Quantifying System Resilience Using Probabilistic Risk Assessment Techniques

Clayton A. Smith and Timothy J. Allensworth

ABSTRACT

Resilience provides a framework to assess a system's likelihood to succeed in its mission even as disruptions perturb the operations. A system's resilience is therefore essentially a risk proposition of the mission succeeding and as such can be quantified using probabilistic risk assessment (PRA) techniques developed over the past three decades. Reliability engineering methods for evaluating hardware are insufficient by themselves, as they do not examine procedural mitigations, system margin, or human training applied to overcome anomalies. A resilient system needs to address failures, regardless of whether their cause is component malfunction, operator error, or external disruption, and continue to operate (perhaps at reduced functionality) within off-nominal operating environments. Resilience approaches look beyond hardware-only solutions to generate additional mitigation concepts to prevent, withstand, adapt to, and rapidly recover from failures or external disruptions. The methods and techniques used to produce PRAs encompass not only the hardware but also the operating procedures, the contingency plans, the software, and the physics of the eventual consequences. This article discusses how PRA is applied to quantifying system resilience and focuses on two aspects: scenario development and uncertainty quantification.

INTRODUCTION

System resilience is essentially a risk proposition in that when we think about it, we seek to understand how a system might be affected by an adverse event, its potential end states, the associated likelihoods, and the impact of mitigations. Reliability assessments of engineered systems began in earnest with the creation of fault tree analysis in the early 1960s within the aerospace industry. By the late 1970s, the nuclear power industry was using fault tree analysis in conjunction with event trees in what is now called probabilistic risk assessment (PRA). The methods and techniques used to produce PRAs apply to not only the hardware but also the operating procedures, the contingency plans, the software, and the physics of the eventual consequences.

Reliability assessments quantify the probability that a system will perform its intended functions for a spe-

cific period. Typically, the definition of a system in the reliability approach is relegated to hardware components operating in their intended environment with all the nominal supporting operations performing and logistics present as specified. A resilient system needs to overcome failures, regardless of whether the cause is component malfunction, operator error, or external disruption, and continue to operate (perhaps at reduced functionality) within an off-nominal operating environment. This article briefly describes the elements of a PRA and focuses on two aspects of PRAs as they relate to resilience quantification: scenario development and uncertainty propagation.

PROBABILISTIC RISK ASSESSMENT

The Nuclear Regulatory Commission has used PRAs since the 1970s to quantify safety risks at nuclear power plants. The commission recognized that hardware reliability assessments were not sufficient to understand how an adverse event might affect these engineered systems. Today PRAs are also being used in the defense, petrochemical, and offshore oil drilling industries. The Johns Hopkins University Applied Physics Laboratory (APL) has used PRA techniques to solve challenges for a variety of sponsors, including NASA, the Missile Defense Agency, Naval Sea Systems Command, the U.S. Air Force, and the Space Security and Defense Program.

PRA studies over the past several decades have pioneered the inclusion of new analysis tools to examine common-cause failures, external events, software reliability, human factors, and consequence modeling. All this activity led Bedford and Cooke to state that "the trend in all areas is for PRA to support tools for management decision making, forming the new area of *risk management*."[1] This methodology actively tied PRA to specific decisions to reduce risk. Kaplan and Garrick[2] see PRA as a method to answer the following questions: What can go wrong? How likely is it? What are the consequences? How credible are the results presented? Leadership needs answers to these questions to make decisions.

Resilience can be improved with implementation of mitigation solutions that enable the system to prevent, withstand, adapt to, and/or rapidly recover from adverse events.[3] With PRA, solutions can be applied and evaluated together at the physical system layer (fault tree) or at the functional system layer (event) within the context of an operation or mission. Each potential solution affects the probability that the scenario will succeed, and the spectrum of resilient solutions can be compared to quantify how each increases the likelihood of mission success.

A PRA is constructed by integrating several elements, as shown in Fig. 1. The master logic diagram represents the list of events that can perturbate the system, be they hardware failures, procedural errors, external events, or attacks. Initiating events (green spheres) are the disruptions (i.e., threats, failures, or adverse environmental impacts). End states (red diamonds) are related to the consequences of performance degradation and recovery time. A model of scenarios

