The Contingency of Success: Operations for Deep Impact's Planet Hunt

Richard R. Rieber
818-304-2831
Richard.R.Rieber@jpl.nasa.gov
Robert F. Sharrow
818-354-5020
Robert.F.Sharrow@jpl.nasa.gov
Jet Propulsion Laboratory; California Institute of Technology
Pasadena, CA 91109

Abstract—The Deep Impact Flyby spacecraft completed its prime mission in August 2005. [1][2]It was reactivated for a mission of opportunity add-on called EPOXI on September 25, 2007. The first portion of EPOXI, called EPOCh (Extra-solar Planetary Observation & CHaracterization), occurred from January 21, 2008 through August 31, 2008. Its purpose was to characterize transiting hot-Jupiters by measuring the effects the planet has on the luminosity of its parent star. These observations entailed using the spacecraft in ways it was never intended. A new green-light, success-oriented operational strategy was devised that entailed high amounts of automation and minimal intervention from the ground. The specifics, techniques, and key challenges to obtaining the 172,209 usable science images from EPOCh are discussed in detail.

TABLE OF CONTENTS

1. INTRODUCTION

The Deep Impact spacecraft was launched on January 12, 2005. It consisted of two adjoined spacecraft which separated just before intercepting the target, comet 9P/Tempel 1, with one doomed to collide with the comet, and the other trailing so as to image the impact plume created by the collision. After the dramatic success of this mission on July 4, 2005 the Flyby spacecraft was decommissioned and placed in hibernation from August 2005 to September 2007 with only basic systems powered. On September 25, 2007, the Flyby was reactivated to support two Missions of Opportunity: **EPO**Ch (Extra-solar Planetary Observation & CHaracterization) and DI**XI** (Deep Impacted eXtended Investigation), which were combined to form the mission's name, EPOXI. DIXI, a flyby mission of the comet Hartley-2 in November 2010 is the primary goal of EPOXI and is similar in concept to the primary mission. The second investigation, EPOCh, started on January 21, 2008 and completed on August 31, 2008 during the cruise to Hartley-2. The purpose of EPOCh was to perform photometric investigations of stars with known transiting hot-Jupiter planets, utilizing the poor focus characteristics of the High Resolution Instrument (HRI) camera to its advantage to increase the point spread function of the measurements. However, since DIXI is the primary scientific objective, EPOCh operations must be constructed such as to pose minimal risk to the primary objective, DIXI.

The key challenges to readapting the flight system for this new operations scenario will be discussed. Also, the evolution of the EPOCh observation strategy as an example of a goal oriented planning strategy that was able to accommodate many significant bumps along the way through advanced planning and prioritized focus to keep the mission lean, yet on track to producing more successful science results.

2. EPOCH DESIGN APPROACH

Success-Oriented

As a Mission of Opportunity, EPOCh was originally proposed and planned with an optimistic, "Green Light", or "Success-Oriented" add-on to DIXI. Traditional mission operation strategies tend towards pessimism and it involves significant amounts of ground interaction to monitor each step. The operation strategy for EPOCh was fundamentally different. The optimistic strategy laid out several activities in sequence and relied heavily on automation and ground based testing. This allowed for the spacecraft team to design operations to utilize the maximum capacity of the spacecraft and a small flight team to deliver results well beyond the baseline science requirements. While the scenario planning was optimistic, the spacecraft flight team compiled a list of the risks to this strategy and established mitigation options to address these risks.

1

[1] 978-1-4244-2622-5/09/$25.00 ©2009 IEEE
[2] IEEEAC paper#1376, Version 1, Updated 2009:01:09

Flexible Operations

EPOCh was planned to be a low-cost, low impact, flexible, and simple to implement Mission of Opportunity. This flexibility was achieved by using a series of on-board relative-timed, sub-sequences or blocks, all spawned from a master, absolute-timed backbone sequence. Each block was engineered to perform a specific task: take images, telecommunication configuration, slew the spacecraft, downlink images, etc. Each of these was spawned by a unique backbone sequence that was built and reviewed on the ground to span 1-4 week intervals. The timing for each block was controlled by the backbone absolute-timed sequence. This strategy allowed a series of standard and unchanged on-board blocks to be used with infinite flexibility. These blocks were reusable throughout the entire EPOCh mission. Having tested and reviewed the reusable blocks prior to being stored on-board the spacecraft, only the unique backbone sequence needed to be reviewed prior to each imaging campaign. This strategy vastly simplified the building and reviewing of the backbone sequences. It provided enormous flexibility to the order and duration of the repeated activities on the spacecraft to optimize the layout of each backbone sequence.

The schedule for operations at the top level is based on the spacecraft's orbit and attitude requirements. These determined the geometric availability of targets. Candidate targets were chosen from ground-based observations. These observations were able to roughly characterize the orbital parameters of the target planet. These orbital parameters determined when the timing of the known desired observations. These included "transits" (when the planet passes in front of its parent star) and "eclipses" (when the planet passes behind its parent star) would occur. In the backbone sequence, these periods were called Keep Out Zones (KOZs).

