

The NASA Phoenix 2007 Mars Lander Thruster Calibration Estimator: Design and Validation

Michael E. Lisano^{*}, Gerhard L. Kruizinga[†], and Brian Portock[‡]
*Jet Propulsion Laboratory, California Institute of Technology
Pasadena, CA 91109*

The NASA Phoenix 2007 Mars Lander mission, launched in August 2007 on its mission to land near the north pole of Mars in May 2008, had a driving need for entry-corridor delivery precision, which parlayed into stringent requirements on deep space navigation accuracy. This, in turn, necessitated in-cruise calibration of the three-axis thrust force vectors produced by each of the vehicle’s four reaction-control system (RCS) thrusters during frequent daily low-catalyst-bed-temperature firings done to maintain the 3-axis attitude deadbands. A novel recursive sigma-point consider-covariance filter was designed, validated and ultimately utilized extensively during flight operations, to estimate the RCS force vectors, per individual thruster. The estimate was achieved through ground-based processing of Deep Space Network (DSN) and telemetered gyroscope data from the spacecraft’s inertial measurement unit (IMU), using a novel sigma-point consider filter (SPCF) formulation. During early-cruise active calibration, the spacecraft was flown in attitudes chosen, using this filter, to maximize observability of all thruster axes, to an extent constrained by vehicle thermal and communication considerations. The design of the Phoenix thruster calibration filter, and its validation through processing of archived Mars Odyssey thruster calibration radiometric data, and simulated sets of data, are discussed in this paper. The paper concludes with the formulation of the thruster calibration campaign and a summary of the thruster calibration campaign results. The SPCF algorithm is summarized in the Appendix.

I. Introduction

On August 4, 2007, NASA launched its Phoenix 2007 Mars Lander spacecraft, commencing a nearly ten-month interplanetary cruise to Mars, to enter the Martian atmosphere, descend and land near the Martian north pole. Targeting the landing site was accomplished by means of ground-based radiometric orbit determination along with ground-based calculation and commanding of trajectory-correction-maneuvers (TCM’s) at intervals along the cruise trajectory. The pre-launch requirement for delivery accuracy at the top of the Martian atmosphere was quantified as 0.28 degrees, $3\text{-}\sigma$, in total entry flight path angle delivery error. The pre-launch error budget for the entry flight path angle accuracy requirement is shown in Table 1.

Table 1. Phoenix Entry Flight Path Angle Required Performance Error Budget (Pre-Launch)

Error Source	Error ($3\text{-}\sigma$) in Entry Flight Path Angle (deg)
Reconstructed Small Force Errors	0.19
Predicted Small Force Errors	0.05
TCM-6 Slew Errors	0.13
Orbit Determination Errors	0.12
TCM-6 Burn Execution Errors	0.09
Total Entry Flight Path Error (RSS, $3\text{-}\sigma$)	0.28

^{*} Senior Staff Engineer, Guidance, Navigation and Control. Email michael.e.lisano@jpl.nasa.gov. Senior Member, AIAA.

[†] Senior Staff Engineer, Guidance, Navigation and Control.

[‡] Navigation Team Chief, Phoenix 2007 Mars Lander.

The largest contributor of entry flight path angle delivery error was the forecasted uncertainty on the reconstructed small forces, i.e. estimates of forces from frequent, automated, cold firings of the reaction control thrusters during attitude-control deadbanding, that act continually on the spacecraft through interplanetary cruise. The reconstructed small force values are used to model velocity change and thruster pulse times in the orbit determination software. Other major contributors to the entry flight path angle error are TCM slewing and burn execution errors, as well as orbit determination errors (which include radiometric measurement errors, solar radiation pressure model errors, and errors due to unmodeled effects such as vehicle outgassing). From this error budget arose the need to precisely calibrate the small forces while the vehicle was in flight to Mars.

In-flight thruster calibration campaigns have been performed for recent prior deep space missions, including the 2001 Mars Odyssey, 2003 Mars Exploration Rover, 2005 Mars Reconnaissance Orbiter (MRO), and Stardust comet sample return missions.¹⁻⁵ The calibration campaigns have differed in methodology from one mission to the next, due to differences in the spacecraft configurations.

For the thruster calibrations of Odyssey and MRO, which were three-axis stabilized using reaction wheels, the spacecraft was maintained in a series of attitudes which exposed x, y, and z thrust axes to direct observation from Earth by Deep Space Network (DSN) Doppler measurements, as the thrusters were fired numerous times in a manner that emulated the vehicle's frequent (hundreds/day) momentum desaturation impulses. For the Mars Exploration Rovers, which were spinning spacecraft, the attitude control system consisted of balanced thrusters, the spacecraft was similarly turned to attitudes during a brief maneuver campaign, to allow any residual thrust resulting from imbalance of the thrusters to be observed via DSN Doppler tracking data from Earth. The Stardust spacecraft, which had no reaction wheels and was three-axis stabilized with autonomous deadbanding pulses from its RCS thruster clusters, also used a DSN Doppler-measurement-only approach to observe the velocity change imparted to the spacecraft due to pitch or roll attitude deadband maneuvers, and also to measure the additional velocity change contributed by the RCS during Stardust TCM's. The small forces calibration methodology for Phoenix was different from all of these earlier campaigns, due in large part to the design of the spacecraft attitude control system.

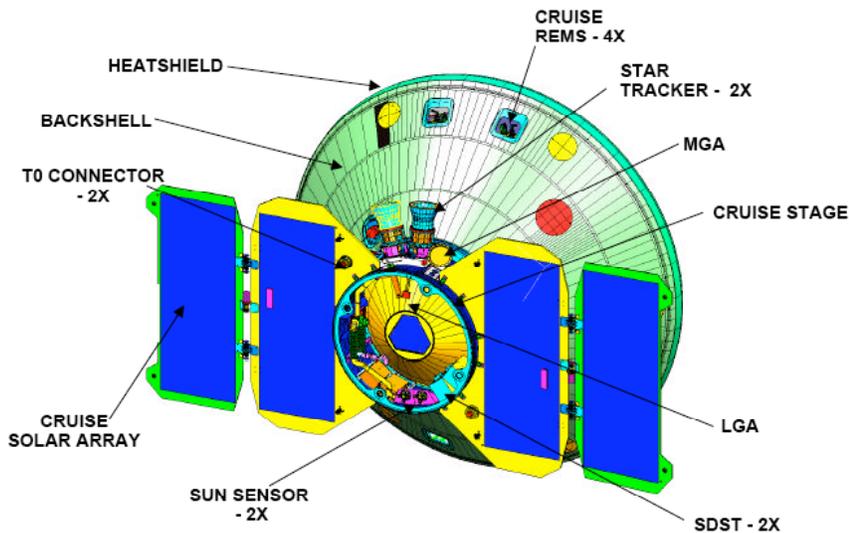
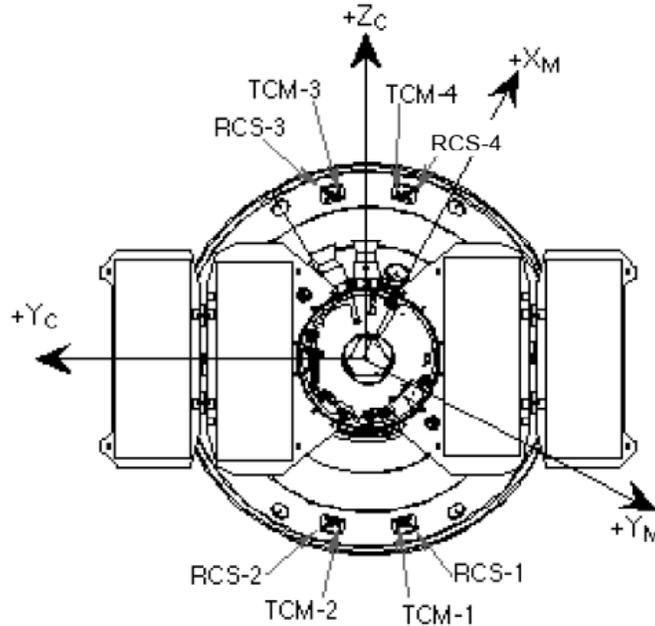


Figure 1: Phoenix Spacecraft Interplanetary Cruise Configuration.
Figure courtesy of Lockheed Martin.