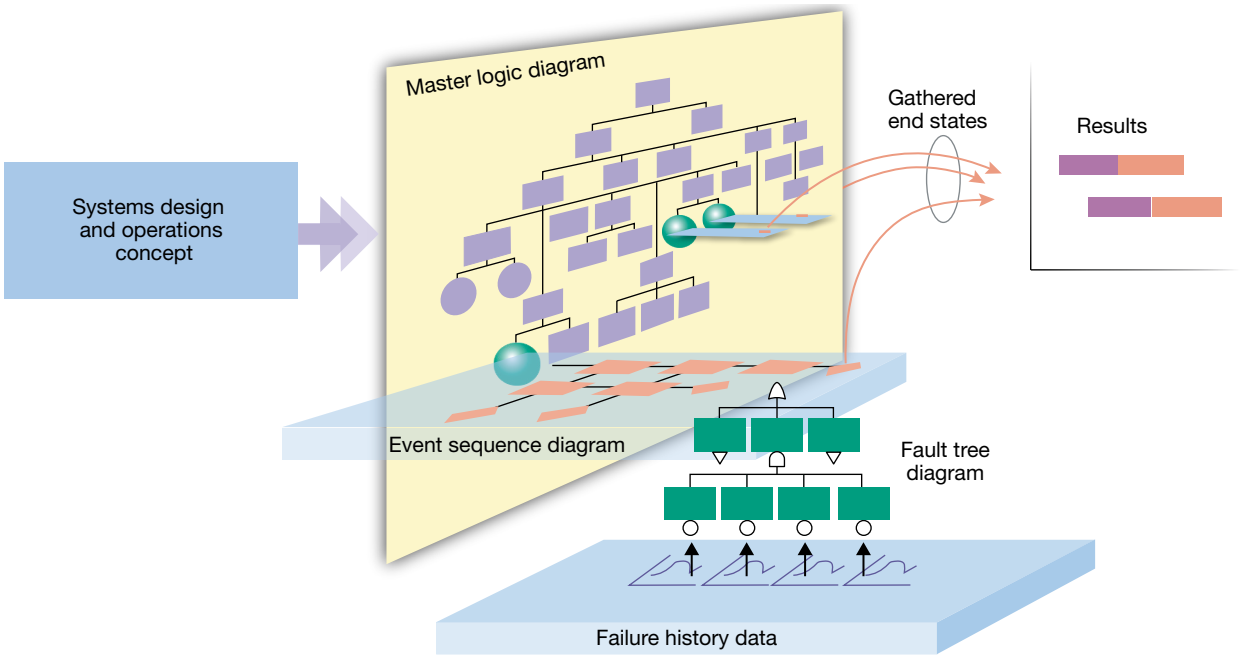

Figure 1. Integrated PRA components—a cohesive model. [Reproduced from Smith, C. A., "Probabilistic Risk Assessment for the International Space Station," in *Proc. Joint ESA-NASA Space-Flight Safety Conf.*, B. Battrick and C. Preyssi (eds.), European Space Agency, ESA SP-486, p. 319 (2002).]

called an event sequence diagram (ESD) connects the initiating event to end states, describing an operation or intended mission.

Each event in the ESD is a question about the system responses. These questions, called pivotal events, are assigned a probability of occurrence either through fault tree analysis, data, or simulation results. Resilience considerations to improve the robustness, redundancy, rapidity, and resourcefulness of the system can therefore be poised as pivotal events.[4] The end state probability and consequences are determined and gathered across all initiating events to produce a profile of values. Mitigation solutions are evaluated to determine the impact of the design change on the consequences affecting the end-state probability of success.

PRAs represent risk with Farmer curves or risk profiles. These curves plot the magnitude of many different consequences against the complementary cumulative probability distribution of the scenario. Figure 2 shows a famous Farmer curve from the *Reactor Safety Study*[5] used to illustrate relative risks of nuclear power plants compared to other risks. The analyst compares the risk of nuclear power plant fatalities with frequencies of fatalities attributable to air crashes, fires, dam failures, explosions, chlorine releases, and air crashes. Henley and Kumamoto further explain the figure by stating:

> Nonnuclear frequencies are normalized by a size of population potentially affected by the 100 nuclear power plants; these are not frequencies observed on a worldwide scale. Each profile in [the figure] is called a Farmer curve; horizontal and vertical axes generally denote the accident severity and complementary cumulative frequency per unit time, respectively.[6]

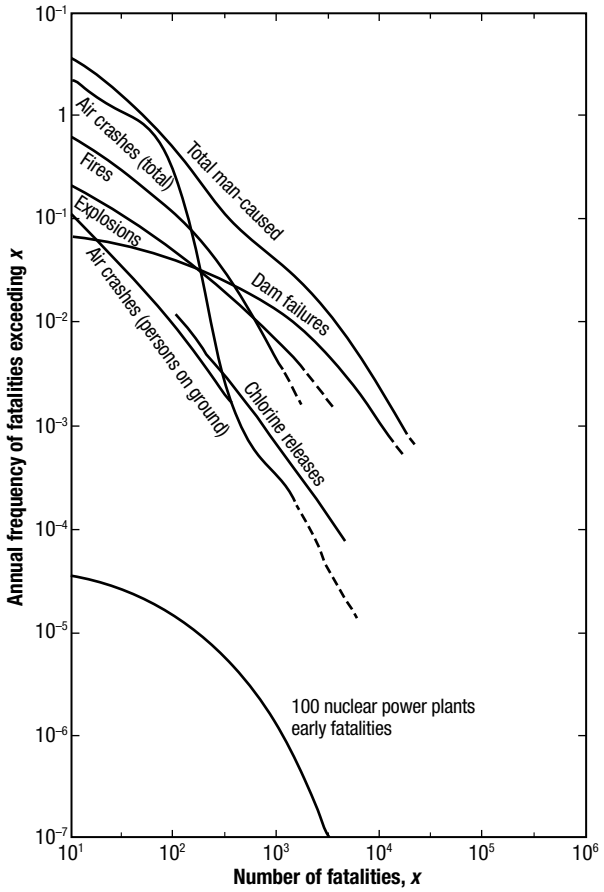

Figure 2. Famous PRA Farmer curves example. (Reproduced from Ref. 5.)

