INL/CON-15-35374
PREPRINT

Learning from Trending, Precursor Analysis, and System Failures

ABRISCO 2015 and Topical PSAM Meeting on Safety and Reliability of O&G Exploration and Production

R. W. Youngblood and R. B. Duffey

November 2015

The INL is a
U.S. Department of Energy
National Laboratory
operated by
Battelle Energy Alliance

Learning from Trending, Precursor Analysis, and System Failures

R. W. Youngblood
Idaho National Laboratory, USA

R. B. Duffey
DSM Associates, Inc., USA

1. INTRODUCTION

This paper makes the following points.

1. Existing data-driven theories of reliability growth are straightforwardly reconcilable with a simple class of simulation models for reliability growth.
2. Based on point 1, a simple simulation model is used to illustrate notionally the potential return on an investment in precursor analysis.
3. There being at present no satisfactory theory of how to winnow out precursor events from a flood of experience data, recent thinking on that topic is recapitulated from the point of view suggested by the illustration in point 2.

2. RELIABILITY GROWTH: THEORY

"Reliability growth" refers to the process by which a system under development improves in reliability, as its failure modes are experienced in system trials, and thereafter eliminated by the system developers through changes in design or operating practice. For a useful review, see [1]. The original idea of reliability growth is generally attributed to Duane, who plotted cumulative failure rate (vertical axis) against a measure of "experience" (operating hours for some systems, cumulative trials for one-shot systems) (horizontal axis). On such a plot, early in development, when there are many failure modes yet to be eliminated, and failure occurs frequently, the cumulative number of failures rises steeply as a function of experience; as the dominant failure modes are eliminated, the curve bends over and flattens out, corresponding to a longer interval between failures. Duane[1] was able to fit a simple power law formula to such a curve,

$$\lambda(t) = Ct^{-\alpha}, \quad (1)$$

where λ is the average failure rate (total failures over total time), t is the operating, testing, or risk exposure time, and C and α are fitted parameters. This construct can be used to forecast expected unreliability for a range of developmental systems [1].

Numerous authors have built on this general idea since Duane's work, and incidentally addressed certain counterintuitive features of the original formula. The present discussion is based on the recent work of Duffey and Saull [2]. For present purposes of illustration, certain features of the Duffey-Saull method (DSM) are especially noteworthy:

- The presence of failure modes decays exponentially with increasing experience,
- Similar parameter values apply to a broad class of systems, provided "experience" is defined appropriately, and
- For a given system and data set, the DSM theory helps the user to focus on the right measure of experience.

The failures are regarded as occurring randomly but decrease systematically with learning, manifested as error correction, skill acquisition, improved training and procedures, system redesign, or fault elimination.

[1] In the original paper, there is reverse titling of the two graphs shown.

For reference, the instantaneous failure or learning rate derived from the DSM Learning Theory is dynamically changing with operating experience, trials, tests or risk exposure, t, and is given by:

$$\lambda(t) = \lambda_m + (\lambda_0 - \lambda_m) \exp^{-kt}, \quad (2)$$

where λ_0 is the initial rate, λ_m the minimum achievable rate, and k the learning rate constant that is fitted to data. The failure rate can be straightforwardly converted to a failure probability variation, which is also exponential in form, including for randomly occurring events.

Illustrative plots applying the DSM learning theory are given later.

3. SIMULATION OF RELIABILITY GROWTH BASED ON FAILURE EXPERIENCE

Given the essential simplicity of the above idea, simulation of a reliability growth history is straightforward. Any system can be regarded as consisting of individual hardware and software components, numerous sub-systems, and/or a series of multiple independent "barriers" that are subject to potential failures, errors or faults. Consider, as an example, development of a one-shot system[2] such as a space launch system, so that experience can notionally be measured in trials (launches) for some purposes. Then a simulation could proceed as follows:

0. Initialize the system state: Specify the failure modes in the system at time zero (before any trials have taken place), their probabilities, and, for each failure mode, the probability that once experienced, that mode will be successfully eliminated immediately (before the next trial, or any continuation of system operation).
1. Perform a trial: sample random numbers to determine which of the current failure modes occur in the present trial.
2. Sample other random numbers to determine which, if any, of the failure modes experienced in this trial are successfully eliminated.
3. Based on Step 2, strike the eliminated failure modes from the current list of failure modes. Record the failure(s) experienced in this trial, the current list of failure modes remaining after elimination, the current system unreliability, …
4. Go to step 1; or, if sufficient trials have been simulated, stop the simulation.
5. Either analyze the data from this time history, or simulate enough time histories to improve statistics, and analyze the ensemble.

