New Horizons Mission Operations and Encounter Planning

Alice F. Bowman and the New Horizons Mission and Science Operations Teams

ABSTRACT

New Horizons was the first mission to Pluto. The spacecraft was launched on January 19, 2006, and flew by Pluto on July 14, 2015, returning historic images and data that revealed new insights about the Pluto system. It then flew by Arrokoth on January 1, 2019. But work on New Horizons began long before 2006, including many years of effort to design and propose the mission, build the spacecraft and its instruments, and develop and implement the mission operations concept. This article describes the mission operations concept for both nominal and encounter planning as well as anomaly resolution. It details the challenges of operating a spacecraft that will fly across the solar system and how the mission operations team met these challenges, including the technical hurdles in implementing science observation requirements into spacecraft commands; the balancing of spacecraft communication and science observation periods; the constraints imposed by the spacecraft subsystems; the distances (and thus time delay) between the mission operations center on Earth and the spacecraft; and the programmatic constraints to control mission operations costs for this long mission.

INTRODUCTION

New Horizons, the first mission to explore Pluto—and to date the only mission to do so—was launched on January 19, 2006, and flew by the Pluto system, including its largest moon, Charon, on July 14, 2015. During the flyby, the spacecraft captured critical data and images revealing an amazingly complex and beautiful world (see Weaver et al., in this issue, for more on the scientific findings). The primary mission took ~10.5 years from launch through playback of the data collected at Pluto. The first Kuiper Belt extended mission (KEM) included the flyby of the Kuiper Belt object Arrokoth in 2019, and New Horizons is continuing its exploration of the outer solar system and beyond during a recently approved second extended mission. If NASA concurs, the mission will continue far into the future (see Figure 1 and the article by Brandt et al., in this issue).

As shown in Figure 1, it took many years of operational designing, planning, testing, and verifying to accomplish the encounters. Not only did the team have to account for the long round-trip light time, or RTLT (on the order of 9 h at Pluto) and the almost decade-long journey, but it also had to consider and account for the

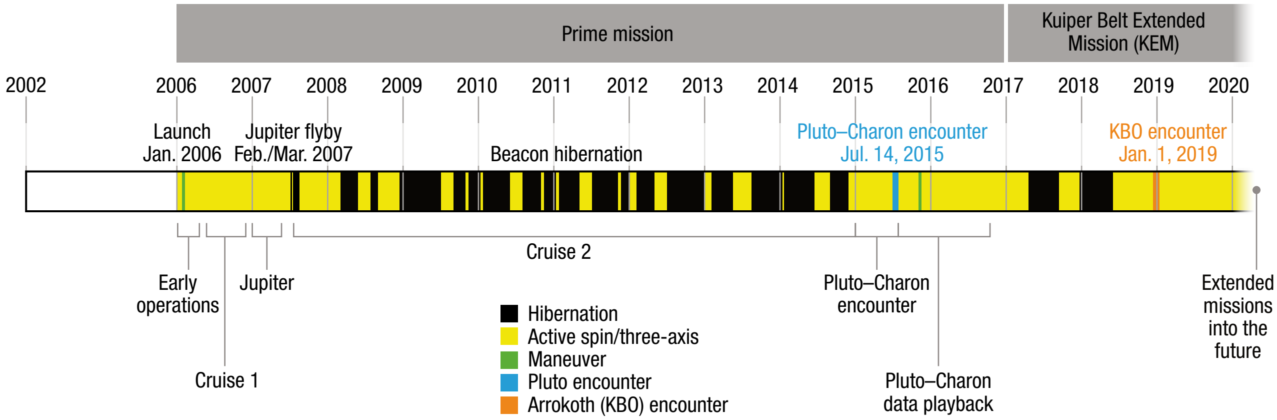

Figure 1. New Horizons mission timeline. This timeline summarizes the major flight activities from launch in 2006 to completion of the first extended mission—Arrokoth encounter and data return. The early operations phase included initial spacecraft checkout and a maneuver to correct for the small launch trajectory insertion errors. Instrument checkouts and preparations for the Jupiter flyby were completed during Cruise 1. The Jupiter flyby provided the additional energy required to reach Pluto by 2015 and enabled important observations of the Jovian system; hibernation phase-in tests were included while Jupiter data were being downlinked. During the second cruise phase (March 2008–December 2014), the spacecraft was primarily in hibernation, with short periods each year for spacecraft and instrument checkout, as well as very short periods of active operation to align the spacecraft antenna to Earth. The Pluto encounter phase began with a "wake-up" from hibernation in December 2014 followed by early observations of the Pluto system beginning in January 2015. The observation activities increased in frequency as the spacecraft approached Pluto; these observations provided navigation (and hazard) information for small trajectory corrections and increasingly valuable scientific observations. The "encounter sequence" covered the period from 7 days before until 2 days after Pluto closest approach (see the article by Holdridge et al., in this issue, for more details). Pluto data playback commenced shortly after the encounter and continued until completion in October 2016. The playback was interrupted while the spacecraft performed a maneuver modifying its trajectory to target the Kuiper Belt object MU69 (Arrokoth) in November 2015 and to perform a post-Pluto instrument calibration in July 2016. The extended mission/cruise to Arrokoth continued in 2017–2018 with the spacecraft returning to hibernation periodically until June 2018 when preparations for the Arrokoth flyby commenced with navigation observations to make final trajectory correction maneuvers (TCMs). The Arrokoth flyby occurred on January 1, 2019; downlink of the Arrokoth encounter data continued into 2022. During the extended mission, the spacecraft made observations of unresolved Kuiper Belt objects, and the resulting data were downlinked concurrently with the Arrokoth data. That process is continuing, along with observations by the heliospheric instruments, which have been on nearly continuously throughout the mission. These observations continue into a second extended mission and may continue beyond if NASA approves.

