

LASER NASA Grant to Columbia University, Astrophysics Laboratory
New Insight from Old Images: Reanalysis of Lunar Orbiter Photographs
NASA Proposal Number 07-LASER07-0005

Analysis and Correction of Artifacts in Lunar Orbiter Photographs

Charles J. Byrne
January 19, 2011

Introduction

In 1966 and 1967, five Lunar Orbiter Missions surveyed nearly the entire Moon. The target sites included high-resolution coverage at 1 meter to select and certify Apollo landing sites and other sites of scientific interest, comprehensive high-resolution coverage at 200 meters or better to survey the entire front side of the Moon, and essentially complete coverage, at lower resolution, of the far side of the Moon.

This coverage is the oldest comprehensive survey of its kind of any solar system planet other than Earth. Together with the Apollo coverage of the equatorial region of the Moon, it will always be the earliest available coverage of any rocky body in our solar system. The Moon, like Mercury, forms a witness plate of the last four billion years of the events of the inner solar system. Its history continues with both natural and man-made events. The degree of modification in the last 40 years and in the indefinite future can be judged by comparison with these photographs.

This is the first of three reports. It covers the algorithms and programs that were prepared for the processing of the photographs to achieve the best possible product for geologic analysis.

The second report will cover the results of the production stage.

The third will be an analysis by Dr. Arlin Crotts of Columbia on the relevance to the search for Transient Lunar Phenomena.

The Lunar Orbiter camera systems

There were two cameras in each spacecraft. A medium resolution camera had a lens with a 3.2 inch focal length and the high resolution camera had a lens with a 24 inch focal length. Each camera had a platen that was moved during an exposure to compensate for image motion.

The photographs were originally exposed on 70 mm fine-grained film (Eastman Kodak SO-243). This was developed in the spacecraft by a Kodak Bimat process.

A CRT scanner illuminated the film in a 0.1 inch (2.5 mm) wide scanning line which was swept across the film by a mechanically driven mirror. Scan lines were separated by 2.3 microns. The reflected spot from the scanner was about 5 microns in diameter on the spacecraft film. Phosphor persistence may have made its effective diameter a little larger in the direction of the scan. Light transmitted through the negative was measured by a photomultiplier tube and transmitted to Earth in analog form. The unit of image in a single scan across the 70 mm width of the film is called a framelet.

The transmitted information was recorded in two ways, both analog. The primary record was produced on 35 mm film by the Ground Recording Equipment (GRE). The GRE demodulated the suppressed-carrier vestigial sideband signal, recovered synchronization, clamped the black level, did some clipping and gain control, and drove a CRT scanner. The film was moved by the scanner.

Primary framelets from the GRE have not been located, but the better copies have been selected by USGS, rescanned at 25 microns (corresponding to 3.2 microns in the spacecraft tape), and digitized at 8 bit (256 shade) precision. The framelets were then reassembled into medium resolution frames and high resolution subframes. The primary purpose of the USGS effort was to provide cartographic information to the Unified Lunar Control Network – 2005. Additional Very High Resolution framelets were scanned, digitized, and assembled into frames for a number of target sites.

The second recording method was by a rotating-head magnetic tape drive. The intermediate signal was recorded on the tape, leaving the suppressed-carrier vestigial sideband signal modulated. The Lunar Orbiter Image Reconstruction Project at NASA's Ames Research Center has obtained primary tapes from the Jet Propulsion Laboratory, restored appropriate tape drives, and rebuilt the demodulation equipment. They are digitizing at a sampling interval equivalent to 1.6 microns in the spacecraft and 16-bit precision (65536 shades). Their method avoids the non-linearity of the GRE equipment and is essentially a linear record of the negative film in the spacecraft.

Until recently, the best widely available Lunar Orbiter images have been in a hard copy format prepared by Kodak (for Mission 1) and by Langley Research Center and the Army Map Service (subsequent missions). These were processed from the original GRE framelets or from new framelets produced at the Langley Research Center by playing the magnetic tapes into GRE equipment. These framelets were assembled into frames and subframes by the Army Map Service. A number of copies of the assembled images (each about two feet by two feet) were distributed to the Deep Space Image Facilities and have been available for study..

One of the DSIF centers is at the Lunar and Planetary Institute. Under the direction of Jeff Gillis (now Jeff Gillis-Davis) the LPI staff photographed a set of the hard copy images there (selected for comprehensive coverage), using a digital camera. These, the first large set of digitized images, were posted on the LPI web page. Because of camera imitations, the images were at reduced resolution; each frame or subframe was about 1000 by 700 pixels, while the highest quality USGS and LOIRP data covers about 23,000 by 16,000 pixels. Now that digitized images were available, modern computers could be used to analyze the artifacts and remove them. The

LPI images were cleaned and published in two books, “The Lunar Orbiter Photographic Atlas of the Near Side of the Moon”, and “The Far Side of the Moon. A Photographic Guide” each by Charles J. Byrne and each published by Springer.

In the past year, the author has been privileged to view new digitized products from the USGS and the LOIRP project. Each of these products closely approximates the inherent resolution of the Lunar Orbiter images. The LOIRP project has recovered a greater dynamic range of shades, some of which have been lost in the film framelets scanned and digitized by USGS.

The work of the current grant has been focused on the USGS product, which is currently more available. Reference will be made to the LOIRP product where it clarifies the nature of the artifacts, especially the synchronization structure.

Artifacts

The available records scanned and digitized by USGS have passed through at least two film processes (spacecraft and GRE), at least three optical processes (spacecraft CRT, photomultiplier, GRE CRT), electronic synchronization, modulation, demodulation, synch recovery, clamping, nonlinear processes, and the mechanical motion of scanning across the film in the spacecraft. All processes were in the state of the art as of the early 1960s. The most mature art was that of the spacecraft film itself, with a granularity that justified examination at 3.2 microns and a film width of 70 mm.

Although many artifacts were introduced, there is sufficient calibration data to reduce many of them and approach the image on the spacecraft film. That is the objective of this study. Further steps are desirable to refer the image to the brightness on the lunar surface, but that step is beyond the scope of this study.

The artifacts are of two types: those that affect the position (sample coordinates) of data in the digitized products and those that affect the brightness (sample shade value) of the digitized product. They are reduced in two programs, Framelet and Assembly. Each program deals with parts of each type of artifact.

Sample position reduced by the Framelet program

The USGS reassembled its frames from framelets by aligning pre-exposed reseau crosses. This approach is valid for cartographic purposes, but because of positional artifacts within the framelets does not produce a good product for geologic interpretation. The following paragraphs discuss types of positional artifacts, their sources, and the remedial techniques. In these descriptions, the X axis is across the framelets (USGS framelets are a uniform 970 samples wide). The Y axis is along the framelets (USGS framelets are a uniform 16,550 framelets long).

The first artifact to be discussed and reduced in the Framelet program is Y skew. This has two components. First, the framelets are randomly spaced within the sample space. Second, the scan lines are not at right angles to the image length. The first effect is due to variations in cutting the films to be scanned and aligning them on the scanning table. The second is mostly or all due to

the vertical component of the CRT scan line and the alignment of the spacecraft CRT with the axis of the mechanical scanning mirror. It varies from mission to mission, but is always in the same direction and is greatest in Mission 5. Y skew has a small curvature component.

