Design & Execution of Dawn HAMO to LAMO Transfer at Ceres

Daniel J. Grebow*, Brian M. Kennedy†, Dongsuk Han†, and Gregory J. Whiffen*

Jet Propulsion Laboratory, California Institute of Technology

4800 Oak Grove Drive, Pasadena, CA 91109-8099

On October 23, 2015, the Dawn spacecraft left the High Altitude Mapping Orbit (HAMO) around Ceres and began its final decent to the Low Altitude Mapping Orbit (LAMO), arriving on December 15. The transfer between the two science orbits, a tight spiraling trajectory with over 100 revolutions, required the operations team to perform weekly maneuver designs for a period of 50 days. While the first six weeks of the transfer executed as planned, unexpectedly the spacecraft incurred a multi-sigma delivery error to the final science orbit that was subsequently clean-up at the first orbit maintenance maneuver. In this paper we discuss the design architecture for the transfer in detail, including challenges the team faced in flying the transfer and lessons learned.

I. Introduction

Dawn is the ninth NASA Discovery mission selected to explore the protoplanet Vesta and dwarf planet Ceres to learn more about the origins and evolution of the solar system. The spacecraft was launched on September 27, 2007 and arrived at Vesta in July 2011. After one year of science operations in orbit around Vesta, the spacecraft departed to travel towards the very different primitive body Ceres. The spacecraft arrived at Ceres in March 2015, and has been in orbit ever since. The nominal plan for the mission is to remain in orbit around Ceres until the spacecraft consumes the propellant resources necessary to control the spacecraft, at which point it will become an artificial moon of the dwarf planet. See Russel et al.[2] and Rayman et al.[1] for more details about the Dawn mission.

Vesta and Ceres are unique bodies in the main asteroid belt. Together the comprise nearly 40% of the total mass of the main asteroid belt. They are relatively large and spherical in shape, unlike most asteroid belt objects. Vesta and Ceres are believed to be vestiges from a time when objects were coalescing together in a process known as accretion that may have occurred at the *dawn* of the solar system. The lack of large bodies between Mars and Jupiter is believed to be caused by Jupiter's gravity, where instead collisions over billions of years have caused a region of many, rather small disparate objects, comprising the main asteroid belt. Vesta and Ceres have remained intact despite continuous bombardment from neighboring asteroids.

To learn more about these two fascinating worlds, the Dawn spacecraft would need to orbit Vesta and Ceres at many different altitudes and over a long period of time. No other mission has orbited two extraterrestrial bodies. This capability was enabled by Dawn's unique inspace propulsion system. The Dawn spacecraft is equipped with three ion engines. The engines provide a small, continuous thrust over a very long period of time. To date, the thrusters have logged over 5.5 years of accumulated thrusting, and they have a total impulse capability of 11.3 km/s.[3]

There are many challenges to flying a spacecraft with ion propulsion. Simply designing a trajectory to go from point A to point B in space while continuously thrusting is a challenging process, as there are infinitely many trajectories connecting the two points. Ultimately trajectories that optimize some objective are desirable. For the designing Dawn's maneuvers we use the low-thrust and optimization software Mystic[4,5] and the statistical maneuver design tool Veil,[6] a Monte Carlo tool wrapped around Mystic. The maneuver is designed on the ground and uplinked to the spacecraft to be flown open loop. The length of time the spacecraft can be flown in this open loop fashion before it diverges from the planned course depends on the

*Outer Planet Mission Analysis Group, Mission Design and Navigation Section, M/S 264-820.

†Outer Planet Navigation Group, Mission Design and Navigation Section, M/S 264-820.

uncertainties in flying the spacecraft and the sensitivity of the dynamical environment. For interplanetary cruise, design cycles can be as long as a month. For the tight, spiraling trajectory from HAMO to LAMO, cycles cannot exceed one week. Other challenges include agility limitations inherent in commanding the spacecraft to follow rapidly changing thrust profiles, and the loss of two of four reaction wheels, which completely changed our strategy for maneuver designs.

In this paper we discuss these challenges in detail. After a brief description of the Dawn spacecraft and mission phases at Ceres, we focus our discussion on the HAMO to LAMO transfer, the longest of the transfers performed at Ceres. We provide a detailed description of the architecture design process for this transfer. A new software tool was developed to analyze a wide range of objectives allowing us to assess spacecraft agility. For the first six design cycles of the transfer, the spacecraft executed the maneuvers as planned. However, in the final cycle of the transfer to LAMO the spacecraft incurred a multi-sigma delivery error. In this paper we discuss the source of this delivery error and potential ways to avoid such problems in the future.

II. Ceres Mission Design

II.A. Dawn Spacecraft

A diagram of the Dawn spacecraft is provided in Figure 1. (For more information about the spacecraft, see Rayman and Patel.[7]) As shown in the diagram, the spacecraft has three ion propulsion system (IPS) thrusters, one thruster on the +z axis, and two others canted 45 degrees from the center-line thruster. Only one IPS thruster is operated at a time. The spacecraft and IPS thrusters are powered by the two large solar panels extending in the ±y directions. Spanning at total of 20 meters in length, the solar panels provide 10 kW of power at 1 AU. The solar panels are always facing the Sun, even during IPS thrusting. Not shown on the diagram are the Reaction Control System (RCS) thrusters, a system of six hydrazine thrusters located at various places on the spacecraft. The spacecraft is also equipped with four wheels for momentum control. However one of the wheels stopped functioning in June 2010, and in August 2012 a second wheel failed shortly after departing Vesta. After losing two wheels the team was uncertain of the reliability of the remaining two, and therefore attitude was controlled entirely via the RCS for the entire interplanetary cruise between Vesta and Ceres, as well as the spiral down to LAMO. The wheels would be used again in LAMO in conjunction with RCS thrusting to save hydrazine. However even without the use of wheels in LAMO, the prime mission could have been completed.[8]

II.B. Ceres Orbits and Transfers

The science orbits at Ceres are similar to the orbits at Vesta.[9] Like Vesta operations, at Ceres there were four science orbits: RC3 (Rotational Characterization orbit 3), Survey, HAMO (High Altitude Mapping Orbit), and LAMO (Low Altitude Mapping Orbit). A description of the orbits is provided in Table 1 and the orbits are shown in Figure 2. See Polanskey et al.,[10] Han et al.,[11] and Whiffen[12] for more details about the Ceres mission. The total mission time planned for operating the spacecraft during the prime mission around Ceres was 1.5 years.

Table 1. Description of transfers and science orbits at Ceres.

Mission Phase (Purpose)	Description	Start to End (Duration)
Approach (Transfer)	End interplanetary cruise and capture to RC3, 5 design cycles, 7-day builds	Dec 27, 2014 to Apr 23, 2015 (117 days)
RC3 (Science)	Circular orbit with radius = 14,000 km, period = 15.3 days, beta angle = 5 deg, inclination = 90 deg	Apr 23, 2015 to May 9, 2015 (16 days)
RC3 to Survey (Transfer)	2 design cycles, 4-day builds	May 9, 2015 to Jun 3, 2015 (25 days)
Survey (Science)	47:6 repeat ground track, circular orbit with radius = 4718 km, period = 3.1 days, inclination = 90.74 deg	Jun 3, 2015 to Jul 1, 2015 (28 days)
Survey to HAMO (Transfer)	5 design cycles, 4-day builds	Jul 14, 2015 to Aug 16, 2015 (32 days)
HAMO (Science)	25:12 repeat ground track, circular orbit with radius = 1943 km, period = 18.9 hours, inclination = 90.52 deg	Aug 16, 2015 to Oct 23, 2015 (69 days)
HAMO to LAMO (Transfer)	7 design cycles + 2-day TCM, 4-day and 3-day builds	Oct 23, 2015 to Dec 13, 2015 (51 days)
LAMO (Science)	56:93 repeat ground track, circular orbit with radius = 849 km, period = 5.4 hours, beta angle = 45 deg, inclination = 90.57 deg	Dec 13, 2015 to Jun 30, 2016 (200 days)

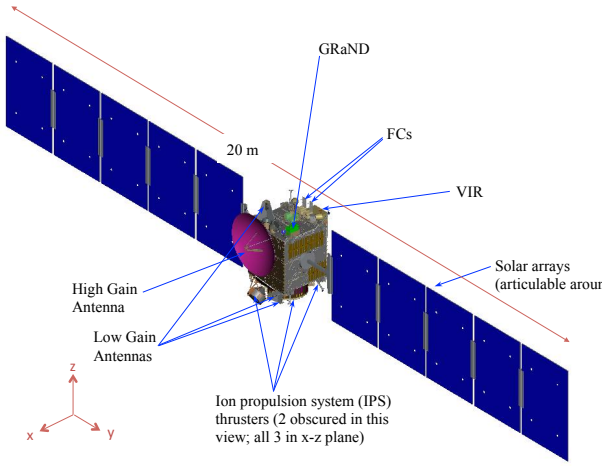

Figure 1. The Dawn spacecraft.