The KOZs drove the schedule for antenna time on NASA's Deep Space Network (DSN). Nearly every target required the spacecraft to point in a direction that placed the Earth outside the range of motion of the high-gain antenna's (HGA) gimbal. This essentially eliminated high bandwith communication with Earth while imaging. Since the spacecraft must be pointed at the target and imaging during each KOZ, it could not point the HGA at Earth to downlink images. Time on the DSN was scheduled around KOZs. When the DSN schedule was finalized, the start and end of each DSN pass could be used to signal events. The backbone sequence called each block at specific times relative to the start and end of each pass. Each backbone sequence lasted from one to four weeks and was a function of the duration of the finalized DSN schedule. This strategy allowed for automation of an extremely flexible schedule with the DSN and maximized the time the spacecraft was imaging.

Different Modes of Operations

The Deep Impact Flyby spacecraft and ground systems were originally designed for a flyby mission. The astronomical observatory mission concept was never considered during the design of the spacecraft or mission. The architecture of the flight system was fundamentally intended for a flyby mission that collected large amounts of data in a short period of time during comet encounter. The new use of the spacecraft for EPOCh collected small amounts of data continuously, storing it for several days and then transmitting it to Earth. However, the flexibility of the spacecraft, along with the creativity and methodical systems re-engineering by the Ball Aerospace and JPL flight team were able to meet the challenges inherent to this new architecture.

3. LIMITING OPTIMIZATION, SIMPLIFIES OPERATIONS

Telecommunication Configuration

Almost all of the targets placed the Earth outside the HGA's range of motion. Since the operations plan was using reusable sequences, the telecommunication configuration sequences had to be designed for the worst-case scenario. Thus, while imaging, the spacecraft was commanded to use the low-gain antenna (LGA) which had a much larger beam-width. When a DSN pass was scheduled, the spacecraft was commanded to turn back to nominal cruise attitude, placing the Earth back inside the HGA's field of view.

A telecommunication anomaly in April 2008 placed a moratorium on moving the waveguide transfer switch. This switch allows the spacecraft to swap between the Low Gain Antenna (LGA) and the HGA. To accommodate this new requirement, the telecommunication sequences had to be redesigned to keep the spacecraft on the HGA while imaging instead of switching to the LGA. Since the Earth would be outside of the HGA's field of view while imaging, the HGA's feature to track the Earth had to be turned off while imaging. This eliminated the risk of the HGA hitting its stops but also eliminated any telemetry signal from being received at Earth while the spacecraft was imaging.

Variable Target Attitudes

Each EPOCh imaging campaign required the delivery of a unique backbone sequence and quaternions for each target to be observed during the campaign. The backbone always called the same block to slew the spacecraft to the imaging target. This block was hard-coded with a command to slew to the "EPOCh Science Target ID". The definition of this ID could be changed, allowing for the same commands to point to different targets. This technique proved to be powerful as updates to a target could be made in an attempt to compensate for poor spacecraft pointing. To account for different slew durations between targets that could be further or closer together, a standard value of 30 minutes was used, which is equivalent to 180-degree slew.

Data Completeness & Retransmission Requirement

With the small flight team and limited budget, there was a limited ability to support real time commanding. Most commands had to be automated, including the retransmission of data. The Deep Impact prime mission had significantly more frequent and longer DSN passes. The playback strategy for prime mission was to downlink data for the first 4 hours of the pass. This data was analyzed in real-time for completeness and errors. The remainder of the DSN pass was then spent building, reviewing, and sending commands to fill in the gaps in data. This strategy was inefficient. There was always a scramble to analyze the data, build the commands, and get them sent before the end of the antenna pass. For EPOCh, this strategy would be difficult. DSN passes were all times of the day and night and only a minimal flight team was available. This strategy was not an option for EPOCh due to schedule and cost constraints.

A leap-frog downlink technique was also considered. DSN passes would be scheduled every other day. Data from days 1-2 would be downlinked on the day 2 DSN pass. It would be removed or overwritten after the DSN pass on day 4, leaving days 3-4 to analyze the data and build commands to retransmit missing data on the day 4 DSN pass. This strategy entails a rigid DSN schedule with the potential for wasting DSN time that is not needed for retransmission, adversely effecting the imaging duty cycle. This strategy also avoided the off-shift, real-time scramble to analyze data and build commands that the prime mission strategy utilized.

The final downlink configuration chosen was a redundant downlink of every file. This allowed more flexibility in the DSN schedule, having shorter more frequent passes, or longer less frequent passes; however, the total DSN time required was slightly larger than the other strategies. It also avoided any real-time commanding as all downlinks were sequenced. The only limitation was to ensure that the DSN station was configured correctly and ready to receive data by the time the files were sent from the spacecraft.

