MESSENGER: Mission Operations in Orbit at Mercury

Andrew B. Calloway, Richard G. Shelton, and Eric D. Melin

ABSTRACT

The orbital operations phase of the MErcury Surface, Space ENvironment, GEochemistry, and Ranging (MESSENGER) mission was stunningly successful, with >4 Earth years or nearly 18 Mercury years of uninterrupted activities without a single safing event in arguably one of the harshest operational environments in the solar system. The spacecraft returned >277,000 camera system images and millions of spectra and laser altimetry measurements, all while safely conducting 19 orbit-correction maneuvers, including a two-part adjustment in orbit period from 12 to 8 h. The team's long list of achievements included actively managing solar-array positions, heater settings, and spacecraft pointing on a weekly basis as well as conducting trailblazing helium-pressurant-only maneuvers at the end of the mission, plus an onboard clock rollover, >200 ephemeris builds, and 180 propulsive commanded reaction-wheel momentum unloading sequences over the 4105 orbits. This remarkable success can be attributed to the combination of substantial preparations and practice during the cruise phase of the mission, close collaboration and continuous communication among all of the project team members, and establishment of a reliable and repeatable cadence of flight activities and deliverables with few deviations except as required to work around geometry-driven events such as superior solar conjunctions and the evolving thermal challenges of each successive Mercury year.

INTRODUCTION

The MErcury Surface, Space ENvironment, GEochemistry, and Ranging (MESSENGER) mission, the first spacecraft orbital mission to the innermost planet, was known for safe and efficient operations from orbit. We provide here an overview of the preparations, processes, tools, and philosophies that enabled successful orbital mission operations. Topics include the proactive readiness approach taken during the 6.6-year cruise phase, ranging from the experience and practice of the six planetary flybys, to the building-block approach of day-in-the-life to week-in-the-life (WITL) to multi-WITL workshops that incrementally expanded objectives, to in-flight proof-of-concept tests. We also focus on the communications that were crucial to mission success, including weekly meetings and the products that were produced, reviewed, and disseminated among the team members, as well as the dual preliminary design review/critical design review approach to critical event planning. Most importantly, we describe what recurring activities were performed from the perspectives of real-time and mission planning operations. The MESSENGER team met all these challenges by maintaining consistency and rigor while not shying away from continually implementing innovations, revising techniques, enhancing reports with content and format, and updating ground system tools and scripts that together increased science return safely as resources were reduced through the extended mission phases.

CRUISE-PHASE PREPARATIONS FOR THE ORBITAL PHASE

The operations team worked closely with the science planning team and the ground system and software development teams to implement the tools and procedures for Mercury orbit, including interface control document revisions. The team also took full advantage of the three Mercury flybys and proactively pursued flight software upgrades for the instruments and the main processor (MP) that would benefit orbit-phase operations. A key component to orbit readiness involved a comprehensive building-block test program that broadened from days in the life to WITLs, culminating in a 4-consecutive-WITL test. These tests confirmed plans and readiness for the weekly cadence, and the orbital phase was in many respects treated as an entirely new mission from the cruise phase. The science team also generated 52 weekly test command loads that were gradually run through the operations faster-than-real-time StateSim constraint checking tool in what the team dubbed year-in-the-life validation testing. Lastly, to ensure readiness after Mercury orbit insertion, a series of orbit-phase concept-of-operations reviews were held for each spacecraft subsystem. These reviews were complemented by detailed planning for the first two command loads after orbit insertion and contingency planning for dealing with potential differences in post-insertion orbit parameters from those planned, including quick-turnaround rebuild options and time-biasing strategies.

Several in-flight tests were conducted in the later stages of the cruise phase to ensure readiness for orbit operations and adequacy of the tools and planned staffing levels. This testing was above and beyond the development and fine-tuning of instrument operations concepts and sequence building blocks on the basis of lessons learned from the three Mercury flybys. Many projects shy away from conducting such flight tests, but the MESSENGER team felt that specific tests would benefit from actual flight conditions in situations where hardware simulator runs could not provide sufficient confidence because of reduced fidelity or emulated behaviors. A good example was a progression of battery performance flight tests that culminated in a “long eclipse/hot-pole flyover/highest charge rate” end-to-end test, modeling an orbit at Mercury. Such tests helped ensure that there were no late surprises to operations planning that could have significantly changed thermal and power management, including instrument power cycling. A new automated strategy for conducting commanded momentum dumps (CMDs) had also been identified, and it was essential to verify that process in flight because the procedure would be used weekly through the orbital phase.

Another orbit concept that benefited from in-flight testing was a new scheme involving downlink rate stepping during Deep Space Network (DSN) tracking passes. When the concept was first envisioned in the cruise phase, it was unclear whether downlink interruptions at the steps would last seconds or minutes. By conducting flight tests in 2010, the team was ready to optimize the downlink during the orbital phase, tracking the DSN antenna elevation profiles. Details were ironed out before orbit operations thanks to those flight tests, including the number of steps needed to institute, how close to the end of a track to place them, and how to tune the CCSDS (Consultative Committee for Space Data Systems) File Downlink Protocol (CFDP) settings to minimize data-gap retransmissions. An Earth acquisition rotation RF characterization test was also conducted to ensure that the team practiced and knew what to expect for demotion contingency situations. The most significant flight test was conducted in the final week of August 2010. This was a composite test compressing WITL activities in time, including solar array, heater, and power setting management. The team opted to “trick” the onboard software with a modified ephemeris to simulate a spacecraft orbiting Mercury in a manner that was realistic yet would not impose undue risk to the spacecraft. The team also opted to institute planned orbital phase states while still in the cruise phase for trending and experience. For example, in November 2010 the accelerometers were powered on and left running for the remainder of the mission, and the cadence of weeklong command loads began before orbit insertion.

Each of the instruments was updated at least once with a flight software load, enhancing the capabilities for orbit-phase operations in some way. For example, the data processing unit (DPU) was updated with new command capabilities for image snaps and pivot steps, considerably reducing the overall number of commands that would be required in orbit. Even the MP software was updated with a crucial load in 2009. The MP memory map was originally partitioned with 256 time-tag bins, 900 small macros, and 124 large macros. The number of time-tag command slots was increased to 512 because testing revealed that 256 would be insufficient. An additional 20 large macros were also added by taking memory space from unused autonomy bin space. These changes increased the effective size of each command load resident in the MP, a step that was essential to the success of the orbital phase.

