Engineering Risk Assessment of Space Thruster Challenge Problem

Donovan L. Mathias[*a], Christopher J. Mattenberger[b], and Susie Go[a]
[a] NASA Ames Research Center, Moffett Field, CA, USA
[b] Science and Technology Corp., Moffett Field, CA, USA

Abstract: The Engineering Risk Assessment (ERA) team at NASA Ames Research Center utilizes dynamic models with linked physics-of-failure analyses to produce quantitative risk assessments of space exploration missions. This paper applies the ERA approach to the baseline and extended versions of the PSAM Space Thruster Challenge Problem, which investigates mission risk for a deep space ion propulsion system with time-varying thruster requirements and operations schedules. The dynamic mission is modeled using a combination of discrete and continuous-time reliability elements within the commercially available GoldSim software. Loss-of-mission (LOM) probability results are generated via Monte Carlo sampling performed by the integrated model. Model convergence studies are presented to illustrate the sensitivity of integrated LOM results to the number of Monte Carlo trials. A deterministic risk model was also built for the three baseline and extended missions using the Ames Reliability Tool (ART), and results are compared to the simulation results to evaluate the relative importance of mission dynamics. The ART model did a reasonable job of matching the simulation models for the baseline case, while a hybrid approach using offline dynamic models was required for the extended missions. This study highlighted that state-of-the-art techniques can adequately adapt to a range of dynamic problems.

Keywords: PRA, Simulation, Dynamic PSA, Space Thruster Challenge Problem

1 INTRODUCTION

The Engineering Risk Assessment (ERA) team at NASA Ames Research Center utilizes dynamic models with linked physics-of-failure analyses to produce quantitative risk assessments of space exploration missions. This paper applies the ERA approach to the PSAM Space Thruster Challenge Problem [1]. The original challenge was presented at the PSAM8 conference in 2006 and has been expanded in scope for the current version. The problem centers around a deep space mission using an ion propulsion system. There are groups of redundant thrusters, as shown in Figure 1, with a time-varying operations schedule. In addition, the required number of thrusters varies with mission duration. The system includes propellant tanks and distribution lines as well.

Figure 1: Schematics of the PSAM Space Thruster Challenge Problem propulsion system [1].

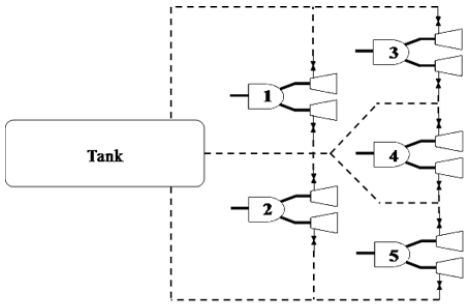

The extended version of the problem adds additional time-varying failure modes, such as time-varying leakage parameters, uncertainty in initial propellant load, and mission options that include the use of a "support station" for potential spacecraft repair. Full details of the problem can be found in Reference [1].

[*] Donovan.L.Mathias@nasa.gov

This paper describes the ERA modeling approach used to assess the baseline and extended versions of the problem. Results are presented in the form of the methodology comparison table provided in [1]. A simplified, static logic model was also created as a means of assessing the importance of dynamic modeling for this problem. These results are also included and conclusions about the applicability of the approach for this type of problem will be discussed.

2 MODEL OVERVIEW

The dynamic mission was modeled using a combination of discrete and continuous-time reliability elements within the commercially available GoldSim software [2]. The top level of the model includes the elements to control the mission timeline, evaluate the loss-of-mission (LOM) metric, and manage the inputs and outputs. In addition, there are specific elements to facilitate the exchange of information with the lower-level elements that represent the specific propulsion components and mission events. Figure 2 shows the layout of the model's top level.

Figure 2: Top-level view of the mission model.