As depicted in Figure 1, the Phoenix spacecraft consists of a lander, contained within an aerodynamic backshell and heatshield used during entry, descent and landing (EDL), and a cruise stage featuring solar panels, a medium-gain communications-and-tracking antenna, and a pair of star cameras. All TCM and

RCS thrusters are located on the lander, and are oriented out of the backshell; that is, the cruise stage contains no elements of the cruise propulsion subsystem.

Figure 2 depicts the definition of the Phoenix spacecraft body frame, also called the cruise frame. The figure also shows locations and orientations of the four rocket engine modules (REMs). Each REM includes both TCM thruster and RCS thruster, which have nozzles that are scarfed to align with the backshell outer wall.



**Figure 2: Phoenix RCS and TCM Thruster Locations on Spacecraft (“c” subscript denotes spacecraft cruise frame, “m” subscript denotes mechanical frame).
Figure courtesy of Lockheed Martin.**

For Phoenix, due to the pointing of the RCS thrusters (primarily along the +/- cruise-frame Y-axis, with components in the +/- Z-axis and components only in the X-axis), reaction-control firings can be fired in a coupled manner in the cruise-frame Y or Z direction, but because of the orientation of the thrusters, all RCS firings result in a net positive-X translational acceleration in the cruise frame. Mis-modeling and uncertainty of RCS thrust levels therefore has a direct impact on knowledge and predictions of the spacecraft translational state, i.e. orbit determination, and results in the prominence of the reconstructed small force errors in the error budget of Table 1.

While characterization of the reaction control thrusters in ground-based laboratories was done for Phoenix, predictions of in-space thrust based on terrestrial measurements is subject to significant uncertainty due to the intractability of exactly replicating thruster firing conditions in the near-perfect vacuum of space during terrestrial thruster firings. Given this, and the importance of reaction-control force modeling error on the Phoenix orbit determination accuracy and entry-corridor targeting accuracy, a comprehensive in-flight thruster calibration campaign was required, based on leveraging all available precision in-flight measurements.

II. Filter Formulation

For Phoenix, the ability to precisely determine the spacecraft’s heliocentric orbit is sensitive to accurately modeling and predicting the accumulated velocity changes caused by the RCS firings. However, the velocity changes caused by RCS thruster firings are not always observed with Deep Space Network precision Doppler radiometric tracking data. This is because the thrust vector is generally not aligned with the spacecraft line-of-sight to any of the Deep Space Network stations that may be tracking the spacecraft at the time of the thruster firing.

Therefore, it was necessary to explicitly calibrate and model the average three-dimensional thrust vector for each of the four Phoenix RCS thrusters during post-launch cruise. The RCS thrust levels were estimated by processing a batch of downlisted measurement data using a sequential-processing consider-covariance Kalman filter (called “TIM” for Thruster Inflight Measurement), the formulation and validation of which is described here. Two data types were processed simultaneously in TIM, to estimate the RCS thruster performance. One data type is attitude rate data telemetry from the Phoenix Inertial Measurement Unit (IMU) 3-axis gyroscope. The second data type is Deep Space Network (DSN) Doppler tracking residual data, formed by differencing raw Doppler measurements with a precision predicted spacecraft trajectory, produced using precision deep-space orbit determination software that does not include modeling of the thruster firing events. This is the first known attempt to estimate three-axis RCS cold-firing thrust vectors by combining DSN Doppler and spacecraft gyroscope data in a single filter.

The Phoenix TIM filter is a sequential-processing sigma-point consider filter (SPCF), outlined in the Appendix of this paper, and characterized in detail in Refs. 6 and 7. To apply the SPCF to state estimation (and parameter-error consideration) of a particular system, one must supply the following information:

- A function “h” which returns the modeling expressions for the measurements processed by the filter, as a function of the a-priori state
- A function “f” which returns the filter state vector as a function of input time, a-priori state and control. This usually involves a time-integration of the equations of motion of the system
- A priori state error covariance (P_0), process noise covariance matrix (Q), measurement noise covariance matrices (R), and information that partitions the state vector and state covariance into estimated, neglected, and considered sub-elements

No partial derivative information for “f” (i.e. state-transition expressions) or for “h” (i.e. a measurement partial derivative Jacobian with respect to the state vector) is required for the SPCF algorithm, as is the case for any sigma-point, or “unscented”, Kalman filter. Given the brief timeframe in which the TIM filter was to be implemented (approximately three months), this was an advantageous aspect of using the SPCF, in that there were fewer expressions to develop and debug as the filter was implemented.

As typically configured, the TIM filter estimates a thirteen-parameter state vector, while considering error uncertainty in ten other model parameters. The filter is not limited to this estimate-and-consider configuration, however, and a user can select to estimate, consider or neglect any parameter in the TIM parameter list.

The vector of thirteen states that are estimated for the Phoenix thruster calibration is:

$$X_{est} \equiv \begin{bmatrix} \bar{T}_1 \\ \bar{T}_2 \\ \bar{T}_3 \\ \bar{T}_4 \\ b_{dop} \end{bmatrix} \begin{array}{l} (3x1) \text{ Thruster 1 } x, y, z \text{ force components } (N) \\ (3x1) \text{ Thruster 2 } x, y, z \text{ force components } (N) \\ (3x1) \text{ Thruster 3 } x, y, z \text{ force components } (N) \\ (3x1) \text{ Thruster 4 } x, y, z \text{ force components } (N) \\ (\text{scalar}) \text{ Post - ACS -clampdown bias in Doppler observation } (m/s) \end{array} \quad (1)$$

The vector of ten parameters that are *considered*, but not estimated, for the Phoenix thruster calibration, is:

$$C \equiv \begin{bmatrix} I_{xx} \\ I_{yy} \\ I_{zz} \\ I_{xy} \\ I_{xz} \\ I_{yz} \\ M_{sc} \\ \bar{\mathbf{r}}_{cg} \end{bmatrix} \begin{array}{l} \text{(scalar) Spacecraft } x \text{ - axis moment of inertia in Cruise Frame (kg * m}^2\text{)} \\ \text{(scalar) Spacecraft } y \text{ - axis moment of inertia in Cruise Frame (kg * m}^2\text{)} \\ \text{(scalar) Spacecraft } z \text{ - axis moment of inertia in Cruise Frame (kg * m}^2\text{)} \\ \text{(scalar) Spacecraft } xy \text{ - plane product of inertia in Cruise Frame (kg * m}^2\text{)} \\ \text{(scalar) Spacecraft } xz \text{ - plane product of inertia in Cruise Frame (kg * m}^2\text{)} \\ \text{(scalar) Spacecraft } yz \text{ - plane product of inertia in Cruise Frame (kg * m}^2\text{)} \\ \text{(scalar) Spacecraft total mass (kg)} \\ \text{(3x1) Spacecraft center of mass coordinates in Cruise Frame (m)} \end{array} \quad (2)$$

Other parameters available in the TIM filter parameter list that may optionally be estimated, considered, or neglected, per analyst choice, are: spacecraft body rate vector, 3-axis IMU gyro measurement bias, thruster throat 3-axis location per thruster, and low-gain antenna phase center 3-axis location in spacecraft cruise frame. Typically, by default the effect of error in these other parameters on covariance is neglected.

There are two types of measurements modeled in the TIM filter, and processed simultaneously in the filter's state-update function. These measurements are the gyroscope data telemetry from the spacecraft IMU, and the residual of the DSN Doppler data (i.e the difference between the Doppler data and a predicted Doppler value based on the propagated spacecraft navigated state, as modeled sans thruster firings). As described below, the gyroscope data are used to infer thrust vector components via the propulsive torques acting on the spacecraft during RCS firings, and the Doppler residual data provide insight into thrust vector components via the velocity changes along the spacecraft-to-Earth line of sight.

As discussed in Section IV, neither of these measurement types would have been able, applied singly, to completely ascertain the average three-axis thrust components of each RCS thruster, when fired in the low-impulse, "cold-catalyst-bed" mode that was used during Phoenix cruise deadbanding. However, the combination of measurements enables complete observability for all RCS thrust components. This complete observability is strongly subject to the alignment of the Earth-spacecraft vector relative to the direction of thrust acting on the spacecraft.

The formulation of the TIM filter is straightforward, and the system dynamics model ("f") for TIM is particularly simple, given the assumption that the estimated system (average thrust vector component values over the calibration data time interval) is a stationary process. That is, the time derivative of the entire state and consider parameter vector is zero. The gyroscope and Doppler measurement models that contribute to the function "h" are discussed below.

