Title:

Automated Coregistration of MTI Spectral Bands

Author(s):

James Theiler, Amy Galbraith, Paul Pope,
Keri Ramsey, and John Szymanski

Submitted to:

<http://lib-www.lanl.gov/cgi-bin/getfile?00852365.pdf>

Automated Coregistration of MTI Spectral Bands

James Theiler, Amy Galbraith, Paul Pope, Keri Ramsey, and John Szymanski

Space and Remote Sensing Sciences, Los Alamos National Laboratory, Los Alamos, NM 87545

ABSTRACT

In the focal plane of a pushbroom imager, a linear array of pixels is scanned across the scene, building up the image one row at a time. For the Multispectral Thermal Imager (MTI), each of fifteen different spectral bands has its own linear array. These arrays are pushed across the scene together, but since each band's array is at a different position on the focal plane, a separate image is produced for each band.

The standard MTI data products (LEVEL1B_R_COREG and LEVEL1B_R_GEO) resample these separate images to a common grid and produce coregistered multispectral image cubes. The coregistration software employs a direct “dead reckoning” approach. Every pixel in the calibrated image is mapped to an absolute position on the surface of the earth, and these are resampled to produce an undistorted coregistered image of the scene. To do this requires extensive information regarding the satellite position and pointing as a function of time, the precise configuration of the focal plane, and the distortion due to the optics. These must be combined with knowledge about the position and altitude of the target on the rotating ellipsoidal earth.

We will discuss the direct approach to MTI coregistration, as well as more recent attempts to “tweak” the precision of the band-to-band registration using correlations in the imagery itself.

Keywords: MTI, multispectral, imagery, coregistration

1. INTRODUCTION

1.1. Multispectral Thermal Imager

The Multispectral Thermal Imager (MTI) is a U. S. Department of Energy research and development project whose purpose is to demonstrate the utility of advanced multispectral and thermal imaging from space.^{1,2} From its near-polar sun-synchronous 600 km altitude orbit, the MTI satellite records images in fifteen spectral bands, ranging from visible to long-wave infrared. MTI is a multispectral pushbroom sensor. Each spectral band has its own filtered linear arrays aligned in the cross-track direction. As the satellite travels, the response to the light incident on each element is integrated for a short time, and then read out, forming a two-dimensional image in each band.

In the MTI focal-plane, there are three sensor chip assemblies (SCAs), each of which contains a line of pixel detectors for each of the fifteen spectral bands. It is the job of the coregistration software to collate the output of these detectors into a spatially consistent multispectral image cube where a given location on the ground corresponds to the same position in the image for each of the spectral channels. Coregistration of the spectral bands on MTI is crucial for virtually any kind of retrieval that includes a spectral component in the signature of the desired quantity:³ vegetation,⁴ columnar water vapor,⁵ surface water temperature,⁶ volcanic mineral identification,⁷ *etc.*

The approach taken with the automated coregistration software has been to align the spectral bands to a common grid which is defined in terms of coordinates on the ground.⁸ Although this requires that the image be resampled, it enables us to remove distortions due to temporal nonuniformity in the motion of the imager and due to spatial nonuniformity in the optics.

The standard coreg product (LEVEL1B_R_COREG) is computed by a direct approach, based on mapping pixels in the focal plane back to their corresponding position on the surface of the rotating ellipsoidal earth. The ground positions of these pixels are then resampled onto a fixed grid to to produce an undistorted coregistered image of the scene. The three SCAs in the focal plane produce three side-by-side swaths of the scene with a small overlap between them. Aligning these SCA images to each other is another part of the coregistration task.

E-mail: {jt,amyg,papope,karamsey,szymanski}@lanl.gov

1.2. Alternative coregistration schemes

MTI coregistration has also been pursued with algorithms which do not attempt the direct mapping to a ground coordinate system. These algorithms do not employ attitude information at all, but attempt to use contrast features in the imagery itself to align the bands to each other. One of these is a translation-only registration developed at Sandia National Laboratory.⁹ In this registration, the user identifies target features and search areas, and the algorithm finds the position in the search area that best matches the target features. Resampling is avoided entirely; all adjustments are by whole pixels. Another approach is the unofficial LEVEL1B_R_SIR (“Simple Interactive Registration”) product that is occasionally used at Los Alamos. Here a user hits arrow keys and interactively adjusts the bands with respect to each other; these adjustments are purely translational, but can be smaller than a pixel. Because the motion of the spacecraft is relatively uniform, these approaches generally produce quite reasonable coregistrations; and because they do not depend on attitude information, they are more robust to glitches in the attitude reporting. Both, however, do require manual intervention. Two requirements for the standard coreg processing are that it be fully automated, and that it produce coregistered images even when there is little or no contrast in the scene.

2. COREGISTRATION

2.1. Coordinate systems

To project pixels from an orbiting satellite to the spinning earth, the time and space coordinates must be carefully defined.

Although time is a continuous quantity, the natural discretization for the coregistration problem is the “tick” – the time between pixel readouts. Although it is adjustable, the nominal time per tick is $715\mu\text{sec}$. We also define a particular time for each image, midway through the overflight, as a fiducial. The coreg product records characteristic quantities such as satellite latitude, longitude, and altitude at this fiducial time.

The importance of making precise and unambiguous definitions for the coordinate systems cannot be overemphasized. Further, having defined these coordinate systems, it is also crucial to reliably transform back and forth between them. Although an early version of our software adopted an Euler angle (roll-pitch-yaw) approach for parameterizing satellite attitude, we found that this produced too many ambiguities (among other things, as noted by Tandon,¹⁰ there are actually twelve distinct conventions for Euler angle coordinate systems). Calculations involving satellite attitude and coordinate conversions are performed using quaternions.^{11,12}

In the end, we defined seven separate spatial coordinate systems, and provided quaternions to transform between them.