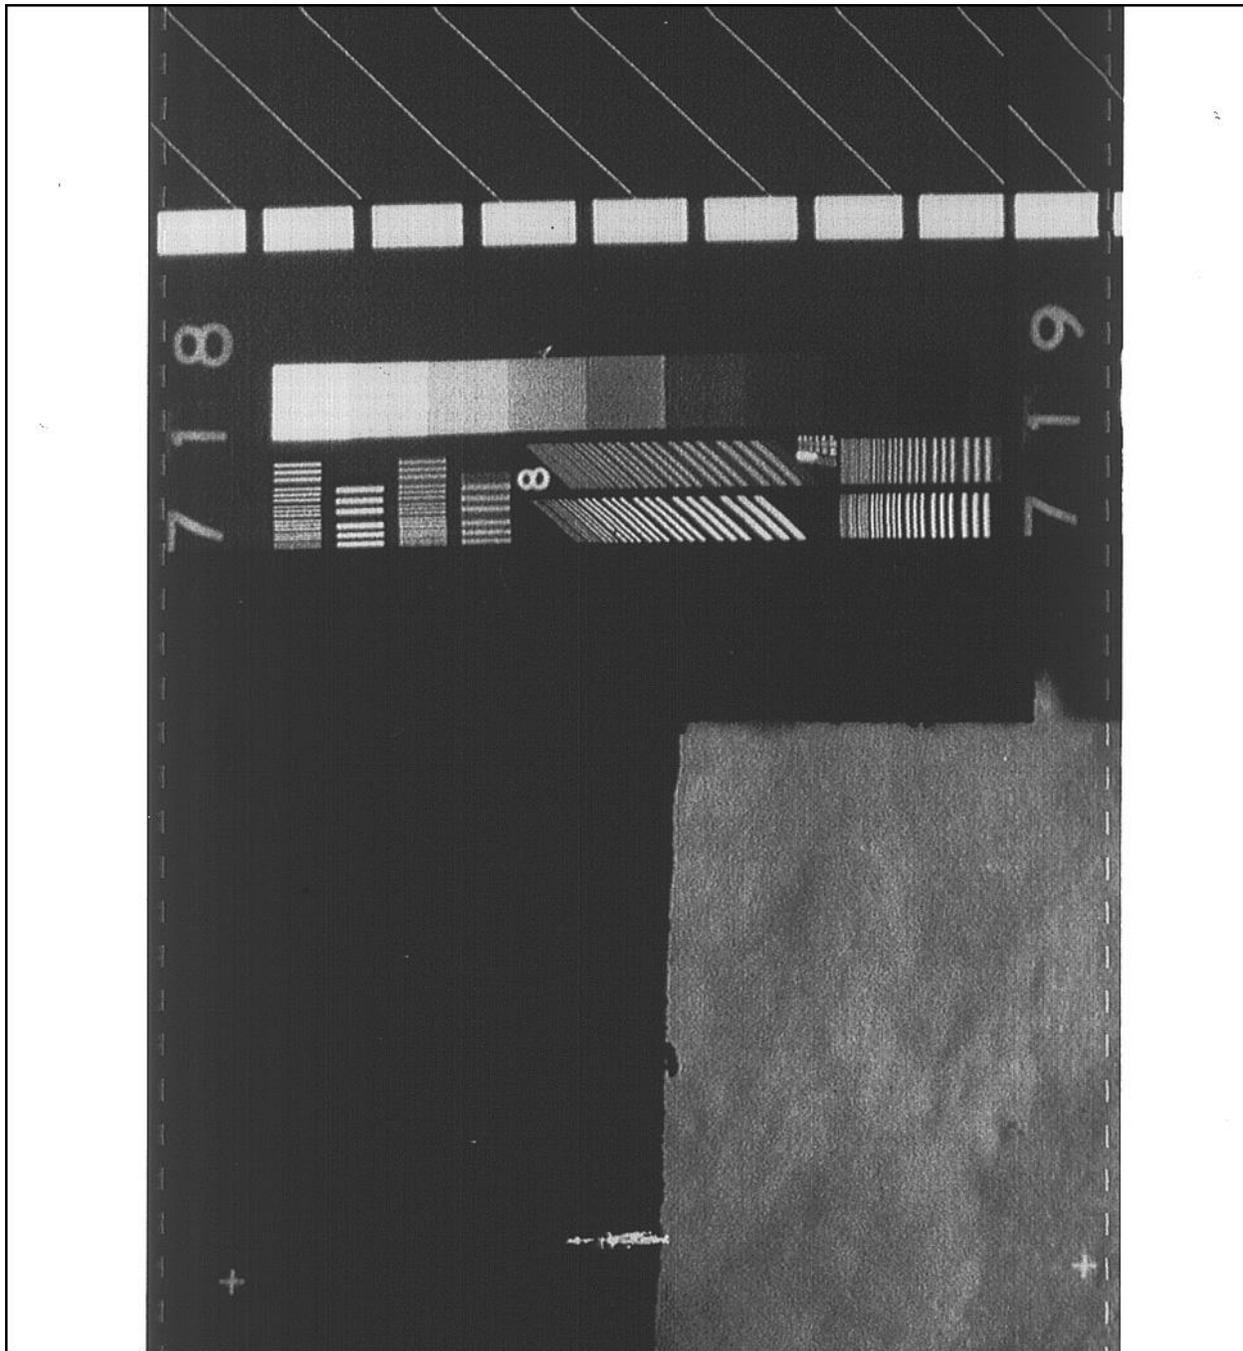


Figure 1: This is an end framelet of LO5-208H, showing the calibration strip and a frame edge. Scanning was by USGS at 25 microns, corresponding to 3.2 microns in the spacecraft film. Each framelet is 970 pixels by 16550 pixels. The white areas on either side are synch pulses inserted electronically at the spacecraft at the black level of the spacecraft negative (white on this positive image). The brightness has been increased in this image.

Figure 1 shows the Y skew of the calibration strip, which is uniformly parallel to the edge of the spacecraft film. The amount of the Y skew is measured by determining the edge of the linearization pattern (white rectangles) at about 100 locations across the framelet. A second-order Taylor series is fit to the measurements. The new image is set so that this second order curve is made a straight horizontal line as it was on the spacecraft film. The top of the framelet is trimmed to a sample number of 200 from the top of the linearization pattern. The rest of the image is derived from this. The offset is to a fraction of a pixel. The brightness of the each new pixels is proportionately calculated from the two old pixels above and below the precise value of the offset. Figure 2 shows the image of the calibration strip after Y skew is reduced.

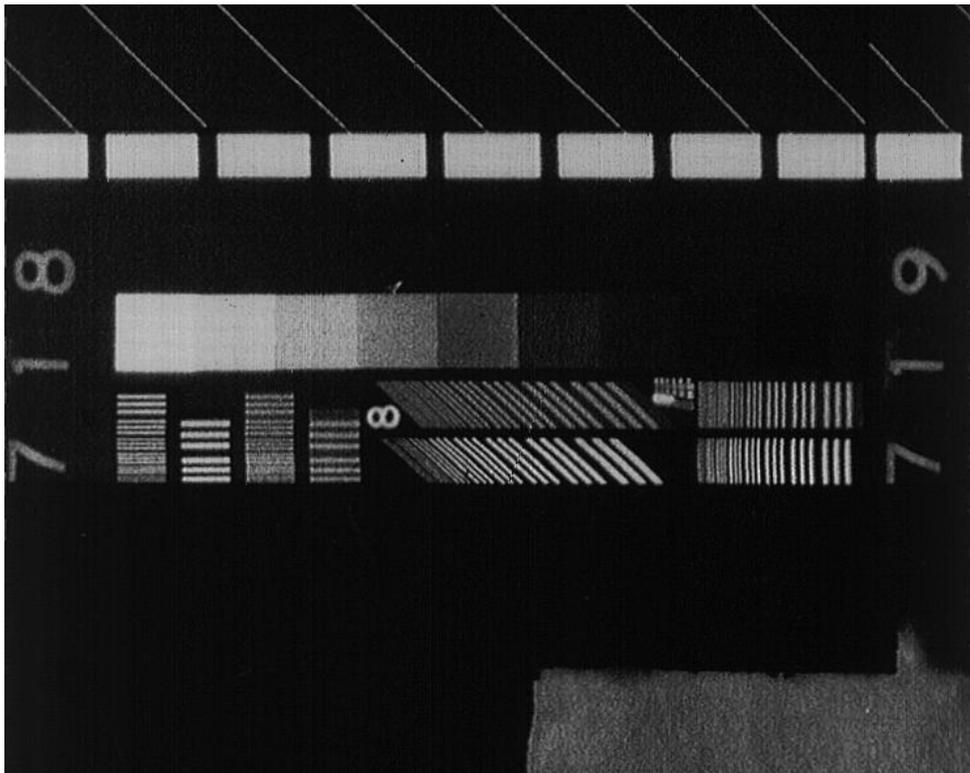


Figure 2: This shows the result of correction of Y skew. The linearization pattern is now straight and horizontal as it was on the spacecraft film. The top of the framelet has been trimmed. The offset (as a function of width across the framelet) has been corrected for the entire framelet.

