

Tuning the MESSENGER State Estimation Filter for Controlled Descent to Mercury Impact

Brian R. Page^{*}, Christopher G. Bryan[†], Kenneth E. Williams[‡], Anthony H. Taylor[§], and Bobby G. Williams^{**}
KinetX Aerospace, Space Navigation and Flight Dynamics (SNAFD) Practice, Simi Valley, CA, 93065

The Mercury Surface, Space Environment, Geochemistry, and Ranging (MESSENGER) mission is the seventh in NASA's Discovery Program. The spacecraft has been orbiting Mercury since March 2011 and after propellant reserves are depleted will impact the planetary surface at the end of March 2015. In preparation for the controlled descent, the MESSENGER navigation operations team has begun the process of updating the filter inputs for the state estimation parameters through a statistical analysis of the orbit determination solutions accumulated to date.

I. Introduction

The Mercury Surface, Space Environment, Geochemistry, and Ranging (MESSENGER) spacecraft is being flown as the seventh mission in NASA's Discovery Program.¹ The MESSENGER mission is led by the principal investigator, Sean C. Solomon, of Columbia University. The Johns Hopkins University Applied Physics Laboratory (JHU/APL) designed and assembled the spacecraft and serves as the home for project management and spacecraft operations. Navigation for the spacecraft is provided by the Space Navigation and Flight Dynamics Practice (SNAFD) of KinetX Aerospace, a private corporation. Orbit determination (OD) solutions are generated through processing of radiometric tracking data provided by NASA's Deep Space Network (DSN) using the MIRAGE software package.

The MESSENGER spacecraft successfully executed Mercury orbit insertion (MOI) on 18 March 2011 to begin a nominal one-year science investigation of the planet.^{2,3} An illustration of the MESSENGER spacecraft at Mercury in Fig. 1 shows the bus shielded from the Sun by the sunshade. Through careful shepherding of fuel reserves and the robustness of the onboard instruments, the mission has been extended and the scientific return enhanced. The mission extension has allowed for variations in orbital period and closest-approach altitudes beyond what was planned for the primary mission. During the second extended mission (XM2), the remaining propellant will be used to accomplish several periapsis-raising maneuvers, recurrently increasing the spacecraft's closest-approach altitude and postponing the eventual Mercury surface impact. The entire sequence of post-MOI orbit-correction maneuvers (OCMs) is depicted in Fig. 2, where OCM-1 to OCM-9 have already been executed and OCM-9 to OCM-



Figure 1. MESSENGER Spacecraft at Mercury.
Artist's conception, color-enhanced planet surface.

^{*} MESSENGER Lead Navigation Analyst, SNAFD, 2050 E. ASU Cir. #107, Tempe, AZ 85284, AIAA Sr. Member.

[†] MESSENGER Navigation Team Chief, SNAFD, 2050 E. ASU Cir. #107, Tempe, AZ 85284.

[‡] Flight Director, SNAFD, 21 W. Easy St. #108, Simi Valley, CA 93065, AIAA Sr. Member.

[§] Flight Director Emeritus, SNAFD, 2050 E. ASU Cir. #107, Tempe, AZ 85284.

^{**} Director, SNAFD, 21 W. Easy St. #108, Simi Valley, CA 93065.

12 are planned as part of the XM2 controlled descent sequence.⁴ Preliminary analysis has verified OCM targeting and design for each of these upcoming burns.

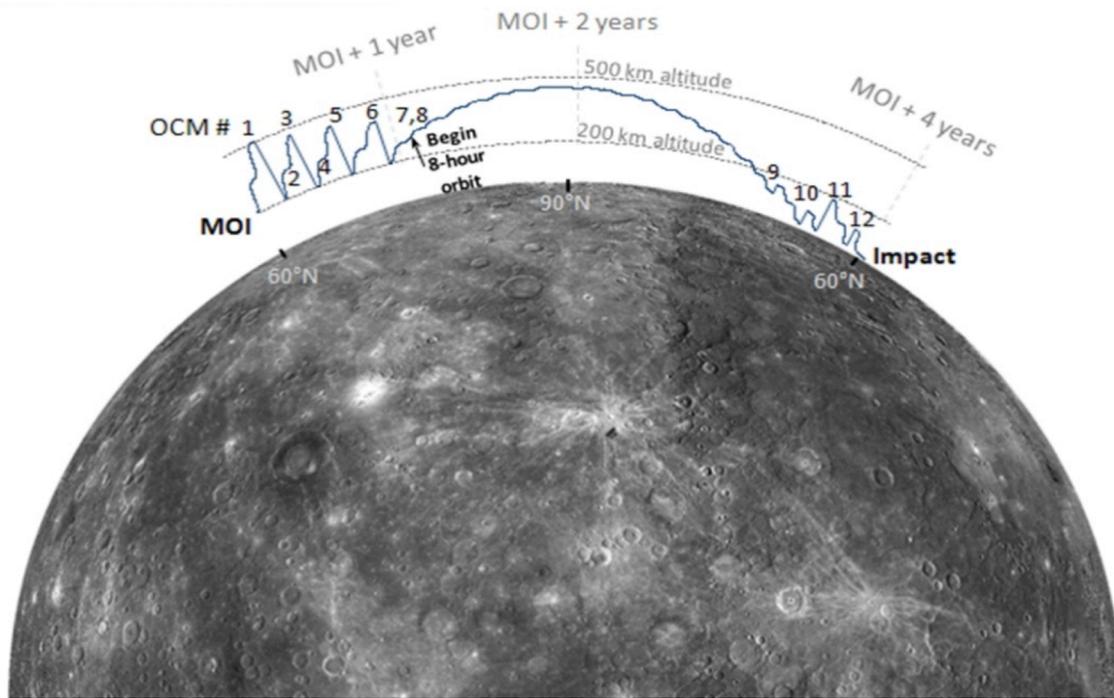


Figure 2. MESSANGER Periapsis Altitude Progression during the Mercury Orbital Mission Phase. Latitude markers are notional in this projection, such that spacecraft periapsis argument of latitude ranged from 60° at MOI, to an 84° peak, and approaches 59° at impact (courtesy of James McAdams, JHU/APL).

Since MOI there have been more than 150 OD solutions delivered by the navigation team to the mission operations team (nominally one per week), a number that is expected to grow to over 200 by the end of the mission. In preparation for the descent into a low-altitude (< 200 km) periapsis regime about Mercury starting in April 2014, a statistical analysis of previous on-orbit estimates was undertaken in order to refine the nominal and *a priori* values and uncertainties used as inputs to the estimation filter. Tuning the filter is an ongoing process, but the identification of certain problematic relative geometries between the spacecraft, planet, and Sun have motivated a redoubled effort to review estimation strategies and plan for worst-case contingencies during the upcoming descent toward Mercury impact currently predicted to take place in late-March 2015. It is the solar third-body perturbations on the orbit of the MESSANGER spacecraft that will eventually induce the collision with the planetary surface, ending the mission.

The estimated state vector for the spacecraft includes more than 500 parameters, far beyond simply the Mercury-centered-inertial (MCI) position and velocity. A 20×20 spherical harmonic planetary gravity model encompasses 438 terms, including the Mercury gravitational constant (GM, where G is the universal gravitational constant and M is the mass of Mercury). Solar and planetary radiation pressure (SRP/PRP) parameters include 60 specular and diffuse reflectivity coefficients, corresponding to a 10-flat-plate spacecraft model, and two scale factors. Mercury planetary ephemeris is estimated with six canonical orbital elements. Commanded momentum dumps (CMDs) and OCMs are incorporated in the filter, as these events occur typically once per week for the CMDs and during trajectory adjustment campaigns for the OCMs, such as the upcoming descent sequence. Velocity change (Δv) or thrust component parameters, three for each OCM burn segment, are estimated for each discrete propulsive maneuver. An itemized listing of these estimation parameters is presented in Table 1, along with the additional parameters included in the sensitivity matrix and used to enhance the accuracy of the associated computational *a posteriori* uncertainties output by the filter.³ The selection of these dynamic and observation model inputs are part of the state estimation tuning process, although any changes in modeling methodology are implemented with the utmost care and only after thorough verification and regression testing.

Table 1. MESSENGER State Estimation Filter Parameters.