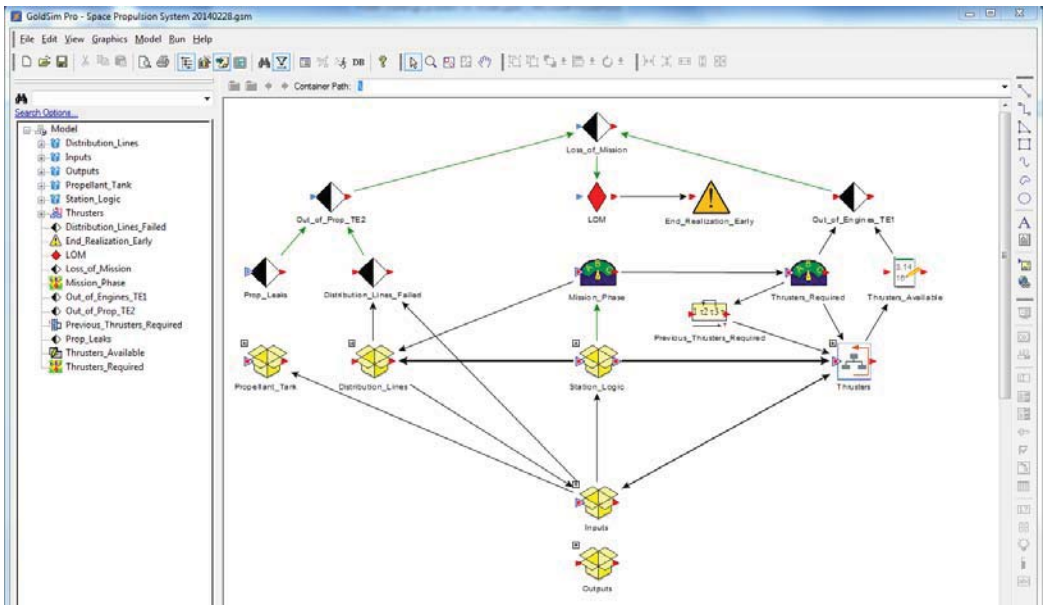

The yellow, box-like icons in Figure 2 represent containers for sub-elements of the model (similar to folders in a file structure). For example, the propellant tank container can be expanded to view the specifics of the tank model, as shown in Figure 3. The tank is represented using a reliability element populated with the appropriate failure rates/modes from the problem statement. Since the problem includes the potential for leakage, as opposed to just a simplified mission-ending failure, elements have been included to generate a leakage rate from the uncertainty distribution and track the propellant lost before the leak can be repaired. A global propellant element tracks propellant used nominally by the thrusters as well as any amount lost due to leaks. The top-level model triggers a loss of mission if the propellant is depleted prior to mission completion.

The propellant distribution lines are also represented using a reliability element, but additional components are required since the lines are also subject to damage accumulation. Figure 4 shows the resulting model, with the tree structure on the right side constructed to model the random walk process [3] per the problem statement [1]. The top of the tree is an integrator that tracks the accumulated damage while the Damage_DL element generates the random "step size" based on the input parameters supplied by the data elements below it. The Phase_Length element manages the timeline so that the appropriate damage rate is applied, including the additional leakage failure rate near the midpoint of the mission. As with the tank, the top-level model monitors the damage accumulation and triggers a LOM if the threshold (80,000 units) is exceeded.

Figure 3: Propellant tank and leakage model view.

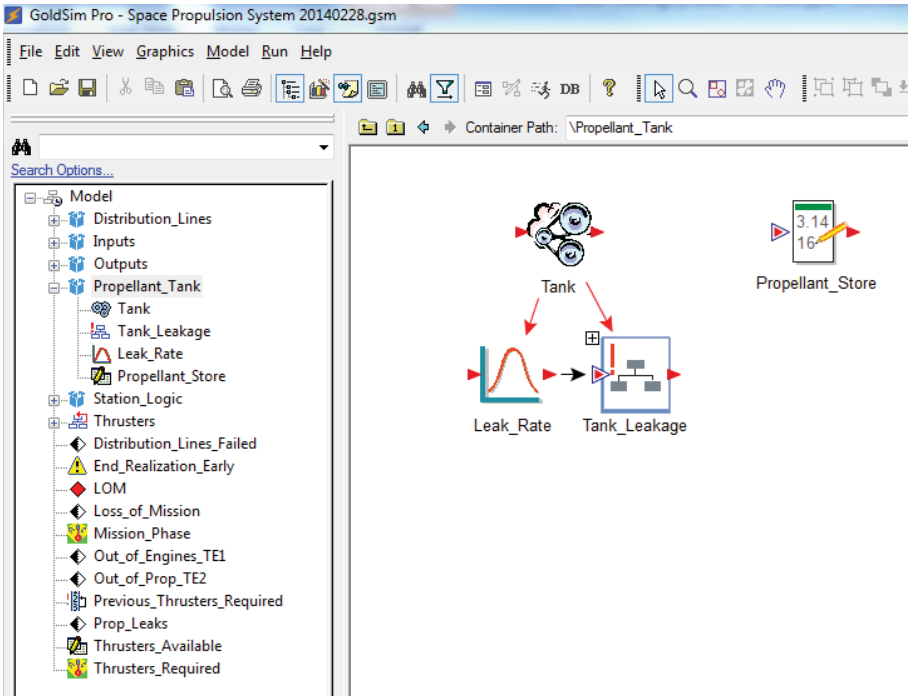

Figure 4: Propellant distribution line model.