X skew is the next artifact that is reduced in the Framelet program. The largest skew is drift from left to right of the image over the framelet length. This amounts to as much as 30 pixels along the length of a framelet. In addition to this major skew is variation of the width of the image between synch pulses. There is no blacker-than-black level on the synch pulses, as in a standard TV signal of the time, so the image brightness actually influences the edges of the image on the GRE film.

A better measure of width is the thin dashes, called fiducial dashes a few pixels inside each vertical edge of the images (see Figure 1). Examination of magnetic tape images from LOIRP show that these come from high frequency spikes that are electronically inserted in the signal by

a counter circuit closely related to the spacecraft CRT scan frequency. Although these dashes were also intended to show the edge of overlap and were in fact used as cut lines in the original frame assembly process, they are not precisely aligned with overlap. In the case of Mission 5, image would be lost if they would in fact be used as cut lines.

To eliminate X skew, the dash lines are identified by scanning with a kernel pattern and looking for the point of maximum correlation for each dash. As with the linearity pattern, the dash positions (about 460 on each side of a framelet) are approximated with a second order Taylor series.

Experimental assembly of the framelets revealed the dash lines were too close to permit being trimmed away. To preserve continuity of the image, it would be necessary to leave at least one dash line in the assembled image. Ultimately, it was decided to balance the images, leaving both dash lines within the adjusted image. This implied suppression of both sets of dashes to avoid distraction in examining the assembled frames. The new values of the dash lines were chosen to be 86 and 829 pixels from the left edge of the framelets, typical values. These values will be revised in the future to correct a 4% excess width in the aspect ratio. The images outside the dash lines were adjusted proportionately to the images between the dashed lines. The white synch pulse areas were trimmed or padded as needed to maintain the framelet width at 970 pixels.

To avoid the distraction of the dashes when examining the frames, they were removed by interpolation from adjacent parts of the images, becoming nearly invisible even at high resolution. A similar process has been applied by USGS.

After the Y skew and X skew are adjusted and the dash lines removed, the framelets are ready to be trimmed and assembled by the Assembly program. However, it was decided to adjust the major problems of variations in sample shade values in the Framelet program by filtering and to prepare the way for further adjustments in the Assembly program.

Logging of X and Y skew

At the suggestion of Arlin Crotts, the coordinates of the Taylor series for the linearity pattern and the left and right rows of dashes are written out as a text file. In a production run, this file is overwritten as each file is finished, so the record will be preserved in case the production program halts.

Sample shade value filtered by the Framelet program

Examination of the raw image (Figure 3) shows a strong cross-hatch pattern from horizontal and vertical streaks. The vertical streaks result partially from granularity in the spacecraft CRT and partially from the GRE CRT. The horizontal streaks are partially from the two scanning patterns (one by the spacecraft and GRE and one by the USGS scan, which could not be synchronized with the first).

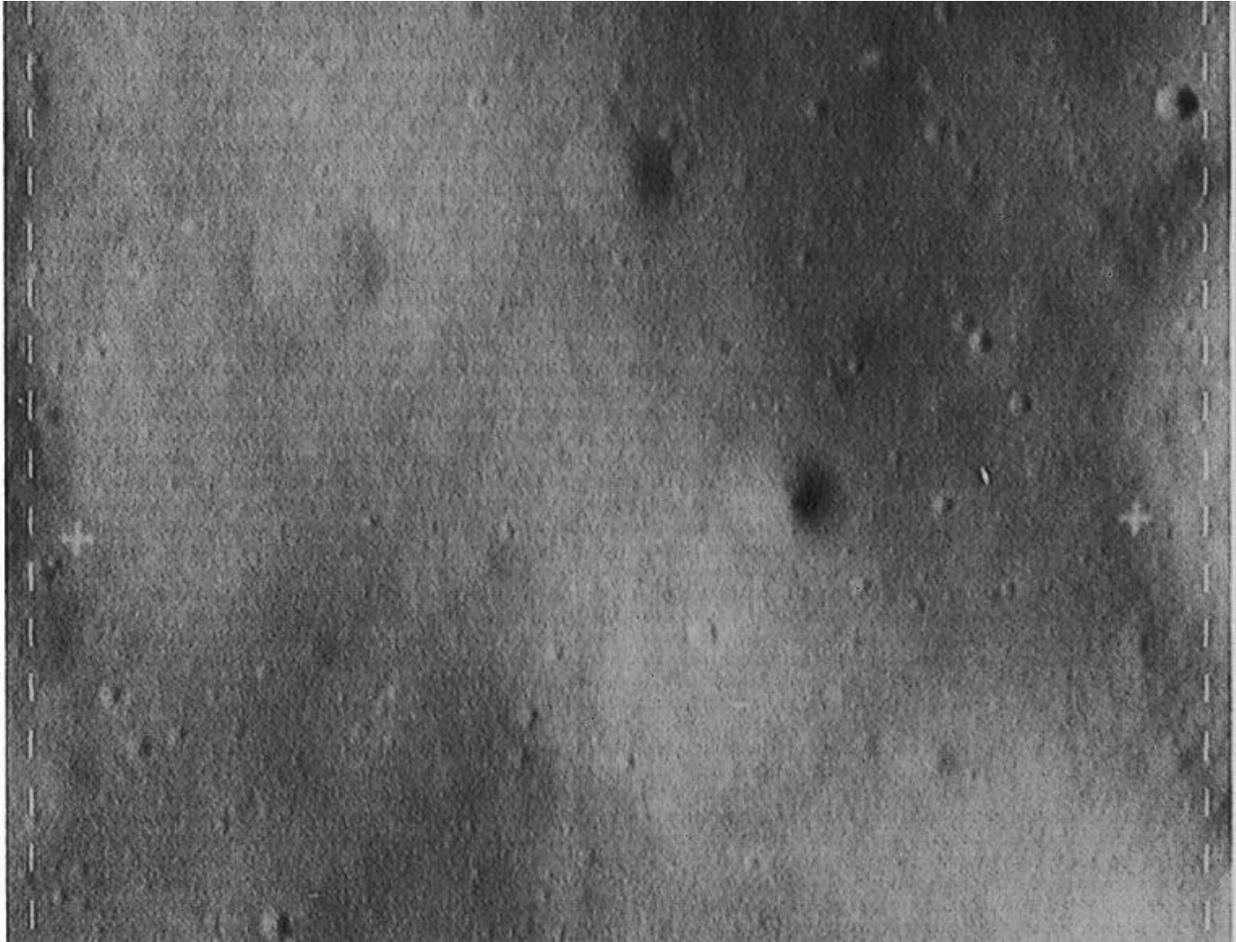


Figure 3: This portion of framelet 015x of LO5-202H1 shows both horizontal and vertical streaking from the repeated scans of the image.

A two-dimension Fourier transform was performed to examine the spectrum of the image. This transform was performed by the fast Fourier transform algorithm (also called the butterfly algorithm) which assumes a sample size that is a factor of two. The sample size for the width of the framelet was taken as 1024, with the extra width taken to be at shade 0. The sample size for the length of each framelet was taken to be 16,384 which is less than the full length, but sufficient to cover the entire image, including the pre-exposed calibration strip and a short distance beyond the lower frame edge. The resulting Fourier transform provides real and imaginary fields each 2048 by 32,768 frequency samples.

The magnitude of the transform field was normalized in frequency, producing v and u axes for display (see Figure 4). The horizontal streaks can be seen as bright areas in the vicinity of the v axis (low frequency in the u direction, high frequencies in the v direction) and the vertical streaks as bright areas along the u axis (low frequency in the v direction, high frequencies in the u direction).