Filter Variables		
Description/Parameter		Model Details
Solved For	Spacecraft Position/Velocity	MCI Cartesian components
	Radiation Pressure	SRP/PRP specular and diffuse reflectivity coefficients and PRP global-scale factors
	Mercury Gravity	20×20 spherical harmonic coefficients and GM
	Mercury Ephemeris	Solar system barycentric-inertial canonical element corrections to DE423
	Propulsive Maneuvers	OCM component thrust magnitudes/directions and CMD Δv components
Considered	Earth Ephemeris	Solar system barycentric-inertial canonical element corrections to DE423
	Earth Polar Motion	Surface position components and UT1 corrections
	Station Locations	Earth-fixed position component corrections
	Atmospheric Media	Refraction corrections for Earth wet/dry Troposphere and day/night Ionosphere

The nominal and *a priori* values and uncertainties for the estimation parameters are used to constrain the filter output in such a way that realistic solutions are produced within expected bounds. Those nominal values that are not static are typically derived from previous filter runs, such as the translational state (position and velocity), which is advanced to the epoch of the next weekly OD. The radiometric tracking data are initially fit to the nominal parameter values, and subsequent adjustments are performed in an iterative manner to minimize the cumulative observed minus computed difference between the radiometric measurements and those predicted by the models.

It is the modification of *a priori* estimation values and uncertainties that constitutes the primary mechanism for tuning the filter. Input uncertainties are one-standard-deviation ($1-\sigma$) values. If the estimates are unbiased and Gaussian, then two-thirds of the time the solutions should fall within this range. The goal of the state estimation tuning process is to induce the entire set of parameters to simultaneously produce unbiased, Gaussian results on a consistent basis. The adjustment of the input values and uncertainties used to seed the MESSENGER state estimation filter is accomplished through analysis of the empirical results corresponding to both operational and experimental OD solutions. This analysis is done to produce optimal trajectory predictions for regularly delivered orbit updates, which in turn facilitates the planning of science observations and enhances the value of the mission.

Variations in the relative weighting of different track data types can also be used to tune the filter. These weights generally correspond to the expected inverse variance of individual measurement points. However, particular circumstances and empirical behavior of the data may dictate that adjustments be made to the weighting scheme during a filter run, even from iteration to iteration, in order to arrive at a smooth path toward convergence. Superior solar conjunctions are a prime example of such problematic line-of-sight (LOS) geometry, where the signal carrying the Doppler and range data may be degraded as it passes through the solar plasma. In such cases *ad hoc* adjustments to data weights are often necessitated on a track-by-track basis, or even on separate intervals of the data within specific tracks. Large solar flares can also introduce undesirable noise into the track data, but this noise can sometimes be mitigated through selective deweighting as well, or ill-conditioned segments of the data can simply be deleted, if necessary. Experience has shown that an especially problematic relative geometry occurs during eclipses of the Sun by Mercury on the MESSENGER spacecraft, and the lower the spacecraft altitude over the planetary surface during the eclipse, the worse the effects are on the data residuals and the estimates. Deweighting is not

always appropriate in these instances though, as the issue here is not LOS signal behavior but rather fidelity issues with the PRP model.

II. Radiation Pressure Tuning

Radiation pressure perturbations on the MESSENGER spacecraft trajectory are treated in the estimation process through the inclusion of specular and diffuse reflectivity coefficients. These parameter sets are typically less well defined than the other dynamic inputs to the filter in terms of relative *a priori* uncertainties. SRP was modeled throughout the cruise phase of the mission, and manipulation of this perturbing force allowed solar sailing techniques to be used to target planetary flybys and MOI without the need for propulsive trajectory-correction maneuvers and the associated depletion of the available fuel margin, which could thus be allocated to extending the primary orbital mission at Mercury.^{5,6} PRP effects include both albedo and infrared re-radiation from the planetary surface as well as optional global-scale factors that can be used to magnify or attenuate these perturbations in the modeling process.

The relatively large uncertainties associated with these parameters are due in part to the complex interdependence of radiation pressure estimates on other filter models and environment variables. This list includes: a simplified spacecraft model, comprised of ten flat plates representing the surfaces on which incoming photons may impinge, along with the material properties of those surfaces; the attitude, or orientation, of the spacecraft with respect to the incoming radiation; the relative geometry between the spacecraft and the sources of radiation, the Sun, and Mercury; and the solar activity level during the time interval of the track data used for a particular estimate. In addition, as mentioned above each spacecraft surface exposed to incident radiation has both specular and diffuse reflectivity components incorporated as filter inputs.

Geometric optics and ray tracing best describe specular reflection, whereby nominally the angle of incidence equals the angle of reflection. Diffuse reflection corresponds to the resultant scattering of photons by a rough material surface, independent of the incidence angle. The diffuse reflectivity distribution model takes the form of a cosine function relative to the surface normal. Combining these two mechanisms allows the process to be modeled in a physically realistic manner, although the difference between distinct perturbing effects may not always be observable within the filter, in the sense that they can be definitively distinguished from one another. Likewise, the effects of SRP and PRP perturbations in general are not always completely observable, and thus the estimates of individual radiation pressure parameters may be subject to biases, aliasing, and cross-correlation. These effects are evident in the associated *a priori* 1- σ values that seed the filter to initiate state updates, and one goal of the filter tuning process is to reduce these uncertainties as much as possible relative to the associated *a priori* values.

It is primarily for these reasons that the radiation pressure parameters are good candidates for initial tuning through the statistical analysis of previously accumulated solutions. These parameters are often where unmodeled perturbations on the orbit tend to be absorbed, the preferred sink in the state estimation filter. Adjustment of the radiation pressure initialization values and uncertainties will have fewer undesirable effects on other estimated parameters than *vice versa*. This situation makes testing for regression and verification of improved results, in the form of fit metrics, more efficient and intuitive.

A. Solar Radiation Pressure Parameters

During interplanetary cruise, SRP parameters were initially estimated as areas, or an effective percentage of each nominal spacecraft surface from which incident radiation was reflected. Specular and diffuse reflectivity coefficients replaced these area estimates prior to the first Mercury flyby, in order to increase the fidelity and utility of the solutions. The original *a priori* values and uncertainties for these reflectivity coefficients were conservative approximations based on the thermal material properties of the various exposed spacecraft surfaces. These *a priori* inputs were adjusted intermittently on the basis of empirical results and then readjusted just ahead of MOI through a rough statistical analysis of the estimation solutions during the tail end of the cruise phase of the mission.

Experience during cruise revealed that within the filter these reflectivity coefficients could sometimes, albeit infrequently, violate the physical bounds prescribed to them by the modeling process. Reflectivity coefficients represent a percentage of the incident radiation that is reflected, either specularly or diffusely, and therefore should nominally assume values between 0 and 1. However, the filter is not hard-coded to enforce these bounds, and when such a violation occurs it is handled through manual intervention, which in turn provides useful diagnostic feedback during the iterative correction phase of the OD. This intervention usually involves reducing the offending parameter uncertainty in order to constrain the flexibility of the estimate relative to the *a priori* value, although the problem can also sometimes be remedied by modification of data weights or by deleting borderline out-of-family track data points.

Immediately subsequent to MOI, the problem with SRP parameters drifting outside of acceptable bounds during OD iterations was exacerbated by the necessary addition of PRP parameters into the state estimation process. Tuning the filter became an urgent issue at that point, as additional force and observation modeling was required to accommodate the orbital phase of the mission and many of these models could be tested rigorously only *in situ*. The Mercury gravity field, planetary ephemeris, spacecraft antenna selection and motion, and PRP were all being implemented operationally for the first time and together. The necessary modifications to operational procedures and adjustments to various aspects of the estimation process were continuously refined until such issues as those encountered with the reflectivity coefficients were minimized. This step was accomplished expeditiously without any adverse impact upon mission or science operations.

The lessons learned during previous filter tuning efforts have been once again applied in preparation for entry into the low-altitude campaign marked by the final series of OCMs and the controlled descent to Mercury impact. A statistically significant accumulation of data points, representing delivered, backup, and experimental on-orbit OD solutions, have been analyzed with the objective of optimally updating the state estimation filter inputs. An example of the history of operationally delivered OD results for a specific estimation parameter, in this case the SRP specular reflection coefficient for one of three component plates that model the spacecraft sunshade (SPEC01), is presented in Fig. 3. Differences between estimation solutions and corresponding nominal *a priori* values (Δ -solutions), post-fit formal $1-\sigma$ filter uncertainties defining the associated errors, and *a priori* $1-\sigma$ uncertainties for the example estimated parameter (red lines), are shown in the plot. Note that this is a plot of the difference between the *a priori* values and the estimated values, and the red lines are centered on zero since the solutions are initially assumed to be unbiased.