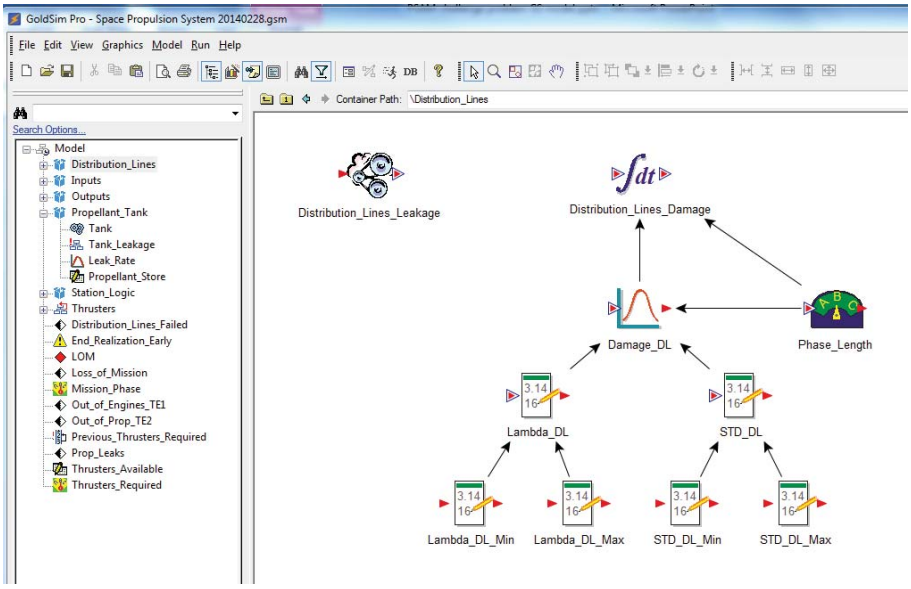

The most complicated part of the model involves the thrusters, not just because there are redundant strings, but because of the sequencing and switching logic used if a thruster fails. Figure 5 shows the top level of the thruster model. The individual thruster components are contained in the Thruster_n boxes and the arrows represent connectivity between model elements. The top elements exist to track the number of failed thrusters and control the thruster operation. In the current implementation, thrusters are utilized sequentially, always starting with Thruster 1 and cycling in order to the next available thruster. The bottom half of the figure shows a number of event triggers, which perform the common-cause failure (CCF) if an active thruster fails. The specific CCF logic was provided in the problem statement [1].

Figure 5: Top view of the thruster model.

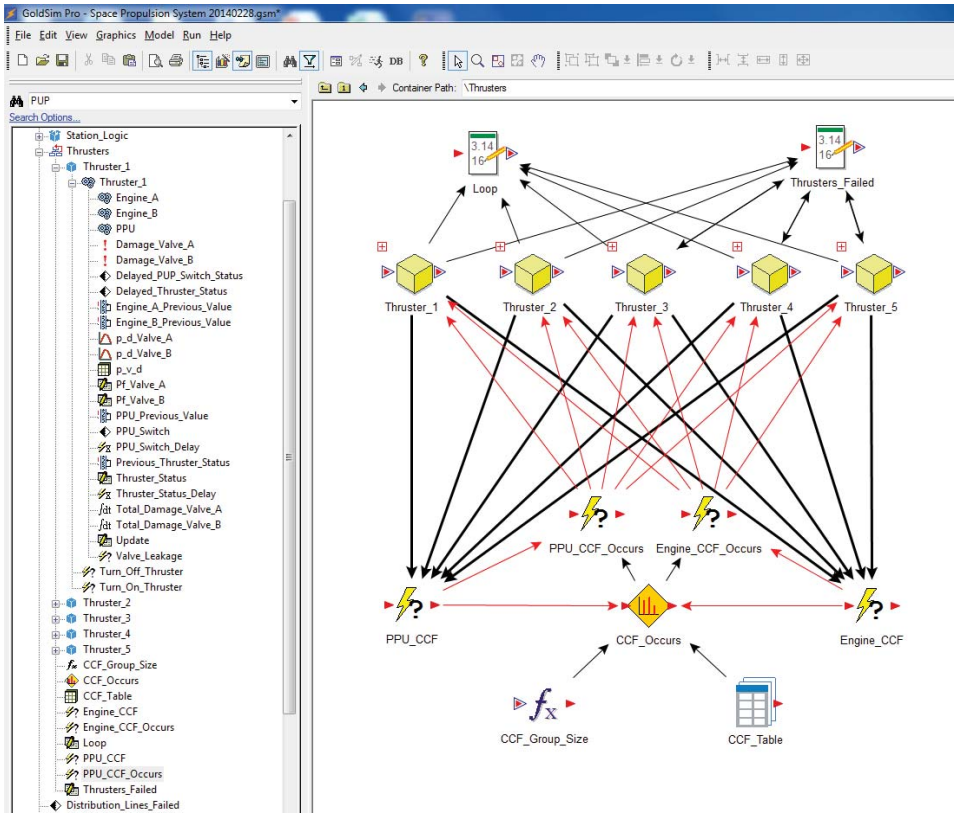

Expanding the Thruster_1 box gives a view of an individual thruster model, as shown in Figure 6. Each of the assembly components—e.g., the power processing unit (PPU), Engine_A, and Engine_B—is modeled as a continuous-time reliability element that accrues risk when it is active. The thruster model contains logic to nominally turn the appropriate engine on and off as well as to switch to Engine_B in the event that Engine_A fails. The valves are treated as discrete events and are exposed to demand risk upon actuation. Each valve actuation also includes additional risk due to a pressure oscillation with an uncertain impact. The bottom elements sample and integrate the pressure oscillation risk. Further description of the model components can be found in the methodology comparison table (Table 1) in the Appendix.

A simplified, static model was created using the Ames Reliability Tool (ART) for comparative purposes [4]. The ART logically combines the component failure rates and event failure probabilities to generate a deterministic loss-of-mission probability. Three ART models were constructed to illustrate the impact of simulation modeling. The basic problem was modeled with no treatment of dynamics by simply applying the failure rates across the appropriate operational modes. Discrete event failure probabilities were included through the number of event exposures. CCF values were applied using a simple beta model. The extended problem was first modeled with no dynamics by using fixed failure rates for the tank and distribution lines. Finally, a hybrid approach was constructed using two offline models to produce effective failure probabilities for the distribution line and tank failures. The extended problem ART models were run for cases with and without the support station.

The offline tank model used the mean propellant load and nominal burn rates to construct an effective propellant margin. Assuming a mean tank repair time and leak rate given a failure, the number of tank failures that would be expected to consume the margin was computed. This represented additional failure tolerance and was included as additional ‘cold spares’ in the reliability calculation.

Figure 6: A single thruster model.

