EPRI

Reliability and Preventive Maintenance: Balancing Risk and Reliability

For Maintenance and Reliability Professionals at Nuclear Power Plants

Technical Report

TEAM

Reliability and Preventive Maintenance: Balancing Risk and Reliability

For Maintenance and Reliability Professionals at Nuclear Power Plants

1002936

Final Report, December 2002

EPRI Project Manager
M. Bridges

EPRI • 3412 Hillview Avenue, Palo Alto, California 94304 • PO Box 10412, Palo Alto, California 94303 • USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

ORGANIZATION THAT PREPARED THIS DOCUMENT

Applied Resource Management

ORDERING INFORMATION

Requests for copies of this report should be directed to EPRI Orders and Conferences, 1355 Willow Way, Suite 278, Concord, CA 94520, (800) 313-3774, press 2 or internally x5379, (925) 609-9169, (925) 609-1310 (fax).

CITATIONS

This report was prepared by

Applied Resource Management
313 Nobles Lane
Corrales, NM 87048-7708

Principal Investigator
D. Worledge

This report describes research sponsored by EPRI.

The report is a corporate document that should be cited in the literature in the following manner:

Reliability and Preventive Maintenance: Balancing Risk and Reliability: For Maintenance and Reliability Professionals at Nuclear Power Plants, EPRI, Palo Alto, CA: 2002. 1002936.

REPORT SUMMARY

This report explains the connections between the reliability and availability of commercial nuclear power plant components and the preventive maintenance (PM) they have been subject to. The scope includes a simple outline of the main characteristics of failure rates, probabilities of failure-on-demand, reliability, and availability of equipment in general, and relates these characteristics to elements of existing testing and preventive maintenance programs. It also covers the processes of optimizing preventive maintenance programs, of achieving balance between reliability and availability, and of monitoring reliability and availability at nuclear power plants. In particular, it relates these topics to the decisions required of personnel who manage and implement preventive maintenance programs. The report is therefore geared towards maintenance professionals who are seeking to understand and apply concepts involving reliability, availability, equipment monitoring, and preventive maintenance.

Background

In recent years, the growing trend toward risk-informed regulation of the nuclear industry has placed increasing requirements on plant personnel to incorporate rigorous quantitative methodology into their maintenance programs. A case in point is 10CFR50.65, the Maintenance Rule, which requires the use of quantitative risk concepts formerly familiar only to specialists in the field of probabilistic safety analysis (PSA). Other new concepts have recently been introduced into the regulatory process, such as the balance between the reliability achieved by improved PM and the equipment unavailability that results from the performance of PM. A significant challenge has been to provide ways to monitor the reliability and availability of equipment over short periods.

Objectives

- To describe the most important results and insights on the relationships between preventive maintenance and reliability, availability, and testing
- To show how these insights relate to maintenance-related decision making in a range of practical areas

Approach

Much of this material was developed over a 27-year period during which the author was involved with applying reliability data and reliability and availability methods to non-nuclear industries, developing PSA methods and data in the decade of rapid application of PSA between 1980 to 1990, and developing insights on the application of RCM to nuclear power plants in the decade from 1985 to 1995. This period also saw the development of the PM Basis database, a compendium of U.S. nuclear power plant preventive maintenance information for 60 major

component types. Some of the material in the report appeared in previous EPRI training courses on maintenance decision making, in earlier EPRI reports, and in the PM Basis Application Guideline.

Results

Sections 1 and 2 are devoted to the fundamental topics of reliability, availability, and testing; the objectives of performing preventive maintenance; the importance of different types of PM tasks; and a critique of various approaches to reliability centered maintenance. Section 3, The Dependence of Failure Rate on Time and Preventive Maintenance, and Section 4, A Balance Between Reliability and Availability, present topics which depend on an understanding of the basic material of Sections 1 and 2. Section 5 addresses the topic of monitoring reliability and availability, pointing out the difficulties in a nuclear power plant environment and some solutions. Finally, Section 6 contains basic guidance on manipulating reliability data, and the report's appendix describes several generic models connecting reliability to PM.

The information will assist maintenance professional with developing a working knowledge of these concepts in order to make decisions reliably in a rapidly evolving regulatory and economic environment. Utility personnel new to PSA are unlikely to find another single source with the same scope and depth of coverage. Furthermore, some of the results represent new insights.

EPRI Perspective

As the nuclear power industry continues to strive to reduce operating and maintenance costs, it becomes increasingly important that maintenance tasks are focused on the right equipment at the right time. Preventive maintenance can be an effective tool for improving plant performance. However, improperly applied PM can be an expensive exercise in futility. EPRI has continued to work to provide processes and techniques that will facilitate cost-effective equipment reliability. This report is a continuation of that effort.

Keywords

Reliability
Availability
Preventive maintenance
Maintenance optimization
Component reliability
Component availability
Testing

ABSTRACT

This report relates equipment reliability to the quality and coverage of the preventive maintenance program. It contains primarily technical rather than programmatic guidance on the design of a PM program to address a wide range of failure mechanisms. It also characterizes the reliability and unavailability which result from adherence to a PM program, and relates these quantities to Maintenance Rule performance criteria, probabilistic safety analysis (PSA) parameters, and testing programs. A major motivation has been the emergence of new attempts to improve PM programs at U.S. nuclear power plants, in order to improve equipment reliability. These improvements incorporate the INPO AP913 Equipment Reliability program, make use of the EPRI PM Basis database, and benefit from maturing Maintenance Rule programs at most plants. A detailed critique of reliability centered maintenance (RCM) methods is included to capitalize on all that has been learned about the performance of RCM-related approaches to PM optimization in the last 15 years. The emphasis is on enabling a maintenance professional to apply the available insights to improve maintenance and component reliability. It is hoped that this report will become a standard reference, because high equipment reliability and availability continue to grow in importance to deregulated utilities as risk-informed regulation expands.

ACKNOWLEDGMENTS

The author gratefully acknowledges review and technical insights on the topics of reliability centered maintenance and criticality assignments by Stephen M. Hess, Sensortex Inc., and by Glenn R. Hinchcliffe, G&S Associates.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1
BASICS OF COMPONENT RELIABILITY

This report explains the connections between reliability, availability, and preventive maintenance (PM), as these quantities and activities are understood and used in commercial nuclear power plants. The scope is limited to the reliability and availability of individual items of equipment, mainly because preventive maintenance is carried out at the equipment level. In this report, the terms *equipment* and *component* are used more or less interchangeably, notwithstanding the fact that *equipment* usually refers to a complex aggregate of smaller components, subcomponents, or piece-parts. The acronym SSC, meaning system, structure, or component, is also used because of its ubiquity in the emerging field of risk-informed regulation governing plant maintenance and other programs which address equipment reliability. However, this report treats an SSC as if it were a single item of equipment, and does not attempt to model its behavior as a function of its parts.

The synthesis or modeling of the reliability and availability of engineered systems (consisting of hundreds of components or pieces of equipment) is usually treated using reliability block diagrams and fault trees, with the reliability and availability of components as basic inputs, but those topics are outside the scope of this report.

1.1 Failures and Functional Failures

A component is failed when it can no longer perform its intended function. When defined in this way the failure is often called a functional failure. Component failures which do not fail important functions of the component do not usually merit the expenditure of significant PM resources. However, essentially all failures must be repaired at some point, whether they are functional failures or not. Because not all functions are important, not even all functional failures may merit the expenditure of PM resources.

Plant personnel responsible for reporting failures to a regulatory body usually prefer to limit the reporting to functional failures, on the basis that functional failures are the only important ones, and the cost of reporting is thereby reduced. When failure events are reported to an industry reliability database, and are used to develop failure rates for equipment, reporting only functional failures would unnecessarily restrict the data collected, and impede the generation of reliability parameters and the value to the industry. This restriction arises because the specific functions required of any component in its given application would prevent the reporting of certain failures which could be important if the component were operating in a different system. Maintenance Rule reporting of failures with respect to performance criteria is restricted to functional failures, whereas event reporting to the Equipment Performance and Information Exchange (EPIX) and to other reliability databases is not confined to functional failures.

Maintenance work orders almost never directly identify functional failures. Definition of SSC functions is obviously essential for identifying functional failures from maintenance work orders or from other event reports. For the Maintenance Rule, usually the system engineer or the Maintenance Rule coordinator decides which failures constitute functional failures. Performance of an unintended function, or spurious actuation of a function, should usually be considered functional failures because SSC functions often have implied or explicit "when" conditions attached to them.

The majority of equipment deficiencies are not functional failures, even when a subcomponent part has actually broken. Work orders are written to correct all manner of defects and only a small fraction of maintenance work orders—even corrective work orders—address functional failures.