Gyroscope data model as a function of average thrust vector components:

The TIM filter treats the change in measured spacecraft body rate across an RCS pulse interval in order to observe the average 3-axis thrust vector components for each of the thrust vectors. The gyro delta-rate measurement is modeled as an array of three scalars:

$$\Delta \bar{\omega}_{sc} = \mathbf{I}_{sc}^{-1} \left\{ \sum_{iThr=1}^4 \left(\bar{\mathbf{r}}_{iThr} \times \bar{\mathbf{T}}_{iThr} \right) \Delta t_{on_{iThr}} - (\bar{\omega}_{sc} \times \mathbf{I}_{sc} \bar{\omega}_{sc}) \right\} \quad (3)$$

where:

\mathbf{I}_{sc} is the Phoenix spacecraft inertia tensor about the center of mass

$\bar{\omega}_{sc}$ is the instantaneous spacecraft body rate, in the cruise frame, from IMU measurement telemetry

$\bar{\mathbf{T}}_{iThr}$ is the average thrust vector (cruise frame) of the i^{th} RCS thruster, from the state vector

$\bar{\mathbf{r}}_{iThr}$ is the cruise-frame, c.g.-relative location of the nozzle of the i^{th} RCS thruster

and

$\Delta t_{on_{iThr}}$ is the on-time of the i^{th} RCS thruster, from the spacecraft small-forces telemetry record.

The measurement model in equation (3) is used to form a residual with the difference between a smoothed pre-pulse and post-pulse body rates, which are the averages of ~ 1 -second-duration intervals of 200-Hz IMU gyro data immediately before and after each 15-millisecond RCS pulse.

Doppler-residual data model as a function of average thrust vector components:

The Doppler measurements processed in the TIM filter are actually *navigation Doppler residuals*, which are formed by differencing the two-way X-band Doppler measurement made by the Deep Space Network using the Phoenix spacecraft's SDST (Small Deep Space Transponder), with a prediction of what the Doppler measurement would be if there were no thruster firings occurring, based on propagating the best current navigated state of Phoenix from orbit determination. The TIM filter directly models this Doppler residual as a change in Doppler due to RCS thrust, in order to observe the 3-axis thrust vector for each of the thrust vectors.

The Doppler-residual measurement is modeled as a scalar quantity, and includes the instantaneous rotational velocity of the spacecraft low-gain antenna:

$$\Delta \dot{\rho} = \left\{ \mathbf{m}_{sc}^{-1} \sum_{iThr=1}^4 \bar{\mathbf{T}}_{iThr} \Delta t_{on_{iThr}} + (\bar{\boldsymbol{\omega}}_{sc} \times \bar{\mathbf{r}}_{LG}) \right\} \bullet \bar{\mathbf{r}}_{E/sc} + \mathbf{b}_{dop} \quad (4)$$

where:

\mathbf{m}_{sc} is the Phoenix spacecraft mass

$\bar{\mathbf{r}}_{LG}$ is the cruise-frame, c.g.-relative location of the low-gain antenna electronic phase center

$\bar{\mathbf{r}}_{E/sc}$ is the cruise-frame, c.g.-relative unit vector in the direction of the Earth

\mathbf{b}_{dop} is the post-ACS-clampdown Doppler bias, from the state vector

and the other quantities in Equation (4) are defined in the subsection above.

III. Validation, Part 1: Processing of Mars Odyssey Archived Thruster Calibration Data

The TIM filter was developed and validated over three months during the summer of 2007, around the time of Phoenix launch. The first part of the validation entailed reproducing, as closely as possible, the results of the 2001 Mars Odyssey thruster calibration, which was based on Doppler measurements only (possible because Odyssey has reaction wheels, and therefore was able to maintain its inertial direction very precisely while pulsing pairs of RCS thrusters).

Figure 3 is a plot of the Mars Odyssey (“ODY”) 2-way Doppler navigation residuals, formed by differencing Doppler measurements with a predicted, no-thrust-modeled, Doppler prediction. Again, these “residuals” are the measurements that are processed in the TIM filter. Over the approximately 7.5-hour active calibration data collection period during its Earth-to-Mars cruise (on May 4, 2001), the Odyssey spacecraft was flown in three different, approximately orthogonal attitudes relative to the spacecraft-Earth vector, while pulsing pairs of its four RCS thrusters, testing each of the six possible combinations three times, twenty times per pair at 15-second intervals between firing time. A “rest time” of 15 minutes was spent between each set of pair firings, in order to let the catalyst beds and thrusters cool to low temperatures as they would have during momentum-wheel desaturation firings. The spacecraft attitude was changed at approximately 10000 sec and 20000 sec after the start of the calibration campaign.

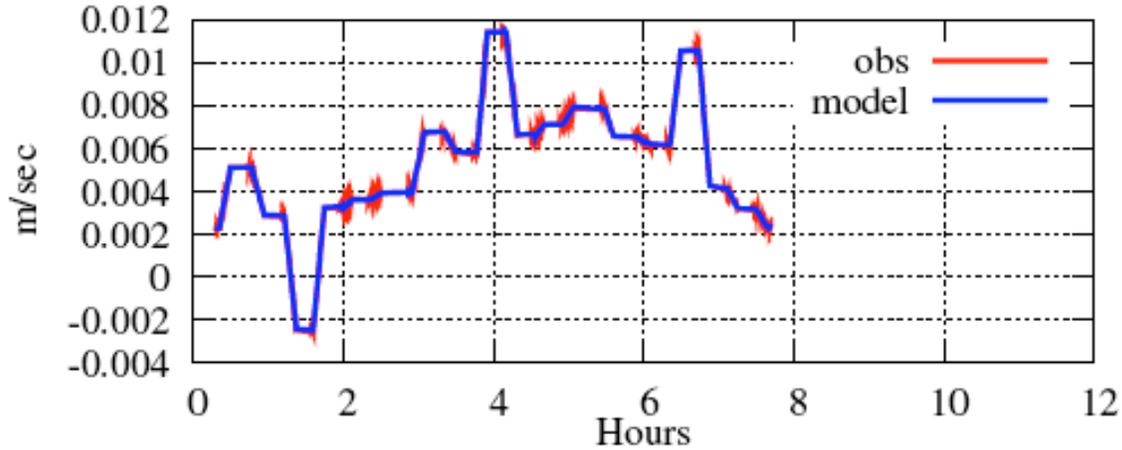


Figure 3. Two-way DSN Doppler, Differenced with Thrust-Free Navigation Prediction of Doppler, for Mars Odyssey Spacecraft. Measurements Made During Odyssey In-Cruise Thruster Calibration.

The TIM filter processed the Doppler data plotted in Figure 3, by initializing the state vector to zero a-priori, with very open covariances (several Newtons per thrust component). The data was processed sequentially through TIM, without any iteration over the batch of measurements. The TIM results were then compared with the results of the original Odyssey thruster calibration estimator (“tcal3”), a simple least-squares batch algorithm which also had zero a-priori state and was not iterated.

Table 2 shows the difference between the estimated per-pulse delta-V per thruster axis, and also the difference between the delta-V magnitude per pulse per thruster (the per-pulse delta-V per thruster was approximately 0.12 mm/s for Odyssey). The results of the two, completely independent filter algorithms (TIM being sequential and sigma-point, tcal3 being a batch least-squares method) match to within 1% or better per thruster in delta-V magnitude.

Table 2. Difference Between TIM Filter Odyssey Thruster Calibration Delta-V Result (2007), And Solution from Odyssey Thruster Calibration Batch Estimator “tcal3” Result (2001)

	DVx diff (mm/s)	DVy diff (mm/s)	DVz diff (mm/s)	DV diff
Thruster 1	0.0012	0.0004	-0.0001	-0.0012 (1% diff)
Thruster 2	0.0006	-0.0004	-0.001	-0.0002 (0.02% diff)
Thruster 3	0.0006	0.0009	0.0001	0.0002 (0.02% diff)
Thruster 4	0.0009	0.0003	-0.0009	0.0011 (1% diff)

Figure 4 is a plot of Mars Odyssey radiometric Doppler tracking post-filtering error residuals formed with with the Phoenix TIM filter (green points), and with the tcal3 estimator (red points). The Doppler noise level being rejected by the two estimation approaches is consistent. Figure 5 shows the TIM filter estimate convergence for the X component of Odyssey RCS thruster 1 (red), compared with the tcal3 solution, between plus or minus 10% bars (green) – similar results were seen for each axis of all thrusters.

