Automation and Process Improvement Enables a Small Team To Operate a Low Thrust Mission In Orbit Around the Asteroid Vesta

Timothy M. Weise[1]
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, 91109, United States of America

NASA's Dawn mission to the asteroid Vesta and dwarf planet Ceres launched September 27, 2007 and arrived at Vesta in July of 2011. This mission uses ion propulsion to achieve the necessary delta-V to reach and maneuver at Vesta and Ceres. This paper will show how the evolution of ground system automation and process improvement allowed a relatively small engineering team to transition from cruise operations to asteroid operations while maintaining robust processes. The cruise to Vesta phase lasted almost 4 years and consisted of activities that were built with software tools, but each tool was open loop and required engineers to review the output to ensure consistency. Additionally, this same time period was characterized by the evolution from manually retrieved and reviewed data products to automatically generated data products and data value checking. Furthermore, the team originally took about three to four weeks to design and build about four weeks of spacecraft activities, with spacecraft contacts only once a week. Operations around the asteroid Vesta increased the tempo dramatically by transitioning from one contact a week to three or four contacts a week, to fourteen contacts a week (every 12 hours). This was accompanied by a similar increase in activity complexity as well as very fast turn around activity design and build cycles. The design process became more automated and the tools became closed loop, allowing the team to build more activities without sacrificing rigor. Additionally, these activities were dependent on the results of flight system performance, so more automation was added to analyze the flight data and provide results in a timely fashion to feed the design cycle. All of this automation and process improvement enabled up the engineers to focus on other aspects of spacecraft operations, including spacecraft health monitoring and anomaly resolution.

I. Introduction

The Dawn spacecraft launched September 27, 2007 on its journey to study the protoplanets Vesta and Ceres. It is the robotic eyes and ears of the Dawn Project, the National Aeronautics and Space Administration's (NASA's) ninth Discovery class project.[1] Dawn will study the composition of these two complementary bodies to help scientists understand the make up of the solar system at its earliest stages. In order to study these bodies in detail, the spacecraft will need to operate at multiple orbital altitudes, starting with a high orbit to acquire visible and near infrared imaging with the Framing Camera (FC) and visible and infrared spectral data with the Visible and Infrared Mapping Spectrometer (VIR) of most of the surface and begin mapping the gravity field, then going to a lower orbit to improve the imaging resolution, and yet lower to achieve the necessary resolution for the Gamma Ray and Neutron Detector (GRaND) instrument.[2] The spacecraft was built by Orbital Sciences Corporation and is operated by NASA's Jet Propulsion Laboratory (JPL).

---

[1] Mission Operations Engineer IV, Flight Engineering Group, 4800 Oak Grove Drive, M/S 264-828, AIAA Senior Member.

One of the most significant enabling technologies of the Dawn spacecraft is the use of solar electric propulsion (SEP) to achieve the required delta-V to transfer from Earth to Vesta, and later from Vesta to Ceres, as well as change orbits at both protoplanets. SEP was first demonstrated on NASA's Deep Space 1 mission, and much of Dawn's SEP design was based on that mission. SEP uses electricity produced from solar arrays to ionize xenon (Xe) gas and then accelerate the ions to produce thrust. This form of propulsion is highly efficient, nearly ten times the efficiency of chemical propulsion. This system will utilize 425 kg of Xe to achieve approximately 11 km/s of orbital velocity change during the mission.[3] While SEP is very efficient, it also has a very low thrust – for Dawn the maximum thrust is around 91 mN, and as the solar distance increases, there is less power available for the SEP system resulting in even lower thrust. Therefore, to achieve the required delta-V, the SEP system must supply thrust for many hours. In fact, during the cruise phase from Earth to Vesta, the mission design utilized an average thrust duty cycle of about 95%. This translates to interrupting thrust for about eight hours each week to communicate with the operations team on Earth (Fig. 1). After almost four years of flight, the spacecraft arrived at Vesta with an accumulated thruster operation time of about 23,000 hours.