Even within the above simple process, numerous variations can be contemplated. In perhaps the simplest version, only the failure that occurs first in a given trial (e.g., earliest in the ascent of the launch vehicle) is eliminated (the system is destroyed before the later failures manifest themselves). In a more complex version, a more complex rule could be contemplated, allowing for elimination of more than one failure mode. In addition to the possibility of unsuccessful remediation of a particular failure mode, one could allow for the possibility of introduction of completely new failure modes during the process of trying to eliminate the old ones. Unfortunately, illustrations of this can be found in previous experience [3].

Reference [1] distinguishes failure modes that will not be corrected even if identified ("A-modes") from failure modes that will be corrected ("B-modes") if identified. A-modes correspond to accepted risks, while B-modes are simply unknown initially. For simplicity, the following discussion neglects A-modes.

4. SIMULATION OF RELIABILITY GROWTH BASED ON FAILURE EXPERIENCE PLUS PRECURSOR ANALYSIS

Looking back after a major accident, one is sometimes able to identify previous events or measurable performance trends that were, in some sense, signaling the potential for that major accident: potential

[2] By "one-shot," we mean that each launch vehicle is used once and not recovered for re-use, whether or not it fails, each shot or launch corresponding to a measure of the experience or risk exposure.

precursors that could have been recognized and acted upon, but were not recognized until the accident occurred. This could be a previously unrecognized cause of accidents, or underestimation of the likelihood that a recognized potential cause would actually operate. It is clear that many major accidents are preceded by precursor events; for example, the Davis-Besse stuck-open PORV (1977) presaged the TMI-2 core melt accident (1979), and both Space Shuttle disasters were preceded by precursor events that seem clear in hindsight [4]. There are also innumerable industrial accidents that have precursors, and are precursors to accidents still in the future, such as train derailments, oil leaks, chemical fires and explosions, shipping losses, as well as accidents traceable to manufacturing defects (e.g. in automobile airbags and switches).

At this point, "precursor analysis" has been practiced in some arenas for decades [5]. However, doing it efficiently while avoiding false negatives (failing to notice important precursors) is still a significant challenge. The problem is that while the significance of outright failures is easy to understand, the significance of mere anomalies is not, in general, obvious. Youngblood *et al.* argued [6] that "apparent risk significance" is not a reliable guide to analysis of operational anomalies; it might well be useful to add a filter based on how "surprising" an anomaly is. The underlying idea is that if an anomaly is "surprising," this means that our model of the world assigns a low prior probability to that anomaly; therefore, occurrence of the anomaly indicates that our model of the world is probably wrong, and it may be important to understand just how it is wrong. In [6], the failure to act on the clogging of containment air filters at Davis-Besse (before the vessel head problem was understood) is cited as an example of how apparent risk significance can be a misleading consideration.

From that point of view, it is reasonably natural to think of some precursor events as being simply milder instances of the system failure modes, having higher prior probability than outright failure but less prior probability than our assumption of normal behavior. For example, suppose that extreme vibration could cause failure of a launch system, but the designers either expect little or no vibration in practice, or have not contemplated that failure mode. Then occurrence of a significant (anomalous, *a priori* unlikely) amount of vibration arguably should occasion a re-examination of the design. To simulate reliability growth in this situation, we posit, in addition to a failure mode and its probability, a threshold value of probability corresponding to occurrence of a "precursor:" an anomaly large enough to be unexpected, but still short of outright failure. To continue the "vibration" example, observing an unexpected (*a priori* unlikely) amount of vibration would be regarded as a precursor. In the simulation, in each trial, we sample each failure mode to determine whether (a) a failure, (b) a precursor event, or (c) neither occurs, and we may sample again to determine whether the failure mode is eliminated before the next flight, conditional on the occurrence of the failure or the precursor, as the case may be.

Clearly, within such a modeling framework, there is the potential for eliminating some failure modes without first suffering the accident associated with that failure mode.