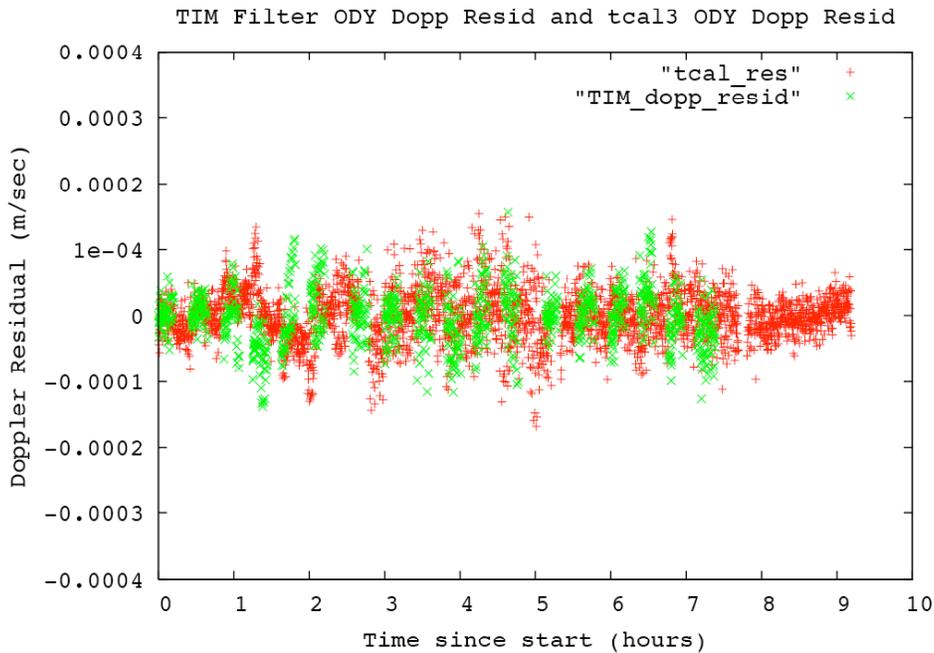


Figure 4. Comparison of Mars Odyssey Doppler Residuals from Odyssey Calibration Filter “tcal3” (red) and Phoenix Calibration Filter “TIM” (green)

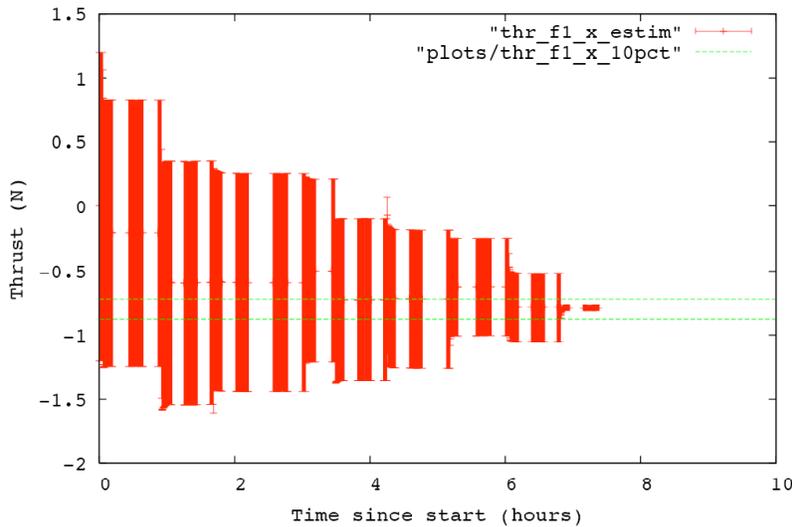


Figure 5. Estimated X-Component of Mars Odyssey Thruster Number 1 Thrust Vector, +/- 1-sigma from Covariance, Based on 7 Hours of Mars Odyssey Doppler Tracking Data, TIM Filter (red) versus Odyssey “tcal3” Filter Final Solution +/- 10% (green dashed lines).

The results described in this section successfully validated the performance of the TIM filter, for processing of actual in-flight DSN 2-way X-band Doppler measurements. The fact that this was done with archived Doppler measurements from the successful thruster calibration campaign of another, formerly-Mars-bound spacecraft, added a measure of confidence to the use of the TIM filter for Phoenix flight operations. Still, an assessment was needed of Phoenix-like processing with TIM, i.e. simultaneous Doppler and gyro rate data processing in the presence of Phoenix like attitude motions and constraints, which led to the validation step described in the next section.

IV. Validation, Part 2: Processing Simulated IMU and Doppler with Phoenix-like Dynamics

To validate the TIM filter accuracy for a rotating (i.e. deadbanding) spacecraft, and also to validate the algorithms for processing Doppler data along with gyro rate data, a combination of Doppler and gyro rate data were synthesized with a simulation of Phoenix spacecraft motion, and processed in the TIM filter.

The simulated 2-way Doppler “residual” measurements shown in Figure 6 (red “obs” curve), as well as the filter’s reconstruction of the signal (blue “model” curve, at this scale a perfect overlay of the observations curve) are for an attitude profile, described in the next section, that would eventually become the as-flown Active Thruster Calibration Campaign attitudes. The Doppler signal level ranged from -0.005 to 0.005 m/s, peak-to-peak, for the entire 7.5 hour calibration campaign. Compared to the Odyssey Doppler values in Figure 3, this is somewhat tighter, but comparable (range of 10 mm/s for Phoenix versus 14 mm/s for Odyssey).

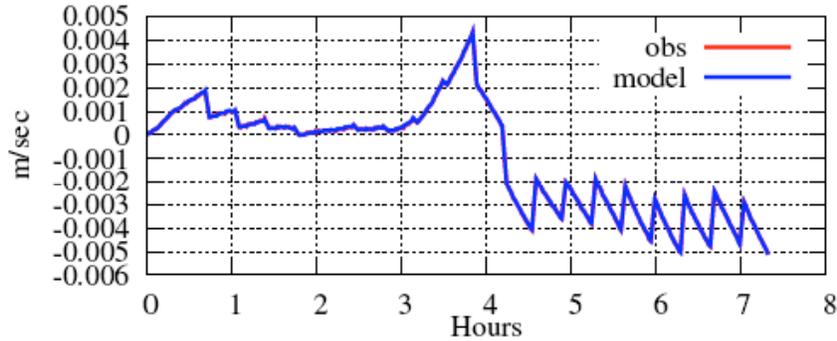


Figure 6: Simulation of Two-way DSN Doppler Differenced with Thrust-Free Navigation Prediction of Doppler, for Phoenix (Using Planned Attitude History for Active Thruster Calibration Campaign).

The TIM filter residuals for the Doppler observations, based on processing Doppler and IMU data simultaneously, are plotted versus time in Figure 7 (red curve), and are seen to be approximately two orders of magnitude smaller than the peak-to-peak signal in Figure 6, and also in good agreement with the filter’s prediction (blue curve) of the 1-sigma residual covariance. The “spikes” seen in the residual covariance curve are due to resetting the filter covariance for the Doppler bias at each time the modeled spacecraft performed an attitude-restoring deadbanding maneuver.

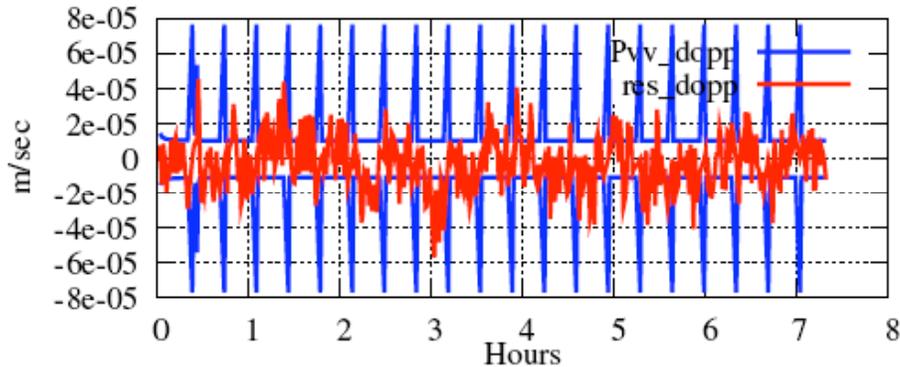


Figure 7. TIM filter residuals (red) for 2-way Doppler “residual” measurements, processed simultaneously with Gyro rate observations. Co-plotted with 1-sigma expected value of filter Doppler residual (blue), from filter residual covariance (standard part of Kalman gain computation).