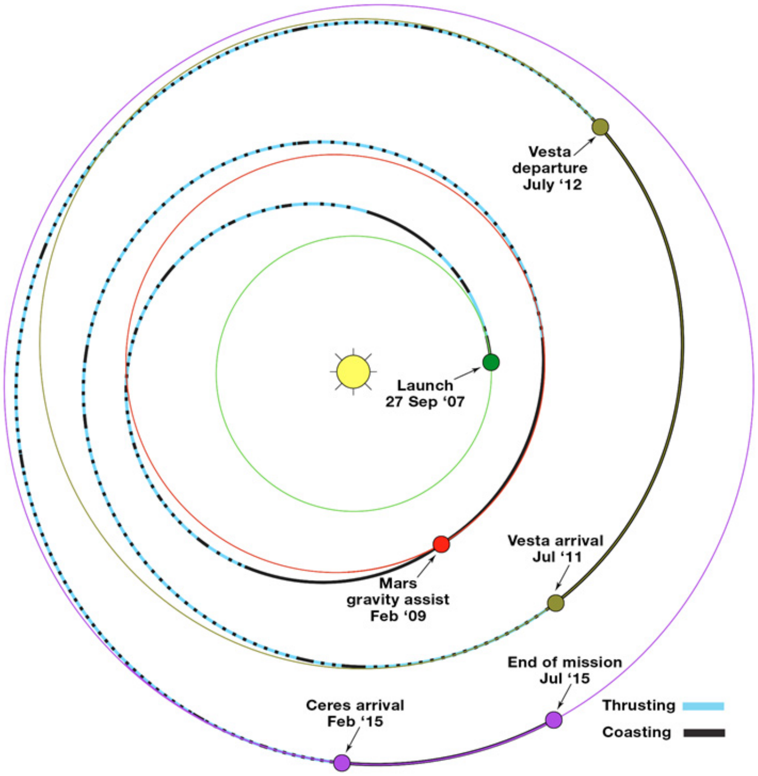

Figure 1. Dawn Mission Trajectory

The thrust duty cycle aspect of the mission design drove a significant portion of the Mission Operations System (MOS) processes, procedures, team interfaces, and software tools. Since the spacecraft attitude for thrusting generally was not consistent with the attitude for high-rate communications, it was commanded, via an on-board sequence of commands, to stop thrusting, turn the High Gain Antenna (HGA) to Earth, and transmit the engineering telemetry recorded since the last communications session. At the end of the session, the sequence would command the spacecraft back to the thrusting attitude and resume thrusting. During the communication session, the operations team would perform a real time health and status check and uplink any required commands. Later, between sessions, the team would review the recorded telemetry for health and performance assessments. These sequences of commands that implement the flight team's design of activities were built, reviewed, and tested (if necessary) to fit this pattern of thrusting and communicating. With infrequent contacts and almost continuous thrusting, the next sequence of activities was being designed before the previous was complete. That means that performance data could not be folded into the next immediate sequence, but one later.

The last significant aspect of the MOS is that the Dawn mission is a Discovery Class mission, which bounds the project costs to a certain amount. The eight-year duration of the mission resulted in a need to spread out the operations costs over the years with the end result being a relatively small flight team. Many of the subsystem engineers supporting flight operations were staffed at a part time level, from one-quarter time to half time for many subsystems. These engineers support several projects to fill out their work load. This presents a couple of challenges. First, the engineers don't have many hours to spend on Dawn each week, so the MOS and Ground Data System (GDS) needed to enable them to spend most of their time in their area of expertise. Additionally, they have multiple project commitments, which required the work schedule on Dawn to have some flexibility to work around their other commitments.

This paper will first describe the initial cruise operations processes and schedule. It will also discuss the significant operations aspects that those processes captured. Next, the early lessons will be discussed and how they led to improvements in the processes and tools to lessen the operations team's workload. With the cruise operations in hand, operations at the asteroid Vesta, and the associated challenges in both the operational pace and the ramifications of orbiting a previously unvisited, uncharted body, will be discussed. The solutions to these challenges will be presented, along with the suitability of those solutions to the remainder of the mission to Ceres. Finally, the benefits of the process improvement and tool automation will be presented.