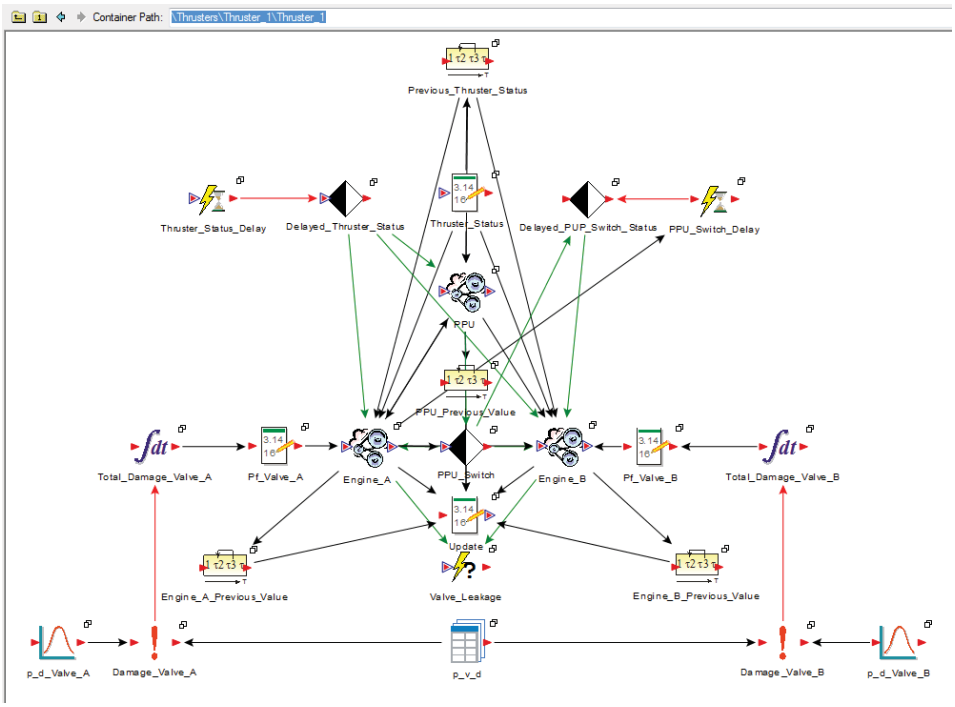

Line damage probability distributions were constructed by applying the random walk parameters (mean and standard deviation) across each of the mission phases. Monte Carlo samples were generated and combined to yield a distribution at the end of the mission along with the corresponding failure probability. This result was used as an input to the ART model.

3 RESULTS AND DISCUSSION

A convergence study was performed to determine the number of realizations required to converge the loss-of-mission estimate. Figure 7 shows the baseline mission LOM results for 100, 500, 1,000, 10,000, and 50,000 realizations with computed credibility intervals. LOM has settled to its final value by 1,000 realizations. However, the 10,000-realization results were utilized throughout the paper since they were available.

Figure 7: Convergence of baseline model LOM.

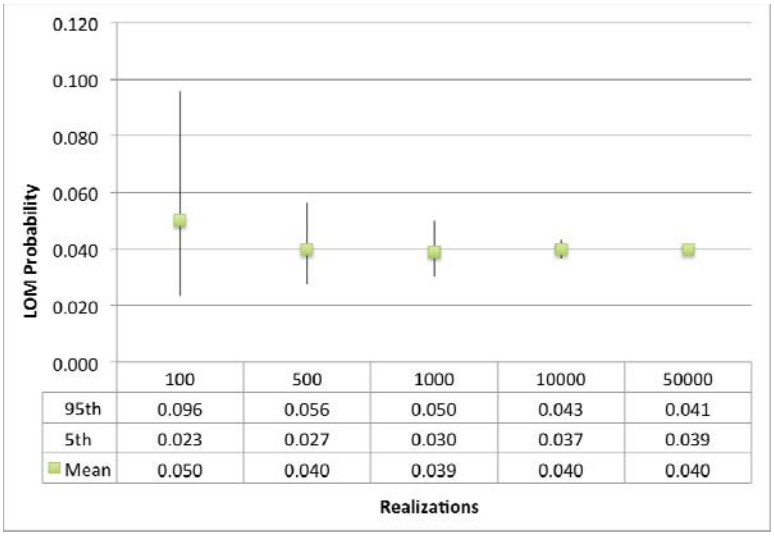

	100	500	1000	10000	50000
95th	0.096	0.056	0.050	0.043	0.041
5th	0.023	0.027	0.030	0.037	0.039
Mean	0.050	0.040	0.039	0.040	0.040

Figure 8 shows the time-distributed failure frequencies based on the number of observed failures for the baseline and extended mission options. The frequencies were computed by averaging the failure counts over the 10,000-realization simulation. The baseline mission risk accumulates relatively evenly throughout the mission, ending in a final value of 0.040. The risks for the extended missions, however, both approximately double around 50,000 hours and exhibit sharp increases beginning at about 70,000 hours. When the support station is not utilized, the risk increases by more than four times near the end of the mission.

Figure 8: Cumulative LOM probability for baseline, extended no-station, and extended with-station missions.