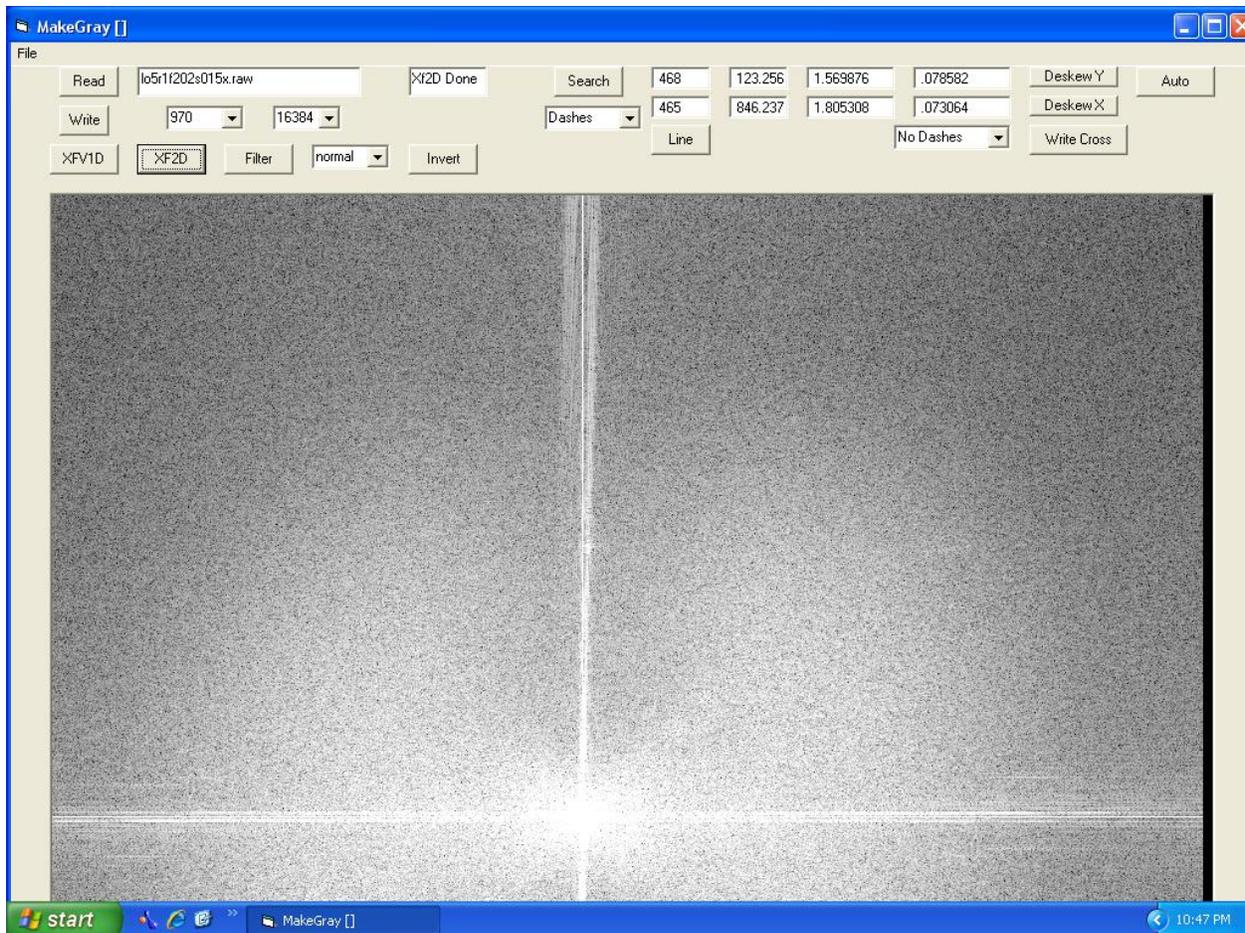


Figure 4: This is a screen capture of the two dimensional Fourier transform of the framelet partially shown in Figure 3. The u coordinate is horizontal and the v coordinate is vertical. The positive quadrant is shown in full. The control and reporting structure of the Framelet program is shown above the transform.

Most of the rest of the broad bright areas are the signal from the interaction of slanting sunlight on the topography, with a contribution from variation in albedo variations as well. There is little signal in the vicinity of the v axis because the sun direction is nearly east to west. Because of the Moon's unique photometric function, slope variations at right angles to the plane of illumination do not create brightness variations. What brightness that is near the v axis may be due to albedo variations, but is mostly film granularity and transmission noise.

The strongest signal, at an angle of about 45° from the origin, is due to craters, which have a constant ratio of vertical and horizontal signal. The smaller the craters, the further from the origin are their signals.

At high frequencies in the u direction, the signal falls off quite uniformly as a function of u , independently of the value of v , leaving noise of similar quality and amplitude to that along the v axis. This unusual pattern can only be due to a filtering operation on the X signal, the signal across the framelets. This filter must have been applied before the addition of most of the noise. It could have been done in the spacecraft to limit the bandwidth of the transmitted signal or in the

GRE. It could also have been due to phosphor persistence in the spacecraft CRT, which would widen the effective scanning spot size.

Three filters are applied to the Fourier transform to reduce noise (see Figure 5). Each transform is based on the low pass attenuation value $1/(1 + (\omega/\omega_0)^2)$ and the high pass attenuation value $(\omega/\omega_0)^2 / (1 + (\omega/\omega_0)^2)$. A low pass filter is applied to the u axis to attenuate the noise, there apparently being very little signal there. A second filter is composed of a low pass filter in v and a high pass filter in u to remove the spectral signature of horizontal streaks. A third filter is composed of a low pass filter in u and a high pass filter in v to remove the spectral signature of vertical streaks.

