The Journal of Reliability, Maintainability, and Supportability in Systems Engineering

Summer 2019

Table of Contents Summer 2019

Editor's Note

John Blyler

The Summer 2019 edition of the RMS Journal is here! This issue starts with an article from two well-known and welcomed contributors, John M. Green and Jerrell Stracener. Their discussion presents a specific framework for the application of systems effectiveness modeling analysis capability for the acquisition process. It describes why the key decision criterion is the probability of mission success and outline an approach to the derivation of the framework. This framework is inclusive of capability, readiness, mission reliability, and survivability analysis which is typically omitted in system effectiveness evaluations.

In the next article, readers are asked to assume the role of maintenance detectives to discover unrevealed failures. These failures occur in both production and protective systems. According to the author, V. Narayan, the actions needed to mitigate these risks are the same in either case. His discussion focuses on protective systems in continuous process industries.

The article by Matthew Hogan acknowledges that reliability verification is a relatively new practice for many IC design companies and intellectual property (IP) suppliers. Still, the practice promises significant process improvements if foundry-qualified and maintained reliability rule decks are followed. Rule decks ensure that designs can be manufactured successfully and are a good first step for ensuring reliable design. But what else should be done? Read the article to find out.

The final article discusses the rationale of aircraft maintainability design principles to help mitigate errors caused in maintenance. Clive Nicholas begins by noting the common belief that errors caused by maintainers are not the concern of the designer. Rather, designers argue that aircraft maintainers should be trained not to make errors. However, recent failure events in the aircraft industry have resulted in an increased awareness by regulators, designers, manufacturers, operators and others of the impact that the design characteristics of aircraft can have on safe and effective maintenance performance. This is particularly important for the avoidance of maintenance error and the mitigation of unavoidable or undiscovered errors.

Finally, we sadly note the passing of one of our own in the reliability industry. Benjamin Blanchard, long time practioner and author of many seminal textbooks in the areas of reliability, maintainability and logics, passed away on July 11, 2019. He will be missed.

I hope you find this issue of value. Please don't hesitate to share your comments and future articles with me via the email below.

Cheers!
John

A Framework for a Defense Systems Effectiveness Modeling & Analysis Capability: Systems Effectiveness Modeling for Acquisition

John M. Green
Jerrell Stracener, Ph.D.

Abstract

The purpose of this paper is to present a response to two current Department of Defense (DOD) initiatives. The first is the DOD National Defense Strategy of 2018 which encourages the adoption of new practices to improve system performance and affordability to meet current and future threats. The second initiative is the DOD Digital Engineering Strategy which outlines five strategic goals in support of the first initiative. The first strategic goal: "Formalize the development, integration, and use of models to inform enterprise and program decision making" is the specific subject of this paper. The response is a conceptual methodology that addresses an analytic deficiency identified by a 2017 congressional commission that examined the capabilities of the DOD civilian staff in their determination of force and weapons systems requirements. Specifically, this paper presents a framework for a "Defense Systems Effectiveness Modeling and Analysis Capability" whose metric is the probability of mission success. The objective is the application of modeling and analysis to guide decisions leading to fielding systems having optimum effectiveness constrained by affordability and reduced development time. While the current U.S. focus is on Systems Readiness, it is an integral element of the more robust Systems Effectiveness. While cable raceway fires, cable bundle severing events, and other common cause cable failures (e.g., rodent damage, chemical damage, fraying and wear-related damage, etc.) are known to be a serious issue in many systems, the protection of critical cabling infrastructure

and separation of redundant cables is often not taken into account until late in the systems engineering process. Cable routing and management often happens after significant system architectural decisions have been made. If a problem is uncovered with cable routing, it can be cost-prohibitive to change the system architecture or configuration to fix the issue and a system owner may have to accept the heightened risk of common cause cable failure. Given the nature of cables where energy and signal functions are shared between major subsystems, the potential for failure propagation is significant.

Introduction

The 2018 National Defense Strategy (NDS) (NDS 2018) makes readiness and warfighter needs a priority with lethality and warfighting the primary objective. The strategy emphasizes affordability with sustained and predictable investment to achieve greater performance through modernizing the military and restoring readiness. Within this context, improvement of readiness involves developing the right systems or systems of systems with alacrity.

To support the goals of the NDS, the Department of Defense's Under Secretary of Defense for Research and Engineering has initiated the Digital Engineering Strategy (DES) that has five goals intended to drive acquisition of future systems (DES, 2018). The five goals promote a model-based, systems engineering (MBSE) wherein systems are digitally rendered. The resulting digital artifacts become the means of communications between stakeholders. The goals are:

1. Formalize the development, integration, and use of models to inform enterprise and program decision making;
2. Provide an enduring, authoritative source of truth;
3. Incorporate technological innovation to improve the engineering practice;
4. Establish a supporting infrastructure and environments to perform activities, collaborate, and communicate across stakeholders; and
5. Transform the culture and workforce to adopt and support digital engineering across the lifecycle.

Purpose

An approach to the first goal of the DES is the purpose of this paper. A crucial element of the formalization process is the development of an effectiveness modeling and analysis framework. The advent of DES is important because recent criticism by a bipartisan congressional commission noted that civilian analytical capabilities for force and weapons development within the Department of Defense (DoD) have severely degraded since their original establishment in the 1960s by Robert McNamara (WSJ, 2018). The truth of this statement is borne out by the lack of an established methodology within DoD for acquiring systems of systems. There is current work underway addressing systems of systems, mission engineering, and capability portfolio analysis but not at the level of the Weapon System Effectiveness Industry Advisory Committee (WSEIAC) study to be discussed shortly.

Specific Contribution of this Paper

The contribution of this paper is twofold. First, it provides clarity of purpose for readiness, an oft used and abused term. Why not readiness? A focus on readiness may lead to sub-optimum system solution because it ignores three other factors important to systems effectiveness and mission success. Mission success is the applicable measure because it drives force projection and war-fighting capability. Second, the paper presents a framework that addresses the role of readiness within the context of mission success. This framework applies to both systems and systems of systems acquisition, providing the stakeholders with quantified results.[1]

Organization of Paper

The paper provides a brief discussion of relevant past work that is foundational to the development of the "Defense Systems Effectiveness Modeling and Analysis Capability" (DSEMAC). Key terms are defined mathematically followed by a brief discussion of the requirements for a framework that provides the needed structure for the DSEMAC which in turn is followed by a description of the proposed framework. A summary and a description of future work conclude the paper.

Past Work

A focus on readiness ignores the larger context of systems effectiveness and the additional attributes of mission reliability, mission survivability, mission capability. It is the premise of this paper that system effectiveness and mission success are the same and the overarching goal. Readiness is a subset of the larger picture that includes mission reliability, mission survivability, and mission capability as shown in Figure 1. This view is not a new concept. The relationships have a long history that starts in the 1950s and was extensively documented in a report published by the Weapon System Effectiveness Industry Advisory Committee in the 1960s (WSEIAC, 1965). Figure 1 is based on the WSEIAC report and illustrates the relationship between overall mission effectiveness and its constituent components of mission readiness, mission reliability, and mission capability. Note that mission survivability is not included in the report and thus omitted from Figure 1. Survivability is included in this paper for completeness as shown in Figure 2.

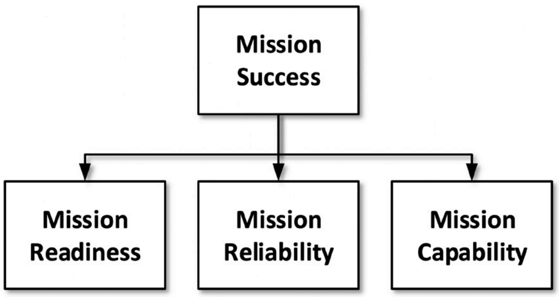

Figure 1 The WSEIAC Systems Effectiveness Hierarchy (Adapted from WSEIAC 1966)

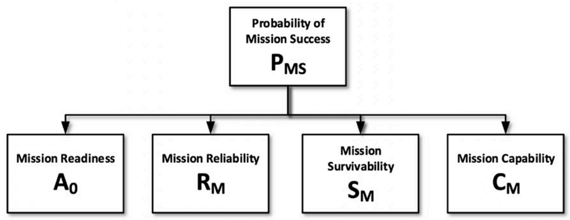

Figure 2 The Revised WSEIAC Systems Effectiveness Hierarchy

As defined by the WSEIAC report, mission readiness (often known as operation availability (A_o) or operational readiness (OR)) quantifies the percentage of time that the system is ready at the start of the mission. Mission reliability (or dependability) quantifies the likelihood that the system will perform its mission essential functions throughout the mission. Both these terms are well represented in the literature. Mission capability quantifies the adequacy of the system to meet the mission goals. Capability is about ways and means. It matters not if the system is available and reliable throughout the mission if it cannot achieve the desired results because the said ways and means were insufficient or incorrect.

Figure 2 presents a complete view of the relationships with the addition of mission survivability. The probability of mission success is a function of the four terms. Therefore, the graphic is a top-level objective hierarchy. As an objective tree, the goal is to maximize

[1] System will be used throughout this paper.

the probability of mission success. The lower-level objectives each describe a specific aspect of mission success and are, therefore, inherently important. The lower-level objectives can be expanded by including another level of detail. For example, mission survivability can be expanded to susceptibility and vulnerability. In this case, the goal is to reduce both to increase survivability.

The systems effectiveness hierarchy and the following equation for P_{MS} provides a quantitative basis for the acquisition of weapons systems and systems of systems. The WSEIAC report provided a general mathematical relationship for mission success as follows:

$$P_{MS}=(P_{A_O})(P_{RM})(P_{SM})(P_{CM})$$

where;

P_{MS} = the probability of mission success for a specified mission.

P_{A_O} = the probability that the system is available at the start of the mission.

P_{RM} = the probability that the system will successfully perform specified mission essential functions by mission phase.

P_{SM} = the probability that the system will survive the mission.

P_{CM} = the probability that the system meets the capability objectives.

Note the probabilistic formulation of mission success. There are several valid reasons for this approach. First, military operations are characterized by random variables for example, probability of detection or probability of kill. Second, probabilities are dimensionless making them easier to work with across diverse system elements such as sensors and weapons.

In systems terminology, Figure 3 is a context diagram that becomes a starting point for the framework requirements discussed in the following section.

