Preliminary Assessment of the Mars Science Laboratory Entry, Descent, and Landing Simulation.

David W. Way
National Aeronautics and Space Administration
Langley Research Center
Hampton, VA 23681-2199
757-864-8149
david.w.way@nasa.gov

Abstract— On August 5, 2012, the Mars Science Laboratory rover, Curiosity, successfully landed inside Gale Crater. This landing was only the seventh successful landing and fourth rover to be delivered to Mars. Weighing nearly one metric ton, Curiosity is the largest and most complex rover ever sent to investigate another planet. Safely landing such a large payload required an innovative Entry, Descent, and Landing system, which included the first guided entry at Mars, the largest supersonic parachute ever flown at Mars, and a novel and untested Sky Crane landing system. A complete, end-to-end, six degree-of-freedom, multi-body computer simulation of the Mars Science Laboratory Entry, Descent, and Landing sequence was developed at the NASA Langley Research Center. In-flight data gathered during the successful landing is compared to pre-flight statistical distributions, predicted by the simulation. These comparisons provide insight into both the accuracy of the simulation and the overall performance of the vehicle.

TABLE OF CONTENTS

NOMENCLATURE

σ	Standard deviation
GN&C	Guidance, Navigation, and Control
ATLO	Assembly, Test, and Launch Operations
BUD	Bridle, Umbilical, and DRL
DGB	Disk-Gap-Band parachute
DoF	Degrees of Freedom
DRL	Descent Rate Limiter
DSN	Deep Space Network
EBMD	Entry Balance Mass Device
EDL	Entry, Descent, and Landing
EPU	EDL Parameter Update
EVR	EVent Report real-time telemetry
IMU	Inertial Measurement Unit
LLWG	Landing Location Working Group

[1] IEEEAC Paper #2755, Version 3, Updated January 22, 2013.

Figure 1. Artist's concept of novel Sky Crane system used to land *Curiosity*. Image Credit: NASA/JPL-Caltech.

MER	Mars Exploration Rover mission
MSL	Mars Science Laboratory mission
OD	Orbit Determination
POST2	Program to Optimize Simulated Trajectories II
RTI	Run Time Interrupt, 8 Hz
sRTI	Sub-RTI, 64 Hz
SUFR	Straighten-Up and Fly Right
TCM	Trajectory Correction Maneuver
TDS	Terminal Descent Sensor – RADAR
TPS	Thermal Protection System
TZERO	EDL GN&C start time, $t = 0s$

1. INTRODUCTION

On August 5, 2012, at 10:32 PDT[2], the Mars Science Laboratory (MSL) *Curiosity* rover successfully landed on Mars. After completing a 245 day, 568×10^6 km, inter-planetary transit, *Curiosity* entered the Martian atmosphere 125 km above the surface and traveling at 5.845 km/s. During the next seven minutes, the rover flawlessly executed a complex sequence of autonomous actions, safely coming to rest just 2.385 km away from the 4.5965 ^oS and 137.4019 ^oE target
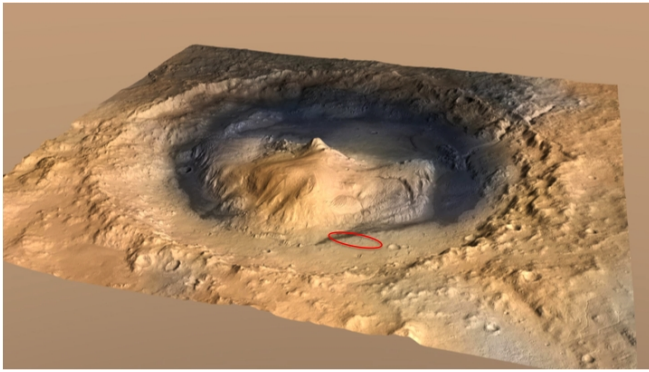
[2]Earth Received Time.

Figure 2. This oblique view of Mt. Sharp (Aeolis Mons), rising approximately 5.5 kilometers above the floor of Gale Crater, shows the MSL landing footprint in red. There is no vertical exaggeration in the image. Image Credit: NASA/JPL-Caltech/ESA/DLR/FU Berlin/MSSS.

inside Gale Crater. These seven minutes were an extremely critical and challenging phase of the MSL mission, known as Entry, Descent, and Landing (EDL). At nearly 900 kg, a full five times heavier than the previous *Spirit* and *Opportunity* rovers, *Curiosity* is the largest and most sophisticated rover to ever explore beyond Earth.

The challenges associated with landing a rover that is the size of a small car, coupled with unprecedented requirements for accuracy, led to the development of a unique EDL system architecture that incorporates both heritage and innovation, while extending the limits of the Viking-derived EDL technologies qualified by the Mars Viking, Mars Pathfinder (MPF), and Mars Exploration Rover (MER) missions [**?**]. Accordingly, several elements of the MSL EDL design were technological firsts for Mars, such as the first guided entry and the novel "Sky Crane" landing system (pictured in Figure 1). These elements, now flight-validated, form an important step forward for the future of Mars exploration, which will likely require the precise landing of even larger payloads.

Because of differences in atmosphere and gravity, end-to-end EDL system verification and validation tests are not possible on Earth. Mars flight projects must, therefore, rely heavily on computer simulation results. Consequently, the EDL simulation is a key element in any successful landing. Simulation predictions are used throughout the project life-cycle: to inform EDL design choices, to compare and certify candidate landing sites, to verify EDL system performance, to select flight software parameters, and to evaluate operational decisions. Considering the importance of these activities, it is crucial for the EDL simulation to accurately, yet conservatively, model and predict the complex flight dynamics of the EDL system. Since each successful mission leads progressively forward in succession to the next, it is vitally important to the EDL community for each mission to provide a critical assessment of the EDL simulation and the models used. This assessment is made by comparing data collected in-flight, along with post-flight reconstructions, to pre-entry simulation predictions. This paper provides an initial assessment of the MSL EDL simulation, using flight data that was available shortly after landing. Additional comparisons will be made as post-flight reconstructions become available [**?**].

2. MSL POST2 SIMULATION

The Program to Optimize Simulated Trajectories II (POST2) is a general Six Degree-of-Freedom (6-DoF) trajectory simulation tool, that solves both the translational and rotational equations of motions for up to 20 independent rigid bodies. It is maintained by the NASA Langley Research Center and has been used to solve a wide variety of flight dynamics and trajectory optimization problems. More germanely however, POST2 has had significant Mars EDL flight heritage. POST2 has been used successfully on the Mars Pathfinder [**?**], Mars Exploration Rover (MER) [**?**], Mars Phoenix [**?**][**?**], and now Mars Science Laboratory (MSL) missions.

A complete, end-to-end, POST2 simulation of Curiosity's complex EDL sequence was developed at the NASA Langley Research Center. This simulation was designated as the prime EDL performance simulation for MSL. Reference [**?**] contains a detailed description of the MSL EDL architecture, which is illustrated here in Figure 3 and Figure 4. This architecture consists of six segments: Exo-atmospheric Flight, Guided Entry, Parachute Descent, Powered Descent, Sky Crane, and Flyaway. Reference [**?**] provides a description of the sensed triggers used at critical transitions.

The POST2 simulation modeled all events of the MSL EDL sequence, beginning 50 s after Cruise Stage separation. Arguably one of the largest and most complex simulations of its kind, it leveraged the versatility and heritage of previous POST2 simulations and added MSL-specific models and flight software. Fourteen independent bodies were modeled, including the Descent Stage, the parachute, the heatshield, the backshell, the rover, and each of the eight ejected balance masses. Multi-body forces, originally developed for MER, were also used to model vehicle configurations containing two or more of these elements during two of the EDL segments. The first of these was the Parachute Descent segment, in which the parachute and backshell were attached via parachute riser lines. The second was the novel Sky Crane segment, in which the Descent Stage and rover were attached via the Bridle, Umbilical, and Descent Rate Limiter (BUD).

The EDL simulation is used to assess the robustness of the EDL system to off-nominal or uncertain conditions by tracking statistics on pre-defined output variables. These output variables are typically instantaneous flight conditions at specific events (e.g. Mach number at parachute deploy), but may also be minima or maxima of flight parameters experienced over specific regions or segments of EDL sequence (e.g. maximum entry deceleration). These outputs are collected, in a Monte Carlo fashion, over thousands of individual instantiations of the simulation, where each case contains random samples of uncertain parameters, such as initial states and environmental variables. The MSL simulation contained 682 random dispersions and 4834 output variables.

