ONBOARD IMAGE CORRECTION

D. R. Martin, A. S. Samulon, and A. S. Hamori
TRW Defense and Space Systems Group

ABSTRACT

This paper describes a processor architecture for performing onboard geometric and radiometric correction of LANDSAT imagery. The design uses a general purpose processor to calculate the distortion values at selected points in the image and a special purpose processor to resample (calculate distortion at each image point and interpolate the intensity) the sensor output data. A distinct special purpose processor is used for each spectral band. Because of the sensor's high output data rate, 80 M bit per second, the special purpose processors use a pipeline architecture. Sizing has been done of both the general and special purpose hardware.

I. Introduction

In performing analyses of imagery produced by earth resource observation satellites, it is frequently desirable that two images of the same scene be registered; that is, each physical part of the scene is in the same location on the two images so that the picture elements (pixels) of the two images can be aligned. Such precision is not easy to attain, mainly because of varying distortions and viewing conditions from one image to the next. Registration can be accomplished, however, by estimating these distortions and processing the image data accordingly.

With the advent of LANDSAT D, currently under development, earth resource observation satellites are coming closer to operational, rather than experimental, use.

Key features of LANDSAT D are:

- Ground sample distance of thirty meters
- Geodetic accuracy to 3 meters (RMS) using ground processing
- Visible, near infrared, and thermal infrared spectral bands
- Swath width of 185 kilometers
- Repeated coverage every sixteen days
- Seven spectral bands having eight bit radiometric resolution

While the features of LANDSAT D are all desirable for an operational system, several elements must be added to create a truly operational system. Chief among these is rapid receipt of corrected imagery by the user. Due to the high data rate (10 million picture elements per second), present plans for LANDSAT D involve geometric correction (on the ground) of only ten per cent of the (over land) imagery. Corrected imagery will be produced within two days of transmission with shipment through the mail adding several more days delay between imaging and availability of the data.

LANDSAT D consists of a Multimission Modular Spacecraft (MMS) combined with an instrument module containing the Thematic Mapper (TM). As the spacecraft passes over a region, the TM scans back and forth, as shown in Figure 1. Each scan contains 16 scan lines spaced at approximately 30 meter intervals and provides coverage of seven spectral bands. In each spectral band, the moving image is swept past an array of detectors by the scan mirror action. Each detector combined with mirror scan motion produces one image line in one spectral band. The scan line corrector compensates for spacecraft motion during the scan, thus yielding straight scan lines perpendicular to the spacecraft velocity vector. As the ground footprint of each detector moves 30 meters cross-track, its output is sampled and converted to an 8 bit digital word. These words are then multiplexed to form an 83.268 Mbps data stream.

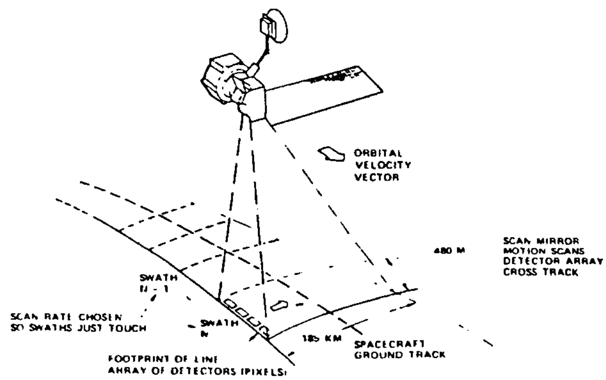

Figure 1. Multispectral Scanning Sensor Geometry

The on-board processing technique described in this paper will provide corrected data in real-time. The registration accuracy, although not as good as that produced by ground processing, will be approximately half the ground sample distance (15 meters). The expected accuracy should be quite sufficient for doing crop yield assessment using multi-date imagery, change detection, and determining progress of environmental disturbances such as crop disease and fire. Further processing on the ground will still be able to provide the extremely high accuracy imagery required relatively infrequently for mapping purposes.

The primary factors that are making on-board image correction viable are: 1) the extremely high accuracy ephemeris information to be available in realtime from the Global Postioning System (GPS), 2) the availability of small, high density, low power memories, and 3) high speed, low power processors.

Even with these technological advances, practical solution of the on-board correction problem requires a subsystem architecture that is a balanced combination of a general purpose computer and special purpose hardware using both parallel and pipeline processing.

More specific details on both the sources of distortion and the geometric correction technique can be found in Reference 1.