III. HAMO to LAMO Transfer

The purpose of the HAMO to LAMO transfer is to transport the spacecraft from the High Altitude Mapping Orbit (HAMO) to the Low Altitude Mapping Orbit (LAMO), the final destination of the Dawn prime mission. The transfer, the longest performed during Ceres operations, reduced the mean semi-major axis of the orbit radius from 1943 km to 849 km (see Figure 3), within two 2× the radius of Ceres. The maneuver required 51 days to complete, totaling 118 revolutions around Ceres. Thrusting for the start of the HAMO to LAMO transfer began on October 23, 2015 and ended on December 13, 2015.

The mean period of LAMO is 5.4 hours. The inclination of the orbit is targeted at 90.57 degrees, slightly off from polar to fix the beta angle at 45 degrees during LAMO by balancing the natural drift with the effects of J_2 (Ceres normalized $J_2 = 0.0012$). The period of LAMO is such that the ground track is a 56:93 repeat cycle (56 Ceres rotations = 93 LAMO revolutions), With this orbit the science team expected to achieve GRaND and gravity global coverage within one repeat cycle or 21.2 days, with a ground track longitude spacing of 3.9 degrees. An example groundtrack for one cycle of LAMO is shown in Figure 4. In LAMO Dawn would obtain GRaND and gravity coverage and framing camera chronology measurements, in addition to looking for evidence of Ceres-derived meteorites. High resolution VIR measurements would also be taken, fulfilling the Level-1 requirements for the mission.

The prime mission ended on June 30, 2016, after 200 days in LAMO operations, roughly 890 orbit revolutions around Ceres or nine complete cycles of the 56:93 repeat pattern. Each cycle had an Orbit Maintenance Maneuver (OMM) that consisted of two 12-hour windows for thrusting separated by one week of coasting.

End-of-mission was planned to occur in LAMO after all hydrazine was consumed by the spacecraft, resulting in loss of attitude control. Planetary Protection requirements stipulate that the team demonstrate the spacecraft would not impact Ceres for at least 20 years after loss of control. To demonstrate the stability of LAMO, the Dawn team performed a Monte Carlo analysis that showed the spacecraft will indeed remain in LAMO for at least 50 years, probably much longer.[13]

III.A. Mystic Veil Statistical Maneuver Design Tool

Designing maneuvers for the Dawn spacecraft is a long and involved process. The purpose of a maneuver is to transport a spacecraft from its current location A to a desired place B. With continuous thrusting there are infinitely many trajectories connecting points A and B in space, and ultimately optimal trajectories are desired. Usually the optimization objective is to minimize propellant usage, although other objectives may be more important for flying a spacecraft (as discussed below). All trajectories for the Dawn mission are designed with the trajectory optimization software package Mystic[4,5] that incorporates an optimization algorithm based on Differential Dynamic Programming.[14]

A potential trajectory from point A to point B is called a *reference trajectory*. The reference trajectory

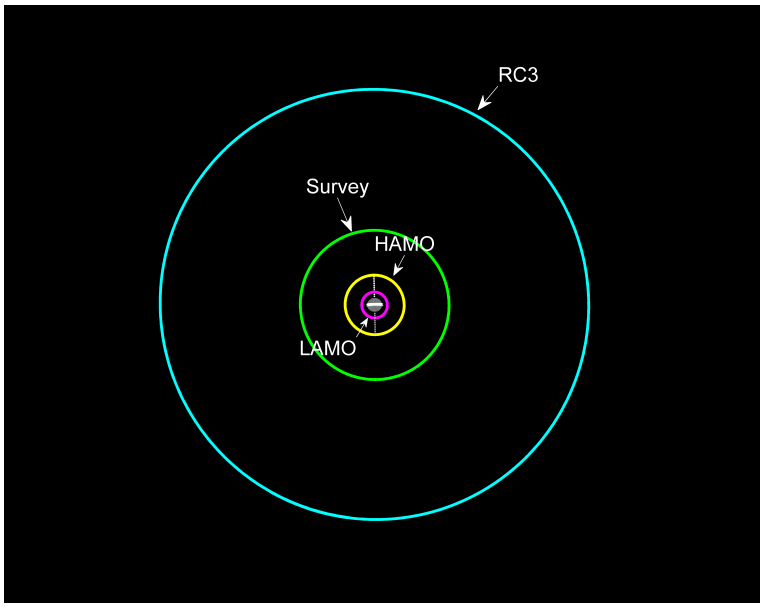

Figure 2. Ceres science orbits.

includes many design parameters such as maneuver build times, design cycles, Trajectory Corrections Maneuvers (TCMs) and Maneuver Expansion Periods (MEPs). (For example, see Figure 5.) The build time is the time from the data-cutoff, where the DSN receives final data from the spacecraft, to the beginning of the design cycle. During this time period the maneuver is designed by the ground team operating the spacecraft. Once the design is built, it is uplinked to the spacecraft and flown open loop for a pre-determined length of time called a *design cycle*. A maneuver expansion period, or MEP, is a period of forced-coasting built into the reference trajectory. The purpose of the MEP (which can also be imposed in the reference trajectory via a duty cycle) is to absorb statistical errors due to imprecision of orbit determination, gravity uncertainty, and thruster modeling inaccuracies, enabling the spacecraft to be transported to the target location with margin. During maneuver design operations, the spacecraft is allowed to thrust during a MEP. A TCM is like MEP except it spans an entire design cycle. TCMs are usually placed at the end of a transfer, after a short quiet period of no thrusting. The purpose of the TCM is to clean-up injection errors. The reference trajectory also includes coast durations where the spacecraft turns the HGA to Earth, or periods where the engine is throttled down; thus, the reference incorporates the team's best knowledge of how the transfer will be flown during operations.

It is sometimes challenging to determine the appropriate length of the design parameters. Longer build times give the team more time to work on building maneuver sequences (reducing stress) while degrading the delivery to the start of the design cycle. Similarly, longer design cycles result in fewer builds, also reducing the workload of the grounds team at the expense of delivery since the time the spacecraft is thrusting open loop is longer. Additionally, human factors such as the work week or holidays might be considered in determining the length of design cycles. The length of MEPs and TCMs are designed to be as short as possible without inhibiting the spacecraft from transferring to the destination target within the desired accuracy.

To help determine the length of the design parameters, the software package Veil was created.[6] Veil is a statistical maneuver design tool wrapped around Mystic that demonstrates confidence in maneuvering the spacecraft in the presence of uncertainty. A schematic for the Veil process is shown in Figure 6. Every Veil run starts with a reference that is broken into design cycles as previously described in Figure 5. A seed of 'truth' samples (typically number in the thousands) are randomly selected from an injection covariance before thrusting begins. Every truth sample has a unique initial state, gravity and pole (represented by purple in diagram in Figure 6). The goal of the maneuver designs is to guide the truth samples to the final target to within a desired accuracy. A sample state, pole, and gravity is randomly selected around each truth sample at the data cut-off using the knowledge covariances, and nonlinearly mapped forward to the start of the cycle. Then Veil solves for an optimal thrust profile from the knowledge state to the reference state target at the end of the cycle (see red trajectory in diagram in Figure 6). The thrust profile is applied to the truth

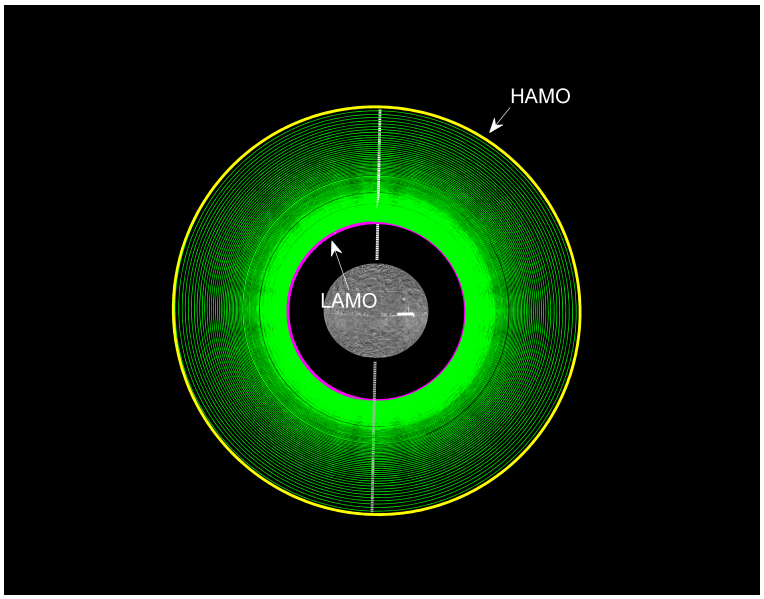

Figure 3. HAMO to LAMO transfer: 118 orbit revolutions.

sample including maneuver execution errors and small impulses from momentum desaturation maneuvers and/or RCS thruster firings. The process repeats for each design cycle, until the transfer is completed.