COMMUNICATIONS: MEETINGS AND REPORTING PRODUCTS

Timely and effective communications were one of the primary keys to the success and safe operations of the MESSENGER mission. This task was accomplished by the institution of several key recurring meetings, a rigorous preliminary design review/critical design review template approach to critical events planning, and quick responsiveness whenever any issue came up that affected the contents or schedule of a weekly command load. The four primary team meetings during the orbital phase were the weekly mission status meetings run by the mission operations manager (MOM), the Science Planning Group meetings led by the payload operations manager (POM), the power and thermal Mercury season planning meetings led by the operations power lead, and the command load review meetings led by the mission planners from the operations team. Secondary meetings included anomaly review and closure meetings co-led by the MOM and mission systems engineer, Configuration Change Board meetings led by the mission systems engineer, and planning meetings between the POM and the development and implementation team for the SciBox science acquisition planning tool (reference the Related Literature section), plus many other intra-team-level meetings. Every nonroutine, first-time, or critical event was assigned a preliminary design review/critical design review schedule, and all element leads were tasked to provide applicable charts for detailed walk-throughs in an internal peer review forum, including actions and test plan review. This adherence to process and schedule was a key component to safely conducting so many nonroutine events in orbit such as propulsive maneuvers, an onboard clock rollover, special calibration activities, and end-of-mission activities, to name a few.

Several important products were used for team-wide communications and information exchange, some of which were manually maintained and some of which were distributed by software such as cron jobs. These products included a daily mission timeline and a time-tag bias notification e-mail with historical trends and look-ahead predictions maintained by the MOM, file directory snapshots and hourly downlinked file lists, DSN track working schedules, as well as weekly reports and monthly presentations. The monthly reports captured information from all of the element leads, including propellant consumption statistics and solar-array degradation trends. Seasonal boundaries for solar-array offset positioning changes and instrument state changes were also publicized in table form, partitioned by dates and by high and low Mercury true anomaly thermal seasons. Other products included clock drift history and prediction plots (which eventually incorporated relativistic effects accounting for proximity to the Sun), Spacecraft Planet Instruments C-matrix Events (SPICE) meta-kernel trajectory update notifications with priority ordering overlay specifications, and attitude prediction and history reports. The latest attitude data status was defined in a SPICE meta-kernel and composed of attitude history from downlinked engineering files, short-term weeklong attitude predictions based on StateSim output reports, and long-term predictions based on the most recent science planning outputs from guidance and control (G&C) and SciBox. The navigation team required the weekly attitude updates for orbit determination.

Every week a predicted solid-state recorder (SSR) file playback list was delivered to and ingested at the Science Operations Center so that alarm messages would go out if a file was flagged as missing 72 h after the predicted downlink time. Orbit event files produced by the mission design team were ingested by the operations planning system to populate the command loads properly. Periapsis crossing time tables were provided by the navigation team for operations time-tag biasing calculations. DSN station allocation files were sent to the science developers for re-optimization runs along with recorder space snapshots to improve predictions of onboard SSR loading. The science developers, in turn, delivered optimal DSN track requests that the MOM tried to match as closely as possible. Command and data handling engineers were notified of onboard correctable single-bit errors. The operations team notified instrument teams when instrument memory dumps revealed comparison mismatches to correct potential single-event upsets via patch or power cycle. E-mail ListServs were created, including an overall operations distribution list and a SPICE users list.

REAL-TIME OPERATIONS AT MERCURY

Mission operations is the hub of any planetary space mission, interfacing with all team members on a project, including navigation, mission design, subsystem engineering, science planning, science operations, ground systems engineering, the DSN link controllers and engineering support, and project management, among others. The MESSENGER operations team was divided into three focus areas: real-time flight control operators, mission planners, and spacecraft/payload systems engineering analysts. With a lean core team of just 15 personnel, several of whom were part-time later in the mission, the operations team members continuously cross-trained so that every role and deliverable procedure had a fully trained backup who periodically practiced the primary cross-training tasks to ensure currency. A steady progression of task streamlining and automation implementation was also conducted over the course of the 4 years to accommodate extended mission downsizing.

The flight control team consisted of a lead flight controller (FC) and individuals who worked in shifts of two, aligned with the DSN track schedule. Support was eventually reduced to single-person coverage for the extended missions, and the operations team analysts served as the second set of eyes for all nonroutine subsystem-led activities. The flight control team interfaced with the DSN personnel for communications and troubleshooting when necessary; to monitor spacecraft health and produce logs, reports, and plots; and to send commands for planned tasks. There was little turnover over the years, but any new flight control team member had to undergo a training and checkout period culminating in a series of proficiency demonstrations charted by a detailed training checklist. The lead FC was responsible for maintaining the shift and coverage schedule and training for the flight control team.

The flight control team was also responsible for setting up the Mission Operations Center (MOC) whenever there was a planned critical event such as a maneuver, and for maintaining the MOC equipment and facilities. This role included performing voice checks of all MOC communications boxes before arrival of all of the support engineers, configuring each workstation for telemetry flow with a load-balancing checklist, and setting up the overhead screens with appropriate displays and timelines. The operations team also relied on ground monitoring software to ensure reliability and send alerts for thresholds on disk space, central processing unit usage, and related tasks. A backup MOC was configured in a separate building and was periodically used for real-time supports in order to maintain currency. As automation tools were incrementally brought online, including limit-violation e-mails and texts, flight control team members eventually developed a rotating schedule that allowed them to observe night and weekend tracks from home. Experience and reliability eventually allowed the MOM to be comfortable with automated commanding on unattended weekend tracks by enabling CFDP concurrent with turning on the transmitter at the DSN so that file delivery protocol handshaking would close out transactions onboard throughout those tracks rather than letting the transactions accumulate until Mondays. Flight control team members were always ready to travel in should any glitch affect the uplink or downlink on a case-by-case basis. Such events happened infrequently, and the benefits of automation far outweighed those rare instances, particularly as flight control team members began to support other projects part-time. All of the flight control team members were cross-trained to fill supplemental roles for the mission analysts. For example, one of the FCs was responsible for command and telemetry database management. Another FC eventually became the command and data handling and CFDP lead for the operations team. Another was in charge of automation development and maintenance. These added responsibilities and cross-training were essential for keeping such a small team experienced and focused.

RECURRING ACTIVITIES DURING THE ORBITAL PHASE

Configuration control was a major factor in the success of the mission. Every new or revised procedure or display page underwent a testing and review process by a second person before MOM approval for placement into the production system. Standardized forms and databases were a key part of this control process, including MOCRs (mission operations change requests). Every command to the spacecraft underwent similar scrutiny, with two-person validation followed by questioning by the MOM. These commands were handled by MURFs (MESSENGER uplink request forms). In addition, a table of preapproved maintenance commands and procedures was maintained in the MOC for routine activities to be referenced as needed.