Image Size & Pointing

To minimize image storage size and maximize the time spent imaging, it was chosen that only the center 128x128 pixels of the total possible 1024x1024 pixels would be stored. This image size had a total field of view of 53.1 arcseconds. The small field of view drove the necessity for a stable imaging platform that was not considered in the design of the spacecraft for the original Deep Impact mission. Calibration data from the prime mission was used to determine if the 128x128 pixel image size would be feasible. However, the Impactor provided a small, but non-negligible amount of structural support. It was unknown what affect the release of the Impactor had on the instrument calibration.

A full recalibration of the HRI to star trackers could have been performed, but cost and schedule constraints, along with the uncertainty of whether the calibration was needed. The uncertainty drove the optimistic decision to forgo recalibration of the instrument until flight data could demonstrate there, in fact, was a pointing problem and that a more simple mitigation technique would not be effective.

Prior to EPOCh operations, several solutions to the pointing stability issue were identified. However, no flight or test performance data existed to show which solution, or a combination of several, if any, would be the best solution. It was decided to not implement any solution until several data sets had been gathered.

4. POINTING STABILITY

EPOCh observations began on January 26, 2008. The first data set yielded only 29.8% of images usable for photometry. The ensuing two weeks of imaging improved this statistic to 83% of images usable for photometric data, still below the baseline requirement. These poor results were due to spacecraft pointing errors causing the target star to wander outside the 128x128 pixel field of view of the instrument. Figure 1 displays a scatter plot of the star

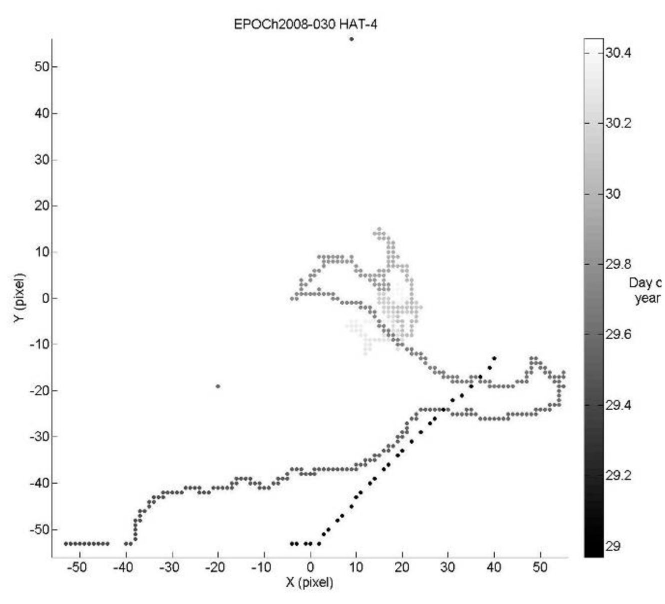

Figure 1 - Target drift in the field of view

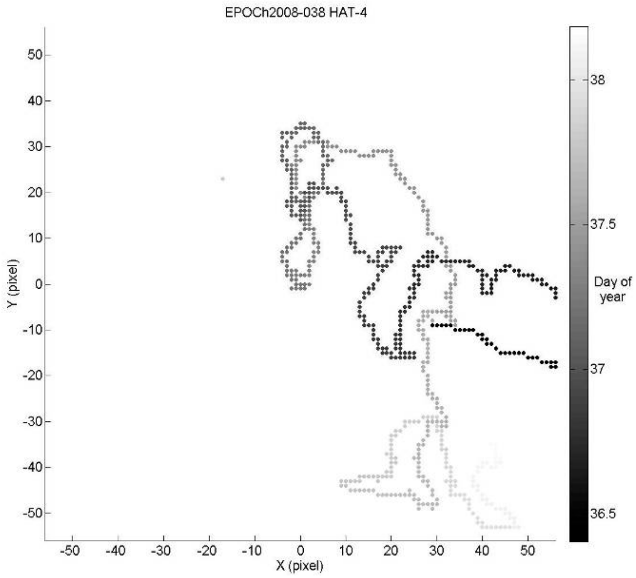

Figure 2- Target position with pointing offset

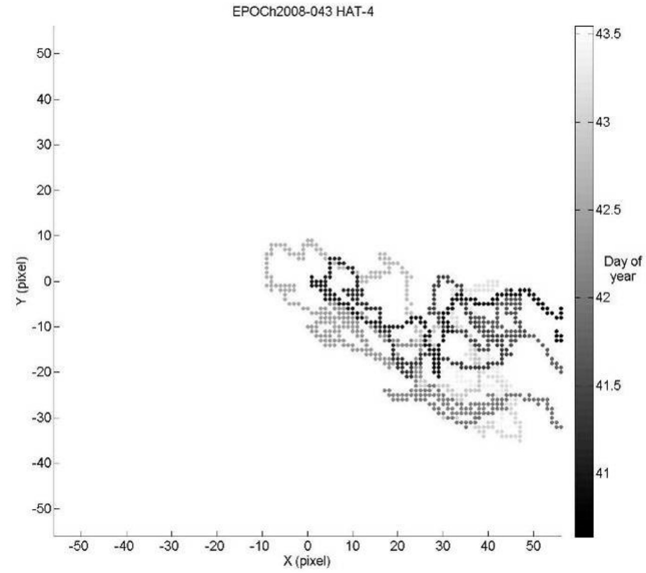

Figure 3- Target drift with offset and gyro bias update

over the course of two days of imaging. There were several hours where the target was outside the field of view in the minus-Y direction and no data could be gathered.