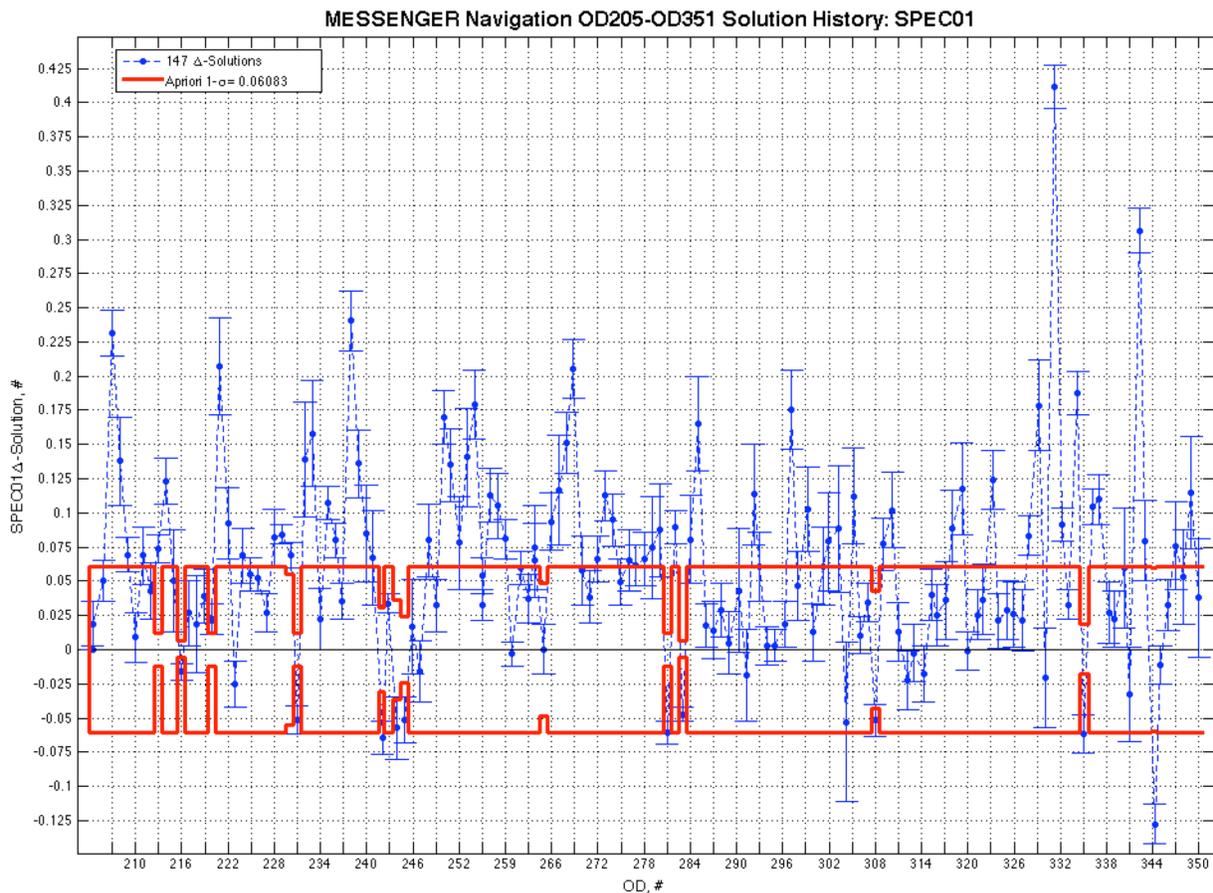


Figure 3. MESSENGER OD Solution History for the Estimated Parameter SPEC01.

Specific SPEC01 estimation solutions for which the *a priori* uncertainties (red lines) were reduced relative to the baseline value is an indication that the estimated parameter required additional constraint, imposed through the aforementioned manual intervention process during filter iterative correction, in order to prevent the results from drifting outside of allowable limits. The solutions for which this was necessary are identified by the nonlinear

segments of the red lines in the plot and were excluded from the subsequent statistical analysis because of their dependence on these external manipulations. The statistical analysis involves a determination of the root mean square average of the included samples, weighted by the output formal $1\text{-}\sigma$ uncertainties represented by the error bars displayed in the plot, along with the calculation of the standard deviation of the ensemble. The results of this analysis were then used to modify the nominal and *a priori* values, and uncertainties, needed to seed the state estimation filter. These changes were then tested to verify an improved fit to the data compared with the previously baselined values.

The results of the statistical analysis for SPEC01 are displayed graphically in Fig. 4. The bias, as determined by the weighted mean, in the delta solutions of Fig. 3 is quantified and the standard deviations for the SPEC01 solution ensemble are overlaid as $1\text{-}\sigma$ through $3\text{-}\sigma$ horizontal lines. This plot provides visual confirmation that the results generally conform to statistical expectations, with the possible exception of one or two outliers.

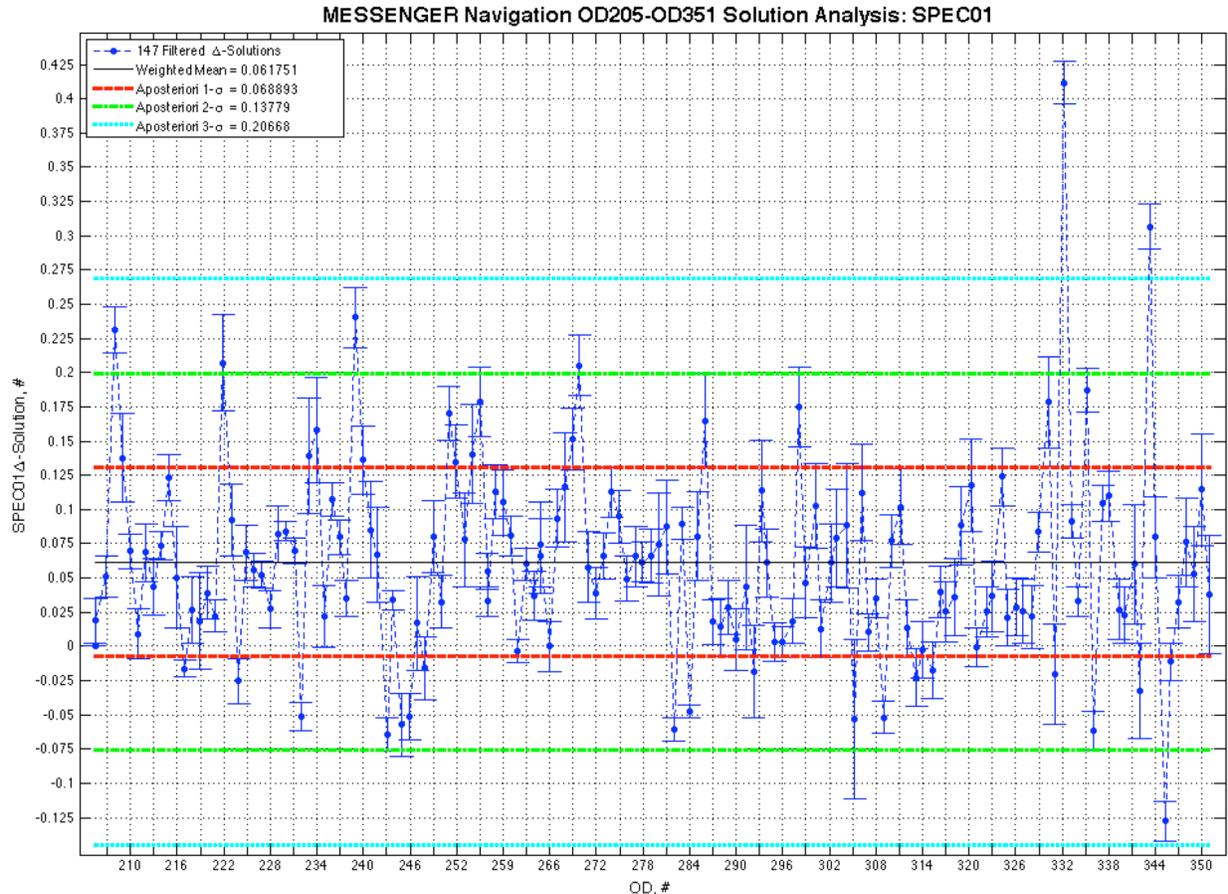


Figure 4. MESSENGER OD Solution Analysis for the Estimated Parameter SPEC01.