inevitable aging of the ground equipment and team members between the official start of the mission (2002) and the Pluto–Charon encounter (2015) and then the Arrokoth encounter on New Year's Day 2019. The success of these encounters can be attributed to the plans and procedures applied throughout the mission, as well as the team's communication, attention to the mission's human elements, dedication, and focus on accomplishing a singular goal: to explore the Pluto–Charon system.

TEAM LEADERSHIP

From the outset of the primary and first extended missions, mission operations team leaders emphasized the importance of fostering and maintaining good team dynamics by making sure that each team member understood that their job was critical to accomplishing the goal. For example, during weekly team meetings, team members were discouraged from sitting in the back of the room and were instead encouraged to sit at the table. This invitation was also extended to interns, who were considered full members of the team just like more experienced members. Mutual respect of everyone working on the mission, including those on the mission operations team as well as those on the other teams, was expected. Everyone needed to be comfortable sharing their ideas and concerns to make the encounters successful.

Traditionally, APL mission operations teams are composed of people with varied backgrounds, including in aerospace, computer science, software development, electrical engineering, astronomy, and physics, and this operations team (some members are shown in Figure 2) was no different. Although the New Horizons team's processes and procedures were based on those of previous missions, questioning and improvement were welcomed. Team members were encouraged to listen with an open mind during discussions, leaving egos at the door and adopting the attitude that there was always something more to be learned. A sense of humor and calm leadership, especially needed during times of high stress, helped team members fully focus on the task at hand. The team's expectations and dynamics contributed to the success of the New Horizons encounters.

Figure 2. Some members of the New Horizons operations team in the mission operations center (MOC). Shown here are flight controllers celebrating in the MOC at APL after they received confirmation from the spacecraft that it had successfully completed the flyby of Pluto on Tuesday, July 14, 2015. (Photo credit: NASA/Bill Ingalls.)

SPACECRAFT COMMANDING

There are two types of spacecraft commanding procedures: real-time and sequenced. Real-time procedures contain commands that are sent directly from the command workstation in the mission operations center (MOC) and executed immediately after the spacecraft receives and accepts them. The New Horizons MOC consists of a set of workstations that accept telemetry from NASA's Deep Space Network (DSN), are used to create the spacecraft commands, and include tools to verify the command sequences (Figure 3). Command sequences (sequences for short) contain commands that are time-tagged to execute at a specific time based on the clock onboard the spacecraft. The duration for a sequence varies depending on the spacecraft mode. Hibernation sequences can be as long as 1 year. Nominal sequences in spin mode and three-axis mode generally span 14 to 21 days. Encounter sequences span 9 days. A sequence is uplinked from the MOC command workstation, received and accepted by the spacecraft, and stored in spacecraft memory until the onboard clock matches the specified command time, at which point the command is executed onboard the spacecraft.

The team realized at the onset of the mission that many spacecraft tasks traditionally accomplished by real-time commanding would be better transitioned to sequenced commanding because of the long light-time delay, the resulting ease of verification, the inability to conduct pointed spacecraft operations concurrent with spacecraft-to-Earth communications, and power management. This was one of the first post-launch challenges for the New Horizons team.

The goal was to rewrite as many real-time procedures as possible as sequenced operations, with the target milestone being early in the second cruise phase or the period between the Jupiter gravity assist and up to 6 months before the Pluto system encounter. From commissioning forward, the operations team began transitioning real-time procedures to sequenced operations. Only a few solely real-time command operations remain for New Horizons.

Because the system of command sequencing was manual, the team needed extra time to design, develop, test, and verify the highly complicated and densely packed sequenced encounter observations. To address this, the mission operations team developed a two-phased approach, called Phase A and Phase B. Every command sequence had a Phase B. Mission-critical, intense, and complicated spacecraft periods also had a Phase A.

Phase B

Except for periods of spacecraft hibernations and encounters, the norm for New Horizons was a 14-day command sequence with a 9-week development period starting 1 week before the Preliminary Design Review (PDR) through the start of execution onboard the spacecraft. The total memory space devoted to holding command sequences was virtually divided in half,