- ECI: The earth-centered inertial coordinate system has its origin at the center of the earth, but is fixed with respect to the stars. The “zenith” vector $\hat{\mathbf{z}}$ points to the north star, and the $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ axes are aligned with the celestial equator. The satellite position and instantaneous velocity that are reported with the downlink telemetry are in ECI coordinates.
- ECEF: The earth-centered earth-fixed coordinate system also has its origin at the center of the earth, and $\hat{\mathbf{z}}$ still points to the north star, but the $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ axes follow the rotation of the earth, and correspond directly to terrestrial longitude. Latitude and altitude can be obtained from (x, y, z) in ECEF coordinates,¹³ using the WGS-84 ellipsoid¹⁴ to model the earth. The conversion between ECI and ECEF is a time dependent rotation around the $\hat{\mathbf{z}}$ axis, given by the quaternion

$$\mathbf{q}_{\text{ECI} \rightarrow \text{ECEF}} = [0, 0, \sin(\phi/2), \cos(\phi/2)] \quad (1)$$

where $\phi = \omega(t - t_o)$, and ω is the earth’s sidereal rotational rate. There is also a (very small) quadratic term to account for the earth’s slowing rotational rate, and a (small amplitude) oscillation due primarily to the varying distances of the earth to the sun and moon.¹⁵

- CF: The spacecraft control frame has its origin in the satellite and is the frame with respect to which the ground alignment and calibration measurements were made. The spacecraft orientation that is reported with the downlink telemetry is a quaternion that converts CF to ECI coordinates.

- CFLS: The spacecraft control frame line-of-sight is pitched approximately four degrees “behind” CF, and has a $\hat{\mathbf{z}}$ axis that corresponds to the centerline of the optical axis. The CF→CFLS boresight correction is given by a single fixed quaternion that was supplied to us by Charles Schira¹⁶:

$$\mathbf{q}_{\text{CF} \rightarrow \text{CFLS}} = [0, \sin(\theta/2), 0, \cos(\theta/2)] \quad (2)$$

where $\theta = 4.3878^\circ$.

This four degree pitch was quite deliberate. The idea was to point the CF $\hat{\mathbf{z}}$ axis at the target of interest; then when the spacecraft “settled down” a few seconds later, the CFLS $\hat{\mathbf{z}}$ axis would have the target in its view.

- FMLS: The focal map line-of-sight coordinate system is essentially the same as the CFLS except that the $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ axes are rotated 90° from CFLS, as different conventions were used by those designing the satellite platform and those characterizing the focal plane layout. The transform is given as

$$\mathbf{q}_{\text{CFLS} \rightarrow \text{FMLS}} = [0, 0, \sin(\pi/4), \cos(\pi/4)]. \quad (3)$$

- GND: There are two ground frames, and both have their origin at the target scene on the earth surface. In both cases, the zenith vector $\hat{\mathbf{z}}$ points directly away from the center of the earth, and the $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ axes are perpendicular to $\hat{\mathbf{z}}$ and therefore approximately tangent to the earth surface – not exactly tangent because the earth is an ellipsoid and not a sphere; see Section 3.2.1.
 - GEO: In the geo-aligned ground frame, $\hat{\mathbf{x}}$ is east and $\hat{\mathbf{y}}$ is north.
 - ORB: In the orbit-aligned ground frame $\hat{\mathbf{y}}$ is aligned to the direction of the satellite; it is always approximately northward with a positive alignment for daytime (when the satellite is moving generally northward), and a negative alignment for nighttime. The final coreg image is usually provided in the ORB ground frame coordinates.

2.2. Producing the coreg image

The coregistration of the different bands is produced indirectly, by registering each band separately to a common coordinate system on the surface of the earth. This “dead reckoning” approach requires accurate knowledge of satellite position, orientation, and timing. It also requires an accurate model of the earth’s surface and its phase and rate of rotation.

The ground coordinate system is a plane whose origin is a position at a fixed elevation h_t above sea-level. This elevation corresponds to the elevation of the ground at the position being imaged, and is obtained from a database. This plane is centered on the scene being imaged; originally, we used the plane tangent to earth surface at the subsatellite point, but we found that it produced too much distortion for all but direct nadir looks.

The first step involves the projection of the telescope boresight to the ground in order to define the ground coordinate system. With that in place, the pixels in the focalmap are projected down to the ground to produce an (x, y) position in ground coordinates for each pixel at each time step. The third step is to transform all the pixels and ground coordinates into a set of images (one image for each spectral band) on a common ground coordinate system. These three steps are described in more detail in the following sections.

2.2.1. Projecting boresight to the ground

The first step in computing the coreg image projects the boresight to the ground. The boresight is the vector that corresponds to a central reference location on the focal-plane as it is projected out the telescope (*i.e.*, the origin in Fig. 2). It serves as the centerline of the optical axis, and it defines $\hat{\mathbf{z}}$ in the focal map line-of-sight (FMLS) coordinate frame.

