Available online at www.sciencedirect.com
ScienceDirect
Remote Sensing of Environment 107 (2007) 22–32
Remote Sensing of Environment
www.elsevier.com/locate/rse

Multi-angle geometric processing for globally geo-located and co-registered MISR image data

Veljko Jovanovic *, Catherine Moroney, David Nelson

Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109, USA

Received 6 March 2006; received in revised form 23 August 2006; accepted 27 August 2006

Abstract

This paper focuses on the high accuracy geo-location and co-registration performance operationally obtained for the Multi-angle Imaging SpectroRadiometer (MISR) data via a robust production algorithm. The MISR instrument (which is part of the payload for NASA's Terra spacecraft) continuously acquires systematic, global, multi-angle imagery in reflected sunlight since it was launched in December 1999, with the objective of supporting ecology and climate studies. The moderate resolution data need to be autonomously geo-rectified prior to being used in subsequent scientific retrievals. This is particularly critical to the unique cloud height/wind retrievals which require sub-pixel co-registration accuracies on a global basis. In order to address this problem, the MISR ground data processing system is based on a fully photogrammetric approach and is complemented with a quality monitoring system used to verify performance globally and over desired time periods. The geo-rectification requirements were developed from a series of data quality investigations based partially on sensitivity requirements for the cloud height/wind retrieval algorithm. An update to the nominal production algorithm has been implemented, and the final operational results are presented in this paper.

Keywords: Calibration; Geometric; Global; Mapping; Orthorectification; Cloud motion winds; Photogrammetry

1. Introduction

Many modern Earth observing sensors are designed to acquire remote sensing data on a global basis for a nominal life time of five to ten years. The Multi-angle Imaging SpectroRadiometer (MISR) (Diner et al., 1998) is part of the payload for the Terra spacecraft, (one of NASA's Earth Observing System (EOS) satellites) launched in December 1999. It continuously provides global, multi-angle, and multi-spectral imagery in reflected sunlight using nine separate push-broom cameras observing the Earth at nine discrete angles up to 70.5° relative to the local vertical and in four spectral bands. The instrument measurements are designed to improve studies of the Earth's ecology, environment and climate.

Powerful science data processing systems have been implemented to keep up with the high data acquisition rate, and to provide the user community with retrieved and validated parameters in a timely manner. These operational systems use autonomous algorithms devised to convert raw instrument data into calibrated and geo-located measurements and subsequently into globally gridded maps of science data products. MISR production algorithms (Bothwell et al., 2002) include an integrated, digital photogrammetric approach with an emphasis on the ancillary datasets generated as part of in-flight geometric calibrations. The initial in-flight calibration datasets were produced during the first year of the mission and the corresponding operational results have been published (Jovanovic et al., 2002). A full set of in-flight calibration datasets has subsequently been completed and is included in standard production operations.

In addition, a geometric quality monitoring system was implemented as the complement to the production system, in order to verify its performance on a global basis over specific time periods. As it happened, this system was particularly effective in regard to identifying a problem with the pointing stability of one of the most oblique cameras, and the performance verification of the other cameras, as they relate to the requirements of the MISR wind and height retrievals.

* Corresponding author.
E-mail address: Veljko.Jovanovic@jpl.nasa.gov (V. Jovanovic).

0034-4257/$ - see front matter © 2006 Elsevier Inc. All rights reserved.
doi:10.1016/j.rse.2006.08.013

This paper begins with an overview of the baseline geo-rectification approach including the description of its complementary quality monitoring system. Section 3 presents a discussion of initial operational results prior to the update. Sections 4 and 5 deal with the update to the baseline approach as motivated primarily by the sensitivity of cloud height/wind retrievals in addition to unstable performance of one of the nine cameras. Section 6 includes results corresponding to the final implementation which has been operational since May 2005, with the goals of assuring the desired accuracy and providing confident data quality indicators as part of the product.

2. Nominal geo-rectification and quality control

2.1. Data acquisition

The Terra spacecraft is in a sun-synchronous orbit, with a baseline inclination of 98.186°. The orbit period of 98.88 min and orbit precession rate of 0.986°/day imply a ground repeat cycle of the spacecraft nadir point of 16 days with an equatorial local crossing time of 10:30 AM. From the orbit altitude of about 705 km, the zonal overlap swath-width of MISR images (i.e., the swath seen by all nine cameras simultaneously along a line of constant latitude) is nominally 380 km providing multi-angle coverage of the Earth in nine days at the equator and 2 days near the poles. The data in 36 channels (9 cameras times 4 spectral bands per camera) are continuously acquired, pole-to-pole, on the day side of the orbit. The nine push-broom cameras are arranged with one nadir-pointing camera (designated An), one bank of four cameras pointing in the forward direction (designated Af, Bf, Cf, and Df in order of increasing off-nadir angle), and one bank of four cameras pointing in the backward direction (using the same convention, but designated Aa, Ba, Ca, and Da). Images are acquired with nominal view angles, relative to the surface reference ellipsoid, of 0°, 26.1°, 45.6°, 60.0°, and 70.5° for An, Af/Aa, Bf/Ba, Cf/Ca, and Df/Da, respectively. The instantaneous displacement in the along-track direction between the Df and Da views is about 2800 km, requiring about 7 min for a ground target to be observed by all nine cameras.

2.2. Required product

In order to satisfy co-registration and geo-location requirements, multi-angle, multi-spectral data are processed onto a common map projection. Space Oblique Mercator has been selected as the reference map projection grid because it was specifically designed for continuous mapping of satellite imagery (Snyder, 1987). The ground resolution of the map grid is 275 m for all cameras. This segment of ground processing chain has been defined as “geo-rectification” and the derived product as the Geo-rectified Radiance Product (GRP). There are two basic parameters in the Geo-rectified Radiance Product depending on the definition of the reference surface: a) ellipsoid-projected radiance, and b) terrain-projected radiance. The ellipsoid-projected radiance is referenced to the surface of the WGS84 ellipsoid (no terrain elevation included) and the terrain-projected radiance is referenced to the same datum including a digital elevation model over land and inland water. The spatial horizontal accuracy requirements are driven by the needs of the geophysical parameters’ retrieval algorithms, especially those designed to simultaneously derive cloud-top heights and winds (Moroney et al., 2002; Zong et al., 2002). The initial goal is to achieve an uncertainty better than ±140 m (1σ) regarding both the absolute geo-location of nadir camera and co-registration between all nine cameras. The 140 m corresponds to about half the size of the 1) instrument ground sampling distance, 2) nadir camera instantaneous field of view, and 3) GRP ground resolution.

2.3. Autonomous geo-rectification and quality monitoring

In order to assure that we meet our specific accuracy requirements and to keep up with the high data acquisition rate, we adopted a production strategy that distributes the overall effort into three main segments (Jovanovic et al., 1998): 1) in-flight geometric calibration, 2) standard production (geo-rectification), and 3) geometric quality monitoring (see Fig. 1).