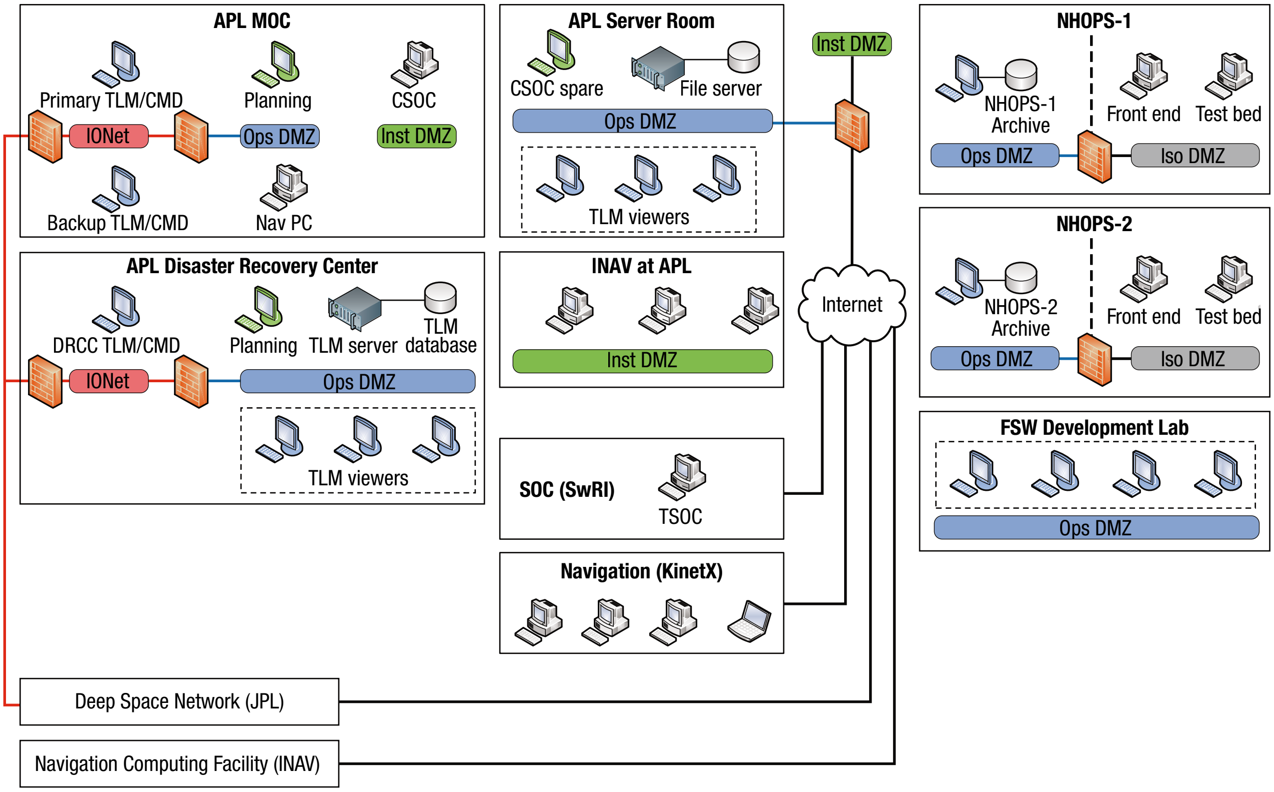

Figure 3. The New Horizons MOC and its interfaces. The New Horizons ground system is distributed across the United States. Its key elements are centered in the MOC at APL. This central hub is the primary control center for operating the spacecraft. The MOC workstations send commands (CMD) and receive telemetry (TLM) via NASA's Internet Protocol Operational Network (IONet) and DSN. The MOC is connected via the internet (with appropriate security interfaces via the demilitarized zone, or DMZ) to the Tombaugh Science Operations Center (TSOC) at Southwest Research Institute (SwRI) in Boulder, Colorado, and to the primary navigation team at KinetX. There is a separate set of navigation workstations at APL for an independent navigation (INAV) team that verifies the critical navigation analyses during encounter. A set of file servers at APL archive all mission data. Outside the MOC, but connected via the DMZ, is a set of workstations the flight software (FSW) development team used during flight software development and to make updates on a few occasions during flight. Two spacecraft flight simulators (NHOPS-1 and NHOPS-2), consisting of engineering models of the spacecraft hardware with computers to simulate the spacecraft dynamics, are used to verify command loads. In addition to the primary MOC at APL, a separate set of workstations in a separate building provides a backup in case an environmental or power failure should occur in the primary MOC during critical times. A separate set of workstations at the Jet Propulsion Laboratory (JPL), with limited command capability, provide a third level of reliability for the critical CORE command load.

with one half containing the currently executing command sequence and the other half holding the next or upcoming command sequence. This allowed a seamless transition from the end of one command sequence to the beginning of the next. A Critical Design Review (CDR) was held no later than 1 week before command sequence uplink to the spacecraft. During this process, science and subsystem objectives and activity placements were confirmed, the sequence command set was verified by software simulators, any first-time or critical events were simulated with hardware, and the command sequence (initial and final) was verified by the subsystem engineers, mission scientists, instrument engineers, and the science and mission operations teams. All command sequences executed on the spacecraft went through this development period, called Phase B. For those command sequences that also had a Phase A development period, the Phase B period started with the command sequence developed in Phase A.

Phase A

Because the Pluto system encounter was a flyby (meaning there was only one chance to collect measurements), an additional period was added for development of command sequences containing critical activities. Called Phase A, this development period started years before Phase B development, specifically 7 years before the Pluto–Charon encounter and just over a year before the Arrokoth encounter. The management team, with input from the various mission teams, determined which command sequences would undergo Phase A development on the basis of the criticality of the activities in the command sequence and the amount of time available in the schedule.