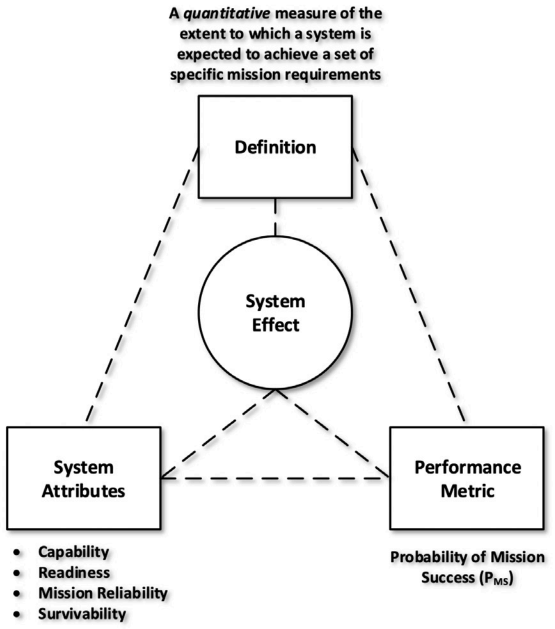
A *quantitative* measure of the extent to which a system is expected to achieve a set of specific mission requirements
Definition
System Effect
System Attributes
Performance Metric
• Capability
• Readiness
• Mission Reliability
• Survivability
Probability of Mission Success (P_{MS})

Figure 3 Systems Effectiveness Relationships

Framework Requirements

A framework is a structured way of relating objects of interest and their resulting interactions. The importance of a framework in the acquisition of systems cannot be understated. First, a framework organizes theory and practice and provides a structure for methods. Second, complex systems and systems of systems are typically not developed as a single architecture. Thus, there are time-phasing and contractual issues. Individual systems are usually single function, and system couplings are interdependent (Luman, 2000).

Third, there currently is no systematic method of measuring systems effectiveness. The literature is devoid of theory and standards. Most approaches center on qualitative methods which are subjective at best.

Basic Requirements

There are four major requirements for the framework: the supporting methods must be quantitative, the supporting methods must present results probabilistically, and the supporting methods must be reliability based. Finally, the framework must support hierarchy and abstraction. The end goal is a framework that supports evaluation of mission success versus cost where the emphasis is on the likelihood of mission success.

Quantitative

One of the first steps in an analysis is to describe the processes involved. Mathematics is precise and explanatory, facilitating analysis and explanation of more complex problems than possible using qualitative methods. The model for the probability of mission success must be based on proven methodology. The challenge is developing and maintaining a model for each mission which will be large and complex for complex systems.

Probabilistic

Military operations are about achieving success and the estimation of event probabilities, typically described as measures of effectiveness (MOE) or measures of performance (MOP). Often parametric values are used incorrectly as measures. For example, detection of a threat is expressed as a probability of detection and is a function of several parameters including range. The outcome is the probability of detection as a function of range.

Reliability-Based

Reliability theory is based on the premise of system success and failure ($P_{success} = 1 - P_{failure}$). Many of its concepts are foundational precepts to quantifying system effectiveness. Further, most of the system variables of interest are reliability related. Figure 3 identifies them as key system attributes.

Hierarchy and Abstraction

Systems are hierarchical by nature with increasing detail at each level of expansion. The framework must support models that describe each level of expansion. This paper suggests a black box approach at each layer.

A Notional Effectiveness Model

Systems concepts are based on a need to meet an operational requirement. The effectiveness of how well this need is met (mission success) is a measure of its tactical utility and its value to the force structure. Figure 4 is a notional model adapted from Figure 2-1 found in the *Reliability Engineering Handbook* (NAVWEPS, 1964). It summarizes the first three figures and is intended to convey several points. How well the system will perform, how long the system will perform, and how often the system can perform.

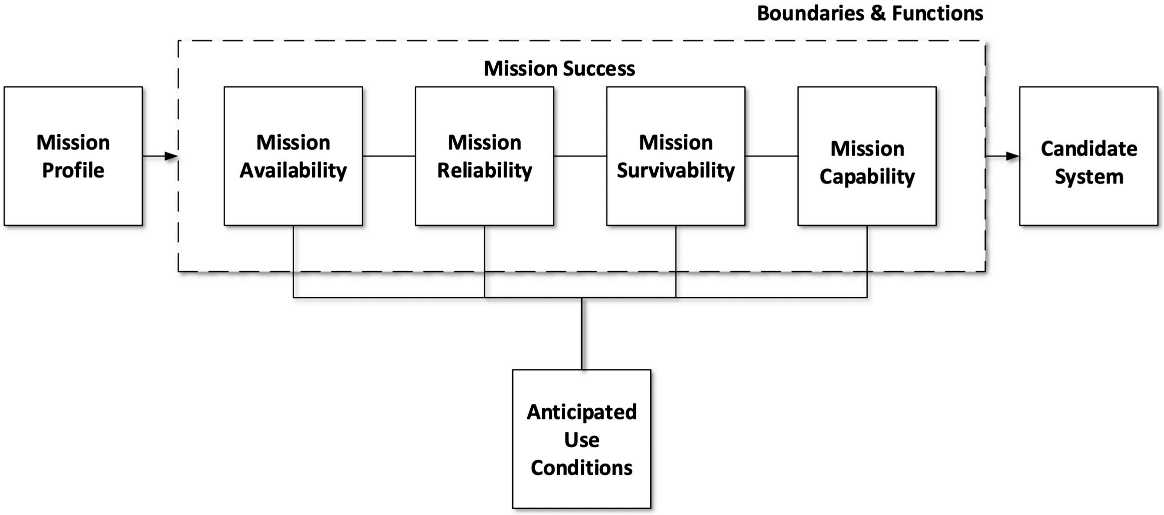

Figure 4 A Systems Effectiveness Model (Adapted from NAVWEPS 1964)

This model, when combined with a decision process, becomes the basis for the overall framework model.

Proposed Framework

Figure 5 is a generic decision process. It serves as a guide to understanding how to incorporate Figure 4 into a larger context. Figure 6 is the resulting proposed framework.

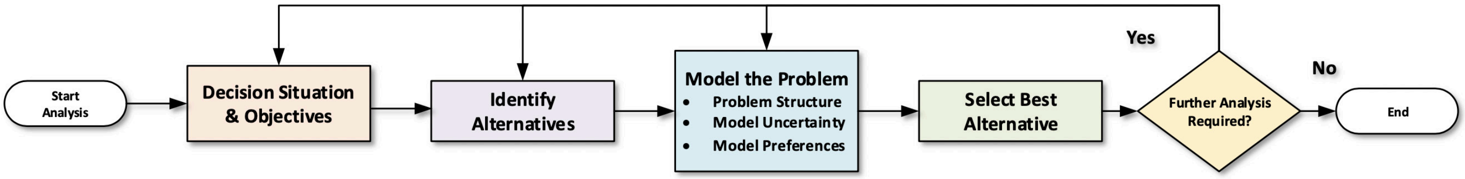

Figure 5 Generic Decision Process

Problem Formulation

With the framework in place, it is appropriate to return to the purpose of the framework to wit: to make decisions about system selection. There are three basic steps to the decision process. First, understand the set of system variables and how they interact quantitatively and accurately. Knowledge of the system is imperative. In the framework, this is represented by the upper five boxes (orange and purple). Second, select a single MOE expressible in terms of the variables represented by the blue box. A premise of this paper is that mission success is that MOE. The final step is to select the method by which the best system is selected represented by the green box.

The decisions involve making choices from a set of candidate solutions in order to find the most desirable solution. Once the decision is made it becomes an irrevocable

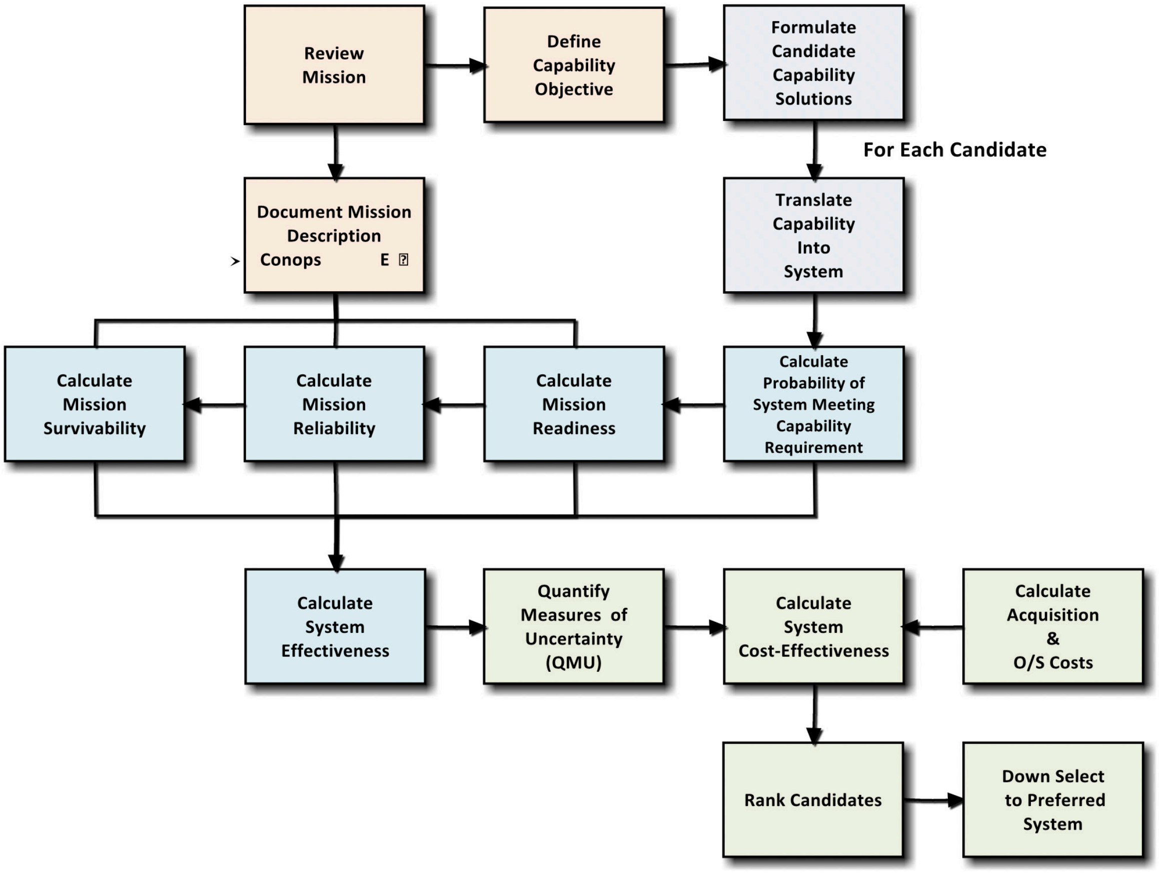
Review Mission
Define Capability Objective
Formulate Candidate Capability Solutions
For Each Candidate
Document Mission Description Conops E
Translate Capability Into System
Calculate Mission Survivability
Calculate Mission Reliability
Calculate Mission Readiness
Calculate Probability of System Meeting Capability Requirement
Calculate System Effectiveness
Quantify Measures of Uncertainty (QMU)
Calculate System Cost-Effectiveness
Calculate Acquisition & O/S Costs
Rank Candidates
Down Select to Preferred System

Figure 6 A Framework for a Defense Systems Effectiveness Modeling and Analysis Capability

allocation of resources. Given the set of candidate solutions, the task becomes one of defining a system such that:

$$P_{MS}=(P_{A_O})(P_{RM})(P_{SM})(P_{CM})$$

Subject to the following constraints

- Specified Mission
- Required Performance
- Budget

This is a basic optimization problem. It is decisive because the result is one system—the best one.