The same analysis protocol was exercised for the other 19 SRP coefficients of the 10-plate model used by the navigation operations team, and the updates to the *a priori* values and standard deviations for those reflectivity coefficients are presented in Table 2. The tuning results for the SRP specular reflectivity coefficients reveal substantial modifications to *a priori* values for the three parameters associated with the sunshade (SPEC01-03) and two that represent the back sides of the articulated solar panels (SPEC09-10). The only parameters that showed a substantial change in *a priori* uncertainties were the front surfaces of the solar panels, associated with the specular reflectivity parameters SPEC07-08, and diffuse reflectivity parameters DIFF07-08. These four groups of parameters: SPEC01-3, SPEC07-08, DIFF07-08, and SPEC09-10, are each strongly correlated in the filter, as are all radiation parameters associated with these specific spacecraft surfaces, because each group shares the same material properties.

Table 2. Tuning Results for SRP Specular and Diffuse Reflectivity Coefficients.

Parameter	<i>A Priori</i> Value		<i>A Priori</i> Sigma	
	Baseline	Update	Baseline	Update
SPEC01	0.07000	0.13	0.06083	0.07
SPEC02	0.07000	0.13	0.06083	0.07
SPEC03	0.07000	0.13	0.06083	0.07
SPEC04	0.07000	0.07	0.06083	0.06
SPEC05	0.07000	0.07	0.06083	0.06
SPEC06	0.07000	0.07	0.06083	0.06
SPEC07	0.2100	0.20	0.0183	0.04
SPEC08	0.2100	0.20	0.0183	0.04
SPEC09	0.07000	0.11	0.06083	0.07
SPEC10	0.07000	0.11	0.06083	0.07
DIFF01	0.26000	0.24	0.03924	0.04
DIFF02	0.26000	0.24	0.03924	0.04
DIFF03	0.26000	0.24	0.03924	0.04
DIFF04	0.26000	0.26	0.03924	0.04
DIFF05	0.26000	0.26	0.03924	0.04
DIFF06	0.26000	0.26	0.03924	0.04
DIFF07	0.0660	0.07	0.0222	0.03
DIFF08	0.0660	0.07	0.0222	0.03
DIFF09	0.26000	0.25	0.03924	0.04
DIFF10	0.26000	0.25	0.03924	0.04

The three sunshade components are correlated at 98% and the front and back surfaces of the solar arrays are correlated at 99.9% each. The SPEC04-07 and DIFF04-07 parameter groups are uncorrelated to other group members or other groups and represent the specular and diffuse reflectivity coefficients of the top, bottom, and sides of the spacecraft bus, which are shielded from incident SRP by the sunshade and are therefore nominally unaffected by the estimation process. Since these bus coefficients are shielded from the Sun, any change in their estimated values during a filter run may be an indicator of a problem with the setup or an unusual attitude orientation, and they are optionally included for diagnostic purposes. The numbering system described above and used to associate the discrete flat plates of the spacecraft model with radiation pressure coefficients carries over to the PRP reflectivity parameters as well, with distinct prefixes identifying different radiation types and reflection modes.

B. Planetary Radiation Pressure Parameters

The two forms of PRP modeled by the MIRAGE software are albedo and infrared radiation pressure. PRP parameters were incorporated into the estimation process immediately subsequent to MOI and initialized with filter inputs inherited from the SRP *a priori* values and uncertainties. However, whereas the flat plates of the spacecraft model representing the bus are shielded from the Sun by the woven ceramic sunshade, they are exposed to PRP on a regular basis. The perturbing accelerations of PRP are typically several orders of magnitude less than that of SRP, although during low-altitude periapsis passages these differences are narrowed appreciably, and when eclipse intervals occur SRP and albedo perturbations are eliminated completely.⁷ Thus during eclipse passages the normal suite of radiation pressure perturbations are reduced to planetary infrared re-radiation effects only. This outcome can cause modeling challenges around low-altitude periapses because of uncertainties in the Mercury surface temperature model.⁸

Analogous to previous descriptions of specular and diffuse reflection modes, albedo radiation can be considered to trace the geometric path of visible light reflected off the multitude of uneven planetary surface features and imparting a resultant momentum to the spacecraft, whereas infrared radiation is the effect produced by thermal emission from the planet when absorbed incident solar energy is re-emitted as heat in an isotropic manner.

1. Albedo Modeling

The state estimation filter models Mercury albedo effects as a project-defined single-valued bolometric albedo, characteristic of the average percentage of incident energy reflected from the visible sunlit planetary surface. Although the capability exists to model albedo with spherical harmonics, this formulation has not been considered necessary since there is little variation in the reflective properties of different terrain types on Mercury and the simpler alternative approach has usually yielded well-behaved estimation parameters. An optional albedo scale factor (ALBCOF) has been carried along in the filter since MOI but is not always utilized, and since its presence or absence in the parameter list doesn't appreciably alter the behavior of the individual reflectivity coefficient solutions, it is generally not estimated. The albedo parameters are normally better conditioned than the infrared and SRP coefficients and only very rarely require adjustments in the form of manual intervention to tighten the associated *a priori* uncertainties during filter runs.

The statistical analysis results used for tuning the albedo parameters are presented in Table 3, where the ALMU prefix is associated with the specular reflectivity coefficients and ALNU with the diffuse coefficients. The largest changes evidenced here are to the *a priori* uncertainties, where for example the albedo specular reflectivity uncertainties are considerably reduced, indicating that these coefficients are generally better determined than predicted by the original baseline uncertainties. The scale factor demonstrates the opposite effect in both *a priori* value and uncertainty, which reveals the limitations of its utility.

Table 3. Tuning Results for PRP Albedo Reflectivity Coefficients.

Parameter	<i>A Priori</i> Value		<i>A Priori</i> Sigma	
	Baseline	Update	Baseline	Update
ALBCOF	1.0	2.0	0.5	1.8
ALMU01	0.07000	0.066	0.06083	0.010
ALMU02	0.07000	0.066	0.06083	0.010
ALMU03	0.07000	0.066	0.06083	0.010
ALMU04	0.07000	0.063	0.06083	0.013
ALMU05	0.07000	0.07	0.06083	0.02
ALMU06	0.07000	0.08	0.06083	0.03
ALMU07	0.2100	0.2100	0.0183	0.0004
ALMU08	0.2100	0.2100	0.0183	0.0004
ALMU09	0.07000	0.09	0.06083	0.04
ALMU10	0.07000	0.09	0.06083	0.04
ALNU01	0.26000	0.24	0.03924	0.04
ALNU02	0.26000	0.24	0.03924	0.04
ALNU03	0.26000	0.24	0.03924	0.04
ALNU04	0.26000	0.26	0.03924	0.04
ALNU05	0.26000	0.26	0.03924	0.04
ALNU06	0.26000	0.26	0.03924	0.04
ALNU07	0.0660	0.07	0.0222	0.03
ALNU08	0.0660	0.07	0.0222	0.03
ALNU09	0.26000	0.25	0.03924	0.04
ALNU10	0.26000	0.25	0.03924	0.04

The most substantial changes to the albedo reflectivity coefficient parameter *a priori* values were produced in the specular components of the solar panel back sides (ALMU09-10), which saw an approximate 30% increase over baseline. Otherwise, these changes were all about 15% or less. The *a priori* uncertainty updates ranged from an approximate 35% increase associated with the diffuse components of the solar panel front sides to a nearly two order of magnitude decrease in the uncertainty of their specular components.

2. Infrared Modeling

Planetary infrared radiation pressure perturbations on the spacecraft orbit are modeled in the filter with a 10×10 spherical harmonic expansion of planetary emissivity. The emissivities are computed from a Mercury surface

temperature profile, which is dependent upon the orbital ground track as well as the heliocentric true anomaly. At Mercury perihelion and aphelion, the inertially fixed MESSENGER orbit approximates a dawn-dusk ground track about the planet, which avoids the most extreme subsurface temperatures. There are two so-called hot poles located along the Mercury equator where the maximum surface temperatures are found. These are the planetary subsolar points at successive perihelion passages and are separated by 180° because of the 3:2 spin-orbit resonance of Mercury with respect to the Sun. The spacecraft traverses the hot poles three times each Mercury year, during or near the two eclipse seasons and around aphelion. This configuration repeats itself because the MESSENGER orbit is stable relative to inertial space while Mercury is revolving about the Sun every 87.97 days and rotating about its axis every 58.65 days. The two eclipse seasons that are encountered each Mercury year differ in character depending upon whether the periapsis point of the spacecraft orbit is over the sunlit or night side of the planet.