2.3.1. In-flight geometric calibration

In-flight geometric calibration activities have been designed to produce specialized datasets, which are then used as inputs to standard production. These datasets not only reduce the overall processing load but also assure the required geo-rectification accuracy. In particular, the Camera Geometric Model (CGM), Reference Orbit Imagery (ROI), and projection parameters (PP) provide facilities to take into account errors in the camera pointing geometry including errors in the spacecraft supplied navigation and attitude.

The CGM dataset is designed to deal with static pointing errors. It consists of a set of parameters used in a mathematical expression that gives the pointing direction of an arbitrary pixel in the spacecraft attitude frame of reference. These parameters represent the geometry of the camera system and account for distortions from an ideal optical system. The calibrated CGM is not sufficiently accurate to reach the required product accuracy or to provide a means for on-line geo-rectification quality assessment. This is especially true in the case of the most oblique angles where a pointing error of 10 arc sec will introduce a geo-location error of about 300 m.

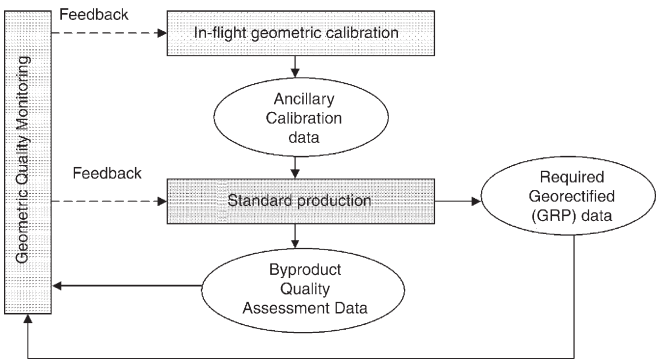

Fig. 1. Overview of MISR autonomous geo-rectification and quality monitoring.

In order to routinely deal with dynamic pointing errors and to facilitate automatic quality assessment, 233 pairs of PP and ROI files were produced, one pair for each of the 233 unique MISR paths. A ROI file consists of cloud free MISR imagery, selected from a number of orbit passes over the same orbit path and mosaicked into a single image. The PP file is produced, using rigorous photogrammetric methods, to provide accurate geo-location data for the corresponding ROI file pair. The process of creating ROI and PP pairs is similar to regular orthorectification of time dependent imagery. A major difference is that the acquired imagery (ROI) is geo-located through PP but not re-sampled. A simultaneous bundle adjustment utilizing multi-angle imagery and ground control information, consisting of a global Digital Elevation Model (DEM) and ground control image chips, is used to model dynamic errors in the supplied spacecraft navigation data. As of October 2002, the final ROI dataset was included in standard processing, providing a global high accuracy ground truth dataset with regards to the overall geo-rectification process. It should be noted that nominal horizontal resolution of the global DEM is 3 arc sec (approximately 100 m in mid-latitudes) with varying height accuracies across the globe (Logan, 1999). These errors will have some impact on the final geo-location and co-registration accuracy. The geo-location and quality monitoring algorithms are using the nadir (An) camera as the reference to take advantage of its low sensitivity to the topography errors. Thus, the absolute geo-location error will be present in the nadir imagery with size depending on the quality of DEM; an error in surface height of ±150 m will cause nadir geo-location error of about 1 m along-track and 45 m across-track. All of the geo-location errors for the other eight cameras are measured and presented in the following sections as the co-registration errors between that particular camera and nadir.

2.3.2. Standard processing

Given the geometric calibration datasets as an input, the geo-rectification process during standard processing has been significantly simplified. In particular, geo-rectification is facilitated by image-to-image registration between new MISR imagery and reference imagery. It has been demonstrated via operational data that this process is robust, since the registration is occurring between images with the same viewing geometry. An image point intersection algorithm is employed; using backward projection based on the camera model and supplied navigation, in order to obtain an initial guess for the tie points to be used during registration (Jovanovic et al., 1998). Final location of the tie points, prior to re-sampling, is obtained through least-square area-based matching. The terrain-projected radiance product generated during geo-rectification is used as the input to Level 2 Aerosol/Surface retrievals and cloud mask generation. Another part of the geo-rectified product, ellipsoid-projected radiance, is used for Level 2 Top-of-Atmosphere/Cloud stereoscopic retrievals.

2.3.3. Geometric quality monitoring

Implemented as an off-line activity, geometric quality monitoring is based on utilization of quality measurement data obtained from two sources: 1) a direct a-posteriori assessment of the final product, and 2) information by-products of standard production. In the first case, an automatic point measurement algorithm (Ackerman, 1984; Forstner, 1987) which uses a combination of interest point extraction and least-square image-matching to localize high fidelity multi-image conjugate points is applied to the final Geo-rectified Radiance Product. It provides a systematic and global measure of the overall co-registration with reference to nadir imagery. In the second case, inclusion of the Reference Orbit Imagery into standard processing created an opportunity to begin a systematic and permanent monitoring of the pointing stability of all nine MISR cameras. In particular, the performance of the image-to-image registration between the ROI and the newly acquired images provides a basis for evaluating the so-called Geometric Data Quality Indicator (GDQI). This GDQI is stored within the final product as well as in separate Quality Assessment (QA) files. Additionally, QA files contain Image Coordinate Corrections (ICC) representing two dimensional transformations between reference and new images as obtained for every X number of lines or Y seconds worth of data.

Due to their high coverage and reliability, both the co-registration and ICC data have been invaluable in detecting and resolving localized spatial accuracy issues as well as permitting long term assessment of data quality prior to deciding on updates to standard production algorithms. In fact, an improved version of the co-registration measurement algorithm has been folded into standard processing in addition to remaining a part of a-posteriori quality assessment. The GDQI algorithm was not as effective as originally designed, because it incorrectly assumed that pointing stability performance would be similar for all nine cameras. However, this was not the case (see Section 3) and a newly devised GDQI was included as part of the final update as described later.

3. Initial operational results

The first operational results that addressed geo-rectification and co-registration requirements were related to the performance of the Camera Geometric Model (CGM) with its pre-launch calibration (Jovanovic et al., 2002). Geo-location errors measured over a globally distributed set of Ground Control Point (Bailey et al., 1997) data were used as the input to the calibration algorithm. After several iterations, a final estimate of the CGM parameters was generated and included into operations on April 2002. The intent of the final delivery was to include static estimates of the overall pointing with the best fit to the data available up to that time. To further deal with any dynamic pointing changes within an orbit or eventual static changes in the overall pointing, we created final ROI files which were included in standard production on October 2002. Once this final delivery of the ancillary datasets was made, focus shifted to completion of the quality monitoring system along with analysis of the geo-rectification performances which are presented in subsequent sections.