The boresight-to-ground projection is performed in the ECI coordinate frame, and in that frame the boresight is a unit vector given by $\mathbf{b} = \mathbf{q}_{\text{FMLS} \rightarrow \text{ECI}}(\hat{\mathbf{z}})$. If we define $\mathbf{R}_{\text{sat}}$ as the satellite position in ECI coordinates (so it is a vector whose origin is the center of the earth), and $\mathbf{R}_{\text{gnd}}$ as the position on the ground where the boresight intersects earth, then we have

$$\mathbf{R}_{\text{gnd}} = \mathbf{R}_{\text{sat}} + \ell \mathbf{b}, \quad (4)$$

Figure 1. This figure illustrates the projection of the boresight in the direction $\mathbf{b}$ from the satellite at $\mathbf{R}_{\text{sat}}$ to the target on the ground at $\mathbf{R}_{\text{gnd}}$ which has elevation h_t above sea-level. Here $\mathbf{R}_1$ and $\mathbf{R}_2$ are the first two iterations of Eq. (5). In the production code, we take two more iterations, and estimate $\mathbf{R}_{\text{gnd}}$ by $\mathbf{R}_4$.

where ℓ is distance along the length of line of sight. If we choose to project to a plane that is tangent to the nadir point of the satellite, we could use $\ell = h/\cos\theta$, but as described in Section 3.2.2, this is an unsatisfactory approximation for MTI. Instead we compute $\mathbf{R}_{\text{gnd}}$ iteratively, using

$$\mathbf{R}_{n+1} = \mathbf{R}_n + \frac{h_n - h_t}{\cos\theta_n} \mathbf{b} \quad (5)$$

starting with $\mathbf{R}_0 = \mathbf{R}_{\text{sat}}$. Here h_n is altitude above sea-level of the point $\mathbf{R}_n$, and h_t is the elevation above sea-level of the target. (By “sea-level” we mean the level of the WGS-84 ellipsoid.¹⁴) We compute $\cos\theta_n$ using

$$\cos\theta_n = \frac{|\mathbf{b} \cdot \mathbf{R}_n|}{\|\mathbf{R}_n\|} \quad (6)$$

See Fig. 1. Note that $\mathbf{R}_1$ corresponds to the flat earth approximation. We use four iterations, and have found that this produces sufficiently precise convergence.

Using this scheme, we can compute $\mathbf{R}_{\text{gnd}}$ as a function of time, during the exposure. The value of $\mathbf{R}_{\text{gnd}}$ at the fiducial midpoint time defines the origin of the ground coordinate frame. The zenith vector $\hat{\mathbf{z}}$ is parallel to $\mathbf{R}_{\text{gnd}}$ (that is, it points away from the earth center), and the $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ plane is perpendicular to $\hat{\mathbf{z}}$ with $\hat{\mathbf{y}}$ pointing either in the direction of the satellite motion (ORB) or to the north (GEO). With the ground frame thus defined, we recompute $\mathbf{R}_{\text{sat}}$, the satellite position in the ground frame as a function of time (in ticks), and $\mathbf{Q}_{\text{sat}}$, the satellite orientation as an array of quaternions (FMLS→GND) for each time (in ticks).

2.2.2. Projecting focal plane to the ground

Fig. 2 shows the x and y components of the unit vector $\mathbf{p}$ associated with each pixel. This is the vector that projects the pixel out of the telescope. The origin ($p_x = p_y = 0$) corresponds to the $\hat{\mathbf{z}}$ axis in the FMLS coordinate system.

While it is in principle possible to measure the vector $\mathbf{p}$ for every pixel, such a task would be arduous, and is not really necessary. Instead, physical measurements of pixel locations were made directly on the focal plane and then these measurements were propagated through the optics by an appropriate choice of Zernike polynomials. The polynomial coefficients were fit by measuring the actual vector $\mathbf{p}$ for a representative selection of pixels.

Figure 2. Focal plane layout, as it is projected onto the unit sphere; this includes optical distortions, etc. The bands, in order from the center outward are: DCBAHEF-GIOKJNML. Note that L-to-L distance is 0.0309 radians, A-to-D is 0.0026, A-to-A is 0.0089, L-to-E is 0.1450, and L-to-D 0.1494 radians.

So, given a vector $\mathbf{p}$ for a given pixel, we rotate it from the FMLS coordinate system to the ground coordinate system, using the element of the quaternion array $\mathbf{Q}_{\text{sat}}$ that is appropriate for the given time. If $\mathbf{b} = q_{\text{FMLS} \rightarrow \text{GND}}(\mathbf{p})$ is the vector $\mathbf{p}$ in the ground coordinate system, then we can project it to the ground using

$$\mathbf{r} = \mathbf{R}_{\text{sat}} - (R_{\text{sat},z}/b_z)\mathbf{b}, \quad (7)$$

where $\mathbf{r}$ is the projected position of the pixel in the ground frame (note that $r_z = 0$). This is done for every (active) pixel on the focal plane, for each time step (or “tick” – the time between successive focal plane readout times).

2.2.3. Resample pixels on the ground

For a pushbroom sensor, the sampling interval in the cross-track direction is determined by the pixel spacing on the focal plane, and the sampling interval in the along-track direction is determined by the ground-track speed and the pixel readout rate (time between ticks). For MTI near-nadir looks, these ground sample distances (GSD’s) are nominally designed to be 20m in the thermal bands, and 5m in the VNIR bands.

Based on the positions that are projected to the ground, the actual (average) GSD is computed for both the IR and the VNIR bands in both the cross and along-track directions. Based on these averages, a grid is defined on the ground with a pixel size that is roughly commensurate with the average GSD and is an even multiple of 5m. For the nadir looks, these nominal grid size is 5m for the four VNIR bands and 20m for the rest of the bands. The off-nadir looks generally have a sampling of 10m for VNIR and 40m for IR. These pixel sizes can also be set by the user. As a convenience, the coreg also supplies, as an auxiliary image cube, a resampling of the VNIR bands to the same pixel size as the IR bands.