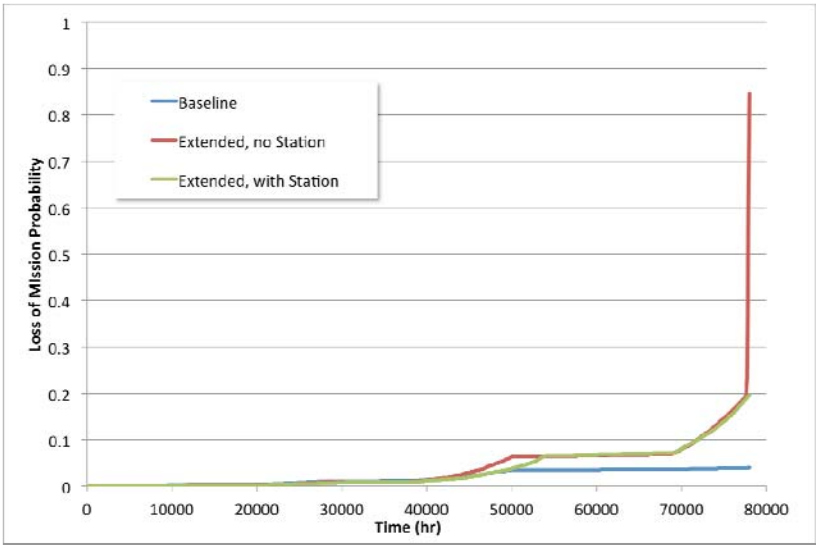

Figure 9 decomposes the curve for the no-support-station extended mission into its main components. In this case, the sharp increases in risk can be attributed to the leakage and distribution line damage. At approximately 38,000 hours, the slope of the risk curve changes—this corresponds to propulsion system operation and represents the region where leaks and nominal usage exhaust the propellant supply, resulting in significant LOM increase. The relatively less risk-intensive region, beginning at 50,000 hours for the no-station case and 54,000 for the with-station case (Figure 8), is due to the propulsion system becoming inactive during the coast phase and only random leakage events causing system failure. The slope increases again at 79,000 hours when the system is again activated and nominal propellant usage resumes. In addition, there is a large LOM spike that occurs in the last mission phase due to distribution line damage. The random walk model used to accumulate damage and determine LOM must exceed the value of 80,000 units before a failure occurs. Side model calculations indicate that the distribution line damage has a mean mission value of 80,100 units with a standard deviation of approximately 118 units. There is uncertainty in the random walk, but it still requires almost the entire mission before the critical level of damage occurs. A large number of missions fail in this way, but such failures always happen in the last 500 hours or so since that is how long it takes the damage parameter to accumulate to failure. Thus, the decision to stop at station and 'repair' the first 18,000 hours of damage effectively reduces the distribution line damage risk to zero.

In Figure 8, the mission utilizing the support station shows a reduced mid-mission risk accrual compared to the no-station case, and the observable difference is delayed by 18,000 hours due to the fuel that would have been replenished in the support-station case. However, due to stopping at the station, the engines are burned for longer durations and eventually the extra fuel consumption increases the risk intensity later in the mission. While propellant leakage and distribution line failures are driven by dynamic models, they are single-point failures and require no additional decomposition. While the "out of engines" end state is a relatively minor portion of the total LOM for the extended missions, the engines are of first-order importance for the baseline model.

Figure 9: Cumulative LOM for the extended mission, with no support station.

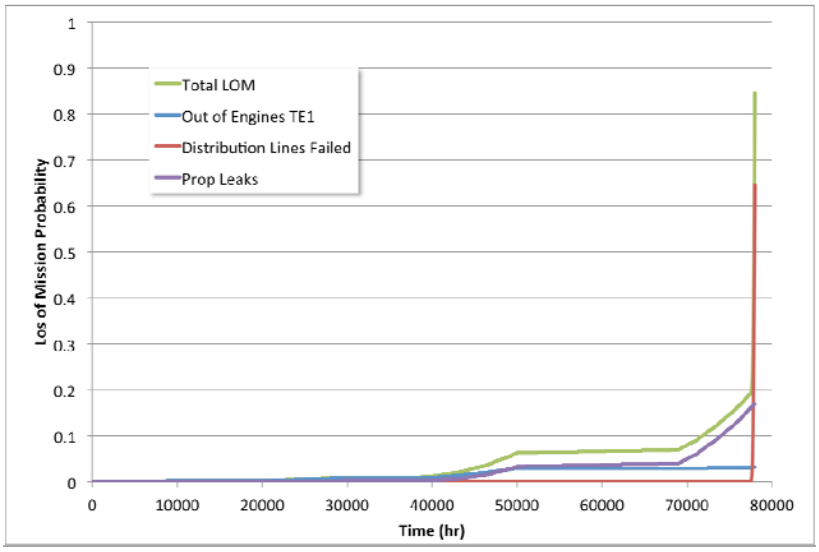

Figure 10 shows the relative importance of the "out of thruster" and "out of propellant" end states—the engines are responsible for almost three quarters of the LOM. This figure also provides the first comparison between the ART (column 1) and simulation models (columns 2 – 5). The ART model misses the complexities of the engine switching dynamics if one engine within a thruster fails. As mentioned, these dynamic effects impact the results less than the simple operating failure modes, so the ART predictions are reasonably close to the simulation results, particularly given the simplicity of the ART thruster model. "Out of propellant" is a simple calculation for the baseline mission, so these results compare quite closely. Figure 10 also shows the risk contribution from each of the end states versus the number of simulation realizations. In this case, the relative contribution remains fairly constant and the overall magnitude converges by 1,000 realizations.

Figure 10: Static ART model of the baseline mission compared with simulation results.