Gyro rate measurement residuals are shown in Figure 8. The RMS residuals (root-sum-square over X, Y, and Z gyro axes) were on the order of 2×10^{-5} rad/sec, or about a tenth of the peak rates experienced during the modeled attitude deadbanding. As with the Doppler measurement residuals, the gyro rate residuals (red

curves) were in good agreement with the 1-sigma predicted residual covariance (blue curves) from the filter, indicating that the filter was appropriately tuned and also that the modeling in the filter was consistent with the attitude dynamics models used to generate the data.

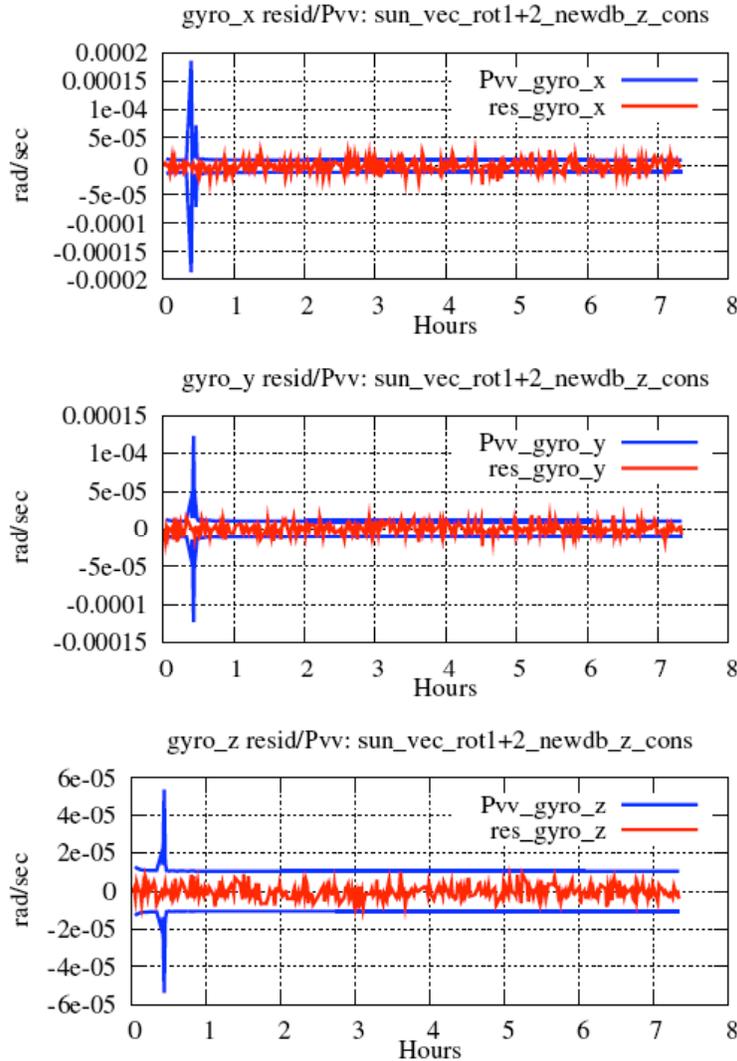


Figure 8. TIM filter residuals (red) for X, Y, and Z gyro attitude rate measurements, processed simultaneously with Doppler observations. Co-plotted with 1-sigma expected value of filter gyro residual (blue), from filter residual covariance (standard part of Kalman gain computation).

Figure 9 shows plots of the x, y, and z components of the thrust vector of RCS thruster 1, as estimated by the TIM filter based on Doppler and gyro rate data (with +/- 1-sigma formal uncertainty bars in red, from filter estimated covariance, and +/- 1-sigma consider-covariance uncertainty shown as green error bars). The TIM filter yielded an estimate of RCS thrust vector that converged to the true solution very quickly in X and Y axes, and that converged quickly in Z after the spacecraft rotated about the spacecraft-sun vector at time $t = 4$ hours. Very similar results were seen for the other three RCS thrusters. Note that the consider covariance – which is responding predominantly to the considered, 1-centimeter-level uncertainty in the location of the spacecraft c.g. – only has a significant effect on the uncertainty in the thrust vector z-axis. This is because the largest uncertainty in the spacecraft c.g. is in the cruise frame z direction – the axis along which the two propellant tanks are centered. Solution and covariance data shown in Table 3 indicated that the performance in thrust magnitude reconstruction was better than the 10% goal, with thrust vector direction accuracy better than 1 degree for each thruster.

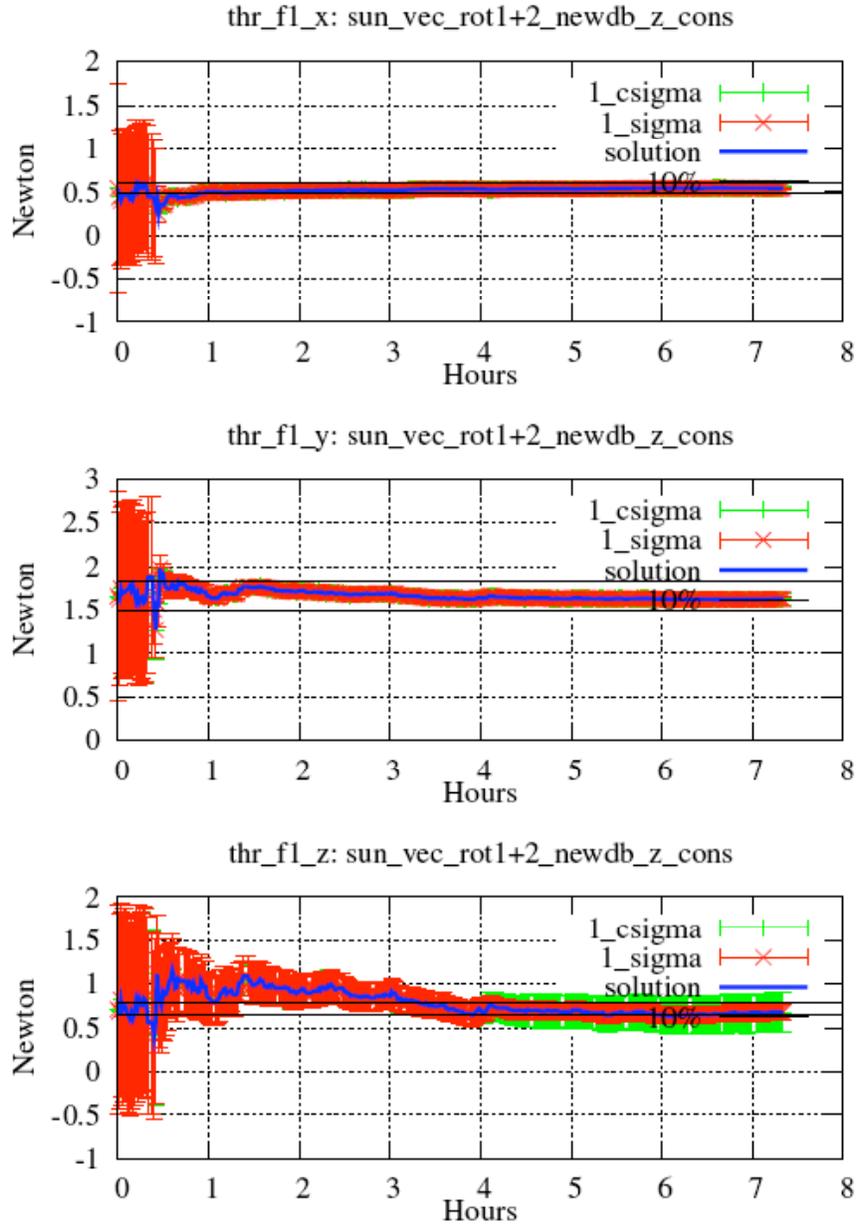


Figure 9. TIM Filter Solution for Phoenix RCS Thruster 1 Three-Axis Thrust Vector (X, Y, Z in Cruise Frame Coordinates), Based On Processing Simulated Doppler and Gyro Rate Data.