All active bands and all three SCAs are resampled to the same grid, but in the SCA-overlap region, the middle SCA dominates the other two. The transition between SCA’s is sharp and occurs at the same place for all bands (to avoid “rainbow” effects at the boundary).

The resampling itself employs a local linear interpolation of the calibrated radiance values. This interpolation and resampling is always saved for last. Re-registering from an already registered image would introduce unwarranted blurring in the final image cube.

2.3. The “quality” or “pixelstate” image

In addition to the four-byte calibrated radiance imagery, a one-byte “quality” or “pixelstate” status image is provided which is coregistered to the actual image. These pixels provide information, on a pixel-by-pixel basis, about the estimated quality of the calibration in the pixel. Additionally, the first three bits in the image specifies which SCA the image data came from (remember the coreg product combines all SCAs into a single image). See Table 1.

Note that a pixelstate of zero corresponds to the parts of the scene that were not imaged. This is the black padding around the image as seen in Fig. 3.

Bit	Interpretation
1	SCA1
2	SCA2
4	SCA3
8	REDUCED_CONFIDENCE
16	SATURATED
32	INTERPOLATED
64	<i>unused</i>
128	<i>unused</i>

Table 1. Each one-byte pixel in the quality image contains bitwise information about where it came from and how it was calibrated. The information is OR'd; for example, a coreg pixel that was obtained by resampling from two calibrated pixels, one saturated and one with reduced confidence, and both from SCA3, would have a quality value of $16+8+4=28$.

2.4. The geo-aligned and georeferenced products

The coreg product can optionally provide imagery that is “geo-aligned”; this means that the grid to which the points are resampled is nominally a standard Universal Transverse Mercator (UTM) coordinate system. See Fig. 3. Although geo-aligned is more convenient (north is “up”), we generally produce orb-aligned coreg images for two reasons: it is more parsimonious with memory (less padding around the edges), and there are fewer resampling artifacts when the ground grid is aligned with the raw data grid.

Compared to the geo-aligned coreg product, the actual geo product (LEVEL1B_R_GEO) more accurately references the image into UTM coordinates. But it requires the manual identification of ground control points. By locating conjugate features in the image and determining (*e.g.*, from an accurate map) their UTM coordinates in the scene, one can fit an affine (*i.e.*, linear) transform that maps the initial ground coordinate grid into one whose position, orientation and scale are consistent with the control points. The affine coefficients are included in the header of the geo product, and the affine transformation is applied to the initially estimated ground positions for all of the pixels. These positions are resampled to the UTM grid to produce the final image in the geo product. A fuller description of georeferencing effort was provided at the 2001 MTI Symposium.¹⁷

Figure 3. Orb-aligned (left) and geo-aligned (right) coreg images for Albuquerque, NM. In the geo-aligned image, north is up, and this is evident from the grid layout of the city streets. Although pixel size is 5m in both images, the geo-aligned image is larger just because there is more padding around the edges.

Figure 4. The cross-correlation algorithm provides estimates of $(\Delta x, \Delta y)$ for many points (x, y) in the image. Shown here is a scatter-plot of the cross-track correction Δx versus the along-track position y . **(a)** The panel at the left shows a kink in the curve, which was traced to discretization in the quaternions obtained from the satellite telemetry. We also observed misregistration in the imagery which changed abruptly at the same y position as the kink. **(b)** The panel on the right shows the same scatter plot after smoothing the quaternion stream; see details in Section 3.3.

2.5. Tweaking based on cross-correlation

We have recently investigated an optional extra step in the coregistration which usually improves the quality of the coreg product. We tweak the estimated ground positions using cross-correlation of the images in the different spectral bands. In some ways this parallels the efforts to produce purely image based coregistrations (*e.g.*, the SIR product, or the translation-only registration), but we have as a starting point images that are already approximately registered and for which (some) distortions have been removed.

Although we carried out some pre-flight experiments¹⁸ for estimating a single time-dependent “jitter” function, we found in practice two things: the first is that the motion of the spacecraft is very smooth, and a high-frequency jitter is not evident; the second is that small errors in the focal-plane layout or optical distortion or spacecraft position (or, possibly due to unaccounted-for effects such as atmospheric refraction or relativistic velocity aberration) manifest themselves in band-to-band misregistration errors that cannot be fixed with a single jitter function.

Therefore we take the approach that a single two-dimensional $(\Delta x, \Delta y)$ translational adjustment is made to each band/SCA. These tweaks are included in the header of the coreg object and are applied before the final resampling. For each pair of bands, we produce a large number of small image “chips” and for each chip at position (x, y) we estimate the $(\Delta x, \Delta y)$ that produces the maximum cross-correlation.¹⁹ As Fig. 4 illustrates, we can test whether the estimated correction is consistent over the image and (if so) we can take medians for a robust estimate of the optimal shifts for each pair of bands. With N bands, we get $N(N - 1)$ estimated shifts, and we use least squares to estimate the best shift for each band. A similar approach is used to match up the SCAs.

The final step, after writing these tweaks to the coreg product, is to recompute pixel positions on the ground and resample to produce an image cube. We emphasize that the tweaked coreg product only involves one resampling of the calibrated data to the final grid.

2.6. Information in the coreg product

The coreg product can contain up to four image cubes: 1. the VNIR cube with high-resolution (nominally 5m) sampling of bands ABCD; 2. the IR cube with low-resolution (nominally 20m) sampling of bands EFGHIJKLMNO; 3. a VNIR cube with the same low-resolution as the IR cube (this has no new information, but makes the simultaneous exploitation of the visible and infrared bands more convenient); and 4. a top-of-atmosphere brightness temperature cube involving only the thermal IR bands JKLMN. Not all bands are active for all images, and if for instance bands ABCD are not active (as is often the case for a nighttime collect), then the VNIR cubes will not be provided. Quality images are provided for whichever of the data cubes are available (except we do not produce a quality image for the brightness temperature cube since that information is in the IR quality cube).