This representation allows an analyst to move away from displaying only a single dimension of risk to displaying the probability of all possible values of consequence. Often, these types of assessments examine a threat or vulnerability and assume a causal link to only one consequence. From a resilience viewpoint, initiating events take many paths to several consequences because of the existence of intrinsic system capabilities and mitigation measures. PRA methods consider thousands of scenarios that involve multiple failures or events, thus providing an in-depth understanding of potential system failure modes. Traditional risk matrices do not investigate such an enormous number of possible scenarios.

Resilience quantification revolves around assigning probability to various events. Since events and scenarios of interest are quite infrequent, gathering data based on observations is not possible. In the context of PRA and resilience quantification, the concept of probability is a measure of the degree of belief that an event will occur. This is not to say these values can be anything one chooses. Jaynes describes it this way:

> Probability theory is an extension of logic, which describes the inductive reasoning of an idealized being who represents degrees of plausibility by real numbers. The numerical value of any probability $(A \mid B)$ will in general depend not only on A and B, but also on the entire background of other propositions that this being is taking into account. A probability assignment is subjective in the sense that it describes a state of knowledge rather than any property of the real world; but it is completely objective in the sense that it is independent of the personality of the user; two beings faced with the same total background of knowledge must assign the same probabilities.[7]

Collecting and analyzing data is critical to supporting the ESD development. The best resources for predicting future events are past experiences and tests. While a system may have no past relevant experience for a large encompassing event, there are usually more accessible data for smaller decomposed events. Hardware, software, and human performance data are inputs to assess performance of triggers and mitigating events. It must be recognized, however, that historical data have predictive value only to the extent that the conditions under which the data were generated remain applicable. Generally, within PRAs, generic data are collected and statistically analyzed for relevance to the project at hand. These data may also come in the form of modeling and simulation results from existing tool suites. Probability distributions

are then generated to account for the uncertainty and variability. Probability distributions may also be generated from expert judgment when interviews are conducted properly.[8,9] With probability distribution data in the model, Bayesian updating techniques can be used, not only making predictions for the current project better but also building a repository of information for future projects.

The use of probability and its uncertainty distributions creates a better picture of what the community of experts knows or does not know about the events. The measure of uncertainty, and identification of the key contributors to that uncertainty, provides an understanding of the quality of results for informed decision-making. Alignment of PRA objectives with those of resilience engineering is an ideal fit as it formalizes the process of identifying and analyzing potential outcomes and determining the associated uncertainty of those outcomes.

MISSION SCENARIO DEVELOPMENT

The initial step in the developing a scenario is to understand and characterize the system. While it appears to be intuitive and trivial, this step is imperative as it increases the credibility of the results and decreases the amount of resources necessary later. Knowledge of the physical and functional layout and concept of operations, as well as fault protection schemes designed to protect, prevent, or mitigate hazard exposure conditions, is necessary to begin the PRA. All subsystems, structures, locations, operating procedures, and activities expected to play a role in the initiation, propagation, or arrest of any adverse condition must be understood in enough detail to enable construction of the models necessary to capture the possible scenarios. Past major failures and abnormal events that have been observed with the same or similar systems should be noted and studied. This information ensures inclusion of important applicable scenarios.

This system information is used to create an operational or mission scenario. The physical systems or subsystems, components, and interfaces are collected in fault tree analyses and then mapped to the functions in the ESDs that are necessary to achieve the mission scenario. Fault tree analyses provide the framework for organizing the physical system information, while the ESDs organize the functions in a logical manner, such as by procedures necessary to complete an operation or mission scenario.

Mission scenarios are the hypothetical sequence of events, constructed for the purpose of focusing attention on causal processes. They are coherent descriptions of alternative images of the future, created from mental maps and models reflecting different perspectives on past, present, and future developments. To be credible as an analytical tool, scenarios must be internally consistent, plausible, and recognizable stories. An analyst develops scenarios by analyzing how failures or disruptions can propagate through a system, leading to adverse consequences. Along the way, various pivotal events can either exacerbate that problem or mitigate it. Cause and effect relationships among triggering and mitigating events or circumstances are investigated, along with the impacts of risk-mitigating actions that may be taken. The path through the scenario is probabilistic, fulfilling the PRA approach to risk as a set of triplets (scenarios, probabilities, consequences). By adding resilience mitigations as pivotal events, the traditional PRA can be leveraged to quantify the improvements the solutions bring to increasing the likelihood of mission success.