II. Registration Problem

Two images of the same region are said to be registered when each physical part of the scene is in the same location in each image. This allows direct comparison of different images of the same region. Unfortunately, unprocessed images do not meet this criterion because of the distortions contributed by the sources discussed in Section III. Ground or on-board processing can be used to correct the imagery. Depending on the amount of on-board correction, the amount of additional ground processing required to complete correction of the imagery will vary. Some of this additional ground processing is very simple and can be done readily by individual users. Therefore, it is important to consider the degree of correction obtainable with different amounts of on-board processing. In this section various levels of registration are defined. To achieve each successive level requires additional on-board capability.

Isodistance Registration

Isodistance registration of different images of the same region requires the scene to have constant interpixel distance and parallel scan lines. Figure 2 illustrates this concept. In the unprocessed images, the distance between pixels is unequal. After isodistance registration has been accomplished, the distance between pixels is

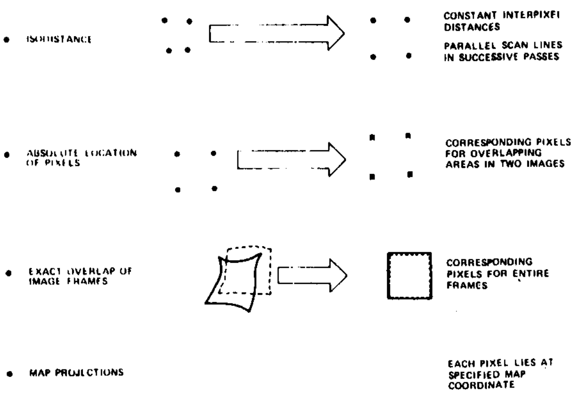

Figure 2. Registration Levels

the same in each image and scan lines are parallel. If the relative alignments of the two isodistance registered images were known, they could be lined up and compared since any two given locations in a scene are separated by the same number of pixels for any two isodistance registered images of the scene. However, knowledge of this misalignment is not required for the on-board isodistance registration. In fact, direct comparison of pixels still will most likely require further interpolation since the relative shift between corresponding pixels is not necessarily an integer number of pixels.

Absolute Location of Pixels

Absolute location of pixels accomplishes everything that isodistance registration accomplishes. In addition, the relative alignment of the two images being compared is always a known number of pixels. This allows direct comparison of the intensity of corresponding pixels without further resampling. This is illustrated in Figure 2 by showing that the square regions can be made to coincide.

In comparison to isodistance registration, the absolute location registration requires a more precise distortion measurement technique. The absolute magnitude of all distortions is important now, not just those which vary during the scene. The corresponding correction technique is comparable for isodistance registration and absolute location registration. Since subsequent resampling on the ground is avoided by absolute location registration, it is clearly advantageous to do it. However, whether or not it can be done depends upon the capability of the on-board distortion measurement technique.

Exact Overlap of Image Frames

The pixels in images which are absolute location registered can be compared by extracting corresponding portions of the two images. Note that the edges of the scene are not required to overlap. Consequently, if subsequent imagery is used to compare with a reference scene, up to four images must be used to reconstruct the same region covered by the original image. This difficulty is overcome by exact overlap registration, which causes the pixels of subsequent images to be in the same position as in the reference images.

This level of registration requires the same distortion measurement capability as absolute location registration. The dichotomy between absolute location registration and exact overlap registration is in the geometric correction technique which must be employed. Relatively little ground processing is saved by this technique compared to absolute location registration. However, if the increase in on-board processing complexity is relatively small, this additional level of registration is worthwhile.

Map Projection Rectification

Map projection rectification requires each pixel in an image to lie at a specified map coordinate; furthermore, the interpixel spacing must correspond to that of the map projection. This requires more than registration, since the repeatability of imagery does not guarantee the image corresponds to any type of map projection.

Production of images which are rectified with respect to a well-known map projection (e.g., Universal Transverse Mercator or Space Oblique Mercator) is not attempted in this implementation. A significant amount of on-board storage is required to produce map projections such as these. However, distortion due to earth rotation is important to eliminate, since this rotation will affect different images of the same region in a different way. This requires some sort of map projection to provide a measure of the effect of earth rotation. Such a map projection does not need to be a conventional map projection.

III. Sources of Distortion

Raw data received from the Thematic Mapper cannot be directly registered with other data scanned on previous passes over the same region because each unprocessed image is affected by a unique set of distortions. The four primary causes of distortion are:

- Sensor Caused Distortions
- Attitude Variation
- Alignment Variation
- Ephemeris Variation

Image distortion will result in a corresponding registration error if the distortion is not estimated and removed. After performing this geometric correction, the resulting registration error is determined by the accuracy with which the distortion is estimated, not by the actual magnitude of the distortion. In this section no estimation or correction is assumed, hence distortion and registration error are virtually synonymous concepts here.