As previously mentioned, Veil needs to generate perturbations to the entire system state during each segment of a Monte Carlo sample run. It samples fully correlated covariance matrices to generate these perturbations. The Orbit Determination (OD) team is responsible for performing the covariance studies to generate these covariances. The covariance study uses the to-be-studied reference trajectory as a model. Errors are assumed for RCS execution, IPS delivery errors, spacecraft location and uncertainties in the Ceres rotating frame, Ceres GM and Ceres gravity harmonics. Using the planned schedule of tracking and imaging, the OD team simulates observations of the spacecraft using Doppler and ranging data, and observations of the Ceres surface using imagery from the Dawn camera. These observation simulations are then filtered by the OD team using the same software tools that will be used in operations. Once filtered, the OD team can assess the formal post-filtered uncertainties of the spacecraft state and Ceres parameters, and export the system error knowledge as a set of fully correlated covariances. These are expected to reflect our actual knowledge of the spacecraft and Ceres parameters at critical points in the orbit transfer. In practice at Vesta, there were several Vesta parameters whose final current best estimate (CBE) value was significantly different from earlier estimates. These are described in Kennedy.[15] Because of this, an engineering judgment was made to increase the formal uncertainties by a factor of five in order to provide conservatism and increase confidence in the results of the Veil Monte Carlo studies. Some CBE Ceres parameters were also found to be significantly different from earlier estimates in the Ceres mission to Kennedy.[16]

The Dawn maneuver design team runs Veil for all transfers several months prior to flying the transfer. For every Veil run the feasibility of the samples (ability of the design to target back to the reference state) and the delivery of the truth samples are assessed. The length of the design parameters are adjusted until 99% confidence is achieved and the delivery is adequate to meet the science objectives. If the uncertainty has been properly allocated, Veil simulations should represent all scenarios the maneuver design team may encounter during normal operations.

III.B. Transfer Architectures

Designing the architecture for the HAMO to LAMO transfer began two years before the transfer executed by the spacecraft. A summary of all architectures studied during this time period is provided in Table 2. Each Veil run takes roughly 5 days to complete, running in parallel utilizing 80 cores of a compute cluster. For the HAMO to LAMO transfer, each cycle targets 6-states from the reference at the end of the cycle, following our anticipated strategy for operations. If a sample is incapable of targeting the six-state because of lack of thrust authority, then that sample is deemed infeasible. Samples that are infeasible early in the transfer may

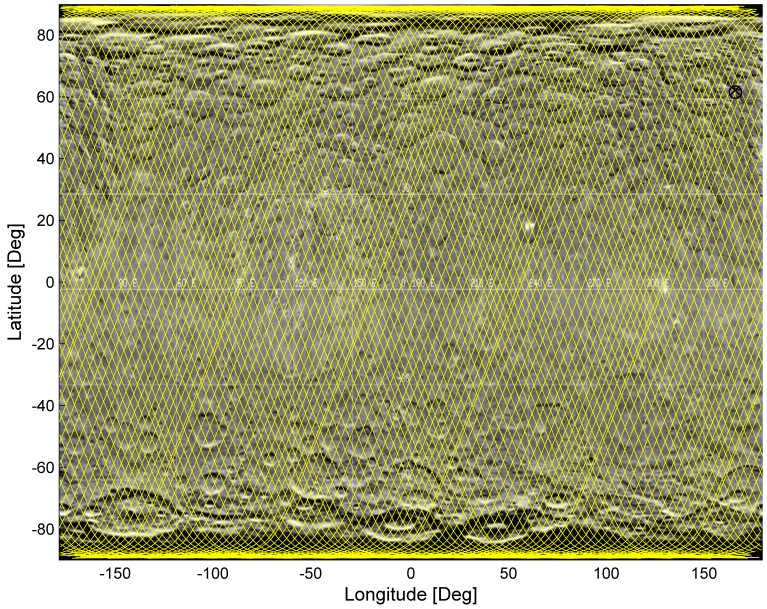

Figure 4. Designed 56:93 repeat ground track for LAMO.

recover downstream in later segments. The percentage of feasibility for the last segment is recorded in the fifth column in Table 2. Another parameter that was used to assess performance of a particular architecture is delivery to the final science orbit, as indicated by the difference in longitude from the beginning to the end of a cycle. This quantity (mean $\pm$ 1σ) is provided in last column of Table 2. A zero value corresponds to a delivery that is perfectly achieving the desired 56:93 repeat pattern.

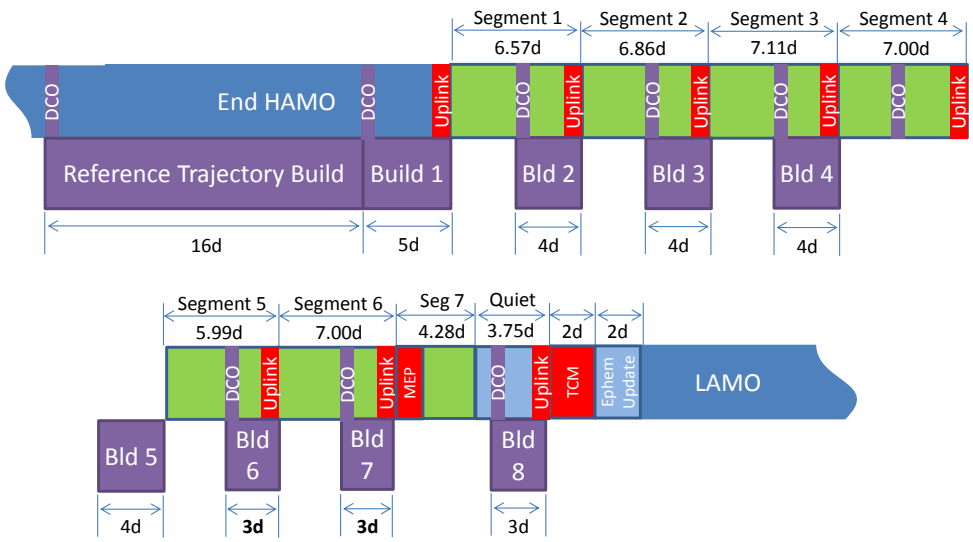

Figure 5. Architecture for HAMO to LAMO transfer, including build times (purple), design cycles (green), and MEP/TCM (red).

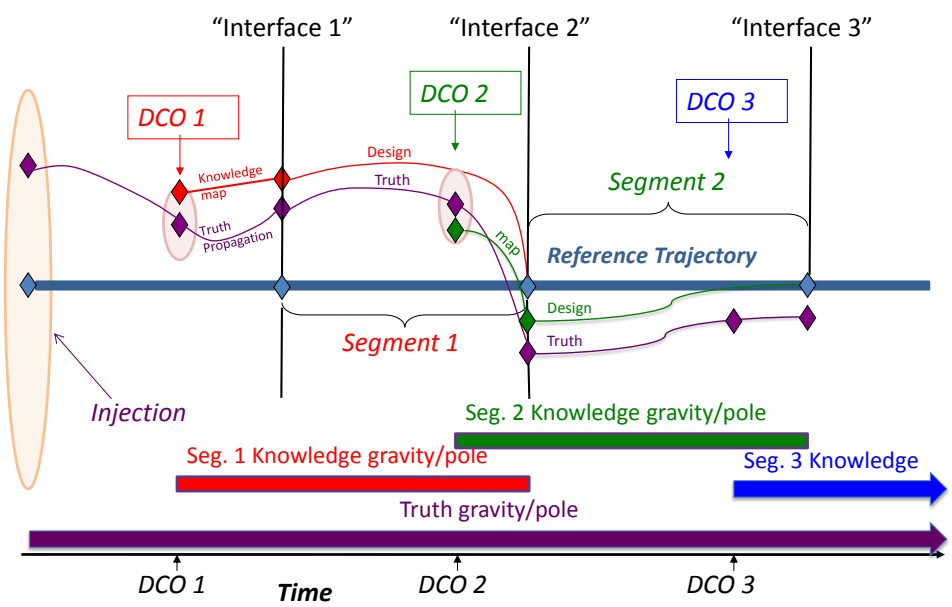

Figure 6. Schematic for statistical maneuver design tool Veil.

Table 2. Summary of Veil runs for HAMO to LAMO architecture study.

