Going Beyond Reliability to Robustness and Resilience in Space Life Support Systems

Harry W. Jones[1]
NASA Ames Research Center, Moffett Field, CA, 94035-0001

The words reliability, robustness, and resilience are often used interchangeably to describe tough and dependable systems but the distinctions between them suggest how to design more serviceable space systems. Reliability is simply the quality of consistently performing well. A system that dependably meets its design requirements in the specified environment is reliable. The designers may not consider themselves responsible for failures under unanticipated conditions. Robustness is the capability of performing without failure under a wide range of conditions, which can go beyond the expected range to include possible off-nominal conditions. Resilience is the ability to recover from or adapt to unanticipated damaging events, such as failures, accidents, external disruptions, and repurposing. Such changes can invalidate the usual operating assumptions and cause system failure.

Reliability, robustness, and resilience describe dependable performance under increasingly difficult conditions, first the specified environment, then a wider possible environment, and finally unanticipated damaging conditions. These three qualities are increasingly desirable and increasingly difficult to achieve.

Engineering for resilience would design systems that can ignore or repair failures, survive accidents, and recover from unanticipated disruptions. Increasing the resilience of space systems would greatly increase space crew safety. Improving reliability and robustness requires dealing with known problems, but improving resilience requires implementing a general approach to reducing the impact of unknown future events.

The need for robustness and resilience has been stated for decades but little has been done. Systems designers often assume that they understand everything they need to know. The potential failures caused by changes, failures, accidents, unknown environments, and unknown unknowns can be ignored. Such overconfidence can lead to neglect of reliability, robustness, and resilience.

Nomenclature

CCF = Common Cause Failure
ECLSS = Environmental Control and Life Support System
ESM = Equivalent System Mass
ISS = International Space Station

I. Introduction

THE need for reliability, robustness, and resilience has been proclaimed for decades and most space engineers readily agree. However, little actual work has been done and results are scarce, especially in life support. The problem seems to be that systems designers plunge into the design process believing that they know everything they need to know to do their job. They assume that they understand the requirements, technologies, designs, architectures, integration, testing, operations, and environments involved. The potential problems of requirements changes, hardware failures, improper operations, accidents, unanticipated environments, and unknown unknowns are easily neglected. Systems designers need how to set requirements for reliability, robustness, and resilience and how to design to implement them. Crew safety depends on this.

[1] Systems Engineer, Bioengineering Branch, Mail Stop N239-8.

II. Definitions of Reliability, Robustness, and Resilience

Reliability, robustness, and resilience are similar concepts but their dictionary definitions indicate some significant differences. Reliability is the quality of being dependable, trustworthy, or of performing consistently well. Reliability requires working as expected in normal, well known circumstances. Robustness is the capability of performing without failure under a wide range of possible conditions. Robustness implies strength and toughness under potential off-nominal conditions. Resilience is the ability to recover from or adjust easily to an unanticipated accident or change.

The technical literature has this definition, "Reliability can be defined as the probability that a device, product, or system will not fail for a given period of time under specified operating conditions." The system performance is defined for specific operating conditions.[1]

A general definition of robustness comes from an ecological perspective: "Robustness is the persistence of specified system features in the face of a specified assembly of insults." [2] [3] The definition is systems engineering is similar. "Taguchi's general approach to 'robust design' is to provide a design that is insensitive to the variations normally encountered in production and/or in operational use." [4] Robustness goes beyond reliability by including an ability to perform satisfactorily, if not exactly according to the stated performance requirements, under conditions that go beyond the expected environment to include know problems and anticipated variations.

Resilience goes further than robustness, requiring some ability to perform after the occurrence of unspecified problems and changes that violate the design assumptions. "These changes were couched in various ways—anomalies, disruptions, discoveries—but they all ultimately had to do with changes in underlying assumptions. Invalid assumptions, whether due to unexpected changes in the environment, or an inadequate understanding of interactions within the system, may cause unexpected or unintended system behavior. A system is resilient if it continues to perform the intended functions in the presence of invalid assumptions." [5] Resilience engineering is concerned with building systems that are able to circumvent accidents through anticipation, survive disruptions through recovery, and grow through adaptation. [6]

III. Advancing from Reliability to Robustness and Resilience

Advancing from reliability to robustness and resilience goes through stages of more realistic assumptions, more inclusive models of failure causes, and wider and stronger approaches to ensure crew safety. This progress requires increasingly better models of operational reality and wider understanding of failure causes. The models are improved by encountering actual failures not predicted by the current models. This process can be described in five stages shown in Table 1.