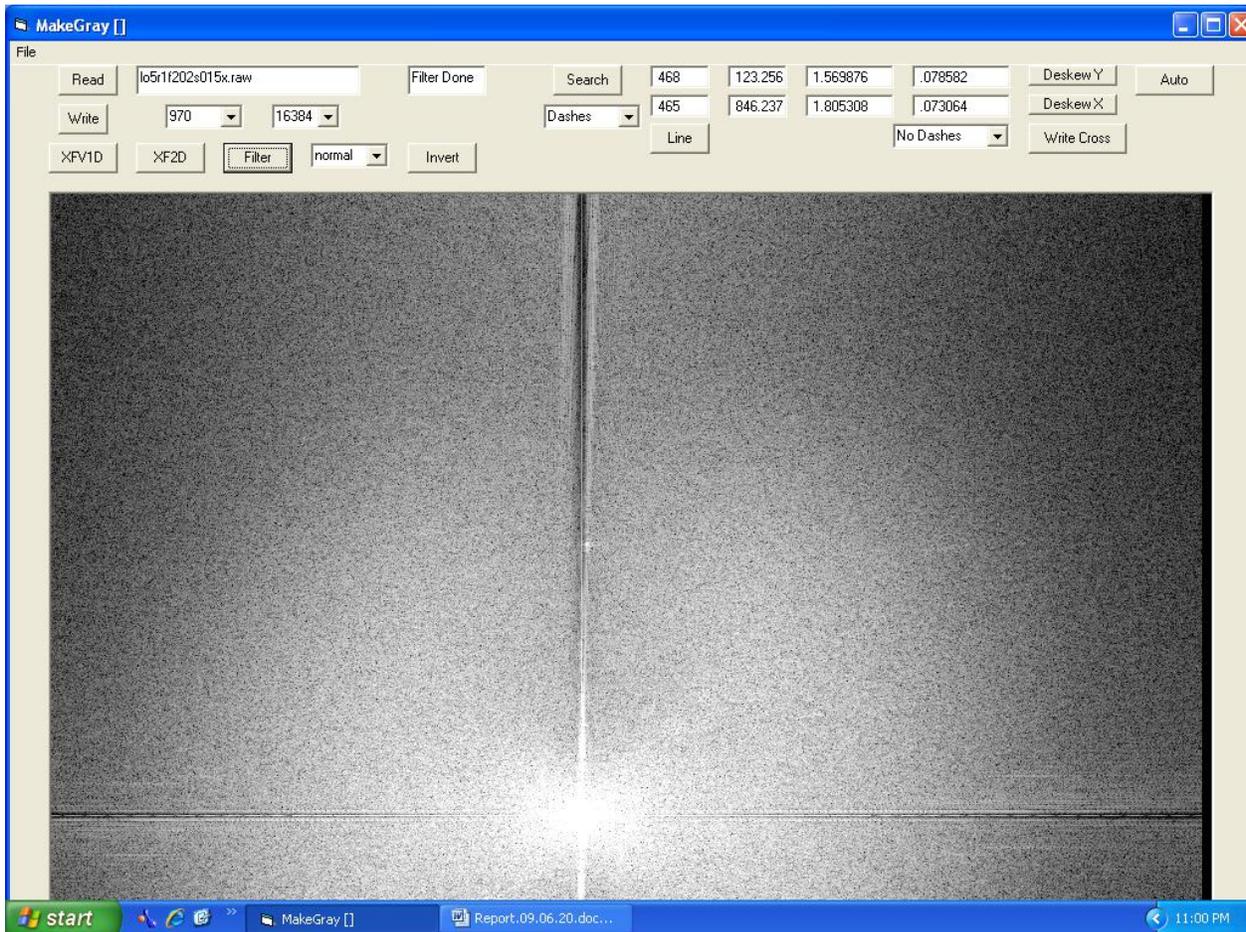


Figure 5: This screen capture shows the three filters applied to the transform. The one along the vertical v axis suppresses horizontal streaks and the one along the horizontal u axis suppresses vertical streaks. Another filter suppresses the high frequency noise in the region of the u axis beyond the signal there, which has been suppressed by a filter in the spacecraft or GRE, as described in the text. The filter parameters have been chosen as a compromise between noise suppression and reduction of the signal.

The characteristic gain characteristics of these filters are those of second order Butterworth filters, but unlike Butterworth filters, there is no phase shift because these filters are not limited by the requirement that they be implemented in passive networks.

The three filters were designed with five values of ω_0 determined by trial and error to remove noise and preserve signal.

The resulting real and imaginary components of the transform are again transformed in two dimensions to recover the framelet image (Figure 6). The filters are designed to preserve resolution, essentially remove the crosshatching, and reduce the random noise somewhat. The resulting images have the quality of high quality film photography even examined at high magnification, as long as the pixels are not resolved. The main remaining artifact is high frequency noise, which is much easier for the eye to ignore than crosshatching. It is now much easier to see subtle patterns of brightness due to low slopes and albedo variations.

After filtering, the Framelet program writes out the processed framelets in the same format as the input framelets were received from USGS, raw format (.raw suffix) with a width of 970 pixels and a length of 16550 pixels.



Figure 6: The filtered image is recovered by inverting the two-dimensional Fourier transform. Subtle shade variations due to minor slope variations are much clearer and shadow edges remain sharp as in the crater in the upper right corner. The reseau marks are clear. Some traces of horizontal streaks are still visible in the brighter areas, a compromise with signal suppression. The fiducial dashes have been suppressed in this image and the width adjusted.

Figure 7 shows the filtered pre-exposed calibration strip. Although there is some horizontal and vertical blurring, an artifact introduced by the horizontal filter in the spacecraft or GRE have been removed. The sharpness of the framelet image in Figures 1 and 2 had been artificially enhanced by a filter (in either the spacecraft or the GRE) which produces overshoot on the right edge of the rectangles of the linearity and gray scale patterns. This artificial sharpness makes the calibration strip look better, but introduces possible errors in the lunar image.

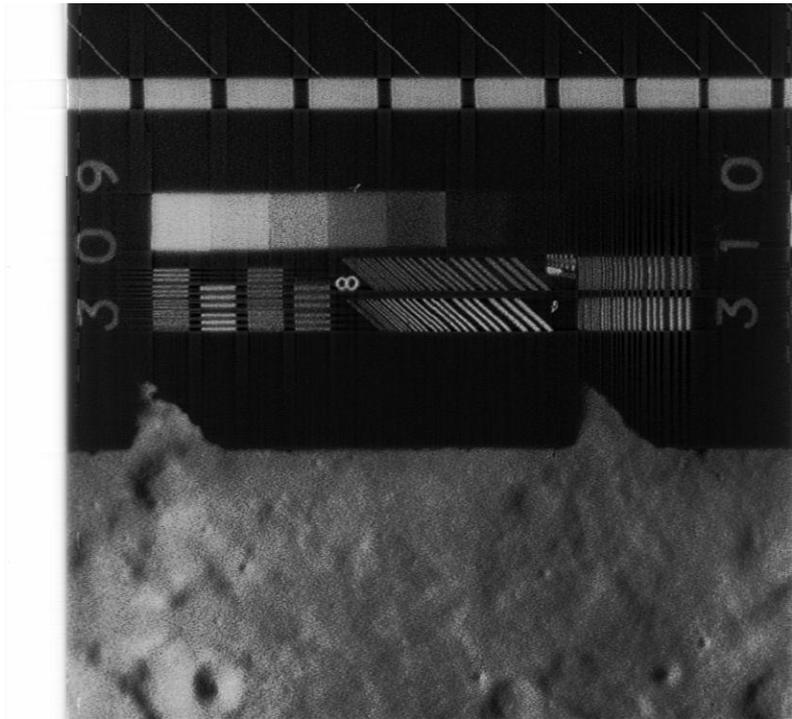


Figure 7: A calibration strip after filtering. The edges of linearity bars and gray scale squares are less sharp than shown in Figure 1, but there is no overshoot. The resolution, as measured by the parallel stripes, is about the same.

Assembly of the framelets

After reading in all of the processed framelets of a frame or subframe, the Assembly program trims them. After experimentation, the best preservation of image was found to be achieved by using three pixels on either side of the line connecting the centers of the fiducial dashes. In this way, the width of the trimmed framelets was set as 748 pixels. This will be reduced to about 720 pixels in the future to correct a 4% error in the aspect ratio of the current frames. The length is still a uniform 16550 pixels, but on the edge away from the edge containing the calibration strip. The length of selvedge (beyond the lunar image) is variable because the Framelet program trimmed only the calibration strip edge. This uneven effect could be trimmed, but it is helpful in identifying individual framelets, so it has been left.

The Assembly program butts the trimmed framelets together to create a medium resolution frame or high resolution subframe. The result is nearly seamless, even examining the continuity of the pre-exposed calibration strip and features which cross the framelet edges (See Figure 8).