Table 1. Primary mission sequences that underwent Phase A development

Sequence Name[a]	Days from Pluto	Comments
15174	–21 to –11	Critical optical navigation inputs (OpNavs), penultimate TCM opportunity
15184	–11 to –7.5	Last TCM opportunity, last solid-state recorder (SSR) erasures
15188	–7.5 to +1.5	CORE-nominal sequence
15188 GIS	–7.5 to +1.5	CORE-alternate sequence
15197	+1.5 to –5.5	Departure science

[a] Sequences are named in year-year-day-of-year, or YYDOY, format based on their start-of-execution date.

Each command sequence that underwent a Phase A development period was assessed during a PDR 1 week after the Phase A start date, just as in the Phase B development period, but Phase A development continued for a duration decided by the science and mission operations teams. This duration depended on the complexity of the observations, the expected receipt of the input files (target ephemeris, etc.), the complexity of the software and hardware simulations, current mode of the spacecraft, and the time needed for team review.

At the end of the Phase A development period, the team held a delta CDR. One of the outputs of this review was a liens list containing items that did not fully meet the design of the science observation or account for the spacecraft constraints or the DSN tracking. The liens list was dispositioned by the principal investigator and the project manager based on team input, schedule, and risk assessment during a Configuration Management Readiness Review generally held a few weeks after the Phase A delta CDR. The Phase B period started with the command sequence developed in Phase A followed by a review of the Phase A liens list and which liens would be incorporated into the Phase B period.

For the Pluto–Charon encounter, four command sequences underwent Phase A development (Table 1). The first period the team tackled was the most scientifically dense and critical part of the encounter—the 24 h surrounding closest approach (July 14 11:49:57 or 2015:195:11:49:57), a subpart of the CORE sequence.

From this point, the operations team developed the commands outward in time to complete the CORE sequence, relying heavily on the guidance and control (G&C) team, the command and data handling (C&DH) team, the science team encounter lead, the mission systems engineer, the autonomy engineer, and the subsystem and payload engineers. Since the command sequence was driven by the onboard clock "time" and the pointed observations required accurate knowledge of Pluto's and Charon's location at the time the observation was made, all observation times were made relative to the time and position of the point of closest approach, as described in the article by Holdridge et al., in this issue.

During development of the CORE sequence, the team discovered that the allocated memory space was filled before the nominal 14-day command sequence could be completed. To address this situation, the CORE sequence was shortened and preceded by a small sequence (15284) comprising 4 days and using only 20% of the memory space. To meet the requirements of the sequence transition and memory allocation software, a similar 4-day 20% sequence (15197) was added at the end of the CORE period. This meant that the CORE sequence used 80% of the command sequence memory space and spanned 9 days.

CORE Sequence

The spacecraft autonomy system (see the article by Hersman et al., in this issue) had a special mode, the encounter mode, that would return the spacecraft to the encounter command sequence should a safing event occur during the critical period of the Pluto encounter, as there was no time to intervene from Earth and accomplish the encounter objectives. The start of this CORE sequence was chosen based on two things: (1) the estimated time it would take the team to recover the spacecraft from a safe mode event before entering three-axis encounter mode and (2) the placement of the Group 1 science objectives in the sequence (the required objectives; see the article by Stern and Krimigis, in this issue, for details). The team believed it would take about 4 days to recover from a fault it had not previously experienced, given the RTLT (8:49:03 h) when at Pluto. Because most of the Group 1 science objectives started at P–3 days, and to allow time for the spacecraft to recover from an anomaly, the start time selected for the CORE sequence was P–7 days. The CORE sequence would continue from P–7 days to P+2 days to ensure that all critical observations were made.

The team developed various recovery plans to speed up the process in the event of a spacecraft anomaly. For example, the Stream-Lined Autonomy Macro (SLAM; see the article by Hersman et al., in this issue) would be used if New Horizons went into safe mode after the CORE sequence had been uplinked, accepted, and burned into flash onboard memory for both C&DH computers. The estimate of 4 days for recovery was for the case where the fault encountered had never been experienced previously in flight, occurred before the CORE sequence was in flash memory, and was detectable by the autonomy software (i.e., it was a fault for which there was code onboard to handle). This condition meant that the team was recovering New Horizons from a spinning Earth-safe mode to a three-axis normal mode before the CORE sequence could engage.

Alternate CORE

Phase A development of the Pluto–Charon sequences started in early 2008. Originally, only four sequences were planned to have a Phase A development period, but this changed with the discovery of Kerberos (2011) and Styx (2012), two of Pluto's moons. These discoveries raised concerns that there might be dust orbiting Pluto that could impact and destroy the spacecraft. Actual assessment of such hazards could not be fully evaluated until a few days before the encounter (see the article by Holdridge et al., in this issue). As a result, a fifth Phase A command sequence, known as the 15188 GIS (Generic Inner SHBOT), was added to the plan. This fifth sequence was an alternate CORE sequence that put the spacecraft in the ram direction, or with the high-gain antenna in the direction of spacecraft movement so it could act as a shield for the 4 h in the part of the trajectory where New Horizons had the highest probability of hitting undetected moons/micro-meteoroids. The 15188 GIS sequence was developed so it could be flown on any of the four trajectories (nominal, SHBOT-1, SHBOT-3 and DIS, or Deep Inner SHBOT; see Figure 2 in the article by Holdridge et al. in this issue).