The flight team had to improve the percentage of usable images in order to satisfy the baseline science requirements. The mitigation strategies that had been identified early in the planning phase were reexamined based on the flight experience and data collected. The results of the mitigation attempts are described below in this section.

Pointing Offset

The first two weeks of data yielded only 83% usable photometric data out of the baseline required 90%. Solutions to this problem had to be implemented. The initial data showed that the target was clustered in one corner of the image. The simplest solution was to skew the pointing by the average position of the target. This brought the star closer to the center of the field of view and allowed the continuation of data collection with only a minor interruption and trivial increase in command complexity. A scatter plot of the first data set with this improvement is shown in Figure 2

Attitude Determination Technique

The gyro bias estimation time constant defines how the attitude determined from the star trackers is weighted versus attitude determination determined from the gyros. The second corrective action to improve the pointing was to update this parameter. Having the flight data to analyze made this a simple activity. The pointing variation was effectively damped providing much longer durations of good pointing. However, the target still tended to drift in and out of the field of view and tended to stayed outside the field of view for longer durations. A scatter plot of data with both the pointing offset and the gyro bias estimation time constant update is shown in Figure 3

Pre-visit

The second target observed proved to have a different pointing offset than the first target. Since the pointing bias was a function of the inertial attitude, a new bias had to be applied for each inertial attitude. A pointing offset was made for each target, effectively mitigating the pointing problem.

Under the assumption that each inertial pointing would have a different pointing bias, it was decided to perform a "pre-visit" of each target several days prior to it being vigorously observed. The flight team could now calculate the pointing bias and generate the necessary commands to implement it for the nominal observations.

The pre-visit strategy drastically improved the pointing performance, but since the target continued to drift over several hours in the small field of view, the target still tended to drift out of view. A new solution had to be developed.

Hybrid Image Size

The spacecraft experienced a serious telecom anomaly in April 2008. The telecom system was only emitting about 33% of the expected power. At this low power, the DSN stations could not lock onto the 200,000 bits per second sub-carrier signal. The telemetry rate had to be reduced to 100,000 bits per second in order for DSN stations to be able to lock on to the signal. All EPOCh observations were put on hold until this telecommunication problem could be solved. The flight team took advantage of this time to make significant improvements to the EPOCh strategy.

Improvements were made to both the pointing stability and the downlink strategy. The downlink strategy will be discussed in detail in the next section. The two months of experience running EPOCh allowed the flight team to assess the performance of the current technique and further optimize the observational strategy.

It was believed that the pointing stability of the spacecraft was as good as it was going to get, but it still was not enough to meet the baseline requirements for EPOCh. A different approach had to be taken.

The 128x128 pixel image size was chosen to minimize image storage and downlink time. However, the limiting factor on the number of images able to be stored was not the total storage size, but rather the number of files saved to memory of the spacecraft. This is because of the flash file system used. In addition, the flexibility of the downlink for EPOCh yielded more DSN resources than necessary.

Since the KOZs were a higher priority and of a known length and known time, it was decided to use a 256x256 image to cover a single KOZ between DSN passes and fill the rest of the time with 128x128 images. This increased the required downlink time and decreased the maximum imaging duration. However, the average imaging duration (time between DSN passes) experienced thus far was less than the maximum imaging duration capable.

This hybrid imaging technique preserved the higher priority data during planetary transits and eclipses and

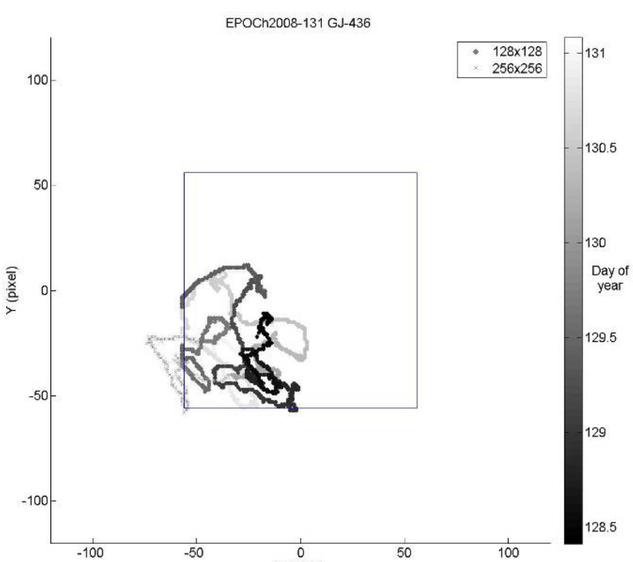

Figure 4 - Target drift with hybrid imaging technique

provided adequate photometry during the remaining imaging time. Figure 4, displays a scatter plot of the target drift over several days of observations. The plot contains data with the pointing offset, pre-visit, gyro bias time estimation constant update, and hybrid image size.