Table 1. Five stages in advancing from reliability to robustness and resilience.

Stage	Requirement	Assumption	Environment	Approaches
1	None	Reliability no issue	Nominal conditions	Rationalize neglect of reliability
2	Basic reliability	Failures all due to parts failures	Nominal conditions	Reliability specification Reduce parts count Use high reliability parts Use spares and redundancy Plan low level repair Fault tolerance
3	Advanced reliability	System level failures Interface problems Design errors Operator errors	Nominal conditions	Reliability growth test Failure analysis Redesign Life test Diverse redundancy
4	Robustness	Anticipated challenges	Off nominal conditions	Expand reliability specification to prevent or tolerate known risks
5	Resilience	Unanticipated failure modes	Unexpected environments	Excess processing capacity Buffering and storage Reconfigurability Multipurpose tools and materials Reduce vulnerability Add monitoring and control

At stage 1 a project may neglect reliability, as was done for many years in NASA space life support. Systems were expected to be compared using Equivalent System Mass (ESM), supposedly having analysts use their expertise to adjust ESM so the systems "satisfy the same life support … reliability and safety requirements." [3] Reliability analysis and even the suggested subjective adjustments are not often done. [7]

Stage 2 is traditional basic reliability analysis as it is routinely practiced. It usually leads to greatly underestimated initial failure rates. These rates can be considered lower bounds on the final failure rate after stage 3. Stage 3 is advanced reliability including reliability growth and life testing. The failures are discovered by test and are removed by redesign and ameliorated using maintenance procedures. Some low rate and difficult to correct failure modes can be accepted. Diversity is used to avoid Common Cause Failures (CCFs). Stage 4 recognizes that some failures are caused by expected internal or environmental variations. The requirements are expanded to require operation in these circumstances. The designers may not consider themselves responsible for failures under unanticipated conditions, such as extreme environments or human error. Stage 5 is based on the recognition that some serious failures are surprising and unpredictable. Recognizing and coping with unpredictable failures requires changing assumptions and the system model built on them. The general approach includes simplifying and strengthening the system and adding capabilities beyond normal operation.

The advance from seeking reliability to adding robustness and resilience is driven by an advance from using basic parts-based reliability improvement models to including first known and later unknown failure modes. As the failure mode models become more realistic and inclusive, the failure reduction actions must become more extensive. Engineers work from a system model that is always limited and partial. Lower-level subsystems and components are represented as black boxes defined by their nominal inputs and outputs. Higher level and external effects are ignored. Both deep internal and external changes can be source of disruption. If the model is expanded to include known possible impacts, the model can be used to design for robustness. If the wider model is further expanded to allow for unknown sources of disruption, this awareness can be used to design for resilience. The breakpoints between designing first for reliability, then adding robustness, and ultimately resilience are determined by the reliability model transitioning from normal failures in nominal conditions first to predictable failures in defined degraded conditions and then to unpredictable failures in actual operational situations.

When is resilience needed? It is when unanticipated disruption is a significant concern, perhaps due to a history of unexpected failures. Unexpected accidents may cause loss of service, damage, and even loss of life. These notably occur in large complex systems, such as communications networks and power grids. Such accidents can be caused by human actions, diverse and unpredictable environmental events, and possible deliberate attack. They require both rapid automatic and thoughtful human response to provide resilience. Small systems such as computers and cell phones also may require resilience. Computers sometimes use a redundant array of reconfigurable memory, but this is robustness to a known failure mode. Cellular phones have multiple communications channels for resilience, including multichannel cellular, internet Wi-Fi, short range Bluetooth, and possible future worldwide satellite links. Although these channels have different purposes, ranges, and bandwidths, they provide alternate methods of communication. Resilience engineering requires accepting ignorance about the future and designing systems to manage unexpected events in unpredictable forms.