Events in an ESD can and should be further decomposed using methods such as fault tree analysis to obtain a credible probability value. In many cases, pivotal events may require equipment to function (hardware and software), people to perform a task (procedural or repair), or testing to succeed, all of which can be modeled. Sometimes, data on failures at these higher-level events are not available, and this necessitates that fault trees be developed down to a level where data do exist to compute a probability at the higher level. Care must be taken to explicitly model dependencies, including common-cause failures.

Modeling risk scenarios begins with a description of the mission success sequence (see Fig. 3). It can be thought of as a trajectory of the system as it proceeds from start to mission completion. At each point in the trajectory, we can ask what can go wrong. These perturbations redirect the system to off-nominal trajectories

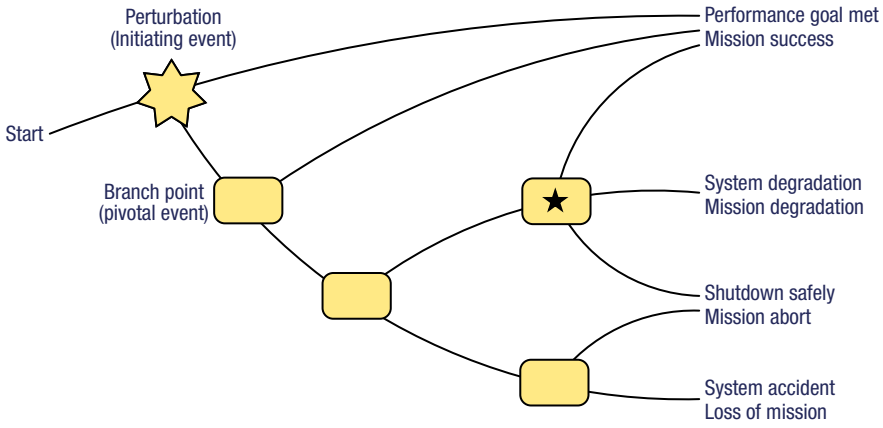

Figure 3. Conceptual mission scenario shows multiple paths from initiating event to various end states.

that could result in performance degradation or mission loss. Many engineering systems include safety or backup systems meant to be activated in response to various events. If backups work as intended, consequences are typically insignificant. However, if an event occurs and corresponding backup systems fail, there could be serious consequences. Resilience approaches can provide additional mitigations through alternative operating procedures or improvement in the rapidity of a failure recovery. If the mitigation is able to restore functionality within the temporal threshold values, the mission can still succeed. For instance, the branch point (marked by a star in Fig. 3) is an event that can swing the system from success to degradation or failure depending on how fast the mitigation restores functionality. The analyst would focus on such an event as a potential place to increase the resilience of the system.

Scenarios are useful tools in articulating key considerations, assumptions, and constraints. They provide a platform to blend qualitative and quantitative knowledge of systems and their interactions. Analysts still need to be cautious to avoid common traps such as narrowly examining a situation, applying assumptions inconsistently, or not fully documenting assumptions, thereby reducing transparency. One thing to note about scenarios in this context is that they are meant to describe a class of situations that can occur. They are not meant to explicitly describe every possible permutation of events, an infinite set of permutations.

Technical risk analysts formulate scenarios using ESDs that are supported with embedded fault trees. ESDs are inductive logical constructs showing the progression of an initiating event through a series of uncertain events, system elements, and procedural steps to end states (consequences). ESDs provide an excellent visualization of scenarios and facilitate communication among analysts, engineers, operators, and managers.

Figure 4 is an illustration of an ESD with a typical structure. The scenario shows events leading to the system staying in a state that fully meets requirements (OK), a state with degraded capability (Degraded), and a state showing a loss of mission (LOM). Scenarios are developed for each threat or adverse condition while the analyst asks about probability of the system avoiding or absorbing damage, about other assets that can fulfill total or partial functionality, and about recovery and reconstitution events. These questions are answered quantitatively with conditional probabilities and their associated uncertainty distributions. End-state probabilities are determined separately for each performance metric, which may result in many scenarios for the same adverse action.

The top row of this ESD examines whether the system will be affected by the initiating event. The second row questions the system's ability to recover from the attack and the level of functionality that can be reacquired. The third row shows paths where the system attack is not going to recover and the system must be augmented by capability from another system.

The probabilities can be determined in many ways depending on the context of the question. Figure 5 shows an example ESD and various models used. This scenario, stemming from a hypothetical laser attack on a spacecraft, explores the response of an optical sensor to meet imaging requirements. We ask questions about the spacecraft bus (could increased heat load on the solar array cause failure?) and then the sensor (could material liberated from the spacecraft deposit on the optics in sufficient quantity to cause failure?). We then examine whether a heater on the sensor can liberate material from the optics to carry out its mission.

