Recovering Lunar Orbiter Framelets from Digitized Magnetic Tape Record

A. Epps, M. Sandler
www.moonviews.com

INTRODUCTION

The goal of the Lunar Orbiter Image Recovery Project (LOIRP) is to digitize and archive the magnetic tape records generated by the five Lunar Orbiter spacecraft in the mid-1960s. The readout scanners utilized onboard the Lunar Orbiter spacecraft employed a phosphor-covered anode bombarded by an electron beam to focus a spot of light on 70mm film developed onboard the spacecraft. This light was modulated by the density of the image and read by a photomultiplier tube. Each individual pass of this scanning procedure across the 70mm film produced a thin strip of a larger image, referred to as a "framelet". The product of the spacecraft's readout system was a video waveform that was modulated and transmitted to three DSIF stations and recorded onto 2-inch magnetic tape via Ampex FR-900 data recorders. This document discusses the process by which these video signals were converted into digital images.

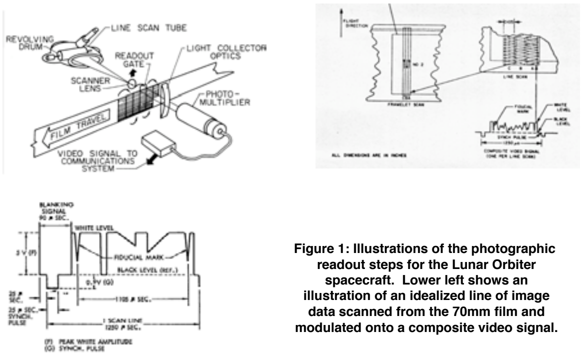

Figure 1: Illustrations of the photographic readout steps for the Lunar Orbiter spacecraft. Lower left shows an illustration of an idealized line of image data scanned from the 70mm film and modulated onto a composite video signal.

PROCESS

In order to convert the analog signal outputted by the FR-900 data recorder into a digital image, a software equivalent of the 1960s-era Ground Reconstruction Equipment (GRE) had to be developed. This process starts with the analog waveform from the tapes digitized at 5×10^6 samples per second at 16-bit resolution, and concludes with digitized framelets in a format that can be readily assembled into complete frames.

The analog waveform (illustrated in **Figure 1 - lower left**) is referred to as the "composite video signal" and contains a synchronization pulse between each line of image data. Although nominally each line should occupy 1250µs of tape, this parameter exhibits acutely accumulating drift. This is evident in **Figure 2**, which reveals the drifting artifacts evident when for a line width of 1250µs is assumed for a section of the data.

An initial attempt was made to correct the drift by extracting, for each line, a specific feature of the synch pulse -- namely the midpoint along the synch pulse's largest falling edge. Surprisingly, the results of such precise per-line synch pulse extraction produced apparent visual artifacts. In particular, the lines in each framelet's calibration marker, intended to be straight, appeared noticeably jagged (see **Figure 3-left side**).

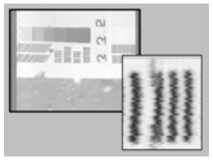

Figure 3: Image created by beginning new lines after each individual synchronization pulse. The inset section consists of 26 lines whose synchronization pulses were visually extracted manually by an operator.

To ascertain if the jaggedness was a result of the how the code was estimating the synch pulse's position per line (versus being inherent to the data), an experiment was carried out in which the two authors manually and independently extracted the synch pulse for the same twenty adjacent lines (via visual inspection of the synch pulse signals using MATLAB). With almost no variance in the feature's localization between the two operators, the two estimates where averaged (for each line), and the results plotted for comparison - as shown in **Figure 3 - right side**. The results where no better than those obtained by applying the result of the automatic method's per-line estimates,

While the non-zero mean bias (versus 1250 µs) in the line-length produces the drift clearly visible in Figure 3, on close inspection (not shown here) there is no visually apparent line-to-line jaggedness at the calibration markers. As a result, it become clear that a long-run-averaging method would be needed to correct the drift, but without introducing any line-to-line jaggedness.

Figure 2: A section of image readout with lines of constant 1250µs length. The image has been compressed vertically by 10x to emphasize the drift. The section highlighted in red occurred when the readout tape player lost lock with the tape's pilot tone.

THE PROBLEMS

At this point, context should be provided on some of the other significant problems that needed to be addressed to maximize the quality of recovered framelets. These problems included:

- The presence of ringing in the data
- Variability in the change in levels
- Loss of playback machine lock with pilot tone, and other playback machine errors
- Erratic and unpredictable variability in needed drift correction, even within a tape

Compounding these difficulties is ringing present from the analog electronics of the playback and demodulation steps. The presence of this ringing complicates any feature detection method that relies on fitting an idealized synchronization pulse. Compare the waveform in **Figure 4** with the illustrated waveform in **Figure 1-lower left**. The substantial ringing present in both the tip and the front porch of the synchronization pulse causes unpredictable behavior in feature detection algorithms that attempt to fit the pulse with a stylized template pattern, or even one fitted from the local data of the tape.

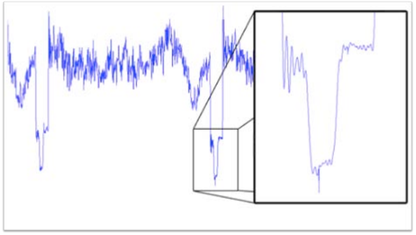

Figure 4: Ringing in the synchronization pulse adds a layer of difficulty when attempting to localize via feature detection methods.