In addition to the data and quality images, the coreg product provides a wealth of auxiliary information about the overflight, and about the coregistration process. Table 2 provides a list of some of these variables and arrays. This information is also available in the header file that comes with the coreg product.

	<u>Information about the overflight</u>
image_id	Unique integer identifier for this image.
activeBands	Array of flags for each band/SCA indicating whether the band is active for this look.
gps.times	Array of times (generally every two seconds) that correspond the position and velocity values.
position	Array of xyz spacecraft coordinates (ECI, meters) as a function of time.
velocity	Array of spacecraft velocity vectors (ECI, meters/second) as a function of time.
quat.times	Array of times (generally every second) that correspond to the quaternions values.
quaternions	Array of spacecraft quaternions (ECI→CF) as a function of time.
	<u>User-specified information about the coregistration</u>
frame	Specifies ground coordinate frame ('orb' or 'geo') to which the data are resampled.
pixSize	Size of pixel, in meters on the ground, to which the IR image is resampled in the coregistered image; nominally 20m.
visPixSize	Size of pixel, in meters on the ground, to which the VNIR bands (ABCD) are resampled in the coregistered image; nominally 5m.
affine	Specifies the affine transform (an array of six floating point numbers) that was applied to the coreg data.
extent	Four numbers providing the coordinates of opposite corners of a rectangle that encloses the image. Note: if frame ='geo' then this will be approximate UTM coordinates. Although the extent can be specified by the user, it is more usual for the coregistration software to compute it as the smallest rectangle that encloses all the data.
	<u>Intermediate attitude determination information</u>
typical_GSD_cube_track	Ground sample distance (that is, distance between centers of adjacent pixels projected to the ground), in meters. <i>cube</i> can be VIS or IR; <i>track</i> can be cross_track or along_track .
typical_IFOV_cube_track	Instantaneous field of view for a pixel projected to the ground, in meters.
pointing	Array of xy coordinates (in the ground frame, units of meters) of position of boresight on the ground as a function of time in ticks.
Rsat	Array of xyz coordinates (in the ground frame, with units of kilometers) for the position of the satellite as a function of time in ticks.
Qsat	Array of quaternion values containing the orientation of the satellite (FMLS→ground) as a function of time in ticks.
dxtweaks,dytweaks	Array of tweaks for each band/SCA; these are by default zero, but if the tweaked cross-correlation has been applied, then those tweaks are written here.
	<u>Derived information about the overflight and coregistration</u>
daylight	Indicates whether the image was taken during daylight ('T') or nighttime ('F').
solar_angle	<i>angle</i> = elevation or azimuth ; degrees from target to the sun at the time of the exposure.
roll, pitch, yaw	Three scalar angles indicate average orientation of the satellite during the exposure.
altitude	Altitude of the spacecraft, in kilometers above sea-level.
grnd_speed	Average ground speed of the spacecraft (meters/second).
zenith2mti	Angle (degrees) between the target zenith and the line from target to satellite.
azimuth2mti	The direction, in degrees east of north, from the target to the subsatellite point.
boresight_latlon	Latitude/Longitude (degrees) of the point on the earth where the satellite points at a fiducial time midway through the image acquisition.
MTI_latlon	The latitude/longitude (degrees) of the spacecraft at the fiducial time midway through the image acquisition.
nadir2scene	Angle (degrees) between the satellite nadir and the line from the satellite to the target (image center), as measured from the satellite.
alongtrack_direction	Angle, in degrees clockwise from vertical, of the average alongtrack motion in the coregistered image. Note: for frame ='orb' images, this angle is very nearly zero.
crosstrack_direction	Angle, in degrees clockwise from vertical, of a line of detector pixels projected to the ground, and observed on the coregistered image.
scene_shear	crosstrack_direction - alongtrack_direction ; note: 90° is a square image.
image_orientation	Angle, in degrees clockwise from vertical, of the North-pointing vector on the coregistered image. Note: for frame ='geo' images, this angle is very nearly zero. For frame ='orb' images in the northern hemisphere, this angle is generally small and positive (of order ten degrees) in the daytime, and of order 170 degrees at night.
utm_latlon_zone	Specifies which UTM zone the image is in.

Table 2. Partial summary of non-image data in the coreg product

Figure 5. Velocity effect. Here $\Delta\theta$ is the angle between a pair of bands on the focal plane ('I' is a leading band, and 'f' is a following band). If the spacecraft altitude above the ground is h , then it will travel a distance $h\Delta\theta$ before the 'I' and 'f' bands line up on the ground. If the effective velocity of the spacecraft is v , then the time for the spacecraft to do this is $\Delta t = h\Delta\theta/v$. The coregistration will then line up the 'I' and 'f' bands with a time delay of Δt . However, if due to a miscalculation, the velocity is assumed to be v but is actually $v + \Delta v$, then lining up the bands using the Δt delay will lead to a misregistration error (in meters on the ground) of $\Delta v\Delta t = h\Delta\theta\Delta v/v$. For MTI, $h \approx 600\text{km}$ and $v \approx 7\text{km/sec}$, so for instance if $\Delta v = 22\text{m/sec}$ then the A-D band-to-band ($\Delta\theta = 0.0026$) misregistration would be 5m, or one high-resolution pixel.