A. All Models are Partial Models

All models are partial, with simplifications and omissions. Good models are useful when their assumptions hold. There is a fundamental split in thinking about models and reality. Some hope that models can ultimately include all the factors needed to provide a complete true picture of reality. Others think that all models are finite and incomplete and can be falsified by omitted factors or future developments. For a true understanding of reality it is necessary to accept the incompleteness of models, the occurrence of fundamental changes, and the truth of assumptions only in limited situations. This understanding is fundamental in achieving reliability, robustness, and resilience.

Critical thinking and challenging assumptions is the key to good engineering. Scenario analysis may be helpful to consider the implications of alternate assumptions. Even if no design or operational changes result, there is a benefit is analyzing and planning coping mechanisms before a change occurs. The widely accepted models of the financial system made the 2007-2008 financial crisis seem impossible, yet it occurred. The general problem is that idealistic models built on favorable assumptions describe normal operations and simply do not include the real-world failure modes. They can be more traps than guides.

Engineers who believe that they can fully understand the design situation and correctly predict the future do not see the need for resilience. Adding resilience increases cost, decreases efficiency, and admits fallibility, which are much less attractive than optimistic or even over-confident assumptions. Some respect for the challenges of reality and even fear of failure are required to build resilient and safe design. Sometimes the margins, buffers, and over

capacity that engineers know are required for resilience must be disguised to prevent their being removed to cut cost. The situation is similar with project cost and schedule where contingency reserves are often not allowed.

B. Known Low Probability Challenges Can Be Improperly Ignored

Low probability events are often not adequately considered. The Challenger disaster occurred in 1986 but it was not until after Columbia was lost in 2003 that the impact of the loss of water and oxygen resupply to the ISS was analyzed and plans were made to reduce the crew until water and oxygen recycling systems could be delivered. One or two tsunamis occur every year and major ocean wide tsunamis occur every 10 or 20 years. A devastating tsunami occurred in the Indian ocean in 2004, but the extent of the 2011 Japanese tsunami was a surprise.

C. Distinguishing Between Anticipated and Unanticipated Failure Causes

There are always many unanticipated potential failure causes. Planning to deal with failures during operations can include providing the parts, tools, procedures, and training to repair a few failures or a few dozen failures. However, there can be hundreds or even thousands of potential failure causes. The process of increasing robustness requires taking focused steps to prevent some specific anticipated challenge. If a risk is accepted because its probability or impact are low, it is still anticipated. Even without specific prevention or mitigation, the property of resilience can help deal with a known risk if it occurs. Resilience can help deal with both unanticipated challenges and anticipated but unmitigated challenges.

It is important to recognize that unanticipated challenges to firmly accepted assumptions do occur but are very difficult to recognize and cope with. Actual unanticipated events are sometimes described as 30 sigma events, black swans, or 100-year hurricanes. All recognized assumptions should be criticized and alternative scenarios developed under the assumption that the assumptions are violated. If it is understood that a certain assumption is being made and that if wrong it would cause a failure, the problem becomes anticipated. Unanticipated challenges are caused by incorrect assumptions that are not recognized as fallible assumptions but rather accepted as undisputed fact. Surprise events create new facts that require a revised model to explain them. The new model may be more complicated, more negative, and less acceptable than the earlier failed model. In extreme cases, believers in the conventional model may be captured in a false narrative based on implausible assumptions that produces a continuing storm of failures.

IV. Methods to Achieve Reliability, Robustness, and Resilience

The increasingly difficult goals of reliability, robustness, and resilience can be achieved by building successively on earlier gains while using increasingly stronger analysis and design techniques. Ideal success seems unattainable. Often simply gaining high reliability is too difficult in real world situations.

A. Methods to Achieve Reliability

Initial reliability analysis usually assumes a plausible but simplified ideal design problem. The system architecture is familiar and proven and all the components are standard with a known low constant failure rate. The system design problem is to trade off performance, reliability, cost, etc. for the customer's optimum product. If the component failure rates are low, the system failure rate is simply the sum of the component failure rates. During operations, components will fail at the known rates and be replaced by spares. If the components have well established failure rates, the number of spares needed to have the required probability of having sufficient spares can be exactly computed. The actual number of failures on the mission is a random variable.