	Date	Name	Description	Feasibility	Predicted Delivery Cycle Longitude Error
Escape Vesta	**Sep 4, 2012**				
Veil Run 1	Aug 14, 2013	CH2L1-02	Noted that LAMO was close to 3:5 resonance, many samples delivering to 3:5	99.60%	$36.5^\circ \pm 0.8^\circ$
Veil Run 2	Aug 28, 2013	CH2L1-03	Same as CH2L1-02 but incorporates optical data from HAMO injection	99.70%	$0.9^\circ \pm 13.5^\circ$
Veil Run 3	Sep 25, 2013	CH2L1-04	Double size of estimated RCS thruster firings	93.60%	$3.4^\circ \pm 39.6^\circ$
Veil Run 4	Sep 25, 2013	CH2L1-05	No RCS thruster firings modeled	100.00%	$0.5^\circ \pm 3.5^\circ$
Veil Run 5	Sep 25, 2013	CH2L1-06	Same as CH2L1-03 except no RCS thrust modeled in LAMO	99.85%	$0.9^\circ \pm 10.7^\circ$
			gave ± 2, ± 1, 0-σ cases to science, science determined all coverage requirements satisfied		
Veil Run 6	Sep 25, 2013	CH2L1-07	Same as CH2L1-06, no TCM	100.00%	$-18.5^\circ \pm 16.5^\circ$
Veil Run 7	Nov 6, 2013	CH2L1-06S	Same as CH2L1-06 but with 5x scaled-up covariances	87.90%	$1.2^\circ \pm 44.7^\circ$
Veil Run 8	Nov 6, 2013	CH2L4-02	Exaggerated asymmetric Ceres	94.70%	-
Veil Run 9	Apr 23, 2013	CH2L1-10	5x covariances, all optical	99.80%	$0.7^\circ \pm 10.2^\circ$
Veil Run 10	Apr 23, 2013	CH2L1-11	Combining first two cycles into one two-week cycle	99.90%	$0.7^\circ \pm 10.0^\circ$
ISB	Jun 18, 2014		Delivered Architecture CH2L1-10		
Veil Run 11	**Oct 1, 2014**	**CH2L10-02**	**Updated spacecraft bus power model (shortened final cycle)**	100.00%	$1.0^\circ \pm 7.0^\circ$
Veil Run 12	Nov 13, 2014	CH2L10-03	Same as CH2L10-02 but with duty cycle MEPs	99.90%	$-0.2^\circ \pm 9.1^\circ$
NAG	**Nov 13, 2014**		**Architecture CH2L10-03 presented**		
Veil Run 13	Jan 21, 2015	CH2L12-01	New architecture with appropriately scaled thruster model	100.00%	$1.6^\circ \pm 8.9^\circ$
Veil Run 14	Feb 4, 2015	CH2L12-02	All four day builds except first and last build	98.50%	$2.8^\circ \pm 10.6^\circ$
Veil Run 15	Feb 4, 2015	CH2L12-03	All three day builds except first	99.80%	$1.4^\circ \pm 8.0^\circ$
Veil Run 16	Feb 18, 2015	CH2L12-04	Same as CH2L12-02 but 1x covariances	99.20%	$2.9^\circ \pm 10.1^\circ$
Veil Run 17	Feb 18, 2015	CH2L12-05	Same as CH2L12-02 but second to last build also 3 days	99.40%	$1.8^\circ \pm 9.2^\circ$
Veil Run 18	Feb 18, 2015	CH2L12-06	Same as CH2L12-05 but 1x covariances	99.59%	$1.8^\circ \pm 7.4^\circ$
Capture at Ceres	**Mar 6, 2015**				
Veil Run 19	Jun 3, 2015	CH2L13-01	Use CH2L12-02 Architecture, Switch to CBE thruster for architecture	97.18%	$-1.8^\circ \pm 7.7^\circ$
Veil Run 20	Jun 3, 2015	CH2L13-02	Same as CH2L13-01 but last three builds are all 3 days	99.20%	$-1.6^\circ \pm 8.7^\circ$
Spacecraft Safe	**Jul 1, 2015**		**Safing occurred at start of Survey to HAMO transfer,**		
			Team decided to switch from FT3 to FT2		
Veil Run 21,22	Aug 26, 2015	CH2L15-01,02	FT2 and updated bus power model, 1x and 5x covariances	-	-
Veil Run 23	Aug 26, 2015	CH2L16-01	Utilize more power available, recognition of sensitivity to $C_{4,3}$ and $S_{4,3}$	99.70%	$-29.69^\circ \pm 8.0^\circ$
Veil Run 24	Sep 16, 2015	CH2L17-01	Same as CH2L16-01 but with more credible 4th degree harmonics	99.92%	$-0.01^\circ \pm 7.99^\circ$
Veil Run 25	Sep 23, 2015	CH2L17-02	Same as CH2L17-02 but with $100\times$ GM and $10\times$ J_2	98.29%	$1.49^\circ \pm 15.65^\circ$
Final Decision on Architecture	**Oct 15, 2016**		**Team accepted to fly architecture CH2L17-01**		
Start HAMO to LAMO Transfer	**Oct 23, 2016**				
Start LAMO	**Dec 13, 2016**				

The first architecture CH2L1-02 was designed based on intuition into the problem from experience. Initially the team planned to have 7 week-long segments followed by a 5.3-day segment and a 2-day TCM. The total transfer time was 60.3 days, and we planned to have 3-day builds for each segment. Models were based on data for Flight Thruster 3 (FT3), because RCS thruster firings were less noisy for the center-line thruster. The reference was designed with a scaled-back conservative thruster model. MEP allocation was one day for each segment and implemented as forced coasts in the reference. Knowledge covariances used for the Veil run were unscaled and based only on radiometric data.

As shown in Table 2, the feasibility results for CH2L1-02 were acceptable but the delivery error was too large. In fact, most of the truth samples were delivering to a nearby 3:5 resonance, which introduced huge gaps in coverage, compromising the science value of LAMO. (See Figure 7 for a graphic showing all resonances with cycles less than 150 days.) The team considered selecting a new LAMO further away from the 3:5 resonance, however, we first investigated the benefits of incorporating optical data in the OD knowledge covariances. At Ceres, the results of the covariance studies using only radiometric data showed a large correlation between the knowledge of the spacecraft altitude and the knowledge of the Ceres GM, J_2 and J_3 parameters. By not being able to distinguish between these parameters, individual knowledge of each parameter was poor. This poor knowledge was then reflected in the poor prediction performance of the trajectory. When adding optical data to the solution, the correlations between orbit altitude, GM, J_2 and J_3 were much weaker, allowing the filter to resolve much better knowledge of the individual parameters. This better knowledge fed forward into better trajectory predictions among the Monte Carlo samples. As expected, incorporating optical data into the state estimation dramatically improved delivery. The Navigation team gave $\pm 2\sigma$, $\pm 1\sigma$, and 0σ cases to the science team from the CH2L1-06 Veil run, and science determined that all coverage requirements would be satisfied.

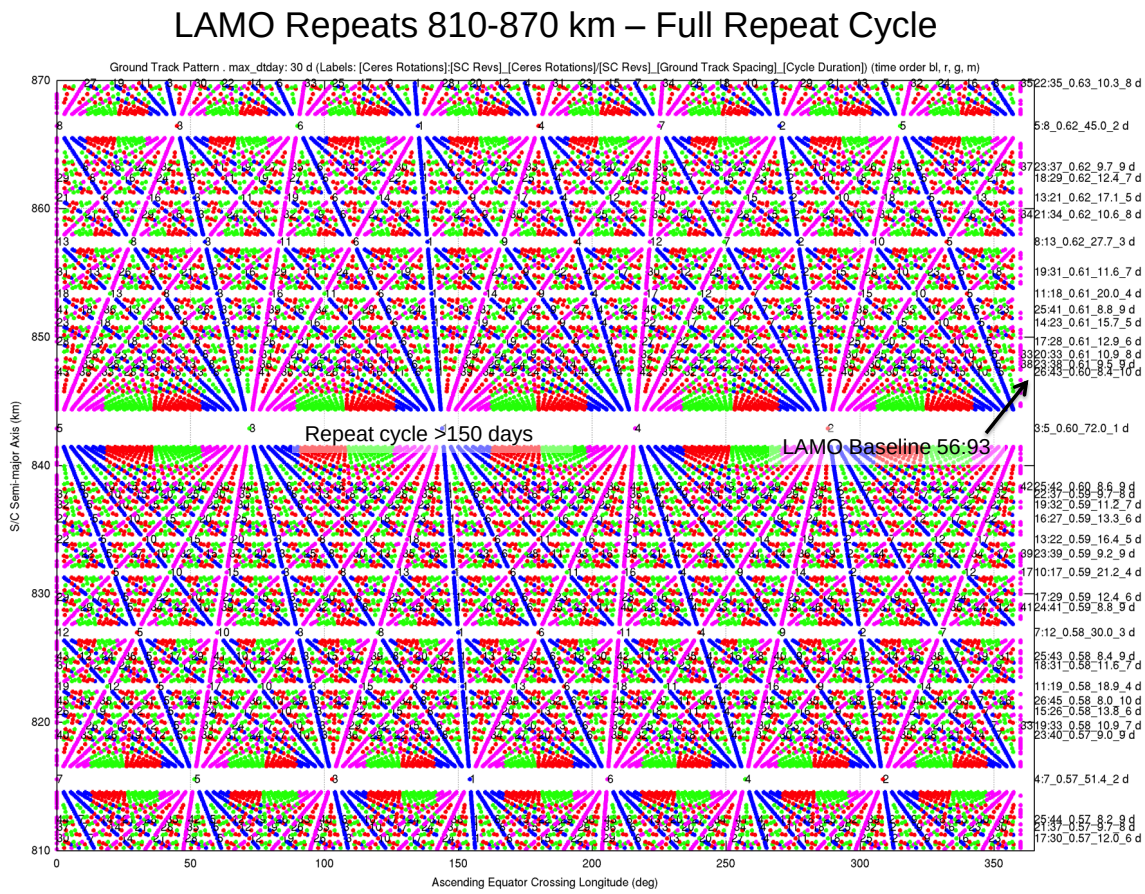

Figure 7. Ground track spacing for all resonances with cycle times less than 150 days. The 56:93 repeat pattern is annotated on the graph, just 5 km above the 5:3 resonance.