The principal objectives of the quality assessment measurements are to evaluate: a) pointing stability of all nine cameras, b) geo-location and co-registration errors in the final product,

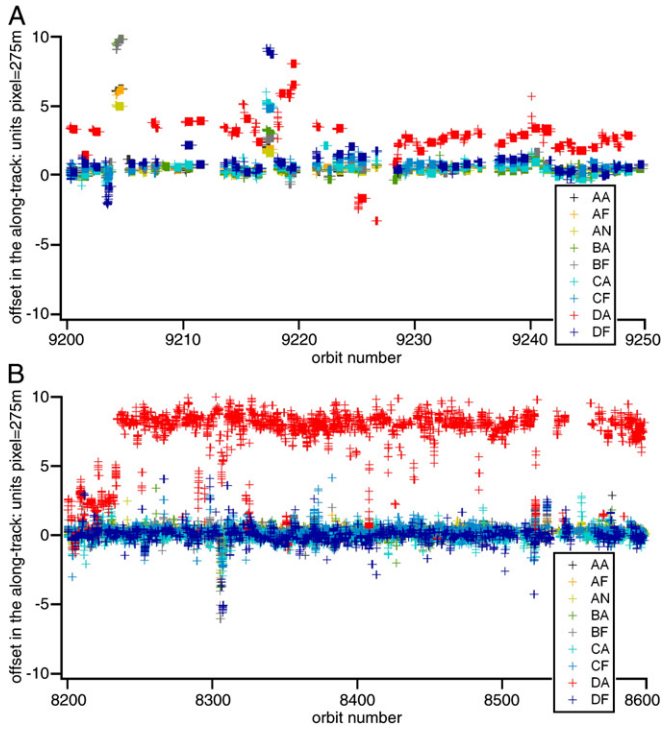

Fig. 2. A: Pointing stability as evaluated in the along-track line direction for nine MISR cameras: ICC (offsets in pixel units, 1 pixel=275 m) plotted against orbit number corresponding to approximately 4 days of data, Sept. 10–14, 2001. B: Pointing stability as evaluated in the along-track line direction for nine MISR cameras: ICC (offsets in pixel units, 1 pixel=275 m) plotted against orbit number corresponding to approximately 25 days of data, July 4–29, 2001.

and c) reliability of the Geometric Data Quality Indicator (GDQI) in the final product. Pointing stability and errors in the final product on a local and global basis will be presented in the remainder of this section. The initial GDQI values were not adequate and are not presented here; the results of the final, updated GDQI implementation are presented in Section 6.

3.1. Pointing stability

The Reference Orbit Imagery (ROI) provides a stable "ground truth" and is used as an integral part of standard production. The nine MISR camera images are compared to the ROI's in order to estimate coefficients of a two dimensional transformation describing the fit between images for a predefined length of a labeled image block. These coefficients, called Image Coordinate Corrections (ICC), were subsequently used in the processing software to optimally "warp" new images to match the reference. The ICC was also collected and summarized in order to give insight into pointing stability over desired time periods during the mission. In the following figures, we report the measure of pointing stability by plotting the ICC's in the along-track (line) direction for every block. The ICC's in the across-track direction were also collected and they are, for nadir and near nadir cameras, of comparable magnitude to the along-track ICC's. They are slightly smaller for the more oblique cameras. Also, impact of these errors on the accuracy of subsequent stereo cloud wind/height retrieval is significantly less dominant due to the fact that stereo parallax is nearly completely aligned with the along-track direction. Nevertheless, mean and rms of ICC's in both directions are given in the table complementing the plots. The ICC's are reported in pixel units, where a pixel is equal to 275 m, which is also the ground sampling distance of the MISR cameras. As an illustration of pointing stability, Fig. 2A and B include plots of ICC's corresponding to two time periods defined by the orbit number range on the *x*-axis. As can be seen from the plot, most of the time eight out of nine cameras are very close to the reference with occasional deviations corresponding to spacecraft orbit maneuver activities. However, the Da camera exhibits an irregular pointing change significantly larger than the other cameras. Corresponding to Fig. 2A and B is Table 1 with mean and standard deviations of ICC's as obtained for 5000 orbits closely corresponding to a one year time period. Again, these are as expected for all but the Da camera, revealing the small biases of the static camera model as well remaining dynamic errors in the supplied spacecraft attitude. The analysis regarding whether these measured offsets are sufficient to assure final geo-location and co-registration accuracy for all cameras including Da is addressed in the next section.

3.2. Geo-location and co-registration accuracy

For the purpose of validation, the final geo-located products were scanned with a version of the image-to-image-matching program (CameraMatch) that evaluated and reported co-registration errors for each camera relative to the nadir (An) using 275 m resolution, red band images. The algorithm performed numerous image-matching operations (see Section 5.1) and computed the mean, fractional pixel offset errors per fixed image segment size of 512 lines (called a block) for both along-track and across-track directions. In the following figures, we plot only errors in the along-track direction. Means and rms of geo-location and co-registration errors in both directions are given in Table 2 complementing the plots. As can be seen in Fig. 3A and B in which errors are displayed for exactly the same orbit ranges as in Fig. 2A and B, most of the static pointing biases are removed and the rms errors are within the limits of the automatic image-matching algorithm. However, it is clear that nominal implementation of the standard production algorithm was inadequate to fully take into account the pointing instability of the Da camera. Improvement certainly has been made but overall accuracy of Da camera data is still notably worse than the accuracy of data from other eight cameras. The slightly larger than expected mean error of the Bf camera (Table 2) should also be noted.

Table 1
Mean and rms of ICC (pointing stability given in pixel units, 1 pixel=275 m) obtained from 5000 orbits during period from December 2000 trough November 2001

Cam		AA	AF	AN	BA	BF	CA	CF	DA	DF
Line	Mean	0.24	0.25	0.20	0.17	0.24	0.03	0.34	1.5	0.13
	Rms	0.38	0.40	0.36	0.50	0.50	0.71	0.84	2.98	1.26
Sample	Mean	0.26	0.21	0.32	0.11	0.09	0.22	0.18	0.51	0.19
	Rms	0.50	0.43	0.42	0.32	0.40	0.43	0.45	0.75	0.48

Table 2
Geo-location and co-registration mean and rms errors (relative to nadir An camera and given in pixel units, 1 pixel=275 m) obtained from 5000 orbits during period from December 2000 trough November 2001

Cam		AA	AF	AN	BA	BF	CA	CF	DA	DF
Line	Mean	0.01	0.00	NA	0.02	0.11	0.00	0.04	0.41	0.01
	Rms	0.17	0.17	NA	0.24	0.27	0.40	0.40	1.5	0.65
Sample	Mean	0.05	0.03	NA	0.02	0.10	0.01	0.02	0.41	0.10
	Rms	0.26	0.25	NA	0.37	0.30	0.50	0.40	0.91	0.69

3.3. Global geometric performance assessment

The previously described geometric performance assessments, performed for certain time periods as data became available, were a good indicator of the potential problems with the nominal implementation of the standard production algorithm. They were also useful when dealing with occasional reduced accuracies in the supplied spacecraft ephemeris and attitude data as occurred during periodic orbit maneuvers. Nevertheless, any decision to update the standard production algorithm needed to be supported by a global analysis. Fig. 4A and B summarize the quality assessment data to illustrate the pointing stability and co-registration performance for a four-year period. Fig. 4A is the plot of mean offsets (ICC) taken for every 500 orbits. As illustrated, pointing stability is fairly good for eight out of nine cameras. Small pointing changes, common to all eight cameras, slowly varying with a phase period of approximately one year can be observed. These are considered acceptable, given the magnitude and rate of recurrence, and are accounted for by nominal standard processing. They seem to correspond to spacecraft inclination maneuver history, which is certainly more dominant at the beginning of the mission prior to the spacecraft reaching a stable orbit. This is not the case for the Da camera which exhibits an independent and very irregular pointing change. Fig. 4B serves as the confirmation that sub-pixel geo-location co-registration accuracies have been achieved for eight cameras during the entire mission. Corresponding mean and rms errors are very similar to those from Table 2.

