A MASSIVELY PARALLEL BAYESIAN APPROACH TO PLANETARY PROTECTION TRAJECTORY ANALYSIS AND DESIGN

Mark S Wallace*

The NASA Planetary Protection Office has levied a requirement that the upper stage of future planetary launches have a less than 10^{-4} chance of impacting Mars within 50 years after launch. A brute-force approach requires a decade of computer time to demonstrate compliance. By using a Bayesian approach and taking advantage of the demonstrated reliability of the upper stage, the required number of fifty-year propagations can be massively reduced. By spreading the remaining embarrassingly parallel Monte Carlo simulations across multiple computers, compliance can be demonstrated in a reasonable time frame. The method used is described here.

INTRODUCTION

The next NASA Mars mission, the Interior Exploration using Seismic Investigations, Geodesy, and Heat Transport (InSight), is scheduled to launch in March 2016 using a United Launch Alliance (ULA) Atlas V 401 rocket with a Centaur upper stage. The project is carrying a requirement, levied by the NASA Planetary Protection Office, that the upper stage of the launch vehicle have a less than 10^{-4} chance of hitting Mars within 50 years after launch. This 10^{-4} probability includes the probabilities of an upper stage anomaly. Considering the nominal mission or a worst-case anomalous scenario is insufficient. The nominal Centaur mission is to separate from the InSight spacecraft at a specific attitude, slew to a second attitude, perform a Contamination and Collision Avoidance Maneuver (CCAM), slew to a third attitude, and dump the remaining LOX/LH2 propellant and burn the remaining hydrazine. This final step is known as the blowdown and hydrazine depletion sequence or simply "the blowdown." All three attitudes are available as controls to allow the designer to comply with the planetary protection requirements as well as the hyperbolic Earth departure targets. The anomalous scenarios considered by the InSight project are: failure to separate, failure to perform the CCAM, and failure to perform the blowdown. The failure of any given event ends the sequence. If the upper stage fails to separate, then neither the CCAM nor the blowdown will occur. Likewise, if the CCAM fails, the blowdown will not occur either.

After the separation-CCAM-blowdown sequence is completed or aborted due to failure, the upper stage will fly toward Mars. The initial uncertainty in the spacecraft state, expressed as a

* Mission Design Engineer, Mission Design & Navigation Section, California Institute of Technology/Jet Propulsion Laboratory, M/S 301-121. Pasadena, CA 91109

6×6 covariance matrix, called an injection covariance matrix, or ICM, is mapped forward to the time of closest approach using the state transition matrix. The resulting hyper-ellipsoid is then projected into the Mars B-Plane[1], illustrated in Figure 1. The Earth departure hyperbola is designed to ensure that the upper stage has a low probability of impacting Mars on that first encounter. The spacecraft must remove this “launch bias” in order to meet its mission objectives. The attitudes for Separation, CCAM, and Blowdown are designed to push the first Mars encounter further from Mars.

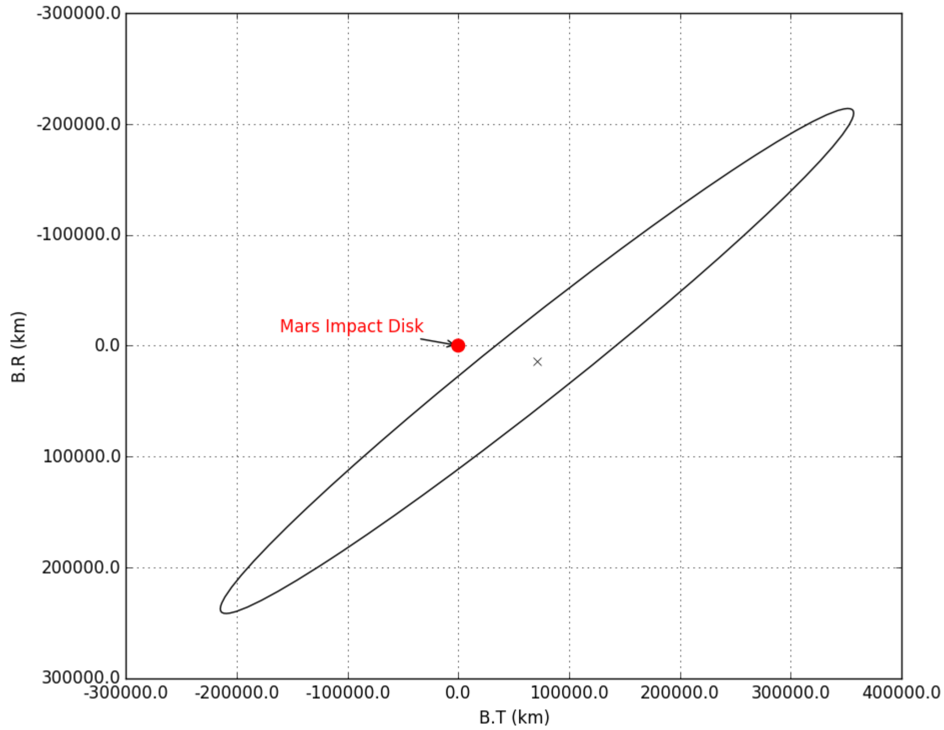

Figure 1: An ICM mapped to the first Mars encounter

After this first encounter, future Mars impacts can occur due to several factors. The first is that the flyby may place the upper stage on a direct resonant trajectory such that its orbital period and the Mars year are a rational fraction of each other. This generally requires that the flyby be relatively near Mars, as the ΔV required is on the order of 100s of m/s. The loci of B-Plane intersections that place the vehicle on a resonant trajectory follow arcs, called isochrones. These flybys result in nearly constant post-heliocentric orbital periods. Figure 2 highlights those points that reach the Mars sphere of influence[2] (574,000 km) sometime after the first flyby in an example Monte Carlo. Note that none of these cases impact Mars. In theory, the bias and attitudes could be designed to avoid these isochrones and eliminate the probability of impact via this mechanism.