The effects of not having a TCM are demonstrated by Veil run CH2L1-07 which yielded a delivery error of $-18.5° \pm 16.5°$. Since the TCM showed a substantial improvement in delivery ($0.9° \pm 10.7°$), the

team decided that the HAMO to LAMO transfer should include a TCM. Veil run CH2L1-10, with optical covariances inflated by a factor of five, was used for the Initial Sequence Build (ISB) on June 18, 2014. By this time, we were also investigated the effects of an exaggerated asymmetric Ceres gravity field. The team was well aware of the challenges of passing through the 1:1 resonance for the Vesta HAMO to LAMO transfer.[9,17] However unlike Vesta, Ceres is believed to be in hydrostatic equilibrium and consequently the gravity field is rotationally symmetric. Therefore we did not anticipate difficulties for passing through the 1:1 resonances at Ceres, which occurs roughly mid-transfer HAMO to LAMO. Still, for confidence the team created and studied the effects of an exaggerated asymmetric Ceres gravity field. The gravity field was loosely based on Hubble observational data and stellar occultation measurements,[18] and Keck observational data of Ceres,[?] corresponding to a triaxial ellipsoid of uniform density with principle axes a = 500 km, b = 475 km, and c = 450 km. The results of Veil runs with the exaggerated asymmetric gravity were similar to Veil runs with nominal Ceres gravity.

The team also performed a study that compared the use of forced coast MEPs with duty cycle MEPs, and found that duty cycle MEPs yielded slightly better feasibility and delivery, perhaps because the duty cycle version spreads the MEP over the entire cycle rather than placing it at potentially undesirable locations. From this study the team decided to design the transfers with duty cycle MEPs. Architecture CH2L10-03, which incorporated duty cycle MEPs, was presented at the internal Mission Design & Navigation Advisory Group (NAG) meeting on November 13, 2014, along with the other Ceres transfer architectures. At this time the architectures for transferring between orbits around Ceres were considered mature and ready to fly.

Between the time of the NAG and Ceres Capture on March 6, 2015, there was a desire to decrease the amount of anticipated coasting in each design to conserve hydrazine which is consumed more rapidly at lower altitudes and during coast arcs. (Reducing the amount of coasting also introduces agility challenges for maneuver design during operation, a topic that will be discussed in more detail below.) To reduce coast times in each cycle, the reference architecture was redesigned with CBE thruster and power models. The total HAMO to LAMO transfer time decreased to 50.2 days, and instead of the 7 week-long cycles followed by a 5.3-day cycle, there were only 6 cycles needed, followed by a 2.5-day cycle. The build times for the cycles were also increased to 4 days whenever possible to allow the grounds team more time for the designs.

The spacecraft safing event at the start of the Survey to HAMO transfer triggered a switch to FT2 due to wear of the FT3 gimbal. Based on ACS software analysis, maneuver execution error models for FT2 were incorporated into new Veil runs. The final decision to fly architecture CH2L17-01, shown in Figure 5, was made on October 15, 2016, and thrusting for the HAMO to LAMO transfer began on October 23. A summary of the Veil thrust error models and the results from the CH2L17-01 Veil run are shown in Tables 3-6. Histograms for the expected coast times for each cycle are shown in Figure 8. Figure 9 is a plot of inclination and beta angle as a function of time for the HAMO to LAMO reference. The effects of the 1:1 can be observed 25 days into the transfer. We also see in Figure 9 that inclination decreases to 89.9 degrees even though the final target is less than a tenth of a degree away from the initial inclination. The change in inclination is optimally designed to leverage J_2 to achieve the beta target of 45 degrees.

Table 3. Thrust Error Assumptions for Veil Run CH2L17-01

IPS Pointing Error	Cycles 1-3: once every 10 min	uniform [0° to 0.25°] 77.2% probability
		uniform [0.25° to 1.9°] 22.8% probability
	Cycles 4-TCM: once every 10 min	uniform [0° to 0.36°] 64.5% probability
		uniform [0.36° to 1.98°] 35.5% probability
IPS Magnitude	All cycles: once every 10 hrs	normal $1\sigma = 0.5\%$
	Cycle 1-2: once every 10 min	normal $1\sigma = 0.5\%$
	Cycle 3: once every 10 min	normal $1\sigma = 0.4\%$
	Cycle 4: once every 10 min	normal $1\sigma = 0.35\%$
	Cycle 5-TCM: once every 10 min	normal $1\sigma = 0.25\%$
RCS Thrust	once every 10 hrs	normal $1\sigma = 10$ mm/s (each axis)

Table 4. Segment Knowledge Errors for Veil Run CH2L17-01

	Data Cut-off Knowledge σ_{pos} [km]	Segment Start Knowledge σ_{pos} [km]	True Delivery σ_{pos} [km]
Injection	-	-	133.48
Cycle 1	0.0558	25.26	59.8
Cycle 2	0.2783	24.75	88.43
Cycle 3	0.2344	23.55	86.81
Cycle 4	0.1437	22.7	99.97
Cycle 5	0.0864	20.32	62.02
Cycle 6	0.0547	11.62	56.85
Cycle 7	0.0597	12.13	35.58
Quiet	-	-	83.42
TCM	0.0415	13.31	17.12

Table 5. Feasibility for Veil Run CH2L17-01: 1272 samples

	Feasibility
Cycle 1	100.00%
Cycle 2	100.00%
Cycle 3	100.00%
Cycle 4	99.84%
Cycle 5	99.61%
Cycle 6	100.00%
Cycle 7	100.00%
TCM	99.92%

Table 6. Statistics for Delivery to LAMO

mean $\pm\sigma$	Reference 17	CH2L17-01
Insertion range (km)	859.25	859.32 ± 0.18
Insertion orbital period (hr)	5.492	5.4910.005
Insertion mean sma (km)	843.54	843.51 ± 0.22
Insertion inclination (deg)	90.549	90.549 ± 0.006
Insertion latitude (deg)	-12.5	-12.5 ± 1.8
Insertion beta angle (deg)	45.011	45.011 ± 0.008
Beta angle rate (deg/yr)	1.24	1.25 ± 0.88
All Consecutive Equator Xing		
Δ-longitude (deg)	143.20 ± 0.05	143.2 ± 0.1
Cycle 1 Δ-longitude (deg)	-0.91	-0.01 ± 7.99
Cycle 2 Δ-longitude (deg)	-3.19	-2.28 ± 8.03
Cycle 3 Δ-longitude (deg)	-2.13	-1.24 ± 8.02
Cycle 4 Δ-longitude (deg)	-1.86	-0.96 ± 8.06