The figure shows that three different methodologies—fault tree analysis, physics modeling, and expert opinion—are combined to determine the probabilities of the pivotal events. The analysts do not need to tie themselves to any methodology. A fault tree of the

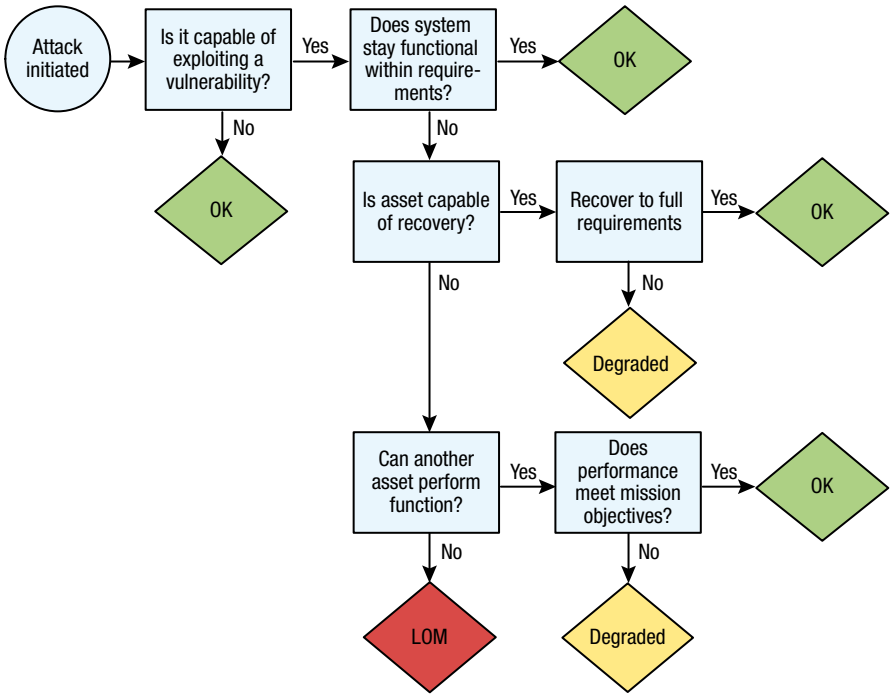

Figure 4. Example ESD suggests the types of questions the analyst needs to address.

spacecraft is used to determine bus failure from damaged solar arrays. The probability of the sensor functioning is the result of several physics models of the relationship between the laser properties and the sensor. Finally, the recovery of the sensor is addressed with subject-matter expert input in the form of a probability distribution signifying the expert's degree of belief in the probability of recovery, which is informed by past experience and/or a modeling exercise.

A tool like this allows for the exploration of alternative mitigations to such an attack. Analysts can compare different materials on solar arrays (option A) or a shutter to cover the sensor (option B). Note, a shutter would temporarily prevent the system from carrying out its mission, necessitating a trade between the need to gather data during the attack or keep the sensor ability safe for later use (see Fig. 6b).

For each of the scenarios, the feasibility, cost, and schedule impacts can be quantified through traditional methods. Through the use of PRA, the impacts to improved mission success can all be evaluated against one another as shown in Fig. 6. This approach provides decision-makers a quantifiable and comparable assessment of a broad spectrum of solutions (e.g., procedural, hardware, software, or maintenance solutions) to inform the best design or mitigation change to achieve the desired mission results. Additionally, after the system is built, if new threats emerge, system or procedural changes can similarly be evaluated to assess the impact to improved mission success.

UNCERTAINTY

Quantitative analyses of system resilience and the phenomena occurring in many engineering applications are based on mathematical models that depend on a number of assumptions and approximations. Systems under analysis cannot be characterized exactly—knowledge of the underlying phenomena is incomplete. This leads to uncertainty in both the values of the