1.2 Failure Modes

A valve failing to open may not be a functional failure even when failing to close is a functional failure, and vice versa. Failure to open and failure to close are *failure modes* of the valve. As used by safety analysts and reliability professionals, failure mode is the tangible effect of the failure, defined by the state of the equipment, as opposed to the effect on the system or plant. Several different failure modes may contribute to a functional failure—for example, the functional failure "failure to isolate system x" may result from the following failure modes of an isolation valve: 1) failure to close; 2) internal leakage; and 3) external leakage. The Institute of Nuclear Power Operations (INPO) has published standard failure modes for EPIX reportable equipment. An example of failure modes for 4kV switchgear could be the following: fails to close on demand, fails to open on demand, fails to break current (that is, extinguish the arc), spurious trip, and fails to trip.

The above definition of failure mode is the traditional one used by the probabilistic safety analysis (PSA) community. The EPRI report *Nuclear Power Plant Common Aging Terminology* also adopts this usage [1]. However, maintenance personnel commonly use "failure mode" to mean the cause-oriented description of the event, such as "binding of the valve stem from heat and corrosion." The maintenance worker's "failure mode" is thus close to the safety analyst's "failure cause." In this report, *failure mode* is reserved for traditional equipment states such as "fails to open," and the terms *failure mechanism*, *degradation mechanism*, or *failure cause* are used to describe how the failure mode came into existence.

Failure rates are often needed for a specific failure mode of the equipment. On the other hand, failure rates for all the failure modes may be added together (or not differentiated in the first place) to create a total failure rate for the equipment. The need to separate failure data into subsets corresponding to particular failure modes is one obvious way in which obtaining statistical data for adequate failure rate quantification is made more difficult.

The omission of information to support failure mode and functional failure determination from corrective maintenance work orders, or from other problem reports, results in the information being unusable for many reliability, risk assessment, and Maintenance Rule purposes.

1.3 Standby SSCs

A standby system or train is one that is not normally operating; it only performs its function when initiated by an automatic or manual demand. However, a standby SSC must be available continuously, in the sense that if its functions were suddenly required it must be able to perform them, starting at any future random instant. Auxiliary feedwater, containment spray, and hydrogen recombiners are examples of standby systems. Operation in which equipment is alternated between operating and standby states, each of considerable duration, (for example, weeks or months at a time) is usually treated as normally operating rather than standby.

If an SSC's function is required to be continuously available, an installed spare (for example, one of three 50% trains which is not in use, but which may be alternated with the other two trains) may not be considered a standby under some circumstances—for example, if the automatic switch-over to the installed spare is not reliable, or if it might result in a transient which could trip the system entirely. A reasonable test is to find out if the plant probabilistic safety analysis takes credit for the installed spare.

Failures of standby SSCs are not usually observable until the SSC is required to perform its function. These are called "hidden" failures, as they are not quickly revealed to the operating crew. This makes clear the purpose of "failure-finding tasks" in a PM program. Surveillance tests are typically tests to find hidden failures. It is good to keep in mind that these tests are mainly intended to reveal what has already happened. They do not prevent equipment failures from occurring, or give advanced warning of an impending failure, but provide the benefit of revealing when the equipment has already failed.

If an SSC fails during the time it is in standby (for example, because a hot environment made elastomer O-rings deteriorate and leak) its function will be unavailable from the time of failure until it is repaired after the failure is discovered at the next surveillance test. This could be as long as the whole interval between surveillance tests, making hidden failures a potent source of loss of system function. The hidden failure increases the chance that additional failures will further diminish the reliability of a set of redundant equipment and thus defeat the redundancy and the function being provided. The chance of additional failures will be proportional to the time the system is exposed to the failed component. The idea that redundancy will be compromised by additional failures before the first failure is discovered and repaired depends on the assumption that the standby component fails as a result of the passage of time during the standby period.

1.4 Degraded States and the Effect of PM

SSCs may degrade progressively and continuously in time, as in cases of wear-out—for example, normal wear of a gear or a switch. On the other hand, failure may occur suddenly and with no warning—for example, from an unusual mechanical load on the gear, or from bending of the switch handle, perhaps from errors in operation. In the case of the gear, the excessive mechanical load may also be introduced by a change in application, or by maintenance, which leaves it in a misaligned condition. The wear-out process may then not normally be present as is the case for normal wear, but initiated at a random point in time by an external event. The

characteristics of wear-out and random failure mechanisms will be examined in more detail in later sections, but it is clear that the development of degraded states toward the failure point as a function of time can be quite complex. It should also be clear that wear of the gear could be caused by other influences, such as a lack of lubrication, by using an incorrect lubricant, by contaminating the lubricant, or even by an inappropriate selection of alloy or surface treatment for the gear.

An SSC (for example, a compressor) is typically composed of a dozen or more subcomponents (for example, gears, pistons, and switches), most of which can degrade by more than one mechanism (for example, normal wear, improper lubrication, fatigue). In many of these cases different influences or stressors can cause progression of the failure mechanism at different rates in time (for example, the influence of contaminated lubricant may have quite a different effect on the failure time compared to the effect of a lack of lubricant).

This section has the objective of simply pointing out that this mixture of degraded states exists for all active components, and is composed of from ten to a few hundred combinations of subcomponents, degrading mechanisms, and stressors, each one with its own timetable for reaching the failure point. For consistency the following nomenclature is used:

- The subcomponents are called *failure locations* to avoid defining what a subcomponent is.
- The process which is occurring at the failure location is termed the *degradation mechanism*.
- The stressors which initiate the process, that drive it and make it progress faster or slower, thus determining the time at which failure occurs, are referred to as the *degradation influences*.
- The combination of the three preceding characteristics is collectively termed the *failure mechanism*. The failure mechanism is approximately equivalent to many levels of failure cause, but avoids singling out a proximal cause, direct cause, indirect cause, or root cause.

For the purpose of preventive maintenance planning and program optimization it is not usually necessary to draw sharp distinctions between these terms, except to insist that each case is adequately described so that a reasonably experienced maintenance professional can understand the cases under discussion. The maintenance professional is also more concerned with protecting against the continued progression of these mechanisms than with determining the point at which an advanced state of degradation becomes an actual failure. The result is that preventive maintenance task intervals are correlated much more closely with the time it takes to reach a severely degraded condition than with mean times between failures.

However, the term preventive maintenance refers explicitly to preventing failures, not preventing the failure mechanisms themselves. Well-designed PM tasks prevent the degradation mechanisms from proceeding all the way to the failure point, but very few PM tasks can prevent the existence of, or even delay the progression of, the degraded conditions themselves. There are a few very important exceptions to this general rule. For example, appropriate lubrication can often delay or even prevent the appearance of wear at sliding surfaces; ensuring correct charging conditions for large stationary batteries delays or prevents a large number of degradation mechanisms. But because most PM tasks do not prevent degradation mechanisms from

occurring, condition reporting schemes should anticipate that a significant level of degradation will be found at the performance of many PM tasks. Furthermore, many PM tasks will include repair work which in other contexts might be viewed as corrective maintenance. A recurring theme throughout this report is that if PM tasks are intended to prevent important functional failures, repair work carried out during the tasks, and often subsequent to them if done in a planned evolution, should properly be considered as preventive in nature, because the most important failures are indeed being prevented.

The failure rate is a strong function of service conditions, as well as preventive maintenance, and may also depend on the duty cycle. This is because the degradation influences, such as high or low temperatures, high vibration, or humidity, constitute the service conditions, and other stressors such as the number of starts and stops and accumulated run time essentially constitute the duty cycle. Thus, in a very broad categorization, failure rates can be stated specifically for high or low duty cycle equipment, and for severe or mild service conditions. The nature of the PM program itself is the third major influence on the failure rate. The most likely of these factors to change over time is preventive maintenance, often in response to the emergence of new failure mechanisms, but also in efforts to improve reliability or to reduce costs. When new failure mechanisms are discovered, PM changes are often made to attenuate failures from the new mechanisms. The initial increase in failure rate caused by the newly discovered failure mechanisms may be significant, but may not be measurable in the short term with any precision. The failure rate over longer periods often remains approximately constant because of the effectiveness of the new PM activities. If new PM activities were not brought to bear on the problem, the failure rate would increase over time.

Less often, PM is inadequate to attenuate the effects of new failure mechanisms. The impact of just one or two new mechanisms on the failure rate can then be considerable unless design changes can be introduced. The history of nuclear power equipment reliability generally shows that PM and design changes have compensated for the emergence of new failure mechanisms, so that reliability has generally improved, or at least held constant over time.

The maintenance influence on failure rate can be represented in an indirect way by stating the importance of the equipment function. Equipment whose failure leads to extremely serious consequences, such as a personnel hazard, a plant trip, or loss of a safety function, is designated as “critical” equipment. Currently, in most nuclear plants this equipment will have a fairly comprehensive PM program. Equipment that has significantly less functional significance is often designated as “noncritical” and therefore has a more superficial PM program. Equipment that has little or no functional significance will have no PM tasks and is designated as “run-to-failure.” The same piece of active equipment with the same duty cycle and service conditions is likely to have very different failure rates depending on its level of PM. For a more complete discussion of criticality, see Section 2.3.