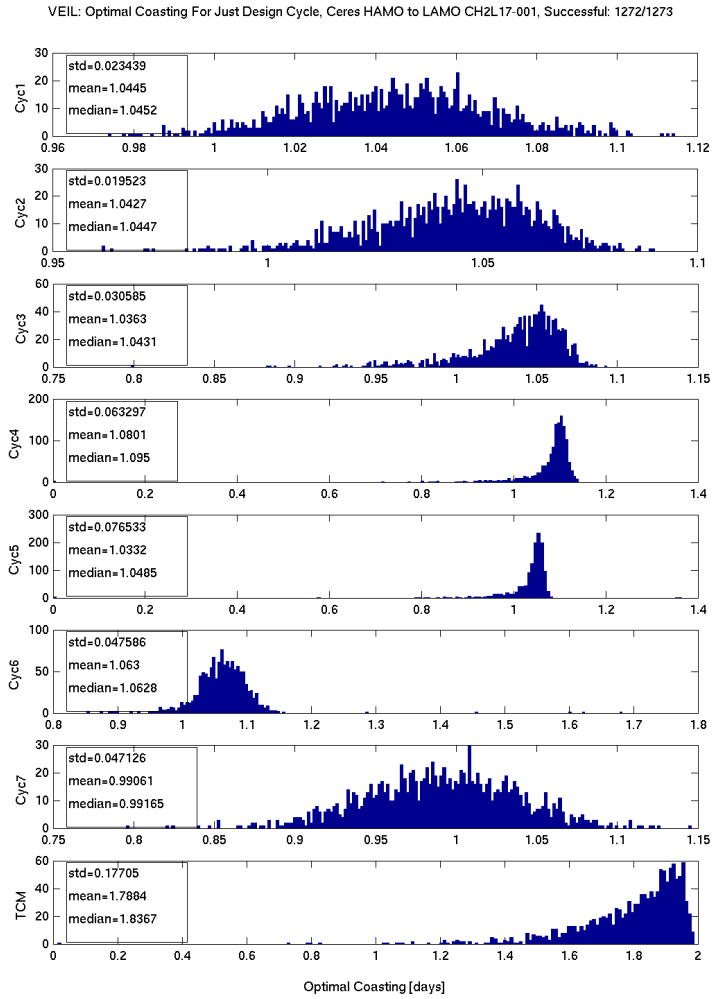

Figure 8. Cycle optimal coast durations: CH2L17-01.

III.C. Maneuver Design

Since controlling the attitude of the spacecraft requires more hydrazine at lower altitudes and during coasting, maneuvers for the HAMO to LAMO transfer were designed to minimize coasting as much as possible. This meant forcing the spacecraft to IPS thrust even during periods where it was optimal to coast, thrusting roughly one day longer than necessary due to the 1-day MEP allocation built into the reference to overcome statistical uncertainties in flying a spacecraft (recall Figure 8). For the maneuver designs of Cycles 1-6, the only time the IPS engine was *not* thrusting was during a 12.75 hour forced coast period to turn the HGA to Earth to uplink the maneuver. Cycle 7 had a 1-day forced coast in place at the beginning to reduce maneuver execution error.

Designing a maneuver to thrust continually introduced many challenges for the maneuver design team. Instead of mass optimal trajectories, maneuvers were designed using thrust direction optimization. In thrust direction optimization the thrust vectors are pulled as close as possible to a user-specified direction. In Mystic, the thrust direction is determined by the inputs right ascension (RA) and declination (Dec). Optimizing with an arbitrary RA and Dec usually introduced high attitude turn rates to follow a thrust profile that is quickly changing, a problem especially challenging for tight spirals where attitude rates are higher. High attitude rates could push the thruster gimbal beyond the operational limits, resulting in a safe mode. The key for the maneuver designs was to find the precise RA and Dec that optimized to a thrust profile within spacecraft agility limitations. Furthermore this process must be completed in one work day, since the total build times for HAMO to LAMO were only 3-4 days.

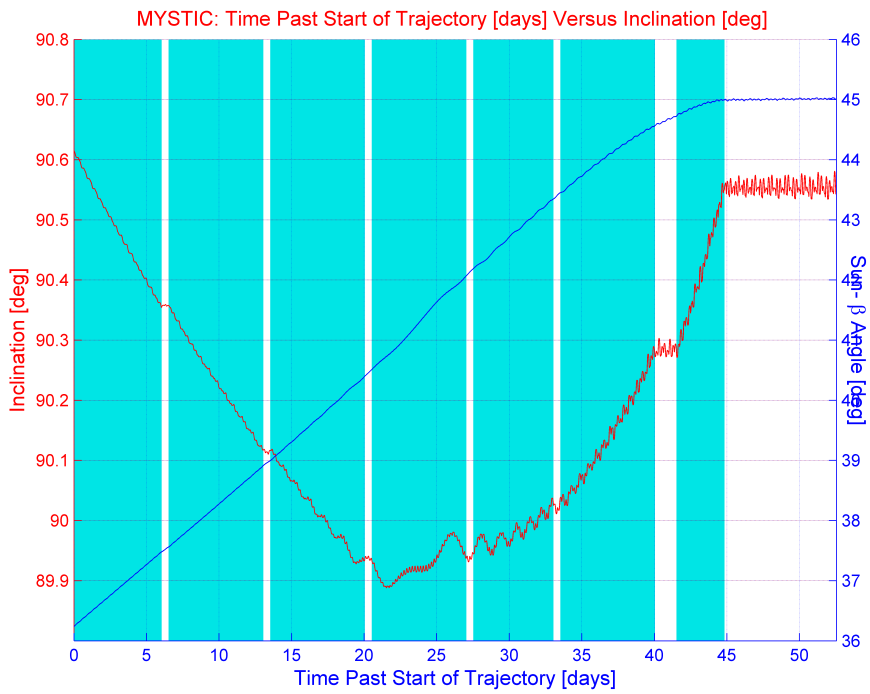

Figure 9. Inclination and beta angle for HAMO to LAMO transfer, shaded regions correspond to thrust-on times.

III.C.1. Mystic-Qstat (MQ) Tool

To quickly assess the attitude rates and gimbal excursions, the software tool Qstat was built for the maneuver design team.[20] The software tool MQ was created by John Smith to help determine the optimal thrust direction for maneuvers. MQ operates by sampling the entire trade space of RA and Dec, optimizing each sample with Mystic, and testing the results with Qstat to see if the agility requirements are met. The tool runs on a JPL compute cluster with 184 nodes. The maneuver design team ran MQ for preliminary OD deliveries a few days prior to the actual design providing the team with a good guess leading into the actual design of the maneuver.

An example output from from MQ is provided in Figure 10. Each square represents a different RA and Dec solution optimized in Mystic. The squares are colored by how well they satisfy certain conditions. Green means 'good', blue 'okay', and red 'unacceptable'. The background color behind the squares indicates whether or not the sample passed Qstat. The upper left corner corresponds to thruster gimbal margin; the upper right corner corresponds to propagation error incurred from the polynomial fit of the thrust vectors; the bottom left corner corresponds to how well the thrust vector polynomials are matching the Mystic solution; and, the bottom right corner corresponds to average angular error from following the thrust vectors. The best solutions are circled. This figure allowed the maneuver design team to quickly visualize the results, and was valuable in designing all the maneuvers for the HAMO to LAMO transfer. For this example maneuver design, the target select was RA = 152° and Dec = −6°. A plot of the thrust vectors on the RTN thrust sphere is shown in Figure 11.

MQ only takes a few hours to run, and results such as the plot in Figure 10 may be analyzed in real time as samples converge. Therefore MQ could be ran during a actual builds if necessary. However the team found that running it a few day prior to the design with the preliminary OD delivery was usually sufficient for operations.

III.D. Execution

A summary of the actual delivery errors for each design cycle of the HAMO to LAMO transfer is provided in Table 7. As shown in the table, the delivery error for the first six cycles was consistently sub-sigma, with many cycles delivering to within 10 km of the reference. In Cycle 7, the maneuver unexpectedly delivered 206.7 km away from the reference, a 6σ delivery error. Recall from Table 4 that the expected 1σ error for

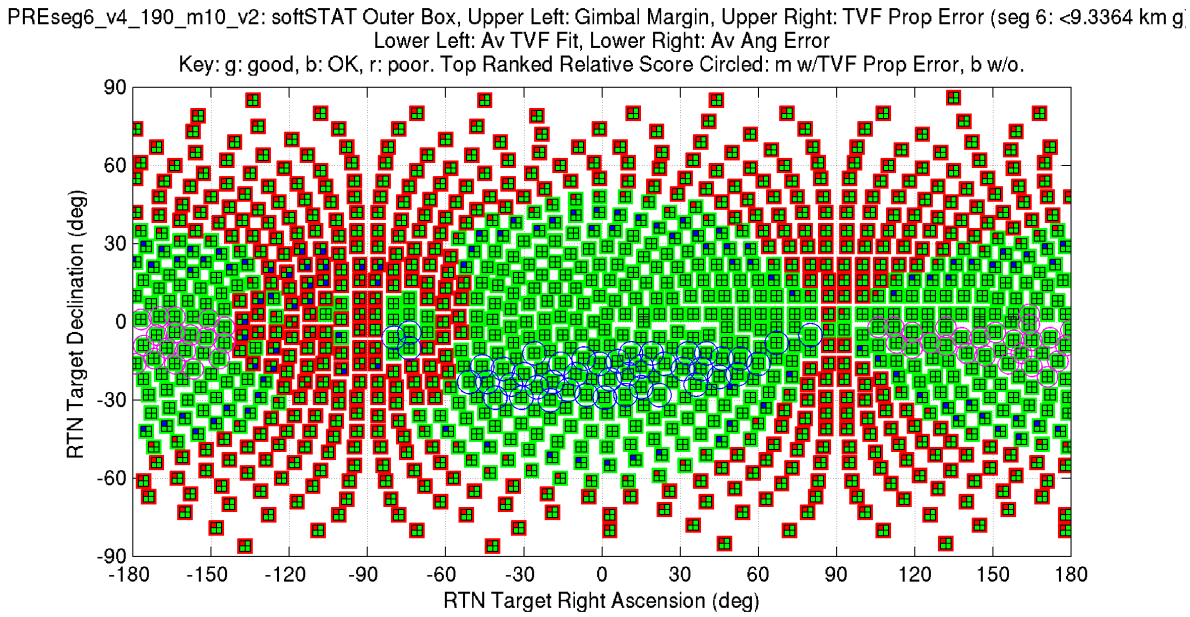

Figure 10. Example MQ output plot for Cycle 6.

