

In-Flight Performance of MESSENGER's Mercury Dual Imaging System

S. Edward Hawkins, III^{*a}, Scott L. Murchie^a, Kris J. Becker^c, Christina M. Selby^a, F. Scott Turner^a,
Matthew W. Noble^a, Nancy L. Chabot^a, Teck H. Choo^a, Edward H. Darlington^a, Brett W. Denevi^b,
Deborah L. Domingue^a, Carolyn M. Ernst^a, Gregory M. Holsclaw^d, Nori R. Laslo^a, William E.
McClintock^d, Louise M. Prockter^a, Mark S. Robinson^b, Sean C. Solomon^e,
and Raymond E. Sterner, II^a

^aThe Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Road, Laurel,
MD 20723, USA;

^bSchool of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA

^cAstrogeology Branch, U.S. Geological Survey, Flagstaff, AZ 86001, USA

^dLaboratory for Atmospheric and Space Physics, University of Colorado, Boulder, CO 80303, USA

^eDepartment of Terrestrial Magnetism, Carnegie Institution of Washington, Washington, DC 20015,
USA

ABSTRACT

The MErcury Surface, Space ENvironment, GEOchemistry, and Ranging (MESSENGER) spacecraft, launched in August 2004 and planned for insertion into orbit around Mercury in 2011, has already completed two flybys of the innermost planet. The Mercury Dual Imaging System (MDIS) acquired nearly 2500 images from the first two flybys and viewed portions of Mercury's surface not viewed by Mariner 10 in 1974-1975. Mercury's proximity to the Sun and its slow rotation present challenges to the thermal design for a camera on an orbital mission around Mercury. In addition, strict limitations on spacecraft pointing and the highly elliptical orbit create challenges in attaining coverage at desired geometries and relatively uniform spatial resolution. The instrument designed to meet these challenges consists of dual imagers, a monochrome narrow-angle camera (NAC) with a 1.5° field of view (FOV) and a multispectral wide-angle camera (WAC) with a 10.5° FOV, co-aligned on a pivoting platform. The focal-plane electronics of each camera are identical and use a 1024×1024 charge-coupled device detector. The cameras are passively cooled but use diode heat pipes and phase-change-material thermal reservoirs to maintain the thermal configuration during the hot portions of the orbit. Here we present an overview of the instrument design and how the design meets its technical challenges. We also review results from the first two flybys, discuss the quality of MDIS data from the initial periods of data acquisition and how that compares with requirements, and summarize how in-flight tests are being used to improve the quality of the instrument calibration.

Keywords: Mercury, Imager, MDIS, MESSENGER

1. INTRODUCTION

The MErcury Surface, Space ENvironment, GEOchemistry, and Ranging (MESSENGER) spacecraft^[1] began its 7.6-year mission to explore Mercury on August 3, 2004. The spacecraft with its science payload will address key questions regarding planetary formation and the evolution of the inner planets. The Mercury Dual Imaging System, or MDIS^[2], is a key component of the payload and will provide new insights into fundamental questions that the mission will address.

Mercury poses unique challenges to an orbital imaging investigation^[3] due both to thermal and mission design requirements. Mercury's close proximity to the Sun (perihelion ~0.3 AU) results in the equivalent of up to 11 Suns of radiant intensity. The spacecraft is protected by a ceramic-cloth sunshade, and spacecraft pointing flexibility is limited by the requirement to maintain shade on spacecraft structures. The planet is locked in a 3:2 spin-orbit resonance, resulting in a solar day equal to two Mercury years (176 Earth days). These factors, combined with the high eccentricity of Mercury's orbit, result in a subsolar temperature of ~460°C at perihelion^[4]. Even though MDIS is protected from

* Send correspondences to S. Edward Hawkins, III; E-mail: ed.hawkins@jhuapl.edu

direct sunlight by the sunshade, pointing at the hot planet during imaging presents a challenge to maintain key components at acceptable operating temperatures. The large change in spacecraft velocity required to enable the probe to be captured by Mercury is accomplished through the use of a lightweight spacecraft, miniaturization of instruments, and multiple planetary flybys: one of Earth, two of Venus, and three of Mercury. To date, MESSENGER has completed five of the six planetary encounters, with the sixth and final flyby to occur in September 2009.

Prior to MESSENGER, much of our knowledge of Mercury had been obtained from the Mariner 10 mission. During three flybys of the planet in 1974 and 1975, Mariner 10 imaged about 45% of the planet, most at a resolution of >1 km/pixel. Images from Mariner 10 revealed a heavily cratered surface, major contractional faults suggesting that Mercury's diameter decreased as the planet cooled internally, and huge expanses of smooth plains of moderate albedo unlike the dark maria of the Moon. To provide an improved understanding of Mercury's geology, the objectives of MESSENGER's imaging investigation are to:

- Distinguish surface units of different spectral character;
- Provide global imaging at hundreds of meters resolution or better;
- Provide global multispectral mapping at kilometer sampling to determine mineralogy and construct spectral maps; and
- Provide high-resolution imaging of key areas.