The large box is the 128x128 pixel field of view. The data points with the cross shape (bottom left) are images taken in 256x256 mode during a KOZ. The hybrid imaging technique saved these images that would have otherwise been lost with the smaller image size.

Summary of Pointing Stability Issues

In summary, the green-light strategy resulted in many small fixes to a large problem. One must keep in mind that the final strategy to the pointing problem contained solutions to a few observed problems rather than solutions to many potential problems. This green light strategy yielded a more efficient use of flight team resources and an earlier start to collecting science data.

5. Retransmission & Data Completeness

The possible downlink strategies considered during initial development were previously described in Section 3. To recap, the techniques considered were:

1. Single downlink with real-time data analysis and custom re-transmission; used in prime mission; inefficient and labor intensive.

2. Leapfrog strategy; store images for two DSN passes, single downlink on the first DSN pass and retransmit the missing data on the second DSN pass; poor duty cycle; some DSN time may be wasted; rigid DSN schedule; labor intensive, but less than real-time retransmission.

3. Redundant downlink of all data; fully automated; no real-time commanding required; poor duty cycle; minimal required labor; extremely flexible DSN schedule; requires more DSN time.

Original 128x128 Downlink Strategy

As mentioned before, the final downlink configuration selected was a redundant downlink of every file. The total DSN time required was slightly larger than the other strategies, but it avoided any real-time commanding as all downlinks were sequenced.

The largest problem with this technique was that DSN passes could be too short to receive all the data transmitted from the spacecraft. This resulted in the first transmission of the 128x128 images to not be received on the ground.

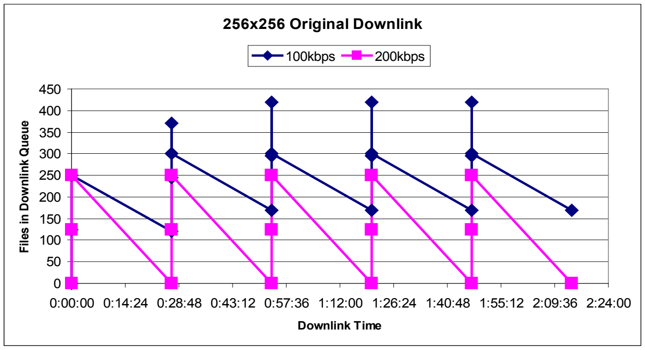

Figure 5 - Unmodified Series Downlink Strategy

Table 1 - Unmodified Series Downlink Strategy

Added to Queue	1-125	251-375	501-625	126-250	376-500
	126-250	376-500	1-125	251-375	501-625
100kbps received	1-130	131-260	261-390	391-445	591-625
				501-590	1-5
					126-215
100kbps lost				445-500	6-125
					216-625
200kbps received	1-250	251-500	501-625	126-375	376-625
			1-125		

Fortunately, the redundant downlink later in the pass re-transmitted the lost images.

Initial 256x256 Downlink Strategy

With the new hybrid imaging strategy and the telecom anomaly in April 2008, a new downlink strategy had to be implemented. The telecom anomaly resulted in about two-thirds of the radiated RF power from the spacecraft being reflected back into the spacecraft bus and only one-third of the power being radiated towards earth. At this reduced power, the DSN could not lock onto the 200kbps signal that was being transmitted. The highest data rate capable of being resolved at the DSN was 100kbps. The initial downlink strategy would not work at this low data rate. Only half of the images commanded to be downlinked made it to the ground. This was because the file downlink queue on the spacecraft was hard-coded to 300 images. Any file added to a full downlink queue would overflow the queue and be lost. The redundant downlink technique utilized did not protect against a lower data rate. At the time the sequences were being designed, it was assumed that the downlink bit rate would remain constant.

The new downlink strategy had to be capable of transmitting all images twice at 200kbps or all images once at 100kbps. The challenges to these requirements are listed below:

- Variable number of 128x128 images
- 256x256 images were of a higher priority
- Maximum file downlink queue of 300 files
- Variable DSN pass length