Figure 6. Altitude effect. If the elevation above sea-level is in error by an amount ε , then the misregistration error will be $\Delta(\varepsilon \tan \theta) = \Delta\varepsilon \sec^2 \theta \Delta\theta$, where θ is the angle off-nadir. For instance, a nadir look with an altitude error of 1.9km would lead to a 5m error in the A-D band-to-band registration.

3. ESTIMATING THE EFFECT OF VARIOUS ERRORS

The “dead reckoning” approach to data coregistration requires accurate, and properly used, information about the satellite attitude and about the target location. Approximations and misestimates of this information are manifested as band-to-band and sca-to-sca misregistration. But some errors have larger effects than others, and it is important to quantify the magnitude of these effects.

3.1. Velocity effects

A number of attitude misestimates have the same effect as a velocity error. Fig. 5 shows how an error in the estimated velocity of the satellite (or, more specifically, the velocity of the satellite's boresight on the ground) can lead to a misregistration between bands. Table 3 shows quantitatively the relation between them. By the same token, small misregistration errors can be adjusted by manually choosing a Δv , and adding $\Delta v(t - t_o)$ to the ground positions computed in Eq. (7). The error in Fig. 7(a) uses this trick to produce the corrected image in Fig. 7(c).

3.2. Elevation errors

Fig. 6 illustrates the effect of a mis-specified elevation on coregistration, and the last column of Table 3 shows quantitatively how elevation errors affect coregistration errors. Fig. 7(a,b) provides an example of a misregistration that is fixed by using the correct target elevation.

Bands	Type of misregistration	$\Delta\theta$ (radians)	Pixel size	Velocity error Δv for one pixel error	Altitude error ε for one pixel error
A-D	band-to-band	0.0026	5m	22.4m/sec	1.92km
L-D	band-to-band	0.0149	20m	16.6m/sec	1.34km
L-E	band-to-band	0.0145	20m	16.1m/sec	1.38km
A-A	sca-to-sca	0.0089	5m	6.5m/sec	0.56km
L-L	sca-to-sca	0.0309	20m	7.5m/sec	0.65km

Table 3. Sensitivity of coregistration to errors in the attitude. Shown are the errors in velocity and in altitude (or target elevation) that lead to one pixel misregistration errors. From this table, it is clear that the sca-to-sca misregistration for band A is the most sensitive, while the band-to-band misregistration of the visible bands is least sensitive (the pixels are smaller, but the bands are much closer together). The altitude error in the table corresponds to nadir ($\theta = 0$) views. The ε that leads to a one-pixel error would be one quarter the size for a 60° off-nadir look.

Figure 7. All three panels show the coregistration image for visible band A at a point on the SCA1 and SCA2 seam. (a) Ground elevation was (incorrectly) assumed to be sea-level, and a misregistration results. (b) This image is of Albuquerque, NM, and by including the 1.6km elevation above sea-level in the coregistration process, a better coregistration results. (c) Again, the incorrect sea-level elevation is used, but a manually chosen $\Delta v = -25\text{m/sec}$ provides a corrected coreg image. Even though the actual error is due to an altitude mismatch, the Δv trick provides a way to fix a coreg error whose source has not been identified.

3.2.1. Ellipticity

Although we do account for the ellipticity of the earth in estimating height above sea-level, our ground coordinate systems (ORB and GEO) use as the zenith $\hat{\mathbf{z}}$ vector the ray emanating from the center of the earth. The discrepancy between this and the line that is perpendicular to the surface of the earth's ellipsoid is given by

$$\delta\theta = \frac{R_e - R_p}{R_p} \sin\theta \cos\theta \quad (8)$$

where R_e and R_p are equatorial and polar radii of the earth, and θ is the latitude. The effect is largest at $\theta = 45^\circ$, and with the WGS-84 ellipsoid, it is a $\delta\theta = 0.1^\circ$ effect. For a 12km image, this error can be interpreted as an elevation error that varies from zero (in the center of the image) to about 10m on the edges; quite negligible, and certainly smaller than any terrain effects that one would likely see.

Figure 8. Ellipticity of earth leads to discrepancy between two candidates for zenith angle: perpendicular to surface and radius from center. The effect is largest at mid-latitudes, and for the earth amounts to about a tenth of a degree.

3.2.2. Earth curvature effects

We use a plane as our map projection surface, but since the earth is curved, this can introduce errors. Our algorithm projects from the satellite to the plane projection surface. As Fig. 9 shows, this too can be treated as an elevation error with $\varepsilon \approx d^2/2R$. Here d is the distance from the center of the image and R is the radius of the earth. For our representative 12km image, this is about a 3m effect. Note that if we use an approximation in which the surface of the earth is represented as tangent to the actual surface at the subsatellite point, then for off-nadir images, we can have $d \sim 600\text{km}$ and the effect is 10^4 times larger.

Figure 9. Curvature of the earth leads to an effective elevation error ε that increases quadratically with the distance d away from the origin of the ground frame R_{gnd}

3.3. Discretization in satellite attitude

The spacecraft position and velocity are reported every two seconds, but due to quantization error (caused by an insufficient allocation of bits) the position is reported only to the nearest 0.25 kilometer. We initially computed spacecraft position versus time with a linear interpolation between reported spacecraft positions, but this led to quite large misregistration artifacts. Since we know that the spacecraft motion itself is quite smooth (to a very good approximation, it is in a Keplerian orbit around the earth), we use the reported positions and velocities to fit the motion to a local quadratic. Since the reported velocity is so much more precise than the position, the fit considers only the velocity for the linear and quadratic terms, and the (imprecise) position is used only for the constant term. See Fig. 10.