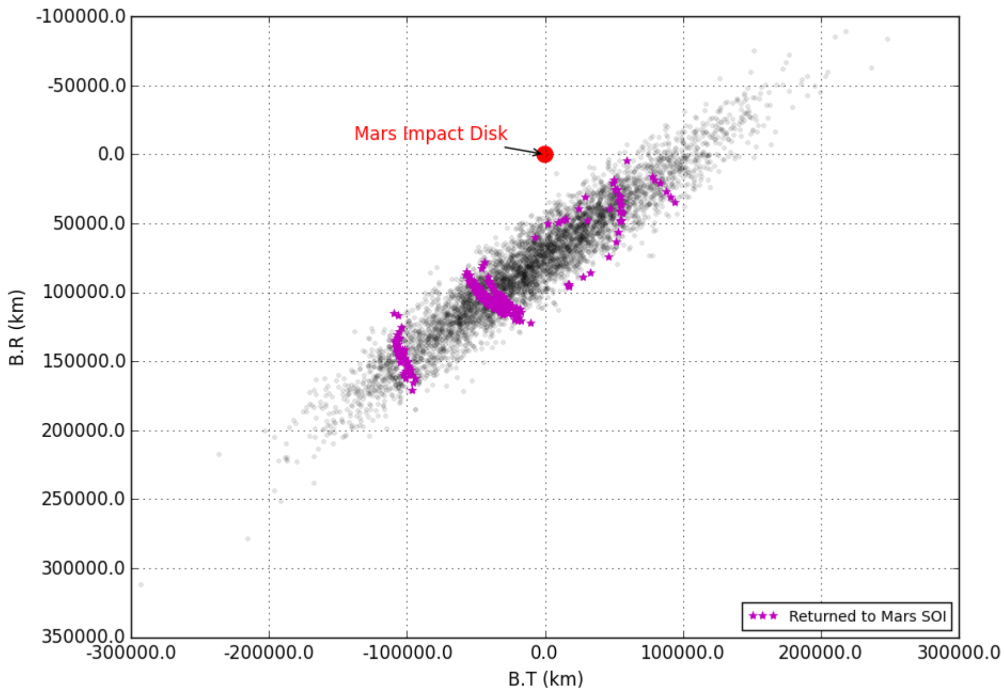

Figure 2: A sample 5000-case Monte Carlo with 256 trajectories returning to Mars

The second type of flyby that can place the upper stage on an impact trajectory is one that places the post-encounter heliocentric orbit in the Mars plane. If the spacecraft and Mars orbits are in the same heliocentric plane, then an impact could theoretically occur any place in Mars's orbit. Like the isochrones of the first mechanism, the locus of B-Plane points that can place the upper-stage on such a trajectory are relatively easy to find and can theoretically be avoided.

Finally, given that the solar system is not comprised solely of the sun and a massless Mars, a third mechanism exists that could place the upper stage on an impact trajectory. A later encounter with the Earth or any other planetary body could change the upper stage's trajectory and lead to a future impact. Exactly that sort of trajectory was observed in one 50,000 case Monte Carlo, illustrated in Figure 3. In this figure, the cyan points returned to the Earth sphere of influence[2] (924,000 km), the magenta stars returned to the Mars sphere of influence as a result, and a single red star impacted Mars. Like the isochrones shown in Figure 2, future Earth encounters could theoretically be mapped and avoided. The net effect of these three mechanisms, coupled with the need to design a single set of attitudes for every day in the launch period, means that the ΔV cost to the spacecraft of entirely avoiding all of these loci is excessive. A rough estimate indicated that just avoiding the direct resonance isochrones would require an order-of-magnitude increase in the launch bias.

A method was required to determine an upper bound on the future probability of impact via Monte Carlo analysis.

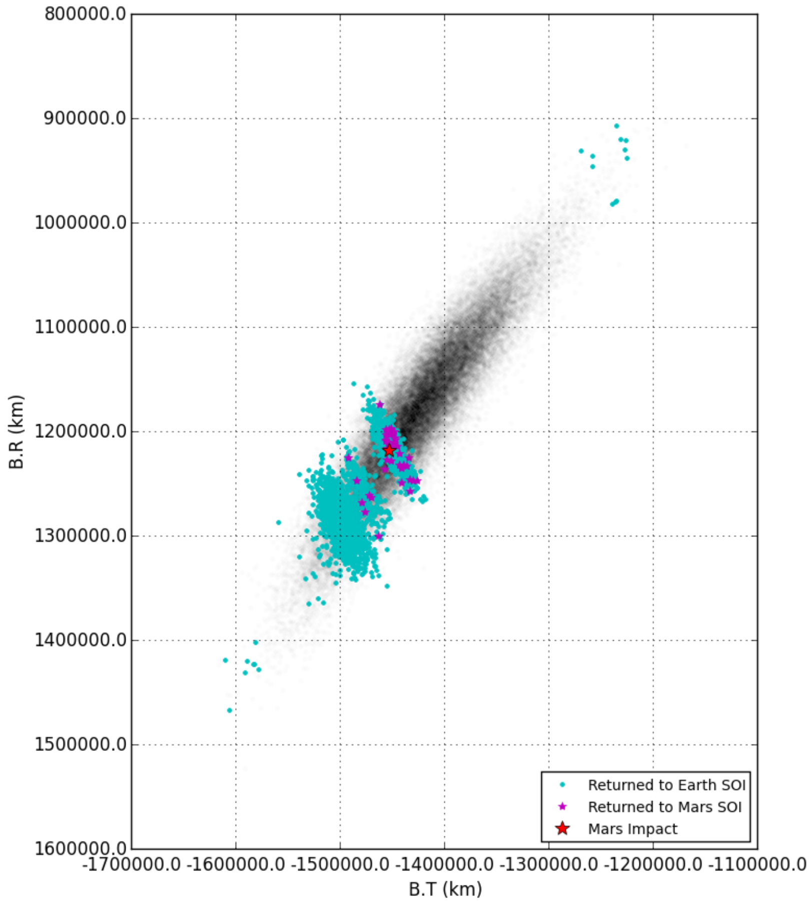

Figure 3: Distant flybys of Earth changed the trajectory enough to cause a future Mars encounter