It is during eclipse seasons when the infrared model most often encounters difficulties. This issue is in large part a result of the problematic nature of characterizing Mercury surface emissivities. During non-eclipse intervals, any emissivity mismodeling is absorbed within the filter by the albedo and SRP parameters, along with the optional infrared scale factor (IFRCOR), which is normally employed to compensate for the tendency of the infrared reflectivity coefficients to drift outside of allowable limits from time to time during iterative corrections. However, during eclipses albedo and SRP perturbations are absent and infrared effects must be modeled more accurately to prevent undesirable behavior of both the estimation parameters and the tracking data residuals. This particular relative geometry between the spacecraft, Mercury, and the Sun tends to generate more troublesome conditions within the filter than others, necessitating the imposition of additional constraints by tightening the *a priori* uncertainties from iteration to iteration in order to assure physically realistic estimates.

One obvious physical effect that isn't directly modeled in the state estimation process is the emission of accumulated internally generated thermal energy by the spacecraft as heat. During eclipse intervals the spacecraft is operating at a very different ambient temperature than usual and naturally attempts to equalize its internal thermal condition with respect to the external environment, and since this condition holds only over the small fraction of the orbit experiencing eclipse, a steep temperature gradient is produced repeatedly over a relatively short timespan. Heat generated by the batteries, actuators, and instruments also produces spacecraft infrared radiation fairly continuously during the normal course of events, and even though this source of heat is mostly innocuous the accumulation of such minor effects are magnified during eclipses. Deleting ill-conditioned residuals and deweighting tracking data segments can often moderate the byproducts of these difficult relative geometries, but additional constraints on specific SRP and PRP estimation parameters are still sometimes necessary to achieve acceptable solutions when dealing with these types of modeling issues.

The statistical results used for tuning the infrared parameters are presented in Table 4, where the IRMU prefix is associated with the specular reflectivity coefficients and IRNU with the diffuse. More and larger changes are apparent within the infrared estimation parameters than seen previously in the analysis results for albedo modeling. This outcome is to be expected given that the starting point for both were the values inherited from the SRP coefficients, and solar radiation peaks in the visible wavelength range. Thus the material properties being modeled by the various plates representing the MESSENGER spacecraft are generally more applicable to the albedo model than they are to the infrared. Even so, no extreme adjustments were revealed and those lesser ones that were of a relatively large extent were somewhat expected.

Once again the scale factor parameter demonstrated the greatest variability, with the infrared IFRCOF results mirroring the ALBCOF change in *a priori* value although displaying a reduced jump in *a priori* uncertainty and resultant overlapping 1- σ error bars. As to the individual reflectivity coefficients, the specular components for the back sides of the solar panels (IRMU09-10) and the bottom of the spacecraft bus (IRMU06) grew larger by approximately 90% and 40%, respectively. The specular coefficients representing the top deck and exposed sides of the bus (IRMU04-05) had their *a priori* values increased by just less than 15%, and all other changes to these infrared parameter baseline values were less than 10%. The *a priori* uncertainties were generally reduced, some quite a bit, by up to an order of magnitude in the case of the specular coefficients for the solar panel faces and by almost 80% for the corresponding diffuse components. This result shows an encouraging trend toward improved estimation accuracy for the infrared model, with the notable exception of the reflectivities associated with the solar array back sides, particularly the specular components, for which the 1- σ uncertainties increased by over 30%. The bus plates all had their *a priori* uncertainties reduced, as did the sunshade components, for both specular and diffuse reflection modes. These results were promising and, as subsequent testing revealed, served to mitigate some of the experiential issues associated with the PRP infrared model.

Table 4. Tuning Results for PRP Infrared Reflectivity Coefficients.

Parameter	<i>A Priori</i> Value		<i>A Priori</i> Sigma	
	Baseline	Update	Baseline	Update
IFRCOF	1.0	2.0	0.5	1.1
IRMU01	0.07000	0.065	0.06083	0.017
IRMU02	0.07000	0.065	0.06083	0.017
IRMU03	0.07000	0.065	0.06083	0.017
IRMU04	0.07000	0.08	0.06083	0.04
IRMU05	0.07000	0.08	0.06083	0.04
IRMU06	0.07000	0.10	0.06083	0.06
IRMU07	0.2100	0.2102	0.0183	0.0018
IRMU08	0.2100	0.2102	0.0183	0.0018
IRMU09	0.07000	0.13	0.06083	0.08
IRMU10	0.07000	0.13	0.06083	0.08
IRNU01	0.26000	0.24	0.03924	0.03
IRNU02	0.26000	0.24	0.03924	0.03
IRNU03	0.26000	0.24	0.03924	0.03
IRNU04	0.26000	0.26	0.03924	0.02
IRNU05	0.26000	0.260	0.03924	0.015
IRNU06	0.26000	0.28	0.03924	0.04
IRNU07	0.0660	0.066	0.0222	0.005
IRNU08	0.0660	0.066	0.0222	0.005
IRNU09	0.26000	0.27	0.03924	0.04
IRNU10	0.26000	0.27	0.03924	0.04

III. Spacecraft Translational Position and Velocity Tuning

The MESSENGER spacecraft position and velocity nominal and *a priori* values for each new weekly on-orbit filter instantiation are typically inherited from the most recently delivered operational OD solution. Ephemeris deliveries to the project are generated so as to produce an hour of overlap between fit spans, which serves to define the new state epoch from week to week. The translational ephemeris is interpolated near the end of the previous fit span to extract a position and velocity, which is input into the filter as MCI Cartesian coordinates. Associated *a priori* uncertainties are generally fixed from one week to the next, although these can be modified at the discretion of the analysts conducting the estimation runs. This is usually based upon the location of the state epoch with respect to the orbital true anomaly and the relative geometry between the spacecraft, Mercury, Earth, and the Sun. Position and velocity uncertainties are periodic functions of the spacecraft orbit true anomaly, being smaller at the slower relative speeds associated with apoapses. The variability of solutions was analyzed statistically using the same methodology as that employed for the radiation pressure parameters, as there is a standard data cutoff time, and the results were used as a guideline to update the associated uncertainties. In this way the statistical analysis of accumulated OD solutions facilitates tuning of the *a priori* uncertainties for position and velocity, as well as other dynamic estimation parameters such as propulsive maneuvers and momentum dumps.

Cartesian components are not well suited for providing intuitive insights into the underlying physical reality influencing the variations in the translational dynamics. Transforming the inertial position and velocity into an orbital reference frame helps considerably in deciphering what is actually going on dynamically from solution to solution in the state estimation filter. If instead of using the traditional Cartesian set, components representing radial, along-track, and cross-track elements of both the position and velocity are analyzed, then the differences between solutions are generalized in a way that makes them more understandable.

The radial direction is represented by the vector from the principal focal point of the orbit, in this case Mercury's center, to the spacecraft. The cross-track direction is parallel to the orbit normal, defined by the cross product of the

translational position and velocity vectors, in that order. The along-track direction is in the plane of the orbit, defined similarly by the outer vector product of the cross-track and radial vectors, making it coincident with the velocity direction at the apsides and completing the right-handed orthogonal triad. This geometry is why radial, transverse, and normal (RTN) are often used as shorthand for the radial, along-track, and cross-track component system, which makes it easier to perceive the underlying dynamics when plotting up solution histories for the purposes of statistical analysis.

As an example of such a plot, Fig. 5 illustrates the same type of analysis methodology presented previously applied to the transverse, along-track component of spacecraft position. Depending upon the conservatism desired in the *a priori* uncertainties, the statistics suggest that the updated uncertainty in the along-track component should be somewhere between approximately 200 and 600 m. This range reflects the fact that epochs for the estimation state vectors are generally chosen near apoapses rather than as a random function of the orbital true anomaly.