Sensor Caused Distortion

The scanning motion of the Thematic Mapper must be precisely the same on successive passes over a region if no distortion is to be introduced. Variation in the active scan duration (i.e., scan velocity) will cause stretching (or compression) of the pixel spacing within a scan line. Variation in the scan period will cause the spacing between scan lines to be different for subsequent images of a region, causing different images of the same region to have a different number of scan lines. Figure 3 illustrates these variations.

In addition to variation in scan period and active scan duration, an additional source of distortion is scan nonlinearity. That is, the angular velocity of the scan mirror does not remain precisely constant during the scan, thus producing irregularly spaced pixels.

Attitude Variation

In the normal mode of operation, the attitude of the spacecraft can be commanded to take on any desired value. Nominally, the attitude is such that the scans are perpendicular to the orbital velocity vector with the Thematic Mapper pointing towards the center of the earth at mid scan. The attitude of the spacecraft is controlled by the attitude control system located in the spacecraft.

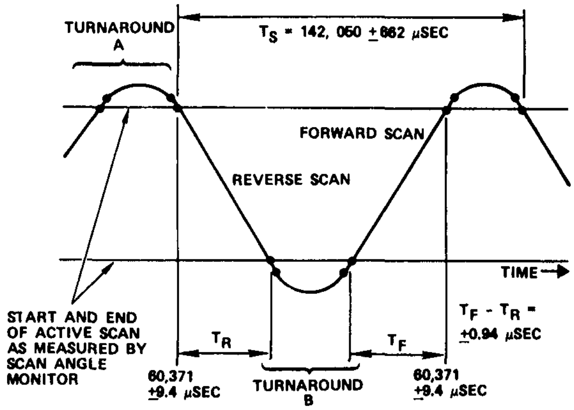

Figure 3. Thematic Mapper Scan Durations

Variation in the absolute attitude of the spacecraft with respect to a previous pass will cause an absolute location registration error proportional to this attitude variation. Such variation is limited by the accuracy of the star tracker. Isodistance registration requires a stable attitude reference during the scene, but is relatively unaffected by the absolute accuracy of this reference. Consequently, isodistance registration is primarily determined by the gyro drift in the stellar-inertial attitude reference system.

Alignment Variation

The attitude of the spacecraft is controlled by the attitude control system which is located in a separate structure than the Thematic Mapper. For the reasons previously cited, the attitude of the Thematic Mapper must be held constant (with respect to the earth-pointing frame of reference) to prevent distortion. However, the coordinate axes of the Thematic Mapper are not the same as those of the Attitude Control System. The difference between these sets of axes exhibits both long term drift and a short term variation due to thermal effects. Consequently, even if the spacecraft's attitude were to remain constant, alignment variation would distort the scanned image.

Ephemeris Variation

The location of the spacecraft with respect to the ground at a given time of day can vary significantly for different passes over a region. Pixels compared at the same time of day for subsequent passes, with no knowledge of spacecraft location, can have a significant offset in pixel location. Absolute location registration requires that this offset be known and corrected and is essential if different images of the same region are to be compared.

Variation in the spacecraft altitude for different passes over a region affects pixel spacing in the cross track direction for both isodistance and absolute registration. Except for the effect of variation in the orbital velocity, the difference in the spacecraft's along track and cross track position has no impact on isodistance distortion. The altitude and orbital velocity of the spacecraft changes very little in comparison to the cross track drift for different passes over a region. Consequently, isodistance distortion caused by ephemeris variation is far smaller than the corresponding absolute distortion.

Correcting these variations requires more than a simple shift of the image. Specifically, perspective differences and earth rotation will combine to distort the image data if the spacecraft is not in precisely the same location as on previous passes. Perspective difference arises in part because the Thematic Mapper samples the detectors at equally spaced angular increments. Because of this, the pixel spacing on the ground increases with distance away from the ground track. This variation in pixel spacing prevents simply shifting the image to produce alignment. Earth rotation shifts the scene during the scan, thus producing an image which is significantly skewed with respect to a conventional map projection of the earth's surface. In addition, earth rotation causes skewing of scan lines for different images of the same region.

IV. Distortion Estimation

Registration is accomplished in two steps: estimation or measurement of the various possible factors which affect registration, and compensation for these factors either through data manipulation or spacecraft commands. The first step of this process is the topic of this section; the second step will be addressed in the next section.