The currently planned measurement campaigns to address these objectives include:

- Flyby multispectral imaging at 2.5 km/pixel followed by orbital multispectral imaging at <1 km/pixel at low incidence and emission angles;
- Flyby monochrome imaging at 200-500 m/pixel followed by orbital imaging at ~ 200 m/pixel, at near-zero emission angles and moderate incidence angles;
- A stereo complement to the orbital, near-global mosaic, plus opportunistic stereo observations during the flybys; and
- Targeted, high-resolution images of type areas, including color imaging at low solar incidence angles, imaging at moderate incidence angles optimized for resolving surface morphology, and stereo imaging of selected targets.

2. INSTRUMENT OVERVIEW

The MDIS instrument (Fig. 1) consists of a monochrome narrow-angle camera (NAC) with a 1.5° field of view (FOV) and a multispectral wide-angle camera (WAC) with a 10.5° FOV. Together the two cameras allow imaging resolution requirements to be met from MESSENGER's highly elliptical orbit. A simplified block diagram of the complete MDIS instrument, which includes redundant Data Processing Units (DPUs) as interfaces to the spacecraft, is shown in Fig. 2. The cameras are co-aligned and mounted on a pivoting platform, which, combined with limited spacecraft slewing, allows imaging at solar phase angles of 28° – 142° at the centers of the FOVs and nadir viewing of most of Mercury. The focal-plane electronics of each camera are identical and use a 1024×1024 charge-coupled device (CCD) detector. The cameras are passively cooled but use diode heat pipes and phase-change-material thermal reservoirs to maintain the thermal configuration (Fig. 3) during the hot portions of the orbit. Only one camera may operate at a time, and switching between cameras is done through the DPU Interface Switching Electronics (DISE). The MDIS instrument provides command and control capabilities, including image acquisition at rates up to 1 Hz, as well as automatic control of exposure times and a menu of pixel-binning and image compression options to manage data volume.

2.1 Optical Design

The WAC consists of a four-element refractive telescope having a focal length of 78 mm and a collecting area of 48 mm^2 . The detector located at the focal plane is an Atmel (Thomson) TH7888A frame-transfer CCD with a 1024×1024 format and $14\text{-}\mu\text{m}$ pitch detector elements that provide a $179\text{-}\mu\text{rad}$ pixel instantaneous field-of-view (IFOV). A 12-position filter wheel provides color imaging over the spectral range of the CCD detector. Eleven spectral filters spanning the wavelength range from 395 to 1040 nm are defined to cover wavelengths diagnostic of different potential surface materials. The twelfth position is a clear filter for optical navigation.

The NAC is an off-axis reflective telescope with a 550-mm focal length and a collecting area of 462 mm^2 . The NAC focal plane is identical to that of the WAC, providing a $25\text{-}\mu\text{rad}$ IFOV. The NAC has a single medium-band filter (~ 50

nm wide), centered at 750 nm to match the center of the corresponding WAC filter for monochrome imaging. Heat rejection coatings on the NAC filter and the first optic of the WAC were designed to reject heat from the planet.

2.2 Pivoting Platform

The MDIS pivot platform is controlled by a stepper motor with a hybrid gear train. The motor phases are controlled directly by software to move the platform forward or backward. The range of motion is constrained mechanically by “hard” stops and in software by “soft” stops. The reference “zero” position is along MESSENGER's +Z axis (co-boresighted with other surface-viewing instruments and normally 90° to the Sun). The total range of motion of MDIS is about 240°, limited by the hard stops at -185° and +55°. The pivot motor drive-train provides precision rotation over the 90° operational range of motion (Fig. 4) about the spacecraft +Z axis. Alternatively the pivot can be used to “stow” the instrument and protect the optics during spacecraft maneuvers, or to view an onboard calibration target inside the spacecraft adapter ring. Among the co-boresighted instruments, the Mercury Atmospheric and Surface Composition Spectrometer (MASCS)^[5] is of particular interest, because its Visible-Infrared Spectrograph (VIRS) component operates at wavelengths that span those of the WAC.

The MDIS pivot actuator is capable of accurately stepping in intervals of 0.01° (~150 μrad) per step. Pointing control is attained by first “homing” the instrument, which is accomplished by driving the actuator into one of the mechanical hard stops. Once the location of the pivot actuator is known, the flight software retains this knowledge and subsequent pointing commands are achieved by counting pulses (steps) to the motor. During ground processing of the images, there are two alternative approaches to reconstructing pivot position: by counting motor steps following homing, as described above, or by using the position returned from a pivot position resolver.

2.3 Thermal System

The high eccentricity of MESSENGER's orbit^[1] results in an approximately 2-hour period near periaapsis during which the spacecraft is subjected to a high thermal impulse from Mercury. A detailed thermal model of the spacecraft was developed early in the design phase of the mission with sufficient fidelity to provide high confidence in the expected thermal environment for the spacecraft, and to capture the worse-case heating conditions on the spacecraft for all orbits. However, the spacecraft thermal model did not account for the independent degree of freedom of the MDIS pivot, and so an independent radiant-energy transfer model was developed by the instrument team.