Investigations focusing on Da camera timing information or engineering temperature data concluded that these are not the cause of apparent pointing change. There are no other obvious hypotheses to be linked with this idiosyncrasy. It is possible that there is a small and occasional physical movement of the Da camera system. However, that can not be concluded with certainty by analyzing available data.

4. Sensitivity of Level 2 cloud winds and height retrievals

An important factor influencing our decision to update the standard processing algorithm is the propagation of the geo-location and co-registration errors and their impact on the subsequent Level 2 retrievals, given their sensitivities to small errors. This is most critical for MISR's novel approach to simultaneously retrieving cloud heights and wind vector data using multi-angle imagery. The algorithm includes the matching

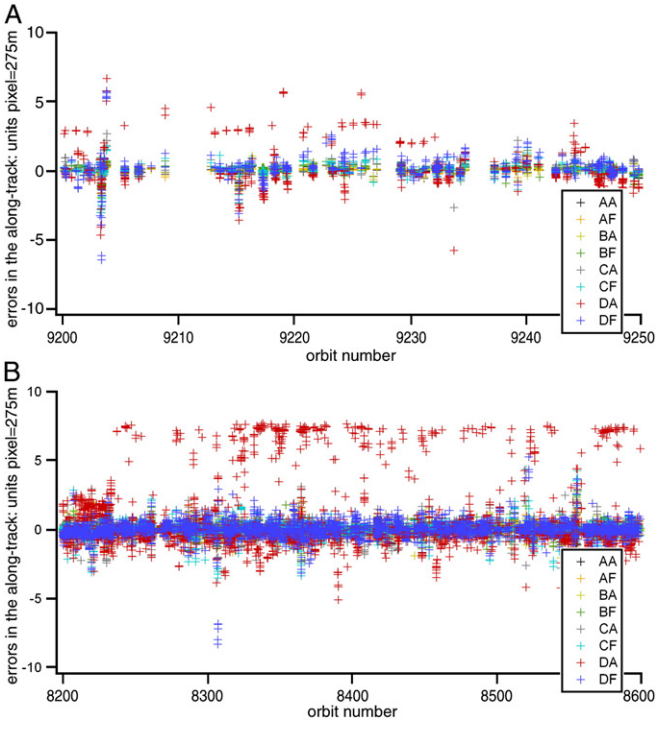

Fig. 3. A: Geo-location and co-registration as evaluated in the along-track line direction for nine MISR cameras: errors (pixel units, 1 pixel=275 m) plotted against orbit number corresponding to approximately 4 days of data, Sept. 10–14, 2001. B: Geo-location and co-registration as evaluated in the along-track line direction for nine MISR cameras: errors (pixel units, 1 pixel=275 m) plotted against orbit number corresponding to approximately 25 days of data, July 4–29, 2001.

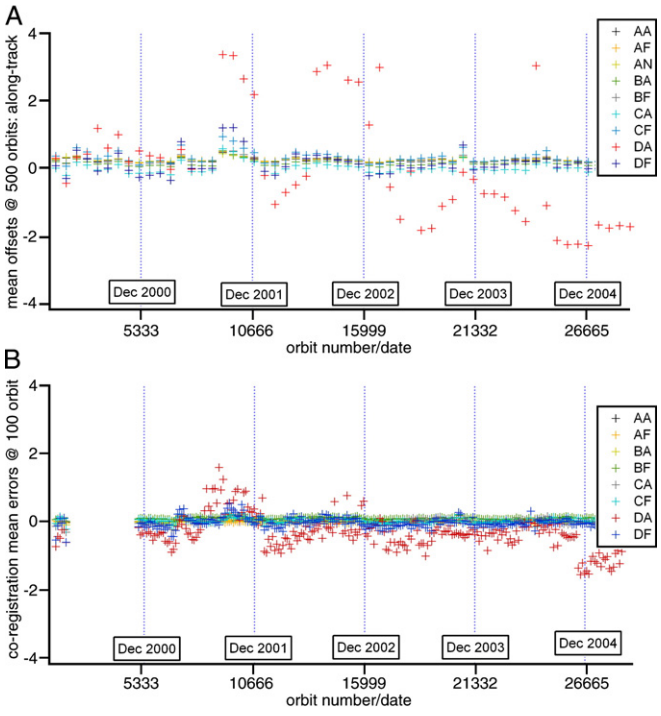

Fig. 4. A: Global geometric performance for nine MISR cameras: Pointing stability illustrated by ICC (offsets in pixel units, 1 pixel=275 m) means at 500 orbits plotted against orbit number. B: Global geometric performance for nine MISR cameras: geo-location and co-registration mean errors (in pixel units, 1 pixel=275 m) at 100 orbits plotted against orbit number.

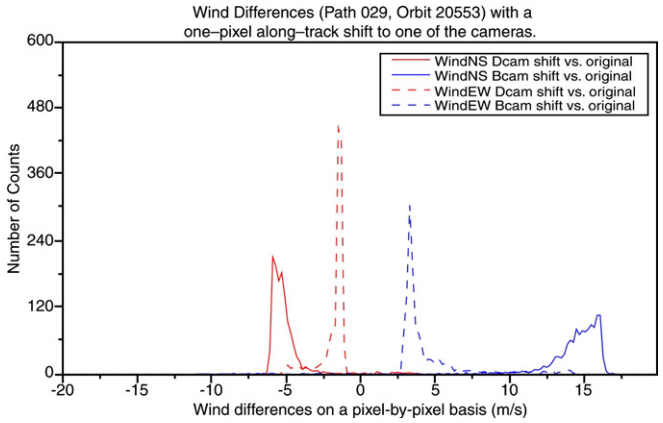

Fig. 5. Simulated sensitivity of wind calculation on co-registration errors — change in wind vector due to pixel shift in the along-track direction.