Most orbit-phase tasks were performed on a recurring weekly cadence. For example, every Monday a G&C parameter block was delivered to the operations team, tested, and loaded to uniquely prepare for and target each Tuesday's weekly CMD. Mondays were also the weekly load transition days on the last hour of each DSN support, and confirmation of that transition was closely monitored. Tuesdays were not only the CMD days for reaction-wheel desaturation but were also the primary opportunities for loading the second half of each weekly command load. Wednesdays and Thursdays were the backup opportunities for loading the second half of the sequences. The weekly ephemerides were delivered on Thursdays, days when the following week's time-tag bias values were chosen by the MOM. Fridays were busy uplink days. The weekly spacecraft ephemerides were tested and loaded, along with the first half of the following week's command load. Fridays were also a second placeholder to conduct a CMD as needed. Each week, the MOM had to determine the order in which these critical long-duration uplinks should occur in order to maximize the use of the available DSN time. Saturdays and Sundays were backup days for the command load uplinks. All days included recorder playback with CFDP and routine maintenance activities such as refreshing the onboard command loss timer and conducting and reviewing engineering dumps and SSR directory listings. There were important activities that were less frequent or asynchronous, led by the mission analysts, such as solar-array off-pointing battery discharge characterization activities, G&C and power system parameter block updates, and activities associated with the orbit-correction maneuvers, such as the external precise oscillator deselections and reselections. Other examples included the memory object cyclical redundancy check and full dumps for ensuring that the hardware simulator was always synchronized with flight.

Command Load Uplinks

The most important flight control task was to uplink each new command load twice a week for >4 years, nominally on Tuesdays and Fridays. Because each segment took 2.5–4 h to radiate and the process had to be started from the beginning if even one uplink frame was dropped because of a transmitter or antenna pointing glitch, this activity was far from routine. This activity became even more complex when the segments were split into parts A, B, and C and had to be loaded across multiple DSN supports, to accommodate shorter tracks associated with the 8-h orbits. Each weekly sequence contained commands to off-point the spacecraft appropriately during each orbit while crossing over the hot planet, to avoid an autonomous safing and command load halting response to a hot-pole keep-out violation from thermal spikes off the sunlit side of the planet. If a command load ran out and the next half-load was not onboard and ready for the automatic hand-over, then the hot-pole keep-out condition would be violated during the very next orbit when the G&C and fault-protection tandem detected the hot planet crossing in the unsafe orientation. Similarly, a weekly ephemeris had to be loaded every Friday (with the weekends as backup opportunities) before the onboard ephemeris expired on Mondays, which would have caused the G&C system to throw a flag and request a safing mode demotion from the fault-protection system.

Time-tag Biasing

Each weekly command load uplink occurred 4 weeks after the science inputs were first compiled for that week. The orbit determination solution over those 4 weeks drifted and changed with perturbations, in particular the timing of predicted versus actual periapsis crossing times. As a result, a time-tag biasing strategy was developed to optimize command timing for improved instrument pointing that accounted for those multi-week drift conditions. The operations team worked with the navigation team to develop and receive a weekly summary of orbit periapsis crossing times spanning 5 previous weeks, parsed in a columnar fashion. By subtracting the most recent periapsis crossing times with that of 3 and 4 weeks prior and then calculating an average change over the course of a calendar week of orbits, two biases were calculated and published on a weekly basis. These biases could be either positive or negative integer seconds. The first bias value covered the Friday through Monday of the currently executing load, which was a week older than the one that would take over on that subsequent Monday. Those Friday biases were loaded in real time once the new ephemeris was confirmed onboard. The Monday bias adjustment would be valid Monday through the following Friday. That bias was always loaded on Friday in the form of a macro to be called on board at the start of each new command load.

Instrument commands executing out of the DPU were not subject to the biasing shift, however. As a result, the DPU memory space was used only when necessary for MP command volume reasons, and there was heightened awareness to place those commands at times when it was unlikely that a real-time bias change could affect an observation in undesirable ways, such as changing the relative timing or order between MP and DPU commands from the same observation activity. Time-tag biasing also had a major impact on management of the DSN key word files. Because the biasing affected only onboard commands, DSN key word files had to be actively managed to ensure that each week's biases were properly captured with the interleaving of ground-relevant commands that were not affected by the bias and those that were directly tied to spacecraft RF commands, to ensure that the proper order was always maintained. This task could be particularly tricky when there were large negative time-tag biases. Large biases did occur around major events such as superior solar conjunctions and orbit-correction maneuvers; otherwise, values were typically only in the tens of seconds rather than minutes.

Telecommunications Planning

Tailoring the RF communications system was also a large part of recurring orbit operations. Placeholder downlink rates were provided well in advance to the mission planners for the command load builds. As actual DSN tracks firmed up, the placeholder downlink rates were replaced with optimized rates tailored for each specific DSN antenna, some of which were significantly better than others. The RF lead would also analyze whether downlink rate stepping should be applied to any tracks in that command load sequence, on a seasonal basis, and would sit down side by side with the mission planners when it came time to insert all of that information into the baseline sequences. CFDP allowed for a more aggressive downlink margin posture, enabling the operations team to reduce the required downlink margin from 3 dB to nearly 0 dB. This change alone nearly doubled the downlink capacity of the mission overall. CFDP was an enabler for the rate-stepping strategy because it automatically requested retransmissions of the dropped data during the brief rate-change dropouts. The RF lead also needed to determine uplink rate and, in particular, the specific track on which that rate should change between 125 and 500 bps. There were also orbit geometries when the 70-m antennas could not be used because of solar aberration, and the RF lead had to keep track of those seasons as well. In addition, there were seasons when the solid-state power amplifiers on board had to be powered off for power and thermal management reasons, and the RF lead worked closely with the power and thermal engineers to coordinate the timing of such changes and ensure the proper configurations.

The RF lead also worked very closely with the radio science team members. The radio science team provided seasonal requirements to the MOM regarding when to collect additional per-orbit data by scheduling DSN tracks in alignment with periapsis crossings and Earth occultation ingress and egress times. The RF lead then had to determine when to place low-gain antenna (LGA) switching commands outside of high-gain antenna (HGA) times as the spacecraft slewed collecting science data, to best ensure a retention of signal lock during the additional LGA DSN coverage. Depending on geometry, some seasons required only one LGA configuration, but other seasons required multiple LGA swaps between a single pair of HGA tracks. Those seasons were more complex not only because of the additional analysis required but also because of the added command volume and the much larger and more involved DSN key word file builds and quality checking reviews.