Cycle 7 was 35.58 km.

In hindsight there are multiple reasons why the errors in Cycle 7 grew to the multi-sigma levels. The significantly large delivery error is believed to be mostly caused by small forces from RCS thruster firings. As the spacecraft period decreases from spiraling inward, the rates for controlling the attitude increase, resulting in greater control needed to follow the thrust profile while keeping the panels facing the Sun. The greater control resulted in more maneuvering by the RCS thrusters, which are unmodeled during maneuver designs. Understanding the effects of the RCS thruster firings on the orbit were not realized until after Cycle 7 was flown. Recall also that a safe mode at the start of the Survey to HAMO transfer caused the team to switch from using FT3 to FT2. For FT3 thrusting, the RCS thrusters used to control attitude are coupled. For FT2 the RCS thrusters are uncoupled and result in a net thrust in +z direction, 70% of which will be along the direction of thrust. Unmodeling the effects of the RCS thruster in the design therefore resulted in an over-thrust. This is corroborated by the fact the Cycle 6 delivered the spacecraft to a much lower than the designed LAMO, below the 3:5 resonance in Figure 7.

The attitude control team also recognized a difference in the location of Ceres in the spacecraft frame from Cycles 1-6 and Cycle 7, which can be observed in Figure 12 as differences in the angles between Ceres and the Dawn spacecraft axes during Cycle 7 as opposed to the preceding Cycles 4, 5 and 6. The different location of Ceres is believed have caused greater gravity gradient torques on the spacecraft, resulting in greater RCS impulses to control the attitude. Telemetry corroborates greater RCS thrusting in Cycle 7 compared to the previous cycles.

Finally, a smaller but still significant effect on delivery is due to inaccuracies in predicting the gravity in Cycle 7. The poor estimation in the gravity is caused partly by unpredictable RCS thruster firings. However it was found that estimating the gravity to order six may have been insufficient for Cycle 7. Analysis has shown that using an eighth order gravity field would have resulted in a better estimate of the state.[16]

The poor delivery in Cycle 7 made the TCM necessary. Fortunately cleaning-up a 6σ delivery error, though unexpected, is not outside the capability of two days of IPS thrusting allotted for the TCM. (From Figure 8, we expected a 6σ delivery error in Cycle 7 to correspond to 1.3 days of thrusting in the TCM.) By this time the team was aware that RCS thruster firings were contributing to the error, and therefore the TCM was designed with a prediction of small forces for RCS. The state was also estimated with an eighth order gravity field. The actual thrusting in the TCM design was 1.5 days, and it effectively reduced the error from 514.8 km to 107.8 km, but the delivery was still on the multi-sigma level. The main cause of the delivery error in the TCM is believed to be from a poor prediction in the OD state used in designing the

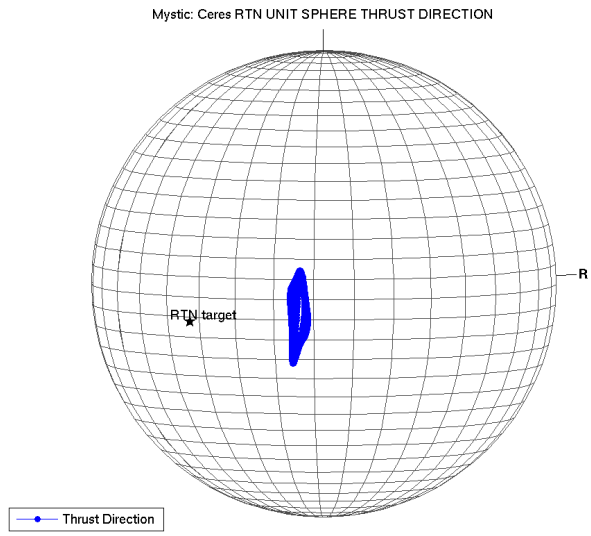

Figure 11. Example thrust profile corresponding to Cycle 6 maneuver design (blue). Compact solutions on the RTN thrust sphere typically satisfy Qstat agility requirements. Thrust direction target RA = 152° and Dec = -6° also shown.

TCM. The state could not be predicted accurately because of poor delivery in the previous cycle, as seen in Figure 12 where the phasing error during the quiet period is 20-30 degrees off from the design. The large phase error equates to an error in pointing the spacecraft and RCS thrusting. Because of this phasing error, the onboard spacecraft ephemeris was no longer correct, and the Dawn +z axis was not pointed towards Ceres. In this attitude, the delta-V from the RCS resulted in a different perturbation than predicted by the earlier ACS models. This resulted in the OD prediction for the Ceres-relative state at the beginning of the thrust being incorrect by several tens of km and over 10 cm/s. In hindsight, this is another reason for having ACS re-predict the effects of RCS activity based on more recent updates to the spacecraft trajectory. We also noticed the predicted small forces used in the design were under-estimating the actual thruster firings, and so again in the TCM there was an over-burn. The altitude of the orbit after the execution of the TCM was slightly higher than the designed LAMO, but at least it was further away from the 3:5 resonance.

In the next section we will describe the impact of these delivery errors on the science planned for LAMO.

Table 7. Actual delivery errors and sigmas for HAMO to LAMO transfer.

	Delivery Error [km]	Fraction of Sigma
Injection	92.82	0.70
Cycle 1	14.77	0.25
Cycle 2	0.70	0.01
Cycle 3	16.50	0.19
Cycle 4	5.83	0.06
Cycle 5	13.36	0.22
Cycle 6	3.69	0.06
Cycle 7	206.70	5.81
Quiet	514.80	6.17
TCM	107.80	6.30

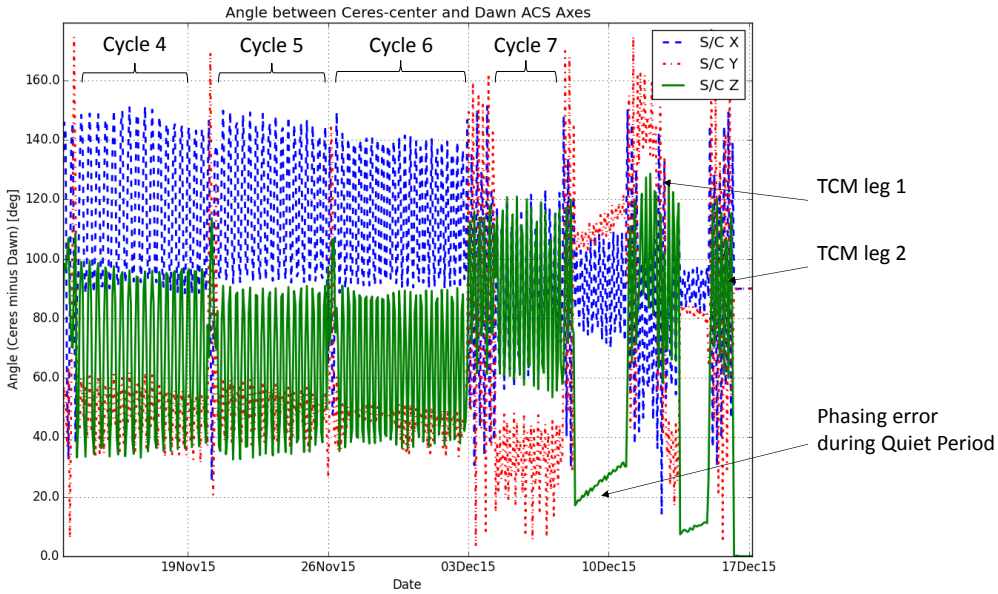

Figure 12. Spacecraft attitude for Cycles 4-7 and TCM.

IV. Summary of LAMO

Recall that Orbit Maintenance Maneuvers (OMMs) were planned to control the orbit more precisely for LAMO operations. Although the deliveries for Cycle 7 and the TCM were on the multi-sigma level, by the time of the first OMM, the causes of the errors were well-understood to control the spacecraft back to the reference. Each OMM is comprised of two 12-hour periods for thrusting separated by one week of coasting. We anticipated using only 3.4 hours of statistical thrusting, with the rest of the 24-hour period open for thrusting as margin. In reality we used 14.5 hours of thrusting in OMM1 to control back to the reference. Unlike Cycle 7 and the TCM, OMM1 delivery was only 3.0 km off from the reference, a 0.05σ error.