Figure 10. Effect of discretization on the position of the satellite. The open circles show reported position of the satellite in ECI coordinates; these numbers are reported only to the nearest quarter-kilometer, and the axes' tick-marks indicate this discretization, which for this overflight is more evident in the x direction. The short solid lines emanating from the open circles correspond to the velocity that is reported with the position. The heavy solid line indicates the satellite trajectory according to the quadratic fit to the reported data, with the squares representing the estimated positions at the times when position is reported (every two seconds). The diamonds (which in this case align exactly with the circles) correspond to the values in the quadratic fit, rounded to the nearest quarter kilometer.

Satellite orientation is reported every second (except for occasional glitches), and we need to interpolate these values to get orientation at every tick. A linear interpolation between quaternions involves a linear rotation about the axis of rotation defined by the quaternion, but the effect is nonlinear in each of the quaternion components. This linear interpolation leads to “kinks” in the attitude solution that can cause noticeable misregistration errors (see Fig. 4). Our solution is a fractional smoothing of the reported quaternions; we take a linear combination of the original solution and a smoothed solution, and we impose a criterion that our smoothing cannot change the quaternion by an amount that corresponds to more than 60m on the ground.

3.4. Other effects

Noerdlinger²⁰ has derived estimates of the effect of atmospheric refraction on pointing; the effect increases with off-nadir angle, and for 30 degrees, a refraction angle of 0.01 degrees is indicated. From a 600km altitude, this would lead to a pointing error of 104m, but in fact the effect is much smaller than this because virtually all of the refraction occurs in the first few kilometers above the surface of the earth. The actual effect is estimated to be 2.2m. At 60 degrees, the effect increases to about 18m, and is increasing at a rate of about 2m per degree. For the visible A-D bands, the effect is about 0.3m (about 6% of a pixel) for band-to-band misregistration. The effect is wavelength dependent; it is somewhat smaller for longer wavelengths.

The earth rotates roughly 0.9m during the time it takes light to travel from the earth to the MTI satellite. This is a small effect, but the effect on misregistration is miniscule, because that depends on the difference in light-travel time for two bands. For the A-D band-to-band it is 0.0026 of that, or about 0.05% of a pixel.

Feynman²¹ describes a velocity aberration phenomenon which is also due to the finite speed of light. The ratio v/c , where v is the relative velocity of the satellite and the earth surface, defines an aberration angle $\sin \delta\theta = v/c$. A nadir angle in the frame of the earth will appear as $\delta\theta$ off-nadir to the satellite. For MTI, $\delta\theta = 2.3 \times 10^{-5}$ radians, and at altitude $h = 600$ km, this leads to a 14m effect in absolute pointing error. The effect diminishes with larger off-nadir lookangles, and the effect on misregistration depends on the rate of change with angle. For the angle differences $\Delta\theta$ shown in Table 3, the difference works out well under a tenth of a pixel for MTI.

4. MEMORY CONSERVATION EFFORTS

The size of the files (in HDF format²²) that contain the MTI image data products are generally quite large. A typical calibrated image product is 160MB, and a typical coreg product is 280MB. For larger images, those numbers can be multiplied by a factor of 4 or even 8. Since the process of making a coreg product involves both products, as well as a number of temporary arrays for intermediate results, the memory usage for the coreg process can be quite demanding. We have taken a number of steps to keep the memory usage in check. For instance, even though the VIS and IR data cubes are stored in the HDF files as large three-dimensional arrays, we altered our image reading and writing code to permit the reading and writing of single bands at a time.

Figure 11. (a) This figure shows memory use as the coregistration algorithm progresses – the horizontal axis can be thought of informally as a kind of “time,” though all it is really showing is sequential order. The two curves illustrate the difference between runs with (lower) and without (upper) the use of IDL’s `assoc` arrays. The three bumps on each curve correspond to the visible bands on each of the three SCAs; since those bands have 5m pixels (compared to the 20m pixels for the IR bands), resampling those bands is more memory intensive. (b) The difference between `assoc` and standard runs is still evident when running the algorithm with tweaking. The first three bumps correspond to separate runs with each SCA; the SCA-to-SCA tweaking that is done after those runs is done with only one band at a time and is not memory intensive at all. The last three bumps then correspond to the usual `coreg` algorithm but incorporating the tweaks that were found in the cross-correlation of the individual SCAs. (c) This run of a much larger image was only possible when the `assoc` arrays were implemented. This run was completed on a computer with 1GByte of memory and 1GByte of swap space.

In Fig. 11, we show the utility of IDL’s `assoc` arrays for keeping memory use under control. An `assoc` array is an “association” between an IDL array and a file – to the IDL code, it looks like an array, but the data is stored in a file. Separate `assoc` arrays were made for the IR and VIS images, one for each of the SCAs for the calibrated data, and one larger array that combined all SCAs for the `coreg` data. We found that this would produce a factor of three improvement in the memory requirements for typical images, and produce a `coreg` product for images of unusual size that would be impossibly large without the `assoc` arrays.

The imagery is archived in both the calibrated and the coregistered image products as thirty-two-bit floating point, although a case could be made for sixteen-bit integer with image-specific scale and offset. The latter approach would provide further gains in the memory conservation efforts, though at the cost of added complexity to the software.

ACKNOWLEDGEMENTS

We have benefitted from the support and advice of the entire MTI team, but we particularly wish to acknowledge our colleagues Bill Clodius (for many thoughtful comments and enlightening discussions, especially on the sources and magnitudes of misregistration errors), Cindy Little (for help with understanding the layout of the focal-plane), Brad Henderson (for various discussions), and Max Decker (for help in understanding the information in the telemetry). This work was supported by the U.S. Department of Energy.