The RF lead worked closely with the G&C team and the DSN network operations engineers as the Sun–probe–Earth angle changed with time. On a seasonal basis, the onboard phased-array model parameter block had to be toggled in coordination with spacecraft rotation to the opposite quadrant for Earth-pointing communications with the electronically steerable phased-array antenna, and a specific pre-track time had to be proposed and agreed upon. The DSN also then had to be notified about that timing and told to expect a different frequency consistent with the opposite transponder characteristics for the next track after the change. The RF lead also worked with the DSN in support of maneuvers, to choose appropriate carrier loop bandwidth settings for the receivers to best maintain lock through the maneuver. The RF lead also notified the DSN when thermal environmental effects changed the best-lock frequency beyond a predefined threshold, typically ±1000 Hz, to warrant an update at the DSN. This strategy was used so that the MESSENGER team did not need to conduct RF sweeps to reacquire signal after occultations or outages. As long as the best-lock frequency was tightly maintained as described, then the no-sweep concept of operations could be used, greatly simplifying operations in support of automation and reliability.

Fault Protection

An operational fault-protection scheme was devised for orbit operations that triggered Fault Protection Processor (FPP) autonomy rules at designated orbit true anomaly points in order to change the spacecraft rotation rate between a fast and a slow rate at known points in each orbit, should an Earth acquisition demotion occur. The fast rate was implemented for thermal mitigation reasons, and the slower rates were the windows where operations could intervene and recover the spacecraft back to operational mode. The placement of those rule triggers had to be modified after each orbit-correction maneuver because the orbit parameters had just changed. This activity was performed throughout the primary mission after maneuvers. Fortunately, the team never had to exercise a recovery from Earth acquisition in orbit, even though the team continually planned for and was ready to do so at any time. Other fault-protection changes were tailored for orbit as well, including automatic star tracker reset recovery commanding and thermal and power threshold monitoring response rules that adjusted solar-array positions, as well as autonomous mid-maneuver fuel tank switching logic.

CFDP AND FILE SYSTEM USE IN ORBIT

The MESSENGER spacecraft was the first developed by the Johns Hopkins University Applied Physics Laboratory (APL) to use an onboard file system for data collection. APL spacecraft prior to MESSENGER stored data sequentially to a random-access memory-based SSR and then downlinked blocks of memory using read/write pointers. Transmission dropouts or communication errors in those spacecraft were manually corrected (when they could not be tolerated) by commanding the recorder to replay the section of missing data. The decision to adopt a file-based system for MESSENGER was driven by predictions of a limited downlink bandwidth from the spacecraft, limited onboard processing power (insufficient central processing unit available to compress images during data capture), and the need to compress images to meet data-return requirements. An additional rationale for using a file system was the need

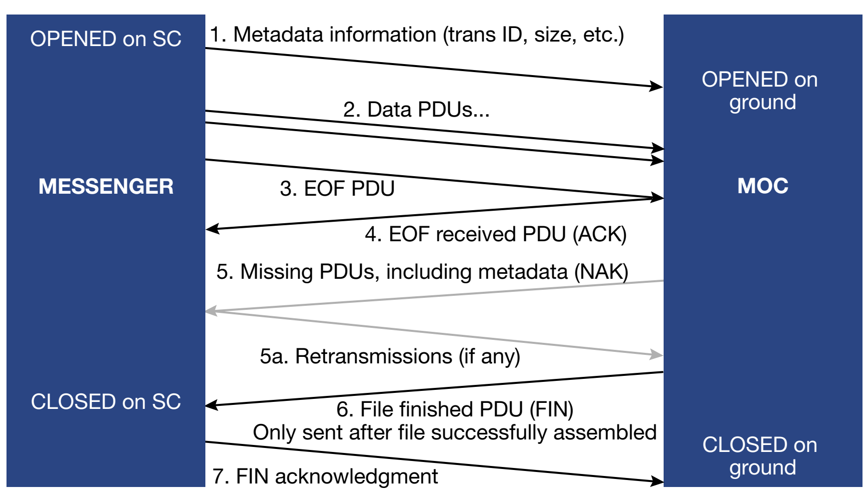

Figure 1. MESSENGER deferred NAK CFDP transaction. ACK, acknowledgment; PDU, protocol data unit.

to record contingency housekeeping information with the intention to rarely or never downlink it.

The file system architecture selection drove the need to define a protocol for downlinking files and managing the file system. In orbital operations, there was typically a range of 5- to 12-min one-way light-time delays between the ground stations and the spacecraft. The CCSDS developed a protocol called CFDP to specifically address single-endpoint file transfer needs for deep-space missions. The CFDP provides capabilities for file transfer and file system management commands such as delete, move, and copy similar to File Transfer Protocol. It is designed to minimize bandwidth usage and round-trips between communicating entities and is fully configurable to support different mission phases or operations concepts. MESSENGER used the reliable (Class 2) communications with deferred negative acknowledgment (NAK) shown in Fig. 1. In this configuration, metadata information about a transaction is sent, followed by the file data and an end-of-file (EOF) indication. If any data were missed, the ground system issued a NAK request to retransmit the missing chunk of data. A transaction was completed when both the flight and ground system confirmed successful delivery. The protocol also includes timers to resend EOF, NAK requests, and finished (FIN) messages to guarantee efficient reliable delivery with the minimal amount of messaging.

Only reliable deferred NAK transactions were implemented, and file system operations were handled externally to the protocol. The MOC ground software team originally chose to incorporate a CFDP implementation from the Jet Propulsion Laboratory. Years later during flight operations, the MOC software team replaced the Jet Propulsion Laboratory CFDP executable implementation with a C language CFDP library provided by NASA Goddard Space Flight Center (GSFC), which was more manageable and conducive to orbit operations.

Some extensions to the base CFDP functionality were made to better support MOC operations as enhancements to the GSFC library. All ground-software modifications were developed by APL and provided to GSFC for potential inclusion in future releases. One key need was the ability to ensure that transaction states across a MOC command workstation reboot or even a failover to the backup MOC command workstation persisted. APL also implemented a pause-and-resume timer functionality to avoid sending unnecessary NAKs due to active timers outside of DSN tracks. This step enabled tighter timer settings for more expedient file deliveries. Lastly, a light-time setting was added to the CFDP library. This addition allowed the base transaction timers to have a consistent value but allowed the actual time between retransmission requests to vary on the basis of current light-time values. The operation of the protocol lent predictability and timeliness to the data downlink and made it easier for the science teams to track their data status.