5. EXAMPLE

Potential random failures of a notional series of components and barriers were encoded in a simple simulation algorithm hereafter called "simulationX," which can, in principle, model random failures in multiple systems and eliminate the failure modes manifesting as precursors and failures. Figure 1 compares two notional examples. In Model 1, failure modes are eliminated only as outright failures occur; in Model 2, failure modes are eliminated upon either failure or the occurrence of a precursor. Both models are initialized with the same failure modes and failure probabilities. For each failure mode, each precursor is assigned a probability equal to 10 times the probability of failure. The actual inputs to the two models are shown below in Table 1. For all three models, the probabilities of precursors and failures are as indicated. In Model 1, no failure modes are eliminated as a result of precursors, but in Model 2, failure modes are always eliminated after a precursor. In Model 3, failure modes are eliminated in all system replicas after a failure or a precursor is experienced in any system replica. The plots of Model 1 and Model 2 show unreliability due to remaining failure modes, as a function of launch number, averaged over 20 time histories for each model.

Because each failure mode has a precursor probability equal to 10 times the probability of outright failure, precursors do not always precede outright failures, but they have a very good chance of doing so; correspondingly, the system unreliability improves at a much, much greater rate within Model 2, i.e., if

precursor analysis is done. Put another way, a given level of reliability is attained much more quickly if the strategy of Model 2 is adopted: learn from precursors.

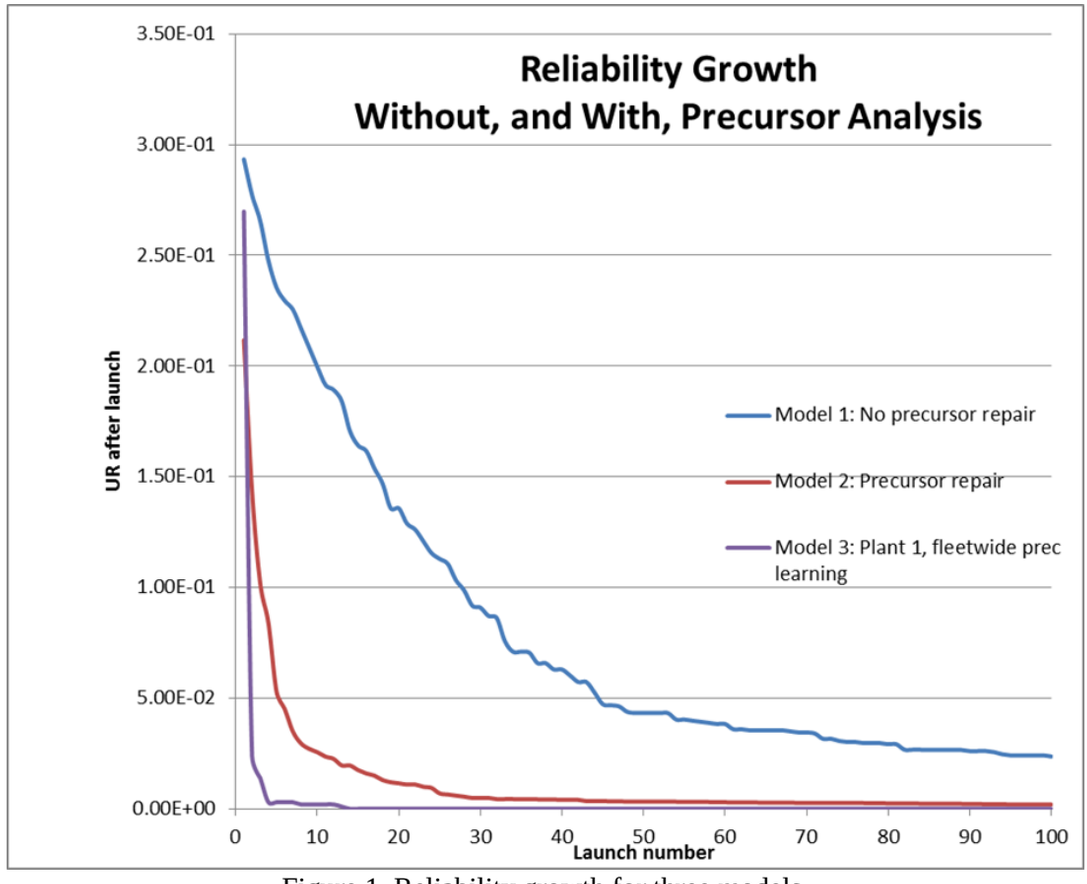

Figure 1. Reliability growth for three models.

In Figure 1, the three models are: (1) elimination of failure modes after they are experienced; (2) elimination of failure modes after either they manifest themselves either as outright failures or in precursor events; (3) elimination of failure modes after failures or precursors anywhere in the operating fleet.