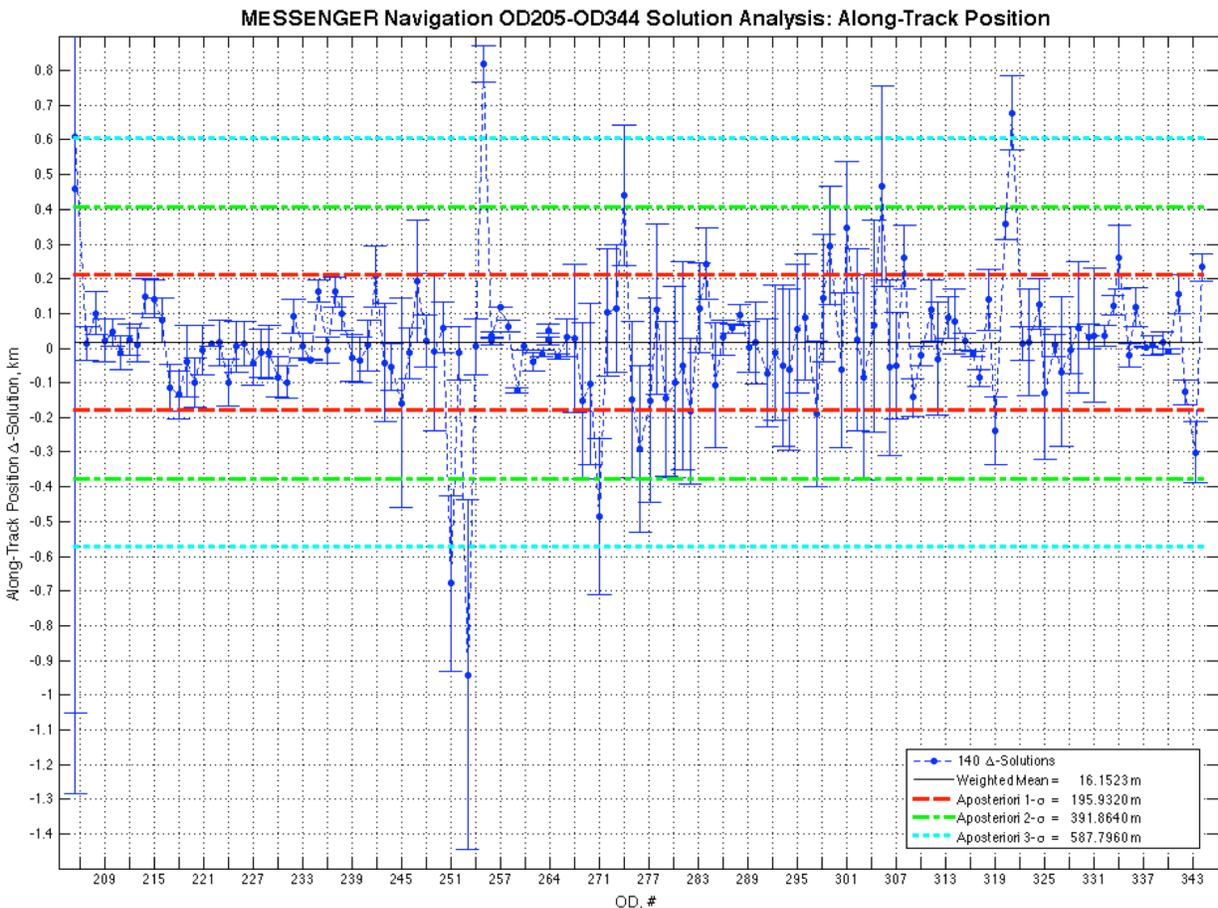


Figure 5. Estimation Solution History for Translational Position in the Along-Track Direction.

During the nearly three years of on-orbit solutions included in the analysis ensemble, the filter *a priori* position component uncertainties generally ranged from 3 to 10 km, although they have on occasion been reduced to 1 km or less. Conservatively large uncertainties may be appropriate for certain relative geometries and events, such as immediately subsequent to a superior solar conjunction or OCM. However, in general these results demonstrate a systematic mismatch between the variability of the position estimates, transformed into the along-track direction, and their associated *a priori* uncertainties. Of course, under-constraining the position and velocity parameters by a moderate amount is not normally problematic, as it would in fact be if they were over-constrained, which is why these uncertainties tend to be conservative. Regardless, tightening *a priori* uncertainties in a reasoned manner may inhibit anomalous variations in the solutions and reorders the priority of parameter adjustments within the state vector.

The *a priori* uncertainties for the position and velocity components represent expected 1- σ deviations from the *a priori* values associated with a previous filter solution and are not directly related to the absolute accuracy of the estimates. The weighted mean result of just over 16 m displayed in the legend of Fig. 5 is not a measure of the along-track solution error, but rather the correction relative to the previously delivered estimate. Analysis of the statistics of accumulated solutions for the radial position component dictated a similar downward adjustment in *a priori* uncertainties to 50-100 m, and an adjustment for cross-track position uncertainties to 600-1700 m. This change reveals that the variability in position estimates from solution to solution is less than what was predicted post-MOI.

The statistical analysis of velocity *a priori* uncertainties indicated that relatively smaller adjustments were warranted than those for position. This analysis involved reducing the range of these uncertainties, from 1-50 mm/s in the radial direction to 10-30 mm/s, and from 5-20 mm/s to 5-10 mm/s in both the along-track and cross-track directions. Therefore the initial post-MOI velocity solution uncertainty predictions were quite good. Even so, tightening these 1- σ values has the potential to improve the quality of state vector solutions. The RTN position and velocity uncertainties are transformed back to Cartesian components for input into the filter, which intrinsically incorporates dependencies on the location of the epoch state as a function of the spacecraft orbital true anomaly at a particular time. Optimizing the position and velocity *a priori* uncertainties in this way permits other state parameters to be estimated more accurately on a consistent basis.

IV. Gravity Field Tuning

Official project estimates of the gravity field of Mercury have been determined by MESSENGER radio science team models (HgM001-005).^{9,10} In parallel, and informed by these efforts, unofficial estimates of the Mercury gravitational field have been determined by the navigation team as a byproduct of ongoing preliminary orbit reconstruction activities (MNG01-04). Tuning of these reconstructed gravity field models for use in navigation operations has generated updated nominal and *a priori* values and uncertainties through the continuing process of trajectory solution refinement. The four iterations of Mercury navigation gravity models, MNG01-04, created to date for use in the production of mission ephemeris deliveries are independent of those produced for the project by the radio science team, HgM001-005. Collaboration between these two efforts has been very fruitful, and the results produced by both teams are generally quite consistent.⁴ Most recently, MNG04 was derived from the latest trajectory reconstruction utilizing track data through April 2012, which includes more than 13 months of orbital data since MOI. The updated nominal and *a priori* values from this analysis have concurrently been made available to the navigation team in time for filter tuning updates for the final spacecraft descent sequence to Mercury impact.

Before MNG04 is incorporated into the navigation operations filter setup, extensive testing will be performed to verify improvement in estimation fit metrics and associated results. The changes to *a priori* values and uncertainties incorporated within gravity field updates are generally subtle, and the effects they produce when folded into the filter inputs nuanced. Reconstruction solutions typically use different input specifications than operational estimates, which purposely employ sundry setup variations themselves on a regular basis. As an illustration of the differences between navigational gravity models, Fig. 6 presents a side-by-side comparison of the most recent reconstruction products, MNG03 and MNG04, showing the contours of gravitational potential at the Mercury surface. It is evident from these plots that MNG04 is the more detailed model and is thus expected to provide additional fidelity to the state estimation process when it is certified for navigation operations in the near future.

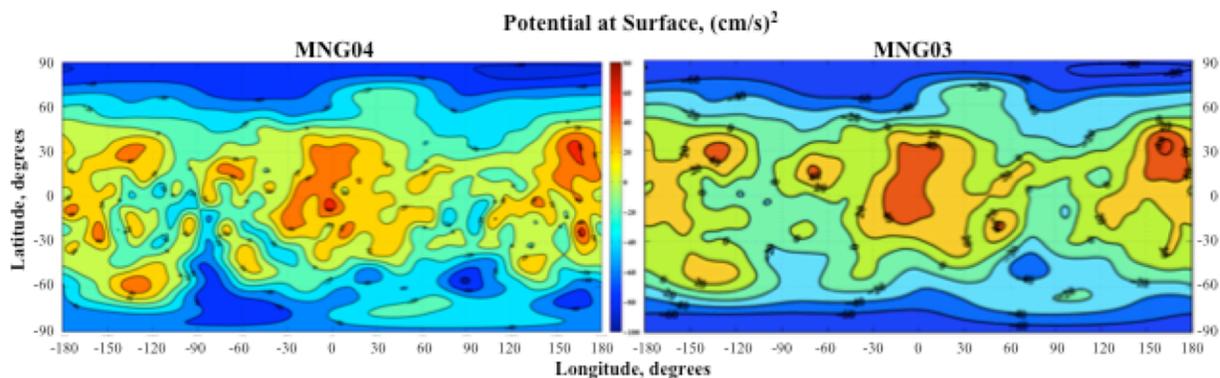


Figure 6. MESSENGER Navigation Team Reconstruction of Mercury's Gravitational Potential.