REFERENCES

1. P. G. Weber, B. C. Brock, A. J. Garrett, B. W. Smith, C. C. Borel, W. B. Clodius, S. C. Bender, R. R. Kay, and M. L. Decker, “Multispectral Thermal Imager mission overview,” *Proc SPIE* **3750**, pp. 340–346, 1999.
2. J. J. Szymanski, W. Atkins, L. Balick, C. C. Borel, W. B. Clodius, W. Christensen, A. B. Davis, J. C. Echohawk, A. Galbraith, K. Hirsch, J. B. Krone, C. Little, P. Maclachlan, A. Morrison, K. Pollock, P. Pope, C. Novak, K. Ramsey, E. Riddle, C. Rohde, D. Roussel-Dupré, B. W. Smith, K. Smith, K. Starkovich, J. Theiler, and P. G. Weber, “MTI science, data products and ground data processing overview,” *Proc SPIE* **4381**, pp. 195–203, 2001.
3. C. C. Borel, W. B. Clodius, A. B. Davis, B. W. Smith, J. J. Szymanski, J. Theiler, P. V. Villeneuve, and P. G. Weber, “MTI core science retrieval algorithms,” *Proc. SPIE* **3753**, pp. 403–415, 1999.

4. P. V. Villeneuve and C. C. Borel, "Vegetation analysis," in *Handbook of Science Algorithms for the Multispectral Thermal Imager*, B. W. Smith, ed., pp. 114–121, Los Alamos National Laboratory Technical Report: LA-UR-98-306, 1998.
5. C. C. Borel and W. B. Clodius, "Columnar water vapor retrieval," in *Handbook of Science Algorithms for the Multispectral Thermal Imager*, B. W. Smith, ed., pp. 106–113, Los Alamos National Laboratory Technical Report: LA-UR-98-306, 1998.
6. C. C. Borel, W. B. Clodius, J. J. Szymanski, and J. Theiler, "Comparing robust and physics-based sea surface temperature retrievals for high resolution, multi-spectral thermal sensors using one or multiple looks," *Proc. SPIE* **3717**, pp. 58–71, 1999.
7. N. R. Harvey, J. Theiler, L. Balick, P. Pope, J. J. Szymanski, S. J. Perkins, R. B. Porter, S. P. Brumby, J. J. Bloch, and N. A. David, "Automated simultaneous multiple feature classification of MTI data," *Proc. SPIE* **4725**, 2002. In this volume.
8. J. Theiler and B. W. Smith, "Interband registration," in *Handbook of Science Algorithms for the Multispectral Thermal Imager*, B. W. Smith, ed., pp. 35–40, Los Alamos National Laboratory Technical Report: LA-UR-98-306, 1998.
9. J. Smith, S. Motomatsu, J. G. Taylor, K. J. Jefferson, and B. R. Stallard, "MTI Translation-only Registration," in *Multispectral Thermal Imager (MTI) Symposium*, U.S. Department of Energy Office of Nonproliferation and National Security, Albuquerque, NM, 12-13 March 2001.
10. G. K. Tandon, "Coordinate transforms," in *Spacecraft Attitude Determination and Control*, J. R. Wertz, ed., pp. 760–766, Kluwer Academic Publishers, 1978.
11. F. L. Markley, "Parameterization of the attitude," in *Spacecraft Attitude Determination and Control*, J. R. Wertz, ed., pp. 410–420, Kluwer Academic Publishers, 1978.
12. L. Fallon, III, "Quaternions," in *Spacecraft Attitude Determination and Control*, J. R. Wertz, ed., pp. 758–759, Kluwer Academic Publishers, 1978.
13. J. Zhu, "Exact conversion of earth-centered, earth-fixed coordinates to geodetic coordinates," *Journal of Guidance and Control* **16**, pp. 389–391, 1992.
14. "Department of Defense World Geodetic System 1984: Its definition and relationships with local geodetic systems," Tech. Rep. TR8350.2, NIMA, 4 July 1997.
15. C. R. Sturch, "Time measurement systems," in *Spacecraft Attitude Determination and Control*, J. R. Wertz, ed., pp. 798–806, Kluwer Academic Publishers, 1978.
16. C. Schira, 29 March 2000. Personal (Email) Correspondence.
17. P. Pope and J. Theiler, "Coregistration and Georeferencing," in *Multispectral Thermal Imager (MTI) Symposium*, U.S. Department of Energy Office of Nonproliferation and National Security, Albuquerque, NM, 12-13 March 2001.
18. J. Theiler, B. G. Henderson, and B. W. Smith, "Algorithms using inter-band cross-correlation for pixel registration and jitter reconstruction in multi-channel push-broom imagers," *Proc. SPIE* **3163**, pp. 22–32, 1997.
19. J. Gao and M. B. Lythe, "The maximum cross-correlational approach to detecting translational motions from sequential remote-sensing images," *Computers and Geosciences* **22**, pp. 525–534, 1996.
20. P. D. Noerdlinger, "Atmospheric refraction effects in Earth remote sensing," *ISPRS Journal of Photogrammetry and Remote Sensing* **54**, pp. 360–373, 1999.
21. R. P. Feynman, R. B. Leighton, and M. Sands, *The Feynman Lectures on Physics*, Addison Wesley, Reading, MA, 1963.
22. National Center for Supercomputing Applications, "HDF: Hierarchical Data Format." <http://hdf.ncsa.uiuc.edu>.