The instrument model included all pointing constraints, momentum management rules, and a model for on-orbit pointing (known to that date) at high time resolution (~1 minute). Both models treated the planet as a blackbody radiator but varied the surface temperature of Mercury from the subsolar point as $T \sim \cos^{1/4} \theta$ ^[1], where θ is latitude. Only the infrared component of heat flux was treated in the models, taking the Bond albedo of the planet to be a conservative 0.05. Five orbits were compared between the two thermal models at a variety of Mercury true anomalies (MTAs), and the predicted heat flux profiles were generally consistent. As expected, the noon-midnight orbits (MTA 240° and 60°) were found to have the highest peak fluxes; however, of primary concern for MDIS was the total integrated energy absorbed by the radiator during each 12-hour orbit period. The worst-case heating orbit was found to occur at MTA ~ 20° (Fig. 5) and corresponded to a spacecraft true anomaly of 50°, with the periaapsis on the night side. This orbit represents a nearly continuous (albeit lower level) heating case for MESSENGER. To account for the fact that heat is radiated away from MDIS when far from the planet, the instrument model incorporated a notched response. An additional factor of safety (20%) was added to the spacecraft thermal model to provide the best agreement between instrument and spacecraft thermal models.

From these two independent models we were able to bound the necessary requirements for the MDIS thermal system. Beryllium radiators are thermally tied to diode heat pipes and phase-change material (PCM) wax packs. Thermal links connect the wax packs directly to each camera's detector as shown in Fig 3. When MESSENGER is far from the hot planet, the heat pipes conduct energy to the radiator, which dissipates it. The wax packs freeze at -10°C, and cooling continues below the freezing point of the wax until the temperature reaches the low-temperature limit of the CCD (-45°C), where it is maintained by compensation heaters.

Near MESSENGER's periaapsis, the MDIS radiators absorb heat from the planet and the diode action of the heat pipes shuts down their high conductivity. When the CCD temperature reaches -10°C, the wax begins to melt, clamping the CCD temperature to its high-temperature limit. The PCM wax packs were sized to 240 g of paraffin on the basis of the spacecraft and instrument-specific thermal models.

3. IN-FLIGHT PERFORMANCE

3.1 Observations Overview: Flybys

MDIS imaging was acquired during the Earth and second Venus flybys to test operations of the instrument and spacecraft. Images of the Earth (Fig. 6a) and Moon used relatively simple spacecraft pointing options and exercised basic MDIS exposure control and compression options. Lunar imaging constituted a major test of radiometric calibration. Two of the lessons learned were (1) how to configure automatic exposure robustly to avoid its being set by cosmic ray hits, and (2) the complicated scattered light behavior of the WAC. Venus imaging (Fig. 6b) was used to further explore scattered light, to validate the flat-field correction, and to practice acquiring image mosaics. The first and second Mercury flybys (M1 and M2) began MDIS's science investigation of Mercury (Figs. 6c and 6d). Each flyby featured global hemispheric WAC multispectral and NAC mosaics, smaller-scale higher-resolution stereo mosaics of equatorial and northern hemisphere regions, and approach and departure navigation imaging. M2 also included a comprehensive set of measurements intended to characterize scattered light in the WAC.

3.2 Image Geometric Calibration

Knowledge of the pointing of each pixel in an MDIS image is required for any type of photogrammetric processing, in particular to construct mosaics and stereo models of Mercury's surface, and to coregister multispectral images from individual frames taken through different WAC filters. The goal for precision of pointing knowledge is $45\ \mu\text{rad}$ (1.8 NAC pixels) image to image, which yields mosaicking errors of 1.8 NAC pixels for nadir-pointed images — slightly smaller than the smallest feature resolvable by the NAC — or misregistration of WAC filters by 0.25 pixels. Three major factors contributing to pointing precision are knowledge of the cameras' focal lengths and radial distortions, knowledge of pointing of the cameras within the pivot plane, and the precision of the reported spacecraft attitude (which controls pointing of the origin of motion within the pivot plane).

Early mosaics of M1 NAC images exhibited significant misalignment between frames (Fig. 9, left), about five times larger than the goal for precision of pointing knowledge. To help diagnose the problem, NAC mosaics were overlaid on projected WAC images. This comparison demonstrated both translations of individual NAC frames relative to WAC images and distortions within the NAC frames, suggesting errors both in the optical models of the cameras and in converting telemetered pivot position to pivot angle. Errors in the optical distortion models can be traced to undersampling of the fields of view during ground geometrical calibrations that allowed focal length to be estimated but not radial distortion. Insufficient knowledge of position in the pivot plane is traceable to usage of telemetry from the pivot. During ground tests, an excessive amount of noise in the pivot position resolver invalidated the resolver test data and led to an interpretation that the resolver was not as accurate as required for determining pivot angle. Instead, a linear transformation of the count of pivot motor steps was used to estimate angular position within the pivot plane. This approach ignored nonlinearity in that relationship. Subsequent flight tests have largely corrected both sources of error and provided improved pointing precision that meets performance goals.