The accuracy with which these distortions can be estimated is of particular concern. After the images have been corrected based on the distortion estimate, the remaining registration error is caused primarily by the error in estimating the distortion. The technique usually employed to estimate distortion on the ground uses ground control points which consist of 32 by 32 pixel subimages with known location. The received imagery is correlated with the ground control point to determine the proper position of one pixel. A dynamic model is used for the distortion, with the correlation information serving as observations of the distortion process. By using Kalman filtering, the distortion at each pixel in the image can be estimated and corrected. Variations which occur at a higher frequency than can be measured by ground control points must be measured by some other technique or else simply ignored if they are sufficiently small.

Unfortunately, the use of ground control points requires a significant processing and data storage capability. In order to make on-board processing viable, the distortion measurement technique described here does not use ground control points. The sensor-caused distortions are measured by the scan angle monitor. Alignment is calibrated in a preoperational mode from the ground station by using ground control points. This alignment is transmitted up to the spacecraft and periodically updated. Attitude is determined by using a stellar-inertial attitude reference system which uses an advanced star tracker design. Ephemeris is determined from the Global Positioning System (GPS).

Although this technique is not as accurate as one using ground control points, it is still capable of producing sub-pixel registration. In fact, in the isodistance sense the registration is nearly as good as can be obtained with ground control points. Table 1 summarizes this performance. The pixel spacing for the Thematic Mapper is 30 meters, which means these registration errors are approximately ½ pixel.

Table 1. One-Sigma Registration Error for On-Board Distortion Measurement Techniques

DISTORTION SOURCE	ALONG-TRACK ERROR (METERS)	CROSS-TRACK ERROR (METERS)
SENSOR	1.6	1.3
MISALIGNMENT	7.3	5.2
ATTITUDE	10.3	10.2
EPHEMERIS	5.0	5.0
RSS	13.7	12.6

Measuring Sensor Distortions

Of all the sources of sensor-caused distortion, by far the largest is variation in scan duration. The Thematic Mapper contains a scan angle monitor which furnishes accurate information both about pixel spacing within a line and pixel spacing between lines. The scan angle monitor (SAM) optically measures when the scan mirror enters the active scan region, when it is at its midpoint, and when it leaves the active scan region. The multiplexer inserts a major frame sync word into the downlink data stream (84 Mbps, interruptible at 8 bit word boundaries) when the SAM indicates the mirror has entered the active scan region. After the end of the scan pulse occurs, the multiplexer inserts an end of scan pattern, line length, calibration and zero restore information.

Scan nonlinearity, if it is significant, will be calibrated for each Thematic Mapper. The extent of this nonlinearity is currently not determined, since Thematic Mapper is not yet operational. A piecewise curve fit can be used to model this nonlinearity, if necessary.

Attitude Determination

The attitude of the spacecraft relative to the true earth-centered inertial frame is determined by 1) approximating the earth-centered inertial frame with an on-board stellar-inertial frame of reference, and 2) commanding the spacecraft to point in a specified direction relative to the stellar-inertial reference.

The three axis attitude reference is derived from integrated gyro data. Attitude and gyro biases are updated periodically from strapdown star tracker measurements which are processed by an onboard algorithm, typically a six-state Kalman filter. If this reference is sufficiently accurate, the attitude distortion is not a matter of concern.

This is possibly the most difficult source of distortion to measure on-board, since the accuracy of the attitude reference system is typically far less accurate than the estimate obtained with ground control points. By using a star tracker of advanced design, a one-sigma attitude reference system accuracy of 3 arc-sec (each axis) is achievable. This is sufficient to attain sub-pixed registration. Even without an advanced star tracker, the isodistance registration will still be excellent.

Alignment Calibration

The attitude of the Thematic Mapper relative to the stellar-inertial frame must be known if the scan is to be of the desired place on the earth. The relative alignment of the Thematic Mapper and the attitude reference system can be determined readily through the use of ground control points. Any other technique for determining this alignment would be extremely difficult. In order to avoid the use of ground control points in an operational mode, this alignment can be performed periodically on the ground and transmitted to the spacecraft. In the absence of significant mechanical stress being placed on the spacecraft, this misalignment should be relatively small. In any event, the effect of this distortion on isodistance registration is minimal because it is slowly varying.

Ephemeris Determination

When operational, the Global Positioning System (GPS) will provide position information accurate to within 15 meters (three-sigma). By using this information to update a Kalman filter model of the spacecraft's orbit, the position, velocity and acceleration of the spacecraft can be accurately estimated. This processing is performed in the GPS receiver, with the results used as inputs to the image processor.