Model 3 has the same failure modes, associated failure probabilities, and associated precursor probabilities, but contemplates a campaign in which each of N systems is launched one after another, and every system benefits from the knowledge of the precursors and failures occurring in all previous launches of all of the systems, not only from its own accidents and precursors. Information is shared across the entire notional "industry." The first launch of the 10th system benefits from experience gained in the first launch of systems 1 through 9; the second launch of the 1st system benefits from experience gained in all of the first launches of all of the systems, and so on.

In Figure 2, we show the best fit of Statistical Error State Theory (SEST), from [2], to the simulationX predictions. The SEST is derived on the basis of randomly observed failures as also assumed in the simulation. The agreement shown strongly suggests that the simulation correctly captures not only the statistical nature but also the "learning-from-precursor" trends with up to a factor of ten potential improvement in reliability. In Figure 2, the fit to "no precursor remediation" is given by

$$p = 0.284*\exp(-n/21.4)+ 0.021, \quad (2)$$

and the fit to "precursor remediation" is given by

$$p = 0.266*\exp(-n/3.3)+ 0.0046. \quad (3)$$

Table 1. Model Specification. In the fourth column, "remediation" refers to elimination of a failure mode once it has been observed. In the fifth column, "remediation" refers to elimination of a failure mode after a precursor has occurred. In Model 1, there is no failure mode elimination following precursors.

Failure Mode	Probability of Outright Failure	Probability of a Precursor	Probability of Failure Mode Elimination After an Outright Failure	Probability of Failure Mode Elimination After a Precursor (Models 2 and 3)
1	0.05	0.5	1	1
2	0.05	0.5	1	1
3	0.05	0.5	1	1
4	0.05	0.5	1	1
5	0.05	0.5	1	1
6	0.01	0.1	1	1
7	0.01	0.1	1	1
8	0.01	0.1	1	1
9	0.01	0.1	1	1
10	0.01	0.1	1	1
11	0.001	0.01	1	1
12	0.001	0.01	1	1
13	0.001	0.01	1	1
14	0.001	0.01	1	1
15	0.001	0.01	1	1