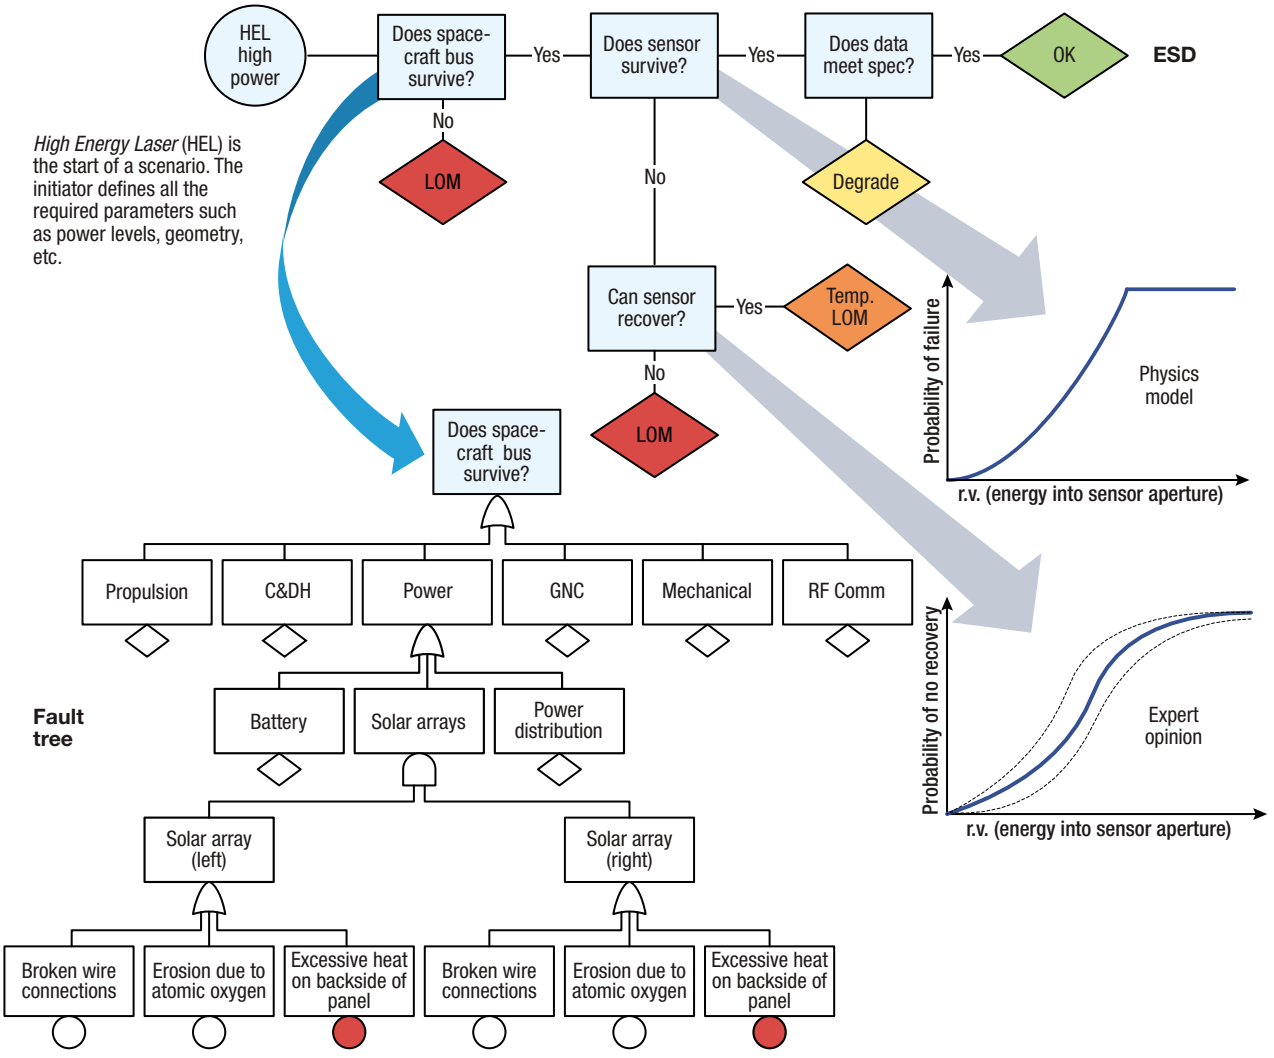

Figure 5. ESD events can be quantified using a variety of models (fault trees, physics models, expert opinion).

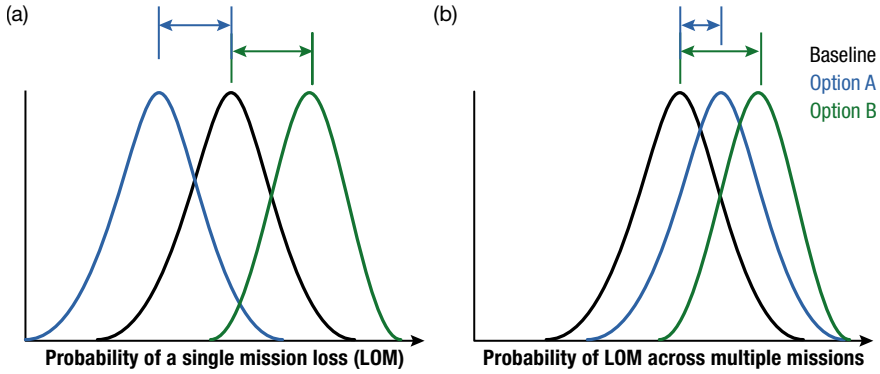

Figure 6. Comparative loss of mission (LOM) results due to alternative resilience options.

model parameters and the hypotheses supporting the model structure defining the scope of the uncertainty analysis. Uncertainty is an unavoidable reality affecting the behavior of systems, particularly with respect to their limits of operation.[10] Despite how much dedicated effort is put into improving the understanding of systems, components, and processes through the collection of representative data, the appropriate characterization, representation, propagation, and interpretation of uncertainty remains a fundamental element of the risk analysis of any system. Following this view, uncertainty analysis is considered an integral part of PRA.