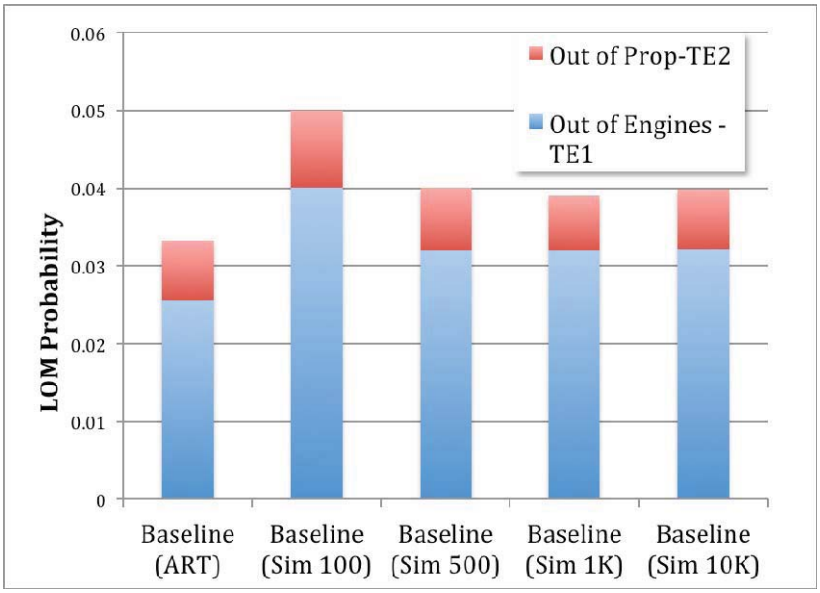

Figure 11 contains the ART results for two versions of the baseline mission and the two extended mission options, along with the 10,000-realization simulation results for the extended missions. Because the dynamic aspects of these models are important overall, several variants of the ART model were created. The first baseline version, *Baseline-1 (ART)* in column 1, simply applied the expected value of the distribution line random walk as well as using the tank failure rate without repair. This is a degenerate case since the failure probability is 1, but it does illustrate that rote application of failure rates to a dynamic problem, as if it were a static equivalent, can result in nonsensical output. The second baseline version, *Baseline-2 (ART)* in column 2, computed a single random walk probability distribution for the distribution lines and again assumes the tank fails at the baseline rate without repair. This improves the results compared to the baseline simulation model, but the risk remains too high and the relative contributions from the drivers are visibly different. The third column, *Extended no Station (ART)*, includes the side models for both the distribution lines and tank, which create equivalent failure probabilities for the ART based on offline dynamics. The results now match the dynamic equivalent (column 4) quite well, but this is no longer a static modeling approach per se. The same side models were used to create the final comparison for the *Extended with Station* case. Again, the agreement is reasonable but the ART results (column 5) are lower than the comparable simulation results (column 6) due to omitting cross-element cut-sets and simplifying assumptions made in the tank leakage side model. While not part of the challenge problem, the ART comparison does give an idea about the importance of considering the dynamic aspects of the extended problem; the answers begin to diverge without dynamics both in LOM magnitude as well as the contributions from the failure modes.

Figure 11: Static and dynamic model comparison for the baseline and extended missions.

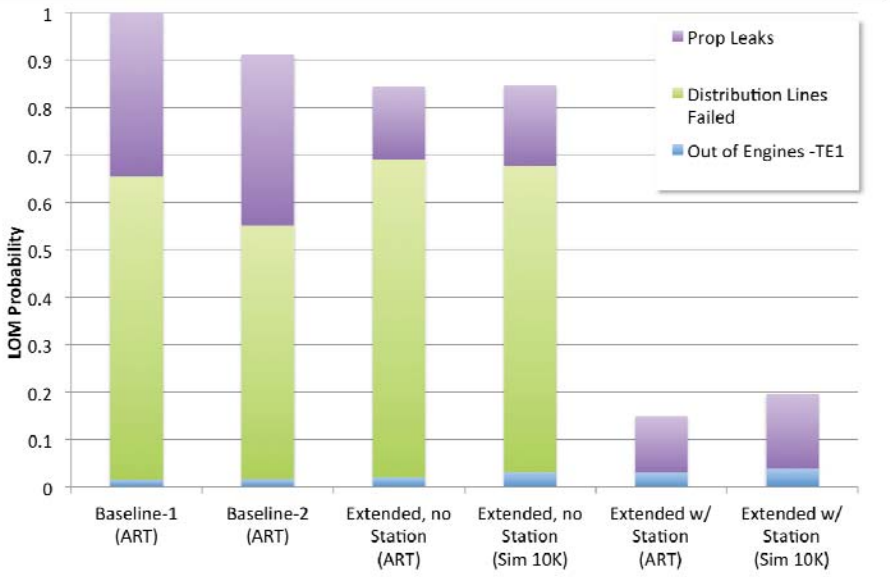

Figure 12 shows a comparison of the "out of engines" failure condition, Top Event 1 (TE1) from Ref [1], for the ART and baseline simulation models. PPU failure frequencies differ between the static and dynamic models by approximately a factor of two, while the two "engines fail to operate" results differ by 20%. Common-cause failure contributions were included in the "fails to operate" bin for both the PPU and engine failures. These differences are due to simplifying assumptions in the ART, which treat common-cause failure modes as either all strings failing through common-cause or all strings failing randomly, with no treatment of combinatorial mixed failure cases. All of the other failure contributors are comparably insignificant and do not impact LOM measurably. Only the thruster failure modes that were involved with LOM were counted in this plot. The cases where the mission continued after a thruster failure, due to operational backup engines, were omitted from the results in Figure 12, though that information is available from the simulation.