Passive components generally receive less PM than active components, but this is because their failure rates are usually sufficiently low even with little or no PM; they are often considered to be “intrinsically reliable.” This is typically because there are far fewer failure mechanisms for passive rather than active components (perhaps ten, rather than hundreds). However, passive

equipment that does receive PM may also differ significantly in failure rate depending on the PM performed. The use of cathodic protection and cleaning for heat exchangers is a good example.

If the PM program is made more comprehensive to account for higher duty cycle and more severe service conditions, the effects of these conditions on the failure rate can be minimized, at least for active components. Compensation by such adjustments of the PM program in response to duty cycle and service conditions is usually applied for critical equipment, but may not be applied to the same degree at most plants for noncritical equipment. Failure rates for critical equipment are therefore likely to exhibit smaller variations with duty cycle and service conditions than failure rates for noncritical equipment.

1.5 Early Life Failures and the Risk of Doing PM

Pre-existing latent defects from the manufacturing or installation processes can cause failures throughout the life of the equipment. Of these, those which occur early in life are the most noticeable, because there tend to be more of them early in life, when wear-out degradation processes have scarcely even been initiated. Those which occur later in life may be less easy to discern among the rising quantity of wear-out failures and other random failures. The effect of sudden failures experienced early in life is often called “infant mortality” and is a common occurrence for many different types of components.

Quality procurement, acceptance testing, “burn in,” and other quality programs can be adopted to limit the number of early life failures, which are otherwise difficult to arrest because they tend to be random and sudden in nature. The number of early life failures has a tendency to decrease fairly rapidly as the component ages. However, the regular replacement of subcomponents during many PM tasks has a marked ability to continue to reintroduce some early life failures throughout the life of the equipment.

An additional effect which is very important is the occurrence of maintenance error. There can be several opportunities for serious error in any PM task, but the opportunity for error is maximized when intrusive PM activities are carried out on complex equipment, with many parts and required adjustments. This risk of doing maintenance is also increased significantly whenever plant personnel have inadequate training, equipment, information, or facilities. Additionally, risks increase when the task is done rarely or under stressful conditions; when control is lacking over contractors on-site or off-site; when susceptible equipment is not recognized; and when intrusive maintenance substitutes for less-invasive techniques or is performed too frequently. The effect can be amplified when a PM task is modified as a result of experiencing a few failures in a small subgroup of the equipment population, and then is applied to the whole population of the equipment type at a plant.

1.6 Random Failures

Random failures are defined by the fact that there is no marked concentration of their times of occurrence around a certain time. These failures occur with the same probability in any given time interval (say per year) either early in the life of the component, in mid-life, or at an

advanced age, or at any point in between. The failure rate from random failures is therefore constant in time. This means that these failures do not exhibit any kind of characteristic life—that is, bunching around such a value. They do, however, possess a value for the mean time between failures (MTBF) because one can always sum the failure times in a sample of failures and divide by the number of failures in the sample. This MTBF is simply equal to the reciprocal of the random failure rate. Obviously, random failures can occur with a variety of different values for the failure rate, depending on what causes them.

Many failure locations in complex equipment can experience failure mechanisms such as a broken switch handle or a mispositioned valve because of maintenance or operator error. Other examples are the plugging of orifices or filters with debris. Such events can occur with an approximately equal chance in a given period at any time in the life of the component. It is not possible to protect against random events if their occurrence is truly sudden, because they then possess no signature of impending failure and their randomness prevents a PM task from being performed at the "best time." Random failure mechanisms are therefore a serious challenge for PM programs.

Fortunately, there are many random failure mechanisms that do provide a signature of impending failure for a short time before failing. This means that protection can be achieved by performing very frequent PM tasks to identify this signature of advanced degradation. To be effective, such tasks must be done very frequently, because the occurrence of the events is random and the time duration of the failure signature is usually short. If a task is to be performed frequently it must be inexpensive and non-intrusive, so that the risk of doing maintenance does not cause more failures than are being prevented. Section 2.2 describes this kind of PM task, which is called a condition-monitoring task. A typical example is vibration monitoring. A very important benefit of condition-monitoring tasks is to provide better protection against random failures than other types of PM tasks can provide.

1.7 Wear-Out Failures

The concept of a wear-out failure applies best to individual degradation mechanisms with specific driving influences. A wear-out failure mechanism is characterized by the useful life, which can be quite long (for example, many years), during which the mechanism will not cause failure, followed by a period of time during which the failure rate from such a mechanism rises. Most components of the same type exhibiting this failure mechanism with similar duty cycles and service conditions will experience a similar useful life, after which failures will become more frequent, bunching around a characteristic value. Wear-out, depicted in Figure 1-1, means that the useful life and the characteristic life exist; the components are very reliable until a certain period of time has passed. Because of the zero failure rate through mid-life and the bunching of failure times around the characteristic life, the failure rate is not at all constant in time, and it cannot be claimed that there is about the same chance of a failure in a given time period at any stage of the component's life. The chance of a failure is zero until the wear-out begins.

Wear-out arises from many processes, such as friction at sliding or rolling surfaces, fatigue, gradual contamination, or change in material properties. Typical failure locations which

experience these effects are bushings, switches, soft goods, and springs. A notional degradation model for the occurrence of wear-out is for the degradation process to proceed more or less continuously, so that the degree of degradation or damage accumulates until it reaches a critical threshold, at which point failure occurs.

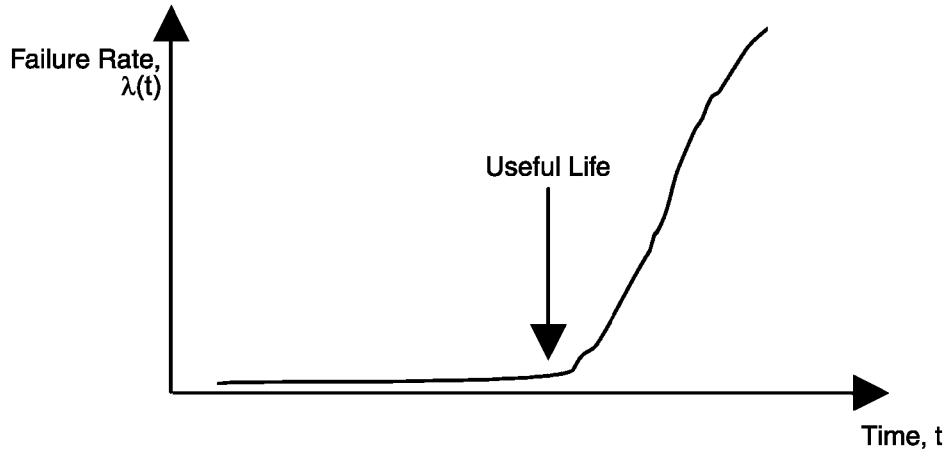

Figure 1-1
A Wear-Out Pattern of Times to Failure

Although wear-out suggests that a PM task could be timed to occur before a marked increase in failure rate sets in, the time duration over which the failures are experienced may be quite long compared to the expected failure-free period. That means that a task performed early in this failure time distribution, while the failure rate is still low, will often be performed on many components which do not yet require it. If the task is a component replacement, a large part of the component population may be replaced unnecessarily after the first few failures have occurred and a long time before most of them will fail.

This is illustrated in the following example. An analysis of Nuclear Power Reliability Data System (NPRDS) data revealed that the average age at failure for 203 failures of ABB HK circuit breakers was 7.3 years. The failure rate of all breakers (all modes) is somewhere in the range 0.01 to 0.001 per year, meaning that the mean time between failures (MTBF) far exceeds one hundred years. The short failure-free period of 7.3 years illustrated by the sample *of the first 203 failures to occur* has to be compared to the MTBF, which is the average of *all* failure times. The total population of breakers numbers in the thousands among all plants reporting; most of these breakers did not fail. Most of these will have failure times much larger than 7.3 years. A PM task performed on breakers at about 7 years will therefore often find the breaker in good condition.

Many wear-out failure mechanisms are randomly initiated in the sense that the stressor which drives the mechanism is not present initially, but occurs as a random event. For example, the ingress of water into electrical conduit initiates deterioration of the insulation, which may proceed for two or three years before it results in a short circuit. A large number of wear-out mechanisms are of this kind, initiated by a random event, but otherwise exhibiting all the characteristics of wear-out. In some of these cases the failure-free period is short, giving the failure events characteristics more like random failures with a short signature of impending failure, rather than like wear-out events. Other wear-out mechanisms involve failure-free periods

which depend strongly on the service conditions, such as the ambient temperature and/or humidity. Corrosion processes are generally of this kind, progressing to failure at markedly differing rates according to local conditions which are not known in advance. These are wear-out processes which appear to be almost random (unpredictable) as a result of the highly variable failure-free period.

For readers who may still be perplexed about the use of the term random, it is worth noting that failure events are always unpredictable, and therefore random in the trivial sense that the exact time of failure is not known. What distinguishes genuine wear-out mechanisms from the *completely* random mechanisms is the existence of a period of time, the expected failure-free period or useful life, during which the chance of a failure is very small.