The OMMs performed as anticipated for controlling the LAMO. Six of the nine planned OMMs were canceled because they were not needed. Despite the initial delivery error, the first and second cycles of LAMO provided 95% coverage of the surface, although the much anticipated central bright region of Occator crater was in the missing 5%.[21] The rest of LAMO operation has been very smooth. The Dawn science team achieved all the planned objectives well before end of the prime mission in June 2016.

V. Lessons Learned

There is a 'chicken-egg' problem between ACS small force estimation and the design of the trajectories to fly the Dawn spacecraft. ACS needs the trajectory to predict the small forces, and those small forces in turn alter the course of the trajectory. This problem was less apparent at higher altitudes where the small forces had a comparatively smaller effect. The team worked to overcome this 'chicken-egg' problem by iterating with ACS in subsequent designs. Otherwise, the only way to avoid such a problem is by simulating the attitude control with the trajectory as it is being designed.

Still, even if the small forces were perfectly modeled, they only accounted for about 60% of the delivery error in the final cycle from HAMO to LAMO. The remaining 40% is believed to be caused by poor gravity resolution. As described in Kennedy,[16] there was an error in the gravity field estimation used for the design of segment 7. This was found to have been caused by not estimating the higher orders of the gravity field, allowing the increased resolution of the gravity field to be aliased into the lower order terms. This was confirmed by re-running the solutions using a higher order field estimation and finding improvement in the predictability of the new estimated field.

All error sources, such as gravity uncertainty or RCS thruster firings, are incorporated into the Veil simulations when designing and testing a particular architecture. The effects observed in the final cycle and

TCM were not seen in any of our Veil runs, which implies that the Veil assumptions were not reflecting real uncertainty. In the future, Veil runs should reflect more accurately uncertainty of unknown parameters.

The final lesson learned is not based on analysis as much as it is a cautionary rule of thumb. From Table 7 errors for Cycles 1-6 were all much smaller than expected. For all of these cycles the maneuver design process was exactly the same. Cycles were designed with thrust direction optimization and the spacecraft was thrusting for the entire duration of the cycle. In Cycle 7 we changed maneuver design strategies. This cycle began with one day of forced coast, because we believed that would help reduce maneuver execution error. While there is no evidence this contributed to the poor delivery in Cycle 7, one cannot help but connecting the unexpected behavior in Cycle 7 to this change in the design process. So a rule of thumb might be *if a design process is working well, don't change it.* Changing strategies also increases the likelihood of making a mistake, although there is no known procedural mistake in the designs of Cycle 7 and the TCM.

VI. Conclusion

In this paper we discuss in detail the Ceres HAMO to LAMO transfer for the Dawn mission. We describe the statistical maneuver design tool Veil, and summarize in detail the design of the transfer architecture. The HAMO to LAMO transfer was one of the longest transfers flown during Ceres operations. Design cycles and build times were much smaller than for the other transfers, a natural consequence of maneuvering the spacecraft closer into the gravity well of Ceres.

A new software tool (MQ) was also introduced in this paper. The tool was instrumental in allowing the maneuver design team to assess the entire design space for selecting thrust profiles that are within the agility limitations of the spacecraft.

The final design cycle of the HAMO to LAMO transfer delivered the spacecraft into LAMO not as precisely as anticipated. Errors were partially corrected during the TCM and fully corrected in the first LAMO maintenance maneuver. Despite the initial delivery error, the first and second cycles of LAMO provided 95% coverage of the surface, and the rest of the LAMO operations has executed nominally.

Acknowledgments

The authors thank Dan Parcher for his aid in learning the maneuver design software tools. John Smith was instrumental in architecture studies and maneuver design processes. We acknowledge his important contributions.

The work described in this paper was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration. Copyright 2016 California Institute of Technology. Government sponsorship acknowledged.

References

[1]Rayman, M.D., et al, "Dawn: A mission in development for exploration of main asteroid belt asteroids Vesta and Ceres," *Acta Astronautica*, Vol. 58, No. 11, 2006, pp. 604-616.

[2]Russel, C.T., et al., "Dawn: A journey in space and time," *Planetary and Space Science*, Vol. 52, No. 5-6, 2004, pp. 465-489.

[3]Garner, C.E., et al., "The Dawn Vesta Science," *The 32nd International Electric Propulsion Conference, Paper No. IEPC-2011-326*, Wisbaden, Germany, September 11-15, 2011.

[4]Whiffen, J.G., "Static/Dynamic Control for Optimizing a Useful Objective," United States Patent No. 6,496,741. Issued Dec. 17, 2002, Filed Mar. 25, 1999.

[5]Whiffen, J.G., "Mystic: Implementation of the Static Dynamic Optimal Control Algorithm for High-Fidelity, Low-Thrust Trajectory Design," *AIAA/AAS Astrodynamics Specialist Conference, Paper No. AIAA 2006-6741*, Keystone, Colorado, Aug. 21-24, 2006.

[6]Parcher, D.W., and Whiffen, G.J., "Dawn Statistical Maneuver Design for Vesta Operations," *AIAA/AAS Astrodynamics Specialist Conference, Paper No. AAS 2011-180*, New Orleans, Louisiana, Feb. 13-17, 2011.

[7]Rayman, M.D., and Patel, K.C., "The Dawn project's transition to mission operations: On its way to rendezvous with (4) Vesta and (1) Ceres," *Act Astronautica*, Vol. 66, No. 1-2, 2010, pp. 230-238.

[8]Rayman, M.D., et al., "Dawn's operations in cruise from Vesta to Ceres," *Acta Astronautica*, Vol. 103, pp. 113-118, 2014.

[9]Parcher, D.W., "Low-Thrust Orbit Transfer Design for Dawn Operations at Vesta," *AIAA/AAS Astrodynamics Specialist Conference, Paper No. AIAA 2006-6741*, Keystone, Colorado, Aug. 21-24, 2006.

[10]Polanskey, C.A., et al., "Architecting the Dawn Ceres science plan," in *Proceedings of the SpaceOps 2014 Conference*, 2014

[11]Han, D., et al., "Orbit Transfers for Dawn's Ceres Operations: Navigation and Mission Design Experience at a Dwarf Planet," *The 14th International Conference on Space Operations, Paper No. AIAAA 2016-2427*, Daejon, Korea, May 16-20, 2016.

[12]Whiffen, G.J., "Dawn's Mission at Ceres," *2016 AIAA/AAS Astrodynamics Specialist Conference, Paper No. AIAA 2016-XX*, Long Beach, California, September 13-16, 2016.

[13]Takahashi, Y., et al., "Forward Modeling of Ceres' Gravity Field for Planetary Protection Assessment," *2016 AIAA/AAS Astrodynamics Specialist Conference, Paper No. AIAA 2016-XX*, Long Beach, California, September 13-16, 2016.

[14]Jacobson, D.H. and Mayne, D.Q., *Differential Dynamic Programming*, Elsevier, Scientific, New York, 1970.

[15]Kennedy, B. M., et. al., "Dawn Orbit Determination Team: Trajectory and Gravity Prediction Performance during Vesta Science phases," *23rd AAS/AIAA Spaceflight Mechanics Meeting, Paper No. AAS 13-345*, Kauai, Hawaii, Feb. 10-14, 2013.

[16]Kennedy, B. M., et al., "Determination of Ceres Physical Parameters Using Radiometric and Optical Data," *2016 AAS GN&C Conference, Paper No. AAS 16-107*, Breckenridge, CO, 2016

[17]Whiffen, G.J., "The Stability of Powered Flight Around Asteroids With Application to Vesta," *AIAA/AAS Astrodynamics Specialist Conference, Paper No. AAS 2011-186*, New Orleans, Louisiana, Feb. 13-17, 2011.

[18]Thomas, P., et al., "Young, Differentiation of the asteroid Ceres as revealed by its shape," *Nature*, Vol. 437, 2005, pp. 224-226.

[19]Carry, B., "Near-infrared mapping and physical properties of the Dwarf-planet Ceres," *Astron. Astrophys.*, Vol. 478, 2008, pp. 235-244.

[20]Smith, B.A., et al., "Understanding Spacecraft Agility for Orbit Transfers on the Dawn Low-Thrust Mission," *35th Annual AAS Guidance and Control Conference, Paper No. AAS 12-097*, Breckenridge, Colorado, Feb. 3-8, 2012.

[21]Polanskey, C.A., et al., "Dawn Ceres Mission: Science operations performance," *14th International Conference on Space Operations 2016*, Deajeon, Korea, May 2016.