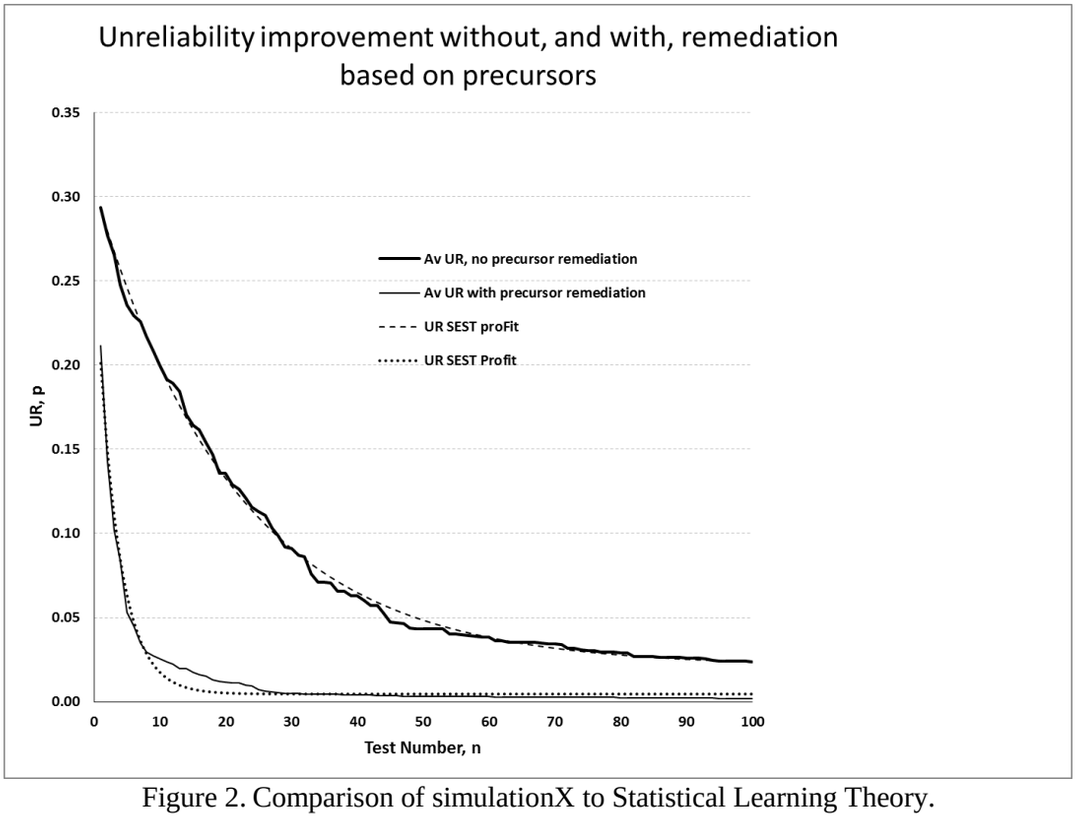

Figure 2. Comparison of simulationX to Statistical Learning Theory.

6. COMPARISON WITH INDUSTRIAL EXPERIENCE

We need to understand and illustrate the impact of learning and precursor data on relative risk. The simplest comparison [2] is by using the well-known information entropy or H-Factor, defined as

$$H = -p \ln p \,, \qquad (4)$$

where p is the unreliability, p(UR), and is a direct measure of the uncertainty and hence can be used as a measure of comparative risk.

The simulations presented earlier gave unreliability as a function of the number of flights or tests, n, where the total experience is for 100 tests representing the maximum depth of experience. In order to compare across industries, we define a non-dimensional depth of experience, or the risk exposure measure, given by N*=n/N. The simulation calculated p(UR), and hence H, for both cases: failure mode elimination only after failures, and failure mode elimination after either failures or precursors, so that the fault does not recur.

Figure 3 shows the comparison of these two cases (the dashed lines, input data given in Table 1) to the Statistical Error State Theory (SEST), which is based on learning from random outcomes, and to a sample of actual data. These data are from [2] for diverse precursors: (a) offshore USA oil spills over 1973-2000 (where risk exposure is in millions of tons shipped); (b) train derailments in the UK during 1988-1999 (where risk exposure is measured in billions of passenger –kilometers), (c) near mid-air collisions for the USA from 1987 to 1998, (where risk exposure is in millions of flying hours) and auto accidents in Australia for 1980 to 1999, (risk exposure in driver-years); and finally (d) US coal mining from 1931 to 1998 (with risk exposure in millions of hours worked) .