II. Cruise Operations

Throughout the mission, the operations team has two primary tasks, and one contingency task. The first primary task is the uplink process. Many sub-teams within the MOS contribute to the uplink process to build, test, and review sequences of commands to be uplinked to the spacecraft. These sequences are designed to achieve particular tasks such as communicating with stations on Earth, performing engineering maintenance on the spacecraft, or firing one of the ion thrusters to propel the spacecraft to Vesta. The second primary task is the downlink process. Again, there are several sub-teams that utilize the data from the spacecraft, but the Spacecraft Engineering Team (SCT) is primarily responsible for assessing the health, safety, and performance of the spacecraft. The downlink process involves getting the data from the spacecraft to the engineers in a format that is useful to them, analyzing the data, and reporting the results. The contingency task is anomaly resolution. This can range from something as minor as losing telemetry data because of bad weather to something as serious as an unplanned computer reset. By making the uplink and downlink processes as efficient as possible, the engineers have more time to spend analyzing the contingencies and recovering the spacecraft back to normal operations.

A. Uplink Process

As was discussed earlier, early cruise operations were defined by near continuous thrusting, with one HGA contact a week. There were several steps required to build a typical sequence, which was four weeks in duration (after several months of cruise operations, the sequences grew to five weeks). Fig. 2 shows the high level steps in the process. First, the Deep Space Network (DSN) schedule must be completed by the Multi-Mission Resource Scheduling Services (MRSS) team, which dictates when the HGA contacts occur. Step 2 shows the SCT providing turn duration and power predicts based on the extended design of the previous sequence. The Mission Planning and Sequencing Team (MPST) used the DSN schedule and a rough prediction of turn durations to determine when to start and stop thrusting in step 3, which were then delivered to the navigation team (Nav). The navigation team designed the thrust activity using these start and stop times and power predictions that indicate how much power was available for the Ion Propulsion System (IPS) thruster. Nav produced a thrust vector file (TVF) that indicated the thrust start/stop times, thrust direction, and throttle level in step 4. An interesting aspect of low thrust missions is that the immediate thrust design (or loss of thrust because of an anomaly) affects thrusting in the future, so when Nav designed the next thrust sequence, they made a detailed design for the current sequence plus two more, as well as a rough design of all the thrusting to Vesta. Then, in step 5, MPST uses the thrust start/stop times from step 3 and the throttle level from step 4 to build the thrust sequence. While that was in progress the SCT processed the TVF into a file that controlled the thrust pointing to be uplinked to the spacecraft in step 6. It is not until step 7 where all of the teams reviewed the products to ensure consistency in the design. If the sequence passes review, it was uplinked in step 8, otherwise it returned to a previous step, which depended on the type of error. This process was spread over several weeks to allow adequate time for each step as well as fit around the schedules of team members with limited hours on the project. As shown in Fig. 2, the Nav design of thrusting and the thrust sequence start with the same inputs, but do not have an automatic process to ensure consistency. This required extra time in the process to review the products, and be able to rebuild the products if an error was found.

B. Downlink Process

On the downlink side, the weekly contacts provided a consistent operations pace. Even so, from the beginning of the mission there was a drive to automate as much of the process as possible to allow the subsystem experts to focus their energy on their area of expertise, not spending time getting the data from the server. There are two main aspects of downlink analysis – real time health and status, and long-term performance. Since the spacecraft only communicated with Earth once a week, during each contact there was real time telemetry with the current status of