Figure 12: Insufficient engines (TE1) comparison.

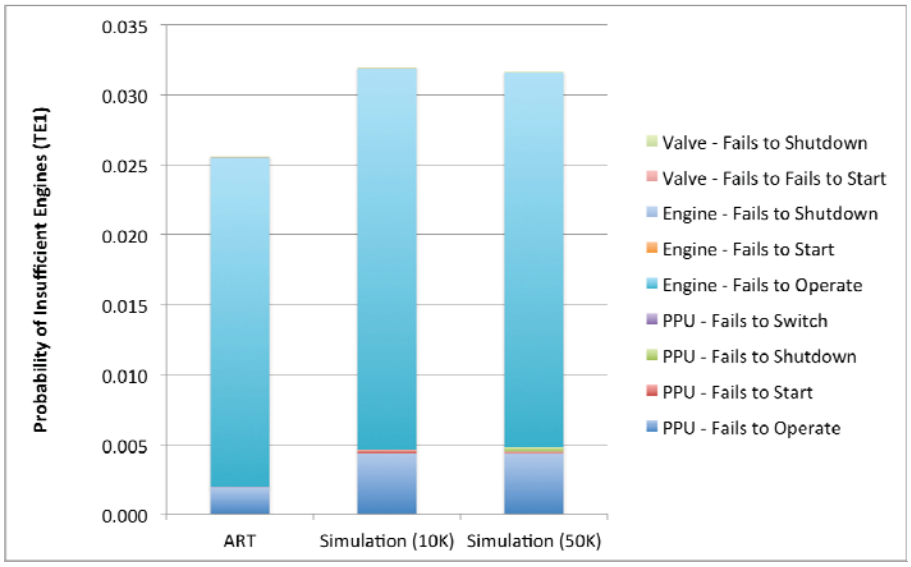

The simulation can also provide information about the observed failure frequency distribution among the thrusters, as shown in Figure 13. Again, only the thruster failures that contribute to LOM are included. Thrusters 1 and 2 have the visibly highest contribution to LOM. This is intuitive because the model draws on thrusters in order from 1 to 5. Therefore, Thruster 1 is utilized on every mission, whereas Thruster 5 only comes into play after 2 – 3 other thrusters have failed. So, even though the thrusters are identical and have inherently the same reliability, the observed failure frequencies differ greatly because of duty cycle differences. By the time Thrusters 4 and 5 are called into action, the remaining mission duration is generally low, thereby reducing their exposure times.

Figure 13: Observed LOM contributions by thruster.

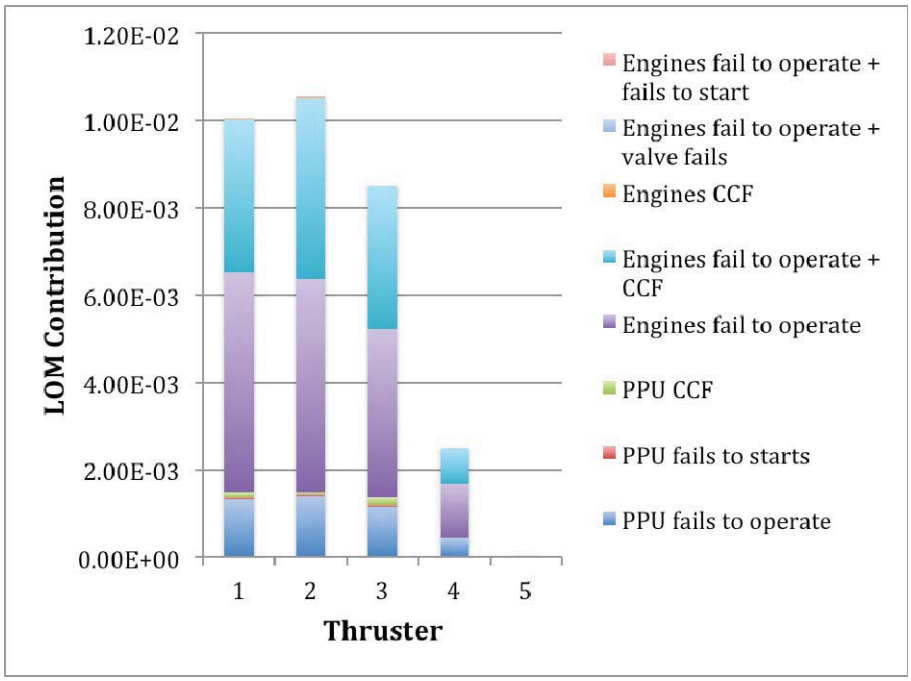

4 CONCLUSION

An Engineering Risk Assessment of the baseline and extended versions of the challenge problem was performed. The dynamic software models were built using the GoldSim software. All aspects of the problem were successfully modeled. Convergence studies were performed that illustrated the mission LOM estimates converged in less than 1,000 realizations, though risk driver resolution required more depending on the level of interest and convergence desired.

The baseline mission contained a dynamic element, but its impact on mission risk was minimal. For the extended versions, however, the dynamic nature of the problem was important to the overall results. Though the propulsion system was the same for all missions, the risk drivers changed based on the mission. Despite the additional thruster demands required when utilizing the support station, all of the scenarios assessed delivered higher levels of mission success with its use. This occurs as a result of resetting the distribution line damage, effectively reducing the damage accrual time by 18,000 hours. Since the distribution lines fail in the last 500 hours of the mission, this reset of the damage effectively removes them as a risk contributor.