1.8 The Aggregate of Failures for a Component

Complex SSCs are subject to a large number of failure mechanisms, of which a significant fraction may be random. The remainder, which are wear-outs—also large in number—will possess a wide range of expected failure-free periods, and will have failure time distributions which have a wide range of durations in time. The failure rate for the component, or even for one of its failure modes, will be the aggregate of these many effects. For the random mechanisms, the sum of numerous constant failure rates is simply a larger constant failure rate. For the wear-out mechanisms, the result may not be a constant failure rate, but it will be a failure rate which varies irregularly over time, with the variations having no discernable trend. The variations will diminish in relative importance as the number of failure mechanisms which are being added increases—that is, for more complex equipment. The sum of the random and wear-out contributions will be even more constant in time. This is the expected result for a large, complex piece of equipment—that is, the overall failure rate is essentially constant, even if there is no preventive maintenance.

When PM is performed, the majority of the wear-out contributions are greatly attenuated, further diminishing the time variations in failure rate. The only failures which still occur are those which "leak through" the PM defenses as a result of any of the following:

- Failure mechanisms of any kind which are simply not addressed by the individual PM tasks.
- Failure mechanisms which correspond to the part of wear-out failure time distributions which precede the task performance.
- Failure mechanisms which correspond to the part of wear-out failure time distributions which follow the performance of the task but where the task was not performed correctly (error of omission) so that the failure mechanism was not discovered and eliminated.
- Random failure mechanisms which were similarly "missed" by a task which was supposed to have detected them.
- Maintenance errors of commission which caused a failure which would not otherwise have occurred.

Keep in mind that replacement of subcomponents and performance of many intrusive PM tasks restores many subcomponents to an almost new condition, thus resetting the internal time frames which govern the failure time distributions discussed above, further scrambling their time dependences.

An idealized representation of equipment that "wears out" can be found in many textbooks. Such a bathtub curve is shown in Figure 1-2A, and displays an early decreasing failure rate as early life failures are gradually removed from the population, a low failure rate from random failures during the majority of the component life, and an increasing failure rate as wear-out occurs.

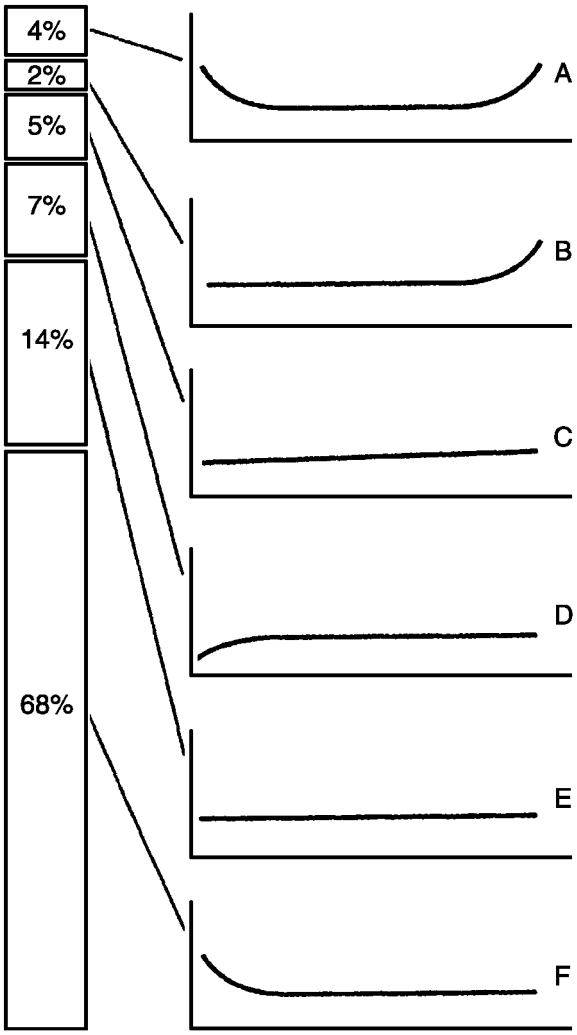

Figure 1-2
United Airlines Time-Dependent Failure Rates

Figure 1-2 resulted from the analysis of a large quantity of data from the airline industry in the 1960s [2], and amply demonstrates the point that complexity and preventive maintenance remove the long-term time dependence of the failure rate in all but a small proportion (6% to 11%; see Figure 1-2A, B, and C) of components, which happen to include aircraft reciprocating

(B) and turbine engines (C). In Figure 1-2, time is plotted along the horizontal axis, failure rate along the vertical axis.

The tendency for failure rates of complex maintained equipment to be constant in time is a distinct benefit in reliability data quantification. Sections 1.8, 1.9, and 6 show that it is difficult to determine failure rates accurately even when they are assumed to be constant in time.

1.9 The Probability of Failure-on-Demand

Some measure is required to represent the probability that equipment in standby will fail to function when required to do so. A probability of failure-on-demand can provide this measure.

The probability of failure-on-demand supposes that the occurrence of failures can be related to experiencing a number of tests or demands to perform a function. The number of demands is then the metric by which the occurrence of failures is measured, and the passage of time is irrelevant. The probability of failure-on-demand, P_f—for example, 0.05 failures/demand—is thus measured as the number of failures, n, divided by the number of component demands, N_d.

The number of failures and the number of demands may be aggregated over a group of like components subject to the same PM program, duty cycle, and service conditions:

$$P_f = n / N_d \qquad \textbf{Eq. 1-1}$$

The probability of failure-on-demand assumes the number of operations or demands is the relevant measure of the operational experience or exposure, but that time is not a relevant variable in measuring the extent of the operational exposure. It was first introduced for military applications such as the probability of failure of missile launches or of other “one-shot” devices where there is no comparable standby operational period except perhaps the shelf life.

For this concept to apply to standby power plant components, you have to imagine that the same number of failures will be experienced on average, among the same group of components whenever the number of demands is the same, regardless of widely varying elapsed time, or time in standby. For example, suppose two components each experience 20 demands, the first over three years in standby, and the second over one year. If they each have the same probability of failure-on-demand, the model requires us to assume that the expected number of standby failures will be the same in each case, even though the first component has withstood the service environment three times longer than the second.

Furthermore, you must believe that the number of failures experienced will be proportional to the number of demands, however widely that number varies, even in cases where the elapsed time, time in service, or run time is held constant. For example, if two components with the same probability of failure-on-demand experience 10 and 30 demands, respectfully, in the same standby period, the second is supposed to experience three times the number of failures as the first.

It is not easy to believe that these results would be true for power plant equipment in general, and it has led to some elaborate rules from regulators regarding which demands can be counted as "true" demands. The rules might exclude, for example, post-maintenance test demands, and demands which are repeated after it is established that the equipment operates successfully on the first demand.

We could assume that the number of demands is important because it is during such demands that the equipment experiences its only run time. However, such failures may be better treated as run time failures, and not confused with the probability of failure-on-demand. In any case, using the number of demands as a metric for the amount of run time ignores the undoubtedly important effect of exposure to the service environment during the standby period. Indeed, both the run time and the standby time can be accounted for by two different failure rates in time, the difference being merely an expression of the dependence of a time-based failure rate on the duty cycle. For example, a continuously running motor can be said to have a high duty cycle and will have a higher (run-time) failure rate than the (standby) failure rate for a standby motor which has a low duty cycle.

Often it is stated that the failure-on-demand model is appropriate because standby equipment fails at or during the demand, seeming to begin to function but then failing, in some resemblance to the failure of a one-shot device. However, in most cases the reason this happens is simply that the demand to start or change state is a stressful event, which causes a fragile degraded condition to exceed the threshold of failure. The degraded condition will have reached the failure threshold by exposure to the service environment during the standby period, either by a wear-out process or by a random process, and the demand is essentially a "probe" which evaluates the condition and finds it inadequate to support normal function.

One other situation which may suggest a failure-on-demand model is where the demands actually wear out the equipment. A good example is a diesel engine in which engine oil drains from the cylinders during time in standby, and is thus not available early in a fast start in sufficient quantity to prevent sliding wear. Another might be the cycling of a horizontally configured motor operated valve with a significant pressure drop across it, where each operation wears the valve guides. In both cases, the number of demands is indeed a metric for the amount of wear, which would not occur at all under other circumstances. Although these examples might justify the use of a failure-on-demand model, it is still not clear that its applicability extends to the whole piece of equipment, beyond the few failure mechanisms which have been described. In both cases, there exist a large number of other failure mechanisms which certainly would not support the use of such a model, and the wear attributed to demands might account for a minority of the failures, except in special cases.

These examples show that although the failure-on-demand model is widely used in PSA, it should be viewed with some skepticism for power plant equipment in general. Except in the particular cases where demands actually cause significant wear (pumps may be another example), a standby failure rate in time can be substituted to advantage.

The advantage of a standby failure rate stems from the fact that the time interval between demands then enters explicitly into the probability of failure. This is essential when optimization

of this interval is being assessed. In contrast, a constant probability of failure-on-demand does not provide the analyst the same opportunity, simply because the probability of failure is then just the probability of failure-on-demand, and is a constant. In fact, the failure-on-demand model predicts fewer failures as the time between demands increases, because any time period will then include fewer demands. Further analysis of this situation is deferred until failure rates in time have been more completely introduced in the next section.