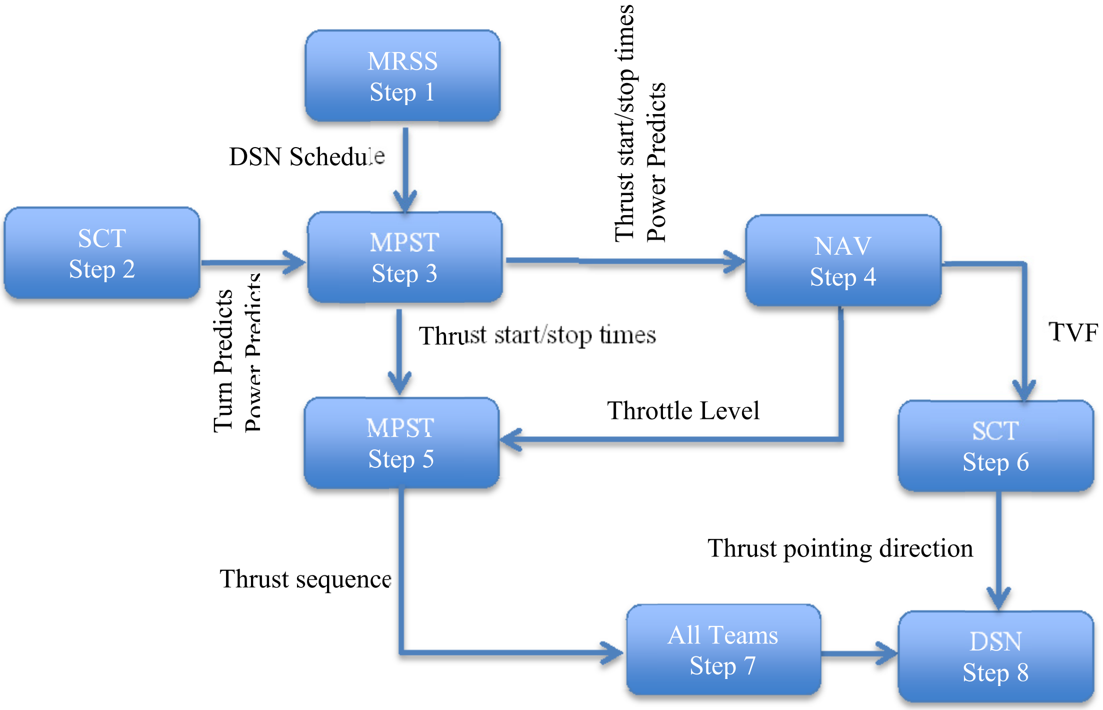

Figure 2. Simplified Uplink Process

the spacecraft and playback data from the previous week. The playback data include all of the data recorded since the previous contact. The immediate health and status was monitored with a process called Automatic Alarm Notification (AAN), which sends a page and an email to the systems engineers on call if any telemetry channel exceeds its red alarm limits. Red alarm limits were set such that they would provide the operations team an alert when a channel is out of range, but before it is critical. For example, temperature alarms are set 5 degrees C inside of the on board fault protection (FP) limits. This provides enough room for the system to operate without triggering false alarms, yet still provides an alert before fault protection takes action. Additionally, engineers would analyze the real time data to spot-check the performance.

To address the need to review the playback data, the GDS team developed tools to query telemetry and plot data where the user provides the time range of interest, a file that defines the list of channels to plot, and an email address to email a PDF (Portable Document Format) of the plotted data. This specific tool was written in the Perl scripting language and handled the interfaces with the telemetry servers and the plotting routine. While this was a good start, it still required each team member to set up their own individual runs of the tool – that means looking up the start and end times of the DSN contact period, waiting until after the contact completes and starting the run. The first improvement on this tool was a wrapper script, also in Perl, that read in a list of plots and who to email them to. This tool was then set up to run after the weekly contacts and every member of the team automatically received plots of their data in their email as they became available. At this point it was still a manual process to setup the schedule to run the process for every contact in a sequence. Finally, these plots were reviewed manually to identify any issues in the data. This can be time consuming and requires diligence to not miss something.

As the engineers developed an improved understanding of how to look at their data and observed how the spacecraft operates, two new tools were developed. The first addressed the need to have running statistics of the plotted data from the beginning of the mission. The SEP nature of the mission resulted in two major operating modes – ion thrusting or not, plus the transition from not thrusting to thrusting. The major differences between these modes are the attitude control mode, the power used by the spacecraft, and the thermal environments. When not thrusting, the spacecraft is typically quiescent with the HGA pointed at Earth and only consuming about 800W of power. Solar pressure provided the only significant external disturbance torque on the spacecraft. When transitioning to thrusting, the spacecraft turns to point the ion thruster in the appropriate direction and the ion thruster is pre-heated which begins flowing xenon from the tank through the feed system to the thruster. After thrusting begins, the power supplies that can provide up to 2,300W to the ion thruster begin to warm, the ion thruster warms, and power system provided up to 3,100W of power for the total spacecraft load. The ion thruster is gimbaled to apply the force through the null torque location, but the ion beam also creates a swirl torque on the spacecraft as it exits the thruster.[3] The team found itself comparing many statistics from week to week trying to answer questions like: What was the hottest temperature seen? What was the average power while thrusting? How do the heater duty cycles compare while thrusting and not thrusting? This resulted in a software tool that provided trend data along with the plots that were already being generated. The trend data consisted of the minimum and maximum values for the week and the mission, as well as the time averaged value and the previous week's time averaged value. All of these parameters were also binned into not thrusting, start up, and thrusting. This provided the engineers a quick reference to see if a temperature spike seen during startup was trending up or not, how the average power consumption while thrusting compared to the prediction, how the heater duty cycles compare between the three phases, or any number of other details.