SPECIAL ACTIVITIES DURING THE ORBITAL PHASE

Although the team strove to maintain consistency and repeatability during the orbital phase as much as possible, MESSENGER was a mission of exploration, and it was understood that new information and new questions drove changes, especially for the extended missions. As such, special requests and accommodations were considered. Some were rejected, but most were approved. The orbit-phase concept-of-operations architecture was designed with flexibility so as to accommodate such requests, including specialized command sequences and changes in recorded data volume handling. Good examples of this flexibility included the ability to observe comets Encke and ISON, to collect Faraday rotation data from radio transmissions through the solar corona during solar conjunctions, and to conduct a delta-differential one-way ranging DSN observation request by NASA Headquarters near the end of the mission. Other accommodations included a substantial modification of the downlink priority scheme for the last extended mission and a request by the radio science team to add a multi-week campaign of short DSN supports aligned with Mercury periapsis crossings, which required broad negotiations with other missions across the DSN community.

MISSION PLANNING IN ORBITAL OPERATIONS

Command Load Building

The ultimate goal each week for the mission planners was to produce a safe and optimized command load that merged spacecraft engineering and housekeeping activities with payload data acquisition. Table 1 depicts a typical command load build schedule, with a three-person staggered rotation for shepherding those loads through the 2-week pipeline from build to merge to uplink. The load names were defined in a "yyddd" format. The command loads were nominally compiled as one Monday-to-Monday week, transitioning daisy-chain style in the last HGA hour of the Monday DSN tracks. Unfortunately, the weekly loads had to be split into two half-week upload segments because there were far too many commands in 1 week to fit in onboard memory, even after pre-orbit software enhancements. The location of the split was defined to be Thursdays or Fridays, chosen to balance, at least approximately, the number of commands in each half and to allow at least two real-time opportunities to load each segment to the spacecraft, avoiding weekends. The mission planners chose the split location on the basis of the actual track lengths and half-load sizes. The split choice was usually straightforward but occasionally had to be made in consultation with the MOM, who made the final call on the basis of whether other activities planned for uplink days could be in conflict or a particular week was very close to uplink availability margins regardless

Table 1. MESSENGER representative 2014 mission planner command load build schedule

		Mon	Tue	Wed	Wed	Wed	Thu/Fri	Fri
Week	YYDDD	Build Load	Finish Load	CLR	SPG	MOps Initials	Re-Run SS	Start Load
6 Oct	14279	14286	14286	14286		14300	14279	14293
13 Oct	14286	14293	14293	14293		14307	14286	14300
20 Oct	14293	14300	14300	14300		14314	14293	14307
27 Oct	14300	14307	14307	14307		14321	14300	14314
3 Nov	14307	14314	14314	14314		14328	14307	14321
10 Nov	14314	14321	14321	14321		14335	14314	14328
17 Nov	14321	14328	14328	14328		14342	14321	14335
24 Nov	14328	14335	14335	14335		14349	14328	14342
1 Dec	14335	14342	14342	14342		14356	14335	14349
8 Dec	14342	14349	14349	14349		14363	14342	14356
15 Dec	14349	14356	14356	14356		15005	14349	14363
22 Dec	14356	14363	14363	14363		15012	14356	15005
29 Dec	14363	15005	15005	15005		15019	14363	15012

Yellow, planner 1; green, planner 2; blue, planner 3. CLR, command load review; SPG, Science Planning Group; SS, StateSim.

of choice. With this vulnerability and concern in mind, the MOM worked closely with the DSN scheduler weeks in advance on a weekly basis to minimize those tough choices. The MOM negotiated with other missions to gain the best combination of daily tracks, track lengths, and station antenna diversity to ensure a successful load segment uplink before the current load segment ran out.

As shown by a representative week in Table 1, the operations team built the initial schedules each week on Wednesdays, incorporating the finalized DSN tracks. The planners also added in engineering activity blackout placeholders for science to work around and defined one of the double Tuesday and Friday tracks to be that week's blackout track for CMDs. Science operations was notified of initials readiness via the JIRA software tool. By Friday of the following week, the science and G&C teams delivered their constraint-checked and error-free instrument and G&C sequences to the planners. Occasionally G&C analysis showed that a science slew should be decomposed into two smaller slews for safety to avoid potential thermal keep-in zone overshoot. The instrument teams and G&C engineers also updated their review and notification steps in JIRA. The planner then started the command load build process by merging the initials and all of the science (one per instrument) and G&C inputs. The following Monday and Tuesday were then used to finish building that command load and to check the results using a suite of reports from mission operations software, including StateSim, a software simulator, and custom scripts. A second planner always performed independent checks on the load. The POM served as the liaison between the operations and payload teams to keep the pipeline flowing and help resolve any delays or issues.

Fridays through Tuesdays provided sufficient margin for the few cases in which the science input deliveries had to be delayed to keep the pipeline moving. When the operations team finalized a load, the science teams and G&C were tasked to review a set of reports, and a mission analyst would perform the uplink on the hardware simulator to ensure there were no issues with the product itself or the load splitting. The following Wednesday, the operations team gathered for the official command load review for that load, led by that load's shepherding mission planner. This meeting was the last opportunity for the team to review the comprehensive set of reports in a consistent way each week and for any final questions or discussions before the MOM would approve and sign the uplink form for that load or, on rare occasions, send it back for a modification. Wednesdays were also Science Planning Group meeting days, led by the POM, where future command load decisions and data-retrieval issues would be discussed and worked out. The operations team was well represented at those meetings. Lastly, on Thursdays or Fridays, the mission planner resynchronized the StateSim SSR model with the onboard SSR state each week. The mission planner reviewed the downlinked directory listing files and made adjustments to the StateSim model to agree with the real SSR when necessary, such as accounting for a DSN pass outage or images that were larger than predicted because of elevated temperature fluctuations in a charge-coupled device.

Planning the Command Loads

The mission operations team used the MESSENGER DSN track schedule to build the sequence generation software (SeqGen, from the Jet Propulsion Laboratory) requests for the HGA passes. These requests included maintenance activities and placeholders for burn events. The maintenance requests were validated by SeqGen flight rules and models and by a second-person manual verification checklist procedure. Once the initial requests were verified, a mission operations initial request file was delivered and a notification was sent to the POM and the science teams via JIRA.

Each science team's SciBox output was run through SeqGen with the flight rules and behavior models in order to validate the science requests. The valid science request files from each instrument sequencer were then passed on to the POM. The POM took the initial request file and all the science request files and merged them in SeqGen, looking for conflicts between the operations team requests and the overall science team requests plus conflicts between the science team requests, as early in the process as possible. The POM was provided preapproved shared-resources details such as SSR allocations and attitude ownership timing. Resource conflicts were resolved by returning requests to the instrument team for corrective action and resubmission. Once all conflicts were resolved and the merged requests were validated using SeqGen flight rules and models, those requests were then reviewed by the instrument engineers for health and safety verification. If any problems were discovered, the instrument requests were returned to the appropriate team for correction. The process repeated until all instrument engineers and the POM were confident they had a safe and clean product. It was seldom required, but the POM and MOM had authority to pull a problematic sequence from a load altogether if a timely resolution could not be achieved, so as not to risk delaying the process to the point of a missed uplink, because that was a spacecraft demotion risk. The approved science request files were delivered to the operations directories.