1.10 Failure Rate

The failure rate in time supposes that the occurrence of failures can be related to the passage of time—that is, time is the metric which "generates" failures. Data on the number of failures, n, occurring to a group of N components in a time T, can be summarized as n/NT, in failures per year or per hour:

$$\lambda = n / (NT) \qquad \textbf{Eq. 1-2}$$

This is a *failure rate*, often denoted by the Greek letter lambda, λ. The failure rate is a frequency, and can be smaller or larger than unity, depending on the unit of time. For example:

$$2 \text{ failures/year} = 2.3E-4 \text{ failures/hour} \qquad \textbf{Eq. 1-3}$$

Frequencies less than about 0.1 can be viewed as probabilities over time periods of up to a few times the basic time unit. So in the above example it could be said that the probability of having a failure in one hour is 2.3E-4. However, probabilities have to be less than unity, so using a time period of one year to express the frequency, as on the left-hand side, prevents us from using the frequency for the probability.

In any rate process with a constant rate λ, (for example, failure rate, scoring rate, radioactive decay rate) the probability of *not* having an event in a given time period is given by:

Probability of *not* having an event in time t $= e^{-\lambda t}$

Therefore, the probability of having at least one event in time t $= 1 - e^{-\lambda t}$

So, if the failure rate is 2 per year, the probability of having at least one event in one year is $1 - e^{-2}$, or $1 - 0.135 = 0.865$. Putting in the numbers for a rate of 2.3E-4 per hour yields a probability of at least one event in one hour $= 1 - e^{-0.00023} = 1 - 0.99977 = 0.0002299$, or 2.3E-4, as stated.

Failure rates are hard to estimate accurately. Consider a single SSC with a failure rate of 0.05/year. Suppose it is observed and the failure rate estimated for many successive years. Most estimates of the failure rate would be zero because no failures would have occurred in the previous year. Suppose a failure occurs after 10 years; the estimate for that year changes abruptly to 1/year over latest year, or 0.1/year over the whole period. If a failure had occurred after only three years, the one-year estimate would also have been 1 per year, but the estimate for the whole period would have been 0.33/yr. Getting a reliable estimate for the failure rate depends on having

an observation period which contains at least a few failures. To contain a few failures, the observation period needs to be as long as a few mean times between failure (MTBF). For a constant failure rate, the MTBF is given by:

$$MTBF = 1 / \lambda \qquad \textbf{Eq. 1-4}$$

So a few (say 6) times the MTBF means $6 / \lambda$ years. This is a long time (~60 years) when the failure rate is in the region of 0.1 per year or less.

Evidently it is important to use groups of many SSCs to get more failures in a shorter calendar time. Aggregating failures among 10 like components means that 60 component years of experience can be accumulated in only 6 calendar years. If these 10 components are in your plant, you can be fairly sure to select them so that they are indeed like components, having the same PM program, duty cycle, and service conditions. But if the failure rate is 0.001 per year, 6000 years of component experience will be required. It will now be necessary to include data from other plants, and probably from the whole industry. Clearly, under these circumstances it becomes progressively more difficult to be sure that the data is homogeneous—that is, that the data comes from like components operated in a similar manner.

The failure rate, λ, can used to give the expected number of failures, n, in a given time, T, among a group of N like components:

$$n = \lambda T N \qquad \textbf{Eq. 1-5}$$

Suppose a continuously running pump has a failure rate of 0.07 per year. How many failures can be expected from 4 such pumps over 10 calendar years? The above equation shows that the expected number of failures is 2.8 over 40 pump years. Obviously, in any 10-year period, only 0, 1, 2, 3, 4 (and so on) failures could occur.

If a group of 24 safety-related air-operated valves (AOVs) used for isolation experienced 12 failures from all modes over a 10-year period, the failure rate over this period for all modes will be 12/(10x24) = 0.05 failures/year. The MTBF would be 1/0.05 = 20 years. On average, one failure would be expected in the MTBF period. If the MTBF is much larger than one operating cycle, as in this case, you should not expect even one failure in a single operating cycle.

If there are four failure modes that can cause a standby cooling train to fail—breaker fails open, suction valve fails closed, pump fails to start, and pump fails to run—and each failure mode has a failure rate of 2.5E-2/year, how likely is a train failure in a 2-year operating cycle? Summing over the four failure modes, the total failure rate per train is 4 x 2.5E-2 = 0.1/year. The train MTBF is therefore 10 years (= 1/0.1). Any 2-year period is therefore unlikely to contain a train failure, but it is not extremely unlikely. The chance is obviously roughly 20%. Using the expression stated above, the probability of at least one failure in 2 years, when the failure rate is 0.1/year, is $1 - e^{-2x0.1} = 1 - 0.819 = 0.181 = 18.1\%$.

1.11 Reliability

For a continuously operating SSC, the definition of reliability, R(t), to be found in textbooks and used in PSA analysis, is the probability that an SSC will perform its required function, under given conditions, over the whole of the stated time period given that it is operating at the start of the period. The time over which its functions must be provided is known as the mission time.

For a standby SSC, R(t) is the probability that the SSC does not fail during the mission, given that it starts on demand. We also loosely refer to the probability that the SSC will start on demand as being an aspect of its reliability, especially in the Maintenance Rule, and in other maintenance applications. The probability that a standby SSC will start when required to do so is obviously a legitimate aspect of reliability, but the mathematicians had first claim on the above, more narrow, definition.

We have already seen that when the failure rate is λ, the probability of not having a failure in a time, T, is $e^{-\lambda T}$. Therefore, the reliability is:

$$R(T) = e^{-\lambda T} \qquad \textbf{Eq. 1-6}$$

When λT is less than about 0.1, $e^{-\lambda T}$ is closely approximated by $(1 - \lambda T)$, which will be true almost all of the time for power plant components.

Unreliability—that is, the probability that an SSC fails in time T—is:

$$Unreliability = 1 - Reliability$$

$$So,\ Unreliability = 1 - R(T) = 1 - e^{-\lambda T}$$

$$= 1 - (1 - \lambda T) = \lambda T \quad whenever \quad \lambda T \quad is << 1$$

For a standby SSC the unreliability over the mission time must be added to the probability of failure to start from standby. The function of the SSC will fail to be provided if either the SSC fails to start or it fails to complete its run-time mission. Therefore, the total probability that it will fail to provide its functions is:

$$Total\ Failure\ Probability = P_f + \lambda T \qquad with\ a\ failure\text{-}on\text{-}demand\ model$$

The equivalent expression for a standby model is deferred until the topic of availability has been addressed, in the next section.

1.12 Availability

Equipment is often taken out of service at power plants for inspection, overhaul, testing, repair, or modifications, or to correct a degraded condition. Equipment which is not available to be started or run during these times, but whose functions are still required, is said to be unavailable.

Unavailable hours are input to the PSA model as a fractional downtime, or "unavailability" (a dimensionless fraction). The fractional downtime is a natural way to define this quantity, because it represents the chance that the SSC's functions will not be available at a future random time, not because the SSC will fail, but because it has already failed or is "down" for some other reason:

$$Unavailability = \frac{Hours\ functions\ are\ unavailable}{Hours\ functions\ are\ needed} \qquad \textbf{Eq. 1-7}$$

Unavailability often depends most on events that are not failures. A large part of unavailability is created by planned preventive maintenance work, and most of the rest is repair time to correct functional failures.

Unavailability is the price paid for preventing, finding, and repairing failures. Poor operations and maintenance practices can greatly increase risk through unavailability. Consider that 70 unavailable hours represents 1%, or 0.01 in unavailability (70/7000) at an 80% capacity factor. Compare this with a typical train failure probability in the range of 0.01% to 5%. The direct comparison is relevant because unavailability has the same effect on risk as does unreliability (failures to start, to change state, or to run), because the summed durations of these activities affect the probability of providing the required functions. A PSA model treats unavailability and unreliability in the same way.

An SSC can have good availability but poor reliability if it fails a lot but is always repaired promptly and returned to service quickly. An SSC can have poor availability but good reliability if it receives a lot of on-line PM and always works well. Unavailability for nuclear plant SSCs primarily involves maintenance. For example, if motor current signature analysis determined that a charging pump should be removed from service in the near future due to a possible impending failure of a motor winding, this would not be considered a functional failure, but a success of condition monitoring. Nevertheless, the event would be captured for Maintenance Rule monitoring purposes by the unavailability of the pump. Such "precursor" events that are not functional failures will often be captured by unavailability.

Unavailable hours caused by the performance of preventive maintenance activities, which are in turn performed to maintain reliability, need to be in some reasonable balance with the reliability which results. The topic of balance is treated in Section 4.