The following sections utilize real-time Event Records (EVRs), as well as recorded telemetry generated by the navigation filter and Inertial Measurement Unit (IMU). The as-flown values are compared to statistical distributions predicted by the simulation. In each case, the Gaussian quantile of the as-flown value is measured, relative to the simulation results. This measurement is expressed as a number of standard deviations (e.g. 3σ). Values within approximately +/-1.5σ indicate good agreement, while values outside of +/-3σ indicate a significant disagreement between actual and simulated results.
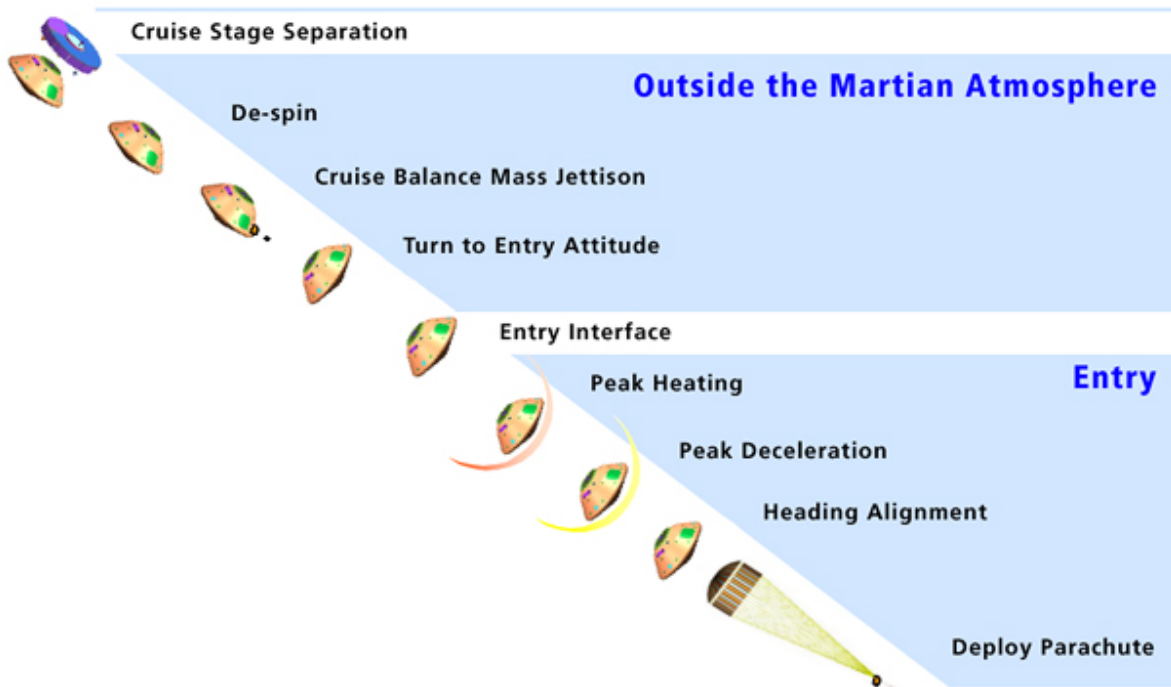
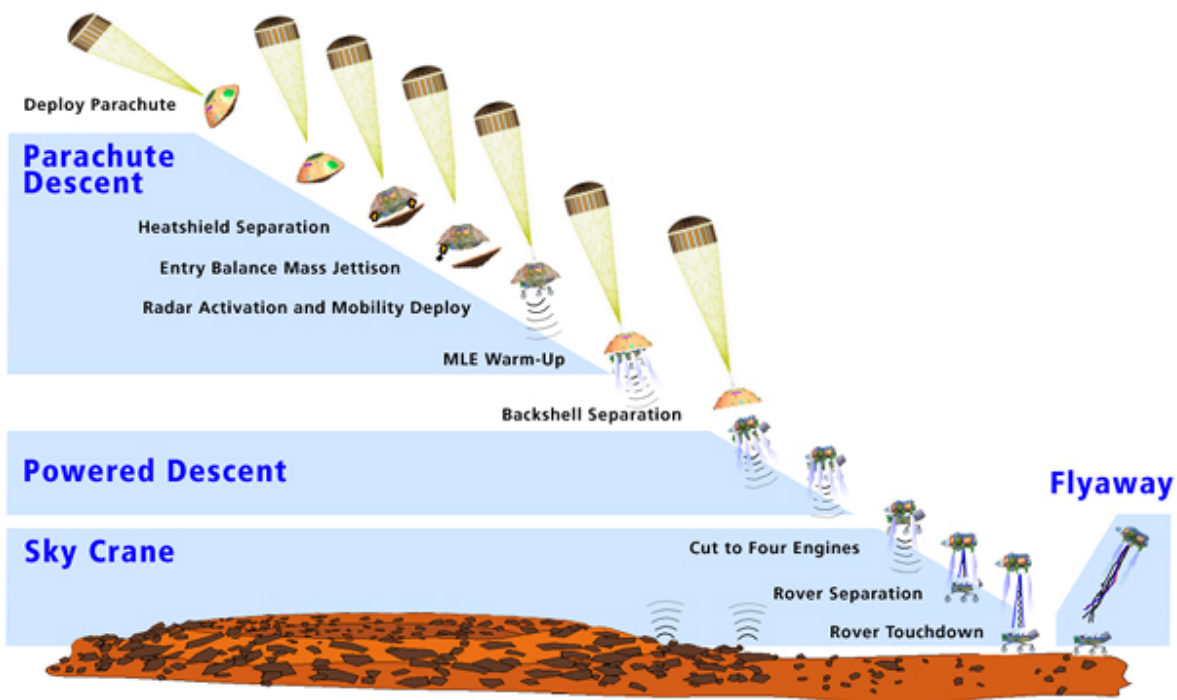**Figure 3**. EDL event sequence, Cruise Stage Separation to Parachute Deploy. Image Credit: NASA/JPL-Caltech.

Figure 4. EDL event sequence, Parachute Deploy through Flyaway. Image Credit: NASA/JPL-Caltech.

3. LANDING FOOTPRINT

One of the more outwardly visible simulation predictions is the landing footprint, or ellipse. Footprint predictions are used throughout the project lifecycle: to inform design choices during design and development, to certify landing sites during the landing site selection process, to verify system performance during ATLO, to place the final landing target during risk balancing, and to inform TCM and EPU decisions during operations. Considering the importance of these activities to the success of the mission, it is evident that the simulation must accurately predict the landing footprint. Or, more precisely, that the simulation must provide a suitably conservative prediction, such that the risk of landing outside of the ellipse is diminishingly small. This was particularly pertinent for MSL, given the proximity of the landing ellipse to hazardous topographic features. Figure 2 presents the topography of *Curiosity*'s landing site in Gale Crater, which is dominated by the presence of Mt. Sharp in the center of the crater. The MSL landing site is indicated by the red ellipse, located on the floor of the crater, near the foot of Mt. Sharp.

Though autonomous entry guidance systems have been used extensively at Earth, and proposed on previous Mars missions, MSL became the first mission to employ a hypersonic guidance algorithm at Mars. Reference [**?**] provides a description of this entry guidance algorithm. This is a major advancement in planetary EDL systems, a first step toward precision landing at Mars, and an enabling element to the "go-to" paradigm in landing site selection. This paradigm is characterized by EDL systems with landing footprints that are smaller than the roving capability of the mobile science platform that they deliver. This situation breaks the restriction requiring the primary science targets to be both scientifically interesting and safe for landing, and makes available primary science targets that were previously unreachable. As was the case for MSL, the primary science target need only be in close proximity to a safe landing zone, within the design range of the rover.

Whereas previous entry systems had relatively large footprints (of order 100 km), the predicted footprint for MSL was only 21 x 7 km. For ballistic entries, the major drivers of footprint size are typically the uncertainties in the initial entry state (flight path angle) and the state of the atmosphere (density and winds). In contrast, because of the guidance's ability to "fly out" the majority of off-nominal dispersions in both entry state and atmospheric conditions, the footprint for a guided entry is typically dominated by navigation errors. For EDL operations, this had the effect of making the predicted size and location of the landing ellipse extremely stable. This stability is evident in Table 1, which lists the sizes and miss-distances of the calculated landing footprints in the final week of cruise leading up-to entry. Updates in the navigation solution generally changed the center of the predicted ellipse by only a few hundred meters.

Table 1. 99%-tile Footprint Predictions.

| OD Number	Date	Landing Footprint (major x minor axis)	Target Miss-Dist.
OD142	28-Jun-2012	17.94 x 6.18 km	437 m
OD152	12-Jul-2012	17.93 x 6.14 km	447 m
OD160	27-Jul-2012	18.73 x 6.45 km	147 m
OD167	28-Jul-2012	18.66 x 6.45 km	162 m
OD176	30-Jul-2012	18.62 x 6.41 km	160 m
OD179	30-Jul-2012	19.09 x 6.92 km	158 m
OD182	01-Aug-2012	19.02 x 6.83 km	215 m
OD187	02-Aug-2012	19.00 x 6.76 km	146 m
OD190	02-Aug-2012	18.92 x 6.68 km	152 m
OD195	04-Aug-2012	18.71 x 6.56 km	143 m
OD202	04-Aug-2012	18.77 x 6.53 km	151 m
OD213	05-Aug-2012	18.71 x 6.44 km	325 m
OD229	08-Aug-2012	18.59 x 6.37 km	297 m
OD230	27-Aug-2012	18.73 x 6.42 km	300 m
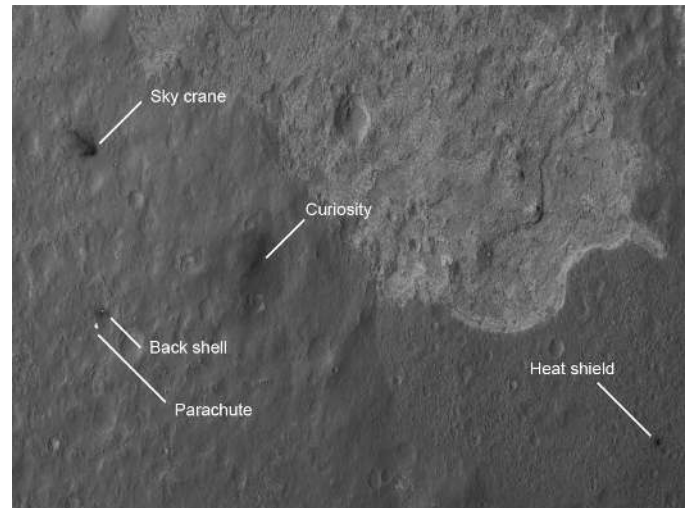
Figure 5. The High-Resolution Imaging Science Experiment (HiRISE) camera captured this image of the MSL landing site about 24 hours after landing. The image scale is 39 centimeters per pixel. Image Credit: NASA/JPL-Caltech/Univ. of Arizona.

During the Approach Phase of cruise operations, entry state estimation is performed on the ground by the mission navigation team. Orbit Determination (OD) solutions for the position and velocity of the spacecraft at atmospheric entry (defined as a radius of 3522.2 km) are based on radiometric doppler and range measurements obtained through the Deep Space Network (DSN). These solutions were used to inform operational decisions for Trajectory Correction Manuevers (TCMs), which are propulsive maneuvers used to re-target the entry state, and EDL Parameter Updates (EPUs), which are changes to on-board stored parameters utilized by the flight software. Prior to each of these scheduled decision points, the latest OD solution, including covariance estimates of the uncertainties, was used to initialize a Monte Carlo analysis of the EDL system performance. The final pre-

entry navigation solution was OD213. The predicted 99%-tile ellipse for this OD solution was 18.71 x 6.44 km. Two OD solutions, OD229 and OD230, were released after landing. These solutions utilized all available radiometric data up to entry and up to landing, respectively, but had little effect on predicted EDL performance or footprint size. Therefore, for the purposes of this paper, the OD229 Monte Carlo will be arbitrarily chosen as the basis of comparison.