Comments on Cost-Effectiveness

In the model described in Figure 6, cost-effectiveness has been chosen as the criterion for the model because it is best used for ranking alternatives that are relatively similar especially when there is a single dominant objective whose attainment can be assessed directly or for which a good proxy value exists (Quade, 1982). It is axiomatic in the world of quantitative analysis that in general, the possibility of selecting between two alternatives based on cost and effectiveness data alone is not possible. It is a choice between specifying performance or cost. If the former then cost is minimized, if the latter then effectiveness is maximized.

Summary

This paper presents the rationale for a framework for a "Defense Systems Effectiveness Modeling and Analysis Capability." It describes why the key decision criterion is the probability of mission success and shows the approach to the derivation of the Framework. This Framework is inclusive of Capability, Readiness, Mission Reliability, and Survivability which is typically omitted in system effectiveness evaluations.

Future Research

As noted, Survivability is not usually included. While availability and readiness have a large literature base, there is very little material on survivability.

A second research topic is Candidate Capability Architecture solution development. There is no literature on performance-based architecture development.

References

1. DES. (2018). DoD Digital Engineering Strategy. U.S. DoD OSD USDRE.
2. Luman, R. R. (2000). Integrating Cost and Performance Models to Determine Requirements Allocation for Complex Systems. Retrieved August 15, 2018, from Johns Hopkins APL Technical Digest: http://www.jhuapl.edu/techdigest/TD/td2103/luman.pdf
3. NAVWEPS. (1964, June 1). DTIC. Retrieved December 12, 2018, from Defense Technical Information Center: https://apps.dtic.mil/dtic/tr/fulltext/u2/a286606.pdf
4. NDS. (2018). National Defense Strategy. DoD.
5. Quade, E. (1982). Analysis for Public Decisions, 2nd ed. New York: Elsevier Science Publishing Co., Inc.
6. WSEIAC, W. S. (1965). Final Report of Task Group 1, AFSC-TR-65-2.
7. WSJ, M. G. (2018, November 14). Study Cites Weak Civilian Control of Military. Wall Street Journal.

Detective Maintenance

V. Narayan

Abstract

Physical assets used in industries, are susceptible to failures, due to natural degradation. The failures are of two types. Of these, revealed failures will be known to operators, but unrevealed failures remain so, till a second event also occurs. In high hazard continuous process industries, these can cause serious harm. This paper addresses the management of unrevealed failures. These failures occur in production and protective systems. Actions to mitigate them, in both production and protective systems are identical. In this paper, the focus is on protective systems in continuous process industries.

Introduction

Industries that provide goods or services use physical assets extensively. With use, and over time, these assets degrade, and cease to perform their intended functions. Failures have consequences, some of which will be severe. If they are very low, no action is required. If not, mitigating actions are required. Faulty protective systems or devices are a special case; there is no consequence till there is a demand on them. They can remain faulty for a long time, unknown to the operator. All the while, the fault threatens the safety of the protected equipment or system. Unrevealed faults pose a challenge, and must be detected in time, before an unexpected demand results in a serious accident. This paper explains the relevant issues. The options available to mitigate risks effectively and economically are examined.

Terminology & Definitions

An *element* is the smallest replaceable part of an assembly, that is normally replaced if it fails. It is also called a *component*.

A *device* is an assembly of elements that perform a desired function.

An assembly of devices that performs a function, is called an *equipment*.

The term *item*, is used to describe an element, device, or equipment.

A *system* has many items and performs a specific function.

A *plant* produces goods or services that customers want and consider valuable. It has many systems.

Failure is the inability to do what is expected of an item, in a stated operating environment.

The term *duty* describes how an item is used. Are there sudden or frequent starts and stops? Is an item loaded and unloaded gradually? Is the load intermittent or continuous? Is the item loaded at, or close to its maximum rated capacity?

Stated environment describes the operating conditions, both internal and external. These may be benign or aggressive, and includes the conditions of exposure or duty.

Revealed failures are those where the consequences occur at the same time as the failure. They are also called *evident failures*.

Unrevealed failures are those where a second event or fault must also occur before there is a consequence. They are also called *hidden failures*.

Production and Protection Systems

The function of most of the assets in a process plant is to produce goods or services for sale. Some assets provide an infrastructure for this purpose. These include buildings roads, computers, and instruments not directly connected to production systems.

There are other systems whose function is to protect the production system. If process parameters get out of control, they bring the production system to a safe state. This may be achieved by starting additional equipment, shutting down or depressurizing affected sections, or by other actions that limit damage to assets, people and the environment. They are rarely called upon to work, as the normal process control systems manage routine deviations. Due to infrequent demands, they are idle for most of the time. In this time, a fault can develop, due to gumming, sticking or calibration drifts, caused by vibrations, heat

Primary Fault	Additional Event/Failure
Flammable-gas detector fails to detect	Ignition source also present
Standby pump does not start	Duty pump fails in service
Relief Valve does not lift at set pressure	Overpressure of protected item
Fault in Emergency Shutdown valve	Unsafe process condition in plant
Rotor axial displacement sensor faulty	Rotor thrust bearing hot+ vibrations
Turbine over-speed trip device faulty	Load drops off suddenly
Flare line drain valve plugged	Liquid in flare line needs draining
Fire water deluge nozzles plugged	Deluge valve opens
Faulty reverse current relay	Reverse current occurs
Tank level trip instrument faulty	Tank level control fails in service
Fighter jet ejector seat faulty	Plane out of control; pilot must eject

Table 1 Examples of Unrevealed Failures

or other environmental conditions. These faults can remain unknown to the operator for months or years. If there is a demand on the system when it has a fault, it will not act in time, or respond at all. A serious accident may then occur.

The examples in Table 1, illustrate this point. Note that in each case the primary fault by itself does not cause harm. It is only when there is a demand on it, due to an additional event or fault, when damage can take place. Normally, the operator cannot know that the primary fault exists. These events occur randomly. Eliminating the pre-existing primary fault is the only option.

Process plants have a number of protective systems. In the event the production system becomes unsafe, they bring it to a safe condition. This can be due to a failure of the normal process control systems, a leak of process fluids, an operator error or an external situation. They are normally dormant, as protective systems are only required to work in abnormal situations. Most of the protective systems are electrical, electronic or programmable logic designs, called Safety Instrumented Systems (SIS). The International Electrotechnical Commission (IEC) standard 61508, is the basis of SIS designs. IEC 61511, based on IEC 61508, is applicable to the continuous process industry. The US ANSI/ISA 84.00.01-2004 mirrors IEC 61511, but has some additional clauses. Other industries, such as railways, nuclear plants, road transport and power drive systems also use standards based on IEC 61508.

The SIS will have one or more sensors. Their signals go to a logic solver unit, that has voting (such as, two out of three or 2oo3), or a decision logic, such as, IF-THEN. Its output signal goes to one or more actuators, such as a solenoid. That is connected to an executive device, such as, a shutdown valve or a switch to isolate the energy source. The action of the executive device makes the situation safe. The protective system succeeds when all the items, cable terminations and mechanical links work promptly and correctly. The scheme is illustrated in Figure 1.

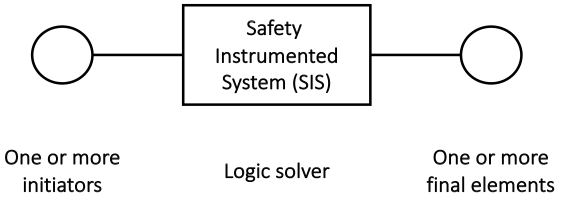

Figure 1 Schematic showing safeguarding elements and links

There are also some mechanical protective systems. These include bursting disks, relief valves, and overspeed trip devices.

Detective Maintenance Strategy

The late John Moubray, author and renowned teacher, used the term detective maintenance, to describe the process to identify unrevealed failures. This is done by simulating a real demand, using a test, under controlled conditions. That verifies how the system would respond to a real demand. The test reveals faults that would otherwise remain unknown. When these tests are done periodically, the time between a fault occurring and its discovery, can be reduced by shortening the test intervals.

Idle systems develop faults for several reasons. Vibration, temperature extremes, fouling, or erosion of sensor elements, gumming of valve stems, wear of mechanical links, and other environmental causes, result in faults.

Revealed or evident failures are known to the operator, so they are easier to manage. The onset of failure would be known, by a change in parameters, such as, vibration, temperature, differential pressure or heat transfer rate. Progress of the fault can be monitored by trending that parameter. That makes prediction of time to functional failure fairly simple. This strategy is called predictive maintenance (PdM), or condition based maintenance (CBM). It is different from detective maintenance, and is only applicable to evident failures with parameters that allow trending.

Levels of Testing

The protective system is designed to take action when there is an unsafe situation. There are two failure modes to manage. The first is when it does not become active in time, leading to an escalation of trigger events. That exposes the plant to danger, and it is termed a fail-to-danger mode. The second case is when it becomes active when the situation is perfectly safe. This is a false alarm, that can lead to a spurious trip. This causes an avoidable loss of production, which is also undesirable. However, safety of the plant is not compromised, so this is called a safe failure mode. Testing is aimed at exposing fail-to-danger modes.

In a real demand scenario, the sequence of events will be as follows.

- Sensor(s) note unsafe condition; signal sent to logic solver
- Decision logic generates a signal; sent to actuator(s)
- Actuator(s) move the executive element(s)
- Executive device takes actions to mitigate the situation

Faults can occur within each item and at interfaces. Any of these can prevent the corrective action. The last step will cause a production loss, which is not acceptable in a test situation. Sensors, logic solvers, actuators and executive devices are, in practice, checked separately, after sectional isolation. If all the items work, it is inferred that the whole system works. That is arguable, but is seen as a pragmatic compromise.

Opportunity-Based Tests

Before accepting that compromise, there is a possibility to use an existing opportunity. If a production system or the whole plant trips, it is useful to download the relevant operating data at the time of the incident. That provides data normally available from a real test. Credit can then be taken for this “test.” Procedures must be in place in advance, to enable this work. If a system or unit has to be stopped for minor planned work, the stoppage can be initiated by a test. In this case, only the initiating signal source is not verified, while the rest of the loop is checked.

When a unit or plant is due for a major planned shutdown (turnaround), a number of tests can be done during the process of shutting it down.

On-Stream Tests

In some plant designs, spare Pressure Relief Valves (PRV) are installed, with isolation valves that are interlocked. One PRV is in use, while the other is blocked off. That enables testing of PRVs while the plant is on-stream.

On-stream relief valve testing methods are also available from commercial vendors. The process simulates the operation of the PRVs under normal operating conditions. An external mechanical lifting force that supplements the valve opening force, is applied gradually. This enables the PRV test to be done without disturbing normal operations. The vendor’s software provides data on the lifting and reseating pressures, and the mass of fluid discharged. One of the risks posed by this method is that sometimes the PRV may not reseat after the test. The other issue is that internal fouling, or other damage will still remain a threat, till an overhaul is done.

Opportunity-based and on-stream tests

can only be applied in a few cases. The rest still need functional tests.