There is also a goal of $180\ \mu\text{rad}$ for long-term accuracy of pointing knowledge to enable the construction of digital terrain models using stereo images. A major factor in controlling accuracy will be knowledge of thermal effects on the relative alignment of the base on the MDIS pivot and the MESSENGER star cameras, and this issue cannot be assessed until the orbital phase of the mission.

Optical Distortion. The optical model of the WAC was updated using flight tests. It consists of a radially symmetric polynomial, with which focal length and radial distortion were fit to a series of star calibration images acquired through the clear filter prior to M1. Each of the color filters utilizes the same distortion model, with focal lengths scaled according to expectations from instrument design parameters. Direct determination of optical parameters for the color filters is precluded by their narrow bandpasses designed for imaging of Mercury, which result in detection of only about five stars even in a long-exposure image. Despite the difficulty of fitting these parameters to calibration data captured in flight, the focal lengths and models for the color filters appear to be working reasonably well, with further improvements expected.

Through M2, the optical distortion model of the NAC was also a simple, radially symmetric polynomial model fit to early star calibration images acquired prior to M1. Although root mean square (RMS) residuals of these regressions were noticeably higher than the residuals from WAC images of the same star clusters (0.8 pixels versus 0.2 pixels), the high NAC residuals were initially attributed to low signal levels. (Even images of the M7 open cluster yielded only 5-15 stars with sufficient signal to centroid properly, whereas an equivalent WAC image of the same cluster would yield 40-

70 such stars.) Misalignment in the M1 mosaics as well as severe ray intersection errors during stereo reconstruction caused the radially symmetric model to be revisited. A ray-trace model of the off-axis NAC camera used design parameters to redetermine the angular relations between pixels projected to infinity (Fig. 7). The radially non-symmetric result was fit with a third-order Taylor expansion, which was then regressed against the actual star positions in the in-flight calibration images to derive an updated focal length. The new focal length estimate and the distortions predicted from design parameters reduced the residuals from 0.8 pixels to 0.2 pixels RMS. Incorporating this new camera model of the NAC's optical distortion into the Mercury surface image registration process resulted in a breakthrough that allowed us to use precise photogrammetric methods for image registration as well as stereo surface reconstructions.

Pivot Pointing Knowledge. Most of the frame-to-frame misalignment of NAC images in early M1 mosaics was found to be in the pivot plane, so the linear conversion of pivot step count to pivot angle was reevaluated. Inspection of specifications of the pivot system suggested that the harmonic drive gear could result in a non-linear relationship, which nonetheless could be rigorously modeled^[6]. To quantify the non-linear relationship of pivot step count and pivot angle, a series of in-flight star images was designed to characterize the relationship in detail at the center and lower and upper bounds of the range of angles expected to be most used during Mercury orbital imaging, the -8° , 0° , and -28° positions, respectively. A single star was imaged at all three of these positions, and two sets of images were obtained at the 0° and -8° positions. In order to sample the entire harmonic drive response at each location, several dozen images were obtained within the 3° -range centered at each position. A model for the harmonic drive was created through Fourier spectral analysis of the test data. The frequencies determined from this analysis were those indicated from the earlier analysis^[6]. Software to determine the MDIS pivot position was updated with the harmonic drive model to obtain the parameters that minimize the RMS difference between the star location within the image and the modeled location. This software was run with one set of images at each of the three locations to obtain the desired parameters, and the resulting model was validated with the remaining two image sets. The result improved the RMS error from 12.6 NAC pixels for the linear model to 6.8 NAC pixels for the harmonic drive model (Fig. 8).

Two effects dominate residuals from the harmonic drive model. The first is a so-called “turn-around effect.” The flight imaging was completed in such a manner that the pivot changed directions halfway through the sequence; residuals exhibit relatively consistent sign and magnitude in each of the “forward” versus “backward” directions. This forward/backward behavior can be seen in Fig. 8, where there is a clear separation between the green diamonds until the turn-around occurs, at which point the two curves meet. The second effect is power cycling of the DPU, which results in homing the MDIS pivot. Each time the pivot is homed, a systematic error of 0 to 4 motor steps (up to $600 \mu\text{rad}$) is introduced. Overcoming this effect would require in-flight calibrations and segmentation of the harmonic drive model at each homing, as well as tracking of the pivot direction. Tracking would be burdensome, and because the pivot direction is not included in instrument telemetry this information would have to be modeled from the uplinked commands.

In-flight behavior of the resolver was assessed in parallel with development of the harmonic drive model. At the locations of the star calibrations, the RMS precision error was within 1 NAC pixel regardless of the direction of pivot motion. With a motor-step-to-angle model in hand, what was missing for calibration of the resolver was a resolver-to-motor-step relationship. Away from changes in direction, the forward and backward harmonic drive models are offset by a constant value. Therefore, a test was designed that would scan the entire pivot plane in order to obtain a resolver-to-motor-step relationship from telemetry. This scan was completed in both directions to provide two samples of the curve. Each curve was line-subtracted and adjusted to have zero mean so that the forward and backward data could be combined. A Fourier analysis of these data was completed to obtain a “predicted” resolver-to-motor-step positional offset as a function of pivot position, and the angle was obtained through the harmonic drive model. Two sets of star images were obtained at the 0° and -22.9° positions to validate the model and to verify absence of noise in the resolver values observed in ground data. The result of these flight tests and analyses is a resolver calibration that is unaffected by DPU power cycling, directional dependencies, or changes in the direction of pivot motion. The RMS error of 1.2 NAC pixels is smaller than uncertainties in reported spacecraft attitude (about 4 NAC pixels).