Immediately following landing, one of the first reconstruction tasks facing the EDL team was to locate the landed position of the rover. A Landing Location Working Group (LLWG) was organized for just this purpose. The group's first estimate (4.59 ^oS, 137.44 ^oE), provided within minutes of the landing, was accurate to within 104.6 m. This estimate combined expected navigation errors from the POST2 simulation results with the estimated landing position reported by the onboard navigation system[?]. Within 24 hrs of landing, the LLWG was able to further improve *Curiosity*'s landed position estimate (to within approximately 5 m, relative to topographic features) by comparing MARDI images of local terrain, taken during the EDL sequence, to previous HiRISE images of the landing ellipse. However, it should be noted that differences in position on the order of 10's of meters are possible due to map-tie error, which is the uncertainty in the absolute inertial position of Mars surface features. This site (4.5895 ^oS, 137.4417 ^oE), shown in Figure 5, was subsequently named "Bradbury Landing" in honor of the late Ray Bradbury.

```
Bradbury Landing Coordinates:
Latitude  = -4.5895 deg
Longitude = 137.4417 deg
```
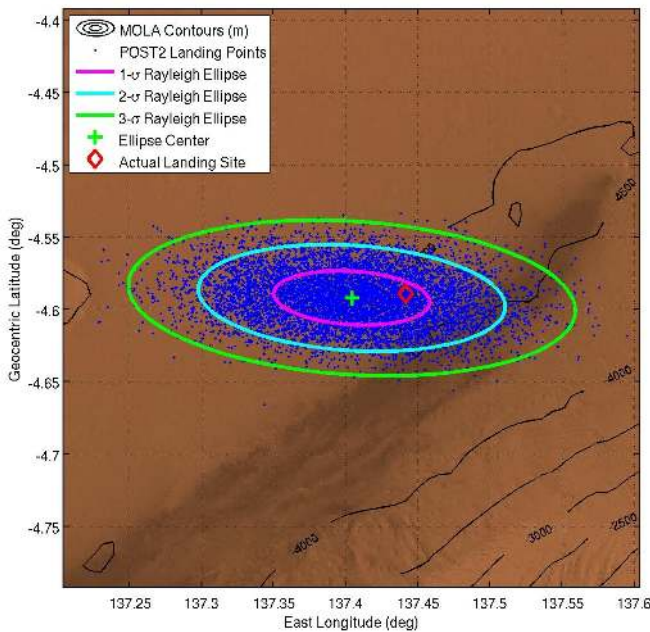

Figure 6. Comparison of actual landing location, indicated by the red diamond marker, to Monte Carlo predictions from OD229. The green ellipse is a 98.89%-tile ellipse, which is the cumulative probability of 3σ for a Rayleigh distribution. Likewise, the cyan and magenta ellipses are 86.47%-tile (2σ) and 39.35%-tile (1σ), respectively.

It is desired to assess the accuracy of the MSL landing prediction by comparing the actual landing location to the statistical distribution of the simulation landing points. Figure 6 compares the Bradbury Landing location relative to the 8001 landing points from the OD229 Monte Carlo, shown as blue dots. Three footprints are drawn for reference (1σ, numsig2, and numsig3): the magenta ellipse is 39.35%-tile, the cyan ellipse is 86.47 %-tile, and green ellipse is 98.89 %-tile. Clearly the landing was well within the predicted 3σ ellipse. However, a more detailed comparison may be derived from the same process used in constructing the landing footprint ellipses for plots such Figure 6. The Cholesky decomposition of the covariance matrix of the simulation results is used to map the correlated latitude-longitude points to an uncorrelated, standard bivariate normal distribution, which has units of standard deviation (σ). The observed landing quantile, determined in this manner and shown in Figure 7, is 0.74725σ. This value represents a 23.747%-tile landing (i.e. fewer than 24% of Monte Carlo cases landed closer to the target). The right-most subplot of Figure 7 is a quantile-quantile plot that has been modified to compare the Monte Carlo results to standard Rayleigh quantiles, rather than standard Gaussian quantiles. The nearly linear data points in this figure indicate that the Monte Carlo data very closely follows a Rayleigh distribution.

4. EDL TIMELINE

Figures 3 and 4 depict the complex sequence of EDL events that must be autonomously performed for a successful landing. This sequence, beginning at Cruise Stage Separation and ending with Flyaway, involves 8 vehicle configuration changes that require the firing of 59 pyrotechnic charges with split-second accuracy. This behavior is controlled by a Timeline Engine, which steps through 145 pre-defined timepoints, and a GN&C Mode Commander, which operates within 39 unique GN&C modes. The times at which these actions occur are captured both within the EDL simulation and the real-time telemetry EVent Records (EVRs), such as the following:

```
"Fired timepoint=TZERONAV_GNC_START
chain=TZERO_NAV_GNC_START
anchor=TZERO_NAV action=GNC_START
(time=397501175.000000 @ srti=0)."
```

Monitoring these events within the EDL simulation allows us to evaluate both the absolute time of each action as well as the amount of elapsed time between successive actions. The former was important to the proper phasing of telecom relay assets prior to EDL. [?] The later was important in verifying the adequacy of Timeline Engine margins under off-nominal conditions. Since many of these events were triggered by sensed conditions, the observed timing of these actions during EDL provides us with valuable insight into the accuracy of the simulation to predict these sensed conditions.

Figure 8 plots the observed absolute time of each Timeline Engine timepoint relative to the simulation mean from the OD229 Monte Carlo. Also plotted for reference are the 1%-tile and 99%-tile from the Monte Carlo. This data shows very good agreement between flight and simulation, with all of the timepoints occurring within 6.9089 s of their predicted mean times.

GN&C Mode Commander state transitions are also compared to EDL simulation predictions. First, all event times for
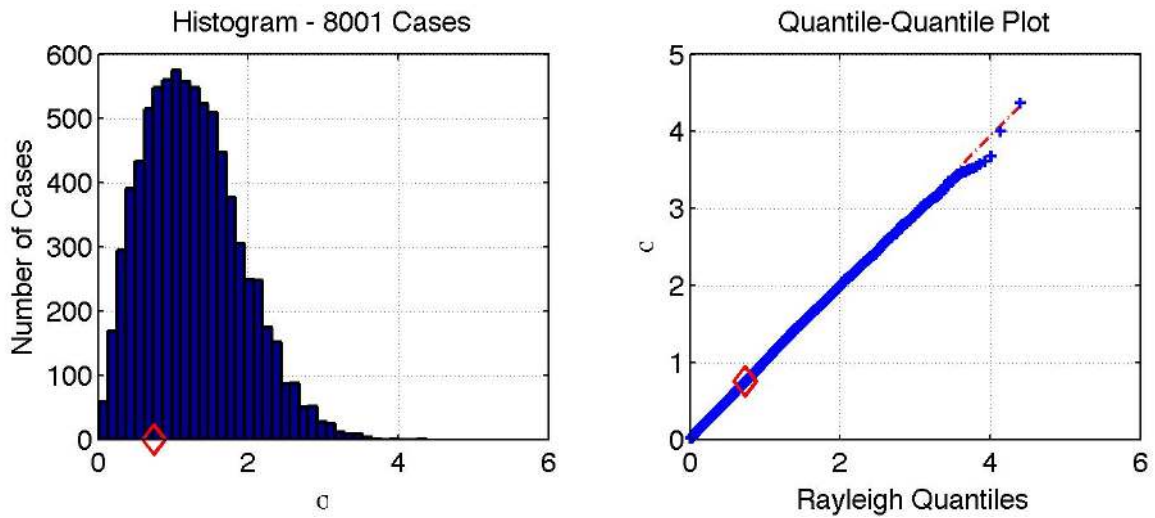
Figure 7. Comparison of landing quantile (σ), indicated by the Red Diamond Marker, to Monte Carlo Predictions from OD229. Modified quantile-quantile plot comparing quantiles of uncorrelated landing distance from OD229 Monte Carlo results to standard Rayleigh quantiles, rather than standard Gaussian quantiles. The actual landing location is indicated by the red diamond marker.
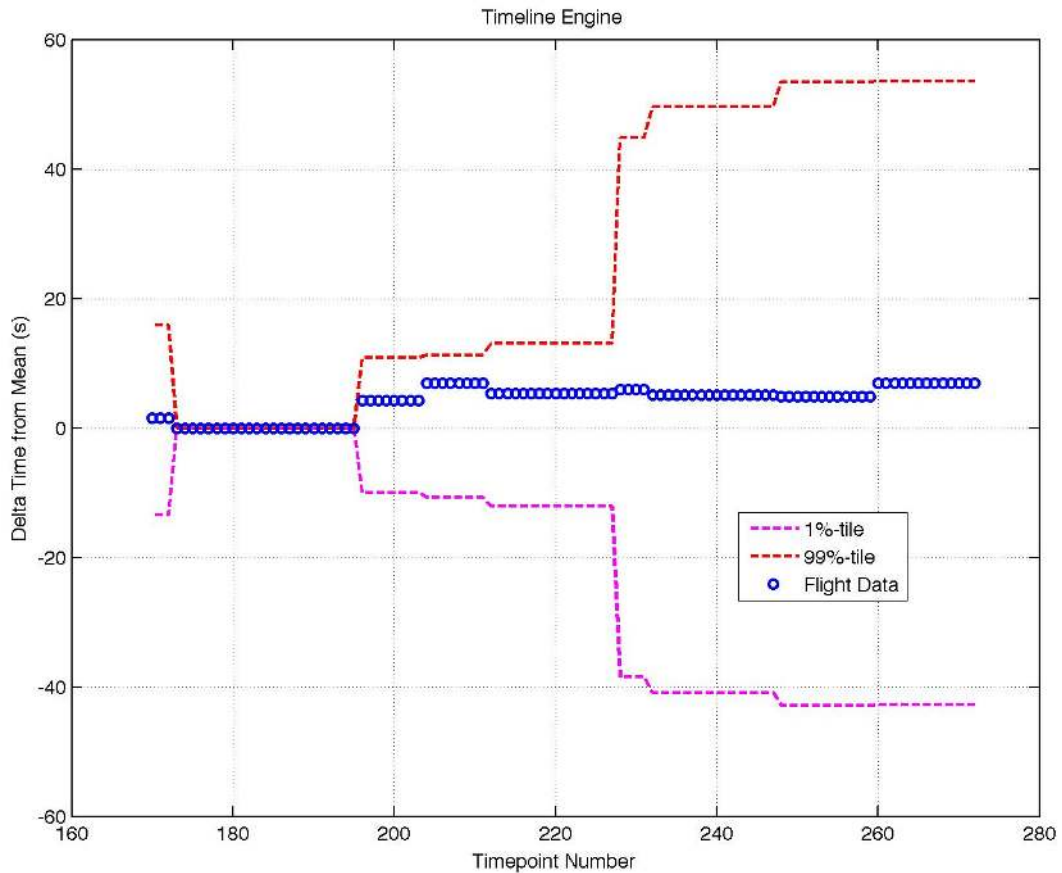

Figure 8. Comparison of Timeline Engine Timepoint Times to Monte Carlo Predictions from OD229.

