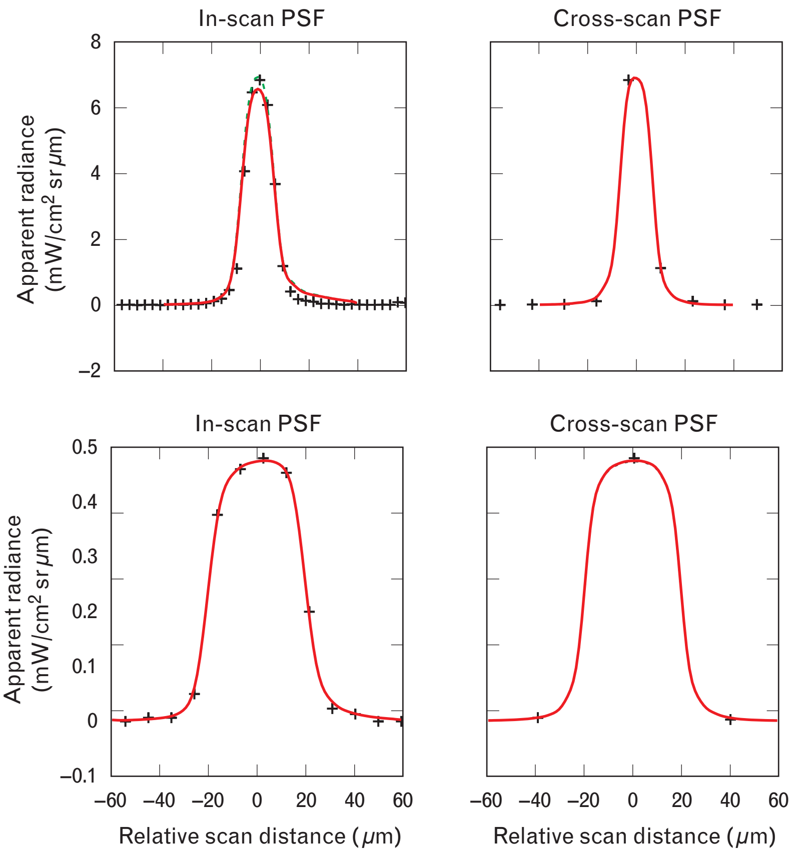

FIGURE 20. Radiance profile of Vega spatial response in panchromatic band (top) and band 4 (bottom). Since the star is a true point source, this is the best test of the point-spread function (PSF) of the ALI. Black crosses are the observed signals, and red curves are the predicted response.

panchromatic band. We apply whole-frame shifts to bring these bands into approximate registration with each other and with the panchromatic band. Next, we up-sample the multispectral bands to the same 10 m resolution as the panchromatic data. By application of a Sobel filter, we find the edge features in each band. We shift each multispectral band by panchromatic pixel increments to achieve the highest cross-correlation of the edge images in the panchromatic and multispectral bands. This method of aligning the bands is not sensitive to contrast reversals between bands.

From the three registered multispectral bands at 10 m resolution, we form an RGB image array, exactly as we do for browse image creation. Using the COLOR_CONVERT procedure in IDL, we now convert the array to the hue, saturation, and value (HSV) color space. We replace the value array with one formed from the panchromatic image. Once we convert back to the RGB color space, we have an array with the spatial resolution of the panchromatic data, but with colors derived from the multispectral bands. These data are written to a file in TIFF format, rather than JPEG, to preserve the full resolution and avoid lossy JPEG compression artifacts. Examples of the images obtained by this process can be seen in Figures 22 and 23, and elsewhere in this issue [9].

Imaging Performance Summary

The Advanced Land Imager was designed to provide finer resolution than Landsat 7 (10 m panchromatic versus 15 m for Landsat), along with equivalent or

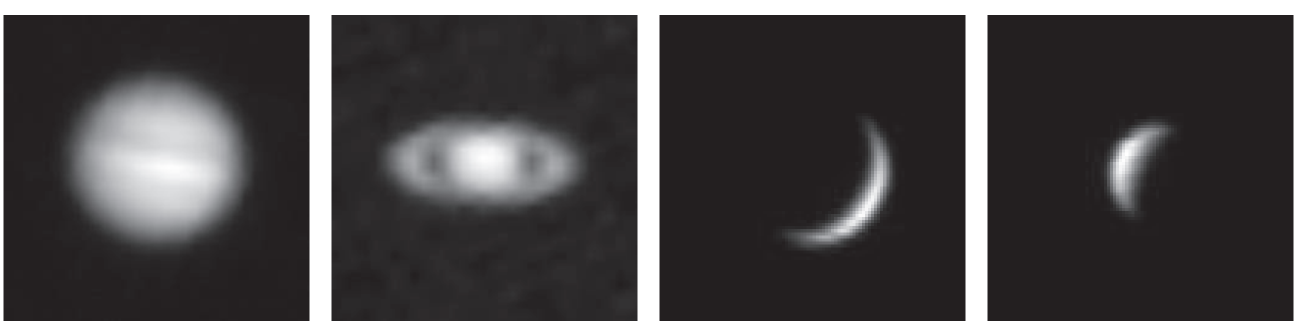

FIGURE 21. Composite panchromatic image of planets scanned by the ALI. Jupiter and Saturn (at left) were scanned on 14 March 2001. Venus was scanned on 7 March 2001 and again on 15 May 2001 (at right).