Uncertainty is split into two different types: randomness due to inherent variability in the system (i.e., in the population of outcomes of its stochastic process of behavior) and imprecision due to lack of knowledge about and information on the system. The former type of uncertainty is often referred to as objective, aleatory, or stochastic, whereas the latter is often referred to as subjective, epistemic, or state of knowledge. Probability models are introduced to represent the aleatory uncertainties, such as a Poisson model to represent the variation in the number of events occurring in a period of time. The epistemic uncertainties arise from a lack of knowledge of the parameters of the probability models. Whereas epistemic uncertainty can be reduced by acquiring knowledge of and information on the system, the aleatory uncertainty cannot, and for this reason it is sometimes called irreducible uncertainty.

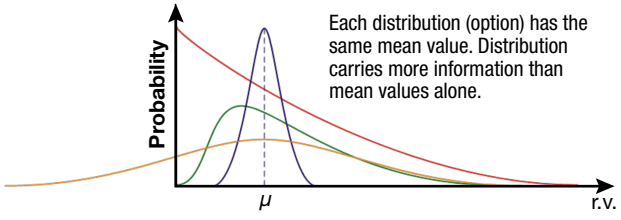

Figure 7. Distributions provide more information than point estimates when comparing alternatives. r.v., Random variable.

Resilience metrics are estimates and subject to the uncertainties discussed above, and therefore analysts must tell the decision-maker what uncertainties exist and how they affect the results. Probability distributions are the mathematically correct way to communicate to the decision-maker the assumptions and approximations and to give a sense of how reliable the numbers are. It is also easier to justify central tendency to reach consensus on a range and distribution than it is on a point value. Take, for example, the four probability distribution curves in signifying the resilience of separate alternative architectures. Each has the same mean value. So if only the number were provided to a decision-maker, all four options would be identical. However, the uncertainty about them tells a different story, both in the amount of spread and the shape.

We often think about uncertainty coming from the data used as input to the various models. The models themselves may be uncertain as well. Figure 8 shows hypothetical relationships based on the ESD example provided earlier, wherein a key laser parameter translates to the probability of success of an on-orbit sensor. The same model is shown in Fig. 9, with 5th and 95th confidence bounds on the curves in addition to the uncertainty about the laser parameter. All propagate through the model yielding a probability distribution for a result.

Once the system PRA is developed, the system's resilience to failures, natural events, and adverse actions can

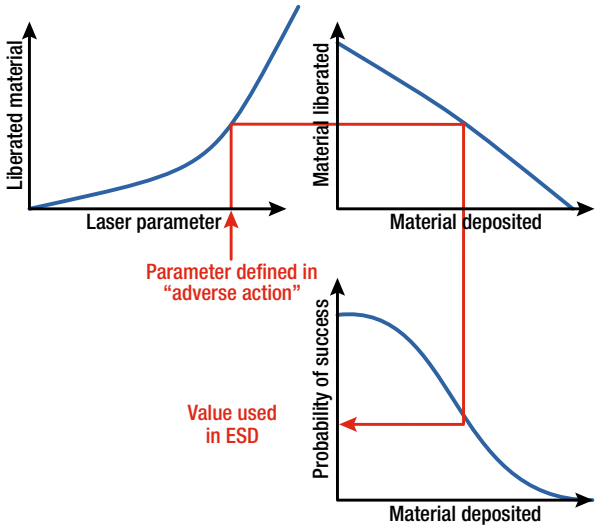

Figure 8. Probability of success is often a function of multiple physics models.

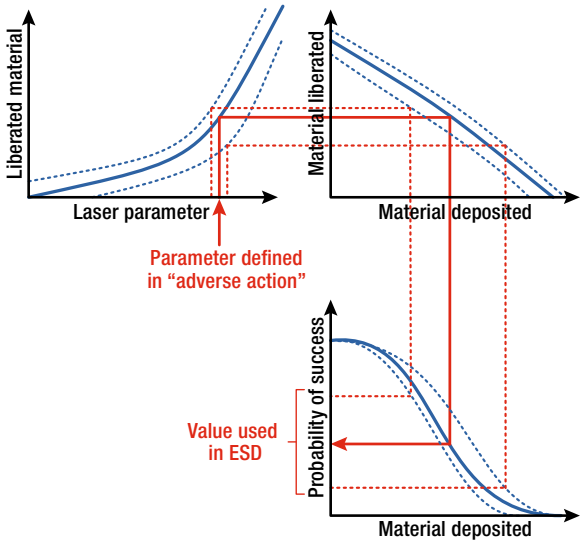

Figure 9. Physics models can show how uncertainty in initial parameters and models can propagate to probability of success.

be quantified. The metric is provided as a probability distribution where various statistics, such as the expected value and quantiles, can be derived and used for planning, engineering trades, and decisions. The ability to rank order elements with respect to their resilience significance is one of the most useful aspects of this methodology. It represents one of the major improvements over current practices. In PRA parlance, values computed to perform rankings are called importance measures.