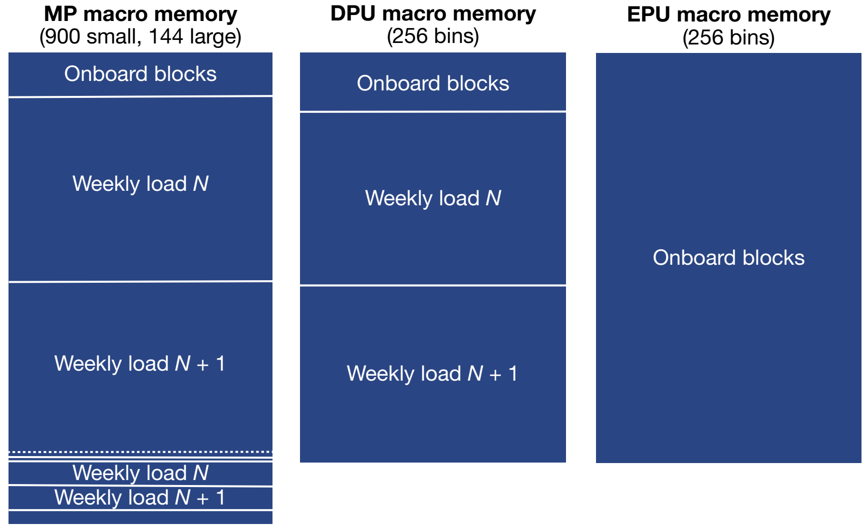

Figure 2. MESSENGER onboard memory allocations for mission planning.

The mission planners merged each week in SeqGen the delivered requests with the latest version of the initial requests. Those updates typically included any late changes to the DSN track schedule and updated burn requests from mission design. The planners occasionally had to adjust science requests in coordination with the POM in order to conduct the activities within the revised spacecraft resources (e.g., memory, power, SSR space, spacecraft attitude). Once a set of requests passed the SeqGen system-level verification, the mission planner built the command load in the uplink format and subsequently performed a second level of verification.

Building the Stored Command Sequences

SeqGen produced the ASCII spacecraft sequence file (SeqGen output, or SeqOut), which was the translation of the canned sequences into individual commands in time order. Next, the APL SeqPost program grouped the commands into transient macros and stored time-tag commands. SeqPost constructed macros on the basis of commands listed in time order for both the spacecraft and the instruments, and it eliminated duplicate macros for efficient use of onboard memory. A command load driver file was created to define the macros and time-tag commands, and a load script was created for use with the binary portion of the command load, which was created later in this process.

MP macros, commonly referred to as onboard blocks, were divided into two groups: semipermanent and transient. Semipermanent macros were loaded once to random-access memory and electrically erasable programmable read-only memory and used many times (with infrequent updates) for routine operations, including DSN track support and SSR operations. Transient macros were used and replaced with each active spacecraft load. Memory allocated for transient macros was divided in half so one range was updated while the other was in use with the active command load. DPU memory had a similar architecture. For each command load, the mission planner decided whether or not to create DPU transient macros. This decision was based on the size of the command load. If the commands would not all fit into MP macros, a large part of the Mercury Dual Imaging System commands would be packaged into DPU transient macro space instead. The planning team had to keep track of the DPU transient memory and clear out the memory if the DPU transient macros were to be reused in a subsequent load. The loaded time-tags executed the transient macros at absolute times. There were 512 available time-tagged command slots for the orbital phase. When the current half-week set of time-tags was exhausted, onboard autonomy loaded the next set of time-tag commands and expanded them into the bins for the next half-week of execution. Figure 2 shows the high-level onboard memory map allocations.

Verifying the Stored Command Sequences

After the conversion to uplink format, the command macros and time-tagged commands were loaded into the MESSENGER software simulator (StateSim). This process permitted an early faster-than-real-time checkout of the merged command sequence as it would be loaded and executed. StateSim modeled the loading of command macros and their time-tagged calls. Any problems in the command processing would be caught and reported via error and warning reports. An autonomy history report showed which autonomy rules tripped. StateSim modeled the spacecraft attitude on the basis of the attitude commands. StateSim computed attitude relative to the Sun line and targets of interest. Such a calculation was useful as a way to assess power and thermal conditions that the spacecraft would experience when executing the command load. The weekly ephemeris had to be loaded into StateSim for the proper Sun position computations.

Command Load Sizes: Uplink Time Limitations and Splitting

Each half-week command load could take 3–4 h to load depending on the Earth-range season, and operators had to start from the beginning if there was a transmission glitch. This happened several times during the cruise phase of the mission, when timing was more forgiving than in orbit. Therefore, the mission planning team designed a utility that could split a half-week further into parts A and B, to mitigate the need to complete a full uplink should there be a dropped frame or DSN error. If a frame was dropped during a part B load, then at least the reload would only have to start at part B and not all the way back at part A, potentially saving 2 h and eliminating the need to risk the uplink on the final load opportunity. This also happened several times in the orbital phase, justifying the utility's development. The real-time team was trained to run a special macro "clear" script for the specific case of loading a part B a second time because part A nominally contained all the precursor macro clear commands for the entire segment. To maximize the chance of success, the MOM published a ratio directive e-mail each week on how large each A/B piece should be; the weekly e-mail was tailored to the specific DSN track durations and placement of flight activities within tracks that week.

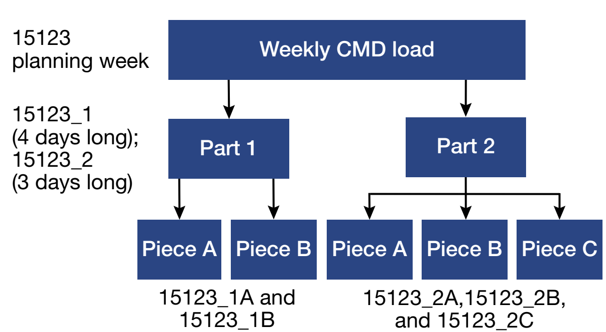

Figure 3. MESSENGER command load splitting strategy.