Functional Tests

This is done at the element level or at the system level. Sensors, such as gas detectors, may be tested individually, at site. Failures will be recorded in the test report, and defective items replaced. Failure rates can be computed as follows:

Failed Items ÷ (sample size × time in service),
and
MTBF = 1 ÷ failure rate

Functional tests can be done at the system (or loop) level, from sensors to executive devices. If the test is successful, there will be a significant loss of production, or of environmental damage. For example, with a successful test, an emergency shutdown valve will close, or a depressurization valve will open. The loss of production can be substantial. That poses a dilemma, as these consequences are unacceptable. To prevent damage or loss due to the test, the movement of the executive device is restricted or stopped completely. This can be done at various levels, as discussed below.

Executive Element Movement Stopped

In some cases, there is a mechanical link between the executive device and actuator. That link can be removed before the test, to stop the movement. In other cases, a physical obstruction prevents the executive element from moving.

Hydraulic pressure is sometimes used to move isolation valves of machinery, such as, compressors or turbines. When the actuator moves the pilot valve to energise the hydraulic cylinder of the main valve, the oil is diverted to a sump. The flow of hydraulic oil to the sump is seen as evidence that the main valve would have moved fully, had the diversion not been made.

In all these cases, the inference that the executive device would have worked satisfactorily, is not always true.

Partial Movement Tests

The purpose of partial movement tests is to limit the movement or stroke the executive device. This is done by physically stopping the stem from moving more than a few mm. If it is a rotary device, like a ball or butterfly valve, the movement is restricted to two or three degrees of arc. Special devices that allow controlled partial movement of large valves that are hydraulically or electrically operated, are available and can be retrofitted.

Such tests confirm that the executive device moves on demand. They also perform another useful function; the small movement breaks off any gumming, allowing free movement thereafter. The time to complete the full stroke is an important requirement, but that cannot be checked by partial movement tests.

Functional tests will not be able to discover internal faults, caused by fouling, fatigue, corrosion or erosion. For that, an internal cleaning and inspection is required.

Test Coverage Factor

When an item is overhauled, internal cleaning and inspection will enable the elimination of any faults identified. After re-assembly and testing, the item is as good as new. In the absence of overhauls, faults such as internal fouling, misalignment of parts, corrosion or fatigue damage will not be known by testing alone. The test coverage factor, is a measure of the level of confidence that can be assigned to the test. It is defined as:

Detectable Faults ÷ Faults Present

The coverage factor can approach 100%, but is always lower.

Computing Test Intervals

Complete loop tests are not practical in most situations for the reasons discussed earlier. Testing the elements separately is relatively easy, as it does not need a system shutdown. Thus, sensing units, such as, gas detectors, will be tested as a set. The formula in this section apply to sets of sensors, logic solvers, actuators and executive devices. From the results the failure rates can be calculated, using the formula given earlier in the section on functional tests.

High criticality of protected systems pose greater risks, and need a better level of protection. That means that the related protective system availability must be high.

When the protective system reliability (measured as Mean Time Between Failures or MTBF) is high, the protective system availability will be high. Similarly, as the test intervals are reduced, the system availability goes up. These relationships can be stated mathematically as:

$$(100\% - \%\text{availability}) =$$
$$(\text{Test interval}) \div (2 \times \text{MTBF})$$

The following information is required for this calculation.

What is the required availability? The range is usually 95% to 99%

What is the MTBF of the relevant protective system? This data comes from the result of the tests on sets of elements. Over a given operating period, of say, 12 months, the element's MTBF in months, is:

$$\text{MTBF} =$$
$$\text{Installed Elements} \times 12 \text{ months}$$
$$\div \text{ Failed Elements}$$

In the absence of test data, default MTBF, values from the vendors can be used (with caution). There are tables and charts in IEC 615111, from which test intervals can be read off for each grade of element and the configuration of the protective loop.

Summary

Protective systems and devices are necessary to operate process plants safely. They are susceptible to faults that will not be obvious to the operator, and remain in a failed state, until they are detected. In their failed state, if a second event occurs, the defective protective device will not work. This will result in high consequences. A Detective Maintenance strategy is the best way to find such faults economically. This paper describes practical solutions, along with the associated risks.

References

1. IEC 61511 Functional safety - Safety instrumented systems for the process industry sector, 2016, at https://global.ihs.com/doc_detail.cfm?rid=Z56&mid=IEC&document_name=IEC%2061511%20SET&item_s_key=00473613&utm_source=google&utm_medium=cpc&utm_campaign=iec&utm_content=IEC_61511_SET&gclid=EAIaIQobChMI3p3N0aum-3wIVQ7HtCh0WRgG-EAAYAiAAEgKoAvD_BwE

Leveraging Foundry Baseline Checks for Robust Reliability Verification

Matthew Hogan

Introduction

At every node, and for every integrated circuit (IC) design, there is one crucial element for success: the quality of the verification rule deck. The requirements defined in rule decks ensure that designs can be manufactured successfully, and that they will perform as intended. Today's IC designers and verification engineers alike rely on the foundry to provide robust design rule checking (DRC), layout versus schematic (LVS), and parasitic extraction (PEX) rule decks that support automated design layout and verification. The complexity of today's designs, combined with the cost of failure (in both delivery time and production expense) provide ample justification for using these foundry rule decks—they have been validated and qualified by the foundry, and define known good solutions for sign-off verification. But what about other verification needs?

In the last few years, the increased demand for product reliability (both in performance and product lifetime) has created a broad need for context-aware reliability verification [1]. Foundries have responded by creating reliability rule decks that deliver a wide range of reliability solutions, from electrostatic discharge (ESD), latch-up (LUP), and interconnect reliability, to power management, electrical overstress (EOS), and other potential reliability impacts [2].

These foundry-qualified reliability rule decks provide an excellent starting point for IC reliability verification flows, and should be the baseline for IC reliability verification in any design company. Electronic design automation (EDA) companies provide verification tools that use these rule decks to automate and standardize IC reliability verification flows, just like they do for DRC, LVS, and PEX. Need more justification? Consider this—customer demand for enhanced product reliability, along with the relentless progress to smaller nodes, is driving the foundries to continuously develop and qualify new and augmented reliability rules, saving design companies countless hours and resources that would otherwise be spent creating custom rule decks. Creating your own rule decks to achieve this desired reliability baseline would require tremendous cost in time and resources, with no guarantee the results would match the foundry's requirements.

Foundry Reliability Rules

Because foundry offerings typically have a specific reliability focus, design companies should understand what each foundry's reliability rule deck offers. While ESD protection is the common thread among all reliability rule decks, they diverge in other areas.

All major foundries now provide reliability rule decks to their customers, each focusing on those reliability issues it considers most critical to its customer base [3-8]. Some address ESD, interconnect reliability, and LUP, at both the IP and full-chip level, with components such as complete ESD/LUP rule check coverage using topology, point-to-point (P2P) resistance, current density (CD), and layout-based LUP rules.

Others address industry-specific needs, such as the analog constraint checks developed for the RESCAR-project for reliability checking, in the form of automotive reliability check templates. These checks enable designers to ensure the enhanced level of reliability compliance that automotive industry standards, such as the international functional safety standard ISO 26262, now require throughout the entire automotive supply chain [9]. What's more, even though these types of reliability checks are targeted towards the analog segment of designs, they can be used to analyze and enhance the reliability of any IC design. Checks for power management, ESD, and charge device model (CDM) protection, as well as analog design constraint checks that address sensitive analog layout requirements, such as device alignment, symmetry, orientation/parameter matching, and more can be useful in a wide range of IC designs.

Some foundries have also implemented IP quality programs, designed to help their customers improve intellectual property (IP) dependability, by using their reliability rules to establish a consistent reliability baseline for both design companies and IP providers across the ecosystem [10].

Establishing a Reliability Verification Baseline

Typically, internal design and CAD support groups leverage the foundry rule deck as their baseline when they need to develop customized rule decks to address their company's specific needs. For DRC, LVS, and PEX rules, this often means starting from the foundry-provided rule decks and, as needed, applying relatively small additions or modifications to ensure the design company's unique or proprietary verification requirements are satisfied. Not only is this more efficient in terms of time and resources, but it also ensures consistency across all designs and nodes, since the foundry rule deck is always the starting point.

The same methodology should be applied for reliability verification. Companies can start by incorporating the foundry-supplied reliability rule deck into their verification flows precisely because those requirements have been thoroughly vetted by the foundry. From IP to full-chip reliability applications, the value of establishing a baseline for reliability acceptance throughout the entire design flow has been established [11]. Whether your company is implementing formal reliability verification for the first time, or you already have a customized in-house reliability checking process, foundry reliability rule decks provide the same benefits as DRC, LVS, and PEX decks—uniform, qualified requirements and consistent foundry maintenance across all projects and process nodes.

What's the best way to implement a foundry reliability rule deck? When design teams evaluate the applicability of the foundry reliability solution, they first must understand what is being verified by the foundry rule deck. This is essential, more so if your company uses different foundries for different projects. Understanding which distinct areas of reliability concerns the reliability rule deck from a specific foundry addresses now becomes a part of that decision and informs areas for supplemental coverage.

Foundry selection can also be affected by mergers and acquisitions. Project teams in a unified company may continue to develop new versions and incremental updates of chips using the same foundries they employed for the original products. Faced with the time and expense of transferring to a new foundry, they adopt the adage of "If it isn't broken, don't fix it."

Whatever the reason, if your company is sending different projects to different foundries, it is essential to make sure you have a clear understanding of the reliability checks each foundry provides, and confirm they align with any internal requirements for the designs to which they will be applied.

Once you know which foundry rule deck(s) you will be working with, you must determine when and how to integrate the foundry rules into your design and verification flows. Whether it's part of transitioning to a new process node, or integrating the rules into an established node, many companies begin by supplementing any internal methodologies and rule checks already in place. Using this incremental approach to gain trust in the process and the results, then transitioning more of the design flow to foundry-provided reliability rule decks, is a typical evolutionary path.

In any case, the first step is to download the foundry's reliability rule deck for your current design process node, and review the contents with your reliability/ESD team. They need to determine how well-aligned the foundry rules are with your own internal requirements, flows, and design practices. For example, the foundry-provided ESD/LUP rules are an excellent starting point for developing a reliability baseline, but depending on what your foundry provides, additions to your full-chip checklist might need to include:

- Validation that all IPs are correctly implemented
- Context/voltage-aware LUP protection verification [12]
- Interconnect robustness analysis
- Stacked devices analysis in the context of the whole chip
- Verification that the correct power ties are used in wells

They'll also need to evaluate the capabilities of any EDA tools you are using. Reliability checking often requires "context awareness," which is the ability to consider both geometrical and electrical information together to determine the correct implementation. If your verification tools don't provide automated context-aware analysis, designers may find themselves spending a lot of time trying to implement these checks with manual annotation and custom code.

Adopting tools that support automated context-aware checking can ensure faster, more accurate reliability checking and debugging.