One other major contribution to unavailability involves failures of standby SSCs. When a standby SSC experiences a hidden failure, the SSC will remain unavailable in the failed state until the failure is discovered at the next demand. During this time the system is exposed to this failed condition of the SSC, and the time the SSC is failed is often referred to as the fault exposure time. In the limit, the SSC could have been unavailable for the whole of the interval between demands, or it could have failed just before the second demand. Either way, its time of failure is unknown. On average, the unavailable hours equal half the interval between surveillance tests, τ (tau), giving an unavailability *in the event of a failure* of $(\tau/2) / \tau = 1/2$. Denoting the standby failure rate as λ_s, the chance that a failure will occur in any given surveillance interval is $\lambda_s\tau$.

Therefore the average unavailability for a standby SSC is:

$$Average\ unavailability = \tfrac{1}{2} \lambda_s \tau \qquad \textbf{Eq. 1-8}$$

This quantity is best viewed as an unavailability, but it is the probability that a standby SSC subject to hidden failures will fail to change state on demand. It is the standby failure rate analogue of the probability of failure-on-demand, P_f. If, as indicated above, we generalize the meaning of reliability to include the probability that an SSC will change state on demand, then in the failure rate model this probability is an aspect of both availability and reliability.

As an example of the magnitude of this term, consider a standby motor with a standby failure rate of 0.05/year and a surveillance test interval of 3 months. For hidden failures, the average unavailability is ½ x 0.05 x 0.25 = 0.00625. This is equivalent to about 44 hours of unavailability per year (0.00625 x 7000 hours). The result can also be stated as the probability that the motor will fail to start on demand equals 0.00625.

We can now complete the discussion of the total probability of failure for a standby SSC which is required to start or change state at a demand, and then successfully complete its runtime mission (of duration, T). As in the case of the probability of failure-on-demand model treated in the previous section, the probability of failure to start under the standby failure rate model is added to the probability of failure to run:

$$Total\ Failure\ Probability = \lambda_s \tau / 2 + \lambda_r T \qquad \textbf{Eq. 1-9}$$

where subscripts of s and r denote standby and run time, respectively.

In principle, reality may be better served by a mixed model where some fraction, α, of the failure to start probability is derived from the standby failure rate model, and the rest from the failure-on-demand model:

$$Total\ Failure\ Probability\ (mixed\ model) = \alpha\lambda_s \tau / 2 + (1 - \alpha)P_f + \lambda_r T \qquad \textbf{Eq. 1-10}$$

However, this now requires three parameters for the failure to start probability, α, λ_s, and P_f, so this approach is rarely used because of the extra demand on data. Beware of any analysis which includes both the standby failure rate term and the failure on demand term without including the fraction, α. Using both terms in an unmodified form is an error, unless the α parameter has been subsumed into the definition of the other two parameters.

Notice that the probability of a standby SSC having a failure in time t, $\lambda_s t$, rises linearly with the time. Immediately after a successful surveillance test it has been ascertained that the SSC was not failed, so the probability of being in the failed state is reset to zero. A chart of the failure probability against time is a saw tooth form factor with the vertical lines at the surveillance test interval, as shown in Figure 1-3.

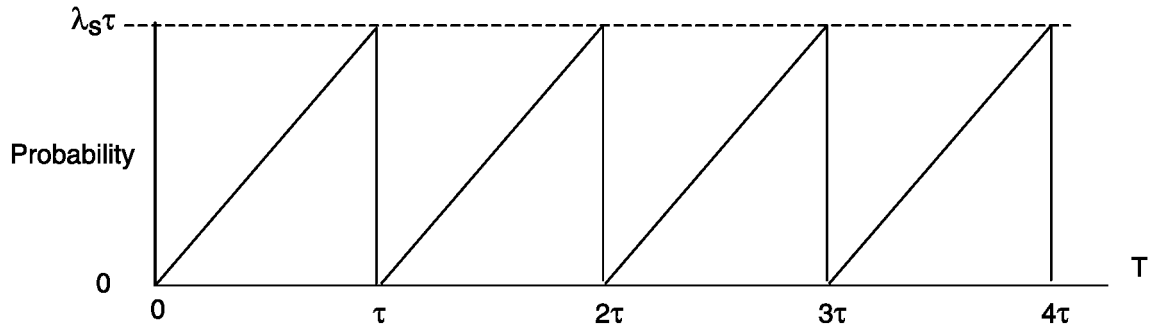

Figure 1-3
Probability of Being in a Failed State as a Function of Time for a Standby SSC Subject to Failure Finding (Surveillance) Tests With Interval τ

1.13 Hidden Failures and Testing

In contrast to the above definition of reliability used in safety analysis, maintenance personnel are more likely to use the term reliability more loosely to indicate either the failure rate or the number of failures to be expected in a certain time, or the probability of failure to start. Most of the time this does not matter very much, because maintenance personnel are not deeply involved with safety analysis. However, in the area of testing a conflict arises. To see how this happens we will examine the effect of changing the test interval on the expected number of failures, and the probability of failure to start, under the assumptions of the standby failure rate model and the failure-on-demand model.

We assume that the parameters of both models, λ_s, and P_f, are sufficiently well established. Consider what happens when the interval between tests, τ, is decreased to $\tau_{shorter}$. The tests therefore become more frequent.

1.13.1 Standby Failure Rate

Probability of failing a demand = $\lambda_s\tau_{shorter}/2 < \lambda_s\tau/2$ i.e., *it decreases*

Expected number of failures in time T

= Expected number of failures in $\tau_{shorter}$ x

Number of surveillance intervals in T

= $\lambda_s\tau_{shorter}$ x $T/\tau_{shorter}$

= $\lambda_s T$ i.e., *it stays the same*

The surveillance test interval *has no effect on the expected number of failures*, provided the standby failure rate stays at the same value. Changing the test interval obviously has no immediate effect on the standby failure rate.

The result for the standby failure rate model is that a safety analyst will say that the probability of failure improves when the test interval is decreased (the $\lambda_s\tau$ /2 effect), but a maintenance person might believe that shortening the test interval has no effect on his interpretation of reliability because the number of failures which will be experienced in any long time interval is unchanged.

1.13.2 Failure-on-Demand

Probability of failing a demand = P_f i.e., *it stays the same*

Expected number of failures in time T

= Expected number of failures in $\tau_{shorter}$ x

Number of surveillance intervals in T

= P_f x $T/\tau_{shorter}$

= $P_f\, T/\tau_{shorter} > P_f\, T/\tau$ i.e., *it increases*

In this model, a safety analyst would claim that the probability of failing a demand has not changed, but a maintenance worker would note that there are now more test demands, and the expected number of failures will therefore increase.

Also, the two models themselves give quite contradictory results, both for the probability of failing a demand, and for the expected number of failures. At the least this causes confusion, and can result in reasonable people disagreeing on the value of increasing the frequency of surveillance tests. Note also that if the test requires the SSC to be taken out of service, the increase in unavailability increases the total failure probability for both models by the same amount.

It is worth questioning the constant probability of failure-on-demand model because the above results do not seem to accord with experience. The expected number of failures is predicted to increase with more frequent tests and will decrease with fewer tests. This may be true if the tests are wearing out the equipment, but such cases are unusual. Furthermore, if test demands are indeed wearing out the equipment, increasing the frequency of such tests is not a good way to improve reliability.

Because increasing or decreasing the number of tests does not change the probability of failure-on-demand, in that model the contribution to the total failure probability is not changed at all by more frequent testing, which raises the question of why more frequent testing is a proper response to a desire to improve reliability. The answer is that more frequent testing can be a proper response to perceived reliability problems, but only under the standby failure rate model, and even then only when it is clear that tests are not wearing out the equipment. Even then, more frequent testing does not decrease the expected number of test failures, but at least it does not increase the number of failures. Increasing the frequency of testing in the failure rate model does

indeed improve the ability to control the hidden unavailability, which is the primary goal of surveillance testing. This result alone should alert us to the fact that changing the test frequency is really about improving the ability to find hidden failures and improving the average availability. It should not surprise us that in the failure rate model it has no effect on the number of failures, and in the demand model it has an undesired effect on the expected number of failures. If reliability in general is identified with the number of failures, as in the maintenance view, it is not surprising that more frequent testing is thought by many to do nothing to improve reliability.

It also seems that the failure-on-demand model seems at odds with the whole reason for surveillance tests, which is to find existing failures. Finding existing failures can only make sense if such failures can have already occurred. This requires a standby failure rate model to account for failures which occur over time in standby. The failure-on-demand approach seems to imply that failures *cannot* occur during standby because the constant probability of failure-on-demand can only generate a failure in response to a demand—that is, at the instant of the demand. The failure-on-demand model does not admit to any hidden unavailability at all. We have already introduced the fact that the reason why so many failures appear to occur at the demand is that this stressful event merely precipitates a failure from a severely degraded condition, which would have become a failure anyway, given a little more time in standby.