An importance measure gives an indication of a certain component's contribution to the total system resilience. PRA analysts have created a set of importance measures to evaluate risk contributions of any given item in a model. Most applications of importance measures aim to provide management insight into three broad areas: design or redesign optimization, test and maintenance strategy development, and daily configuration control.[11] Several risk assessment texts contain detailed explanations and derivations of these measures.[1,12] They can also be applied to provide insights for a particular risk item and constituent elements.

- Risk reduction worth measures change in risk assuming that an event of interest is perfect (will not fail). In other words, it measures how much improvement can be made to a system by fixing one event.

- Risk achievement worth is the inverse of risk reduction worth in that it measures improvement possible if no credit is taken for a given component. It is the change in risk assuming that the component is not there.

- The Fussell-Vesely is a fractional contribution of a component to total risk.

CONCLUSION

Resilience design or system change considerations are able to be quantified and assessed by the increased likelihood of mission success generated. A cross collaborative team of engineering, operators, analysts, and managers may conceive numerous mitigations to improve mission success. Resilience approaches help to generate more mitigation concepts so that a system can prevent, withstand, adapt to, and rapidly recover from failures or external disruptions. Often after the ideas are generated the question is asked: What solution is the best to achieve the desired results? PRAs have proved to be an invaluable tool for industry and projects at APL. As applied to quantifying resilience, they provide a systematic, traceable, and defendable set of metrics with uncertainty.

REFERENCES

[1]Bedford, T., and Cooke, R., *Probabilistic Risk Analysis: Foundations and Methods*, Cambridge University Press, Cambridge (2001).
[2]Kaplan, S., and Garrick, J., "On the Quantitative Definition of Risk," *Risk Anal.* **1**(1), 11–27 (1981).
[3]Cutter, S. L., Ahearn, J. A., Amadei, B., Crawford, P., Eide, E. A., et al., "Disaster Resilience: A National Imperative," *Environment* **55**(2), 25–29 (2013).
[4]Jackson, S., and Ferris, T. L., "Resilience Principles for Engineered Systems," *Sys. Eng.* **16**(2), 152–164 (2013).
[5]U.S. Nuclear Regulatory Commission, *Reactor Safety Study: An Assessment of Accident Risk in U.S. Commercial Nuclear Power Plants*, U.S. Nuclear Regulatory Commission, Washington, DC (1975).
[6]Henley, E. J., and Kumamoto, H., *Probabilistic Risk Assessment and Management for Engineers and Scientists*, IEEE, New York (1996).
[7]Jaynes, E. T., quoted in Kaplan, S., "The Words of Risk Analysis," *Risk Anal.* **17**(4), 407–417 (1997).
[8]Clemen, R., and Winkler, R., "Combining Probability Distributions from Experts in Risk Analysis," *Risk Anal.* **19**(2), 187–203 (1999).
[9]Mosleh, A., Bier, V., and Apostolakis, G., "A Critique of Current Practice for the Use of Expert Opinions in Probabilistic Risk Assessment," *Reliab. Eng. Syst. Safe.* **20**(1), 63–85 (1988).
[10]Aven, T., and Zio, E., "Some Considerations on the Treatment of Uncertainties in Risk Assessment for Practical Decision Making," *Reliab. Eng. Syst. Safe.* **96**(1), 64–74 (2011).
[11]Van der Borst, M., and Schoonakker, H., "An Overview of PSA Importance Measures," *Reliab. Eng. Syst. Safe.* **72**(3), 241–245 (2001).
[12]Rausand, M., and Høyland, A., *System Reliability Theory: Models, Statistical Methods, and Applications*, John Wiley & Sons, New York (2004).

Clayton A. Smith, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Clayton Smith is a member of the Principal Professional Staff at APL and has more than 30 years of experience analyzing systems from risk, reliability, and safety perspectives. These systems included NASA and DoD missions, payloads, ground communication systems, air traffic control systems, and missile systems. He is developing approaches to assess intentional threats against space assets using probabilistic risk analysis and game theory techniques. He created and managed NASA's International Space Station Program probabilistic risk assessment specifically geared toward quantifying the safety risk during operations. Clayton is currently the reliability engineering lead for APL's PSP mission. He received a B.S. in aerospace engineering, an M.S. in engineering management, and a Ph.D. in reliability engineering all from the University of Maryland. His e-mail address is clay.smith@jhuapl.edu.

Timothy J. Allensworth, Force Projection Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Timothy J. Allensworth is a member of the Senior Professional Staff working in the Ocean Engineering Program at APL. He is the assistant program manager for advanced capabilities within the Ocean Systems and Engineering Group. His interests include submarine communications and navigation, undersea sensor technologies, and advanced processing techniques such as machine learning. He received a B.S. in aerospace engineering from Purdue University, qualified as a submarine officer aboard USS *Annapolis* and as a nuclear engineering officer through Naval Reactors, and holds an M.S. in applied physics from Johns Hopkins University. His e-mail address is timothy.allensworth@jhuapl.edu.