STATISTICAL FORMULATION

In statistical parlance, the each step in the Centaur sequence and any impacts can be denoted as events which have a probability of occurrence. In this document, a successful Centaur/InSight separation event shall be denoted as s, and a failure shall be denoted as $\bar{s}$. Similarly, a successful CCAM, c, and a successful blowdown, b, with failures in either being indicated by $\bar{c}$ and $\bar{b}$, respectively. The probability of an event, *e,* is denoted by *P(e)*. The probability that an event, *e*, occurs given that event, *y*, has already occurred is *P(e|y)*. Similarly, the probability of *e* occurring given that both *y* and *z* have occurred is *P(e|y,z)*. By definition, the probability an event *e* does not occur is equal to one minus the probability it does occur, or $P(\bar{e}) = 1 - P(e)$. Thus, the total probability of impact, event *I*, can be expressed as:

$$P(I) = \begin{matrix} P(\bar{s})P(I|\bar{s}) + \\ P(s)P(\bar{c})P(I|s,\bar{c}) + \\ P(s)P(c)P(\bar{b})P(I|s,c,\bar{b}) + \\ P(s)P(c)P(b)P(I|s,c,b) \end{matrix} \tag{1}$$

The individual probabilities of impact given some sequence of events (z) can be calculated in two parts: the probability of impact on the first encounter plus the probability of impact during the ensuing 50 years given that no impact occurred on the first encounter, or Equation 2, where I_0 is a first-encounter impact and I_{50} is a later impact:

$$P(I|z) = P(I_0|z) + P(\overline{I_0}|z)P(I_{50}|\overline{I_0}, z) \tag{2}$$

Equations (1) and (2) can then be combined for the total probability of impact:

$$P(I) = \begin{array}{c} P(\bar{s})[P(I_0|\bar{s}) + P(\overline{I_0}|\bar{s})P(I_{50}|\overline{I_0}, \bar{s})] + \\ P(s)P(\bar{c})[P(I_0|s, \bar{c}) + P(\overline{I_0}|s, \bar{c})P(I_{50}|\overline{I_0}, s, \bar{c})] + \\ P(s)P(c)P(\bar{b})\left[P\left(I_0|s, c, \bar{b}\right) + P\left(\overline{I_0}|s, c, \bar{b}\right)P\left(I_{50}|\overline{I_0}, s, c, \bar{b}\right)\right] + \\ P(s)P(c)P(b)[P(I_0|s, c, b) + P(\overline{I_0}|s, c, b)P(I_{50}|\overline{I_0}, s, c, b)] \end{array} \tag{3}$$

Thus stated, the problem of determining the total probability of impact for a given design is reduced to determining each of the individual terms.

FIRST-ENCOUNTER IMPACTS

The problem of determining the impact probability on the first encounter is a solved one. With the ICM mapped into the B-Plane as an ellipse, the probability that the spacecraft state then lies within the impact radius of Mars has an analytical solution.[3] If the B-Plane ellipse is described by the 2×2 covariance $\mathbf{\Lambda}$, the impact radius of Mars has radius R at the location (u_0, v_o) relative to the center of the ellipse, and the vector $\boldsymbol{u}$ is described by the elements (u, v), then the impact probability is the double-integral:

$$P(I_o|z) = \frac{1}{2\pi\sqrt{\det(\mathbf{\Lambda})}} \iint \exp\left(-\frac{1}{2}\boldsymbol{u}^T\mathbf{\Lambda}^{-1}\boldsymbol{u}\right) dA \tag{4}$$

over the region:

$$(u - u_0)^2 - (v - v_0)^2 = R^2 \tag{5}$$

where the impact radius, R is defined in terms of the hyperbolic semi-major axis, a, and the physical planetary radius, r, as:

$$R = r\sqrt{1 - \frac{2a}{r}} \tag{6}$$

So long as the initial uncertainty can be mapped forward to a time of closest approach linearly, this method is sufficient. It is a standard practice in deep space navigation to use this linear mapping through a deep space cruise, so all of the first-encounter impact terms of Equation (3) can thus be determined analytically, as implied by the notation in Equation (4). However, this method fails if the linear mapping is insufficient to capture the evolution of the uncertainty.

Unfortunately, the linear mapping fails if a sufficiently large and sufficiently uncertain ΔV is applied to the trajectory, such as one which occurs during a planetary encounter. Even absent a subsequent flyby, the simple passage of time leads to non-Gaussian uncertainties with a Cartesian mapping. For example, the resulting semi-major axis uncertainty leads to an along-track position uncertainty that eventually grows such that the initially elliptical position uncertainty becomes distinctly nonlinear, as illustrated in Figure 4. In this figure, an initial uncorrelated 10 km and 1 m/s uncertainty is applied at the perihelion of a simple coplanar Earth-Mars Hohmann transfer ellipse. The resulting trajectories are then propagated both with the state transition matrix (blue ellipses) and with a simple Keplerian propagator (red crosses). After a half-revolution (left panel), the linear propagation remains a reasonable approximation of the non-linear results. However, after a single revolution (right panel), the linear approximation is clearly no longer able to capture the uncertainty in the spacecraft position as the heliocentric semi-major axis uncertainty has spread out the locus of trajectories along the orbit. However the Keplerian time to periapsis remains Gaussian effectively indefinitely as illustrated in Figure 5.