V. Geometric Correction

In order that all images of the same area on the earth be registered with one another, it is necessary to have a reference coordinate system against which to compare each image as it is generated. The goal of the registration procedure, then, is to generate an output image whose pixels correspond to specific locations in the reference coordinate system. The intensity value of each output pixel must be estimated from the data actually scanned by the Thematic Mapper. Thus the generation of each output pixel requires two steps: 1) determination of the location in the actually scanned data corresponding to the specific output pixel, and 2) estimation of the output pixel intensity value from the neighboring scanned values.

Determination of the location in the scanned data corresponding to a specific output pixel requires relating the scanned data to the reference coordinate system. This relation is computed using the GPS ephemeris data, the attitude sensor and control system output, Thematic Mapper scan monitor outputs, and occasional alignment updates. Appropriate selection of the reference coordinate system used is crucial to practical on-board implementation of geometric correction because it affects the amount of data that must be buffered. Because of the complexity of the computation relating scanned data to the reference coordinate system, the complete calculation is performed only for a selected subset of points in the coordinate system. The location in the scanned data corresponding to other points in the reference coordinate system is estimated using an interpolation polynomial.

Once the correspondence has been established between locations in the output frame and the scanned data, neighboring scanned values are used to interpolate an estimate of the intensity value of the output pixel. Because some of the neighboring values are produced by different photodetectors, each with its own nonlinear response to the incident illumination, correction of the detector responses precedes the interpolation process.

Correction Calculation

The output point corresponding to a given input pixel can be computed by using the pierce point calculation. The pierce point calculation uses the ephemeris, attitude, and scan information to determine the latitude/longitude of the input pixel on the earth's surface. This is converted into a point in the output space by using an appropriate map projection. The input point corresponding to a given output point can be determined by iteratively estimating the point in the input space based on the resulting pierce point calculation. This has been shown in ground processing to require at most three iterations.

Although the pierce point calculation can be performed for each output pixel, this requires an enormous computational load. The solution to this problem is the creation of an interpolation grid consisting of a subset of the output picture elements. The distortion is calculated only at the grid points with interpolation used to evaluate the distortion at the other output pixels. Note that this interpolation (which is used to evaluate the geometric distortion) has no relation to the interpolator used to calculate the output pixel intensity (cubic convolution interpolator).

The correction calculation must be performed once each scan (0.07 seconds). This calculation must consequently be made as simple as possible to minimize the on-board processing requirements. The reference map projection (coordinate frame) is of particular concern since the latitude/longitude of each pierce point must be converted to this coordinate frame. Properties desired of the map projection include the following:

- Scan lines nearly parallel to the X-axis of the projection (reduce buffering)
- Valid over the entire orbit
- Simple computationally
- Use elipsoidal model for the earth's radius (allow registration)
- X and Y axes nearly perpendicular (two-dimensional resampling)

Although none of the four projection used for ground processing of LANDSAT D data satisfies these properties, a slight variation of the oblique Mercator projection does. Unlike the space oblique Mercator projection, this projection is not swath continuous and must have a different transformed equator for each image frame.

The use of this projection facilitates the four following simplifications in calculating the distortion at the grid points:

- A simple cross track distance expression to calculate distance relative to known pierce points
- Linearity of the vertical distortion across the scan line
- Avoidance of inverse mapping iterations by making a good initial estimate of the distortion at each grid point
- Distortion calculation at a reduced number of a grid points, with quadratic interpolation used to calculate the distortion at the remaining grid points

By using the first two techniques, only two pierce point calculations are required per scan. The last two techniques significantly reduce the number of evaluations of the cross track distance expression. A computer program was developed which compared the combination of these four simplifications with inverse mapping of pierce points. It showed that at most 0.03 pixel error results.

Using these techniques, the distortion is calculated at grid points spaced once each 64 output pixels. The distortion is assumed to be the same for all sensors in each of the 16 lines, except for fixed delays associated with the time sequence at which the sensors are sampled. The distortion at the remaining pixels is performed using piecewise linear interpolation. The piecewise linear interpolation is extremely simple computationally; consequently, it can be implemented in special purpose hardware along with the resampling of the imaged data. This is extremely important, since the linear interpolation dominates the grid point calculation in terms of number of operations required.