A deterministic risk model was built for the three mission alternates using the Ames Reliability Tool. The goal of this comparison was not to make a static model dynamic, but to evaluate the relative importance of the mission dynamics. In the baseline case, the ART model did a reasonable job of matching the simulation models. However, offline dynamic models were required for the extended missions, creating a hybrid approach. Though the LOM estimates were comparable in these cases, the simulation model provided much additional information, such as failure time history, insight into observed failures of similar systems, etc.

This study highlighted that state-of-the-art techniques can adequately adapt to a range of dynamic problems. The need to use such techniques does not depend on the dynamics, per se, but on the impact the dynamic portions of the mission have on the risk parameters of interest.

References

[1] Mandelli, D., Smith, C., and Rabiti, C., "Space Propulsion System, A PSAM Benchmark Problem for Safety Analysis Algorithms," Benchmark Problem #1_extended_ver3_1, November 2013.
[2] www.goldsim.com
[3] Jonathan Goodman and Keith Lewis, Derivative Securities, Courant Institute, Fall 2008: http://www.math.nyu.edu/faculty/goodman/teaching/DerivSec08/notes/Section3.pdf
[4] B.F. Putney, E. Tavernetti, J.R. Fragola, and E. Gold, "*Reliability Tool for a Preliminary Quantified Functional Risk and Hazard Analysis*", Proceedings of the Reliability and Maintainability Symposium, 2009, Fort Worth, TX.

A APPENDIX

A.1 Methodology Employed

Name: Engineering Risk Assessment
Software tool: GoldSim
Software parameters: 10^4 realizations
Computational resources: Intel i7 quad core laptop running Windows OS
Computational time: 45 minutes with 100-hour time step size
Risk-metric: Observed failure counts/frequencies and temporal failure time frequencies

A.2 Benchmark Problem

Table 1: Benchmark problem methodology comparison table.

Component	*Hypothesis/ Approximations*	*Modeling*	*Notes*
PPU	Time-stepping exponential failure model for time-based elements.	PPUs were modeled using reliability elements, which include discrete switching risk and continuous operation failure modes that accumulate risk when operational.	
Ion engine	Discrete event failure model used for demand failures.	Ion engines were also modeled using reliability elements with the appropriate failure modes per the problem description.	
Valve	Same as above	The valve damage accumulation was modeled using an integration element combined with a random-choice element that produced a sample from the prescribed distribution with a table lookup of an offline numerical integration table.	
Distribution lines	Same as above	Distribution lines were modeled using reliability element with the appropriate failure rate.	
Phase mission	N/A	Mission timeline parameters are managed using a selector element that maps the proper requirements, duty cycles, and parameter to the elements	
CCF	Explicitly modeled by conditional discrete event.	The CCF parameters are included in data elements, and a random-choice element is used to evaluate CCF in the event of an initial failure.	

Results

Table 2 gives the ranked risk drivers from the baseline mission simulation results. Engine and PPU values represent failures to operate, failures to start, and CCFs.

Table 2: Ranked risk drivers from baseline mission simulation.

Baseline Mission	*Failure Rate (failures per mission)*
Engines	2.70E-02
Distribution lines	7.60E-03
PPU	4.00E-03
Valves	1.00E-04

A.3 Extended Problem

Table 3: Extended challenge problem methodology comparison table.

Component	*Hypothesis/ Approximations*	*Modeling*	*Notes*
Tank initial conditions	110,000 + random sample from uniform distribution 5500-16500.	Initial propellant load was represented with a "stochastic" element that selects from the input distribution.	
Distribution line leaks	Random walk mean of 1 (1.3), sigma = 0.4 (0.6) nominally (gray region).	An integration element sums the "steps" of the Brownian motion walk with each step size resulting from a sample from the stochastic element populated with the problem parameters.	
Tank leaks	Time-stepping exponential failure model to determine if leak occurs (and when). Sample from normal dist (mean = 500, sig = 100) for leak rate. Repair time exponentially distributed with lambda = 24 hours.	In the event of tank leakage, the leakage rate is determined by a stochastic element sample. The repair is performed using the built-in repair functionality within the tank element. The propellant lost is removed from the tank reservoir.	
Support station and mission alternatives	User defined input determines if support station will be used.	The reliability elements that define the hardware elements can be restored to their original state within the model. This feature is used to repair any damage to the system if the support station is used. The propellant in the tank is restored to the original value.	The model reliability without the support station was very low due to the distribution line leakage implementation. Therefore the station would be used in every case.

Results

Table 4 and Table 5 give the failure drivers for the extended mission options. The inclusion of failure modes (fails to operate, CCF, etc.) is the same as for the baseline mission.

Table 4: Ranked failure drivers for extended mission with no station.

Extended Mission, No Station	*Failure Rate (failures per mission)*
Distribution lines	6.45E-01
Tank Leakage	1.70E-01
Engines	2.73E-02
PPU	4.74E-03
Valves	1.00E-04

Table 5: Ranked failure drivers for extended mission with station.

Extended Mission, with Station	*Failure Rate (failures per mission)*
Tank Leakage	1.60E-01
Engines	3.32E-02
PPU	5.78E-03
Valves	1.00E-04
Distribution lines	0.00E+00