Those science observations planned for the 4-h period when New Horizons was in the protective ram orientation were either deleted or moved to a less optimum time period. The decision on whether to fly the nominal or alternate sequence was based on analysis of the last set of hazard images, conducted on July 2, 2015. In all, 20 versions of the CORE-nominal sequence were built. There were two in-flight tests, one in 2012 that tested the 22 h of the closest approach and one in 2013 that tested the full 9-day encounter minus the observations looking back to the Sun after the encounter. The Phase A development period was completed by December 2014.

Solid-State Recorder Usage

Solid-state recorder (SSR) management was one of greatest difficulties during the Pluto encounter. In addition to making sure the commands fit into the allocated memory onboard New Horizons, the team had to carefully plan the use of the SSRs. The two SSRs could each hold 8 GB of data. In addition to holding the science observations, this data volume also had to hold spacecraft attitude, small forces, and housekeeping (spacecraft and instrument) data, as well as accommodate data compression and packetization before data downlink.

To minimize risk, backup observations were planned for Group 1 (highest-priority) science observations, and both the prime and backup observations were recorded on different SSRs. The CORE-nominal sequence required/recorded on the order of 6.5 GB. SSR space can only be freed up in chunks of 0.5 GB. Days before the start of the CORE sequence, commands executing from onboard memory compressed data forward into a different location on the SSR to make room for the data collected during the 9-day encounter period (July 7–16, 2015).

Additional Constraints

Other Phase A development objectives were to ensure that the power being drawn stayed within the power margins established by the mission systems engineer, solidify the science observations period, and define the needed DSN tracking. At the time of the Pluto–Charon encounter, the radioisotope thermoelectric generator was projected to be producing an available power of 202 W, so the mission systems engineer and power subsystem engineer closely watched the projected power draw of the sequence and recommend action if needed.

DSN TRACKING

The DSN tracking periods needed to be of sufficient duration to send the required command sets to spacecraft memory, collect navigation data, and bring down the optical navigation (OpNav) and hazard data sets. This would ensure that New Horizons had all needed commands in spacecraft memory, was on the correct trajectory to the aimpoint with no known hazardous objects in its path, and would give the team enough knowledge to determine which CORE sequence to upload within the time allotted.

To negotiate and secure the needed DSN tracking, the mission operations manager (MOM) worked closely with the science team encounter lead, the NASA DSN scheduler, and NASA DSN management. Specific time blocks were chosen to support science observations, DSN tracking, and slew duration to and from Earth to minimize propellant use whenever possible. Since the DSN track schedule was not fully negotiated until 2 or 3 months before any actual track, the MOM made agreements with DSN management for how to accommodate the multiyear planning process for the New Horizons Pluto encounter.

Many discussions with DSN management and stakeholders helped reduce the impact to other DSN missions while ensuring that the New Horizons team had the data necessary to accomplish a successful encounter. Approximately 9 months before the Pluto encounter, the MOM, DSN scheduler, DSN mission planning manager, and scheduling process owner met in person with representatives for the missions whose DSN tracking would be affected by the New Horizons encounter. During this meeting, attendees developed alternative plans that helped mitigate the impact—for example, they worked out plans for other missions to avoid the intense periods of New Horizons' DSN usage by rescheduling activities or moving them to other antennas or networks (i.e., European Space Agency antenna complexes).

OPTICAL NAVIGATION

OpNav was required to target the aimpoint (also referred to as the closest approach time) and had to be intermingled with the approach science observations. Three OpNav campaigns preceded the daily OpNav campaigns, which started at P–40 days (June 4, 2015). Those periods were July 19–26, 2014, January 25–March 6, 2015, and April 5–May 15, 2015. For the 2015 OpNav periods, New Horizons entered a spin period for 2–3 weeks when both transmitters could be used (2TWTA mode) to bring down the OpNav data more efficiently. The original plan was to stay in three-axis mode from January 2015 through July 2015, but the team determined that, despite the propellant cost to spin up and then down, transitioning into spin mode was necessary to ensure that the OpNav data got to Earth as soon as possible since these data would inform whether or not a TCM was needed.

The team used OpNavs taken during the CORE sequence to determine whether a time shift offset needed to be applied to the sequence during closest approach. The operations team had to account for the possibility of a time shift offset on the order of ±450 s. All DSN tracks had to be scheduled to accommodate this type of shift. Invoking the time shift offset, depending on the direction, would have either delayed the beginning of the DSN track (plus time shift offset) or accelerated it (negative time shift offset), similarly affecting other track activities, except for SSR playback and uplink windows. Details of the navigation process and the accuracy of meeting the position and knowledge requirements are discussed in the article of Holdridge et al.

HAZARD MEASUREMENTS

The spacecraft searched for unknown moons, dust, and other obstacles on approach to the Pluto system and again on its approach to Arrokoth. These searches, known as hazard measurements, determined whether the alternate CORE-GIS sequence had to be used to protect the spacecraft while still allowing it to accomplish most of the Group 1 science objectives. For the Pluto system, these measurements were taken in 2015 on May 11–12, May 29–30, June 5, June 15–16, June 22–23, June 26, and July 1. The LORRI instrument took many images, which revealed dim, small objects when "stacked up" or added together by the hazard detection team. No hazards were detected, so the CORE-nominal sequence for the Pluto flyby was uploaded to New Horizons on July 4, 2015.