An additional complication was foreseen when the MESSENGER orbit period was reduced from 12 to 8 h. The HGA portion of the DSN tracks became shorter because of thermal constraints for Earth pointing in proximity to periapsis. Splitting a half-week into parts A and B was no longer sufficient for all cases. A modification had to be coded and tested to split a half-week load into three parts, A, B, and C. Using custom utilities and subjective experience, the MOM determined when a part C was warranted, as shown in Fig. 3; the MOM also determined the size of each segment and the times across multiple tracks during which the segments should be loaded, how they should be customized to fit within those tracks, and how to work around flight activities where commanding was not allowed, such as DSN station hand-overs and CMDs.

Command Volume Efficiency

Depending on the Mercury year and season, certain instruments, such as the Mercury Atmospheric and Surface Composition Spectrometer (MASCS), had to be powered off and on for thermal reasons during each orbit. Therefore, the mission operations team developed semipermanent MP and DPU macros for powering each instrument on and off. Each power on/off would cost only one command (an MP macro execute command) rather than bundling the equivalent number of commands per orbit for uplink within each command load. The operations team also developed a variable looping science macro strategy for two instruments. Those onboard macros could be repeated multiple times from one command, thereby realizing considerable uplink command volume savings. Custom reports were created to strip out the redundant loops for easier human review. The operations team also reviewed all of the subsystem commanding to determine whether repetitive commands could be identified and potentially replaced by semipermanent macros in the MP. For example, a macro was created to open the suite of daily files with just one execute command instead of 20 file-open commands. The operations and instrument teams similarly reviewed instrument command sequences to build macros, which worked well for the DPU and MASCS commanding. The instrument teams also looked for other ways to reduce command volume by small amounts, in situations where the sum total of small contributions was quite helpful. For example, several of the spectrographic instruments changed data-collection rates routinely at discrete parts of each orbit. Those teams were able to reduce the number of changes per day by having their instruments remain in coarser or finer data-sampling modes throughout an orbit, thereby reducing the number of commands without degrading the science and without recording significant amounts of additional data.

Telemetry Pipeline Management

Onboard packet and image files were stored in 10 directories, P0–P9. Files were closed and new files were opened once per day via the command loads to keep file sizes manageable for reasonable playback durations. These openings were tied to tracks but were manually inserted in the rare cases when there was no track on a calendar day. Files that had been successfully downlinked were moved to the TRASH directory automatically. To recover space for new data, the command loads deleted trashed files that were older than 5 days, along with un-promoted P7–P9 files, daily. A file filter table managed by the operations team was used to map the spacecraft application IDs (APIDs) to the files. Each APID represented a different data type for the instruments and the spacecraft subsystems. The file filter table maintained a row for each APID, and each APID could be routed to two files. For each file, the priority, instance, and data-sampling rate were defined by operations. The two-file definition allowed for a small subset of data to be given high priority to support health and safety spot checks, while the larger data set would be given lower priority.

The file filter table changed over the course of the mission, with major changes occurring when transitioning from the cruise phase to the orbital phase and from primary orbital phase to extended mission phases. File directory priority P0–P9 determined the order of file playback to the ground, as shown in Fig. 4, and was managed by the operations team in coordination with the science planning team. Default image locations were changed from P4 and split into P5 and P6 when data-collection science priorities changed. Files from within the same priority directory played back by "oldest onboard creation time first" protocol. Mission planners were able to model the order and amount of SSR playback via StateSim and passed this information to the Science Operations Center for file downlink tracking purposes.

Managing the data volume on the SSR was a continual process. The operations team (using StateSim) and the science team (using SciBox) independently modeled the volume of data on the SSR. The science allocation was conservative in the prime mission phase at 70% but was relaxed to >80% for the extended missions, when the risk of recorder saturation could be better tolerated. When the StateSim model predicted that the SSR might exceed 90%, the MOM conferred with the science team to obtain a second opinion and work out the mitigation strategy. If the onboard SSR ever reached 95%, it would stop recording until the percentage dropped again below that threshold. For long solar conjunctions of ≥5 days, when there would be no reliable contact with the spacecraft because of Sun interference, mission operations proactively worked with the science team and POM weeks in advance to reduce the number of images that would be taken and/or to selectively throttle the rate of the nonimage data collection in order to avoid the chance of recorder saturation. The science team implemented these data-reduction strategies into the science inputs for the loads that would execute during the conjunction periods.

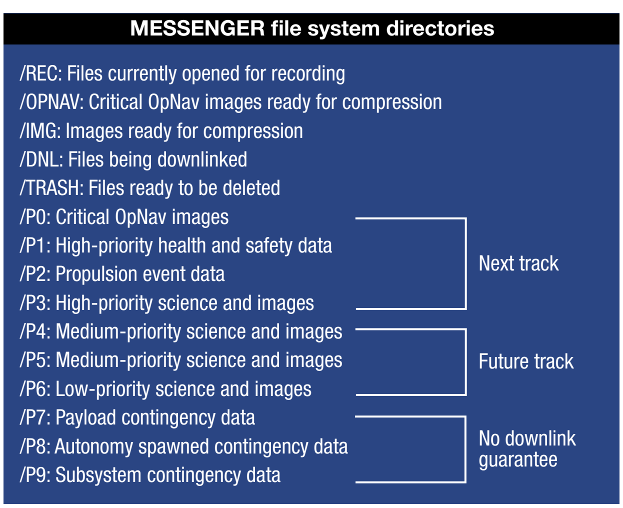

Figure 4. MESSENGER onboard file priority scheme.

CONCLUSION

Practice, communication, continual vigilance, optimization, and dedication were the hallmarks of each team within the overall MESSENGER project. In some respects, operating in such a hostile environment with an orbit that progressively changed in altitude and periapsis latitude in both sunlight and shadow worked to the team's benefit, because the team lived under heightened awareness at all times and complacency was never an issue. Because no two Mercury years were alike, all team members regarded each 88-day Mercury year as a new challenge, to be scrutinized and planned uniquely and in great detail, while maintaining consistency in processes and communications. That is a far different paradigm than operating in Earth orbit, for example, where environmental conditions tend to be similar from orbit to orbit, day to day, and year to year. As a result, the MESSENGER mission is considered by many to have one of the best returns on investment of the Discovery-class missions. It is indeed hard to argue with that, given the data-return metrics, on-time Planetary Data System deliveries, the level of public engagement, the revised textbooks, and even a Mercury globe, none of which would have been possible without MESSENGER and the team's operational safety record.

ACKNOWLEDGMENTS: We acknowledge all of the MESSENGER teams that supported operations. We extend a special thank you to our core operations team, including our mission planners, mission analysts, and superb real-time flight team. They all contributed in some way to the content of this article. We further extend a thank you to our ground system and network engineers, our navigation and mission design engineers, our mission systems engineer and subsystem engineers, our APL and NASA project management, our principal investigator and our science planning and instrument operations colleagues, our DSN partners, our media and public outreach members, safety and mission assurance team members, and of course the public; we truly accomplished something special—the exploration of our innermost planet—together as one family.