Reduction of Spacecraft Pointing Uncertainties. The remaining major factor contributing to pointing precision is knowledge of the attitude of the spacecraft itself. Similar uncertainties in the time phasing of spacecraft orientation have plagued image projection quality in past missions. The U.S. Geological Survey's Astrogeology Science Center has developed software – Integrated Software for Imagers and Spectrometers 3 (ISIS3) – that minimizes image registration issues by reducing pointing errors. The ISIS3 jigsaw application applies an iterative least-squares bundle adjustment of the projected locations of control points in multiple overlapping images to solve for errors in camera pointing. It also has the ability to solve for other parameters that affect mosaic quality, such as camera orientation (right ascension,

declination, and twist of boresight), spacecraft position (inertial Cartesian coordinates), and body (Mercury) radius. Other parameters could be determined by an analysis of jigsaw results, such as a spin rate and pole position (right ascension, declination). This approach corrects for uncertainties in spacecraft orientation as well as any uncorrected error in pivot pointing but cannot correct for errors in models of optical distortion.

An example of the overall improvement in an MDIS NAC image mosaic from M1 that results from more accurate knowledge of pivot position and optical distortion and use of the jigsaw application is shown in Fig. 9. The left view is an uncontrolled mosaic using the original model of pivot angle based on a linear transformation of pivot step counts and the original radially symmetric model of NAC optical distortion. Arrows highlight the most severe frame-to-frame mismatches in the mosaic. The right view uses an updated model of pivot angle, the newer radially non-symmetric model of NAC optical distortion, and control by means of the bundle adjustment.

3.3 Radiometric Calibration

Radiometric calibration of NAC and WAC images consists of the correction of raw data numbers (DNs) to I/F, or radiance divided by solar flux that would be scattered from a white, Lambertian surface. Key variables in the equation to perform this calibration are a model of the response to zero input light (dark level), the relationship between DN per unit exposure time and radiance (responsivity), pixel-to-pixel responsivity variations in each camera or filter (the "flat field"), and the bandpass used to model solar flux in each filter. Both responsivity and dark level are strongly dependent on CCD temperature. Accurate radiometric calibration is necessary but not sufficient to reconstruct surface spectral properties and to produce image mosaics from data taken at different illumination or viewing geometries. The other requirement for these objectives, a model of the light-scattering or photometric properties of Mercury's surface as a function of solar incidence, solar phase, and emission angles, is being determined from M1 and M2 images taken at phase angles of 30° - 130° ^[7]. The goals for radiometric calibration accuracy are 5% accuracy in absolute radiance; 2% precision filter-to-filter in the WAC, between the NAC and the corresponding WAC filter, or between images taken in the same filter at different CCD temperatures; and 0.1% precision of the flat-field correction.

Flat-field Correction. At the time of ground calibrations, two major issues were recognized regarding the flat-field correction. First, glint from the Be radiator created patterns in the window of the environment chamber where most calibrations were taken, so flat-field measurements were taken instead at room temperature with the cameras at the aperture of an integrating sphere. The resulting warm CCD temperatures led to artifacts from dark current noise (vertical streaks in Fig. 10, left). Second, the high f-number of the WAC creates shadows of particulates on the CCD cover glass approximately 25 pixels in diameter. It was expected (and later observed) that these shadows or "dust donuts" would move during launch, so the flat field was remeasured in flight using the onboard calibration target at low CCD operating temperatures. The target was illuminated slightly non-uniformly due to glint from the radiators and thermal blankets. To remove this effect, images of the calibration target were high-pass filtered and "draped" over a low-pass-filtered version of the ground flat fields to maintain the integrity of low-spatial-frequency variations. This approach produced greatly improved flat-field measurements (Fig. 10, right).

Responsivity. MDIS and VIRS spectra acquired during planetary flybys were used to validate the filter-to-filter precision of the WAC responsivity correction (Fig. 11). VIRS is a point spectrometer whose spatial coverage is built by scanning its FOV, so direct comparison of absolute radiances measured by the two instruments is still under study; agreement appears to be within ~10%. The Moon and Venus were seen at different phase angles and orientations by VIRS and MDIS, whereas the Mercury measurements were nearly simultaneous. Most differences in the normalized spectra represent random scatter of a few percent, except for consistently high and low radiances measured by the WAC's 430- and 480-nm filters, respectively. These two filters lie on a steep slope in the spectrum of the integrating sphere used to calibrate the WAC, so systematic error in the WAC calibration at those wavelengths is plausible. Hence empirical corrections to the two filters' responsivities of 1.07 and 0.93, respectively, were applied.

Intercalibration of the WAC and NAC was assessed using near-simultaneous image sets acquired during M2. Taking the WAC as "truth" from its intercalibration with VIRS, the NAC's responsivity was corrected empirically by 0.82.