Another new software tool addressed the need to automatically review the data for expected values – this is similar to the alarm system, but applies limits that help assess the state as opposed to indicating an urgent problem. Since the engineers found themselves reviewing data that behaved in predictable ways, this tool was developed to apply a set of rules to the telemetry data and report the results. This tool accepted the same data file that was used to plot and trend the data, as well as a rules file. The rules file contains definitions of rules that have a name, severity level, and a constraint. The constraints can be comparisons like greater than, less than, equal to, not equal to, change, change by n. For example, a rule can be defined to determine if a channel does not equal a specific value, or changes by more than 5. The output is written to a file and can be configured to only output data that violates rules. Both of these tools were then integrated into the batch process that queries and plots data so now the engineers receive an email containing attachments of the plotted data, trend data, and rule checks, as well as the rule summary in the body of the email. This allows them to quickly scan their email to see if there are any unexpected data values that need investigation instead of opening each plot file and manually reviewing each trace to check simple parameters. This increase in efficiency allowed the engineers to focus on other tasks.

III. Vesta Operations

A. Uplink Process

1. Thrust Sequence Development

Leading up to operations at Vesta, it became apparent that several improvements would be required for both the uplink and downlink process. To begin with, the thrust design cycle time needed to be shortened to achieve more accurate maneuver delivery. While in cruise, each individual thrust sequence design addressed the very well known gravitational effects of the solar system and did not require great accuracy because the trajectory around the sun was changing slowly and errors could be made up in the next design. In fact, with the long sequence design cycle (approximately four weeks) and the long sequence duration (four to five weeks), the thrust designs were being completed before most of the previous sequence had even executed which meant that the current design could not account for maneuver execution error. However, in order to perform science data acquisitions, very specific orbits needed to be achieved. In addition to maneuver execution error, Vesta's gravity field had not been measured prior to arrival, and needed to be determined while maneuvering around the body. As the spacecraft approached the body, the gravity field of Vesta started having a larger influence on the spacecraft's orbit, but the navigation team only had a very preliminary model of the gravity field based on perturbations on other asteroids and Mars and relatively low resolution images from the Hubble Space Telescope. While determining the spacecraft's orbit, errors could not be isolated from uncertainty in the gravity field or execution error. By shortening thrust sequence design timeline, the time from the orbit determination to the maneuver execution would also be shorter, and the error in the actual orbit could not grow large enough to become problematic to achieve the required science orbits. A by-product of this change was more frequent DSN contacts during the thrusting periods.

In order to compress the thrust sequence build timeline, several process improvements were made. First, the sequence designs were split into thrust windows and non-thrust windows. Then, the interfaces between these windows were rigorously defined. For example, an HGA contact period would start with commands to turn to point the HGA at Earth, followed by data playback commands. These contacts were in the non-thrust windows and collectively known as the background sequence, which lasted about three weeks. The end of the contact was marked by a comment in the sequence. This allowed the non-thrusting part of the sequence to be built early with plenty of time to review. Then, during the accelerated thrust build, the thrust sequences were added. These sequences were designed to start from a known state (HGA pointed at Earth), and return the spacecraft to a known state, and were guaranteed to not overlap the non-thrust windows. The thrust sequences were only two to five days in duration and were modeled with the longer background sequence as each thrust design was completed.