Radiometric Correction

As mentioned previously, the Thematic Mapper produces sixteen image lines in each of seven spectral bands with each mirror scan. (Actually, the thermal infrared band produces only one-fourth as many lines.) Each of the simultaneously scanned lines is produced by a different photodiode. Ideally the response of each photosensor is linear so that its output is proportional to the intensity of the illumination in the specific spectral band. In practice these sensors do not respond linearly. In fact, each sensor has its own unique response curve that can vary gradually over a period of weeks or months.

It is necessary to correct the response to make it linear before performing the resampling operations used to accomplish geometric correction. This is because the resampling process requires interpolating intensity values between scanned lines. The different responses of the sensors cause discontinuities in the scanned image intensity from line to line. The sensor caused discontinuities between lines will produce incorrect interpolated values. Once the interpolated value is produced, compensation for the radiometric distortion is not possible.

Thus an essential part of the geometric correction process is an initial radiometric correction. This radiometric correction is accomplished as follows: The Thematic Mapper has a calibration procedure by which the response curve of individual detectors can be determined when requested from the ground. We propose to approximate these curve by piecewise linear functions. The breakpoints and slopes of the piecewise linear functions will be stored on the spacecraft. As each new sensor output value is produced, the value will be compared with piecewise linear function for that sensor to obtain a corrected intensity value.

Resampling

After the distortion has been estimated, the location of the pixel centerpoints of the Thematic Mapper imagery is known relative to the pixel centerpoints of the reference image. This is illustrated in Figure 4. The regular grid in solid lines represents the set of output pixels to be generated. Intersections in the grids represent the centerpoints of the individual pixels. The task of resampling is to calculate a set of intensity values for the output pixels, based on estimates derived from the intensity values of the input pixels plus calculated distortions. There are different resampling techniques, but all make use of the values of the input pixels in the vicinity of the output pixel to be calculated. This process is called interpolation.

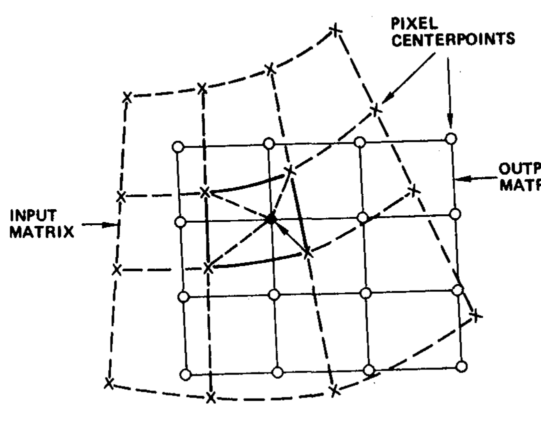

Figure 4. Resampling with Equally Spaced Output Matrix

Ideally, a two-dimensional bandpass process can be interpolated by passing the imagery through an ideal two-dimensional lowpass filter. The reconstructed image can then be "resampled" at the desired output pixel locations. This is mathematically equivalent to weighting the input pixels according to a two-dimensional (sin x)/x function (the impulse response of an ideal two-dimensional lowpass filter). The (sin x)/x interpolation requires an infinite number of points. However, practical interpolation is accomplished by approximating the (sin x)/x weighting with a relatively small number of input pixels.

Three resampling techniques are in common use today: nearest-neighbor, bilinear interpolation, and cubic convolution. In the nearest-neighbor procedure the value of the nearest input pixel to the desired output pixel is used as the value of that output pixel. Nearest-neighbor resampling is computationally simple, but generally produces distortions in the form of small discontinuities at the edges and borders in an image. It also results in an extremely blocky image.

Bilinear interpolation uses the values of the four pixels surrounding the output pixel to be calculated. The intensity of these pixels are bilinearly averaged to yield the intensity of the output pixels, with the relative weighting depending upon the location of the output pixel. The resulting averaging moves the blockiness of the nearest-neighbor technique but introduces small-scale smearing that results in loss of resolution.

The cubic convolution technique uses the values of the sixteen pixels surrounding the desired output pixel (Figure 4). The weighting function in this case is a two-dimensional cubic spline function which approximates the optimal (sin x)/x interpolator. The one-dimensional cubic spline interpolator (shown in Figure 5) is a piecewise cubic polynominal which is the same as (sin x)/x at the breakpoints and is required to be twice continuously differentiable at the breakpoints. Two-dimensional interpolation is accomplished by performing one-dimensional interpolation within each of the four closet rows to obtain four pixels vertically aligned with the desired output pixel. One-dimensional interpolation in the vertical direction is then performed to obtain the desired output pixel. Interchanging the rows and columns in this procedure yields the same result.