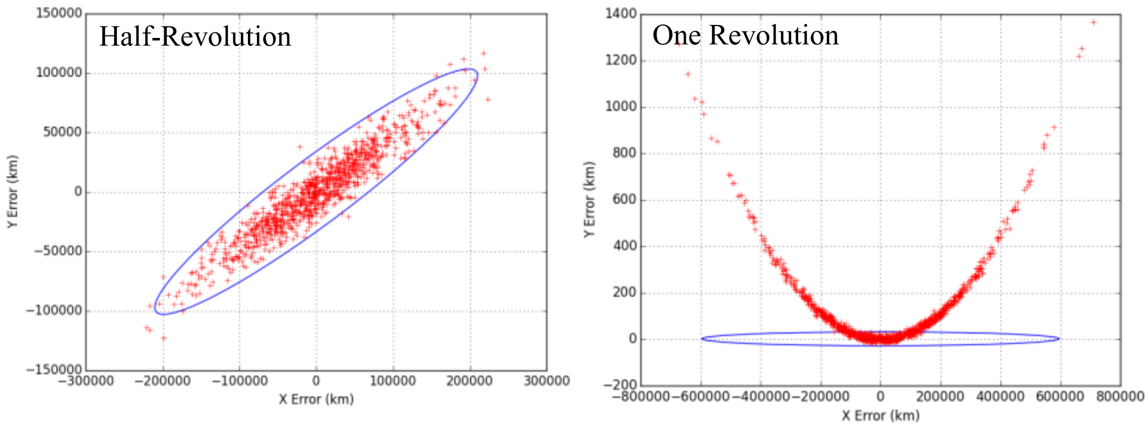

Figure 4: The passage of time causes the initially Gaussian distribution to become non-linear

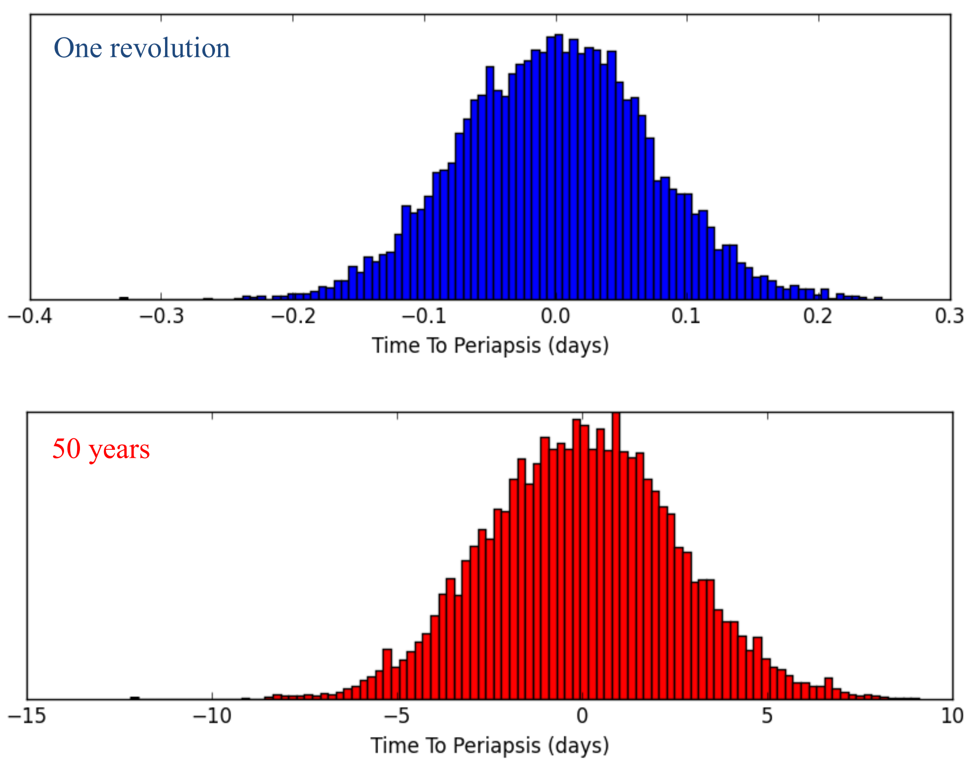

Figure 5: After a revolution (top, blue) and 50 years (below, red), the Keplerian time to periapsis remains Gaussian in the simple two-body system

However, when real dynamics are used, the variances of the Keplerian elements grow dramatically and even the Keplerian covariance matrix is no longer sufficient to capture the uncertainty in the future position. Consider the example case from Figure 3, where a future Earth encounter resulted in a Mars impact. The time history of the ranges to Earth and Mars are shown in Figure 6 in blue and red, respectively, for the nominal trajectory. As you can see, the first Mars encounter is a very distant one, 1.9 million km, and it never gets closer than 2.5 million km from Mars afterwards. Figure 7 illustrates the failure of the linear covariance mapping to accurately represent the uncertainties. In this plot, the red lines are the linearized 3σ bounds on the heliocentric semi-major axis, while the blue lines are 100 samples of a Monte Carlo. Both are taken from the same ICM. Through the 10 million km Earth encounter 31 years after launch, the linearized 3σ bounds contain the non-linear Monte Carlo results. Unfortunately, that Earth encounter causes the non-linear results to diverge greatly. That the heliocentric semi-major axis would be perturbed at this time was predicted by the linear approximation, but not the scale of it. This encounter, in a few cases, perturbed the trajectory enough that subsequent encounters further perturbed the heliocentric semi-major axis at times not predicted by the linear method. The distribution of heliocentric semi-major axes after 50 years is clearly out of family with the linear distribution, leaving a Monte Carlo analysis as the only remaining method of estimating the future state of the upper stage.