All of this was enabled by developing a timing file that all of the teams used to build their products. So, early in the process, MPST built the initial background sequence based on the timing file. This background sequence contained all of the activities except for the thrust activities, leaving empty windows for the maneuvers. The teams reviewed these products in about four days. Then Nav would start the thrust design about three days before the start of the maneuver, again using the timing file to dictate when thrusting was allowed. Next, Nav delivered the TVF, which the SCT used a Perl script to process and build the thrust sequence. This script encoded all of the interfaces between thrust and non-thrust windows so that the thrust sequence would start and end consistently with the background sequence. MPST then modeled the thrust sequence with the background sequence and all of the teams reviewed the products. Fig. 3 outlines the process. Since the only new aspect of the model was the thrust sequence, the review process was streamlined and completed in four hours. The entire thrust sequence design could now be completed in three workdays, or even 48 hours by utilizing non-work hours. By developing rules to control the interfaces with the thrust windows, encoding those rules in a script, and then testing the script rigorously when development time was available before arrival at Vesta, it ensured that the thrust sequence could be built correctly in a very quick turn around. The reliability of this process was essential because there would not be time to rebuild the thrust sequence if errors were detected. This process was used on about 25 thrust sequences so far, and no errors have been encountered.

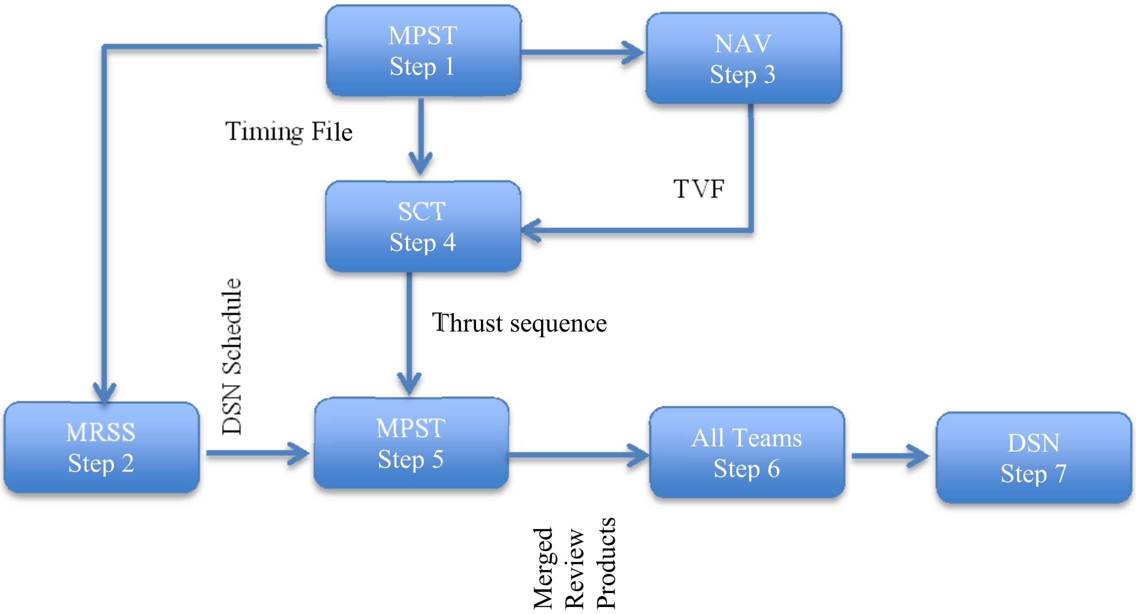

Figure 3. Thrust Uplink Process

2. Science Sequence Development

Furthermore, science operations, which involved pointing the science instruments at Vesta to acquire data and then turn the HGA to Earth and play it back, made the sequencing process more complex. Not only were there more commands being issued, but three separate instrument teams provided inputs to the sequences, involving more people in the process. Also, the science observations in two of the science phases were timed such that the data were acquired when the spacecraft was over the lit surface and played back when the spacecraft was over the dark side. This resulted in the sequences being sensitive to the phasing of the orbit. Again, with the uncertainty in the gravity field coupled with maneuver execution error, the final timing of the sequences needed to be adjusted as late as possible. This was accomplished by constructing the sequences relative to epochs. For example, the time to issue a command was not defined as 4pm on Tuesday, but instead was five minutes after the 15th descending equator crossing. The sequences were then built and reviewed with a reference trajectory over a span of about two weeks. Then, three days before the sequences were to begin executing, the epochs were updated to refine the timing based on the latest orbit knowledge. The process to change the epochs was completely software driven, reducing the risk of introducing errors. Also, the epoch definitions were stored in a separate file from the sequence commands, providing a clear mechanism to prevent errors from being introduced at the last minute. This step concluded with a full team review in less than eight hours. Without using the epoch process, this last minute timing shift could not have been completed reliably in such a short time – it would have required five to seven workdays, instead of three, which would have resulted in larger timing errors relative to the actual spacecraft orbit. While these errors would not have been unsafe, they would have wasted resources because the sequence would image the dark side of Vesta and the overall science return of the mission would have been reduced.

B. Downlink Process