6

Table 2. Error in Fixed-Time GN&C Modes.

MC State	GN&C Mode Name	Timing Error (ms)
16	ALLOW CHUTE DEPLOY TRANSIENTS TO SETTLE	0.375
18	RCS CONTROL INHIBITED	109.375
19	WAIT FOR TDS NAV INIT	-109.375
23	WAIT FOR BACKSHELL SEP	0.625
24	FREE FALL	16.25
25	MLE WARM UP	-16.25

Mode Commander state transitions were extracted from flight telemetry EVRs, such as the following:

```
"===== Gnc mc new state=
MC_RCS_WARM_UP at time=397501175.000000
===================================="
```

These times were then differenced to determine the cumulative time in each of the 38 GN&C Mode Commander states. A complete listing of the cumulative time in each GN&C mode is provided in Appendix A, Table 5.

Five GN&C modes were not used in flight: mode 1, 'OPEN LOOP STANDBY'; mode 2, 'OPEN LOOP FIRE DRCS THRUSTERS'; mode 3, 'OPEN LOOP SET MLE THROTTLES'; mode 31, 'PD CONSTANT DECELERATION SATURATED AT MAX'; and mode 32, 'PD CONSTANT DECELERATION SATURATED AT MIN'. Six additional modes were fixed-time modes (i.e. their transitions were triggered by timers and thus had no uncertainty, such that $\sigma = 0$). However, because of the non-flight manner in which the Timeline Engine and GN&C Mode Commander are executed in the simulation (sequentially in the simulation, versus asynchronously on the flight computer), slight differences exist for these modes. Table 2 lists the fixed-time modes along with the observed timing error. The largest of these errors is just under $110 \ ms$, which is difference between triggering on the following 8 Hz RTI, rather than the next 64 Hz sRTI. Differences at these time-scales, do not affect overall simulation performance, but should be noted when comparing to real-time systems, such as flight or the test-bed.

Omitting those modes discussed already, Figure 9 presents the observed times in each GN&C mode as compared to the simulated statistics from the OD229 Monte Carlo predictions. This data shows that, for the majority of the 39 GN&C modes, the flight data falls within +/-1.5σ, or 86.64% of Monte Carlo cases. However, there are several notable exceptions: Modes 8, 9, 15, 20, 21, and 35. Each of these will be discussed in the following sub-sections.

GN&C Mode 8 & 9: High Altitude Drag

GN&C Mode 8, 'WAIT FOR GUIDANCE START', is active from the Entry Interface epoch, which is a timed trigger of $540 \ s$ relative to TZERO_NAV, until the beginning of entry guidance range control, which is a sensed trigger of 0.2 $g's$ of deceleration. Thus, the duration of Mode 8 depends on the deceleration experienced at very high altitudes. The observed time in this mode was $45.875 \ s$, which is -3.836σ and less

than the 46.0 s minimum time predicted by the simulation. The extremely short time spent in this mode is indicative of much higher than expected deceleration in the upper atmosphere (above $50 \ km$). This could be explained by either a higher atmospheric density in this region and/or a higher drag coefficient. While it is not possible to separate the effects of drag coefficient and atmospheric density, it is estimated that their product would need to have been approximately 20% higher than nominal to produce the observed deceleration.

The net result of the quick deceleration through Mode 8 was to initiate active range control early, at a higher altitude and a farther range to target. The entry guidance responded to this situation by commanding an initial bank angle that was more lift-up than the pre-entry estimate by 8.594817 deg. In order to prevent large control errors, the initial guidance command is profiled. This is the purpose of GN&C Mode 9, 'RANGE CONTROL SLEW TO COMMANDED BANK ANGLE'. The observed time in Mode 9, which is proportional to the size of the profiled turn to the commanded bank angle, was 2.51σ. Therefore, the long time in Mode 9 is consistent with, and a direct result of, the short time in Mode 8.

Whatever the cause for the higher deceleration during Mode 8, the situation quickly corrected itself. By the time the vehicle completed Mode 9 (approximately $50 \ km$ altitude), the deceleration had returned to within 10% of the nominal value and the entry guidance began converging the range. Consequently, there were no lasting effects on entry performance, which is not surprising given the small fraction of deceleration that occurs above $50 \ km$.

GN&C Mode 14 & 15: Supersonic Drag

The time between Straighten Up and Fly Right (SUFR) start and parachute deploy was nominally targeted to be $17 \ s$. However, both events were triggered by navigated velocity, which allows this time to vary. The value of $17 \ s$ was chosen to allow sufficient time to safely eject all six Entry Balance Mass Devices (EBMDs) while completing the roughly 180 deg roll maneuver to the desired parachute deploy attitude. The trigger setting is determined from a nominal trajectory simulation by simply noting the velocity 17 seconds prior to reaching the desired parachute deploy condition. The actual timing of these two events depends on the expected supersonic aerodynamic deceleration from approximately Mach 2 to Mach 1.7.

Throughout the SUFR maneuver, the GN&C Mode Commander operates in two successive GN&C modes: Mode 14, 'SLEW AND SUFR SLEW TO RADAR ATTITUDE'; and Mode 15, 'SLEW AND SUFR WAIT FOR CHUTE DEPLOY'. The observed time in these two GN&C modes was 14.0 s and 5.25 s, respectively, with a combined time between SUFR and PD of 19.25 s. This represents a 3.23σ event, relative to the simulation predictions, and indicates very low supersonic deceleration. Though low deceleration could be indicative of a high ballistic coefficient, due to a low drag coefficient, correlations with cases in the Monte Carlo that also had long SUFR times favored lower, rather than higher, ballistic coefficients. This observation, coupled with stronger correlations with low dynamic pressures, is indicative that low density or tail winds are more likely the culprit.

Though the low supersonic deceleration likely contributed to landing downrange of the target, the extra 2.25 s spent waiting for parachute deploy had very little impact on the rest of the EDL system. During this time it is estimated
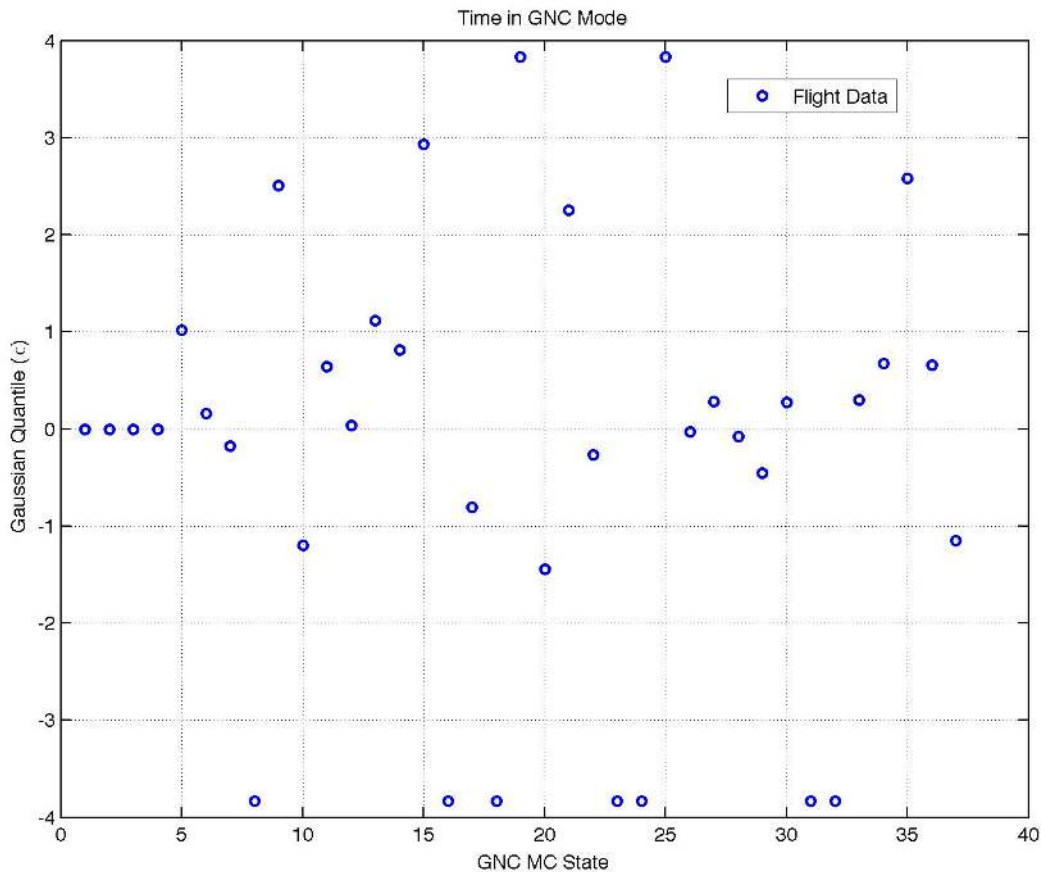
Figure 9. Comparison of Times in GN&C Modes to Monte Carlo Predictions from OD229.

that the vehicle decelerated approximately 5.6 m/s, traveled approximately 840 m downrange, and lost approximately 340 m of altitude. This altitude loss was of little consequence due the ample altitude margin available at Gale crater. However, future missions may consider replacing the second sensed trigger with a timed trigger, effectively making the parachute mortar fire a part of the SUFR timepoint sequence, and removing any possibility of trigger collisions. Even so, additional investigation into this anomaly is warranted to better understand the expected aerodynamic performance in this critical flight regime.