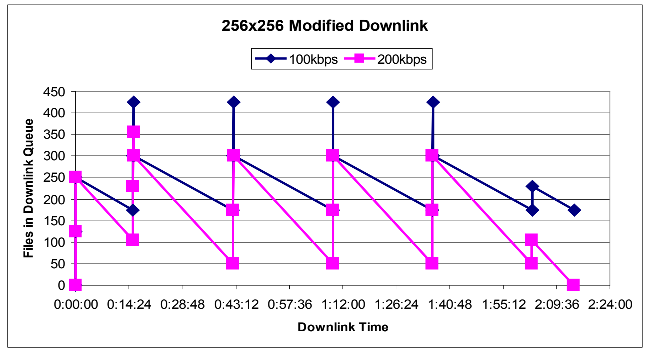

Figure 6 - Modified Series Downlink Strategy

Table 2 - Modified Series Downlink Strategy

Added to Queue	1-125	501-625	376-500	126-250	251-375	196-250
	251-375	126-250	1-125	501-625	376-500	
100kbps received	1-75	76-125	326-375	576-625	451-500	201-250
		251-325	501-575	376-450	126-200	251-255
100kbps lost		126-250	1-125	501-625	376-500	256-375
200kbps received	1-125	271-375	146-195	75-125	576-625	451-500
	251-270	501-625	376-500	126-250	251-375	320-375
		126-145	1-75	501-575	376-450	

Two downlink strategies were investigated: series redundant downlink and parallel redundant downlink. Parallel downlink would place blocks of 125 images in the downlink queue twice, back-to-back. This method was susceptible to data loss if there was a problem with the downlink for greater than 0:03:36, the transmission time of a set of 125 images. However, this strategy is useful for transmitting a variable number of images. The sequence can be halted before completion and it does not terminate the redundancy.

The series redundant downlink would place all blocks of 125 images in the downlink queue in order, twice, one after the other. This strategy was not feasible for a variable number of images since terminating the sequence early would eliminate a portion of the redundancy. In addition, much time would be wasted attempting to transmit images that were not there. However, this method was quite useful for a known, fixed, number of images. This worked quite well at 200kbps. However, it was soon discovered in testing that at 100kbps, it was not transmitting all images once at 100kbps, but rather transmitting the same incomplete set of images twice. See Figure 5 and Table 1 for details. The images are transmitted in five blocks of 125, labeled A, B, C, D, and E. Notice how much of block D is lost.

Final 256x256 downlink strategy

The final solution utilized a parallel downlink technique for the 128x128 images accepting the risk of losing data if the DSN dropped lock for over 0:03:36 along with a modified