of cloud points viewed from different angles and using the retrieved disparities along with known pointing angles to calculate desired parameters. It is possible to separate the effects of cloud motion and cloud height in the image disparities, because multiple cameras view the same scene from different angles, from the vertical to the highly oblique (±70°.) (Moroney et al., 2002). The wind retrieval (using either the An-Bf-Df or An-Ba-Da triplet of cameras) algorithm is applied first. The triplet of cameras forms a set of equations that solves for the along-track and cross-track components of the wind vector as well as for the height of the feature simultaneously. The result of this calculation is a wind vector (speed and direction) and associated height for each 70.4 km domain in the swath. Then, the stereo-matching algorithms are again applied on the An-Af and An-Aa cameras to solve for the cloud-top heights. The already-obtained winds are used in this process to differentiate the image disparity due to cloud motion during the elapsed time from the disparity due to the height of the cloud above the reference height (the surface ellipsoid). However, calculated winds are very sensitive to any errors in co-registration. A 1 pixel error in either of the D cameras results in a 5 m/s wind error, and a 1 pixel shift in either of the B cameras will introduce a wind error of 15 m/s. Therefore, the winds (not the final wind vector, but rather the difference between the vectors derived from the forward and aft pointing cameras separately, see Section 6.2) can serve as a very good indicator of the registration quality, provided one has confidence in the stereo matchers. The sensitivities mentioned above can be determined by artificially shifting the position of one of the cameras by a fixed amount and looking at the change in the winds. Fig. 5 shows the results of this simulation. The significant effects caused by even small co-registration errors in the Ba camera are apparent.

We originally intended to correct only the Da camera, since that was known to have the most serious problems, but testing showed that stereo retrieved superior results (smallest wind difference between the forward and backward cameras) when all four cameras (Df, Bf, Ba and Da) were corrected (see Table 3). The wind is calculated separately for the forward (An-Bf-Df) and backward (An-Ba-Da) cameras, and the difference between these two values serves as both a quality check on the individual wind vector as well as a good indicator of the overall registration quality. The errors in the B cameras are small compared to Da, but the sensitivity of the wind retrieval to the B cameras is so large that we have to take any B camera errors into account as well when generating the final Geo-rectified Radiance Product. Table 3 is used to summarize the co-registration test results to confirm that desired performance relative to cloud height/wind retrieval has been achieved with the updated standard production algorithm.

5. Updates to the nominal geo-rectification process

The analysis of overall geometric performance, along with the wind retrieval sensitivity studies, made it clear that at least two cameras (Da and Bf) had systematic co-registration errors large enough to significantly affect the accuracy of Level 2 stereo wind retrievals. The Da camera problem was particularly severe with variable co-registration errors on the order of more than 2 pixels. The software solution that was adopted to correct this problem was to add two new components to the standard processing stream after nominal geo-located products were generated, followed by a second pass of geo-rectification as originally implemented. The first new component was an enhanced version of the CameraMatch program to identify and quantify co-registration errors. The second generated a corrected set of coefficients for "warping" the images and quantifying the uncertainties in these values. The final pass of geo-rectification then applied the co-registration corrections to selected cameras.

5.1. Camera matching

The core of the CameraMatch program is based on a combination of standard pattern-matching algorithms used in image processing. First a region of pixels in the An camera

Table 3
Test results between preliminary (nominal) and final standard processing algorithm as obtained from processing 40 randomly selected orbits

CAM.	Preliminary				Final			
	Co-registration		Wind diff		Co-registration		Wind diff	
	Mean	STD	Mean	STD	Mean	STD	Mean	STD
Da	0.99	1.41	8.09	10.5	-0.04	0.26	3.85	7.98
Df	-0.05	0.33			-0.03	0.25	3.49	7.96
Ba	0.02	0.17			0.01	0.16	1.61	5.53
Bf	0.11	0.23			0.01	0.18		
Ca	0.00	0.25			0.01	0.25		
Cf	0.03	0.23			0.03	0.24		
Aa	0.01	0.14			0.01	0.15		
Af	0.00	0.14			0.00	0.14		

Preliminary and final co-registration mean and standard deviation errors given in pixel units (1 pixel=275 m). Wind motion vector differences given in m/s. Right most column wind difference shows wind motion vectors improvements given the improvement of D and B cameras. Highlighted rows for B cameras indicate significant improvement in the precision of the wind motion vectors given relatively small reduction of the co-registration errors. The wind statistics quoted here apply to all winds, rather than only the good quality ones, so they differ from similar measures as quoted in Davies et al. (2006-this issue) which only consider those wind vectors that pass the quality tests.

reference image is selected for analysis. Next, a feature extractor identifies the most strongly contrasting feature in the pixel field (Forstner, 1987). Then a small patch of pixels centered on the feature is isolated from the image and is used for matching with the corresponding region in the comparison camera's image. The matching process finds the optimal, spatially-adjusted match of the reference image relative to the comparison image, returning the offset between the two images where the match occurred. In practice, the procedure is able to estimate co-registration errors to within about 1/10 pixel by fitting a set of normal equations describing geometric and radiometric relations between the two image windows (Ackerman, 1984; Rosenholm, 1987).

This matching procedure, including feature extraction for candidate points, is performed at over 400 uniformly-spaced areas in each block for each of the 8 An-X camera pairs. In order for the offsets at a location to be included in the block average, all 8 camera pairs must successfully match. However, this criterion is not usually satisfied, either because there is too little contrast in a scene for one or more cameras (such as over water or homogeneous land), or because the scene is cloud-contaminated. Cloudy scenes can usually be identified and excluded, because their heights above the terrain produce large along-track offsets that decrease monotonically as one progresses from camera Df to Da.

In the initial version of CameraMatch used for validation purposes only, the algorithm was applied only to Level 1 terrain-referenced, radiance imagery. In this product, blocks that are entirely over water do not contain data. The absence of land over large parts of most orbits and the prevalence of clouds resulted in many orbits with no successful matches. Consequently, in order to increase the number of blocks with matches, it was decided to run CameraMatch against Level 1 ellipsoid-referenced, radiance imagery at high latitudes where sea-ice might be present, in addition to over land. Floating ice remains essentially stationary over the 7 min required for MISR's 9 cameras to image a scene; floating ice is essentially at sea-level, so its projection to the ellipsoid is not a concern; and the patterns on their surfaces provide excellent contrast for the matcher.

The latest version of the CameraMatch algorithm was also enhanced in several other ways. Its cloud-detection capability was improved by accommodating the fact that low-level clouds moving rapidly in a southerly direction could create larger along-track offsets than the parallax caused by altitude above the surface ellipsoid. Cloud-detection was further improved by exempting the Da camera from obeying the monotonicity rule, since it was known to have significant co-registration errors that could negate the effects of parallax. Finally, an outlier-detection algorithm was added to remove match locations where the along-track or across-track offsets for any camera deviated from the mean by more than one standard deviation.