GN&C Mode 20 & 21: TDS Performance

When the parachute has been safely deployed, the vehicle has slowed to subsonic speeds, and the heatshield has been jettisoned, the Terminal Descent System (TDS) attempts to acquire the ground. Once it does, the navigation filter may then determine the vehicle's altitude and velocity, relative to the surface of Mars. GN&C Mode 20, 'TDS NAV INIT', begins when the vehicle is ready to process TDS data (approximately 5 seconds after heatshield separation) and ends when the navigation filter converges on an altitude solution. At that point, the mode changes to Mode 21, 'MLE PRIMING LOGIC ENABLED'. The observed times in these two modes were 13.125 s (-1.45σ) and 62.5 s (2.25σ), respectively. When considered together, however, the combined time in both Mode 20 and Mode 21 was 75.625 s, which is very close

to the predicted mean and just 0.105σ. This indicates that the total time on parachute was very nominal. Therefore, the shorter than nominal time in Mode 20, followed by a longer than nominal time in Mode 21, is due to a better than modeled maximum range capability of the TDS.

This result is not unexpected, however. During the MSL development, two TDS models were created. The first was a low-fidelity engineering model, referred to as the "level-1" model. The second was a high-fidelity, physics-based model, referred to as "Sulcata" [**?**]. The level-1 model, which was used in the baseline simulation runs, was kept intentionally conservative. This conservatism was appropriate, since the new-development TDS had yet to be flown at Mars. In Contrast, the Sulcata model results were anchored to actual TDS field test measurements and was, therefore, the best estimate of TDS performance. A direct comparison of the two models, within the EDL simulation, was performed prior to launch. This comparison confirmed that the level-1 model was conservative in its predictions and, more importantly, that no unforeseen performance issues were uncovered by the higher fidelity physics included in the Sulcata model. Figure 10, from this comparison, shows that the Sulcata model correctly predicted longer times in Mode 21.
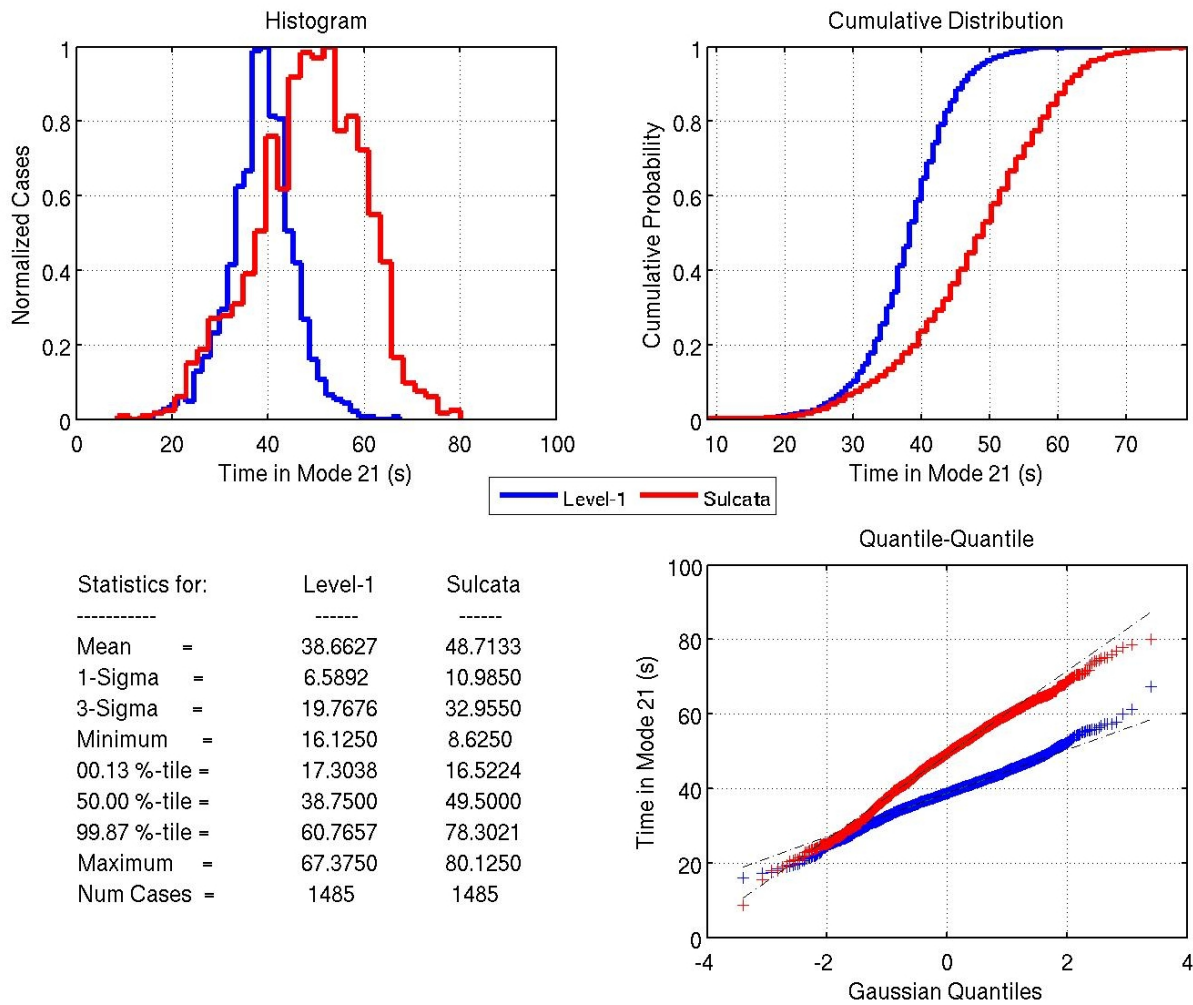
Figure 10. Sulcata vs. Level-1 TDS Model.

Statistics for:	Level-1	Sulcata
-----------	------	------
Mean =	38.6627	48.7133
1-Sigma =	6.5892	10.9850
3-Sigma =	19.7676	32.9550
Minimum =	16.1250	8.6250
00.13 %-tile =	17.3038	16.5224
50.00 %-tile =	38.7500	49.5000
99.87 %-tile =	60.7657	78.3021
Maximum =	67.3750	80.1250
Num Cases =	1485	1485

GN&C Mode 35: Gravity Anomaly

The final, and potentially most dangerous, of the timeline deviations was the time in GN&C Mode 35, 'PD READY FOR TOUCHDOWN'. This is the final mode prior to touchdown. The rover has been lowered on the BUD, the mobility has been deployed, the control system is maintaining a constant descent rate of $0.75 \ m/s$, and the system is waiting to sense contact with the ground. The observed time of $9.656250 \ s$ in this mode was a 2.58σ event with respect to the OD229 predictions.

The reason for this long time in Mode 35 was determined to be the effect of a navigated velocity error that arises from a mismatch between the on-board gravity model, a J2 model, and the actual local gravity field at the landing site. This mismatch created an acceleration error on the order of $450 \ \mu g's$, which resulted in a vertical velocity that was approximately $0.1 \ m/s$ slower than the desired $0.75 \ m/s$ [**?**]. In flight, the sign of this error (local gravity less than modeled) resulted in a softer than expected touchdown at the expense of a small amount of additional fuel use. However, an equivalent gravity error of the opposite sign could have resulted in a touchdown velocity at or beyond the $0.85 \ m/s$ limit, which would have placed a significant structural load on rover mobility components.

This sensitivity to small differences in the gravity model was not fully appreciated before flight. The gravity model utilized in the EDL simulation was a J3 model, even though a higher order, 85 by 85, spherical harmonic model was available. Post-landing, a comparison of these models showed a $1.32 \ s$ increase in the 99%-tile time in Mode 35 and a $0.08 \ m/s$ decrease in the 1%-tile low touchdown velocity, when using the higher order gravity model. The magnitudes of the gravitational acceleration predicted at *Curiosity*'s landing site by these models are provided in Table 3, along with an estimate from a topography-implied gravity model, called the Mars Gravity Model 2011 (MGM2011). This gravity field, developed at Curtin University's Western Australian Centre for Geodesy and shown in Figure 11, is a high resolution gravity field model for Mars that resolves features down to kilometer scales. Table 3 also provides an estimate of the local gravity as measured in-flight by the Rover Inertial Measurement Unit (RIMU), just after touchdown. This estimate was obtained by calculating an average acceleration over 467 frames of real-time RIMU telemetry, shown in Figure 12.

The cause for the EDL system sensitivity to gravity is well understood. Prior to Sky Crane, the navigation filter is able to process data from all six TDS beams in determining the

9

Table 3. Comparison of Mars Gravity Esitmates.

Source	Value	Units
J3 Model	3.7350	m/s^2
85 x 85 Model	3.7328	m/s^2
MGM2011 Model[3]	3.7133	m/s^2
RIMU Telemetry	3.7180	m/s^2

Table 4. EDL Requirements and Constraints (Partial).