CCD Temperature Effects. During ground radiometric calibrations, only limited sampling of the MDIS CCDs' -10°C to -45°C operating range was achieved. Measurements at $+23^{\circ}\text{C}$, -31°C , and -34°C supported only a linear model of the non-linear relationship of the CCDs' responsivity to temperature (Fig. 12). Temperature-dependence is greatest at the WAC's longer and shorter wavelengths, and weak to absent in the NAC and near 750 nm in the WAC. Application of this linear model to M1 and M2 images resulted in systematic over-correction of responsivity at low CCD temperatures, particularly during both approaches, making those parts of the flyby image mosaics appear "too red" (Fig. 13, top).

The correction for the effect of CCD temperature on responsivity was improved empirically using the M1 and M2 measurements, which span the lower half of the CCD temperature operating range. The procedure used leveraged the overlap of CCD temperature at M2 departure with ground calibrations, the nearly temperature-independent responsivity near 750 nm, and the prevalence on Mercury of "intermediate terrain"^[8]. We assumed that on average intermediate terrain is globally homogeneous. Spectra of multiple spots of intermediate terrain were extracted from the mosaics, normalized to the I/F values at 750 nm, and again normalized to the M2 departure spectra and joined with the ground measurements of responsivity relative to those at -34°C. Each filter showed an expected non-linear relationship of responsivity with temperature (e.g., Fig. 12). Application of the revised temperature dependence of responsivity to the flyby images resulted in a more uniform distribution of spectral properties between hemispheres (Fig. 13, bottom).

3.4 Scattered Light

The off-axis design of the NAC led to early concern about scattered light from out-of-field sources that reaches the CCD via a bounce off the inner wall of the NAC housing. Tests conducted during the second Venus flyby indicated a 2-7% contribution of scattered light from out-of-field sources for images of a uniform target several or more times the diameter of the FOV. Scattered light has not been obvious in NAC mosaics from M1 or M2, suggesting that the actual level is nearer to 2%.

In contrast, WAC images exhibit a complex pattern of scattered light around extended sources, reaching more than 5° from a radiance source. For large extended sources, such as Mercury from a range of several tens of planetary radii or less, the scatter appears diffuse and is highly wavelength-dependent in intensity (Fig. 14a). Approximately 1° from the limb of Mercury where Mercury just underfills the FOV, scattered light is about 1% of the brightness of Mercury near 430 nm, increasing with wavelength to 7% near 1000 nm. For extended sources only a few pixels in size, the scatter is resolved as a pattern of out-of-focus spots arranged into a quasi-rectilinear pattern, increasing in number and intensity with wavelength (Fig. 14b-14d). For larger extended sources, this pattern is produced from all illuminated parts of the scene and merges into a diffuse distribution. Ray-trace analyses indicate that the pattern originates from periodic structures on the CCD having the scale of the 14-μm pitch pixels. Light diffracts from the CCD's grating-like surface and then is scattered back onto the CCD from any of the 12 optical surfaces of the four lenses, the filter, or the CCD cover. The ray-trace modeling correctly predicts its strong wavelength dependence.

Scattered light in the WAC presents the greatest problem to scientific interpretation of MDIS data of any of the known instrument artifacts, because it mimics curvature of the spectrum due to a variety of Fe- or Ti-containing minerals or glasses^[8]. Briefly, longer-wavelength light tends to scatter from brighter to darker areas of an image, with the scatter having greatest effects near sharp brightness boundaries. Longer-wavelength light is preferentially lost from bright regions, creating a downward curvature in WAC spectra at wavelengths greater than 560 nm that resembles the curvature due to a 1-μm absorption feature in an Fe-containing silicate such as olivine. Longer-wavelength light is preferentially added to dark regions, creating an upward curvature in the spectra that resembles effects of some low-reflectance oxides. The magnitude of the effects locally approaches about 3-4% and is most severe at high-albedo features several to a few tens of pixels across, similarly sized occurrences of low-reflectance material, and near the limb. Spectral effects can be minimized by extracting spectra from areas at least several tens of pixels in diameter.

In-flight imaging of the scatter pattern in each WAC filter was acquired about one week after M2 when Mercury was 5 pixels in diameter (e.g., Fig. 14b-14d), and these data will provide the basis for a scattered light correction that follows published procedures^[9]. In the frequency domain, a WAC image afflicted by scattered light can be expressed as

$$F = I \cdot (H + \eta)$$

where F is the observed image, I is the same image but without the scattering, H is the distribution function of the scatter or the extended point-spread function, and η is a noise term. The problem of image restoration can be mathematically stated as: given F, find a best estimate of I. The approach being developed uses an optimal filter and can be presented as

$$I = F \cdot H^2 / [H (H^2 + K)]$$

where H is the norm of a complex number \mathbb{H} (the square root of the sum of the squares of the real and imaginary parts of \mathbb{H}) and K is a noise constant. H will be derived from the post-M2 Mercury images. K will be determined by an iterative process for each filter, selected to produce an image with reduced scatter and the least residual ringing in images after restoration.