With the more frequent contacts during the thrust periods and the science phase contacts being coupled to the orbit period, the downlink analysis process needed to be streamlined further. The three orbital altitudes that were used in the mission had periods of 69 hours (Survey Orbit), 12.3 hours (High Altitude Mapping Orbit (HAMO)) and 4.3 hours (Low Altitude Mapping Orbit (LAMO)). Therefore, the first altitude had a 34 hour contact every 69 hours (almost three days) and the second altitude had a 5.5 hour contact every 12.3 hours (two contacts per day!). The third altitude changed the observation strategy from observing on the lit side and playing back on the dark side every orbit to collect over 2000 hours of nadir pointed data with the GRaND instrument, resulting in contacts three times a week, on Monday, Tuesday, and Friday. This seemingly odd pace was chosen to fit in the thrust design cycle for orbit maintenance so the work occurred during the workweek.

The first process improvement was to create conditional red alarm limits. Normal red alarm limits compare the value of the channel to fixed limits. Conditional limits were designed to change the limits based on a particular state. For example, telemetry for a particular unit may be produced, even if the unit is off and not producing any valid data. The system telemetered this data, even though it was not realistic and not consistent with the operational state of the unit. Since the red alarms are designed for the operational state, the non-realistic values would have tripped red alarms. Earlier in the mission, the SCT manually changed the red alarm limits when the status of various units changed. This was both tedious and introduced a lag in the system because the on board FP system could autonomously change the state of the system. Additionally, the science instruments utilized many more state dependent limits where the state was changed regularly by the sequences. The conditional limits used a function in the ground software to set a red alarm limit, or enable or disable red alarms altogether, based on the value of a different channel. For example, if the nominal range of channel A is from 4.5 to 5.5, but only if channel B is equal to ‘ON,’ then the code would monitor channel B and when it changes to ‘ON,’ it set the red alarm limits of channel A to 4.5 to 5.5, and when channel B changes to ‘OFF,’ it disables the red alarm for channel A. This update both reduced workload and increased the robustness of the health and safety monitoring.

The increased frequency of contacts drove several more evolutionary improvements. First was development of an automatic scheduler that scheduled the telemetry processes based on the sequence products. While the scheduling task didn’t take a lot of time to do during cruise, the increased pace at Vesta would have increased the time required to perform it because one, there were more contacts, and two, those contacts were irregularly scheduled. This scheduling tool then read from the products the actual times of all of the contacts and playback times to automatically schedule the appropriate tasks at the correct times. Since some of the health and status monitoring required a human to review data in real time, a more automated check was required to reduce the amount of data that an engineer needed to review in real time, resulting in increased robustness and reduced workload. For example, the HAMO orbit had two contacts every day, which would have required staffing off of normal work hours every day for about a month. This was addressed by developing a script that would query telemetry near the beginning of a contact and create a status email containing the status of critical channels and results of the rule checking tool that was developed in cruise. While it still required some review, it was much more streamlined and didn’t require the on call engineer to log in until after the automated check found an issue.