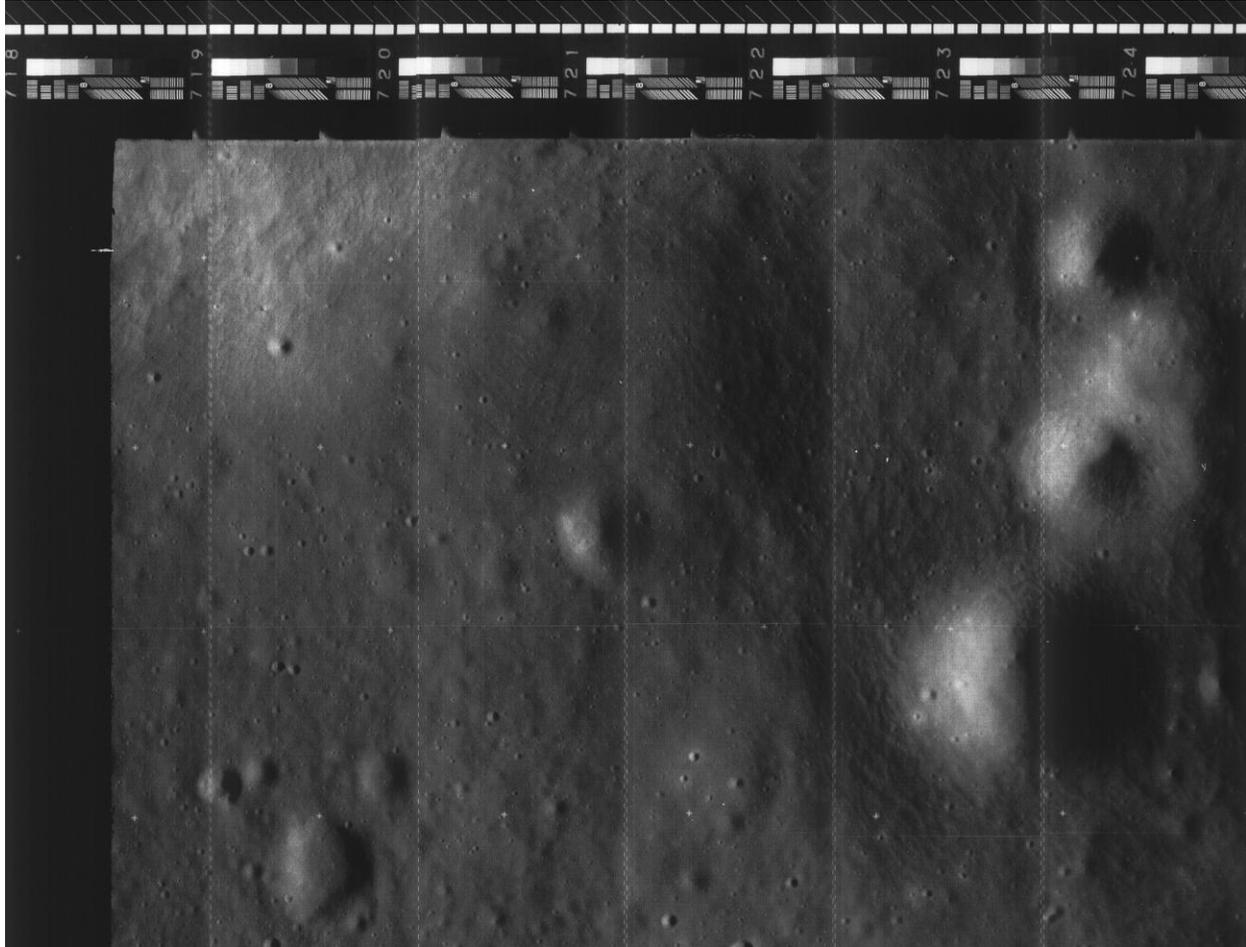


Figure 8: A series of frames after assembly. Note that there is still a significant Venetian blind effect, caused primarily by the spacecraft CRT signature.

Remaining Venetian blind effect

Systematic shade variations, a part of the “Venetian blind” effect persist (see Figure 8). These effects are due to the variation in spacecraft CRT brightness across each scan line. In principle, the CRT signature can be estimated by averaging the pixel shades across all framelets in the frame or subframe, as a function of X sample number. The sample shade values can then be divided by the CRT signature.

Linearization of shades by the Assembly program

To prepare for this compensation of the CRT signature, the sample values must be converted to values that are linear with the attenuation of the spacecraft negative, since that is where the brightness is proportional to the CRT signature times the image attenuation. In other words, the nonlinearity of the GRE signature must be estimated and reversed. The better the estimation, the better the CRT signature can be removed. Fortunately, the gray scale is of help here, along with calibration data found and supplied by Dennis Wingo of LOIRP (Figure 9). The most difficult part of the estimation was for parts of the signal that are brighter than the brightest gray scale square. Originally, the template that was used to expose the gray scale had brightness levels in steps of 0.15 db. The characteristic H-D curve of the SO-243 film, developed by Bimat, established a new set of attenuation figures for each gray scale. These have been measured from pre-exposed strips that were not flown (see Figure 9).

<u>Gray Step No.</u>	<u>SO-243 R/O Density Range</u>	<u>Reassembled Record Density Range</u>
1	0.21 - 0.29	Clipped by GRE
2	0.26 - 0.34	1.89 - 2.11 (clipped if R/O density below 0.30)
3	0.34 - 0.42	1.64 - 1.82
4	0.45 - 0.57	1.33 - 1.49
5	0.61 - 0.73	0.99 - 1.13
6	0.82 - 0.94	0.66 - 0.80
7	1.05 - 1.21	0.50 - 0.60
8	1.32 - 1.48	0.42 - 0.50 (may be clipped)
9	1.40 - 1.56	Clipped by GRE

Figure 9: This is a list of density of the pre-exposed gray scale on the developed spacecraft film, as a function of the gray scale step (personal communication, D. R. Wingo of the LOIRP project). The second column, which gives density ranges for the GRE product, is variable with GRE settings and is not used in this study. The purpose of the nonlinearity seems to be to increase gain of the signal to the 35 mm framelet film on the ground at the expense of dynamic range.

An initial attempt was made to infer the nonlinear function from the data in Figure 9 and extend it beyond the range of the gray scale. Although this was successful for shades well within the gray scale range, it often left bright and dark regions with strong stripes. Further, there was evidence that the region between GRE clipping levels was variable from image to image and even within images.

Accordingly, a new strategy was adopted late in the program. As a result of experience with the magnetic tape images digitized by LOIRP, new understanding led to a significant improvement in the linearity correction function in the Assemble program. An example of the improvement is shown in Figure 10.

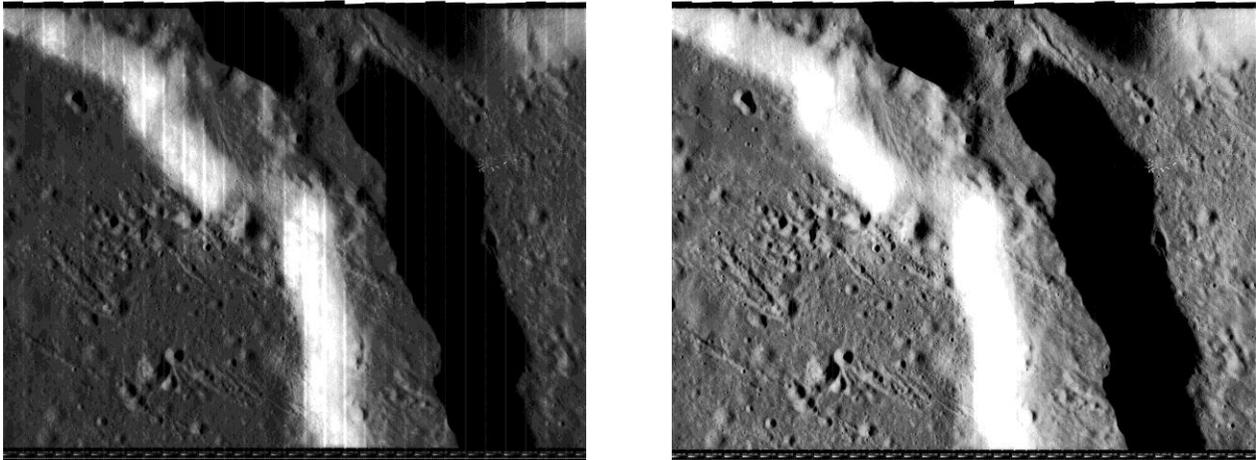


Figure 10: High resolution subframe LO5-203H2 of Vallis Schroteri (Lunar Orbiter Site V-49). These images are 2250 pixels in height, reduced from the 16,550 pixels of the .raw original. The image on the right used an earlier correction function, while the image on the left was corrected by the new process described in the text.

Note that the residual stripes at extreme bright and dark shades have been nearly removed, and the average shade is both brightened and contrast enhanced to fill the 0 to 255 range of shades. In that range, the shades are inversely proportional to the transparency of the spacecraft negative.