Table 3: TIM Filter Thrust Solution Based on Simulated Doppler and Gyro Rate Data Processing (from Active Calibration Campaign Validation Run). Error in Estimated Thrust Vector Direction (vs Simulated Truth) was < 1 Degree for Each Thruster.

	TIM Solution Thrust Magnitude (N)	Solution Magnitude Error - Sim Truth (N)	RSS Thrust Magnitude 1- σ Uncertainty from Covariance (N)	RSS Magnitude Uncertainty 1- σ / True Magnitude (%)
RCS Thruster 1	1.85	-0.03	0.09	4.9
RCS Thruster 2	1.86	-0.02	0.14	7.5
RCS Thruster 3	1.88	0.0	0.11	5.8
RCS Thruster 4	1.92	0.04	0.11	5.9

Finally, to develop an understanding of the relative sensitivity and importance of the two measurement types, a study was done to evaluate the quality of the reconstructed thrust vector, given the planned Active Thruster Calibration attitudes, based on either gyro rate data only or Doppler data only, versus using both data types together. Figure 10 shows plots of the Y-components and Z-components of thrust for RCS thruster 1 (results were similar in nature between the Y and X components, and between thruster 1 and the other thrusters). The two plots (for Y and Z thrust components) in the top row show the solution based on Doppler and gyro data combined in TIM is within 10% of the true value, with 1-sigma covariance levels also at the 10% level or better. This level of accuracy is not obtained by either using gyro data alone (middle row) or Doppler data alone (bottom row). Note that in the Doppler-only and Gyro-only cases, the same measurement noise values were used in the simulation and estimation as in the combined-measurement case.

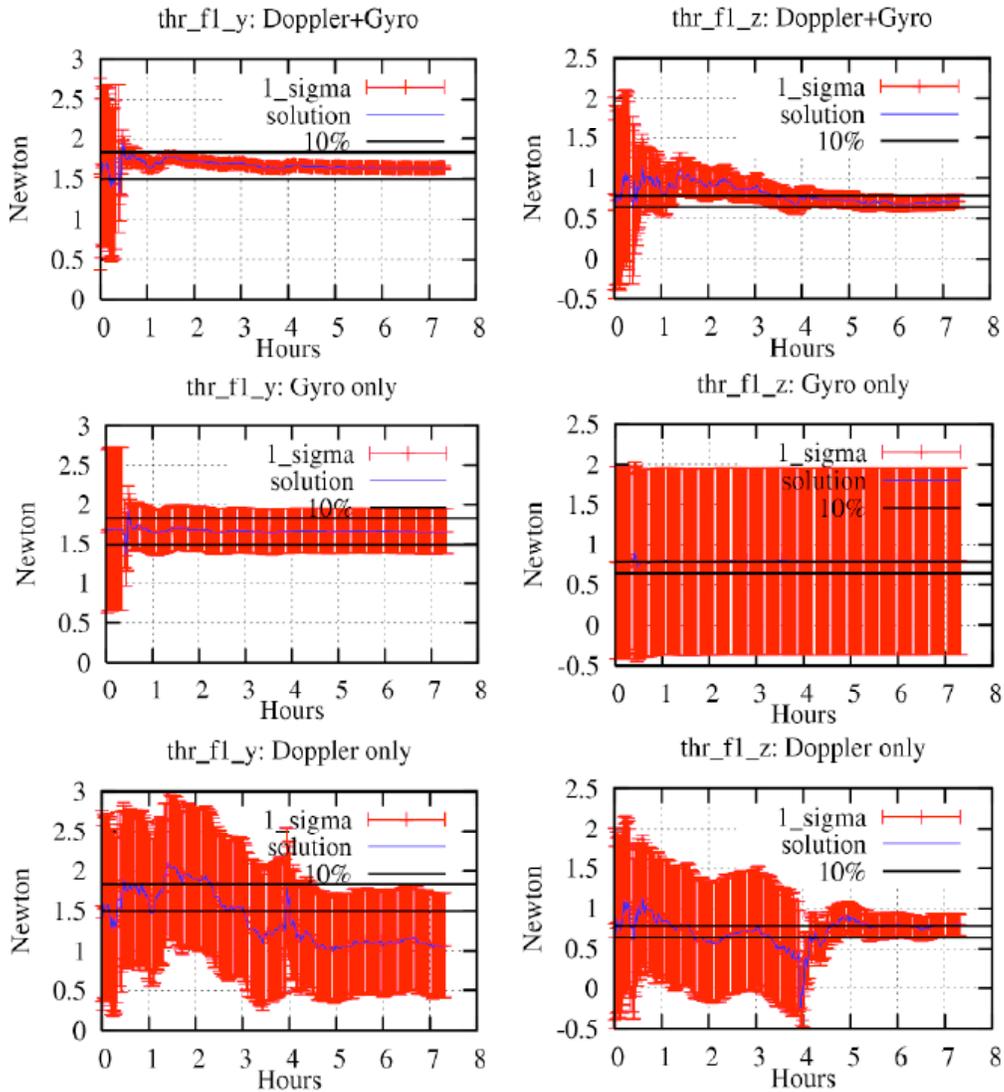


Figure 10. Comparison of Gyro-Only, Doppler-Only, and Doppler-and-Gyro Estimates of Phoenix Thruster 1 Y and Z Components

In Figure 10, the gyro-only results show that performance for the Y-axis force was much better than for the Z-axis force (X-axis force, not shown, was estimated somewhat better than the Y-axis force). This is because the cruise-frame X-coordinate of the spacecraft c.g. is a few centimeters less than 1 meter (compare with thruster locations in Table 2), so that firing RCS thrusters 1 and 2 together, or thrusters 3

and 4 together so that Z-direction impulses do not cancel, imparts very little torque about the Y axis – to such an extent that the Z-force is poorly observed by the gyros. Also in Figure 10, the Doppler-only accuracy was poor of the estimated X and Y force axes, and only moderately good for Z in the spacecraft’s second attitude, maintained over the final 3 hours, because the thermally- and communications-constrained attitude that Phoenix would undergo during its thruster calibration did not align the X and Y RCS thrust axes as closely as the Z axis with the spacecraft-Earth vector.

Hypothetically, had constraints on the Phoenix attitude permitted an Odyssey-like campaign in which the spacecraft underwent large attitude slewing, it is likely that the Doppler-only results could have been somewhat more Odyssey-like, differing to the extent that Odyssey’s attitude was maintained by its reaction wheels and Phoenix would still have had more deadbanding impulses “breaking up” the continuity of the Doppler passes. However, and importantly: the Phoenix limited-attitude-slewing approach to active thruster calibration was enabled by combining the two data types simultaneously in the TIM filter to estimate the average thrust vector states; the combined-data method yielded considerable improvement in solution accuracy over gyro-only or Doppler-only.

In summary, the two-part validation analyses showed that the TIM filter was able to meet its objective of estimating the three-axis average thrust vector for each Phoenix RCS thruster, given a combined set of Doppler “residual measurement” data and also IMU telemetry. The estimate would depend roughly equally on both data types in order to produce an estimate that was accurate to within 10% of the truth, with an accurate consider covariance treatment of formal uncertainty, obtained through its sigma-point consider filter formulation. This accuracy level was deemed to be satisfactory to allow using the new TIM filter for the thruster calibration during Phoenix flight operations.

V. Choosing the Active Thruster Calibration Attitudes, and Conclusions

With a validated filter and a working simulation, an activity commenced to design the actual spacecraft attitude motions needed to provide a high-quality measurement of Phoenix the RCS thruster forces during the Active Thruster Calibration activity. The spacecraft attitudes during the slews were subject to spacecraft thermal and communications constraints, and therefore the design activity jointly involved spacecraft thermal, communications, navigation and GN&C team members, to determine the effectiveness of various thermal-and-communications-constrained slewing strategies.

The accuracy of the thrust vector estimates was – not unexpectedly - sensitive to spacecraft orientation with respect to the spacecraft-Earth line, as this would determine the observability of each thruster’s impulse in the (line-of-sight) Doppler measurements. Figures 11 and 12 contain plots of “Doppler-Thrust-CG (TCG) Observability” - defined as the dot product (expressed as a percentage) of the Earth-spacecraft unit vector and the unit vector in the direction of each RCS thruster’s “non-torque-ing” force, i.e. force acting through the spacecraft c.g. - for two different Phoenix candidate thruster calibration attitude-history scenarios.