The products incorporated into the state estimation filter from these Mercury gravity field updates are the nominal and *a priori* values and uncertainties for the 437 coefficients of the 20×20 spherical harmonic gravity model, along with the planetary GM. The MNG04 solution is not very different from the navigation gravity deliveries that preceded it, even though it incorporates more data with complete and fairly uniform longitudinal coverage corresponding to seven full inertial rotations of Mercury plus the results for the three flybys preceding MOI. The *a priori* filter information was modified for this latest reconstruction to place more emphasis on the earlier flyby determinations of Mercury GM, since the flybys were more sensitive to GM than was the case for the orbit phase.

During the preliminary design and verification phase of the controlled descent maneuver campaign to Mercury impact, a series of tests was conducted with the express purpose of comparing an earlier baseline navigation operations gravity model, MNG02, with the latest radio science team gravity model, HgM005. The mission design team uses HgM005 for the long-term ephemeris propagation that supports the designs of OCM-9 to OCM-12, the final sequence of periapsis-raising maneuvers of XM2. The test results confirmed that there is no substantial difference between these gravity models in terms of the cumulative effect of gravitational accelerations and associated trajectory propagations, and that they both produce equivalent maneuver outcomes and impact predictions.

The final OCMs of the mission have been verified by comparing results achieved with the same maneuver designs in independent ephemeris propagators utilized by the navigation operations and mission design teams. The convergence of these test results provides a high level of confidence that the maneuvers will execute as planned. OCM-9 was successfully executed on 17 June 2014 and raised MESSENGER's periapsis altitude from 115 to 155 km. This maneuver established an orbital geometry with a relatively stable minimum altitude profile over the Mercury surface for an interval of a few weeks. Subsequent OCMs will likewise be performed at apoapses, but when spacecraft periapsis altitudes have drifted down to 15–25 km. Tuning the state estimation filter in preparation for these mission-critical events is a high priority for navigation operations, and Mercury gravity field updates are a major part of this process.

Since the errors on Mercury gravitational field coefficients remain a contributor to force modeling uncertainties, well-constrained and relatively small changes to these parameters are estimated every week and integrated into the operational ephemeris deliveries. The current baseline MNG model is the starting point for each of these solutions. Corrections are made to the 438 gravitational parameters by the filter, along with all of the other solved-for parameters of the state vector, using a nominal one-week fit span. A five-week propagation span is appended to the estimation span, and the resultant trajectory is delivered for mission operations and science planning activities. The gravity solutions are constrained to within statistical expectations of the MNG coefficient values by the *a priori* uncertainties, and each subsequent week the process is repeated in a similar manner. As periapsis altitude decreases during XM2, the terms of higher degree and order in the spherical harmonic gravity field model become increasingly important to the accuracy of the operational estimates.

There had been earlier discussion about whether it would be necessary to extend the 20×20 MNG modeling to a larger degree and order, but the insignificant differences resulting from the long-term ephemeris propagation comparison with the 50×50 HgM005 model have largely dispelled the notion that this level of effort is required for the XM2 descent sequence. However, as the spacecraft orbital periapsis approaches the planet surface, the tracking data progressively contain additional corrections to the higher-order terms of the gravity field that were not previously available. Whereas before the information content of the HgM 50×50 and MNG 20×20 gravitational fields were approximately equivalent, this situation will inevitably change during the descent to Mercury impact.

The state estimation gravity tuning process demonstrates that there are a variety of complementary mechanisms available for accomplishing the same goals. Whereas the methodology of a statistical analysis of accumulated on-orbit solutions could also be applied to the filter *a priori* values and uncertainties associated with the baseline MNG model, it would be a Herculean task given the number of parameters to scrutinize and difficult to justify given current resource limitations and other priorities at this final juncture of the MESSENGER mission. Time on task is an important metric for maintaining the efficiency of a streamlined navigation operations team, and the battles to be fought must be chosen thoughtfully to maximize output utility while delivering the highest quality products possible. Filter tuning is a continuous process of improvement during any operational space mission, and it must be approached with the respect it deserves in order to learn the appropriate lessons that the experience teaches.

V. Future Work

In addition to the testing and implementation of MNG04 operationally, there are several other estimation parameter sets that still await rigorous study for filter tuning purposes. The Mercury planetary ephemeris solutions

can be analyzed statistically in much the same manner as that used for the spacecraft translational position and velocity, and the associated *a priori* uncertainties updated accordingly. Position and velocity estimates for Mercury are represented in the filter by canonical element inputs, and these inputs are updated weekly with each new OD delivery. Statistical characterization of these solution histories would provide useful information, which could then be used to define appropriate $1\text{-}\sigma$ bounds for *a priori* uncertainty updates. Improving the Mercury ephemeris estimates will become increasingly important as the spacecraft periapsis continues to descend towards close proximity to the planetary surface and eventual impact.

OCM and CMD execution parameters are not amenable to persistent *a priori* value updates, since these are defined externally on the basis of the design for their planned execution, which includes the thruster sets to be utilized. The CMDs are accomplished with unbalanced attitude control thrusters, and small changes in the spacecraft translational velocity on the order of several mm/s are typically induced. A fixed *a priori* uncertainty of 0.5 mm/s has been used for all of the on-orbit CMDs to date, but this number could possibly be reduced through a statistical analysis of the approximately 150 momentum dump solutions accumulated from operational estimates.

Propulsive events such as the weekly CMDs and upcoming OCMs are typically estimated only once or twice before they fall outside the fit span of the filter. OCM estimates could also be analyzed, similarly to the CMDs, to fine-tune the associated *a priori* uncertainties. However, this procedure would be a more difficult task since there have only been nine OCMs executed since MOI and the associated thruster sets often varied from maneuver to maneuver. Multiple burn segments within individual OCMs are estimated separately, and as there can be up to four or five segments involved, the sample size is not quite as small as it might seem at first. However, the problem of correlating analysis results for comparable OCMs still exists, and therefore making inferences from a statistically insignificant data set is extremely problematic, especially when predictions for the execution errors are already provided and the process for defining related OD knowledge uncertainties is well established.

Several challenges associated with tuning the state estimation filter in the face of modeling deficiencies still need to be addressed. Whether this is due to unmodeled forces such as thermal emissions from the spacecraft discussed previously, or the intermittent mismodeling of planetary infrared re-radiation that affects the reflectivity coefficients during eclipse geometries, various compensation and correction mechanisms are continually being explored. The inclusion of stochastic accelerations as part of the estimation process is one possibility. This option would effectively transition the filtering mode from batch to sequential processing, which is rather a substantial change to undertake at this stage of the mission, but experimentation with this idea is being pursued. Another approach might be to include periodic accelerations during eclipse intervals, outgassing perturbations, or even small impulsive statistical maneuvers, although initial investigations into the latter have indicated that this direction may not be a successful line of pursuit.

Code modification of the modeling processes within the state estimation filter is generally a last resort during the operational phase of a space mission. However, experimentation with software changes can sometimes be done in parallel without disrupting the status quo if time and resources allow, and internal verification and validation of any such enhancements can provide certain assurances that the new models will improve the estimation process without inducing any unforeseen consequences. An alternative approach would be to merely modify the model inputs, for example increasing the degree and order of the spherical harmonics used for gravity, planetary infrared surface emissivity, or albedo models. These types of extensive efforts may be more than can be reasonably accomplished given the limited amount of time left in the MESSENGER mission, but the possibilities they represent must be kept in mind for the purposes of contingency planning.

VI. Conclusion

State estimation filter tuning is a continuous process that attempts to optimize the accuracy of OD solutions to the maximum extent possible. This objective is an implicit requirement for the navigation, guidance, and control of spacecraft during operational missions. There are a variety of methodologies for the refinement of filter inputs and outputs. Some are accomplished in real-time during the execution of OD runs by editing and weighting adjustments of input track data sets to tame the residuals and smooth iterative state corrections, whereas others are pursued offline with longer-term purposes. All such protocols need to be rigorously tested and proven to be effective in order to justify their repeated use. The best way to do this is to perform a number of state estimation runs in parallel on identical data sets using strictly controlled variations of filter inputs, which will then have identifiable effects in the solution differences.