FIGURE 22. Panchromatic-sharpened color image of part of Yuma, Arizona, taken 23 January 2001.

better multispectral imaging. We thoroughly calibrated its spatial resolution and detector lines of sight under simulated on-orbit environmental conditions. While the reconstruction of images acquired by the ALI has been complicated by the sparse arrangement of the detectors on the focal plane, we have overcome those difficulties through careful design and calibration. Our analysis of images taken on-orbit shows that the images returned are consistent with the system transfer functions and lines of sight derived from the calibrations performed prior to integration with the spacecraft.

Some subtle features of the spatial response that would have been difficult to detect in the laboratory have been revealed through the lunar limb scans. Apart from a few weak optical ghosts, the ALI is essentially

FIGURE 23. Panchromatic-sharpened color image of Kabul, Afghanistan, taken 26 September 2001.

free of image artifacts. The quality of the images fully validates its design as a prototype for future Landsat imagers.

Acknowledgments

The author wishes to thank the many people—too numerous to list individually—who have contributed to the success of the Advanced Land Imager. The calibration work relied heavily on the contributions of Bert Willard, Patrick Quinn, Herbert Viggh, Ronald Sayer, Scott Stuart, Herbert Feinstein, and Travis Hein. I thank David Nathanson for assistance in determining the thermal effects in the vacuum window. Alex Parker supported the optical testing at Goddard Space Flight Center. From SSG, Gerry Perron and Bryan Norman supplied optical subsystem measurements and data, and Marcelo Benech assisted in the theodolite measurements. Paul Bryant and Cliff Nichols of Santa Barbara Remote Sensing provided focal-plane measurement data. Particularly helpful during the on-orbit operations were Jenifer Evans and Margaret Gibbs at Lincoln Laboratory, and Steve Ungar, Lawrence Ong, and Jim Storey at Goddard Space Flight Center. John Kerekes of Lincoln Laboratory supplied MODTRAN calculations of expected radiance levels. I especially thank Jeffrey Mendenhall for his work on system calibrations, and helpful discussions during the analysis of the data. The support and advice of Donald Lencioni and Costas Digenis were helpful throughout the project.

REFERENCES

1. D.E. Lencioni, D.R. Hearn, C.J. Digenis, J,A. Mendenhall, and W.E. Bicknell, "The EO-1 Advanced Land Imager: An Overview," *Linc. Lab. J.*, in this issue.
2. G.D. Boreman, *Modulation Transfer Function in Optical and Electro-Optical Systems* (SPIE Optical Engineering Press, Bellingham, Wash., 2001).
3. B.C. Willard, "Wide Field-of-View Schmidt-Sphere Imaging Collimator," *SPIE* **3750,** 1999, pp. 286–296.
4. D.R. Hearn, "Vacuum Window Optical Power Induced by Temperature Gradients," *SPIE* **3750,** 1999, pp. 297–308.
5. H.E.M. Viggh, J.A. Mendenhall, R.W. Sayer, and J.S. Stuart, "An Automated Ground Data Acquisition and Processing System for the Advanced Land Imager," *Linc. Lab. J.*, in this issue.
6. D.R. Hearn, "EO-1 Advanced Land Imager Modulation Transfer Functions," Technical Report 1061, Lincoln Laboratory (22 Mar. 2000), DTIC #ADB-252266.
7. D.R. Hearn, "Earth Observing-1 Advanced Land Imager: Detector Line-of-Sight Calibration," Project Report EO-1-4, Lincoln Laboratory (29 Dec. 2000), DTIC #ADA-410322.
8. J.A. Mendenhall, D.E. Lencioni, and J. Evans, "Spectral and Radiometric Calibration of the Advanced Land Imager," *Linc. Lab. J.*, in this issue.
9. D.R. Hearn, "A Gallery of Images from the Advanced Land Imager," *Linc. Lab. J.*, in this issue.
10. R.A. Schowengerdt, *Remote Sensing: Models and Methods for Image Processing*, 2nd ed. (Academic Press, San Diego, 1997).
11. D.R. Hearn, "Earth Observing-1 Advanced Land Imager: Spatial Imaging Performance On-Orbit," Project Report EO-1-12, Lincoln Laboratory (31 Dec. 2002).

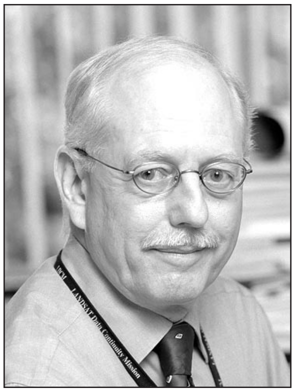

DAVID R. HEARN
is a staff member in the Advanced Space Systems and Concepts Group. He received a B.S. degree in physics from the California Institute of Technology in 1964, and A.M. and Ph.D. degrees in physics from Harvard University. At the Harvard-Smithsonian Center for Astrophysics, he did thesis and post-doctoral research in gamma-ray astronomy. In 1970, he joined the Center for Space Research at MIT, where he developed the low-energy x-ray sensor for the Small Astronomy Satellite 3 (SAS-3). With that instrument, he made several discoveries, such as the soft x-ray binary AM Herculis. In 1979 he developed a novel x-ray CT scanner at Elscint, Inc. He joined Lincoln Laboratory in 1984, where his work initially dealt with adaptive optics for high-energy laser systems. Since that time, he has primarily worked on the development and analysis of optical remote-sensing systems. Those sensors, both ground-based and space-based, have operated at wavelengths from the visible to the very long-wave infrared. He was involved from the outset in the development of the Advanced Land Imager for EO-1, under the NASA New Millennium Program. This work included performance analyses, instrument calibrations, and analysis and presentation of the on-orbit results.