Parameter	<=>	Value	Units
Landing accuracy	<	12.5	km
Peak entry deceleration	<	15	$Earth\ g$
Peak parachute opening load	<	65,000	lb_f
Peak attitude rate	<	120	deg/s
Peak angular acceleration	<	37.1	rad/s^2
Time in GN&C Mode 21	>	5	s

Gravity acceleration (MGM2011)
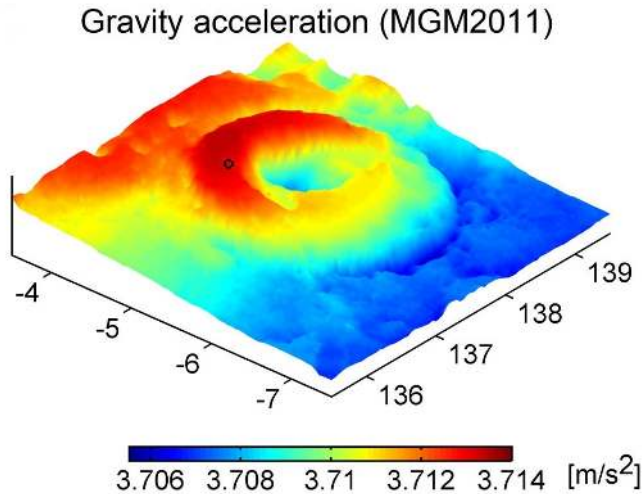
Figure 11. High resolution gravity field estimate at Gale Crater from Mars Gravity Model 2011. The estimated acceleration magnitude at Curiosity's landing site is 3.7133 m/s^2. Image Credit: Western Australian Centre for Geodesy, Curtin University.

spacecraft's 3-axis velocity vector, relative to the terrain. At rover separation, however, four of the six TDS beams are no longer used because of concerns that the TDS could lock-up on the suspended rover. With only the two remaining 50^o beams measuring velocity during Sky Crane, the navigation filter must infer the vertical component of velocity from measured IMU accelerations and an assumed gravitational acceleration. This sensitivity may, therefore, be designed-out of future missions by ensuring that three or more TDS beams are providing valid measurements all the way to touchdown.

5. WATERMARKS

Throughout mission design and operations, the EDL simulation is used to track statistics on pre-defined Monte Carlo output variables. These output variables are typically instantaneous flight conditions at specific events, but may also be minima or maxima of flight parameters experienced over specific regions of flight. The later are referred to as "watermarks" and they help ensure that design requirements and operational limits are not exceeded. By way of example, Table 4 provides a partial set of these requirements. For each Monte Carlo case, over 4800 output variables were captured. These outputs were used to auto-fill a Monte Carlo scorecard containing approximately 700 statistics covering all aspects of vehicle performance.

---

[3]http://geodesy.curtin.edu.au/research/models/mgm2011/

It is desired to evaluate the effectiveness of the EDL simulation, by comparing telemetered flight data to these 700 or so simulation predictions. However, many of these scorecard items are not directly observable in the flight telemetry (e.g. Mach number at parachute deploy, dynamic pressure at heatshield separation, attitude error at rover separation, etc.). Those parameters must be obtained through detailed flight reconstructions. As a quick-look into simulation performance, the following sub-sections present a few scorecard items that can be directly observed in the flight telemetry: peak entry deceleration, peak parachute deceleration, peak attitude rate, and peak angular acceleration. Additional comparisons will be made as reconstructed data becomes available. [**?**]

Peak Entry Deceleration

Entry vehicles, entering the atmosphere at hypersonic velocities, can experience heavy decelerations. The structure of these vehicles must be rated to withstand the loads that are generated. As noted in Table 4, the MSL system was rated for a maximum of 15 $Earth g's$ of deceleration. Figure 13 plots the magnitude of the sensed acceleration vector, as measured by the Inertial Measurement Unit (IMU). This data, taken from the NAV filter telemetry, and recorded at 64 Hz, was low-pass filtered on-board the spacecraft using a Tustin bilinear digital filter. [**?**] This figure indicates two deceleration peaks. The first, and highest peak, indicated by the green circle, is the peak entry deceleration experienced by the blunt aeroshell. The second peak, indicated by the red diamond marker, occurs at full inflation of the parachute. It is desired to compare these values to simulation predictions.

Peak entry deceleration occurred at approximately 620.26 s following TZERO (80.26 s after entry interface). Figure 14 compares the measured value of 12.609 $Earth\ g's$, indicated by the green circles, with EDL simulation predictions from the OD229 Monte Carlo. It can be seen on the histogram in the left half of Figure 14 that the observed flight data is only slightly lower than the simulation mean. The quantile-quantile plot in the right half of Figure 14 further shows that the actual flight value was well within one standard deviation at -0.471σ.

Peak Parachute Deceleration

The peak parachute deceleration, shown in Figure 13, occurred at approximately 800.96 s following TZERO (260.96 s after entry interface). Figure 15 compares the measured value of 6.068 $Earth\ g's$, indicated by the red diamonds, with EDL simulation predictions from the OD229 Monte Carlo. It can be seen on the histogram in the left half of Figure 15 that, similar to the peak entry deceleration, the observed flight data is only slightly lower than the simulation mean. The quantile-
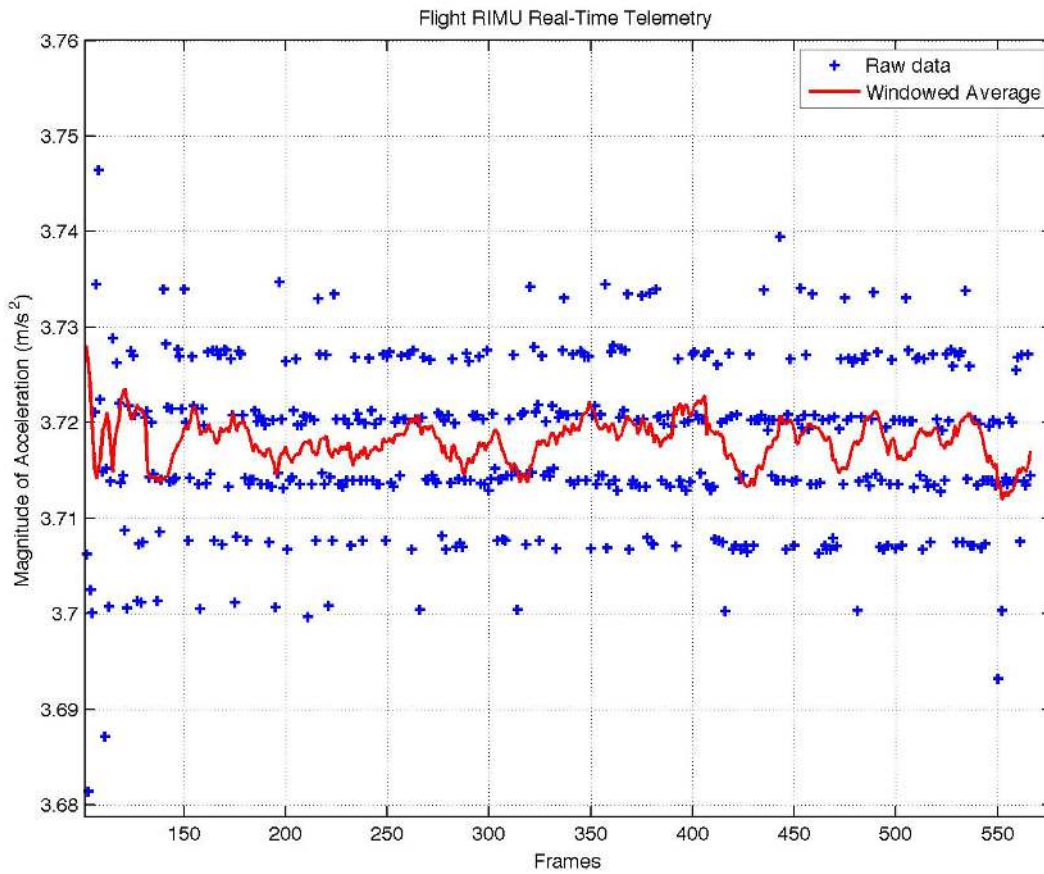
Figure 12. Real-Time RIMU telemetry, magnitude of measured acceleration immediately following touchdown. Average acceleration over 467 RIMU frames is 3.7180 m/s^2.
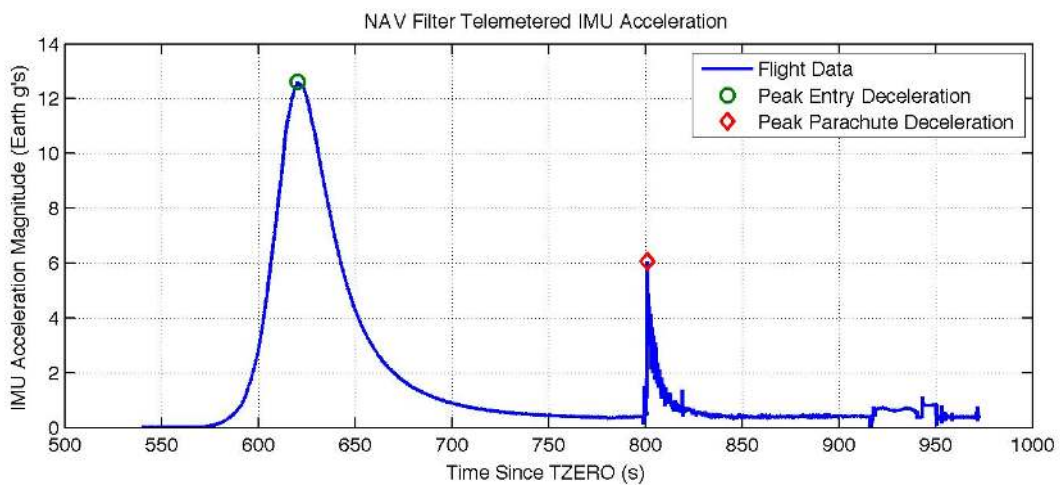

Figure 13. Sensed acceleration magnitude, at the IMU location, vs. time since TZERO, from NAV Filter telemetry.
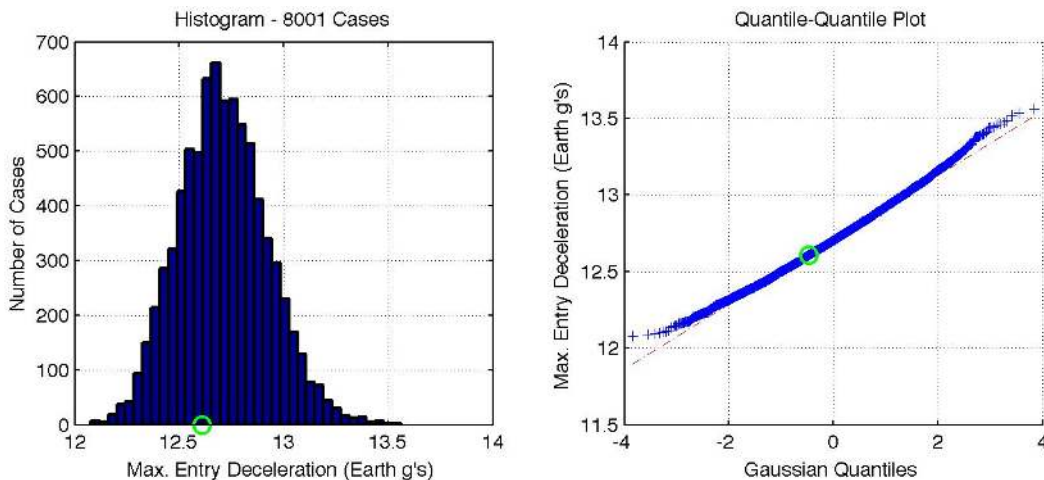
11

Figure 14. Comparison of Peak Entry Deceleration to Monte Carlo Predictions from OD229.

quantile plot in the right half of Figure 15 further shows that the actual flight value was again within one standard deviation at -0.806σ.