A new simple model of the GRE function worked quite well. A smooth symmetric function was chosen that would clip the extremes of the dynamic range as indicated by Figure 11

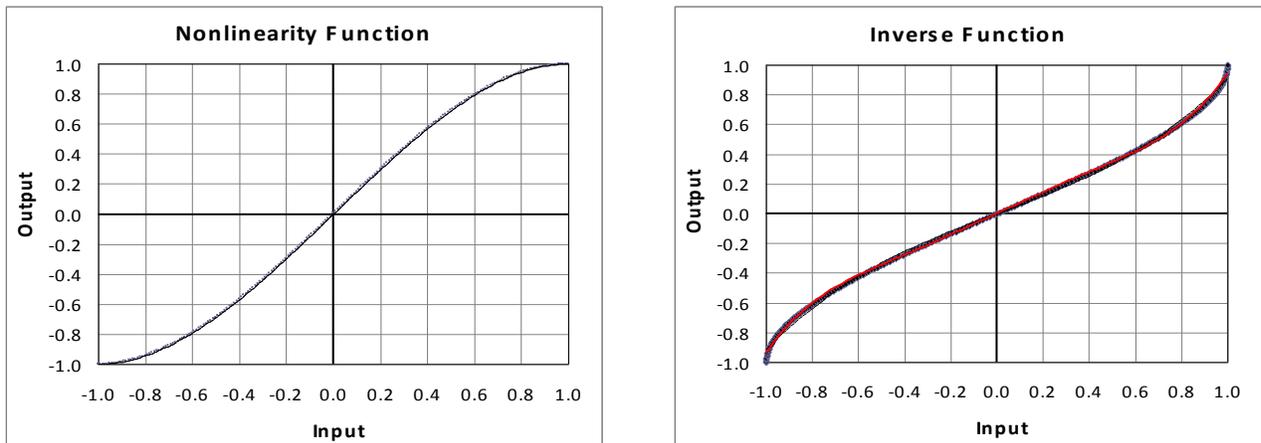


Figure 11: The postulated normalized GRE function is $Y = 1.5X - 0.5 X^3$. This is the simplest odd polynomial with the value at ± 1.0 equal to ± 1.0 and the derivative at ± 1.0 equal to 0. To linearize this function, the approximate inverse formula (shown in red) is: $Y = 0.5798 X + 0.3302 X^3$.

With the normalized function of Figure 9, the remaining parameters are the input shade value of its mean (0) and the input range corresponding to the normalized value of 1. Experimentation with the high-resolution subframes of site V-49 (LO5-202H, 203H, 204H, and 205H) gave the values of the mean shade as 104 and the range as 87. In other words, the GRE was assumed to clip shades of 104 ± 87 . The values of the mean and range have to be changed for other Lunar Orbiter missions, and other parts of Mission 5.

After expanding the values of input according to the inverse function, the mean and range was adjusted to shades of 0 to 255, to make the best use of the available shades, before computing the CRT function and compensating for it.

This procedure was quite successful in suppressing the GRE function. If it had not been, more complex functions would have been trialed. All of the frames and subframes of sites V-46, V-48, V-49, and V-50 were linearized by this algorithm. The mean and range of the nonlinear function were set as needed and recorded in the FrameLog.exl file. The results can be seen in the Gallery.doc file.

Estimation and removal of the CRT signature by the Assembly program

The CRT signature is a compound of the spacecraft CRT, the CRT in the GRE, and the CRT used to scan the film framelets. The composite CRT signature is estimated by averaging the shades as a function of pixel number across all of the framelets, including nearly all of the image pixels, but excluding those pixels outside of the exposure boundaries (for example, the pre-exposed calibration strip). The CRT signature has been found to vary from frame to frame, possibly because of variations in the spacecraft CRT, but also because framelets may have been processed by different GRE units. Attempts to use a uniform approximation to the CRT signature have been unsuccessful, so the averaging technique (which was also used on the two books based on the LPI digitized images) was found to be the best available. One example of the CRT signature is shown in Figure 11.

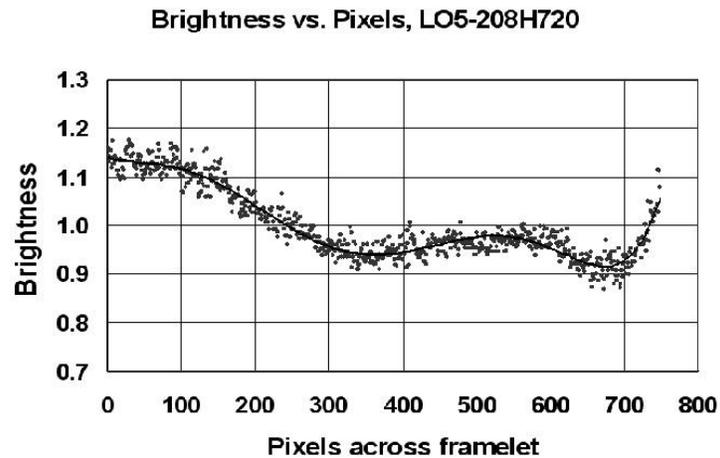


Figure 11: An example of a CRT signature derived by averaging shades across each framelet in a frame.

Each pixel of the image (including that of the pre-exposed calibration strip) is divided by the derived CRT signature. The noise relative to the trend line is partly due to granularity in the spacecraft CRT, which seems to correlate within each frame, so the noisy average is better than the trend line for correction.

After removal of the CRT signature, the signal is left linearized and stretched (to the developed spacecraft film).

The result of the linearization process and the estimation and removal of the CRT signature is shown in Figure12. This image shows that the quality is similar to photographic film, with few distracting artifacts.



Figure 12: This shows the assembled frame of LO-202M (Cobra Head and Vallis Schroteri). The resolution is reduced in this image: the full image is 17,204 by 16,550 pixels.

Remaining artifacts

Currently, the most distracting artifact set is the various patterns (bubbles, curved stripes) that were due to the Bimat developing film sticking to the negative during the development storage. This was due to a conscious compromise between the number of exposures that could be made and observation of rules on the time the pair of films could be left in a storage buffer. To follow these rules would have required intermittent moves of the film past the cameras without exposure. The development artifacts are caused by the negative emulsion sticking to the development film and being physically torn off the negative. Consequently, the information is lost. Airbrushing would reduce the distraction but could be misleading.

The Framelet and Assembly programs have been developed to reduce distracting artifacts that interfere with the science of geologic interpretation. Some artifacts have been left that should be reduced for other purposes.

For photogrammetry, the linearization of shades could be extended from the developed spacecraft film to the lunar surface brightness. The first step would be to invert the H-D curve of the SO 143 film, as developed by the BiMat process. Then the vignetting effect of the appropriate lens should be corrected.

For precise cartography, the reseau crosses could be used to correct some remaining distortion of the scanning process. This step would also support the construction of mosaics of the frames. Measurements of the reseau crosses (Appendix A) suggest that this step may be performed generically, at least for each mission.

For photogrammetry, reseau crosses should be reduced by interpolation, as have the fiducial dashes. As reported in Appendix A, identification of the reseau crosses is feasible, and in fact identification of them has been included in the Framelet program as an option, but has not been applied so as to leave the reseau crosses in the images for further measurement and processing if desired.

Appendix A

Analysis of Geometric Pattern of Crosses on Lunar Orbiter Images

For Lunar Orbiter Missions 2 through 5, a geometric pattern of crosses was pre-exposed on the film as an aid to cartographic analysis. These crosses were used by USGS with their ISIS program to guide their assembly process. In the later missions (3, 4, and 5) the pattern assured that at least two columns of crosses would appear in each framelet, 12 crosses in one and 11 crosses in the other. There were alternate rows of crosses, with the rows offset. In this note, the rows are labeled “rank” with ranks 1.0 to 12.0 in one column and ranks 1.5 to 11.5 in the other. A figure showing the pattern is in the introduction to “Lunar Orbiter Photographic Atlas of the Moon” by Bowker and Hughes.

The assembly process I am using is based on the fiducial dash lines that come from the electronic scanning structure in the spacecraft. This provides a more satisfying assembly for photogeology than the USGS method. However, further improvement may result from using additional information from the crosses.

A program was written to search for and measure the coordinates of the crosses. Identification of the crosses was done by using a template of the shape of the crosses, correlating the template with the image, and seeking the maximum correlation coefficient. A search strategy was designed to find the first cross in each of two vertical columns of crosses. Once one was found, the search was narrowed, based on the pattern. The search strategy is shown in Figure 2. A typical cross is shown, enlarged, in Figure 3.

The program searched for two columns of crosses, although in a few cases there are three within a framelet. The measured center coordinates were output in a text file for each framelet.