This approach is standard, widely accepted, and often all wrong. The assumptions can be violated and the prescribed diagnosis and repair processes not followed or simply not possible. The International Space Station (ISS) Environmental Control and Life Support System (ECLSS) is a useful example. Its system architecture is familiar, having been used in human closed chamber tests since the 1960's, but it is not sufficiently developed for future missions since the operational ISS ECLSS still has an unacceptably high failure rate. Some components such as water pumps were new designs for reduced mass, had limited testing, and have had high failure rates on ISS. The ISS system design process deemphasized building in reliability to save cost while relying on the ability of the crew to repair failures. The crew time required for troubleshooting and repair has been much higher that initially estimated. This is clearly not optimum from the crew's point of view. Generally, and on the ISS, most of the failures that occur during operations are not due to component failures, but to systems level issues, interfaces, environmental problems, human error, etc. It is easy to trace a failure due to a single component, since these deletions simplify system behavior, but it can be nearly impossible to find faults such as shorts, interference, and intermittent failures that expand and complicate system behavior. Some long and difficult trouble shooting processes have occurred in ISS ECLSS. Finally, the initial estimated number of spares has been too low for components with unexpectedly high failure rates.

Each part of the ideal reliability story can be wrong and cause problems. Investigating such reliability problems has led to better approaches to achieving reliability and increased interest in robustness and resilience as ways to get past the limitations of a narrow focus on reliability.

Better reliability analysis starts with more realistic assumptions. Any system architecture is at least partly new and may have undiscovered problems. Some components are new without a proven low failure rate. The design problem is still to obtain the optimum cost-effective reliability as determined by system trade-offs. As before, the first step is to estimate the system and component failure rates, but with the understanding that some need to be verified by test. Some components will be standard with a known low constant failure rate, but others will be new and must be tested. If the new components have a low constant failure rate, the failure rate can be determined by testing multiple units for a shorter time, but in order to confirm that the new components do have constant failure rates without wear out problems, life tests probably should extend to several times the expected mission duration. More usually, new components will have a very high initial failure rate, sometimes called "infant mortality," due to design and manufacturing errors, requirements mistakes, etc. In the initial "reliability growth" phase of testing, design mistakes are found and fixed. Reliability growth can be terminated when the current failure rate is satisfactory, perhaps because it is close to the initially estimated failure rate. Life testing can continue to more accurately determine the failure rate, and this rate can be used to estimate the number of spares required on a fixed length mission. After the component failure rates have been reduced and determined, a similar reliability growth and life test approach can be applied to the full integrated system.

B. Methods to Achieve Robustness

Since reliability requires that a system provide a defined performance under specific operating conditions, the requirements and conditions for reliability demonstration tests are clearly determined. An additional requirement for robustness can expand this envelope. Minimal robustness would require undegraded performance under normal operational variations, which should be part of a reliability specification. Strong robustness might require some minimal performance under significant or perhaps even extreme external failures, environmental challenges, or human error. Normally, systems would be expected to fail under such drastic and unanticipated challenges.

The needs for system robustness must be anticipated and included in the system requirements, by specifying the needed performance and the expected challenges. Since the difficult external conditions are different than usual, they would typically have limited duration, so possible responses include buffer storage, excess capacity, and demand reduction. Additional testing would be needed to verify robustness requirements.

C. Methods to Achieve Resilience

Achieving resilience is more difficult than achieving reliability and robustness. Requirements can be written and tests devised to achieve reliability and robustness, but resilience is a much less defined, more open ended capability.

Resilience is the ability to survive unanticipated damaging events. The word "unanticipated" means that there is no list of damaging events to include in a specification. Resilience is the capacity to endure unexpected change. Change cannot be predicted. The future is unknowable. Resilience is essentially some ability to isolate, recover, or adapt to change.