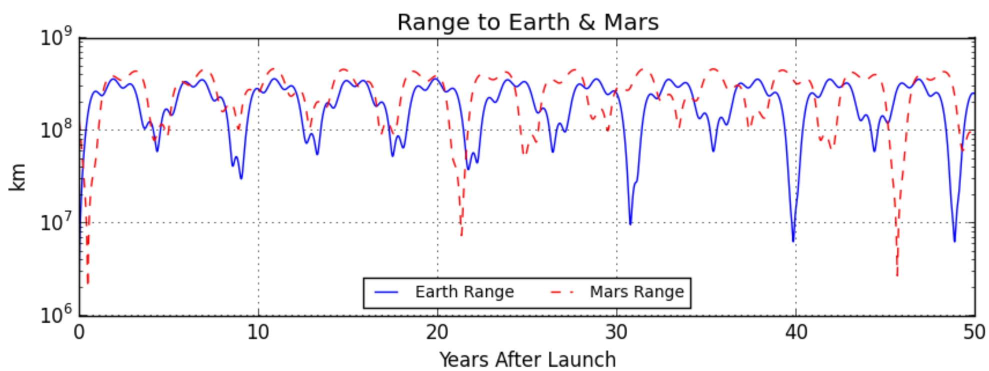

Figure 6: Time history of the nominal trajectory from Figure 3

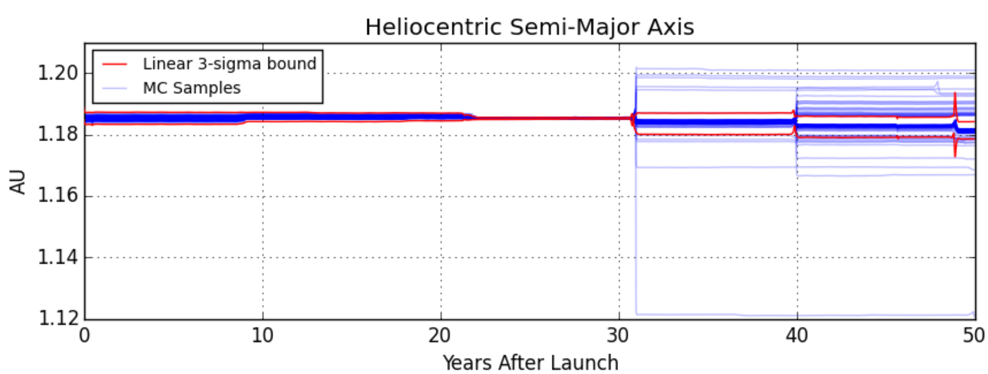

Figure 7: Comparison of the linear 3σ bounds (red) and Monte Carlo results (blue) for heliocentric semi-major axis

LATER IMPACTS

Because a Monte Carlo analysis is a simulation of a number of trials, the calculated probabilities are estimates subject to uncertainty. By treating any given propagation of the Monte Carlo as a Bernoulli process in which the propagation results in either an impact or a flyby, a confidence interval for the estimate of the likelihood of an impact can be determined. The method of calculating the Bernoulli confidence interval (*CI*) presented in textbooks is the normal approximation or Wald method[4], described in Equation (8), where the estimate of the probability of the event, $\hat{p}$, is the number of observed events, *x*, divided by the number of trials, *n*.

$$\hat{p} = {}^{x}/_{n} \tag{7}$$

$$CI = \left\{\hat{p} \pm z\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}\right\} \tag{8}$$

In Equation (8), the parameter *z* is the number of standard deviations between which some percentage of the normal distribution lies. For a 99% confidence interval, $z = 2.576$. Using that value, one can state that they have 99% confidence that the true probability of the event, *p*, lies within the confidence interval.

The Wald interval has two serious shortcomings, however. If no events are recorded, and $\hat{p}$ is thus 0, the confidence interval is always {0,0}. Further, if $\hat{p}$ is small enough, the lower bound of the confidence interval can be negative. The former is not a useful interval and the latter is non-physical. Either additional cases must be run until the confidence interval is useful and physical, or a different confidence interval must be used when $\hat{p}$ is small. One recommended but conservative interval is the Wilson interval[5]:

$$CI = \left\{\frac{n}{n+z^2}\left[\hat{p} + \frac{z^2}{2n} \pm z\sqrt{\frac{4n\hat{p}(1-\hat{p})+z^2}{4n^2}}\right]\right\} \tag{9}$$

If $\hat{p} = 0$, then the Wilson interval reduces to a convenient form:

$$CI|_{x=0} = \left\{0, \frac{z^2}{n+z^2}\right\} \tag{10}$$

Unlike other confidence intervals, this lower bound is identically zero, regardless of the desired level of confidence (99%, 95%, etc.). Since an impact may indeed be physically impossible, the lower limit of zero is desirable.

Using Equation (10), it can be determined that at least 66,352 Monte Carlo simulations would be required to meet a 10^{-4} probability of impact with 99% confidence. Unfortunately, this only holds if no impacts are found. If a single impact is detected in 66,352 runs, then Equation (9) yields a confidence interval of $\{0.0177 \times 10^{-4}, 1.28 \times 10^{-4}\}$. This would violate the requirement. To be robust against a single impact, at least 85,169 runs would be required under this metric. Without accounting for the probability of one of the four scenarios (three anomalous and one nominal) occurring, a brute-force approach is to perform 85,169 fifty-year propagations for any scenario likely to find direct-resonance isochrones and 66,352 fifty-year propagations for the others. If each scenario meets the 10^{-4} requirement, then the total requirement is met regardless of the probability of the upper stage events. However, this is a truly brute-force approach. A simple speed test of the Monte Carlo software used by JPL indicates that a single 50-year propagation-and-search requires about 12 seconds. To analyze the full suite of scenarios required in this manner would require almost 10 computer-years.

Examining Equation (3), the number of fifty-year propagations could be reduced by determining the probability that the Centaur successfully completes each step in its sequence. The reliability of the Centaur can be estimated from the historical flight record. The relevant data, as of early 2015, was 43 successes in 43 attempts. The 99% confidence intervals for success and failure are thus {0.866, 1} and {0, 0.134}, respectively. Since $P(\bar{s}) = 1 - P(s)$, the end-points of the confidence interval can be used (e.g. $P(s)$ = 0.866 and $P(\bar{s})$ = 0.134 or $P(s)$ = 1 and $P(\bar{s})$ = 0); however, it is not necessarily obvious which should be the worst-case, though they could be combined until a maximum value of $P(I)$ is determined. Even then, since the worst-case ends of the confidence intervals were stacked together, the resulting value would be larger than the 99% confidence interval and thus be excessively conservative, leading to more propagations and more computer time than necessary.