RELATED LITERATURE

Bedini, P. D., Solomon, S. C., Finnegan, E. J., Calloway, A. B., Ensor, S. L., McNutt, R. L., Anderson, B. J., and Prockter, L. M., "MESSENGER at Mercury: A Mid-term Report," *Acta Astronaut.* **81**(1), 369–379 (2012).

Berman, A. F., Domingue, D. L., Holdridge, M. E., Choo, T. H., Steele, R. J., and Shelton, R. G., "Orbital Operations Planning and Scheduling for the MESSENGER Mission," in *Proc. Sixth International Workshop on Planning and Scheduling for Space*, Pasadena, CA, pp. 1–10 (2009).

Choo, T. H., Murchie, S. L., Bedini, P. D., Steele, R. J., Skura, J. P., Nguyen, L., Nair, H., Lucks, M., Berman, A. F., McGovern, J. A., and Turner, F. S., "SciBox, an End-to-End Automated Science Planning and Commanding System," *Acta Astronaut.* **93**, 490–496 (2014).

Cooper, S. B., "A Practical Application of Relativity to Observation and Control of MESSENGER Timekeeping," *Johns Hopkins APL Tech. Dig.* **32**(1), 388–394 (2013).

Cooper, S. B., "From Mercury to Pluto: A Common Approach to Mission Timekeeping," *IEEE Aerosp. Electron. Syst. Mag.* 21, pp. 18–23, (2006).

Flanigan, S. H., O'Shaughnessy, D. J., and Finnegan, E. J., "Guidance and Control Challenges to the MESSENGER Spacecraft in Achieving and Operating from Orbit at Mercury," in *Proc. 35th Annual AAS Guidance and Control Conf.*, Breckenridge, CO, paper AAS 12-096 (2012).

Holtzman, G. A., Begley, S. M., and Ercol, C. J., "Return to Mercury: An Overview of the MESSENGER Spacecraft Thermal Environmental Predictions and Analysis Process," in *Proc. International Conf. on Environmental Systems*, Bellevue, WA, paper ICES-2015-307, pp. 1–15 (2015).

Hudson, J. F., and Colwell, E. J., "Spacecraft Clock Maintenance for MESSENGER Operations," in *Proc. AIAA Space 2011 Conf. & Exposition*, Long Beach, CA, paper AIAA-2011-7183, pp. 1–10 (2011).

Kochte, M. C., Sepan, D. M., and Shelton, R. G., "Getting the Message to MESSENGER: Overview of the Weekly Planning and Sequencing of MESSENGER Orbital Activities," in *Proc. SpaceOps 2014 13th International Conf. on Space Operations*, Pasadena, CA, paper AIAA 2014-1915, pp. 1–10 (2014).

Krupiarz, C. J., Heggestad, B. K., and Carper, R. D., "Flying CFDP on MESSENGER," *International Telemetering Conf. Proc.*, San Diego, CA, paper 04-14-02, pp. 1–8 (2004).

Krupiarz, C. J., Burleigh, S. C., Frangos, C. M., Heggestad, B. K., Holland, D. B., Lyons, K. M., and Stratton, W. C., "The Use of the CCSDS File Delivery Protocol on MESSENGER," in *Proc. SpaceOps 2002 Conf.*, Houston, TX, paper T5-35, pp. 1–8 (2002).

Krupiarz, C. J., Artis, D. A., Calloway, A. B., Frangos, C. M., Heggestad, B. K., Holland, D. B., and Stratton, W. C., "File-Based Data Processing on MESSENGER," *Acta Astronaut.* **59**, 1071–1078 (2006).

Mackey, D. P., and Calloway, A. B., "Planning and Operations for Superior Solar Conjunctions During the MESSENGER Mission," in *Proc. Institute of Electrical and Electronic Engineers (IEEE) Aerospace Conf.*, Big Sky, MT, paper 6.0102, pp. 1–6 (2012).

McAdams, J. V., Solomon, S. C., Bedini, P. D., Finnegan, E. F., McNutt, R. L. Jr., Calloway, A. B., Moessner, D. P., Wilson, M. W., Gallagher, D. T., Ercol, C. J., and Flanigan, S. H., "MESSENGER at Mercury: From Orbit Insertion to First Extended Mission," in *Proc. 63rd International Astronautical Congress*, Naples, Italy, paper IAC-12-C1.5.6, pp. 1–11 (2012).

Ord, K. J., and Hibbard, K. E., "MESSENGER Power and Thermal Systems Operations," in *Proc. SpaceOps 2010 Conf.*, Huntsville, AL, paper AIAA 2010-2137, pp. 1–10 (2010).

Page, B. R., Bryan, C. G., Williams, K. E., Taylor, A. H., Stanbridge, D. R., Wolff, P. J., and Williams, B. "MESSENGER Navigation Operations During the Mercury Orbital Mission Phase," *Adv. Astronautical Sci.* **148**, 2825–2844 (2013).

Andrew B. Calloway, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Andrew B. Calloway is a member of the Principal Professional Staff in APL's Space Exploration Sector. He has 24 years of experience in space operations as of 2017, including geostationary Earth orbit, low Earth orbit, and planetary missions. He was the mission operations manager for MESSENGER from 2007 through end of mission in April 2015 and a core team member of New Horizons through the Pluto flyby encounter. He continued exploring the solar system when he joined the OSIRIS-REx (Origins, Spectral Interpretation, Resource Identification, Security, Regolith Explorer) team in April 2016, supporting the mission to return samples from asteroid Bennu as the mission operations manager. His e-mail address is andrew.calloway@jhuapl.edu.

Richard G. Shelton, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Richard G. Shelton of APL's Space Exploration Sector has been a mission operations analyst for the MSX (Midcourse Space Experiment) spacecraft and NASA's NEAR (Near Earth Asteroid Rendezvous), MESSENGER, and New Horizons missions. For MESSENGER, he was responsible for the operations for four of the seven science instruments, including the Mercury Dual Imaging System, and maintained key components of the mission operations software. His e-mail address is richard.shelton@jhuapl.edu.

Eric D. Melin, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Eric D. Melin of APL's Space Exploration Sector is currently the ground software lead for NASA's Parker Solar Probe mission. He was previously the ground software lead for NASA's New Horizons and MESSENGER missions. Eric has also developed ground software for NASA's Van Allen Probes mission. His e-mail address is eric.melin@jhuapl.edu.