When a system behaves in an unexpected way, when it fails for reasons not explained by the model of the system, the system must be understood in a new way. The assumptions it is based on are not completely valid. Assumptions may have changed or many never have been valid. For example, the usual but false assumption that most failures are due to parts failures means that complex system level failures are often unanticipated and difficult to recognize.

It seems that the opposite of resilience, system fragility, is due to designs being based on narrow, rigid, and obsolete or invalid assumptions. Unexpected system failure causes demonstrate that there must be some error in the assumptions. You would expect fundamental errors such as neglecting reliability or over reliance on crew time to be identified and become published lessons learned, but sometimes this does not happen. Failure analysis tends to be a political process of assigning blame, with a tendency to focus on low level, last minute individual actions rather than questionable management assumptions such as relying on ISS on-board repair to compensate for limited reliability effort. This means that chronic problems such as insufficient preflight testing that can actually be anticipated often are not. It is difficult to challenge basic assumptions because they tend to form part of a complete world view. Although it is difficult, the process of designing for reliability and robustness should be based on the soundest possible assumptions. Even if all are valid, some may change. Some key assumptions may be unrecognized and possibly invalid. Designing for resilience should provide capability for survival even if the assumptions are incorrect.

V. Past Work on Resilience in Life Support and Space Engineering

Some interesting work has been done on resilience in life support and space engineering. These papers share many useful ideas and are summarized below.

A. Measuring the Resilience of Advanced Life Support Systems

Bell, Dearden, and Levri investigated the resilience of life support systems in 2002. [3] They claimed that resilience is the most crucial property of a life support system. Dynamic simulation of a water revitalization system was conducted with component failures of known probability. Bell et al. noted that the terms resilience and robustness were used with multiple meanings. Their usage differs from that in the current paper. Their paper title mentions measuring resilience, but the actual simulation includes only component failures with known probability, which is termed robustness in the current paper.

It is observed that the resilience or robustness of a system is not only a hardware quality, "but also of the way that hardware is controlled." Analysis of the resilience or robustness therefore requires dynamic simulation including the control system. A specified set of faults occurs with defined probability and the system reacts overtime. The system may perform improperly and recover or ultimately fail. The specific system outputs, in this case including the quality and quantity of water, can be used to describe the system resilience.

Equivalent System Mass (ESM) "is unable to reflect the improvement in a system that is made by simply changing the controls approach." "The ability of a system to respond to off-nominal events is not captured in ESM." "One of the main points in this paper is that resilience is a dynamic property of a system as a whole, and as such, is not adequately measured by metrics of nominal operation via static computation."

A dynamic model of a simplified a water revitalization system was developed to analyze its resilience. The technical performance under different failure scenarios helps to indicate resilience. A Markov chain model was also used to measure resilience. Crew safety depends on the resilience of life support. Dynamic simulations can investigate resilience early in the design process. [3]

B. Defining ECLSS Robustness for Deep Space Exploration

Escobar, Nabity, and Klaus review the literature on robustness with emphasis on deep space life support. [7] They consider robustness, reliability, resilience, and survivability and note the overlapping use of these terms. They propose using the single term robustness to characterize and improve life support performance, but they expand robustness to include reliability, resilience, and survivability. "Robustness includes the ability to function during expected circumstances, as well as in a changing environment, or during unanticipated events." Response to unanticipated events is called resilience.

Robustness is proposed as a life support optimization objective, defined as "the ability to maintain habitable conditions for crew survival and productivity over the mission lifetime under a wide range of conditions." Design considerations are discussed. Rather than minimizing cost, they suggest instead establishing an acceptable cost threshold and designing towards the maximum possible performance. Uncertainties include probable, possible, and plausible but rare events. The cost of uncertainty is both the cost of failures and the cost of measures to reduce uncertainty, such as redundancy. "(T)he need for improved reliability" has been "strongly emphasized," but "the number of reliability based trade studies in the literature is small." *"*A review of systems design literature reveals a variety of characteristics that contribute to system operability during unanticipated conditions, such as redundancy, modularity, failure mode isolation, network connectedness, simplicity, passive control, wide design margins, maintainability, accessibility, interface commonality, capacity for lower levels repair, functional diversity, etc."