New process nodes present both new design opportunities and a learning experience for designers, as they must become familiar with the nuances of the flow, and understand the potential reliability concerns of new devices and interconnect. New design starts, particularly those on new process nodes, benefit greatly from leveraging foundry-provided rule decks, especially since the acquired knowledge and experience from previous nodes and foundries may no longer apply.

Best Practices for Reliability Verification

While creating consistency in reliability verification across your entire design flow takes time, foundry reliability rule decks provide a proven baseline from which to begin. Once you have implemented reliability verification, there are a number of ways to optimize its use.

IP Validation, Early and Often

Getting to full-chip sign-off quickly and efficiently can be greatly enhanced by verifying standalone IP blocks and larger blocks during chip assembly. Foundry rule decks typically provide internal switches that allow designers to run either full-chip checks or IP-based and block-specific checks. IP-based checks allow the verification process to begin while design teams are still implementing and assembling IPs from internal groups and/or 3rd party IP vendors. Just like DRC, ensuring the IP being delivered to a design meets the baseline criteria in foundry-provided reliability rule decks should be a given. As with DRC, validating reliability at each level as you select and build up the design provides a deterministic path to success, allowing you to consider these design elements in the context of the whole chip.

A substantial percentage of today's designs consist of IP re-use, whether it is IP developed internally for previous projects or sourced externally. However, while the physical layout of an IP block used in a previous design may remain unchanged, the reliability context of how that IP block is used in a new design must be validated. Figure 1 shows proven and trusted IP placed in multiple power domains in a new design, with unified power format (UPF) power state tables (PST) controlling their activity. While each IP may work well in a standalone context, rigorous validation of how they all interact with (and are physically connected within) the new IC design as a whole must be completed, particularly when validating interactions between multiple power domains.

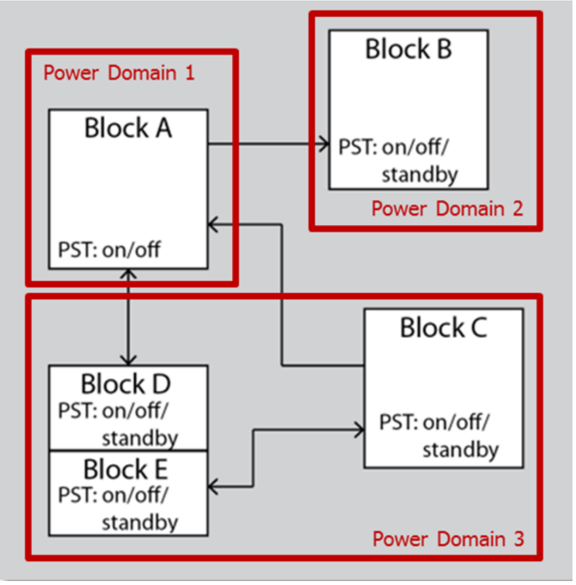

Figure 1 When placed in a new design with multiple power domains, even well-trusted IP must be validated for correct connectivity and interaction.

IP provided by multiple sources are created using a variety of design styles and techniques. Identifying IP design differences early in the design process helps eliminate last-minute issues during IP integration and assembly. Using foundry reliability rules to enforce design consistency and best practices can also simplify long-term maintenance and reduce cost of ownership. For example, one reliability-related design decision is the choice of which common ESD techniques to use for I/O pin protection. Are all of your I/O IP blocks developed for distributed ESD protection (often common in ball grid array designs), or not? Using reliability verification to ensure design consistency across multiple IP can save your company both time and money.

Validation of existing IP becomes even more challenging when it is part of a process node or foundry change. Retargeting IP can be especially difficult during a process shrink, because close attention must be paid to those parts of the design that should not shrink, such as interconnect robustness and device sizing for ESD protection. This is where foundry reliability rule decks are particularly helpful. While shrinking interconnect, transistor dimensions, and spacings across most of the design may be suitable for the new node, maintaining correct geometrical dimensions where energy must be shunted (as is the case for ESD protection circuitry) is essential, and requires careful validation.

While new nodes may offer new opportunities to improve device performance, they may also present new design reliability considerations. For example, when transitioning from planar bulk transistors to finFET or FD-SOI, designers must educate themselves on the differences in reliability characteristics between the old and new devices and processes.

Full-Chip Integration

Verification of individual IP blocks provides the foundation for verification of your chip assembly, but standalone IP verification does not address the overall context of how the IP blocks are incorporated into the larger design. Comprehensive reliability verification at the full-chip level is an essential requirement, with some reliability checks that must be performed at both the IP level and in a full-chip context. As shown in Figure 2, overall chip context is crucial when validating critical reliability applications such as ESD and EOS protection, voltage-aware DRC (VA-DRC), and interconnect robustness (particularly critical for avoiding CDM issues by ensuring low resistance between ESD clamps).

Foundry reliability rule decks used for both IP and full-chip runs often have settings or modes that allow the engineer to define the desired verification level. In the case of device-level EOS, long-term reliability issues will arise if the bulk is tied to a higher voltage than the voltage at which the gate switches. This scenario creates gate-oxide stress that will cause failure over time. These types of failures are challenging to recognize, as they are subtle design errors not easily identified by traditional SPICE simulations. To ensure that time-dependent

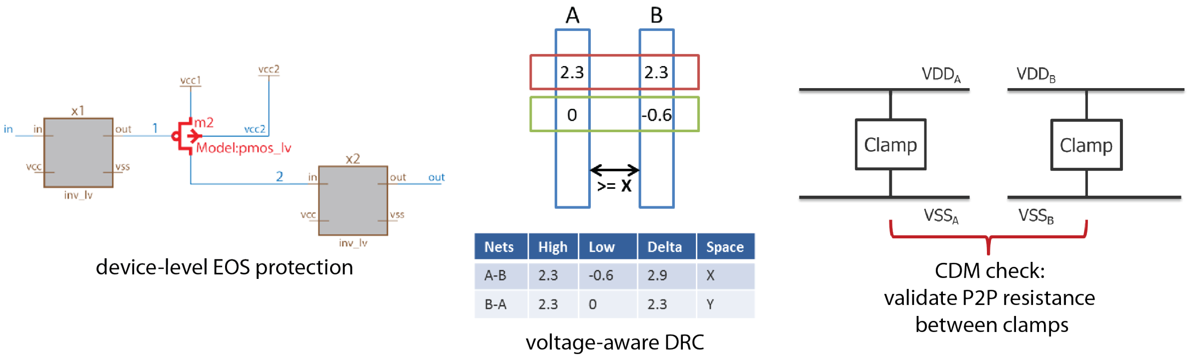

device-level EOS protection

Nets	High	Low	Delta	Space
A-B	2.3	-0.6	2.9	X
B-A	2.3	0	2.3	Y

voltage-aware DRC

CDM check:
validate P2P resistance
between clamps

Figure 2 Reliability verification of critical reliability design issues at the full-chip level.

dielectric breakdown (TDDB) does not lead to premature oxide breakdown of interconnect, VA-DRC spacing checks must be performed in a manner that considers the voltages on these interconnects [13].

While most designers expect basic ESD checking from their automated tool flows, more complex reliability checks, like interconnect robustness verification with P2P and CD analysis, are becoming critical to product success. For example, due to the shrinking of gate oxide thickness and use of multiple power domains at advanced nodes, CDM checking is required to protect gates that are directly connected to power/ground. When active clamps are used, engineers must also validate resistance between global power nets (of different domains) to avoid CDM issues.

Conclusion

For many IC design companies and IP suppliers, reliability verification is a new practice, one that comes with intense visibility and many different demands. Adoption of new process nodes provides design companies an excellent opportunity to reevaluate their entire ecosystem, from IP provider to final chip assembly. Implementing foundry-qualified and foundry-maintained reliability rule decks enables both design and IP companies to incorporate proven baseline robustness and reliability criteria into their verification flow, while eliminating the time and resources needed to create and support proprietary verification solutions. However, a thorough understanding of the coverage provided by a foundry's reliability rule deck is essential to ensure that your company's internal criteria is covered by that rule deck, especially when multiple projects source different foundries. As with other foundry rule decks, companies may need to work with their chosen foundry to expand rule deck coverage as new reliability needs arise.

Reliability verification tools provide a wide range of automated checking capabilities that ensure consistent and accurate reliability checking using the baseline foundry rule deck. These tools help verification engineers quickly and efficiently find and resolve even the most complex reliability issues from the block/IP level through full-chip sign-off verification.

Ensuring consistent, complete, and accurate reliability verification solutions is a critical step for ensuring long-term device performance and product lifetime in today's demanding and expanding markets. Beginning with a foundry-qualified reliability rule deck provides a solid baseline, and offers a proven path for future growth.

References

1. Semiconductor Engineering. "Reliability Verification." Last modified June 4, 2016. https://semiengineering.com/knowledge_centers/eda-design/verification/reliability-verification/
2. Matthew Hogan, "Jumpstart your reliability verification with foundry-supported rule decks", Mentor, a Siemens Business. August, 2018. http://go.mentor.com/521lt
3. Mentor Graphics Corporation. Mentor Graphics Calibre PERC Reliability Checking Solution Used for IP Quality Program by TSMC, Oct. 29, 2013. https://www.mentor.com/company/news/mentor-tsmc-calibre-perc
4. Mentor, a Siemens Business. "Mentor extends solutions to support TSMC 7nm FinFET Plus and 12nm FinFET process technologies." Mentor press release, Sept. 13, 2017. https://www.mentor.com/company/news/siemens-mentor-tsmc-oip-7nm-12nm
5. Mentor Graphics Corporation. "Mentor Graphics Design and Verification Tools Certified for TSMC 16nm FinFET Production." Mentor press release, April 15, 2014. https://www.mentor.com/company/news/mentor-tsmc-16nm-finfet-production
6. Mentor Graphics Corporation. "Mentor Graphics Announces Collaboration with GLOBALFOUNDRIES on Reference Flow and Process Design Kit for 22FDX Platform." Mentor press release, Nov. 9, 2015. https://www.mentor.com/company/news/mentor-collaboration-globalfoundries-22fdx-platform
7. Mentor, a Siemens Business. "Mentor Announces Availability of Tools and Flows for Samsung 8LPP and 7LPP Process Technologies." Mentor press release, May 24, 2017. https://www.mentor.com/company/news/mentor-availability-of-tools-flows-samsung-8lpp-7lpp-process-tech
8. Mentor Graphics Corporation. "Mentor Graphics Announces Availability of Qualified Calibre PERC Rule Decks for UMC 28nm Technology." Mentor press release, Oct. 19, 2016. https://www.mentor.com/company/news/mentor-availability-qualified-calibre-perc-rule-decks-umc-28nm-tech
9. Mentor, a Siemens Business. "Mentor and TowerJazz provide first commercial comprehensive suite of analog constraint checks for enhanced automotive reliability offering." Mentor press release, Nov. 2, 2017. https://www.mentor.com/company/news/siemens-mentor-towerjazz-first-commercial-compre-

hensive-suite-analog-constraint-checks-enhanced-autoreliability-offering

10. Taiwan Semiconductor Manufacturing Co. TSMC9000 Program. http://www.tsmc.com/english/dedicatedFoundry/services/tsmc9000.htm
11. Hogan, Matthew. December, 2017. “Leveraging Baseline Checks for Robust Reliability Verification.” Mentor, a Siemens Business. http://go.mentor.com/4WKLE
12. Hogan, Matthew. June 2017. “Automated and Context-Aware Latch-Up Checking with the Calibre PERC Reliability Platform.” Mentor, a Siemens Business. http://go.mentor.com/4T5a9
13. M. Hogan, S. Srinivasan, D. Medhat, Z. Lu and M. Hofmann, “Using static voltage analysis and voltage-aware DRC to identify EOS and oxide breakdown reliability issues,” 2013 35th Electrical Overstress/Electrostatic Discharge Symposium, Las Vegas, NV, 2013, pp. 1-6. http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6635948&isnumber=6635896

Maintainability Design Principles for Aircraft Maintenance Error Avoidance[1]

Clive Nicholas

[1] First published in 2016 Annals of MIRCE Mechanics, (page 17-28), published by MIRCE Akademy, Exeter, UK

Abstract

There has long been a philosophy in aircraft design that errors by maintainers are not the concern of the designer—maintainers should be trained not to make errors. That philosophy is rapidly changing. There is an increasing awareness by regulators, designer/manufacturers, operators and other organisations in the aircraft industry of the impact that the design characteristics of aircraft can have on safe and effective maintenance performance and, in particular, on the avoidance of maintenance error and the mitigation of its consequences.