Testing filter input variants is an ongoing activity, and with each new weekly MESSENGER OD instantiation two or more parallel runs are regularly conducted with what are usually marginally distinct setups, but which can incorporate important differences if desired. This procedure not only allows for a quick cross check on the solution

intended for operational delivery but also facilitates the implementation of enhancements into the estimation process through innovation and experimentation. Verification runs for proposed process improvements, such as the updated nominal and *a priori* values and uncertainties determined by the statistical analyses above, are conducted in just this way using previously archived operational, backup, and experimental OD solutions.

Test results for the latest series of verification runs are shown in Table 5, where case names listed with an appended x correspond to previously executed experimental or backup OD solutions. Each test case represents a pair of filter runs in which the only variant is the nominal and *a priori* (AP) value and uncertainty inputs. A comparison of quality of fit metrics among converged estimates provides confirmation of the systemic improvement in results. The chosen figure of merit was the sum of squares (SOS) of the tracking data residuals for the estimation span. The noise level of the SOS for defining the statistical significance of the fit to data improvement is less than 0.1%, and these test runs all demonstrate results consistently well above that threshold. This filter tuning methodology can easily be extended to include similar analyses for other state parameters including those associated with the Mercury planetary ephemeris, Mercury gravitational harmonic coefficients, and propulsive events such as OCMs and CMDs. The implementation of MNG04 harmonic coefficient updates will be verified using this same verification testing methodology.

Table 5. Test Results for Nominal and *A Priori* Values and Uncertainties Tuning. *Updates to radiation pressure nominal a priori values and uncertainties and position and velocity a priori uncertainties are included.*

Case	AP Baseline SOS	AP Update SOS	Fit to Data Improvement
OD340x	43.856	43.320	1.2%
OD341	63.790	60.974	4.1%
OD342x	126.89	122.48	3.5%
OD343	144.40	139.21	3.6%
OD344x	107.18	103.22	3.7%

As mentioned previously, not all state estimation filter tuning involves the adjustment of nominal and *a priori* values and uncertainties. There are a variety of other complementary mechanisms available for accomplishing the same goals. There are also different ways to determine appropriate adjustments to the *a priori* inputs other than the type of statistical analysis presented above, a prime example being the reconstruction of the spacecraft trajectory, which is also the responsibility of the navigation operations team. The trajectory reconstruction is an improved estimate of the manifold of previous estimates, representing the accumulated operational OD deliveries. This reconstruction is the sum total of acquired navigation knowledge from throughout the MESSENGER mission brought to bear on the history of state estimation solutions, and one byproduct of this process is the progressive refinement of the Mercury gravity model.

As the MESSENGER orbital mission at Mercury approaches its end, some four years after MOI, it is instructive to look back and reflect on what has made it such a great success. Much of this success had to do with the preparation and planning that went into the effort, even prior to the launch of the spacecraft in August 2004, as well as the care with which the most complex interplanetary cruise trajectory ever flown was executed. The nearly seven years it took to transit from Earth to MOI included six large deep-space maneuvers and six planetary flybys: of Earth once, Venus twice, and Mercury three times before planetary capture. The nominal one-year primary orbital mission has been extended twice, and the initial 12-hour orbital period was reduced to 8 hours after the first 15 months to provide additional opportunities for a wealth of new scientific measurements that can be performed near the surface. Now the final phase of the mission has begun as the spacecraft passed below the 200 km periapsis altitude floor in April 2014, followed by the controlled descent to eventual impact in late-March 2015.

The OCM descent campaign has been designed to forestall the inevitable Mercury impact as long as possible while providing a bountiful opportunity for scientific data acquisition in close proximity to the planetary surface. Several temporary islands of stability in the perturbation environment will be exploited by the planned OCMs, such that periapsis altitudes will remain relatively steady for a few weeks at a time on the way down. This mission plan should allow for instrument calibration and optimization of data collection at increasingly lower periapsis altitudes. In order to facilitate all of this work the navigation operations team will provide a flexible and tunable orbit estimation process. This process will be accomplished by the techniques presented here, although past procedures do not necessarily limit the future scope of associated potential activities. New challenges will be met with new

solutions, as they were after launch, and after MOI, and after the adjustment to an 8-hour Mercury orbit. Advancing into the unexplored territory of orbital periapsis altitudes as low as, and lower than, 15 km is the latest trailblazing experience for a pioneering spacecraft that has already gone where no robotic emissary has gone before. Mercury will soon provide a resting place for its first resident visitor from Earth as MESSENGER fulfills its destiny, to be joined once and for all with its purpose for being.

Acknowledgments

The MESSENGER mission is supported by the NASA Discovery Program Office under contracts to the MESSENGER Principal Investigator, Sean Solomon of Columbia University, and JHU/APL. This work was carried out by the Space Navigation and Flight Dynamics Practice of KinetX Aerospace, under a subcontract from the Carnegie Institution of Washington.

The authors specifically wish to thank James McAdams, MESSENGER Mission Design Lead Engineer, for his permission to use Fig. 2 in this paper, along with the rest of the personnel supporting the MESSENGER mission and science operations and planning at JHU/APL and elsewhere. Appreciation is also due to the NASA DSN facilities personnel who work so diligently to provide the tracking data upon which this navigation operations work is based.

References

- ¹Domingue, D.L., and Russell, C.T. (eds.), *The MESSENGER Mission to Mercury*, Springer, New York, 2007.
- ²Page, B.R., Williams, K.E., Bryan, C.G., Taylor, A.H., Stanbridge, D.R., Dunham, D.W., Wolff, P.J., Williams, B.G., O'Shaughnessy, D.J., and Flanagan, S.H., "Applying Experience from Mercury Encounters to MESSENGER's Mercury Orbital Mission," *Advances in the Astronautical Sciences*, 142, 2012, pp. 2251-2270.
- ³Page, B.R., Bryan, C.G., Williams, K.E., Taylor, A.H., Stanbridge, D.R., Wolff, P.J., and Williams, B.G., "MESSENGER Navigation Operations during the Mercury Orbital Mission Phase," *Advances in the Astronautical Sciences*, 148 (Part III), 2013, pp. 2825-2844.
- ⁴McAdams, J.V., Bryan, C.G., Moessner, D.P., Page, B.R., Stanbridge, D.R., and Williams, K.E., "Orbit Design and Navigation Through the End of MESSENGER's Extended Mission at Mercury," *Advances in the Astronautical Sciences* (to be published).
- ⁵O'Shaughnessy, D.J., McAdams, J.V., Bedini, P.D., Calloway, A.B., Williams, K.E., and Page, B.R., "MESSENGER's Use of Solar Sailing for Cost and Risk Reduction," *Acta Astronautica*, 93, 2014, pp. 483-489.
- ⁶O'Shaughnessy, D.J., McAdams, J.V., Williams, K.E., and Page, B.R., "Fire Sail: MESSENGER's Use of Solar Radiation Pressure for Accurate Mercury Flybys," *Advances in the Astronautical Sciences*, 133, 2009, pp. 61-76.
- ⁷Scott, C.J., McAdams, J.V., Moessner, D.P., and Ercol C.J., "Modeling the Effects of Albedo and Radiation Pressures on the MESSENGER Spacecraft," *Advances in the Astronautical Sciences*, 142, 2012, pp. 2309-2324.
- ⁸Vasavada, A.R., Paige, D.A., and Wood, S.E., "Near-Surface Temperatures on Mercury and the Moon and the Stability of Polar Ice Deposits," *Icarus*, 141, 1999, pp. 179-193.
- ⁹Smith, D.E., Zuber, M.T., Phillips, R.J., Solomon, S.C., Hauck II, S.A., Lemoine, F.G., Mazarico, E., Neumann, G.A., Peale, S.J., Margot, J.L., Johnson, C.L., Torrence, M.H., Perry, M.E., Rowlands, D.D., Goossens, S., Head, J.W., and Taylor, A.H., "Gravity Field and Internal Structure of Mercury from MESSENGER," *Science*, 336 (3078), 2012, pp. 214-217.
- ¹⁰Mazarico, E., Genova, A., Goossens S., Lemoine, F.G., Neumann, G.A., Zuber, M.T., Smith, D.E., and Solomon, S.C., "The Gravity Field, Orientation, and Ephemeris of Mercury from MESSENGER after Three Years in Orbit," *Journal of Geophysical Research* (to be published).