A quantitative robustness metric would facilitate optimization. Life support metrics are reviewed, including mass, life cycle cost, closure, reliability, and the probability of mission failure, but a robustness metric is not yet defined. The paper provides a broad well researched foundation for future investigations. The first step of establishing definitions and a frame of reference was the objective of this paper. Future needs include quantifying life support robustness, developing methods to assess robustness, and providing methods to improve life support robustness. [7]

C. Conceptualizing resilience in engineering systems: An analysis of the literature

Wied, Oehmen, and Welo conducted a survey of the idea of resilience in engineering systems states, "It is now widely recognized that many important events in the life cycle of complex engineering systems cannot be foreseen in advance." And, "(R)esilience theory prescribes presuming ignorance about the future, and designing systems to manage unexpected events in whatever form they may take." [8] This survey analyses 251 definitions of forms of resilience, which include recovery, assimilation, adaptation, graceful degradation, acceptable deterioration, and survival. The main distinction between forms of resilience is between passive absorption and active recovery.

Absorption includes an ability to resist, withstand, or tolerate unexpected adverse events to achieve stability and sustain the original state. Recovery can be reacting or flexibly adapting to achieve an adjusted new normal state. Recovery can stretch to improvement based on recognizing the surprising new situation and adding capability to deal with it. Adaptation may be needed for survival. Sometimes impaired performance or graceful degradation is all that can be achieved by the most vigorous adaptation. Resilience can vary from passive resistance to active countermeasures and may achieve normal or even improved operation, with the possibility of degraded operation or even failure. The forms of resilience suggest the principal design approaches for any resilient engineering system.

D. Designing resilient systems-of-systems: A survey of metrics, methods, and challenges

Uday and Marais present principles to guide the design of resilient systems. [9] [These include both familiar lower level strategies and broader concepts. The basic ideas include providing redundancy, diversity, robustness, and repairability. They further suggest providing localized capability so that failures do not cascade, continually monitoring status and performance, providing full status communication between systems, controllers, and users, having humans in the loop to provide context based understanding and create options, and attempting to prevent disruption as well as provide resilience.

E. Towards a Conceptual Framework for Resilience Engineering

Madni and Jackson provide useful concepts for resilience engineering. [6] Resilience engineering is proactive, attempting to build systems that can reduce the disruptive effects of unanticipated changes and accidents. In the ideal case, resilient systems are able to anticipate disruption, survive, recover, and even improve through long term adaptation. Most unanticipated disruptions are caused by human error or external events, rather than internal component failures. External disruptions include operational problems such as supply chain failure and natural disasters such as hurricanes, tsunamis, earthquakes, and fires.

Systems designed to survive previously encountered problems are often not criticized if they fail unanticipated disruptions. This suggests that systems should be made resilient, beyond merely reliable. Claiming that the failure probability is acceptably low is not enough. The system needs to be able to recover from unexpected perturbations, disruptions, and degradations of the operational environment. The system design and verification must be based on known potential disruptions, but attempts must also be made to make the system resilient for unknown threats.

Some design principles that go beyond verifiable requirements can enhance system resilience. Useful resilience heuristics include physical redundancy, functional or diverse redundancy, safety margins, ability to reorganize and adapt, human-in-the-loop, design simplification, graceful degradation, drift monitoring and correction, and monitoring the system and environment. [6]

F. Engineering Resilient Space Systems

Murray et al. found that the essential definition of resilience is “adaptation in the presence of changing circumstances.” [5] The changes “all ultimately had to do with changes in underlying assumptions.” “Invalid assumptions, whether due to unexpected changes in the environment, or an inadequate understanding of interactions within the system, may cause unexpected or unintended system behavior. A system is resilient if it continues to perform the intended functions in the presence of invalid assumptions.” [5]

Any system design requires making trade-offs between performance, risk, and cost that help determine system resilience. The use of resilience technologies should reduce risk but typically increases cost and reduces the system efficiency or benefit/cost ratio. It is suggested that “the real objective should be to maximize the expected value,” but this is difficult when human life is at stake. An acceptable risk to crew could be used instead.