The best way to think about this confusing situation is that in the standby failure rate model there are two quite distinct mission times. The run time mission has already been described, and applies after the SSC has started or changed state at the demand. The new feature is the standby mission time, which is the period in standby between failure finding tests. A standby component truly has a mission during this standby period, the mission being to remain available to change state at a moment's notice. The probability of failing to start on demand is then the probability of failing during the standby mission—that is, it is the *standby unreliability*. Although this terminology is not in common usage, it accounts for why the hidden unavailability is really an unreliability and justifies the common reference to reliability in this connection.

Regardless of the model used to quantify them, if surveillance tests are really to influence the underlying failure processes, the tests must possess some predictive capability, because otherwise there can be no effect on either the standby failure rate or the probability of failure-on-demand.

Example: Historical data show that auxiliary feedwater pumps failed 6 surveillance tests in 60 pump years of experience. The test interval is 6 months, so assume there were 120 test demands. Compare the unavailability contribution from fault exposure time to the probability of failure on demand.

Answer: With a standby failure rate the rate is 6/60 = 0.1 standby failures per year. The probability of failure in a six-month interval is $\lambda_s\tau = 0.05$. Then the average unavailability = $\lambda_s\tau/2 = 0.025$. In the alternative failure-on-demand model the probability of failure on demand = $P_f = 6/120 = 0.05$ failures per demand.

The two models always differ by a factor of two if there is just one demand per test and no additional demands. However, it is common to include additional test demands. These will clearly reduce the probability of failure-on-demand, which could then become equal to or less than the average hidden unavailability.

Monitoring the number of surveillance failures is a way to account for the hidden failure time in the Maintenance Rule, because fault exposure time cannot be monitored (the failures are hidden). Except for under unusual circumstances, the time of failure and hence the time unavailable are not known. The $\lambda_s T/2$ term is only an average value, which is therefore of no use in tracking the actual hidden unavailability. Note that the term is stated here using T to represent the whole Maintenance Rule monitoring period, not the test interval, because the hidden unavailability over the whole monitoring period would be the quantity of interest in the Maintenance Rule situation.

1.14 A Maintenance View of Testing Models

The previous section has shown that the constant probability of failure-on-demand model for standby failures ignores likely time-related causes of standby equipment failure, and overweights the importance of demands in causing failures. The failure-on-demand approach became popular during the 1980s, when large numbers of PSA studies were being done, and was introduced more for the convenience of analysts than for the validity of its application. The alternative is to model standby failures as a failure rate process in which failures develop over time during standby. However, this ignores the (possibly) legitimate effects of the number of demands on the number of failures. Many PSAs make extensive use of standby failure rate modeling for some standby components, as well as the probability of failure-on-demand for others, although strong claims are not made in the industry for the appropriateness of these choices. IEEE data on reliability often provide both forms of failure data for the same equipment. The industry has not worked to eliminate one or the other approach because, given the same data, they each approximate reality and typically give results that are the same within a factor of about two. Uncertainties of at least this magnitude are acknowledged in both models.

The use of a failure rate in time to describe failures of continuously operating components is supported by the knowledge that such equipment suffers from accumulating wear, material property degradation both mechanical and electrical, fatigue cracking, corrosion and erosion, diminishing effectiveness of lubricants, the accumulation of contamination, the loosening of fasteners, deformation or settling of frames, mounts and other parts, and the random failure of some electronic items. These degradation pathways are influenced to a great degree by environmental factors such as heat, moisture, and radiation. It is a matter of common experience that the damage accumulates over time until the failure point is approached, whereupon an increase in stress, such as a temperature transient, or a sudden pressure pulse precipitates a failure. Central to all of these processes is the passage of time. Even when the rates of degradation in time are extremely variable, or when random events cause the failures, a large enough population of components and plant applications gives a more or less consistent average rate of failure occurrence. The time rate model of failure occurrence simply means that the passage of more time unquestionably leads to more failures. The passage of less time does not allow as much damage to accumulate, and the number of failures is less.

Standby equipment also fails from the same causes as operating equipment, but usually at a lower rate because the amount of wear is less, heat and mechanical stresses are lower, and standby rotating equipment is not pulling in environmental influences such as dust or salt-laden air to the same extent as when it is operating. Failure rates in time for standby equipment may therefore be lower than for normally operating equipment, although equipment which is alternated between running and standby will not experience as much difference. On the other hand, some influences can be more detrimental during standby, such as the relocation of lubricant, or the sagging and bowing of a horizontal shaft. Whether these effects subsequently accelerate wear and other damage in an amount proportional to run time during tests or by the action of starting or cycling depends on the equipment. In either case, however, the deleterious effects of a demand are added to those of the time-driven influences. The same is true of random influences such as damage or misalignment by personnel, or miscalibration; their influences add to the others which are always present.

At least a major part of the influences on failures of standby equipment such as pumps, motors, switchgear, check valves, modulating valves, and heat exchangers comes from processes occurring over time. Only in very specific circumstances (for example, pressure relieving valves) may demand-caused failures be a sufficiently dominating addition to this picture to justify the adoption of the failure-on-demand treatment at the expense of ignoring the time-driven effects. Most cases are in between.

The extent to which standby equipment should be treated as having a constant probability of failure-on-demand rather than a constant failure rate in time is a matter of which influences are dominant. For standby equipment of a particular type it is instructive to conduct two separate thought experiments. The first imagines two identical pieces of equipment which are subject to the same total number of demands, but over very different standby time periods. Are approximately the same number of failures likely to be experienced on each piece of equipment? If not, which has the most failures? The second experiment imagines the same two pieces of equipment but now one is subjected to many more demands (tests or operational cycles) than the other, but over identical time periods. Are approximately the same number of failures likely to be experienced on each piece of equipment? If not, which has the most failures?

For power plant equipment which is alternated between being run and being in standby (for example, switchgear, motors, pumps, check valves, modulating valves such as AOVs, or heat exchangers), the answers are likely either to favor the failure rate model (no appreciable change with demands, but an increase in failures with time in standby), or the situations are too uncertain or difficult to decide upon. A quick survey of the degradation mechanisms and influences encountered in the EPRI PM Basis project, which has prepared a technical basis for preventive maintenance programs, resulted in roughly three times more degradation influences on electric motors that were associated with time in standby than those associated with the number of demands.

1.15 Basic Events and Probabilistic Safety Analysis

The purpose of this section is to connect the above results on reliability and availability to the inputs to a probabilistic safety assessment. These inputs are called basic events. Each basic event

represents a way in which the important functions of an SSC fail to be provided when required. The quantities already described above and listed below constitute most of the basic events for an SSC:

- The probability of failure to start or change state (failure-on-demand model)
- The average hidden unavailability or probability of failing a demand (standby failure rate model)
- The probability of failure to run through the duration of a mission
- Unavailability due to repair of functional failures
- Unavailability due to performance of preventive maintenance

There may be multiple different probabilities of failure to change state for the same SSC, because different failure modes may be represented by such a basic event. For example, for medium voltage switchgear there could be a probability of failure to trip, failure to open on demand, and failure to close on demand. Other basic events may have the same characteristics. For example, an emergency diesel generator may have a probability of failure to run for 5 hours and also a probability of failure to run for 24 hours, depending on the accident sequences it is involved in.

The total probability of failure from all of these basic events can be found simply by adding up the individual basic event probabilities. This is because they all represent mutually exclusive events. Further details of the connection between reliability parameters, maintenance, and PSA parameters can be found in [3].

2
PREVENTIVE MAINTENANCE OBJECTIVES AND STRATEGIES

2.1 Definition of Maintenance

In the Federal Register, maintenance in general is defined as:

> The aggregate of those functions required to preserve or restore safety, reliability, and availability of plant structures, systems, and components. Maintenance includes not only activities traditionally associated with identifying and correcting actual or potential degraded conditions, i.e. repair, surveillance, diagnostic examinations, and preventive measures, but extends to all supporting functions for the conduct of these activities [4].

The EPRI PM Basis database defines preventive maintenance as:

> Planned tasks, either scheduled or done as a result of unacceptable equipment condition, performed to predetermined criteria and prior to failure, with the purpose of preventing unanticipated failure by:
>
> - Monitoring or inspecting equipment condition
> - Replacing or refurbishing prespecified parts
> - Functional testing to determine the ability to function [5]

Under this definition, preventive maintenance therefore includes:

- Periodic tasks such as inspection (clean, inspect, and adjust), calibration, overhaul
- Periodic repair or replacement
- Planned repair or replacement to correct a degraded condition (not periodic but prior to functional failure)
- Periodic tasks carried out by other departments, such as operator rounds or engineering walkdowns
- Condition monitoring tasks such as oil sampling or vibration analysis
- Surveillance, operability, or functional tests

Corrective maintenance is generally accepted to be maintenance actions taken to repair or replace a *failed* SSC, loosely analogous to repairing a functional failure, although the functional failures may or may not be important ones. This means that repairing SSCs which are not

functionally important and which are run to failure (that is, for which no PM is performed) is a part of corrective maintenance, but replacing failed subcomponents, even in a functionally important SSC would not be, if their failures did not cause functional failure of the SSC and were a part of planned PM activities such as refurbishment.