The solution also has to deal with the problems of tape-to-tape variability in levels, since the range of amplitudes can be significant depending on the calibration of the recording machine as well as the average picture level of the image transmitted.

Another problem that manifests itself is the occasional loss of lock between the tape and the playback machine (this can be observed inside the red outline in **Figure 2**). The FR-900 tape recorder maintains synchronization with the tape by the use of a 500 kHz pilot tone embedded in the signal on the tape. If this signal is lost, the picture momentarily disappears while the machine attempts to resynchronize.

Further complexity associated with applying a drift correction is that the magnitude of the correction required not only varied from tape-to-tape, but *moreover* varied within a tape, and in a manner that lacked any consistent predictability.

These issues create several compounding problems. For instance, one could imagine simply taking occasional synchronization pulse localizations along the length of an image tape and applying an averaging effect to compensate for inexactness in the location of the synchronization pulses. However, if the playback machine had lost lock with the tape's pilot tone an entire section of the image could be processed incorrectly, and there would be no indication of such so that attempts could be made to reacquire the missing data.

THE SOLUTION

It was determined that the best way to address the various problems inherent in the structure of the data would be to employ an iterative algorithmic method known as RANSAC , or RAndom Sample CONsensus. The RANSAC method can be differentially effective needing to fit a pattern on inliers while simultaneously separating out outliers. In our case the inliers are the valid synchronization pulses, and outliers result from portions of the data that don't contain pulses due to loss of playback machine synchronization or from ringing causing any particular synchronization pulse to be poorly localized. The pattern being fit here was parameterized, primarily via a local estimate of the *phase (*meaning the drift rate -- as again, a 'perfect' estimate for each line produced the jaggedness shown in Figure 3), and secondarily, by an *offset* value that ensures, given a correct drift rate, that the resulting straightened synch signal runs down the edge of the framelet (versus running, for instance, directly down the center of the framelet).

The steps of the algorithm are outlined below:

- Attempt to localize the synchronization pulses for a sample set of (e.g. 700) lines
- Take a random draw of a subset of the samples (e.g. 3-4), and look at the relative fraction of the rest of the estimates that are consistent with that random draw (e.g. within ~0.75 µs)
- Repeat these steps, tracking the best random draw, as measured by the fraction of the datums which are consistent with the draw
- At convergence in that fraction associated with the best random draw, compile a list of candidates which deviate significantly from that consensus estimate (affected by ringing or other localized deviations in the signal, loss of playback machine lock, etc). Classify these candidates as outliers, and remove them from the pool
- Over the remainder of the pool (the now assumed inliers) perform a least-squares estimate for the *phase*, and with the value of the associated inputted local *offset*, rescale and realign the image data in accordance
- Advance approximately 300 lines and repeat

Over the vast majority of the data, where the playback machine is locked to the pilot tone contained in the signal on the tape, this approach provides excellent recovery of synchronization without introducing any jagged visual artifacts. In those cases where the playback machine lost lock with the pilot tone from the tape, the consensus remains extremely poor despite extensive sampling. In such a case, the section of data is automatically identified as being incoherent, and the algorithm does not attempt to rescale it. Instead, it advances along the tape and continues to look for consensus. When sufficient consensus is again found, the software recognizes the location at which playback synchronization is regained and continues to rescale and synch-pulse-shift align the image data.

Unlike the archival film record, the digitized record produced by LOIRP eliminates the film-based Ground Reconstruction Equipment (GRE). The GRE's consistent inaccuracy in recovery of synchronization led to the introduction of a horizontal skew of a fraction of a degree along the length of an entire framelet. When scanned at high resolution by USGS (970 px width by 16,550 px height) this skew exhibited deviations of as much as 30 pixels between the top and bottom edges of a framelet. This skew has been completely eliminated by the LOIRP process's more sophisticated recovery of synchronization. Additionally, the LOIRP process preserves the full dynamic range of the composite video signal, unlike the GRE process that clipped density ranges below and above certain thresholds, allowing LOIRP images to better capture the full quality of these images for scientific posterity.

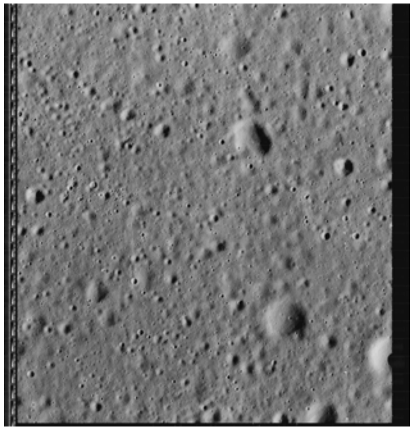

Figure 5: Image assembled from framelets generated using the automatic synchronization recovery algorithm. These framelets were additionally processed by Charles Byrne to remove other artifacts introduced onboard the spacecraft by the film readout system.

References: [1] **C. J. Byrne (2009)** ***Analysis and Correction of Artifacts in Lunar Orbiter Photographs, 07-LASER07-0005 [2] The Boeing Company (1967) Lunar Orbiter I Photography, NASA CR-847, 112 [3] The Boeing Company (1968) Lunar Orbiter III Photographic Mission Summary NASA CR-1069.***