Designers of aircraft, systems and components cannot influence all of the many factors that might influence maintenance performance and maintenance error. However, designers have an important role to play because design characteristics have a significant impact on the form, frequency and duration of the maintenance task and have important implications for the possible occurrence of maintenance error.

From a design perspective, there are a number of complementary and integrated strategies that can be adopted to effectively address the relationship between design characteristics and maintenance error including— i) to specify design requirements for aircraft, system and component design that directly address the possibility of maintenance error, ii) to integrate into design general principles that can be applied by the aircraft, system or component designer to assist -them in designing to prevent maintenance error or, if this is not practicable, to reduce its negative effects, and iii) to analyse design solutions for maintenance error through formal evaluation processes such as human hazard or human error analyses.

This paper examines the second of these strategies. It identifies and discusses the rationale of general design principles that can be adopted by designers as part of an overall design effort for maintenance performance.

It is based on the author's experience in developing design principles for maintenance performance and in developing and delivering practical training for designers of commercial aircraft.

Key Words
Maintainability, Maintenance Performance, Maintenance Error, Human Error, Design Principles

1. Maintenance Error

The aircraft maintenance process consists of a flow of tasks designed to maintain the safe and effective operation of the aircraft in service. Maintenance tasks typically include removal, installation, servicing, rigging, inspection, cleaning and other maintenance activities.

The execution of any maintenance task involves the possibility of error. Error in aircraft maintenance is the consequence of a complex interaction of many factors including system and maintenance task design, maintenance personnel and other resources, maintenance organisation, and the physical environment in which the maintenance occurs.

Maintenance error can be formally defined as the unintentional act of performing a maintenance task incorrectly that can potentially degrade the performance of the aircraft. For example, if a maintainer working in limited conditions of visual access fails to connect a component correctly the resulting maintenance error could be an incorrect installation leading to potential failure of the component.

Human behaviour is variable and is determined by a considerable range of factors that can vary significantly in different conditions and environments. Common factors can produce different responses and effects. Individual behaviours do not display uniformity and the designer would find it difficult to generate a design solution that would be applicable to the individual behaviours of maintainers. However, when designing an aircraft system or component the designer can address common patterns of behaviour manifest in reasonably foreseeable maintenance errors.

Empirical evidence indicates that there are common maintenance errors that tend to reoccur. Frequently occurring maintenance errors include:

- Wrong part installed
- Fault not found by inspection
- Incomplete installation
- Cross connection
- Fault not detected
- Wrong orientation
- Access not closed
- Wrong fluid
- Servicing not performed
- Fault not found by test
- System not deactivated
- Material left in aircraft

Most errors in aircraft maintenance are the result of unintentional or inappropriate actions that lead to maintenance error in a particular set of circumstances. There are also intentional actions on the part of the maintainer when, for some reason, it is either considered to be the correct action or a better way of performing a maintenance task. It should be recognised that maintenance error does not necessarily result in degradation of the aircraft.

An error can be recovered or corrected, before it results in consequential degradation. The consequence of maintenance error may be relatively insignificant or largely economic and recoverable. However, maintenance error can potentially lead to catastrophic consequences with loss of both aircraft and of life.

2. Design Impact

The correct completion of an aircraft maintenance task depends upon the interaction and interrelationships of the design characteristics of the aircraft and its operation in a particular environment. Design characteristics of the aircraft include technical systems and components. They also include the consequent design of maintenance tasks, procedures, manuals, tools, equipment and initial training of maintainers. Operation will include the characteristics of maintenance personnel, the maintenance organisation and the physical environment within which they work.

The potential for maintenance error arises where the maintainer and the aircraft interact through the maintenance task. The purpose of the aircraft is to provide a set of functions that enable its operation to deliver a safe flight that departs and arrives on schedule. The aircraft's ability to deliver safe and effective flights is sustained through maintenance to ensure that it functions as and when required.

The operation, maintenance and support of an aircraft are made up of related processes, which consist of tasks carried out by humans using physical resources.

A maintenance task can be described in the following terms:

- A maintenance task is any specified set of maintenance actions that is performed to maintain the required function of an aircraft component or system.
- The set of maintenance actions is related by their task requirement and their sequential occurrence in time.
- The execution of maintenance tasks involves human actions that comprise of some combination of cognitive ("thinking") and physical action ("doing").
- Each task requires an expected level of maintenance performance to complete each action and the task as a whole.

The successful completion of a maintenance task as specified therefore involves:

- The human: performance and limitations (e.g., vision, hearing, physique, perception, memory, fatigue, etc.).
- System and process design: the demands placed on human performance that are the result of design (e.g., operation, maintenance and support task and resource demands).
- System and process operation: the demands placed on human performance that are a result of operation (e.g., organisation, procedures, etc.).
- Physical environment: the demands placed on human performance that are a result of the physical environment in which the task is performed (e.g., climate, temperature, noise, illumination, etc.).

Aircraft designers are not in a position to control or directly influence all these factors. Nevertheless, the design of aircraft systems and components can have a significant impact on maintenance performance. System and component design characteristics can promote correct performance of the maintenance task. Importantly, design characteristics can potentially reduce the likelihood and consequences of maintenance errors and hazards to the maintainer safeguarding both the aircraft and the maintainer.

As previously stated, the maintainer and the aircraft interact through the maintenance task. It is through the maintenance task that the aircraft affects the performance of the maintainer and the maintainer affects the performance of the aircraft. The design of the system or component will influence the type, frequency and duration of maintenance tasks carried out in operation.

Key questions for the designer to consider are:

- Shat types maintenance tasks does the design generate and what actions do they involve?
- How often is the maintenance task needed and how long will it take?
- What demands does the design place upon the capabilities of the maintainer to complete maintenance task?
- Can the demands of the task exceed the possible limitations of the maintainer?

The complexity of design configuration, physical form, weight, location, access, method of installation, visual information and similar factors play an important part in determining the demands placed upon the level of maintenance performance required to successfully complete a maintenance task. Different designs will have different effects on maintenance performance. For example, the use of fewer parts may influence how easy it is to do the task—improving maintenance performance and reducing the likelihood of maintenance error.

Aircraft maintenance often involves complex processes that place considerable demands upon the maintainer to perform at the level required by the maintenance task. Maintenance often occurs in environments that also often place considerable demands upon the maintainer.

It is important to recognise the human capabilities and limitations of the maintainer and the capabilities and limitations that are inherent in any aircraft design. It involves the design of aircraft so that the relationship between the aircraft design and the maintainer effected through the maintenance task will result in optimal maintenance performance that minimises demands on maintainers that could lead to maintenance error.

The design of aircraft systems and components and the operational environment in which that design functions will influence the behaviour of the maintainer—for example, how easy it is to complete the task. Design characteristics can generate tasks that are within the capabilities and limitations of the maintainer that have a potentially positive effect on maintenance performance. Equally, design characteristics can challenge the capabilities and limitations of the maintainer and have a potentially negative effect on maintenance performance. Amongst other consequences, such as decreased maintenance efficiency, this could lead to error or personal injury during maintenance.

Design can therefore affect the vulnerability of an aircraft to maintenance error and the consequences of that error. By actively integrating general principles that address maintenance error into the design process, it is possible to create design characteristics that can possibly prevent or reduce maintenance error (e.g., sealed units or colour coding), or, eliminate or mitigate the consequences of maintenance error (e.g., isolation or partial operation).

3. General Design Principles

In developing design strategies and principles that enable the practical realization of these strategies through physical design characteristics, it is important to recognize that error is an integral and important part of fundamental human behaviour—it is part of the normal cognitive and learning processes of the human. Indeed, error in itself is not inherently problematic. It is only problematic when its consequences bring about unwanted or negative consequences. Design strategies should therefore attempt to avoid errors or to contain the consequences before they become negative. Error in maintenance is a normal part of maintenance operations that can be addressed during the design process.

Design strategies may revolve around two basic approaches. The first is avoidance of error. Here the error may be completely avoided by prevention. Examples of this type of strategy include designing out operation significant maintenance tasks, the design of components that are physically impossible to assemble or install incorrectly and the use of staggered part positions that require a specific configuration or sealed units that do not require intervention.

It is also possible to reduce the frequency of occurrence of error. Examples of error frequency reduction include the use of different part numbers, colour coding, shaped switch tops, locking switches, standard display formats, standard direction of operation, convenient access panels, reduction of servicing frequency, protection against accidental damage, or lubrication points that do not require disassembly.

The second is tolerance of error. Here mechanisms to detect error, to reduce the impact of error, and to recover error may be employed. Mechanisms to detect error may include built-in tests, functional tests, illuminated test points, functionally grouped tests or warning lights. Detection error can also include initial training of the maintainer for system state recognition. Reduction of the impact of error can be achieved through strategies such as isolation of the consequences of error, the ability for partial operation or the use of redundancy in systems or components. Recovery of error may be achieved through self-correction, the development of recovery procedures or specific training for error recovery.

Specific design objectives can be summarised as follows:

- Design that absolutely eliminates any possibility of an identified maintenance error or eliminates its consequences.
- Design that reduces the size of an identified maintenance error or reduces the extent of its consequences.
- Design that reduces how often an identified maintenance error, or how often its consequences, are likely to occur.
- Design that ensures that the maintenance error or its consequences is evident under all maintenance conditions, easy and rapid to detect, and is detected before flight.
- Design that ensures that following a maintenance error the means to return a system to its correct state are evident, easy, and timely.