Murray et al. suggest that resilience should be measured by changes in cost, risk, and system performance. Principled architectural approaches should be used to manage complexity. A literature review found little work on architecting systems for resilience and almost no work on architecting resilient space systems. Lessons must be learned from other disciplines, including biology, ecology, and computer science. Some general resilience heuristics include diversity, functional and physical redundancy, graceful degradation, and adaptability. Diversity is the incorporation of somewhat different capabilities in dissimilar hardware, hardware, software, and systems. Functional redundancy is similar to diverse redundancy that duplicates a function in different hardware. Physical redundancy involves providing duplicate hardware as spare systems. Graceful degradation is the ability to continue to perform through unexpected faults and environmental interactions but at a less than nominal level. Adaptability goes beyond repairability to include increased performance or greater robustness to the previously unanticipated change.[5]

VI. Unanticipated Challenges - Black Swans and Radical Uncertainty

In addition to the engineering analysis of the need for resilience in response to unanticipated challenges, there is extensive work on the need to consider radical uncertainty and improbable events in economic analysis and organizational decision making.

A. The Black Swan: The Impact of the Highly Improbable

Nassim Taleb's best selling book describes a black swan as an unpredictable event with a massive impact. [10] [The terrorist attack of 9/11 was a black swan but the COVID-19 pandemic was not. Pandemics were mentioned in Taleb's 2007 first edition and pandemic planning is required for some regulated businesses. Known but improbable events such as pandemics require planned robustness, unimaginable black swans like 9/11 require resilience but their occurrence transforms them into known, improbable, but usually planned for challenges. Taleb's 2010 second edition includes an essay, "On Robustness and Fragility," about coping with and adapting to black swans. The fundamental approach to resilience is redundancy, including spare parts, excess processing capacity, and functional redundancy. It is also useful to isolate the effects of harmful changes and to provide the ability to use existing capabilities to perform additional functions.

Taleb explains why unpredictable events can have a massive impact. Hardware, software, and even social systems are constructed using a model and narrative describing how they are expected to operate. The models are always based on a set of accepted assumptions, some explicit and others unstated. Most systems and their models will fail if a key assumption no longer holds. Surprising events such as unanticipated failures directly challenge the working model and its assumptions. Over confidence in the model encourages optimizing efficiency and building in excessive risk that ignores the fragility of the resulting design. Redundancy and slack produce resilience that reduces risk and enhances survivability.

The future is unknown. We do not know today the challenges and opportunities of tomorrow. Disasters are caused by people who do not understand reality and use bogus assumptions and limited models to instill false confidence. The objective of resilience is to confine the impact of human mistakes and miscalculations. But it is possible to benefit from unpredictable change if a system is designed with adaptability and optional responses. [10]

B. Radical Uncertainty: Decision-Making Beyond the Numbers

Kay and King consider decision making under radical uncertainty. [11] Decisions and plans of action are usually based on narratives, stories that help us explain the operative reality. A system description is usually based on a reference narrative which relies on optimistic but hopefully realistic expectations. Shared rational narratives are needed to persuade others to cooperative action. The reference narrative should be explained and challenged. Designing and optimizing a system usually requires making simplifying assumptions, but in a world of radical uncertainty, assumptions can change, and wrong assumptions produce wrong results. The belief that an accurate and complete model can be developed is simply a denial of the fact of radical uncertainty. One way to cope with an uncertain future is to develop and prepare for alternate narratives, another is to rely on resilient adaptation to the unimagined. Good strategies for an uncertain world avoid dubious assumptions and over simplified models. Narrowly optimized systems reduce the options needed for resilience. Options must be built in to deal with future challenges, known and unknown. The value of options is increased by uncertainty and potential for change.

Dominant narratives develop and are defended primarily by repetition and secondarily by attacks on any who disagree with them. They are repeated as if there are no alternatives or any possibility of error. A predominant system failure mode occurs when respect for leaders and deference to authority displaces objective fact as the determinant of truth. The willingness to challenge the accepted narrative is key to progress and to effective decision making. Challenging narratives tests the weaknesses of the foundational assumption, of the decisions, and of the systems based on them.