The question then becomes, what is the reliability of Centaur, and how certain are we of that estimate? A perfect record to date is no guarantee of future success, and a confidence-interval approach has been shown to be conservative. If a Bayesian approach is taken, instead of the frequentist approach of confidence intervals, then this question can be answered. The Beta distribution captures the distribution of the estimate probability of an event based on previous experience. The shape parameters are (α, β). It is recommended that $\alpha+\beta$ equal the number of trials, n, and α equal the number of events, x, recorded in those trials.[4] Both terms must be greater than zero, however. If is either parameter would be 0, then both parameters should be incremented by 0.5 (see Reference 5). The Beta distribution is defined on the domain $0 \leq p \leq 1$ and has the probability density function (PDF)[4]:

$$f(p; \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1} \tag{11}$$

and the cumulative distribution function (CDF):

$$F(p; \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \int_0^p t^{\alpha-1}(1-t)^{\beta-1} dt \tag{12}$$

where Γ is the gamma function:

$$\Gamma(t) = \int_0^\infty x^{t-1} e^{-x} dx \tag{13}$$

and the shape parameters are defined as:

$$\begin{aligned} \alpha &= \begin{cases} x, & 0 < \hat{p} < 1 \\ x + 0.5, & \hat{p} = 0 \; or \; 1 \end{cases} \\ \beta &= \begin{cases} n - x, & 0 < \hat{p} < 1 \\ n - x + 0.5, & \hat{p} = 0 \; or \; 1 \end{cases} \end{aligned} \tag{14}$$

Figure 8 illustrates the PDF of a Beta distribution for the probability of failure derived from the Centaur flight record, and Figure 9 illustrates the CDF. Close examination of the CDF reveals that the 99% confidence value of the probability of failure is 0.0738 by this metric, and not 0.134 as derived from the Wilson interval. This suggests that the brute force method briefly considered above was even more conservative than necessary. If 33,173 cases are run without any impacts, the 99th percentile value of the probability of impact is 10^{-4}. With a single impact, 46,051 cases would need to be run (refer to Figure 10) resulting in a dramatic reduction compared to the Wilson-interval derived values of 66,352 and 85,169. This alone is sufficient to reduce the total computer time from 10 years to 5.3 years.

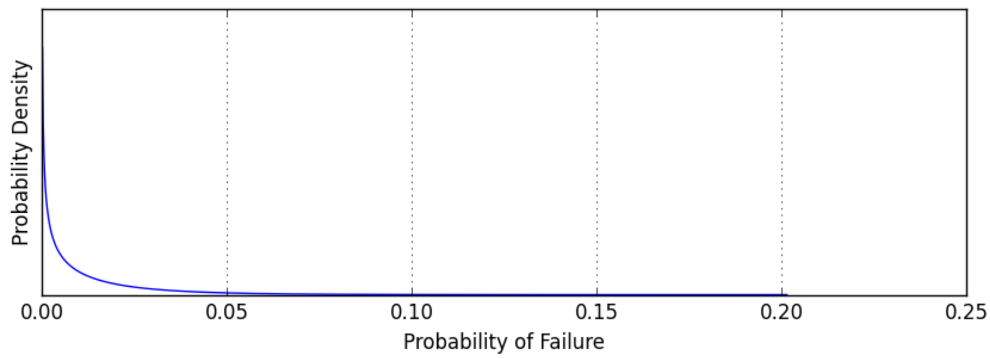

Figure 8: PDF of the estimate of the probability of a Centaur failure given no failures in 43 attempts

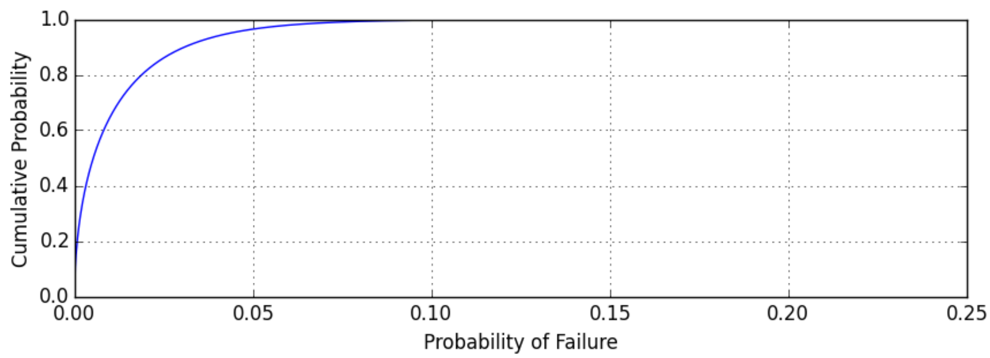

Figure 9: CDF of the estimate of the probability of a Centaur failure given no failures in 43 attempts