The results are reported in ASCII product files containing the means and standard deviations for the along-track (Lin=line) and across-track (Smp=sample) directions measured in fractional pixels. It also contains the block number, whether the match was from a terrain (T) or ellipsoid (E) image and the number of successful matches. A portion of a CameraMatch product file is shown in Table 4.

5.2. Computing corrections

An analysis of the CameraMatch co-registration errors for hundreds of MISR orbits suggested that these errors were slowly varying with respect to the duration of a single orbit. Therefore, it was decided to use the statistics from all the successful matches in an orbit to find constant, along-track and across-track offsets to apply to that entire orbit. Further, since there were Image Coordinate Corrections available from the ROI data in addition to the co-registration corrections from CameraMatch, the results from both sources could be merged together to improve the statistics.

Merging the corrections was accomplished in two steps. First, along-track and across-track vectors representing the 180 blocks in an orbit were populated with corrections at those blocks (nodes) where either an ROI or a CameraMatch correction was available. Next, corrections were found for all empty blocks by interpolating between the nodes. If four or more blocks in an orbit are nodes, an initial mean value is determined, and any points falling more than 1.5 standard deviations from the mean are excluded as outliers. The remaining points are averaged, and the result is the offset that is applied to the orbit.

Table 4
A portion of the CameraMatch program output file providing means and standard deviation of co-registration offsets between nadir (An) camera and other eight cameras

Matcher results by block, line/sample and camera																			
				Camera offset means								Camera offset standard deviations							
Blk	T/E	Dir	Num	Df	Cf	Bf	Af	Aa	Ba	Ca	Da	Df	Cf	Bf	Af	Aa	Ba	Ca	Da
43	E	Lin	15	0.56	0.66	0.22	0.19	−0.18	−0.43	−0.72	−0.44	0.23	0.15	0.12	0.02	0.04	0.14	0.20	0.19
43	E	Smp	15	−0.13	−0.06	−0.07	−0.01	0.04	−0.10	0.11	−0.16	0.16	0.11	0.08	0.04	0.02	0.06	0.04	0.09
48	T	Lin	35	0.42	0.54	0.12	0.19	−0.09	−0.26	−0.51	−0.17	0.24	0.16	0.11	0.03	0.05	0.11	0.24	0.17
48	T	Smp	35	−0.24	−0.08	−0.07	−0.03	0.03	−0.10	0.08	−0.14	0.17	0.07	0.05	0.04	0.03	0.08	0.05	0.13
53	T	Lin	14	−0.67	−0.17	−0.31	0.02	0.08	0.15	0.19	0.78	0.26	0.13	0.08	0.05	0.04	0.08	0.13	0.17
53	T	Smp	14	−0.19	0.00	−0.01	−0.02	0.04	−0.04	0.06	−0.13	0.16	0.09	0.07	0.03	0.04	0.10	0.09	0.22

First four column from the right specify, as follows: 1) block number within an orbit path for which co-registration offset is derived, 2) whether the match is from terrain block (mostly land surface including coast lines) or ellipsoid block (includes also ocean surface), 3) direction of measured co-registration offset, and 4) number of successful matches. The two segments on the right give mean and standard deviation per camera in pixel units, 1 pixel=275 m.

The derived offset is used to modify the ICC coefficients used in "warping" the images to fit the ROI during the first geo-rectification. This was accomplished by simply adding the constant offset values to the 0th-order ICC coefficients for all blocks in the orbit. Subsequently, these new coefficients are used in final geo-rectification processing of MISR data.

5.3. GDQI

The Level 1 ellipsoid- and terrain-referenced products have always carried a field named GDQI (Geographic Data Quality Indicator). The intent of the field was to indicate, for each camera and block in an orbit, whether the geo-location of the image was acceptable for use. Prior to this final version, the GDQI was only marginally useful due to unexpected pointing issues with the Da camera. This latest update to standard processing provided an opportunity to quantify, with much higher confidence, the likelihood that geo-registration (relative to the nadir) camera was good.

The procedure for calculating GDQI was implemented in the same code where determination of the corrections to apply to the data was implemented. Three different cases are identified:

1) When there are no CameraMatch corrections and no ROI corrections in an orbit, then the GDQI for all blocks in all cameras except the Da camera is set to 0.0. The Da camera GDQI is set to −0.5 in this case.
2) When there are no CameraMatch corrections or fewer than 2 camera match plus ROI corrections, then the GDQI for all blocks in all cameras is set to 0.0.
3) Whenever the above conditions are not met, the following algorithm is used. It is applied to each block using the interpolated corrections or offsets described in Section 5.2. First, the "diagonal_difference" is computed between the mean along-track and across-track offsets derived for the orbit and the corresponding offsets for the block. The GDQI for the block becomes 1.0–1.5*"diagonal_difference". Any value less than −1.0 is set to −1.0.

The "diagonal_difference" in 3) above is always positive and approaches zero as the block offsets approach the mean orbit offsets. Therefore, a block for which its offsets are identical to the orbit offsets will have a GDQI of 1.0; a block whose offsets create a "diagonal_difference" of 2/3 will have a GDQI of 0.0; and a block whose offsets create a "diagonal_difference" greater than 4/3 will have a GDQI of −1.0. A 2/3 pixel offset thus becomes the threshold for determining whether a block is adequately registered. An evaluation of GDQI confidence is given in the subsequent section along with other operational results for the final implementation of standard processing.

6. Final geo-rectification performance and its effect on the cloud wind/height estimates

The final update to the geo-rectification component of MISR standard processing has been operational since May 2005. In addition to being applied to data acquired since then, this new version of the algorithm has also been applied to data from other acquisition periods as a result of the regular reprocessing schedule and special reprocessing requests from shorter time periods. Final geo-rectification performance metrics, including co-registration errors and GDQI evaluation, has been obtained from all data processed since May 2005, including the reprocessed data for year 2000. The cloud wind/height algorithm updates, which include the forward–backward wind difference that allow us to measure the co-registration quality, has been operational only since December 2005. Consequently, only data specially reprocessed for this paper from December 2001, are used for the evaluation of the final performance of that algorithm.

6.1. Final geo-location and co-registration performance

Figs. 4B and 6 show the global quality of the geo-rectified product by summarizing co-registration errors. However, the Fig. 6 uses data obtained from the final version of the product. The available data from 2000 and 2005 shows obvious improvement when compared with data in Fig. 4B. Da camera performance is no longer significantly different from other cameras, and virtually all of the mean errors overlap within a ±0.1 pixel range.

6.2. Operational wind vector accuracy improvement

One of the best ways to control the quality of wind vector estimates is to calculate them separately using the forward pointing triplet of cameras (An-Bf-Df) and the backward triplets (An-Ba-Da) and difference these two values. This serves as both a quality check on the individual wind vectors, and when combined into a histogram, a good indicator of the overall registration quality of the orbit. To clarify: an individual wind vector by itself does not contain information about the quality of the orbit, but the cumulative distribution of the wind forward–backward differences is a very good indicator of the overall registration quality. If there is a systematic registration error in one of the cameras (say Da is off by 1 pixel with respect to the others), all the backward winds will be off by 5 m/s. This will shift the modal value of the wind differences histogram by 5 m/s which is very easy to detect.