STAFFING

During the primary mission (the Pluto–Charon encounter), mission operations staffing varied from a low of ~7 full-time equivalents (FTEs) to a high of ~20 FTEs. Prelaunch planning called for staffing to drop to a low of ~3 FTEs 2 years after the Jupiter gravity assist until about 14 months before the Pluto–Charon encounter. With this plan, the operations team would have had only ~6 months to develop the encounter command set before Phase B planning started for the Pluto–Charon approach 1 phase (January 2015), and new team members would have to learn a great deal in a short time. The Jupiter flyby experience (which occurred just 13 months after launch) convinced the team to apply the lessons they learned to develop the Pluto flyby loads, a critical, longer, and more continuous process. Fortunately, after launch, additional funding was secured to add ~4 FTEs, divided between mission operations planners and flight controllers, to the post-Jupiter period.

The additional mission operations planners allowed the team to start Phase A development earlier. Phase A development periods coincided with New Horizons hibernation periods. The workloads of the additional flight controllers ranged from 0.1 FTE to 1 FTE. Those planned at less than 1 FTE ramped up as needed during busy periods before the encounter and to 1 FTE for the encounter. In this way, those supporting at even 0.1 FTE had little to no learning curve when they ramped up to 1 FTE. In addition, ensuring that staff members supported the mission at a low level brought more stability to the team and mitigated the knowledge management risk when team members needed to be added quickly during critical events, anomalies, and encounter periods.

ENCOUNTER TEAM TRAINING

In addition to two onboard flight rehearsals, the team participated in other critical tests. Program-wide operational readiness tests (ORTs) involved all members of the team, including mission operations. Once the ORT progressed, teams stood down if they were not needed. During many of the ORTs, teams practiced the process that would be followed during the encounter, and most ORTs were timed, since many, if not all, of these activities were time critical. ORTs involved the use of simulated data and the actual system tools and workstations. At the end of each ORT, areas for improvement were discussed, accepted, and added to the process. An ORT usually took about a week.

Mission operations team members also participated in green-card exercises. During this type of exercise, which was shorter than an ORT, usually ~4 h, cards describing an anomalous condition were handed out to participants (hence, these were also sometimes called roundtable exercises). Some participants were also given a card that directed how they should respond to the situation. For example, if the exercise focus was to ensure that backup staff members could carry out a particular task in the event of an anomaly, the primary staff members might be given a card saying they were vacationing on an island that did not have internet, meaning they could not assist the backup team in the anomaly response.

Mission operations team members also routinely participated in the testing of any software upgrades (e.g., C&DH flight, G&C, instrument, and autonomy software). This participation ensured that team members were familiar with how software was supposed to behave as well as how it was not. Mission operations team members also participated in testing DSN network upgrades, which included both software and hardware.

ENCOUNTER PROCESSING AND TEAM SHIFTS FOR KNOWLEDGE UPDATES

One of the critical activities for the Pluto–Charon encounter was the knowledge update (KU) process, described in more detail by Holdridge et al, in this issue. The LORRI instrument collected data on the spacecraft and then it was immediately queued for Earth downlink. Once the set of images reached the ground and became available for analysis, the navigation, science, science operations, and mission operations teams determined whether a time shift offset needed to be applied to the onboard encounter commands. Because ~12 h elapsed from the time data were received on Earth for analysis to the first uplink opportunity to send the KU to New Horizons, the planning team was divided into three shifts of two planners each. This ensured a rest period for each shift and that team members had the required training and expertise to process, evaluate, test, verify, and uplink the data. The MOM and deputy MOM (DMOM) worked 12-h shifts.

During normal operations, two flight controllers were on a shift. During the encounter, a third person was added to minimize work backlogs and handle extra activities. The shift change for these three flight controllers was staggered to ensure robustness and knowledge handover between shifts.

LONGEVITY PLANNING

Significant technology changes happen quickly, in a matter of years, and components on Earth age much faster than components in space. For New Horizons, the long flight time to reach the Pluto system, combined with the need to maintain a small team over the course of some 10.5 years (for the primary mission) suggested that the mission required a longevity plan. This plan recognized the need to refresh ground system hardware and to have aids to retain corporate memory.

The New Horizons team developed the longevity plan, and NASA accepted and funded it in 2002. The primary mission encounter of the Pluto system would be in July 2015 at the earliest if Jupiter was used as a gravity assist and as late as July 2020 if the spacecraft followed a direct trajectory from Earth to Pluto.

Given the amount of time between development of the supporting ground station and the Pluto system encounter and subsequent 18 months of data return, a refresh of the ground system equipment was scheduled. The ground system computers and associated hardware and software were replaced ~3.5 years before the encounter. However, even with this planned refresh, some components could not be replaced because they were incompatible with new technology. To solve this challenge, a longevity technology archive was set up to house working ground-system components that were retired on other missions and could be used on New Horizons.

Most programs have one hardware simulator that mimics the full capability of the spacecraft. Because New Horizons was a flyby mission with a long cruise to the Pluto system, in 2008 a second hardware simulator with leftover parts and test beds was built to ensure that a hardware simulator was always available to test and verify spacecraft commands and support spacecraft recovery.