In the next step, the text files for a set of 6 adjacent framelets (LO5-208H718x through 724) were imported into a spreadsheet for analysis.

The X values (across the framelets) showed a horizontal drift in each of the framelets. The average X value of each column was subtracted from that of the individual crosses to produce Figure 1.

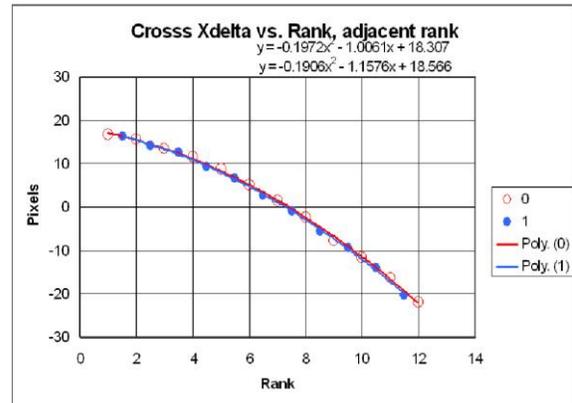


Figure 1: Drift of the columns of crosses in the 0 (12 crosses) and 1 (11 crosses) columns, as a function of rank (row index). The points are within a pixel of the curve.

The drift could be, in principle, either an artifact of the cross pattern or of the scanning mechanism. Similar drift can also be seen in the frame edge in framelet 718x, so the expectation is that all or nearly all of the drift can be attributed to the scanning mechanism, probably a misalignment of the axis of the scanning mirror.

The situation in the Y direction (along framelets) is more complex. A plot of the Y measurements is shown in Figure 4.

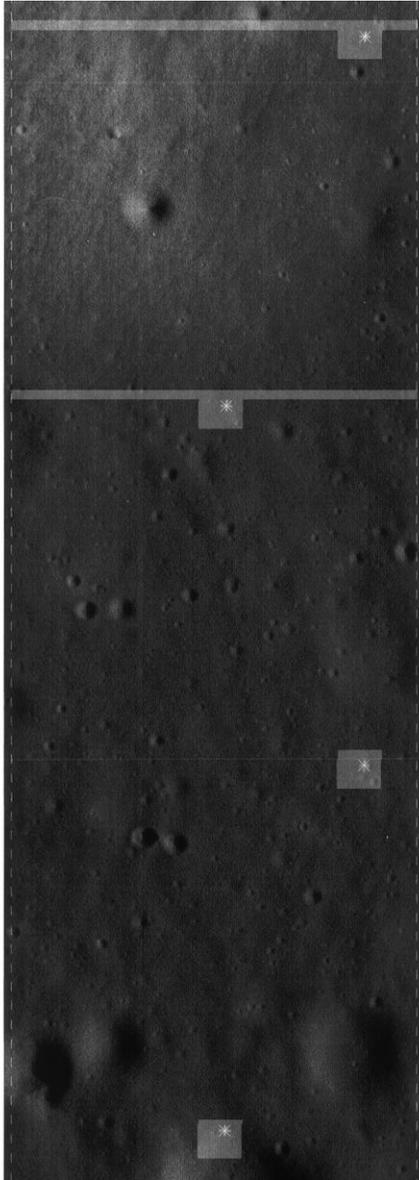


Figure 2: A section of the image, showing a part of the geometric pattern of crosses. The bright areas were placed by the program to show its area of search. Initially, it simply searched across the image. When a correlation coefficient exceeded a threshold, it narrowed the search pattern. After each area was searched, the program jumped to the probable area of the next

cross. As each cross was detected, the search area of the next cross was adjusted to follow drift in the pattern.

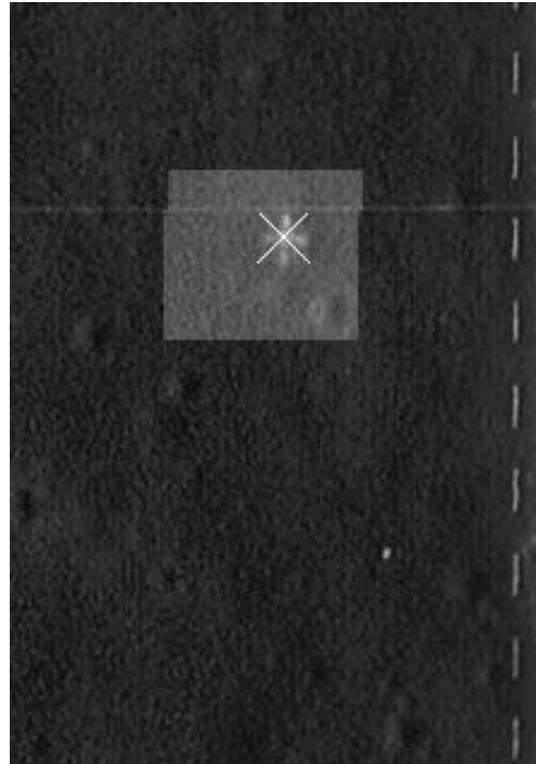


Figure 3: A typical cross. The gray area shows the search area, which was slightly shifted to the left when the cross was first detected. The diagonal lines were inserted by the program to show the measured center of the cross, the point of maximum correlation coefficient.

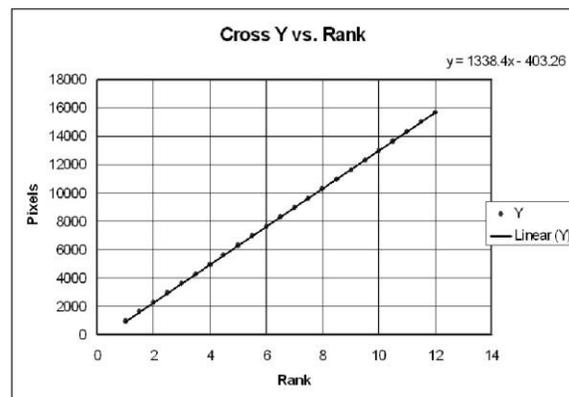


Figure 4: The Y measurements (along framelets) of crosses, as a function of rank.

At first look, it appears that Figure 4 is similar to figure 1, and indeed the percentage accuracy of the cross positions is similar; about 0.1 percent. But while that is very good for the horizontal direction, where the image is only about 750 pixels wide, it is not so good for the vertical direction, which is about 15,000 pixels long. The deviation from the linear trend line is shown in Figure 6.

in the upper part of the framelets. Although one cannot be sure of the separation of pattern from scan deviations, the initial approach will be to assign the trend line to scanning and correct for it. If the remaining errors are due to something like vibration of the scanning mirror as it moves, they cannot be corrected.

The deviations “Ydelta” from linear could again be due to either the cross pattern or to the scanning mechanism. There is clear evidence that there is some contribution from the crosses, because the average deviation of the ranks with integer values are systematically offset from the average deviation from the ranks with half value.

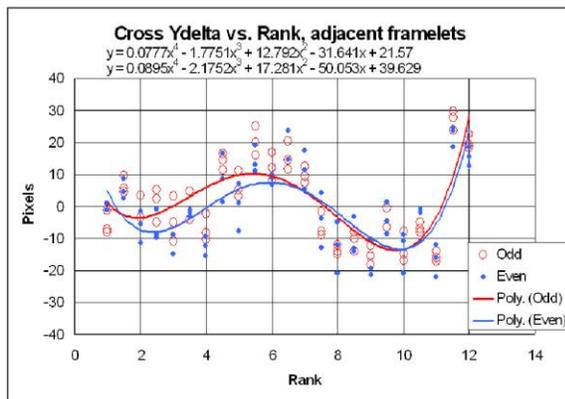


Figure 5 Deviations of the vertical measurements of crosses from the linear trend line of Figure 5, as a function of rank.

To try to separate pattern deviations from scanning deviations, the framelets were grouped into odd and even sets. It is known that scanning deviations result in offset of image patterns like craters across assembled framelets, even if optimally aligned and stretched. This could result from systematic variation with each direction of scan. Indeed, examination of Figure 5 shows that the trend lines for odd and even framelets depart significantly (approximately 8 pixels)