In practice, the strategies of avoidance and tolerance are complementary and it may be felt necessary to design using a combination. An error tolerant design may be combined with error avoidance mechanisms to produce a robust design. Total avoidance of error may be considered to be an ideal given the nature and variability of human performance—error tolerance will capture and contain errors that fail avoidance mechanisms.

The general design principles discussed below provide practical means by which these strategies can be realised.

3.1 Appreciate the Maintainer's Perspective of the Aircraft

Designers design systems or components to deliver their required functionality. Maintainers are responsible for maintaining that functionality over the life of the aircraft whilst ensuring safety standards and operational requirements are met.

As a consequence, maintainers have a very specific perspective of an aircraft that will focus on the efficiency and safety of maintenance. Maintainers look for "maintainer friendly aircraft" whose design characteristics enable them to achieve good maintenance performance that deliver the aircraft back into service when required by the operator and that will complete the flight in safety.

From the maintainer's perspective therefore questions such as:

- How long will the task take?
- Is the task complicated?
- How often is the task required?
- Do I need special training?
- Do I need special tools or equipment?
- Could I make an error?
- How will I know if things go wrong?
- Where is the item located on the aircraft?
- Is there enough space to work in?
- Can I see the item?
- Can I reach the item?
- Where will I carry out the maintenance?

are of critical importance in achieving the necessary standards of maintenance performance to achieve these objectives.

It is particularly important that the design of a system or component does not infringe normal maintenance practices and the reasonable expectations of the maintainer based on training and experience. Maintainers might reasonably expect, for example, that on a dial values will increase clockwise.

3.2 Understand and Design for the Aircraft Maintenance Environment

To fully appreciate the impact of design on maintenance performance it is important to understand the environment in which aircraft maintainers work. Aircraft maintenance generally takes place under conditions that are complex and very demanding.

Line maintenance, for example, is generally performed outside the hanger working on the airport ramp or apron area in all types of weather and climate, often at night with limited visibility. The environment is extremely busy with aircraft loading and servicing vehicles moving around. There is considerable noise and there are fumes from aircraft engines and APUs (auxiliary power units) running. Above all there is constant pressure to complete maintenance activities as quickly as possible to turn the aircraft around on time for departure. Operators are in the business of transporting passengers. Aircraft on the ground cost money and lose revenue for the operator.

Similarly, base maintenance that is generally carried out in the hanger involves an environment where there is a considerable amount of activity and pressure to get the job done. Having to meet exacting work schedules while still observing standard procedures and safety standards can be stressful. The hangar is generally noisy from the use of power tools and there are many fluids and substances (hydraulic fluids, cleaning compounds, fuel, paints, etc.) that are potentially dangerous. Maintenance is often carried out at night when the aircraft are not in use. This means the work requires regular shift working. Requirements for overtime working and call-outs are common. Maintenance tasks can be physically demanding involving lifting, working in uncomfortable positions or working at height on scaffolds or cherry pickers (lifts).

The aircraft maintenance environment places considerable demands upon the maintainer and upon maintenance performance. The physical environment has an impact on maintenance performance through:

- lighting
- climate (dry or humid climates)
- temperature (hot or cold temperatures)
- weather (rain, wind, ice, snow, etc.)
- fumes and toxic substances
- noise
- motion
- vibration

Clearly designers cannot directly influence the many factors present in the working environment that will affect maintenance performance. However, they can have an impact on maintenance performance by taking them into consideration during the design process and reflecting this in design solutions. For example, where maintenance tasks are carried out in extremely low temperatures it is important to consider whether a maintenance task generated by a particular design could be carried out whilst wearing gloves or other protective clothing. On aircraft lighting can be used where there are light limitations for critical tasks such as inspection.

Design solutions that consider the physical environment in which maintenance is conducted can reduce the potentially negative impact that it can have on maintenance performance.

3.3 Protect the Aircraft and Protect the Maintainer

Design solutions can actively influence both the impact that the maintainer has on the aircraft (e.g., through maintenance error or routine violation of procedures) and the impact that the aircraft has on the maintainer (e.g., through the health and safety effects of aircraft design).

Examples of design features that are tolerant to the consequences of maintenance error or resistant to the effects of maintenance activity and maintenance environment include:

- Designing out safety critical maintenance tasks
- Items physically impossible to assemble or install incorrectly
- Staggered part positions
- Partial operation or redundancy
- Shaped switch tops, display formats, direction of operation, etc.
- Warning lights and illuminated test points

Examples of design considerations to protect maintenance personnel from risks, hazards, incidents, injuries or illnesses include:

- electrical isolation and protection from high voltages
- adequate circuit breakers and fuses
- rounded corners and edges
- warning labels
- hot areas shielded and labelled
- hazardous substances and radiation not emitted

Protecting the maintainer is important not only from a health and safety perspective—demands placed on the maintainer that can be potentially injurious can also

The Boeing 777 Refuelling Panel

Boeing didn't think of the fact that existing fuel stands only reached a certain height to fuel under the wings of the airplane. The 747 was about as high as the fuel stands could go to reach that fuelling panel, and the panel designed on the 777 was thirty-one inches higher than the 747.

Fuellers got very upset. "Have you ever fuelled an airplane in a high wind at O'Hare?" they said, "It's really uncomfortable."

To go any higher without additional stability would be a safety issue. Unless the operators hired personnel who are eight feet tall it wouldn't work.

Boeing agreed to move the panel down the wing, closer to the fuselage, and, because the wing is slanted up, by moving it inboard it also came closer to the ground – within six inches of reaching the panel. Safety specialists allowed a stool to be put on the top of the fuelling platform to reach the panel.

Adapted from Sabbagh K. (1996) 21st Century Jet – The Making of the Boeing 777, Pan Books, pp. 73-74.

lead to the occurrence of maintenance error. Design can place undue physical stresses on the maintainer. The maintainer may be required to wear cumbersome protective equipment to work in particular areas of the aircraft such as fuel tanks. The fatigue that can result could generate error. Other stressing design characteristics are those that, for example, involve inadequate lighting, vibration or noise, undue strength requirements for maintenance activities, unusual positions in which to carry out maintenance, or proximity of hot surfaces. A maintainer who must work close to heat generating components in a humid environment may rapidly lose body fluid, through perspiration as a result of increasing body temperature, which will seriously affect the ability to function correctly. If working close to a hot component, the maintainer must continuously avoid being burned whilst undertaking the maintenance task. The presence of such psychological and physical stressors can potentially lead to error.

Airbus A320 Flap Rotary Actuator

There are four rotary actuators on each wing of the A320. The function of these actuators is to translate the rotary motion of the flap drive shaft into movement of the flaps. Following flap lock events, it was reported in several cases that the flap rotary actuators had recently been removed for re-greasing.

Investigation revealed that during accomplishment of removal or installation slight mis-rigging in the flap transmission had been induced. This was found as a contributing factor in the reported flap locks.

Existing flap rotary actuators filled with grease needed removal for re-greasing approximately every 5 years. A new type of actuator introduced is filled with semi fluid and is serviceable on the wing.

The design solution simplified the maintenance task by eliminating the need for removal/installation of the actuators and thereby removing the opportunity for mis-rigging.

Adapted from Airbus FAST Technical Magazine A320 Special May 2005.

3.4 Avoid Complexity of Maintenance Tasks

The design of a system or component will impact upon both the cognitive (thinking) and physical (doing) demands of the maintenance task. Complexity in design can generate complex maintenance tasks that are difficult to understand and difficult to do.

However, the avoidance of complexity in design need not compromise or constrain the technical design solution. The design principle is concerned with the effect that the design on the maintenance task—an advanced design solution does not necessarily generate complexity in maintenance.

3.5 Enable Adequate Maintenance Access

Accessibility means having adequate visual and physical access to perform maintenance safely and effectively. Adequate physical and visual access is needed not only for repair, replacement, servicing, and lubrication but also for troubleshooting, checking and inspection.

Examples of physical access considerations include:

- adequate access to frequent maintenance areas
- openings of adequate size
- avoidance of the need to remove large numbers of components, fittings, etc. to reach a component
- replacement of components with the least amount of handling
- workspace for manipulative tasks, body and tools positions and movements

B-1B Engine Visual Access

Each engine on the B-1B bomber has an accessory drive gearbox (ADG). A hinged access door with four thumb latches is provided on each compartment panel for servicing. The access door permits checking of the ADG oil without having to remove the compartment panel. However, the oil level sight gauge requires line-of-sight reading. The way it is installed, the gauge cannot be read through the access door, even with an inspection mirror. The entire compartment panel, secured with 63 fasteners, must be removed just to see if oil servicing is needed.

Adapted from Worm CM. (1997) The Real World – A Maintainer's View, Proceedings IEEE Reliability and Maintainability Symposium, IEEE.

Examples of visual access considerations include:

- avoidance of unnecessary obstructions to the maintainer's line of sight
- lighting level and direction

Some components by their function or requirements have to be located in poorly accessible areas—a design solution in such cases might be the use of integrated access platforms or other aids to access.

3.6 Positively Standardize and Positively Differentiate

Aircraft maintenance tasks are largely repetitive and standardised. Maintainers rely on pattern recognitions that are determined by their training and experience to identify system and component type properties and the form of the maintenance tasks that are required.

Commonality in design enables such pattern recognition and enhances maintenance performance. If, for example, a part has commonality in function and properties (and, of course, fully meets all requirements of the design specification) then it makes sense from the maintenance perspective to use common parts.

Similar systems or components with variations in configuration can reduce the effectiveness of maintenance and can be a cause of maintenance error. Re-enforcement of pattern recognition can also be applied to commonality in maintenance activities.

If a part does not have commonality with the function and properties of other parts then it makes sense from the maintenance perspective to make the differences obvious. This will provide a clear and unambiguous signal to the maintainer that there are differences in maintenance actions.

3.7 Build Error Detection Into the Maintenance Process

Design solutions can assist in the detection of maintenance error before aircraft dispatch. Design can determine how maintenance error is detected and by whom. Ide-

Boeing 777 Door Hinges

Early in the design process it was realized that there were three separate hinges that are complex parts. In addition, if the hinge came into the door at a different place on each door all the mating, parts would be different. It was recognized early on that the key to make all the parts common was to make the hinge common, notwithstanding the fact that the shape of the body was different.

As a result, not only is the hinge common but also all the mechanism is common. Indeed, 98 per cent of all the mechanism of the door is common.

Adapted from Sabbagh K. (1996) 21st Century Jet – The Making of the Boeing 777, Pan Books, p. 89.

JSF Landing Gear Sensors

The Joint Strike Fighter team has broken new ground by the use of landing gear sensors purely on the basis of improving maintenance performance.

Landing gear present many maintenance problems – one particular problem is the measurement of the amount of hydraulic fluid by observation. This maintenance task has led to damaged landing gear due to overfilling.

The JSF programme, on the recommendation of its prognostics team, has agreed to embed sensors in the landing gear in order to report the exact level of hydraulic fluid, and in doing so has avoided maintenance error and saved cost.

Adapted from A Prognosis Sensor Victory on the Joint Strike Fighter (JSF), DSI International, November 2004.

ally, maintenance error should be detected before the aircraft is handed back to service after maintenance has been completed. In practice, however, the flight crew often detects error either before take-off or, worse, in flight.