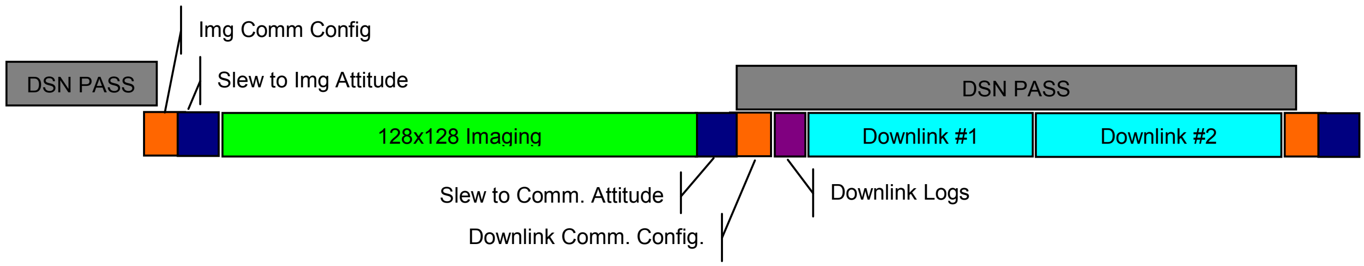

Figure 7 - Initial strategy with only 128x128 images

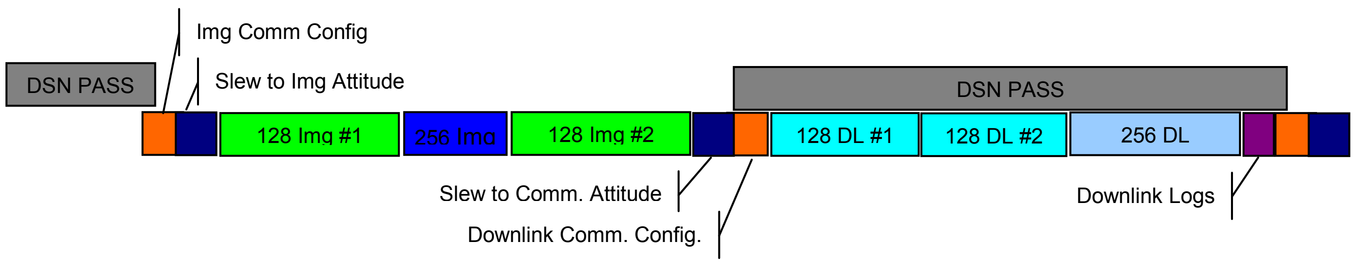

Figure 8 - Final strategy with 128x128 & 256x256 images

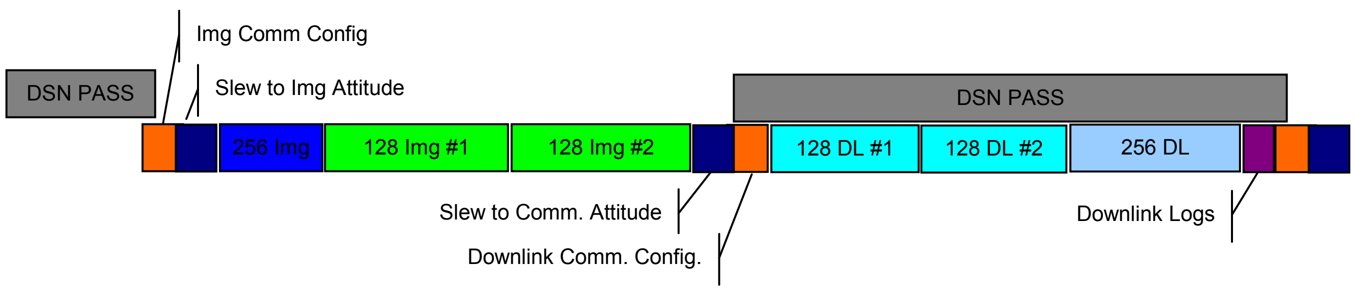

Figure 9 - Final strategy with 256x256 images at the start

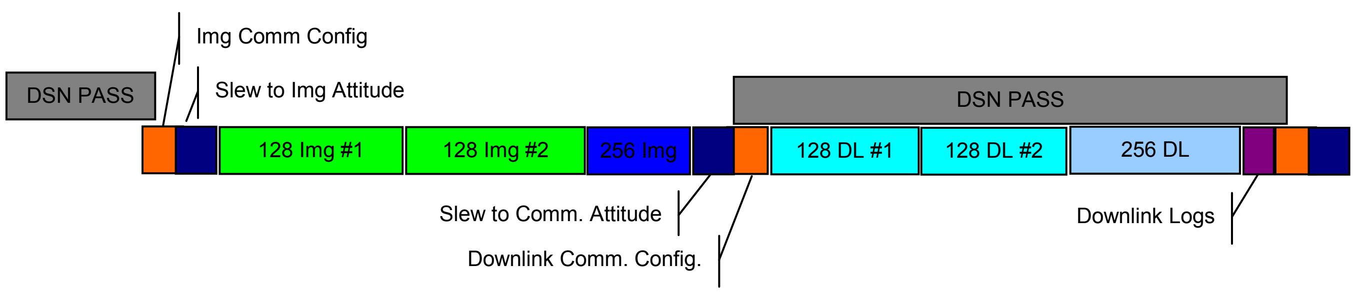

Figure 10 - Final strategy with 256x256 images at the end

series downlink for the 256x256 images. There were two sequences to downlink the 128x128 images. Each sequence transmitted the images taken in the corresponding 128x128 imaging sequence. The first command in these sequences was to halt the previous downlink sequence, effectively creating infinitely variable transmission times for either of the 128x128 downlink sequences. The modified series downlink, used for the 256x256 images placed in blocks of 125 images into the downlink queue in a rather strange order that ensured a single downlink at 100kbps and a redundant downlink at 200kbps. This strategy specifically controlled which images were lost due to the downlink queue overflow. See Figure 6 and Table 2 for details.

6. Block Layout

EPOCh operations were required to be quick, easy, and cost-effective. This requirement leads to the strategy of a backbone sequence and reusable blocks.

Initial Configuration

The initial method used consisted of eight blocks with the order of operations beginning notionally at the end of a DSN pass. They are:

1. Configure communications for imaging
2. Slew to imaging attitude
3. Image
4. Slew to communication attitude
5. Calibration images (during the slew)
6. Configure communications for downlink
7. Downlink #1
8. Downlink #2

This strategy was simple to implement. Ground software would ingest the DSN schedule and autonomously schedule the backbone sequence. Since each block was rigorously tested, each new backbone sequence did not require a test, but rather a visual inspection of the timing. See Figure 7 for a visual depiction of the block layout.

Final Configuration

The modifications to both the layout and the contents of the blocks were necessary due to the problems with pointing stability and the telecom anomaly experienced in April 2008.

The new imaging strategy has been previously described. However, the layout of the imaging sequences has not been discussed. Three sequences of imaging were used: with image sizes of 128x128 and one with image sizes of 256x256. This method allowed for a large amount of flexibility. A single KOZ could be covered with the five blocks of 125 images at a size of 256x256 and the rest of the imaging duration could be covered with 128x128 images. See Figure 8 for a visual aid.

Each imaging sequence had the capability of halting the previous imaging sequence, effectively allowing for infinite adjustment of each imaging sequence's duration. Depending on the location of the KOZ between DSN passes, the 256x256 imaging block could be shifted sooner or later to ensure it was fully imaged. If the KOZ was located at an extreme beginning or extreme end of the imaging duration, the 256x256 imaging block could be shifted before (Figure 9) or after (Figure 10) both 128x128 imaging blocks yielding full coverage of the imaging duration and coverage of the high priority data with the larger images.