FACILITY REDUNDANCY

Because this primary mission was a flyby, the team took additional measures to mitigate the risk of a problem that would prevent data collection during the flyby. All missions have a backup MOC where the team can quickly relocate to continue communication with the spacecraft. Located on the APL campus, the New Horizons backup MOC had its own commercial power feed, hardware, generator, and interface connections to the DSN. Additionally, a remote backup MOC located at the NASA Jet Propulsion Laboratory in California could be used in the event of a catastrophic event on the East Coast affecting both the primary and backup MOCs at APL. This remote MOC was stood up ~1 year before the encounter, and a three-person team visited quarterly to test and confirm that the remote MOC was able to send commands and receive telemetry from all the DSN complexes. This ensured that if the primary and backup MOCs were unable to upload the set of encounter commands by P–7 days, a truncated set of encounter commands could be uploaded from the remote MOC by P–4 days. Thus, a multilayered system for commanding the spacecraft minimized the risk of a hardware failure impacting success during those critical days just before the Pluto encounter.

AN UNEXPECTED CHALLENGE

The approach to Pluto was going as planned until an unexpected spacecraft reset occurred. The last of three TCMs was completed on June 30, 2015, and subsequent analysis showed that no additional TCMs were needed. There had been a placeholder TCM on July 4 in case analysis revealed a need for one more, but by July 2, 2015, the hazard team had completed analysis of the final hazard observation data and had found no detectable hazards in New Horizons' path. Therefore, the management team gave direction to fly the CORE-nominal command sequence. On July 3, 2015, the CORE-nominal command sequence was uplinked and burned into flash memory on the remote terminal C&DH computer (backup computer). When the needed command memory space on the bus controller C&DH computer (main computer) became available early on July 4, 2015, the CORE-nominal command sequence was uplinked to the main computer.

The 2-h uplink radiation of the CORE-nominal command sequence started at 4:21 a.m. EDT on July 4, 2015. At 1:10 p.m. EDT (an RTLT later), the mission operations team watched as the commands were received by the main computer. Unexpectedly, at 1:54 p.m. EDT, communications with New Horizons were lost. The ground system and DSN antenna complex configurations were verified as correct, and redundant hardware was requested to ensure that there was no hardware failure. The MOM notified teams of the anomaly, and an anomaly review board (ARB) convened at 4 p.m. EDT. Over the many years of cruise, the team had experienced a handful of anomalies requiring the ARB to be stood up, resulting in a well-honed and successful anomaly resolution process with team buy-in. This process was followed without exception.

At 3:11 p.m. EDT, communications were reestablished with New Horizons. The spacecraft had experienced a bus controller swap (the main computer was autonomously demoted to serve as the backup computer and the backup computer was promoted to serve as the main computer) and went into safe mode. As a result, New Horizons was spinning at 5 RPM with no active command sequence, and because of the main computer swap, the uplink polarity was reversed. Subsequent analysis showed that the bus controller swap was caused by the main computer processor being unable to service two very intense operations—compression of SSR data and burning the CORE-nominal command sequence to flash—at the same time. The autonomy software detected that something was wrong with the main computer (i.e., the software could not detect the main computer's heartbeat) and, in response, it made the backup computer the main computer, the nominal response.

The ARB meeting began with summaries on the state of New Horizons and known information about the anomaly, followed by a question-and-answer period. The management team conveyed that the priority was to start the 9-day CORE-nominal command sequence as planned on July 7, 2015. Science observations scheduled for July 4–7 were not high priority and did not require rescheduling or an attempt to keep them on the initial timeline. With these timeline constraints, the recovery/mission operations team developed a recovery plan.

Allowing for the durations of the various recovery procedures, the RTLT, and procedure test and verification, the recovery/operations team determined that there were four RTLT chunks of time between the current time and the start of the CORE-nominal command sequence. Three of these were used for recovery and the fourth was kept as margin. At approximately 10:30 p.m. on July 4, a detailed recovery plan to promote the spacecraft out of safe mode was presented to the stakeholders (defined by the ARB) and approved. The first set of recovery procedures was sent to New Horizons on July 5 at 3:40 a.m. EDT. Although it was not part of the spacecraft recovery plan or a requirement, the team was able to reschedule some data recovery commands in progress at the time of the safe mode demotion, thereby preserving the data. Despite this, part of an observation data set that was being compressed at the time of the safe mode demotion was lost and unrecoverable.

A third ARB meeting was held on July 5, 2015, at 4 p.m. EDT to summarize the activities and status of New Horizons since the previous ARB. During this ARB, the previous primary (now backup) computer was reported to be healthy. Because of the operational time and confidence, along with uplink polarity, the decision was made to switch back to this computer for the 9-day flyby. The switch was confirmed onboard the spacecraft at 2:43 a.m. EDT on July 6, 2015. At ~9:30 a.m. EDT, the operations team confirmed that New Horizons had entered three-axis mode and the six SSR 1 segments needed for the 9-day encounter were erased, completing the last recovery steps for transition into the CORE encounter sequence.