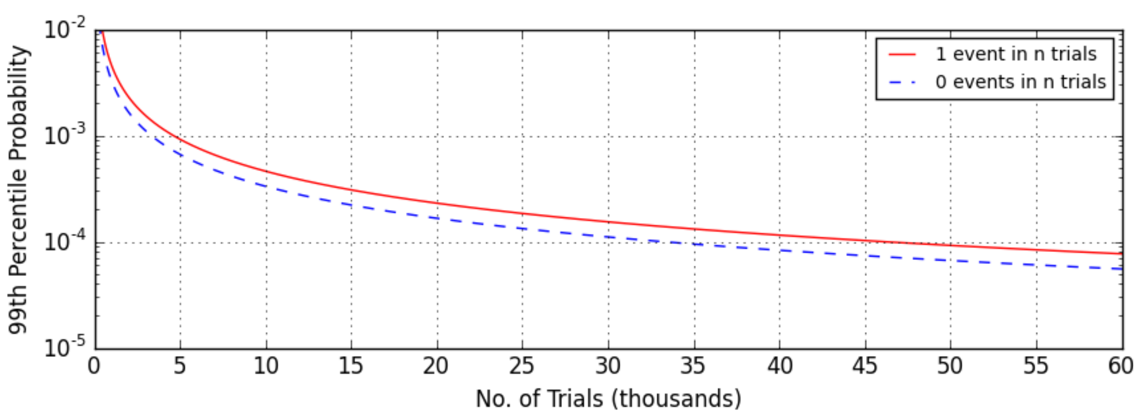

Figure 10: 99th Percentile probability of an event in a given number of trials

The only question remaining is how large of a Monte Carlo is needed for each scenario. If we assume that the Centaur reliability is 0.9, the first-encounter probabilities of impact are 10^{-4} for any of the anomalous scenarios, and that the probability of impact is zero for the nominal mission, then Equation (3) can be re-written to give an order-of-magnitude estimate of the maximum allowable 50-year probabilities:

$$P(I) \approx \mathbb{O}\begin{pmatrix} 3 \times 10^{-5} + 0.7[P(I_{50}|\overline{I_0}, s, c, b)] + \\ 0.1\left[P(I_{50}|\overline{I_0}, s, \bar{c}) + P(I_{50}|\overline{I_0}, \bar{s}) + P\left(I_{50}|\overline{I_0}, s, c, \bar{b}\right)\right] \end{pmatrix} \tag{15}$$

Equation (15) suggests that the anomalous-scenario probabilities of impact should be no higher than 10^{-3} and that the nominal mission probability of impact should be no higher than 10^{-4}. Such values would yield a probability of impact of 4×10^{-4}. If robustness against a single impact is desired, then 4,604 and 46,051 cases would be required. If not, then only 3,316 and 33,173 cases would be required. These are very rough estimates for the required number. The Centaur reliability is very likely to be greater larger than 0.9, putting more emphasis on the number of nominal-mission runs. The first-encounter probabilities of impact are likely to be much smaller than 10^{-4}, which reduces the relative importance of the anomalous-mission runs. Finally, this back-of-the-envelope estimation is a factor of four greater than the requirement, which suggests that more cases may be required.

DETERMINING THE TOTAL PROBABILITY OF IMPACT

Equation (3) can now be solved through a combination of analytical calculations of first-encounter impact probabilities, the Centaur flight record, and 50-year Monte Carlo propagations. These later two are used to feed a random number generator to produce many Beta-distributed values for the component probabilities. The method for analyzing a given set of targets and attitudes is thus a five-step process:

1. Determine the first-encounter probabilities of impact for each of the four scenarios (no-separation, no-CCAM, no-blowdown, and the nominal mission).
2. Perform 50-year Monte Carlos for each of the four scenarios and record how many cases impact Mars during the integration.
3. Generate one million samples of each of the seven Beta distributions representing the probability of each of the scenarios and the resulting probability of impact.
4. Solve Equation (3) one million times with each sample from the seven sets of Step 3 and the results from Step 1.
5. Sort the resulting values and determine the 99th percentile value.

For example, at the opening of the launch window on the first day in the InSight launch period, the first-encounter probabilities of impact were as in Table 1. Given the results of the Monte Carlos in Table 2, then the total distribution of the probability of impact would be as in Figure 11 and the cumulative probability would be as in Figure 12.

Table 1: Example First-Encounter Probabilities of Impact

Scenario	First-Encounter Impact Probability
Failure to separate	9.0×10^{-6}
Failure to perform CCAM	9.3×10^{-6}
Failure to perform the blowdown	1.4×10^{-8}
Nominal Mission	1.3×10^{-300}

Table 2: Example Results from Monte Carlos

Scenario	Monte Carlo Size	Number of Impacts	Beta Distribution
Failure to separate	5,000	0	β(0.5, 5000.5)
Failure to perform CCAM	5,000	0	β (0.5, 5000.5)
Failure to perform the blowdown	5,000	1	β (1, 4999)
Nominal Mission	50,000	1	β (1, 49999)