Cubic convolution does not suffer from the blockiness associated with nearest-neighbor interpolation or from the resolution difficulties which plague bilinear interpolution.

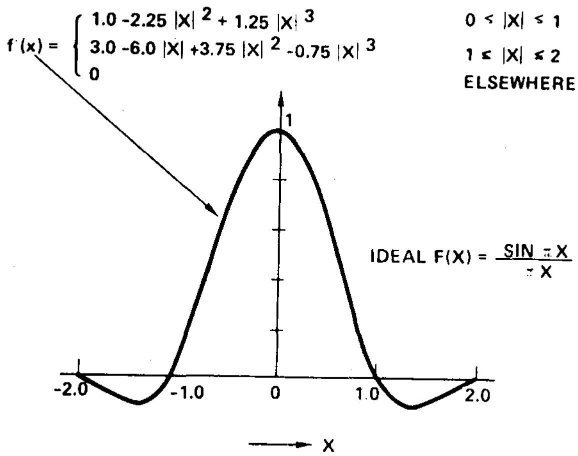

Figure 5. Cubic Spline Interpolation

VI. Implementation

The on-board image processor is functionally divided into two major units: a general purpose programmable processor, and a custom designed resampling processor. A functional block diagram of the entire system is shown in Figure 6. The general purpose processor calculates the recursive distortion coefficients required by the resampling processor and acts as the controller for the resampling processor. The resampling processor performs the along scan and across scan resampling algorithms. In order to perform this resampling, this processor must also perform radiometric correction and skew buffering.

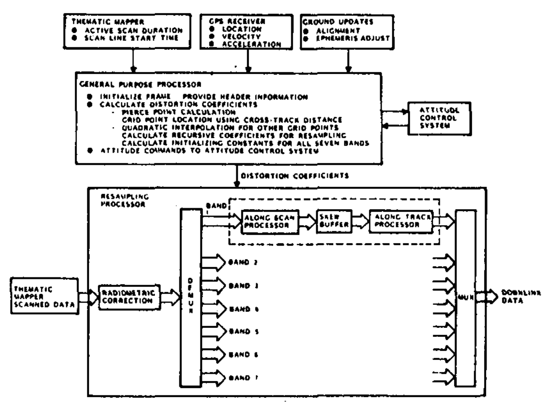

* DOUBLE LINES REPRESENT IMAGE DATA.

Figure 6.. On-Board Processing Functional Block Diagram

Computer for Distortion Calculation

During the time required for one scan (0.07 second) the general purpose processor must calculate the distortion coefficients and perform the required control functions for the resampling processor. Table 2 summarizes the number of operations required to perform this calculation. Since virtually all operations require 32 bit accuracy, Table 2 also shows how many single precision (16 bit) operations are required to achieve the required accuracy.

Table 2. Operations Required for Distortion Calculation

Processing Segment	Number of Operations (Double Precision)				
	Add	Multiply	Divide	Square Root	Trig
Calculate reference points	90	96	10	9	8
Evaluate cross track distance	220	400		20	
Quadratic interpolation	246	336			
Other along scan distortion calculations	202	3			
Cross scan distortion calculation	9	2	1		
Recursive equation initialization	280	140			
TOTAL	1,047	977	11	29	8
Total single precision add/multiply	2,534	4,341			
Total add/multiply per second with factor of two margin	70,952	121,548			

The first processor considered was the NASA Standard Spacecraft Computer - I. Unfortunately, this processor is approximately five times too slow to calculate the distortion coefficients.

The NASA Standard Spacecraft Computer - II was also considered. Its full parallel floating point structure reduces the double precision multiply time to 33.5 microseconds. Consequently, this computer may be capable of performing the distortion calculation, provided the factor of two margin is not required. Its power consumption (110 watts for 8192 words of core memory) is at least twice that required of a processor employing hardware multiplication.

It is estimated that a processor could be developed consuming approximately 30 to 35 watts which has the required capability. For example, a 16-bit version of the 8-bit Payload Signal Processor (PSP) built by TRW and described in Reference 2 would be in this range and would be capable of meeting the performance requirements. The 8-bit version of the PSP is to be space qualified by mid-1979.

This estimate is based on using a 4096 by 40-bit program control memory and a 2048 by 16-bit RAM working memory. Since the processor is not time constrained, extensive use of branching to "subroutines" can be used to keep the program within these limits. Each of these memories requires approximately 10 watts. Combining these memories with the 10 watts required for the CPU yields the 30 watt estimate.