3.5 Thermal System

Temperatures of the MDIS CCDs during cruise have remained within expected design tolerances. During the M1 and M2 flybys, insufficient time was spent near the hot planet to validate the full operation of the thermal system or the predictions from the thermal models. Mercury orbit insertion will be the first opportunity to fully validate the thermal design.

3.6 Engineering Performance

Overall instrument performance has remained nominal during nearly five years of flight. All telemetry points and alarm limits have remained unchanged since launch, with the exception of the filter wheel current monitor. Out-of-range alarms were determined to be related to noise induced by clock transitions occurring during the sampling interval. The MDIS flight software was revised once (Version 8) since launch, with a second revision (Version 9) to be loaded prior to M3 that includes enhancements in command capabilities to optimize the command loads to the spacecraft.

4. SUMMARY AND UPCOMING ACTIVITIES

MESSENGER and its science payload after nearly five years in flight continue to operate nominally. Nearly 2500 images have been acquired during the first two of three Mercury flybys, expanding the global coverage of the planet's surface from only 45% by Mariner 10 to approximately 90% by the two missions to date. On 29 September 2009, the third and final Mercury flyby will occur, and 18 months later the spacecraft will be inserted into orbit about Mercury.

REFERENCES

- [1] Solomon, S. C., R. L. McNutt, Jr., R. E. Gold, and D. L. Domingue, "MESSENGER mission overview," *Space Sci. Rev.*, 131, 3–39, 2007.
- [2] Hawkins, S. E., III, et al., "The Mercury Dual Imaging System on the MESSENGER spacecraft," *Space Sci. Rev.*, 131, 247–338, 2007.
- [3] McAdams, J. V., R. W. Farquhar, A. H. Taylor, and B. G. Williams, "MESSENGER mission design and navigation," *Space Sci. Rev.*, 131, 219–246, 2007.
- [4] Hansen, O. L., "Surface temperature and emissivity of Mercury," *Astrophys. J.*, 190, 715–717, 1974.
- [5] McClintock, W. E., and M. R. Lankton, "The Mercury Atmospheric and Surface Composition Spectrometer for the MESSENGER mission," *Space Sci. Rev.*, 131, 481–521, 2007.
- [6] Nye, T. W., and R. P. Kraml, "Harmonic drive gear error: Characterization and compensation for precision pointing and tracking," in 25th Aerospace Mechanisms Symposium, NASA CP-3113, pp. 237-252, 1991.
- [7] Domingue, D. L., B. W. Denevi, C. M. Ernst, G. M. Holsclaw, N. R. Izenberg, W. E. McClintock, S. L. Murchie, M. S. Robinson, and F. Vilas, "The photometric properties of Mercury's surface: Implications for color and spectral interpretations," in Programme with Abstracts, The Surface Composition of Mercury from Ultraviolet-Visible-Infrared Spectroscopy: State of the Art and Future Strategies, Università degli Studi di Parma, Parma, Italy, pp. 68-69, 2009.
- [8] Robinson, M. S., S. L. Murchie, D. T. Blewett, D. L. Domingue, S. E. Hawkins, J. W. Head, G. M. Holsclaw, W. E. McClintock, T. J. McCoy, R. L. McNutt Jr., L. M. Prockter, S. C. Solomon, and T. R. Watters, "Reflectance and color variations on Mercury: Regolith processes and compositional heterogeneity," *Science*, 321, 66–69, 2008.
- [9] Li, H., M. S. Robinson, and S. Murchie, "Preliminary remediation of scattered light in NEAR MSI images," *Icarus*, 155, 244–252, 2002.



Fig. 1. The fully assembled MDIS instrument showing key components. All electrical connections pass through a rigid-flex assembly rotary interface (twist capsule).




Fig. 2. Simplified functional block diagram of the MDIS instrument and redundant DPUs. FPU denotes focal plane unit, and S/C denotes spacecraft.




Fig. 3. Primary components of the MDIS thermal system. Diode heat pipes and wax pack thermal reservoirs clamp the CCD temperature during the hot pulse near the periapsis of each orbit.

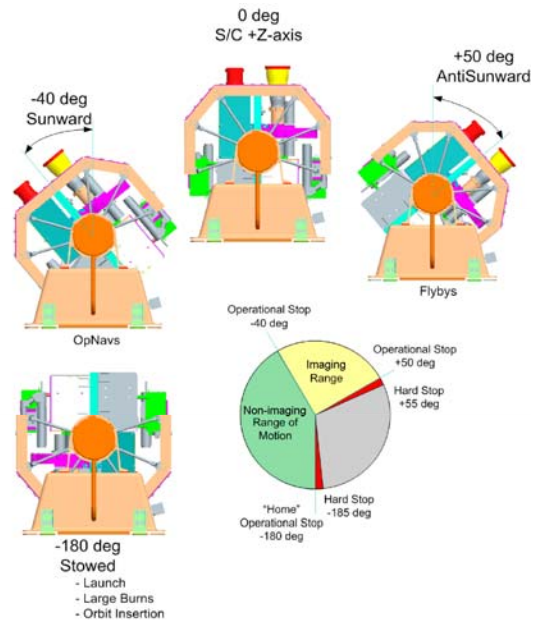


Fig. 4. Graphical representation of the pivot's scanning range.