The final block timeline consisted of ten different elements shown below and shown visually in Figure with variations in Figure 9 and Figure 10.

1. Configure communications for imaging
2. Slew to imaging attitude
3. Image with 128x128 images
4. Image with 256x256 images
5. Image with 128x128 images
6. Slew to communication attitude
7. Configure communications for downlink
8. Downlink 128x128 images #1
9. Downlink 128x128 images #2
10. Downlink 256x256 images

This strategy was more complex than the original plan that yielded time and cost penalties. Since the new strategy was implemented with only about two months of EPOCh observations left, the ground tools were not updated with logic to ingest the KOZ schedule. As a result, the 256x256 imaging block could not be autonomously placed. The times for scheduling the downlink and imaging blocks in the backbone sequence had to be manipulated by hand for optimal placement. This added some risk to loss of science data, but was mitigated by multiple visual inspections of the sequence

7. RESULTS

The end result of EPOCh was a total of 198,434 images taken of eight star systems. Of the images taken, a total of 192,624 images were received on the ground. The missing images were due to various issues with the DSN schedule. The principal issue resulting in image losses were DSN passes scheduled using the initial 128x128 algorithm, but occurred after the hybrid 128x128 and 256x256 imaging technique was first implemented. This resulted in passes that were too short causing images being transmitted to Earth, but with the DSN station not yet configured to receive them. Of the 192,624 images received, 172,209 images were considered photometric and usable for science. This yielded an image acceptance rate of 89.6%.

This acceptance rate was 0.4% below the baseline requirement. But, due to the safing event in February and the telecommunication anomaly in April, two months of replacement observations were implemented. When the original planned number of images of 177,920 is compared to the number of photometric images of 172,209, the image acceptance rate now jumps to 96.8%, well above the 90% baseline requirement for EPOCh.

8. CONCLUSIONS

The EPOCh observation campaign was designed to be low-cost both in terms of spacecraft resources and budget. The operational strategy assumed everything would go according to plan, but critical risks were identified and mitigation strategies were discussed prior to the beginning of observations. As a result, only real problems encountered were fixed and anticipated problems that did not occur did not have to be solved. This saved on both cost and time at a slight risk to science data.

Many lessons can be learned from this operational strategy. Early planning and key risk identification are key to developing an observational technique that is both flexible to schedule and can adapt to the real problems, and not anticipated problems.

For more information, visit the official EPOXI website at http://epoxi.umd.edu/.

The University of Maryland is the Principal Investigating institution, leading the overall EPOXI mission, including the flyby of comet Hartley 2. NASA Goddard leads the extrasolar planet observations. NASA's Jet Propulsion Laboratory, Pasadena, Calif., manages EPOXI for NASA's Science Mission Directorate, Washington. The spacecraft was built for NASA by Ball Aerospace & Technologies Corp., Boulder, Colo.

9. ACKNOWLEDGEMENTS

The authors would like to recognize the following people: The entire JPL and Ball Aerospace spacecraft team for their incredible skills at operating the spacecraft and their critical involvement in the EPOCh observations; Dr. Michael A'Hern, the Principal Investigator for the EPOXI mission; Drake Deming, The EPOCh Principal Investigator and EPOXI Deputy Principal Investigator; Dr. Dennis Wellnitz, a CO-Investigator for EPOXI and the Deputy Principal Investigator for EPOCh; Greg LaBorde and Hal Uffelman, fellow Flight Team Activity Leads who have been amazing resources for providing sound advice and ideas for EPOCh; and Steve Wissler, the EPOXI Spacecraft Team Chief who's incredible knowledge in both the spacecraft and the ground data systems have made the entire EPOXI mission possible.

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

10. BIOGRAPHIES

Robert Sharrow was the developmental Activity Lead for EPOCh and devised the initial strategy. He was the Deputy Proposal Manager for the extended Mission and established many of the green light strategies to keep the concept simple and within a tight cost cap.

He graduated from the University of Michigan in 2002 with Bachelors degrees in both Mechanical and Aerospace engineering and in 2004 with a Master's degree in Space Systems Engineering. While at the Jet Propulsion Laboratory Rob has worked as a Mission Planner on the Mars Reconnaissance Orbiter, Deputy Proposal Manager and Activity lead on EPOXI, Mission Planner for the Aquarius/SAC-D, and is currently working as a Payload Systems Engineer on the JUNO mission to Jupiter.

Richard Rieber began his role on EPOXI as the test bed operator and after having worked closely with the flight team on the testing, began to take a bigger role in the building of the commands. He eventually took over the role as Activity Lead for EPOCh and led the development of the final hybrid imaging and downlink strategies. He graduated from the University of Colorado – Boulder in 2007 with a Masters degree in Aerospace Engineering. He joined the Jet Propulsion Laboratory in 2007. In addition to his work as an Activity Lead and Flight Director, he continues to run the testbed that are used to model all commands, sequences, and software before they are sent to the spacecraft.