Another area of concern was the science data recording and playback volumes. Since the FC and VIR instruments could collect much more data than could be played back in a single contact, managing the data on board was paramount. While there were sequence design tools to address the planning, it was difficult to see how the actual performance matched the plan. Since the telemetry was very context sensitive and didn’t require immediate review, an automated rule checking tool was not appropriate, and the human eye/brain computer can analyze a plot much quicker than being able to define an algorithm to do the checking. However, the behavior of the telemetry data did not lend itself to auto-scaling on a plot for a variety of reasons, which made the automated plots almost useless. Once again, a script was developed to query the telemetry after every pass and intelligently set up the plots with appropriate scales to be able to visualize the data management. This script became another entry in the automatic scheduler, along with the status report, and the batch plots developed in cruise. Without the automatic scheduler, too many hours would have been spent scheduling all of these various reporting tools.

The last, for now, improvement was to automate a thrust execution report based on the IPS telemetry. This was again driven by the need to improve the maneuver execution as much as possible while at the asteroid. Throughout cruise and even early Vesta operations, the navigation team did not require a reconstructed thrust performance report, but in order to meet the requirements to get to the lower orbits, Nav requested a reconstruction after every engineering data playback. Up until then, the reconstruction was manual and relied on engineering judgment to assess the impact of transients, which would have required an engineer to do the analysis at any time of any day. An effort was made to be able to perform this function 'lights out' and specific algorithms were defined that could be coded in Perl. This tool was then developed, and automatically scheduled to provide the reconstructed report without a human in loop within two hours of the data playback. This was especially useful to prevent requiring the IPS engineer to be available to process the data at any time of day, and provided reliable reconstructions helping to improve the maneuver deliveries.

IV. Conclusion

Making these improvements in both uplink and downlink processes enabled the small operations team on the Dawn project to achieve the tremendous science return, which has exceeded the requirements. These improvements greatly improved the robustness of the processes and allowed the team to focus on the new and unexpected issues that arose transitioning from cruise to science operations. There were many spacecraft configuration changes that were new experiences in flight that required the attention of the team to assess and adjust based on performance. Many of these have already been documented by the project as lessons learned and will form the building blocks for the remainder of the mission. For example, the improved thrust sequence build process will be incorporated into the cruise phase from Vesta to Ceres, as well as the automated scheduling tools.

Appendix A
Acronym List

AAN	Automatic Alarm Notification
DSN	Deep Space Network
FC	Framing Camera
FP	Fault Protection
GDS	Ground Data System
GRaND	Gamma Ray and Neutron Detector
HAMO	High Altitude Mapping Orbit
HGA	High Gain Antenna
IPS	Ion Propulsion System
JPL	Jet Propulsion Laboratory
LAMO	Low Altitude Mapping Orbit
MOS	Mission Operations System
MPST	Mission Planning and Sequencing Team
MRSS	Multi-Mission Resource Scheduling Services
NASA	National Aeronautics and Space Administration
Nav	Navigation Team
SCT	Spacecraft Engineering Team
SEP	Solar Electric Propulsion
TVF	Thrust Vector File
VIR	Visible and Infrared Mapping Spectrometer
Xe	Xenon

Acknowledgments

This paper was prepared at Jet Propulsion Laboratory, California Institute of Technology, under contract to the National Aeronautics and Space Administration. T. M. Weise would like to thank Dr. Marc Rayman for his brainstorming and review assistance with this paper. I would also like to acknowledge the tremendous work of those implementing the tools described in this paper: Kristina Larson, Eric Marcin-Cuddy, and Hal Uffelman.

References

---

[1] Rayman, M.D., Patel, K.C., "The Dawn project's transition to mission operations: On its way to rendezvous with (4) Vesta and (1) Ceres," *Acta Astronautica,* Vol. 66, 2010, pp. 230-238.

[2] Rayman, M.D., Fraschetti, T.C., Raymond, C.A., and Russell, C.T., "Dawn: A mission in development for exploration of main belt asteroids Vesta and Ceres," *Acta Astronautica*, Vol. 58, 2006, pp. 605-616.

[3] Smith, B.A., Vanelli, C.A., Lee, A.Y., "Understanding Spacecraft Agility for Orbit ,Transfers on the Dawn Low-Thrust Mission, *35th Annual AAS Guidance and Control Conference,* AAS 12-097, Breckenridge, Colorado, 2012.