Fig. 5. Results of two independent radiant-energy transfer models developed to bound the thermal environment in orbit. The left panel corresponds to a single 12-hour orbit for a Mercury true anomaly of 20°. The right panel shows on an expanded time scale the heat impulse from ~17:00 to 19:00. The red trace (no symbols) is the instrument model, the blue trace (squares) corresponds to the original spacecraft model, and the green trace (circles) is the spacecraft model with 20% margin.




Fig. 6. Images from planetary flybys conducted to date. (a) Approximately true color image of western South America and the eastern Pacific Ocean taken during the Earth flyby, constructed from 480-, 560-, and 630-nm images in the blue, green, and red image planes, respectively. (b) Same for Venus from the second flyby of that planet. No images were scheduled for the first Venus flyby because the event occurred during superior conjunction. (c) False color image from M1 departure, constructed using 430-, 750-, and 990-nm images in the blue, green, and red image planes. This combination enhances subtle color differences due to slope of the spectral continuum. The Caloris basin is the light-toned feature at upper right. (d) Approximately true color image from M2.




Fig. 7. Ray-trace analysis of the angular relations of corners of 64x64-pixel boxes projected to infinity, compared with the pixel locations on the CCD. The square grid represents the locations of the corners of the pixel boxes, and the "X" symbol shows their projected locations. The distortion of the projected locations has been exaggerated 20 times for clarity; the actual maximum distortion is 0.7%.




Fig. 8. Errors in pivot position resulting from the original linear transformation of motor step counts (black plus symbols), the non-linear conversion of pivot step counts that models behavior of the harmonic drive (green - diamonds), and the resolver as calibrated in flight (red triangles). Each NAC pixel is 25 μ rad, and each telemetry count of pivot position is 95 μ rad.




Fig. 9. Combined effects of the suite of improvements to pointing knowledge as applied to a NAC mosaic acquired during M1. (a) Uncontrolled mosaic using the original model of pivot angle based on a linear transformation of pivot step counts and the original radially symmetric model of NAC optical distortion. (b) Controlled mosaic using the updated model of pivot angle and the newer radially non-symmetric model of NAC optical distortion. Arrows are at the same latitude and longitude in each mosaic.




Fig. 10. Comparison of ground-derived (left) and flight-derived (right) flat fields in the 480-nm WAC filter. The vertical streaks in the ground-derived flat field result from inaccurate corrections to dark current in the room-temperature calibration images.




Fig. 11. Comparison of MASCS-VIRS and MDIS-WAC spectra of planetary targets. In each case the WAC spectrum is divided by the VIRS spectrum, and the ratio is normalized to have an average value of unity to account for differences in spatial sampling of the target.

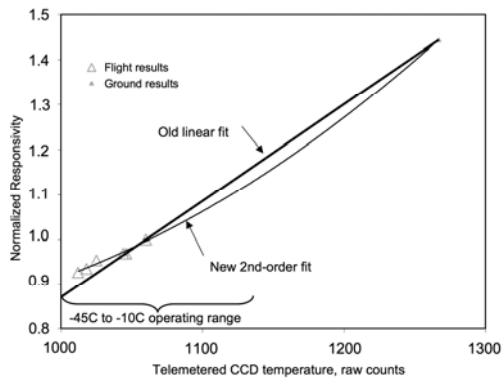


Fig. 12. Comparison of the original linear and newer non-linear models for temperature effects on CCD responsivity for the 900-nm filter. Note that only the lower half of the expected CCD operating temperature range is well-sampled prior to Mercury orbit.




Fig. 13. Effects of CCD temperature on the calibration of WAC multispectral imaging of Mercury. Both mosaics are global WAC mosaics with the 430-, 750-, and 990-nm filters shown in the blue, green, and red image plains, respectively. Saturation of the colors has been enhanced to illustrate calibration artifacts. (Top) Results from application of the original linear correction for effects of CCD temperature on responsivity. Parts of the mosaic taken at lower CCD temperatures are overcorrected for temperature effects at longer wavelengths, making them appear too red. (Bottom) Results from application of the newer non-linear correction for effects of CCD temperature on responsivity. Continuity between different parts of the mosaics is greatly improved. The residual red areas near the terminators result from photometric effects.




Fig. 14. Examples of scattered light in WAC images from the M2 flyby. (a) Whole-disk M2 departure image taken through the 900-nm filter. The bottom half of the image is stretched between $1/F$ values of 0 and 0.15, and the top half is hard-stretched between values of 0 and 0.01. (b) Distant 430-nm image of Mercury taken with a 10-s exposure. Mercury is approximately 5 pixels in size and surrounded by a minimal halo of scattered light; part of the image on the disk of Mercury saturated and is shown in black. (c) Distant 750-nm image of Mercury taken with a 10-s exposure. Scattering effects of the CCD and optics are apparent and form a pattern of diffuse spots. (d) Distant 900-nm image of Mercury taken with a 10-s exposure, with a more extensive pattern of diffuse spots.