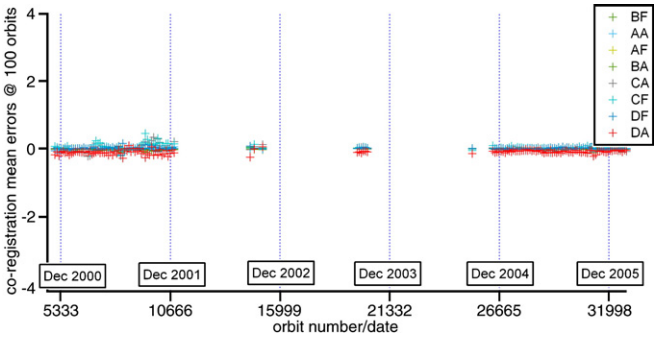

Fig. 6. Final global geometric performance for nine MISR cameras: Geo-location and co-registration mean errors (in pixel units, 1 pixel=275 m) at 100 orbits plotted against orbit number.

Fig. 7 and Table 5 show the differences between the registration quality of the "old" and "new" Level 1 data. The same wind algorithm (Davies et al., 2006) was used for all these analyses, so the only difference is in the co-registration quality. If one is looking at the forward–backward wind difference distribution for a single orbit, a histogram with a small standard deviation and a noticeable bias indicates a constant error in the geo-location of one or more cameras. The individual forward–backward wind differences can also be caused by blunders in the stereo-matching which leads to long tails in the distribution (not shown), so one should not expect a noticeably smaller standard deviation, or an obvious bias (nonzero modal value) in the results when many orbits are combined.

The key thing to note in Fig. 7 is that the corrected data from December 2005, is much more strongly peaked around zero than is the uncorrected data, indicating that much of the registration errors present in the old data have been eliminated. Table 5 shows the percentage of counts that have a wind difference less than or equal to 2 m/s in the north–south component (corresponding to a "VeryGood" wind in the MISR Stereo product), less than or equal to 5 m/s (a "Good") and the percentage of wind retrievals with a wind difference below 10 m/s (any difference more than this value labels the wind as "Bad"). The most recent version of the Stereo software has been modified to take into account the differences in the east–west wind component and the associated heights, but this does not make a big difference in the results. The number of wind retrievals meeting these standards has roughly doubled in all three categories with the inclusion of the geo-location corrections. December 2001 was chosen for the comparisons for reasons of expediency — we already had the necessary stereo data in hand. This month is representative of the entire "before correction" era. Reprocessing of the MISR data will eventually bring up the entire dataset to the most recent version of the software.

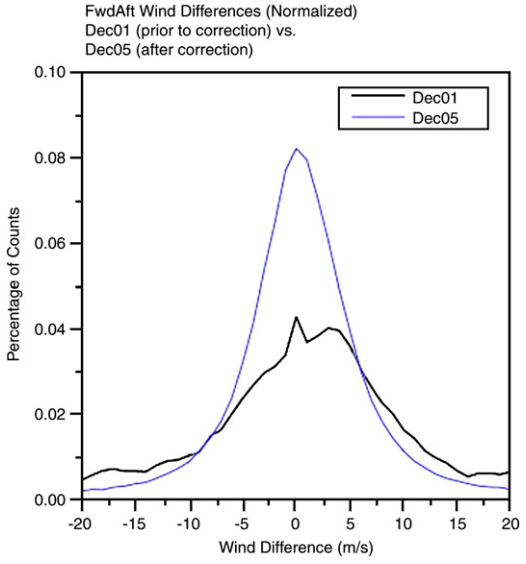

Fig. 7. Wind vector difference histogram illustrating improvement of quality in geo-rectified radiance data and its impact on wind vector calculations.

Table 5
Percentage of winds that had a forward–backward difference less than or equal to the indicated amount

Forward–backward Wind difference threshold (m/s)	December 2001 (uncorrected)	December 2005 (corrected)
2.0	18%	38%
5.0	38%	64%
10.0	57%	82%

Compared with the data for the same time of the year using as the input two versions of the Geo-rectified Radiance Product, the latest version for December 2005 and one version prior to last for December 2001.

6.3. Final GDQI results

In order to assess the degree to which the co-registration correction process and the GDQI determination were successful, the CameraMatch program was run against the output of the new L1B2-corrected products. Almost 10,000 orbits between orbits 1024 and 32859, corresponding to period between February 2000 and February 2006, were analyzed. The data consist of along-track and across-track residual co-registration errors (in

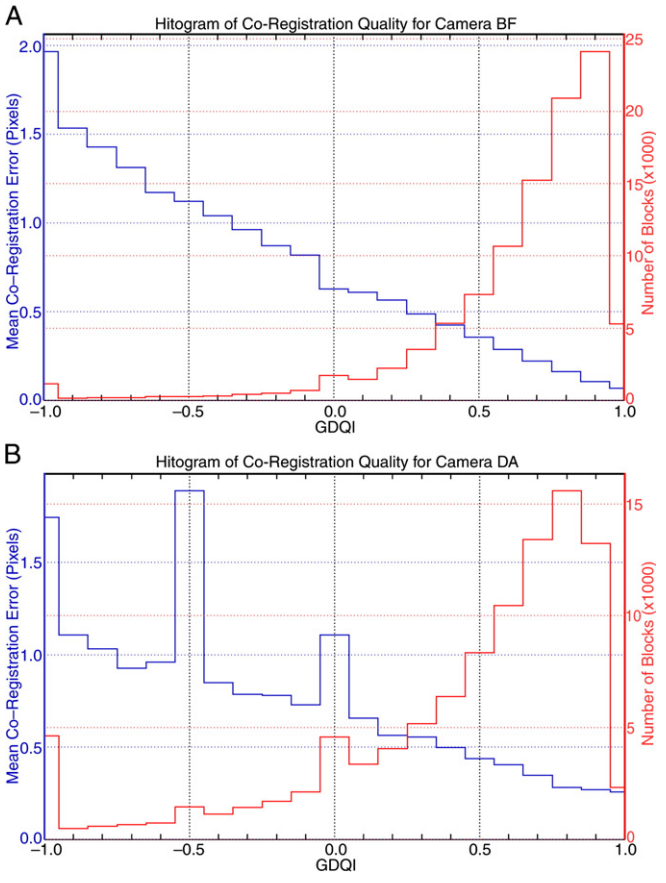

Fig. 8. A and B: Summary of GDQI performance for Da and Bf cameras. The blue histograms represent the mean combined line and sample residual errors in pixels after applying the co-registration correction. The red histograms represent the number of blocks over many orbits where the analysis was performed. Both are plotted against the GDQI, which is estimated per block during primary processing.

units of 275 m pixels) for each block in each orbit where a CameraMatch match was found along with the corresponding block-GDQI values. About 100,000 blocks are represented.