Good strategies for an uncertain world avoid using pretended knowledge contained in simplified models based on untested assumptions. What is needed is a robust and resilient capability to deal with many contingencies through broad diversification. With false assumptions, there is a risk that the system will fail to fulfill important elements of the narrative. The beliefs embodied in a failing narrative can change abruptly when many people see the facts produced by an encounter with reality. The collapse of a false narrative can produce a feeling of loss, even desolation. Rather than indulging in fantasy, it is better to acknowledge that we do not know what the future will hold. [11]

VII. Conclusion

Deep space systems, especially life support, should provide the best crew provision and highest safety possible. This requires designing first for reliability and then adding effort to achieve robustness and resilience. Reliability provides the required performance in the specified environment. Robustness extends this to include satisfactory operation under known internal and external changes beyond normal conditions while resilience is the ability to cope with unanticipated challenges.

The actions taken to design to prevent failure depend on the failure model used by the designers. Although it is possible to severely discount the possibility of failure, the basic reliability approach assumes that all failures are due to parts failures and strives to reduce parts count and increase parts reliability. Advanced reliability allows for design errors, system level interactions, requirements errors, and operator error and conducts initial testing and failure analysis to improve the design. Robustness extends the reliability specification to deal with anticipated deviations from nominal conditions and requires expanding system performance to cope with them. The realization that designers do not know all the potential challenges and failure modes leads to designing for resilience by providing general purpose methods such as excess capacity, buffering, and reconfigurability.

Partial models assume a limited set of failure causes and prevent action to mitigate out-of-model failures. System designers should use increasingly wider and more inclusive failure models. Starting with ignoring the possibility of failure, the models add parts failures, then system design, system interface, operator error and other failures, next anticipated violations of expectations such as deficient parts or off nominal environment, and finally, unanticipated failure modes. Each of these models can be held with conviction and used as the basis of design. Some systems are actually extremely reliable and others including space systems have been unexpectedly found to be much less reliable that expected. The correct model to use depends some on the actual failure history but much more on the cost of failure. Despite their very different failure rates, both jet planes and launch rockets have more engines than they need. Over optimistically using a simpler failure model that ignores major failure modes permits more complex and innovative designs and cuts cost, but sometimes accepts excessive risk and leads to surprising failures. The key to success is realizing that we do not know what may happen in the future.

References

[1] Shishko, R., *NASA Systems Engineering Handbook*, SP-610S, June 1995.

[2] Allen, C. R. "Ecosystems and immune systems: hierarchical response provides resilience against invasions," Conservation Ecology 5(1): 15. 2001.

[3] Bell, A. M., Dearden, R., and Levri, J. A., "Measuring the Resilience of Advanced Life Support Systems," 2002, https://scholar.google.com/scholar?cluster=8372612629185790914&hl=en&as_sdt=0,5, retrieved 1/14/2021.

[4] Blanchard, B. S., *System Engineering Management*, John Wiley & Son, Hoboken New Jersey, 2008.

[5] Murray, R. M., Day, J. C., Ingham, M. D., Reder, L. J., and Williams, B.C., "Engineering Resilient Space Systems, Keck Institute for Space Studies, 2013.

[6] Madni, A., M., and Jackson, S., "Towards a Conceptual Framework for Resilience Engineering," IEEE SYSTEMS JOURNAL, VOL. 3, NO. 2, JUNE 2009.

[7] Escobar, C. M., Nabity, J.A., and Klaus, D. M., "Defining ECLSS Robustness for Deep Space Exploration," ICES-2017-280, 47th International Conference on Environmental Systems, Charleston, South Carolina, 2017.

[8] Wied, M., Oehmen, J., Welo, T., "Conceptualizing resilience in engineering systems: An analysis of the literature," INCOSE Systems Engineering, Volume 23, Issue 1, January 2020.

[9] Uday, P., and Marais, K., "Designing resilient systems-of-systems: A survey of metrics, methods, and challenges," Syst Eng. 2015;18(5):491–510.

[10] Taleb, N. N., *The Black Swan: The Impact of the Highly Improbable*, Random House, New York, 2010.

[11] Kay, J., and King, M., *Radical Uncertainty: Decision-Making Beyond the Numbers*, W. W. Norton, New York, 2020.