Peak Attitude Rate

Throughout entry, up to parachute deploy, an active control system is used to maintain attitude control. [**?**] After parachute deploy, however, large parachute forces (up to 65,000 lbs) are transmitted through the parachute suspension system to three attachment points on the backshell. These forces can generate large torques on the entry body, which in turn, can pump significant amounts of rotational energy into the capsule's wrist mode, which is the pitch and yaw rotations of the capsule about its center of gravity. Analytically estimating and bounding the resulting oscillatory behavior of capsule, while suspended underneath a parachute, is an extremely dynamic and complex problem. Even though wrist mode oscillations decay passively with diminishing parachute forces, the sensitivity of subsequent critical EDL events to high attitude rates led to the inclusion of active attitude control while on parachute [**?**].

Figure 16 plots the magnitude of the angular velocity vector, as measured by the Inertial Measurement Unit (IMU). This data, taken from the NAV filter telemetry, and recorded at 64 Hz, was low-pass filtered on-board the spacecraft using a Tustin bilinear digital filter. The 20 deg/s profiled slews at the three bank reversals (at times 612.875, 633.875, and 663.375 s) and SUFR (779.875 s) can be seen. Also visible are the smaller maneuvers at guidance start (585.875 s) and heading alignment (675.625 s). It is desired to compare the value of peak attitude rate, indicated by the magenta square, to simulation predictions.

Peak attitude rate occurred at approximately 804.14 s following TZERO (5.01 s after parachute deploy). Figure 17 compares the measured value of 69.39 deg/s, indicated by the magenta squares, with EDL simulation predictions from the OD229 Monte Carlo. It can be seen on the histogram in the left half of Figure 17 that the observed flight data is higher than the mean simulation prediction (53.391 deg/s). However, the quantile-quantile plot in the right half of Figure 17 shows that the actual flight value was still within one standard deviation at 0.964σ. This figure shows very good agreement between the simulation and flight data given the

difficulty in modeling parachute dynamics.

Peak Angular Acceleration

Figure 18 plots the magnitude of the angular acceleration vector. This data was derived by finite differencing the 64 Hz NAV filter measurements of the attitude rate. Figure 18 shows a peak angular acceleration, indicated by the magenta triangle, of 10.881 rad/s^2, which occurs at approximately 801.25 s (2.15 s after parachute deploy). Figure 19 compares the maximum observed attitude rate to EDL simulation predictions from the OD229 Monte Carlo. As seen in the quantile-quantile subplot this angular acceleration is within 0.216σ of the 50%-tile simulation prediction (10.370 rad/s^2). Again, this is an extremely good agreement given the difficulty and complexity of the problem.

Together, Figures 17 and 19 show that the wrist mode oscillations in flight were very nominal as compared to EDL simulation predictions. Both the peak attitude rate and acceleration were well below their limits (120 deg/s and 37.1 rad/s^2, respectively). A significant factor in keeping these wrist mode oscillations low was the EDL system risk-balancing that was possible at the Gale landing site. Because of the low landing site elevation, the nominal Mach number at parachute deploy was lowered from 2.05 to 1.70. This lowering of Mach number both reduced the expected severity of areal oscillations and the time at which the vehicle was exposed to them, above Mach 1.4.

6. CONCLUSIONS

The POST2 MSL end-to-end EDL simulation played a critical role in the successful landing of the *Curiosity* rover on Mars, August 5, 2012. Simulation predictions were used throughout the project lifecycle: to inform EDL design choices, to compare and certify candidate landing sites, to verify EDL system performance, to select GN&C parameters, and to evaluate TCM and EPU operational decisions. Considering the importance of these activities, it is evident that the EDL simulation must accurately, yet conservatively, predict EDL system behavior. It is desired, therefore, to evaluate the effectiveness of the EDL simulation by comparing flight data to pre-entry simulation predictions. This paper provided a quick-look at this comparison, using data available shortly
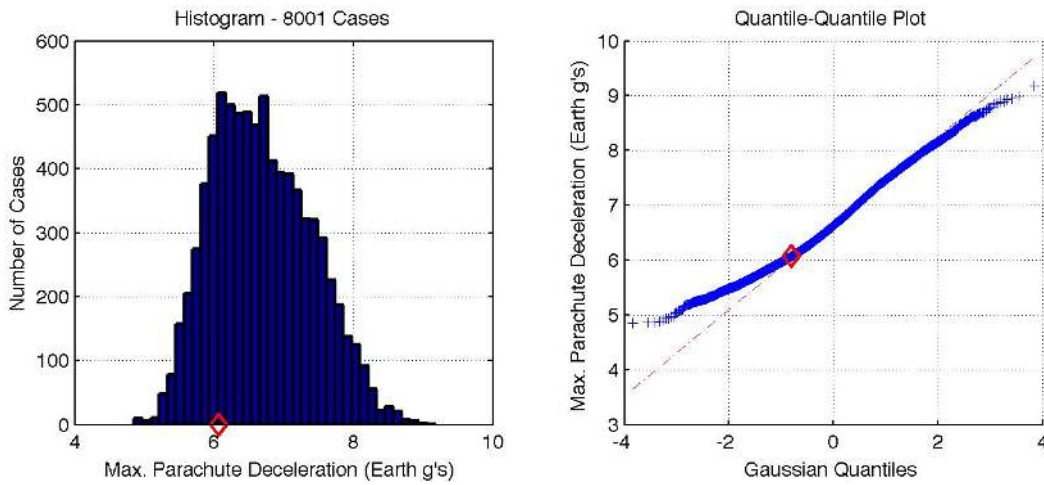
Figure 15. Comparison of Peak Parachute Deceleration to Monte Carlo Predictions from OD229.
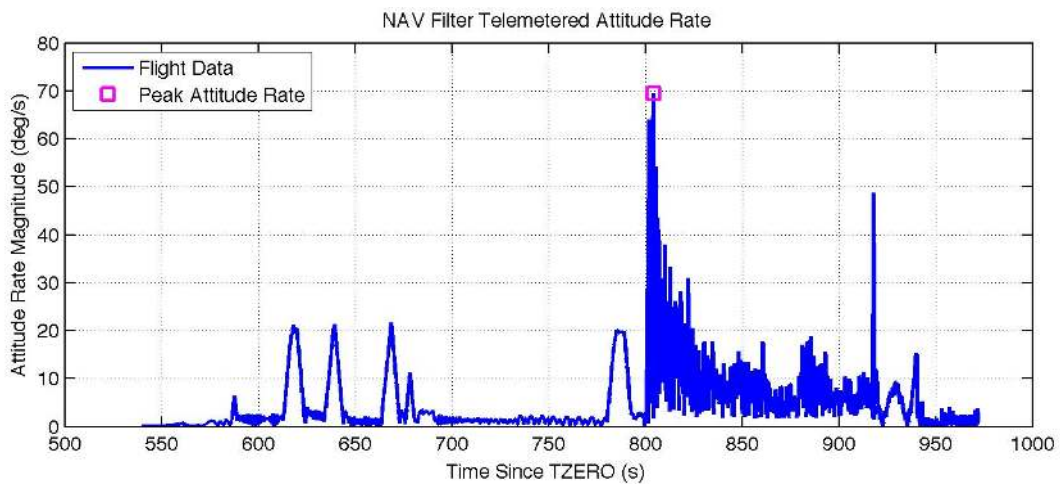

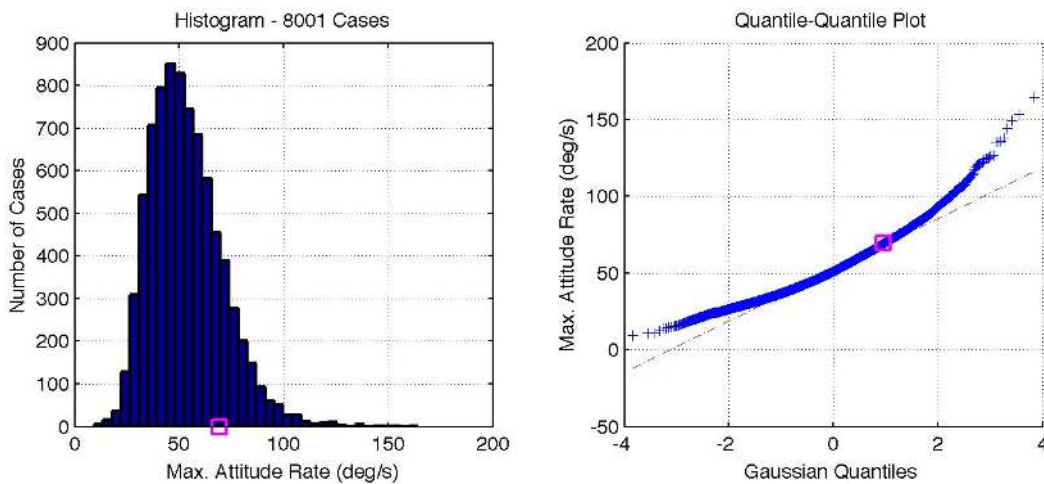**Figure 16**. Magnitude of attitude rate vs. time since TZERO, from NAV Filter telemetry.