The results are summarized in the bar-charts of Fig. 8A and B. Only cameras Bf and Da are shown as representative of all nine cameras. In the charts, the GDQI range of [−1, +1] is divided into 21 bins. The along-track and across-track residual errors are accumulated for the all the blocks whose GDQI values fall into the bin. An estimate of the overall co-registration error for each block is made by finding the "diagonal" error, i.e., the square root of the sum of the squares of the along-track and across-track errors. The mean "diagonal" error of all the blocks in each bin is then computed and plotted in blue at the GDQI bin location. The number of blocks that were used in computing this mean residual error for each bin is plotted in red.

The results are as expected. For each camera, the mean, residual co-registration error has a high, negative correlation with the GDQI value. The number of blocks has a high positive correlation with the GDQI value. At the left where the GDQI indicates poor registration, the mean error is large and the number of blocks that are poorly registered is small. At the right where the GDQI indicates good registration, the mean error is small, converging to about 0.2 pixels for most cameras, and the number of blocks that are well registered is large.

The distribution of number of blocks in the GDQI bins indicates that the bulk of the data are well registered. This trend is modified slightly by the observation that the more oblique cameras have a larger fraction of blocks with poor GDQI. Nevertheless, the Da camera, which had mean co-registration errors of 1 to 2 pixels before the programmatic fix, now has errors less than 0.5 pixels for 70% of the blocks.

The large peak in the Da camera bar-graph at GDQI equals −0.5 is a consequence of the special attention given this camera, whereby the GDQI for all blocks is set to −0.5 when there are no CameraMatch corrections and no ROI corrections in an orbit.

7. Summary

A nominal design of MISR geo-rectification processing as implemented and evaluated during the first four years of the mission fully met geo-location and co-registration accuracy requirements for eight out of nine cameras. Dynamic pointing instability of one of the most oblique cameras, namely Da, could not be corrected with the original algorithm. These conclusions have been made based on results from a global geometric quality monitoring system which continuously acquired and summarize performance assessment data as a complement to the production cycle.

During this assessment period, improvements were also made to the cloud height/wind retrieval algorithm, and sensitivity studies were conducted, both with the goal of producing cloud-resolved wind vectors with speed accuracies below 3 m/s standard deviation (when looking at the good quality winds only (Davies et al., 2006)). As a result, it was concluded that it is necessary to remove even small sub-pixel co-registration errors in the B cameras to reach the desired accuracy and coverage of wind retrievals. Consequently, a revised version of the geo-rectification production algorithm, in addition to correcting the Da camera, also updates the processing chains for the three other off-nadir cameras (i.e. Ba, Bf, Df) used in cloud height/wind retrievals. Final results show that the co-registration performance of the Da camera meets the expected goals. Also, the desired effect on wind accuracies and coverage has been evaluated using operational data. Finally, the revised algorithm includes a new method for generating the GDQI metadata product which now is fully suitable for quality control purposes.

There are no anticipated future updates to the current implementation of the algorithms underlying operational geo-rectification and geometric quality monitoring systems. Current operational work includes nominal system maintenance, regular data transfer and summaries, and periodic interactive data analysis. Future modifications may be needed only in support of special case studies such as independent data validations using a digital photogrammetric workstation, accurate 3D extraction of deep convective cloud morphologies, or fusion with data from other sensors.

Acknowledgments

The authors gratefully acknowledge the efforts of the MISR science data system team at JPL who worked diligently along with the Langley Distributed Active Archive Center (DAAC) (NASA Langley Research Center's Atmospheric Sciences Data Center) staff to make large amount of data available for analysis. Thanks also to Mike Bull and Jason Matthews for making possible the quality assessment data summaries, and to Roger Davies for his valuable suggestions in improving camera co-registration and most critical insights into the sensitivities and requirements of cloud wind retrievals.

References

Ackerman, F. (1984). Digital image correlation: Performance and potential application in photogrammetry. *Photogrammetric Record*, *11*, 64.

Bailey, G. B., Carneggie, D., Kieffer, H., Storey, J. C., Jovanovic, V. M., & Wolfe, R. E. (1997). Ground control points for calibration and correction of EOS ASTER, MODIS, MISR and Landsat 7 ETM+ Data. *SWAMP GCP Working Group Final Report, USGS, EROS Data Center, Sioux Falls, SD*.

Bothwell, G., Hansen, E. G., Vargo, R. E., & Miller, K. C. (2002). The MISR science data system, its products, tools, and performance. *IEEE Transactions on Geoscience and Remote Sensing*, *40*.

Davies, R., Horvath, A., Moroney, C., Zhang, B., & Zhu, Y. (2006-this issue). Cloud motion vectors from MISR using sub-pixel enhancements. Remote Sensing of Environment. doi:10.1016/j.rse.2006.09.023.

Diner, D. J., Beckert, J. C., Reilly, T. H., Bruegge, C. J., Conel, J. E., Kahn, R. A., et al. (1998). Multi-angle Imaging SpectroRadiometer (MISR) instrument description and experiment overview. *IEEE Transactions on Geoscience and Remote Sensing*, *36*, 1072−1087.

Forstner, W. (1987). A fast operator for detection and precise location of distinct points, corners and centers of circular features. *ISPRS Intercommission Workshop, InterLaken*.

Jovanovic, V. M., Bull, M., Smyth, M. M., & Zong, J. (2002). MISR in-flight camera geometric model calibration and achieved georectification performances. *IEEE Transactions on Geoscience and Remote Sensing*, *40*.

Jovanovic, V. M., Smyth, M. M., Zong, J., Ando, R., & Bothwell, G. W. (1998). MISR photogrammetric data reduction for geophysical retrievals. *IEEE Transactions on Geoscience and Remote Sensing*, *36*(4).

Logan, T. L. (1999). EOS/AM-1 Digital Elevation Model (DEM) Data Sets: DEM and DEM auxiliary datasets in support of the EOS/Terra platform. *JPL*

D-013508, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA.

Moroney, C., Davies, R., & Muller, J. P. (2002). Operational retrieval of cloud-top heights using MISR. *IEEE Transactions on Geoscience and Remote Sensing*, *40*.

NASA Langley Research Center's Atmospheric Sciences Data Center, http://eosweb.larc.nasa.gov

Rosenholm, D. (1987). Empirical investigation of optimal window size using the least squares image matching method. *Photogrammetria (PRS)*, *42*, 113−125.

Snyder, J. P. (1987). Map projection — A working manual. *United States Geological Survey Professional Paper 1395, U. S. Government Printing Office, Washington*.

Zong, J., Davies, R., Muller, J. P., & Diner, D. (2002). Photogrammetric retrieval of cloud advection and cloud-top height from the Multi-angle Imaging Spectro-Radiometer (MISR). *Photogrammetric Engineering and Remote Sensing*, *68*.