Resampling Processor

The preliminary hardware sizing described in this section employs off the shelf components and is straightforward in design . It does not assume use of yaw control to reduce the number of scan lines stored. This possibility is discussed in the next sub-section. The total number of parts is estimated to be 810, with a total power consumption of 135 watts (20 percent margin is included). The board area is estimated to be 2.2 square feet without redundancy. (Since much of the resampling hardware is identical for each spectral band, reliability considerations will require far less than 100 percent redundancy.) By careful design and use of custom device fabrication, the power consumption might be reduced by a factor of two.

The resampling processor (Figure 7) employs seven separate along scan and cross scan processors, one set for each of the seven spectral bands. A "skew buffer" memory is used to interface the along scan and cross scan processors. It stores 32 scan lines of data (262144 bytes) in each of the six high resolution bands. A single radiometric correction processor precedes the seven along scan processors. There are two microsequencers, one holding the control code for the radiometric and along scan processors and the other holding the control code for the skew buffer and across scan processors. Both microsequencers drive a delay line so the processors for each band receive a delayed version of the same code. The input and output are loaded into high speed First-In-First-Out-Stacks (FIFOS) for the purpose of resynchronizing the data to the processor rate.

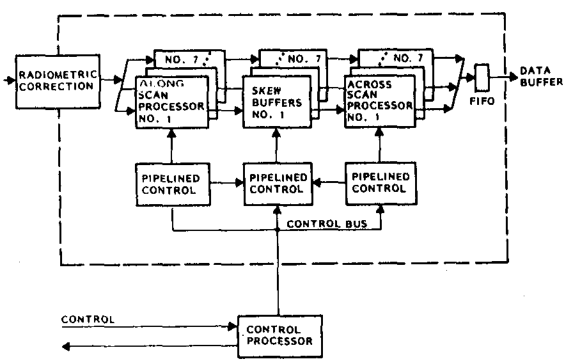

Figure 7. Resampler Block Diagram

The radiometric processor operates at 10 megasamples/second. The along scan and cross scan processors operate at 600 nsec per pixel with subcycles of 150 nsec. This is near the limit of their capability with presently available parts. The parts which limit the speed of the processors are the 150 nsec multiplier and the 64k memory chips in the skew buffer. It is anticipated that faster parts will be available in the near term which will increase the speed margin. In addition, a custom-designed multiply/accumulate chip might be employed to decrease the complexity of the processors. The parts and power could also be reduced by using an alternate memory configuration which saves approximately five scan lines of data instead of 32. This would require increased addressing complexity, but results in a factor of six reduction in skew buffer memory. This coupled with the multiply/accumulate chip could potentially reduce power by as much as one-half. The development cost may be greater, however. This discourages their use in a prototype ground version of the resampling processor.

Attitude Control

As described above, ephemeris variation results in scan lines being skewed with respect to the X-axis of the coordinate frame. One technique of compensating for this skew is by using yaw commands. Small, infrequent commands are capable of compensating for ephemeris caused skew. This skew is virtually zero at the equator and increases to as much as six pixels at high latitudes. However, the change in skew is approximately 0.3 pixel during an image frame with the amount of skew being consistent to within a fraction of a pixel at image frame boundaries. This corresponds to a yaw command of 100 μrad given once each 30 seconds, which is well within the capability of the attitude control system. The dynamics of the attitude control system are measured and compensated in the pierce point calculation, so the commands do not adversely affect the registration.

The amount the yaw should be changed is determined by observing the slope of the scan line at some consistent time within each image frame. This can be directly translated into an attitude command and passed to the multimission modular spacecraft computer for implementation. This calculation adds virtually no burden to the general purpose distortion calculation computer but can reduce the memory required in the resampler to six lines.

VII. Conclusions

We have shown that on-board correction of LANDSAT D imagery to subpixel accuracy is feasible using currently available technology. Specific methods to accomplish this goal have been described. Estimates of required size and power have been provided for both the special and general purpose hardware used. On-board realtime correction offers the potential of vastly increasing the percentage of images corrected and makes direct readout to users a valuable option.

REFERENCES

[1] "On-Board Image Registration Study", prepared by TRW Defense and Space Systems Group for Goddard Space Flight Center under Contract No. NAS5-23725, January 31, 1979.

[2] S. W. Houston, D. R. Martin, L. R. Stine, "Microprocessor Bit Synchronizer for Shuttle Payload Communications, IEEE Trans. Comm., Vol. COM-26, No. 11, Nov. 1978, pp. 1594-1603.