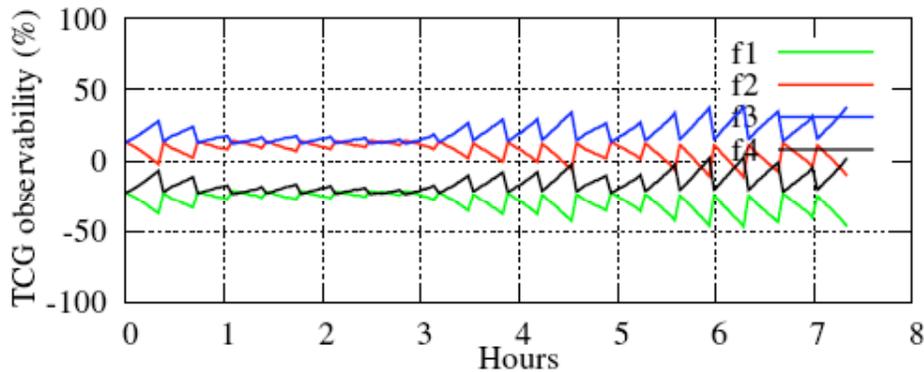


Figure 11. Doppler-Thrust-CG Observability for Nominal Phoenix Inner-Cruise Attitude – Corresponds to Thrust Magnitude Uncertainty (1-sigma from Covariance) of 11% – 33% for All RCS Thrusters

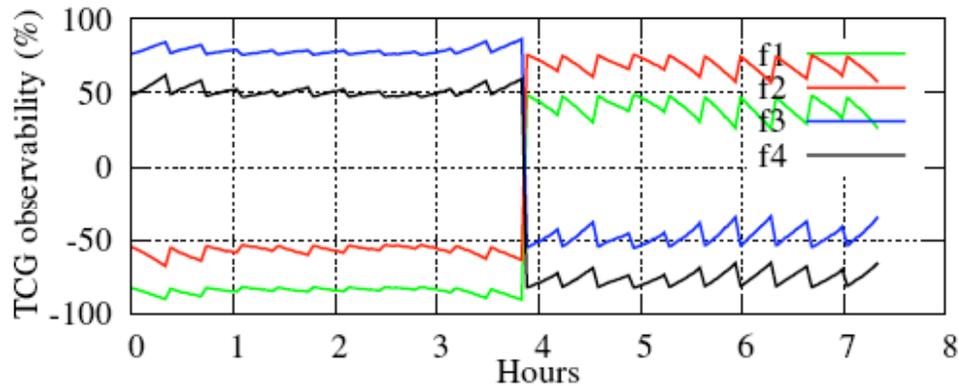


Figure 12. Doppler-Thrust-CG Observability for Active Thruster Calibration Attitude – Corresponds to Thrust Magnitude Uncertainty (1-sigma from Covariance) < 8% for All RCS Thrusters

The plots in Figure 12, compared with those in Figure 11, suggested strongly that slewing about the spacecraft-sun-vector to two different (power-, thermally- and communications-constrained) orientations would improve Doppler observability for all thrust axes, versus performing the thruster calibration while remaining in the normal daily “inner cruise” attitude of Phoenix. The Doppler observability of the portion of each thruster’s thrust vector that acted through the c.g. would range from 50 to 100% during portions of the 7.5-hour calibration activity for the two-attitude-slew scenario, versus less than 20% observability of thrusters 2 and 4, and less than 50% observability of thrusters 1 and 3, by staying in the inner cruise attitude.

Moreover, analyzing the accuracy and uncertainty using the TIM filter for both attitude strategies showed the formal thrust magnitude uncertainty from filter covariance to be reduced to less than 8% for the two-attitude-slew strategy, versus much larger uncertainties (11% to 33%) arrived at by staying in the inner cruise attitude for the thruster calibration. After consideration of numerous attitude and slewing scenarios, using observability plots and also simulated processing with the TIM filter, the team selected two attitudes which would first orient the positive Z-axis, then, with a slew at 4 hours into the calibration sequence, orient the negative Z-axis, to within about 47 degrees of the spacecraft-Earth line-of-sight vector, in a rotation about the spacecraft-sun vector (in order to keep the solar array pointing unchanged). The active thruster calibration campaign implementation is discussed further in the paper by Portock (Ref. 9).

The Phoenix Active Thruster Calibration took place on September 14, 2007, 41 days after Phoenix launched. Over a 7.5-hour period, the sequence of RCS thruster firings was executed without any issues, resulting in a set of high-rate gyro rate data and Doppler tracking data that was complete per the plan. The processing of this data with the TIM filter would be the first time a Sigma-Point Consider Filter has been used in deep space flight operations[¶]. The results, obtained with some additional effort to correct for elevated catalyst-bed temperatures, are described by Portock et al, and indicate that a level of accuracy was achieved in measuring the RCS thrust vectors that was better than 10%-level, consistent with expectations that arose from the work described here.

Acknowledgement

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

[¶] An SPCF was used in 2005 to perform experimental orbit determination runs based on MRO telemetered IMU data during its aerobraking phase at Mars - this “shadow navigation” was done in a non-operational capacity, using the filter described in Ref. 8.

VI. References

1. Antreasian, P., D. Baird, J. Border, P. Burkhart, E. Graat, M. Jah, R. Mase, T. McElrath, and B. Portock. "2001 Mars Odyssey Orbit Determination During Interplanetary Cruise." *AIAA JSR*, V. 42, No. 3, May-June 2005.
2. Wawrzyniak, G., D. Baird, E. Graat, T. McElrath, B. Portock, and M. Watkins. "Mars Exploration Rovers Orbit Determination System Modeling." Paper AIAA 2004-4983, *Proceedings of the AIAA/AAS Astrodynamics Specialist Conference*, Providence, RI, Aug. 2004.
3. You, T., E. Graat, A. Halsell, D. Highsmith, S. Long, R. Bhat, S. Demcack, E. Higa, N. Mottinger, and M. Jah. "Mars Reconnaissance Orbiter Interplanetary Cruise Navigation." *Proceedings of 20th International Symposium on Space Flight Dynamics*, Annapolis, MD, Sept. 2007.
4. Kennedy, B. M., E. Carranza and K. Williams. "1-AU Calibration Activities for Stardust Earth Return." *Proceedings of the AAS/AIAA Space Flight Mechanics Meeting*. Maui, HI, Feb. 2004.
5. Nandi, S., B. Kennedy, K. Williams, and D. Byrnes. "On-Orbit Maneuver Calibrations for the Stardust Spacecraft." Paper AIAA 2006-6409 *Proceedings of the AIAA/AAS Astrodynamics Specialist Conference*, Keystone, CO, Aug. 2006.
6. Lisano, M. E., "Nonlinear Consider Covariance Analysis Using a Sigma-Point Filter Formulation." *Proceedings of the 29th AAS Guidance and Control Conference*, Breckenridge, CO, Feb. 2006.
7. Lisano, M. E., "Comparing Consider-Covariance Analysis with Sigma-Point Consider Filter and Linear-Theory Consider Filter Formulas", *Proceedings of the 20th International Symposium on Space Flight Dynamics*, Annapolis, MD, Sept. 2007.
8. Jah, M., M. Lisano, G. Born and P. Axelrad, *Mars Aerobraking Spacecraft State Estimation by Processing Inertial Measurement Unit Data*, J. Guidance Control Dynamics AIAA, Accepted Dec. 2007.
9. Portock, B. "Navigation Challenges of the Mars Phoenix Lander Mission." AIAA Astrodynamics Specialist Conference, Honolulu, HI, Aug. 2008.
10. Tapley, B., B. Schutz, G. Born. Statistical Orbit Determination. Elsevier Academic Press, 2004. pp. 387-437.

VII. Appendix - The Sigma-Point Consider Filter Algorithm

The Sigma-Point Consider Filter (SPCF) algorithm, which was the basis of the TIM filter, appears in other sources, particularly Refs. 6 and 7, and is included in this paper as an appendix for the convenience of the reader. The SPCF was developed in 2005 to enable sigma-point filters to be used in conditions when a conservative, consider-analysis approach to estimation is needed. As discussed in Ref. 6, the SPCF algorithm is obtained in the same manner of partitioning the filter state vector into estimated and considered parameters as used by Tapley, Schutz and Born¹⁰ to obtain the LTCF, but proceeds along the lines of the derivative-free, Sigma-Point Kalman Filter algorithm. A brief description of the Sigma-Point Kalman Filter algorithm is provided in Ref. 6 or in the appendix of Ref. 7, among numerous other references in the literature.