Mechanisms to detect error may include built-in tests, functional tests, illuminated test points, functionally grouped tests or warning lights but equally they can be very simple such as the use of physical indicators.

Ambiguous, difficult, complex or lengthy means to detect a maintenance error can affect the likelihood of detection being successful. Detection means should ensure that the maintenance error is evident under all maintenance conditions, easy and quick to detect, and detected before flight.

4. Conclusion

There is a growing awareness of the vital role that design has to play in influencing maintenance performance and, more specifically, the avoidance or mitigation of maintenance error and its negative effects on safe and effective maintenance activity.

The maintainer interacts with aircraft systems and components through maintenance tasks that are generated by design characteristics. Design will determine the characteristics of the maintenance task and influence the possibility of occurrence of error – it will also determine the possibility for error avoidance and tolerance. The purpose of this paper has been to put forward general design principles that can be practically adopted and implemented by designers to develop practicable solutions that address reasonably foreseeable maintenance errors.

The design principles have been developed from extensive investigation of maintenance error, its causes and consequences to specifically encourage the designer to consider the impact of physical design on the behaviour of the maintainer. The approach is not intended to prescribe design practice, to teach designers how to design, or to advocate further constraints to the design process but rather to add a vitally important dimension to existing knowledge and skills that will enhance maintenance performance and aviation safety.

References

1. Airbus (2005) FAST Technical Magazine A320 Special, May.
2. Courteney H. (2001) Human Centred Design for Maintenance, UK Civil Aviation Authority.
3. DSI International (2004) A Prognosis Sensor Victory on the Joint Strike Fighter (JSF), November.
4. Hessburg J. (2001) Air Carrier MRO Handbook, McGraw Hill.
5. Sabbagh K. (1996) 21st Century Jet – The Making of the Boeing 777, Pan Books.
6. Worm CM. (1997) The Real World – A Maintainer's View, Proceedings IEEE Reliability and Maintainability Symposium, IEEE.

Eulogy for Benjamin Blanchard

Benjamin Seaver Blanchard Jr. passed away on July 11, 2019. He was 89 years old. He is survived by his wife of 62 years, Dorothy H. (Dot) Blanchard. He leaves behind 3 children and their spouses, Becky and Merritt Beaver of Alpharetta, GA, Ben and Patti Blanchard of Denver, CO, and Lisa and Dan McCade of Charlottesville, VA, and 7 grandchildren, Alison (Jason), Briana, Brooke, Adam, Devon, Shelby and Laina, and 2 great-grandchildren, Kylen and Nathan.

Born in New York City on July 20, 1929, the son of the late Benjamin Seaver Blanchard Sr. and Eleanor Phillips Blanchard, Ben spent much of his early childhood with his maternal grandparents at their estate in Peterborough, NH. There he would shadow his grandfather, a civil engineer and likely inspiration for his own future career path. At age 11, Ben was enrolled in Dublin School, an all-male boarding school in New Hampshire, and a place he remembered fondly throughout his life for teaching him responsibility and springboarding his education. Ben attended the University of Maine for college, thoroughly surprising his Harvard educated father and grandfather, graduating with a B.S. degree in Civil Engineering in June 1951. Three days later, Ben was drafted into the U.S. Air Force to serve in the Korean War.

Following his honorable discharge as a first lieutenant and electronics maintenance officer, he accepted a position as a design engineer for Boeing in Seattle, WA, where he worked on radar systems and landing maintenance for B-52s (1954-1960). It was in Seattle where Ben met Dorothy Hardt, a vivacious elementary school teacher from Chicago, IL. From the beginning, her large and boisterous family welcomed him as one of their own, a point of pride for only child Ben. Following subsequent jobs at Sanders Associates (1960-1961) and Bendix Corporation (1961-1963), he worked on defense contracts as the Design Assurance Department Manager for General Dynamics in Rochester, NY (1963-1970). After 17 years in industry and earning an MBA from the University of Rochester, Ben began his academic career as Director of Engineering Extension at Virginia Tech.

During his 26-year tenure at his beloved Virginia Tech, he helped establish the Virginia Cooperative Graduate Engineering Program and a graduate degree in Industrial and Systems Engineering (ISE), simultaneously authoring or co-authoring nine textbooks and publishing over 250 technical papers on the subject. He eventually achieved the rank of Professor and chair of the graduate program in ISE. Through work and personal travel, he visited 52 countries, leading continuing education workshops and presenting papers in most of them. He retired from academia in 1997 as Assistant Dean of Engineering for Public Service and Professor Emeritus. He was a recognized champion of, and leader in, his field, receiving the International Council on Systems Engineering (INCOSE) Pioneer Award in 2000 and the International Society of Logistics (SOLE) Founder's Medal in 2001, among many other accolades.

When asked what kept him working so hard throughout his multifaceted career, Ben simply said, "family." He adored his children and grandchildren, taking them on trips, sharing his love for sailing, history, and board games, and doting on them at his home on Claytor Lake. He was eternally gracious and poised, a joy to be around, and would most definitely be embarrassed by this collection of his accomplishments.

About the Authors

A Framework for a Defense Systems Effectiveness Modeling & Analysis Capability

Jerrell Stracener, Ph.D. is a Senior Research Associate in the Southern Methodist University (SMU) AT&T Center for Virtualization with a research focus on U.S. defense applications and systems effectiveness. He was Founding Director of the SMU Systems Engineering Program, and as Professor of Practice, taught graduate-level courses in engineering probability and statistics, systems reliability and availability analysis, integrated logistics support (ILS), and performed/directed systems engineering research and supervised PhD student research. He served in the U.S. Navy and earned both PhD and MS degrees in Statistics from SMU and a BS in Math from University of Texas at Arlington.

John M. Green is currently a Senior Lecturer in the Systems Engineering Department at the Naval Postgraduate School. He holds an MBA and MS in Computer Science from University of New Haven, an MA in International Relations from Salve Regina College, and a BS in Physics from Saginaw Valley State University. He is also a graduate of the Naval War College, College of Command and Staff. Mr. Green is a Senior Member of IEEE and AIAA. He is also a member of MORS, ASNE, INFORMS, the Association of Old Crows, and INCOSE.

Detective Maintenance

V. Narayan has graduated in mechanical engineering and obtained part-qualifications in electrical engineering as well as industrial management. He has worked in light engineering, oil and gas, pharmaceuticals and automobile ancillary industries for almost 40 years and across several countries and cultures. Narayan headed the Royal Dutch/Shell's Maintenance and Reliability Center of Excellence. As head of maintenance strategy at Shell UK, he managed the successful introduction of Reliability Centered Maintenance in their offshore operations. Narayan has also trained several hundred students in maintenance strategy selection, reliability engineering, asset integrity and related topics His first book was Effective Maintenance Management (second edition Industrial Press, NY. ISBN-13: 978-0831132491). His next book, written with two associates, describes 42 real life case-studies, illustrating the "how-to" aspects, was published in 2007 (Industrial Press, NY. ISBN-13: 978-0831133238).

Leveraging Foundry Baseline Checks for Robust Reliability Verification

Matthew Hogan is a product marketing director for Calibre Design Solutions at Mentor, a Siemens Business, with over two decades of design, field, and product development experience. Matthew is an active member of the International Integrated Reliability Workshop (IIRW), is on the Board of Directors for the ESD Association (ESDA), contributes to multiple working groups for the ESDA, and is a past general chair of the International Electrostatic Discharge Workshop (IEW). Matthew is also a Senior Member of IEEE, and a member of ACM. He holds a B. Eng. from the Royal Melbourne Institute of Technology, and an MBA from Marylhurst University.

Maintainability Design Principles for Aircraft Maintenance Error Avoidance

Clive Nicholas was a Rendell Scholar and holds a Bachelor's Degree in Economics, a Postgraduate Certificate in Education, and a Master's Degree in Engineering. he was the Deputy Director of the Research Centre for Managing Industrial Reliability, Cost and Effectiveness, (MIRCE) at the School of Engineering, in Exeter University (1988-1999). In 1999 he left the Exeter University and became the Director of Operation of the MIRCE Akademy, at Woodbury Park Exeter, UK 91999-2010). During this period Clive closely collaborated with Operability Department of Airbus (Bristol and Tolouse0 on several projects related to aircraft maintainability and maintenance. He also developed a few training courses for the design engineers at the Airbus, which clearly exposed intricacies of maintenance reality and design perceptions of it.

Colophon

The Journal of Reliability, Maintainability, & Supportability in Systems Engineering

Editor-in-Chief: John E. Blyler
Managing Editor: Russell A. Vacante, Ph.D.
Production Editor: Phillip S. Hess

Office of Publication: 9461 Shevlin Court, Nokesville, VA 20181
ISSN: 1931-681x

Instructions for Potential Authors

The Journal of Reliability, Maintainability and Supportability in Systems Engineering is an electronic publication provided under the auspices of the RMS Partnership, Inc. on a semi-annual basis. It is a refereed journal dedicated to providing an early-on, holistic perspective regarding the role that reliability, maintainability, and supportability (logistics) provide during the total life cycle of equipment and systems. All articles are reviewed by representative experts from industry, academia, and government whose primary interest is applied engineering and technology. The editorial board of the RMS Partnership has exclusive authority to publish or not publish an article submitted by one or more authors. Payment for articles by the RMS Partnership, the editors, or the staff is prohibited. Advertising in the journals is not accepted; however, advertising on the RMS Partnership web site, when appropriate, is acceptable.

All articles and accompanying material submitted to the RMS Partnership for consideration become the property of the RMS Partnership and will not be returned. The RMS Partnership reserves the rights to edit articles for clarity, style, and length. The edited copy is cleared with the author before publication. The technical merit and accuracy of the articles contained in this journal are not attributable to the RMS Partnership and should be evaluated independently by each reader.

Articles should be submitted as Microsoft Word files. Articles should be 2,000 to 3,000 words in length. Please use ONE space after periods for ease of formating for the final publication. Article photos and graphics should be submitted as individual files (not embedded into the article or all into the same file) with references provided in the article to their location. Charts and graphics should be submitted as PowerPoint files or in JPEG, TIFF, or GIF format. Photos should be submitted in JPEG, TIFF, or GIF format. All captions should be clearly labeled and all material, photos included, used from other than the original source should be provided with a release statement. All JPEG, TIFF, or GIF files must be sized to print at approximately 3 inches x 5 inches with a minimum resolution of 300 pixels per inch. Please also submit a 100-125 word author biography and a portrait if available. Contact the editor-in-chief, John Blyler, at j.blyler@ieee.org for additional guidance.

Please submit proposed articles by October 1 for the Spring/Summer issue of the following year and April 1 for the Fall/Winter issue of the same year.

Permission to reproduce by photocopy or other means is at the discretion of the RMS Partnership. Requests to copy a particular article are to be addressed to the Managing Editor, Russell Vacante at president@rmspartnership.org.

The Journal of Reliability, Maintainability, and Supportability in Systems Engineering

Summer 2019