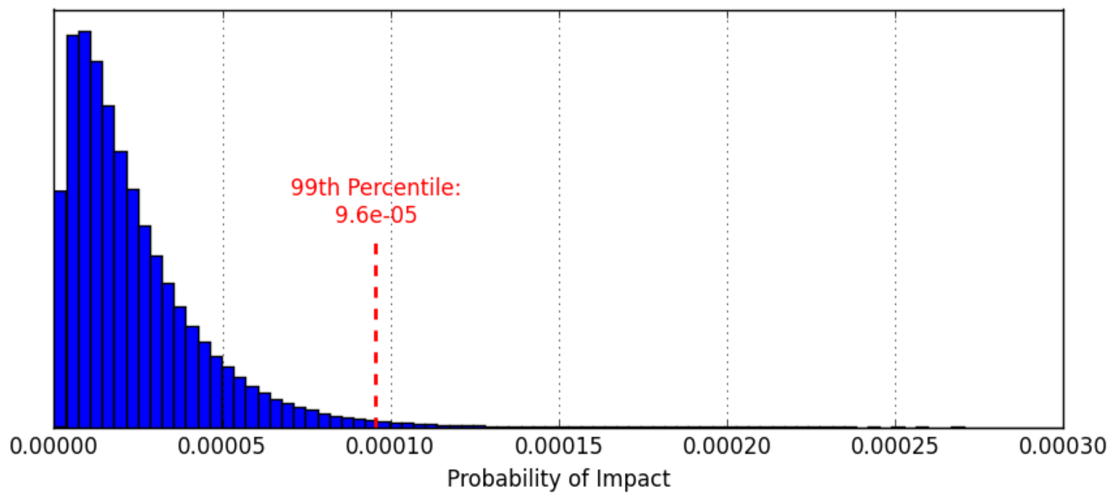

Figure 11: Example Distribution of the Total Probability of Impact

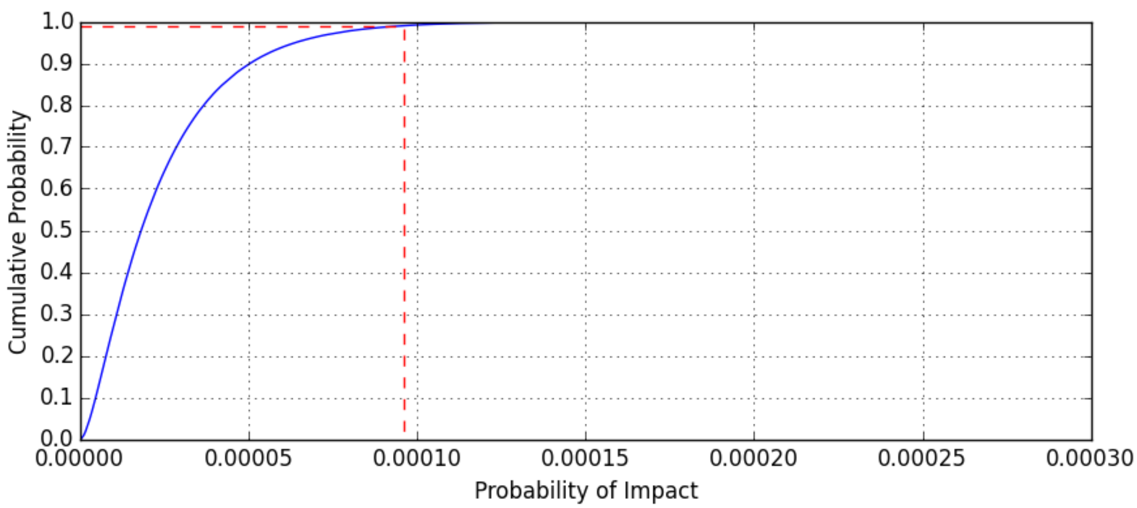

Figure 12: Example Cumulative Probability of Impact

APPLYING PARALLEL COMPUTING

The InSight mission's planetary protection plan states that compliance has been demonstrated if the 99th percentile estimate of the probability of impact is less than 10^{-4} for the open, middle, and close of the launch window on each day of the 27-day launch period, for a total of 81 cases. Within each case, there are four scenarios involving the Centaur upper stage successfully completing its mission. The Centaur must successfully separate from the InSight spacecraft, perform the collision and contamination avoidance maneuver (CCAM), and dump the remaining propellants. Thus, 324 separate cases must be examined via Monte Carlo.

Performing 81 50,000-case Monte Carlos and 243 5,000 case Monte Carlos, as was done for the example above, is a total of 8,265,000 propagations. At 12 seconds per propagation, it would require 731 days to demonstrate compliance with the planetary protection requirement. Fortunately, such a Monte Carlo application is the very definition of an embarrassingly parallel problem.[6] Each Monte Carlo initial condition is an independent draw from the injection covariance matrix. No propagation depends on any other propagation. By distributing the propagations across 776 CPUs, the run time can theoretically be reduced to just under 24 hours. There are other users on those systems, and this theoretical performance remains theoretical. In the real world, the run time was approximately 80 hours. Larger clusters would permit further reductions in the run time.

CONCLUSION

The Planetary Protection Office's requirement that the upper stage of all Mars-bound spacecraft have a less than one in ten-thousand chance of hitting Mars within 50 years after launch presents a unique challenge to the mission designer. While it may be possible to entirely eliminate the possibility of a future impact, doing so would come at an unacceptable cost to the mission propellant budget. The probability of impact must therefore be estimated.

Space is big and Mars is small; the probability of impact within 50 years is accordingly tiny. Textbook statistical methods for estimating the probability of an event given a random sample require that a relatively large number of events be recorded. The number of Monte Carlo samples necessary in such an approach is effectively unbounded. Using the more sophisticated Wilson interval places an upper bound on number of Monte Carlo runs required, but the problem remains intractably large. By taking advantage of the demonstrated reliability of the Centaur upper stage and applying the Beta distribution in a Bayesian approach, the problem of demonstrating compliance with the planetary protection requirement is reduced to a manageable size.

ACKNOWLEDGMENTS

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, under contract to the National Aeronautics and Space Administration.

The author would like to thank his colleagues on the InSight team and within the Mission Design & Navigation section at JPL, particularly Fernando Abilleira, Allen Halsell, Ralph Roncoli, Tomas Martin-Mur, Mike Wilson, Bill Taber, and Ted Drain. Thank you for your guidance, your encouragement, your tough questions, and your help. Gratitude is also owed to Caley Burke and the KSC team, as well as Marc Ginsburg, James Doehring, and the rest of the ULA team for their support in understanding the Centaur sequence of events.

REFERENCES

[1] Tapley, Schutz, and Born, Statistical Orbit Determination, Elsevier Inc., New York, 2004.

[2] Prussing and Conway, Orbital Mechanics, Oxford University Press, New York, 1993.

[3] Michel, J. R. "PQ/PQPLOT User's Guide, with Mathematical Explanations of Algorithms," JPL Interoffice Memo 312.G-03-017, July 30, 2003

[4] Hogg and Tanis, Probability and Statistical Inference, 6th ed., Prentice-Hal Inc., New Jersey, 2001

[5] Brown, Lawrence D., Cai, T. Tony, DasGupta, Anirban, "Interval Estimation for a Binomial Proportion" *Statistical Science*, 2001. Vol. 16, No. 2, pp 101-133

[6] Foster, Ian. Designing and Building Parallel Programs, Addison Wesley, Boston, 1995.