The SPCF algorithm is as follows: a $(p \times 1)$ list \mathbf{C} of constant, non-estimated parameters whose errors are to be considered augments the $(n \times 1)$ estimated parameter list \mathbf{X}_{est} , to form a partitioned consider state vector, \mathbf{X}_{cons} :

$$\mathbf{X}_{cons} \equiv \begin{bmatrix} \mathbf{X}_{est} \\ \mathbf{C} \end{bmatrix} \quad \begin{matrix} (n \times 1) \\ (p \times 1) \end{matrix} \quad (A1)$$

As with the linear-theory consider filter derived in [10], the associated consider covariance \mathbf{P}_{cons} is partitioned according to \mathbf{X}_{cons} , i.e.

$$\mathbf{P}_{cons} \equiv \left[\begin{array}{c|c} \mathbf{P}_{cons(n \times n)}^{xx} & \mathbf{P}_{(n \times p)}^{xc} \\ \hline \mathbf{P}_{(n \times p)}^{cx} = (\mathbf{P}_{(p \times n)}^{xc})^T & \mathbf{P}_{(p \times p)}^{cc} \end{array} \right] \quad (A2)$$

Now, adapting the UKF formulation approach to computing a-priori sigma points in an unscented filter, first we factorize P_{cons} (best estimate of covariance from previous epoch t_{k-1}) into $S_{\text{cons},k-1}(S_{\text{cons},k-1})^T$. In order to preserve the constancy of P^{cc} , the lower-triangular square matrix $S_{\text{cons},k-1}$ is computed using a block-Cholesky decomposition, with the initial pivot starting in the P^{cc} partition.

$$S_{\text{cons},k-1} = \begin{bmatrix} S_{\text{cons},k-1}^{\text{xx}} & 0_{(n \times p)} \\ S_{k-1}^{\text{cx}} & S_{k-1}^{\text{cc}} \end{bmatrix} \quad (\text{A3})$$

where

$$S_{k-1(p \times p)}^{\text{cc}} \equiv \sqrt{(\lambda_c + p)P^{\text{cc}}} = \text{const.} \quad (\text{A4})$$

$$S_{k-1(p \times n)}^{\text{cx}} \equiv (\lambda_c + p) \left(S_{k-1}^{\text{cc}} \right)^{-1} P_{k-1}^{\text{cx}}, \text{ and} \quad (\text{A5})$$

$$S_{k-1(n \times n)}^{\text{xx}} \equiv \sqrt{(\lambda_x + n)P_{\text{cons},k-1}^{\text{xx}} - \left(S_{k-1}^{\text{cx}} \right)^T S_{k-1}^{\text{cx}}}. \quad (\text{A6})$$

and in which $\lambda_c = 3 - p$, and $\lambda_x = 3 - n$. Equation (A5) for S^{cx} will always be analytic, because sub-matrix S^{cc} is a square root of constant positive definite matrix P^{cc} .

As with the linear-theory consider algorithm, the sigma-point consider filter requires that X_{est} and P^{xx} be estimated with the standard sequential estimator algorithm, in this case the Sigma-Point Kalman Filter (also called ‘‘Unscented Kalman Filter’’ in some references) algorithm.

Now, in addition to estimating an n^{th} order system in the ‘‘usual way’’, to compute the additional uncertainty from consider the non-estimating parameters, at epoch t_{k-1} we generate $(2p + 1)$ consider sigma-points, which are a set[#] of samples of system state realizations at its a-priori expected value and also at $2p$ select points on the uncertainty ellipsoid:

$$\chi_{\text{cons},k-1,i} = X_{k-1}^+ + \sigma_{\text{cons},k-1,i}, \quad i = 0, 1, \dots, 2p \quad (\text{A7})$$

where $\sigma_{\text{cons},k-1,0} = 0$, and $\sigma_{\text{cons},k-1,i} = \pm \text{columns } (n + 1) \text{ through } (n + p) \text{ of } (S_{\text{cons},k-1})^T$, containing the square-root of the dispersed values of consider parameter list C plus expected variations in a-priori estimated parameter list $X_{\text{est},k-1}$ due to cross-correlations with C .

These $(2p + 1)$ consider sigma-points are then propagated from epoch t_{k-1} to epoch t_k , using the $((n + p) \times 1)$ vector of nonlinear expressions for the time-evolution of the augmented state vector, e.g. the equations of motion for the dynamical system, augmented with a $(p \times 1)$ list of zeros to maintain the constancy of the consider parameters. In this manner, we obtain the $(2p + 1)$ predicted sigma-points at epoch t_k :

$$\chi_{\text{cons},k,i} = f(t_{k-1}, t_k, \chi_{\text{cons},k-1,i}, u_{k-1}) \quad i = 0, 1, \dots, 2p \quad (\text{A8})$$

Per the sigma-point Kalman filter formulation, the predicted mean consider state is computed using the predicted sigma-points as:

$$X_{\text{cons},k}^- = \frac{1}{p + \lambda_c} \left(\lambda_c \chi_{\text{cons},k,0} + \frac{1}{2} \sum_{i=1}^{2p} \chi_{\text{cons},k,i} \right) \quad (\text{A9})$$

Then, the predicted consider covariance is computed as:

[#] For parameter list $X \in \mathbb{R}^L$ belonging to a multivariate probability distribution, $2L + 1$ samples is the minimum needed to capture the first two moments (expected value and covariance) of the distribution.

$$P_{cons,k}^- = \begin{bmatrix} P_{cons,k}^{xx-} & P_k^{xc-} \\ P_k^{cx-} & P^{cc} \end{bmatrix} = \frac{1}{p + \lambda_c} \left\{ \lambda_c (\chi_{cons,k,o} - X_{cons,k}^-) (\chi_{cons,k,o} - X_{cons,k}^-)^T + \frac{1}{2} \sum_{i=1}^{2p} (\chi_{cons,k,i} - X_{cons,k}^-) (\chi_{cons,k,i} - X_{cons,k}^-)^T \right\} \quad (A10)$$

The primary utility of matrix $P_{cons,k}^-$ is that it yields the $(n \times p)$ sub-matrix P_k^{xc-} of predicted cross-correlations between consider and estimated parameters. At epochs for which measurements are available, the change due to the measurement information on the cross-correlation sub-matrix is calculated the following way:

$$P_k^{xc+} = P_k^{xc-} - K_k (P_k^{xy})^T (P_k^{xx-})^{-1} P_k^{xc-} - K_k (P_k^{cy})^T \quad (A11)$$

in which P_k^{cy} is the cross-correlation matrix between consider-parameter space and measurement space, computed by dispersing the measurement model equation on the $2p + 1$ consider-parameter sigma-points and forming a discrete covariance matrix in the same manner as one obtains P_k^{xy} (c.f. Equation A9 in the appendix of Ref. 7).

The optimal gain K_k , residual covariance P_k^{vv} and a-priori state covariance P_k^{xx-} are all obtained from a standard sigma-point filter (that estimates \mathbf{X}_{est} and neglects errors in \mathbf{C}) operating simultaneously in time with the consider filter algorithm. Finally, the additive uncertainty from considering non-estimated parameters is calculated as:

$$dP_{cons,k} = P_k^{xc+} (P^{cc})^{-1} (P_k^{xc+})^T \quad (A12)$$

and the updated consider covariance matrix, post-measurement, for epoch t_k is given by:

$$P_{cons,k} \equiv \begin{bmatrix} P_{cons}^{xx+} & P_k^{xc+} \\ P_k^{cx+} & P^{cc} \end{bmatrix} \quad (A13)$$

in which the state covariance augmented for consider-parameter uncertainty ($P_{cons,k}^{xx+}$) is related to the standard state covariance P_k^{xx+} by:

$$P_{cons,k}^{xx+} = P_k^{xx+} + dP_{cons,k} \quad (A14)$$

If no observation is available to be processed at time t_k , Equations (A12) and (A13) are evaluated using P_k^{xx-} and P_k^{xc-} instead of P_k^{xx+} and P_k^{xc+} . This algorithm is recursive, so that after calculating $P_{cons,k}$, if filter processing is to continue, the time index k is set to $(k-1)$, and the steps above are repeated at the next processing epoch, until all measurements have been applied or the desired end time has been reached.