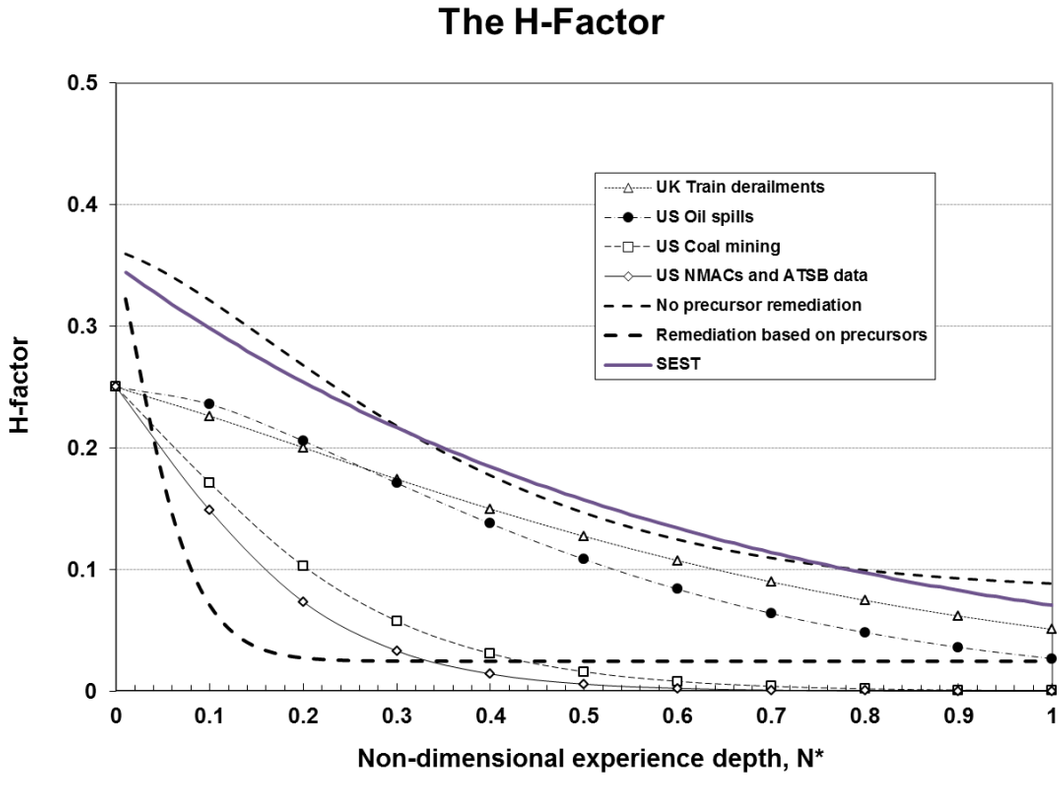

Figure 3. Comparison of simulationX with industrial experience.

The SEST theory line shown is given by H=0.35exp-1.6N*.

The simulationX model generally reproduces the trends of both the theory and the data, and shows the potential of major risk reduction by using precursor information as well as failure information. The two cases

shown (learn from failures, learn from both failures and precursors) generally bracket the other cases shown, and the comparison suggests that prevention of oil spills at sea and train derailments has not significantly advanced compared to aircraft and coal mining. The latter two have introduced significant automation to reduce the human error contribution: reduced manual labor using auto mining equipment, and reduced near misses using auto collision warnings. The other two have not yet made all the advances and learning that is possible, although as of recently, automatic warnings are now being introduced into trains, and additional backup safety measures added to offshore oil transport and drilling.

While the plot is suggestive, we do not wish to overinterpret these results. The simulation data in Figure 3 ("No precursor remediation" and "Remediation based on precursors") given in Table 1 were chosen for illustration, and had we gone beyond Flight 100, the shape of those curves would change somewhat. Nevertheless, the figure strongly supports the idea that something fundamentally simple is going on.

7. SUMMARY

We have shown that indeed failures and the impact of learning can be successfully modeled using statistical analysis and error correction. The agreement between theory and model is shown, plus the insights on comparative risk from the learning and precursor correction trends. The ideas presented here are in agreement with the data trends from diverse industries.

The above discussion notionally illustrates the potential value of precursor analysis geared to the elimination of failure modes before they occur, based on observation of related anomalies which were unexpected *a priori*, and which therefore signal a possible deficiency in our model of the system. Given an observed anomaly that satisfies certain screening criteria, the recommendation in [4] is to perform sufficient investigation to understand the nature of that possible deficiency (in particular, to understand the causal mechanism of the observed anomaly, and how that mechanism might have operated differently to cause an accident), and deal with it accordingly: modify the system, the model, or both.

We cannot quantify the expected benefit of precursor analysis in a particular system without making assumptions about the B-modes that are present. But history has shown that in many systems, important B-modes are present. The present notional illustration proceeds from the heuristic argument that although interesting precursors are, by nature, unlikely *a priori*, they are more likely than accidents, and are correspondingly likely to occur sooner in operating history than their corresponding accidents. That reasoning was hard-wired into the simple simulations performed here; the basic assumption may not be universally valid, but is arguably quite reasonable for many important failure modes, including those that have presaged some well-known accidents. In particular, both space-shuttle disasters had precursor events, as did the Three Mile Island core melt accident. Precursor analysis should be less costly than major accidents, and, for many systems, will be highly net-beneficial.

8. ACKNOWLEDGMENT

This work was performed under Contract No. DE-AC07-051D14517.

9. REFERENCES

[1] Department of Defense Handbook, Reliability Growth Management, MIL-HDBK-189C (DoD, 2011).
[2] DUFFEY, R. B., & SAULL, J. W., *Managing Risk: the human element*, Wiley (2008).
[3] HAMLIN, T. *et al.*, Shuttle Risk Progression: Use of the Shuttle Probabilistic Risk Assessment (PRA) to Show Reliability Growth, (AIAA, 2010) (downloadable from http://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20110004917_2011004008.pdf)
[4] *NASA Accident Precursor Analysis Handbook*, NASA/SP-2011-3423 (NASA, 2011).
[5] *Risk Assessment of Operational Events*, Rev. 2 (USNRC, 2013).

[6] YOUNGBLOOD, R., MAGGIO, G., EVERETT, C., & HALL, A., Value of Analyzing Operating Events, Proceedings of the 9th International Probabilistic Safety Assessment & Management Conference (PSAM-9), May 2008, Hong Kong.