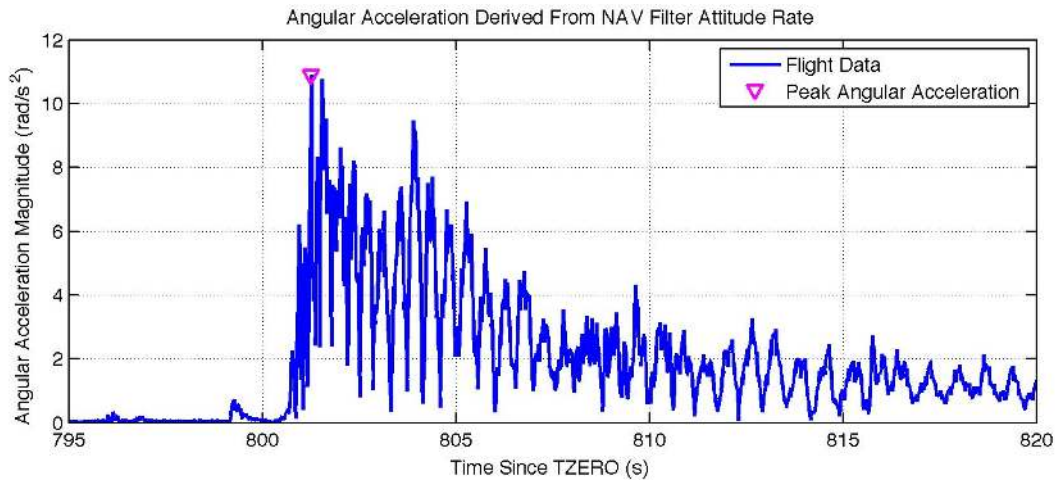**Figure 17**. Comparison of Peak Attitude Rate to Monte Carlo Predictions from OD229.

13

Figure 18. Magnitude of angular acceleration vs. time since TZERO, derived from NAV Filter attitude rate telemetry.
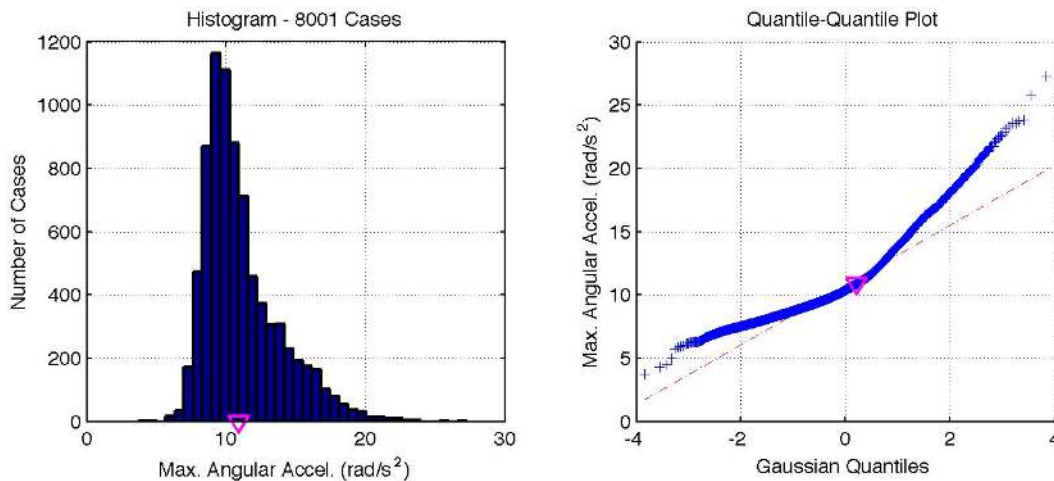

Figure 19. Comparison of Peak Angular Acceleration to Monte Carlo Predictions from OD229.

after landing. The following specific observations were discussed:

1. The predicted 99%-tile footprint for the OD229 navigation solution was 18.59 x 6.37 km. *Curiosity*'s actual landing site (4.5895 ^{o}S, 137.4417 ^{o}E) was only 2.385 km from the target, which represents a 0.74725σ or 23.747%-tile landing.

2. All Timeline Engine timepoints occurred within 6.9089 s of their predicted mean times. However, the cumulative time in several of the GN&C Mode Commander states exceeded +/- 1.5σ of the simulation results. These modes were investigated in more detail.

3. A very short -3.836σ time in GN&C Mode 8, 'WAIT FOR GUIDANCE START', was observed due to higher than expected deceleration at altitudes above 50 km. The immediate effect was for the guidance to command an initial bank angle 8.594817 deg away from the expected pre-bank angle. There were no lasting effects, however, as the vehicle recovered rapidly from this anomaly.

4. Low supersonic deceleration during the SUFR maneuver resulted in a combined 3.23σ time in GN&C Modes 14 and 15. During this time it is estimated that the vehicle lost approximately 340 m of altitude, which was of little conse-

quence due the ample altitude margin at Gale crater. Additional investigation into this anomaly is warranted, however, to better understand the expected aerodynamic performance in this flight regime.

5. TDS maximum range capability exceeded the conservative estimates of the level-1 TDS model. This lead to -1.45σ and 2.25σ times in GN&C Modes 20 and 21. However, the longer time in Mode 21 was correctly predicted by a more detailed physics-based model of the TDS, known as Sulcata.

6. A 2.58σ long time in GN&C Mode 35, 'PD READY FOR TOUCHDOWN', was determined to be due to an error in estimating the local gravitational acceleration at the landing site. In-flight, this error resulted in a softer than expected touchdown. However, had the local gravity been underestimated, instead of over-estimated, a maximum touchdown velocity of 0.85 m/s could have been exceeded. This EDL sensitivity to small errors in the gravity field may be designed-out of future missions by ensuring that three or more TDS beams are providing valid measurements all the way to touchdown.

7. As-flown values for peak deceleration during entry (12.609 $Earth\ g's$) and on parachute (6.068 $Earth\ g's$) were within -0.471σ and -0.806σ of simulation predictions,

14

respectively. Indicating very good entry and parachute model performance.

8. The peak attitude rate (69.39 deg/s) and angular acceleration (10.881 rad/s^2) watermarks were predicted well by the simulation at 0.964σ and 0.216σ, respectively. These indicated that the vehicle experienced very nominal wrist mode oscillations while on parachute. This is extremely good agreement, given the difficulty and complexity of modeling parachute dynamics in an EDL simulation.

APPENDIX

1. CUMULATIVE TIME IN MODE COMMANDER STATES.

Event times for Mode Commander state changes were extracted from 'MC_STATE_CHANGE' EVRs. These times were then differenced to get the cumulative time in each of the 38 GN&C Mode Commander states. In most cases, each GN&C mode was entered only once and the reported value is simply the difference in time between successive mode changes. However, mode 10, 'RANGE CONTROL TRACK', and mode 11, 'RANGE CONTROL BANK REVERSAL', were entered 4 and 3 times, respectively. For these two modes the difference in time between successive mode changes is summed over each time these modes were entered. These results are provided in Table 5.

ACKNOWLEDGMENTS

BIOGRAPHY

David Way received his Ph.D. in Aerospace Engineering from the Georgia Institute of Technology in 2001. Prior to that, he received a M.S. degree in Aerospace Engineering from Georgia Tech in 1997 and a B.S. degree in Aerospace Engineering from the United States Naval Academy in 1991. Following his commission, he served as an officer in the U.S. Navy's nuclear submarine force. Currently assigned to the Atmospheric Flight and Entry Systems Branch of NASA's Langley Research Center, His areas of expertise are flight mechanics, trajectory simulation, and entry system performance. Recently he served as LaRCs technical lead for the Mars Science Laboratory mission that successfully landed the Curiosity rover on Mars. In this role, he was responsible for MSLs prime Entry, Descent, and Landing simulation.

Table 5. Cumulative Time in GN&C Mode Commander State

MC State	GN&C Mode Name	Time in Mode (s)
0	IDLE	527.046875
1	OPEN LOOP STANDBY	0.000000
2	OPEN LOOP FIRE DRCS THRUSTERS	0.000000
3	OPEN LOOP SET MLE THROTTLES	0.000000
4	RCS WARM UP	20.500000
5	SPINDOWN	6.500000
6	TURN TO ENTRY	23.625000
7	WAIT FOR ENTRY	489.375000
8	WAIT FOR GUIDANCE START	45.875000
9	RANGE CONTROL SLEW TO COMMANDED BANK ANGLE	3.000000
10	RANGE CONTROL TRACK	54.250000
11	RANGE CONTROL BANK REVERSAL	32.500000
12	HEADING ALIGNMENT SLEW TO COMMANDED BANK ANGLE	5.125000
13	HEADING ALIGNMENT TRACK	99.125000
14	SLEW AND SUFR SLEW TO RADAR ATTITUDE	14.000000
15	SLEW AND SUFR WAIT FOR CHUTE DEPLOY	5.250000
16	ALLOW CHUTE DEPLOY TRANSIENTS TO SETTLE	10.015625
17	HEATSHIELD SEP LOGIC ENABLED	9.734375
18	RCS CONTROL INHIBITED	3.015625
19	WAIT FOR TDS NAV INIT	2.109375
20	TDS NAV INIT	13.125000
21	MLE PRIMING LOGIC ENABLED	62.500000
22	POWERED DESCENT START LOGIC ENABLED	14.875000
23	WAIT FOR BACKSHELL SEP	1.421875
24	FREE FALL	0.968750
25	MLE WARM UP	0.218750
26	PD DETUMBLE	0.156250
27	PD TURN TO INITIAL ATTITUDE	1.109375
28	PD POWERED APPROACH	21.656250
29	PD CONSTANT VELOCITY ACCORDION	3.046875
30	PD CONSTANT DECELERATION NOMINAL	7.281250
31	PD CONSTANT DECELERATION SATURATED AT MAX	0.000000
32	PD CONSTANT DECELERATION SATURATED AT MIN	0.000000
33	PD THROTTLE DOWN AND DAMP TRANSIENTS	2.531250
34	PD DEPLOY ROVER AND DAMP TRANSIENTS	8.968750
35	PD READY FOR TOUCHDOWN	9.656250
36	PD STOP VERTICAL MOTION	0.593750
37	PD ALTITUDE HOLD	0.203125
38	DONE	