During this recovery process, the team tried to maintain a working environment as similar as possible to that of past recovery efforts. At times this was hard because of the magnitude of the external attention on the mission and the stress the team faced. It took fortitude to establish and enforce boundaries on the team workspace. Equally, it took restraint of those not normally involved in the detailed recovery process to step back and let the spacecraft and operations teams (recovery team) work through the established process. Rest periods were built into the recovery plan, but they were hard to enforce, so the compromise was to establish rest areas where recovery staff members could step away to rest and relax. Those who were not directly working the recovery process made sure that recovery team members had food, allowing them to continue their work. Additionally, recovery team members were sequestered from the public media, which was handled by mission and APL management—this is just one way that management supported the recovery team during an intense time.

REACHING THE GOAL

The CORE command sequence started as planned on July 7, 2015, at 12:00 spacecraft time. The anomaly had resulted in a loss of 2 days, 22 h, and 38 min of nominal New Horizons activities; no losses were Group 1 science observations. From July 7 through July 12, critical OpNavs were downlinked, processed, and evaluated by the KU team. Each day, all the KU stakeholders met to determine whether to uplink the KU. Since all the analysis indicated that New Horizons was within the error margin of the aimpoint, no KU was uplinked.

The most treacherous part of New Horizons' journey through the Pluto system was the 22 h surrounding closest approach; it was also the most scientifically rich. To mitigate the risk of something unexpected happening to New Horizons during this time, sample data sets representing the best observations from each instrument, called the "failsafe" data sets, were downlinked to Earth on July 12 and 13. New Horizons could not contact Earth while observing the Pluto system, so on July 13, the MOC received the last bit of telemetry from New Horizons at 11:17 p.m. For the next 22 h, New Horizons made numerous scientific observations using all the instruments onboard. From 1:20 a.m. through 6:21 a.m. EDT on July 14, mission operations, Radio Science Experiment (REX) and DSN personnel directed uplink signals to New Horizons, at one point training 7 of the 13 DSN antennas toward New Horizons at Pluto. These signals would reach New Horizons a one-way light time, or approximately 4 h and 25 min, later, supporting the onboard REX observation commands (see the articles by Fountain et al. and Weaver et al., in this issue). On July 14, 2015, at 8:53 p.m. EDT, communications were reestablished with New Horizons as planned. The MOM conducted a subsystem status poll during which all subsystem engineers reported nominal status. Recorder status bits indicated a nominal data collection. New Horizons had successfully transited the Pluto system. Sixteen months were required to bring down all the data the spacecraft collected during the encounter.

SUMMARY

A relatively small team accomplished the first reconnaissance of the Pluto system on a budget of ~$720 million from mission award (2002) through retrieval of the Pluto system science data (October 2016). With only 170 lb of hydrazine and an extremely accurate launch injection by the Atlas V, only eight TCMs and one gravity assist (Jupiter) were required. Over the course of the primary mission, the team successfully conducted 18 hibernation periods during which they prepared the five Phase A sequences, including at least 20 versions of the nominal CORE sequence, and recovered from six red beacon modes (denoting an off-nominal condition) and four safe modes, with the last occurring just 3 days before the start of the 9-day encounter. During the final approach to the aimpoint, the team flawlessly processed six critical OpNavs that each fully met the tight time constraint. Even though the last critical OpNav showed the need for a time adjustment, the team decided not to uplink a KU because the measurement was still within the error margin. Post-encounter ephemerides reconstruction showed that, at a distance of 4.7 billion kilometers from Earth, New Horizons missed the aimpoint by just 45 km and 88 s (early), which was well within the error margin. Pluto system data revealed unimagined science. The process developed for the Pluto flyby enabled a successful flyby of the newly identified Kuiper Belt object Arrokoth in January 2019. The Arrokoth flyby was the most distant encounter of a solar system body by a spacecraft, at 6.5 billion kilometers, and had an accuracy of 40 km of the desired aimpoint and was within 23 s of the desired time. The spectacular success of New Horizons is a testament to the careful planning of the management team, the vision of the science team, and the excellence of the engineering and operations teams, all working cohesively to do whatever it took to reach this one goal.

ACKNOWLEDGMENTS: This article highlights the years of hard work of the New Horizons mission and science operations team members. Work on the New Horizons mission was performed under NASA contracts NAS5-97271/TO30 (APL) and NASW-02008 (SwRI).

Alice F. Bowman, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Alice F. Bowman is the New Horizons mission operations manager in APL's Space Exploration Sector. She has a degree in chemistry and physics from the University of Virginia and has more than 30 years of experience in space operations. In addition to being the mission operations manager for New Horizons, she supervises APL's Space Mission Operations Group, leading approximately 50 staff members who operate deep-space and Earth-orbiting spacecraft, including NASA's TIMED (Thermosphere, Ionosphere, Mesosphere Energetics and Dynamics), STEREO (Solar TErrestrial RElations Observatory), New Horizons, Parker Solar Probe, and, recently, DART (Double Asteroid Redirection Test). Alice's experience also includes national defense space operations, systems engineering, program management, space systems, and space instrument development. Asteroid 146040 Alicebowman, discovered by Marc Buie in 2000, is named after her. She is a member of the Society of Women Engineers, which recently presented her with its prestigious Resnick Challenger Medal, is an AIAA associate fellow, and has served on the International SpaceOps Committee since 2009. Her email address is alice.bowman@jhuapl.edu.