Calibration is a corrective action if an instrument has drifted outside its functional set point limits, whether determined by design basis or technical specifications. However, it is likely that recalibrating a drifted instrument which had not drifted outside specifications would not be considered to be corrective maintenance.

Within these guidelines, it is clear that the majority of maintenance work orders are preventive in nature, as they do not address repair of functional failures.

2.1.1 The Distinction Between Preventive Maintenance and Corrective Maintenance

The definitions for the terms *preventive maintenance* and *corrective maintenance* require careful consideration to state with precision, and the operational implementation of these definitions could vary somewhat between industries. When considering the quality and effectiveness of the maintenance program as a whole, and the proportion of maintenance resources spent to prevent failures compared to repairing them, it is important to decide which maintenance work orders are of each kind.

In a typical process for partitioning all work orders that address degraded or failed components, work orders which address regularly scheduled PM tasks are labeled “Regular PM.” These are the work orders which implement traditional time-directed PM tasks such as inspections and restore or replace activities, failure finding tasks such as surveillance tests, and condition monitoring, performance monitoring, and other predictive maintenance activities.

Continuing with the example, the category “On-Condition” refers to work orders in which degraded subcomponents, which were discovered during the execution of regular preventive maintenance tasks on the main SSC, are repaired or replaced. If this restorative work is carried out at a later date, it is typically performed under what most facilities regard as “corrective maintenance” work orders. But these degraded subcomponents (some may even be failed) are usually fully anticipated by the PM program, so the subcomponent degraded conditions or failures do not constitute the larger impact functional failures which the PM program is designed to prevent. An example would be tightening the packing on a pump after a leak is discovered during a routine inspection, providing the leak does not limit the function of the pump. A second example would be the planned changing of a motor bearing after high vibration is discovered during vibration monitoring. In this case, the motor could have a very important function, but the emergent condition is corrected by *planned* intervention before failure occurs. In total, there are a large number of these activities where the work of correcting degraded conditions (which were implicitly anticipated) may not be performed during the PM task in which they were discovered.

The insistence that the emergent on-condition work be planned before being considered to be the on-condition part of PM places significant constraints on the effectiveness of condition

monitoring tasks. If the emergent work is so urgent that it forces a high impact outage, it obviously has to be interpreted as true corrective maintenance. The word *planned* implies there is adequate time to properly plan the work so that the outage can be taken at a time when it still prevents loss of function but also minimizes economic impact.

"Expected CM" work includes the run-to-failure cases, which require corrective maintenance work orders to repair them. But these failures are expected to occur, and they are an anticipated aspect of the PM program. As introduced above, it is not proposed that these work orders should be classified as anything other than corrective maintenance work orders, but they form a class of expected corrective maintenance that does not necessarily indicate a poor-quality PM program, a class which could indeed be increased rather than decreased by maintenance optimization. In a similar way, the "Expected CM" work should also include work orders which repair failures of the components which receive only minimal PM. To the extent that some PM is indeed performed on this equipment, some of these failures are, in fact, unexpected, but the majority will be associated with failure modes which are not by choice protected by PM. It will not be cost-effective to separate the two types of work orders for this category of equipment whose failures have minimal impact. Classing all of these failures as "Expected" also emphasizes that they have been planned and anticipated by the PM program.

Finally, there are the true functional failures which constitute the more important or costly events which PM tries to prevent. These can claim to be "Unexpected," since PM is almost certainly performed to protect the system from them. Their repair can be labeled as "Unexpected CM." These categories of PM and CM are depicted in Figure 2-1.

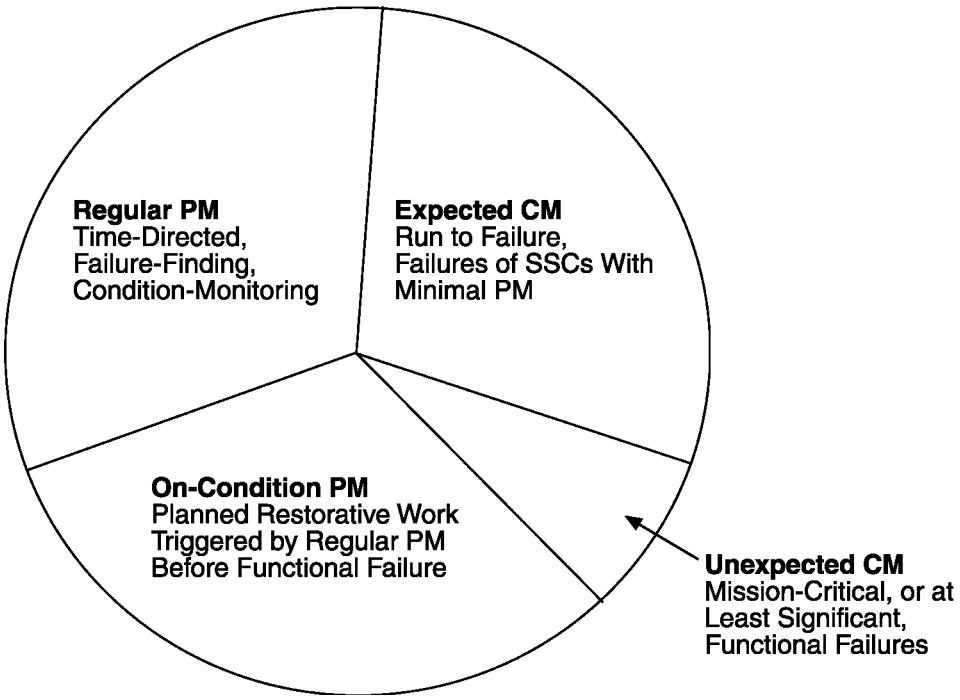

Figure 2-1
Major Classes of Preventive and Corrective Maintenance

In any application where the PM and CM distinction is relevant, such as the estimation of the costs of unreliability, it is important to classify work orders properly so that those addressing the on-condition work are included with the regular PM events on the PM side of the costs. Only part of this requirement can be met by careful process design. Training is also required, as inadequate personnel training on data reporting will result in incorrect classifications. For example, an issue in some plants is the reporting of true corrective work on a preventive work order because the opportunity is taken to perform a pending PM task. It also seems to be true that even if someone is assigned to review all work orders, some PM/CM categorization decisions require considerable experience, usually because of uncertainty over the level of functional impairment, or the degree to which on-condition work was really planned and was able to avoid a forced outage.

Even in an unattainably perfect PM program which eliminates all unexpected CM, there will therefore remain a significant CM cost, consisting of the expected contributions from running to failure the functionally unimportant components, and repairing those failures with minor economic impact.

The result is that we should anticipate that there will always be a significant CM cost, even in a perfect PM program, even when the on-condition costs are properly allocated to the PM program. The issue of whether to treat the expected CM costs as CM or PM is illuminated by this discussion. Treating them as CM acknowledges the fact that they are repairs of failures, albeit anticipated and relatively inconsequent ones. Adding their cost to the other CM costs does not distort the effectiveness of the PM program, because the PM program should be designed to minimize the total cost, by providing an appropriate balance between preventing failures and allowing them to occur. This is an important distinction to make: in helping to assure the safety of the plant, the PM program should minimize the total maintenance cost, not just the corrective maintenance cost. This concept is investigated in further detail in Section 4, on the balance between reliability and availability.

2.2 Types of Preventive Maintenance Tasks

2.2.1 Time-Directed

Time-directed PM tasks are scheduled tasks and are usually performed without knowledge of whether they are needed or not.

They include some of the most intrusive of all PM activities (that is, they require significant disassembly), and are usually timed to address prominent wear-out failure mechanisms, although they undoubtedly also address many random failure mechanisms as well. However, time-directed tasks are not usually very effective against random failure mechanisms, because the chance that such a random mechanism will be present when the task is performed is small.

The primary goal of a time-directed task is to improve the condition of the equipment, not just to diagnose it, although diagnosis is typically involved. In place of calendar time, the number of operations or cycles may replace time as the metric by which the task is scheduled. This is done

in cases where there is a significant effect of duty cycle on the failure rate. In general, time-directed tasks are done to replace, clean, adjust, and so forth, well before wear-out begins at the end of an expected failure-free period.

The wear-out failure mechanism with the combination of highest likelihood of occurring and the shortest time to the first expected failures will often drive the scheduling of a time-directed task. However, there may be other wear-out mechanisms with shorter times to first failures but which are not encountered as frequently. These can then cause occasional failures which are not protected by the task.

Examples of time-directed PM tasks are:

- Detailed Clean and Inspect – at 3R
- Calibrate – at 5Y
- Replace Filter – at 3Y
- Replace Packing – at 6Y
- Overhaul – at 10Y
- Rotate Tires – at 10,000 miles

2.2.2 Condition-Monitoring and Predictive

Condition-monitoring tasks are PM tasks which perform some kind of monitoring to discover the condition of the equipment. They are usually scheduled, and must be performed frequently to have a good chance of detecting a short-term signature that is precursor to a random failure or a short-term wear-out mechanism which may be initiated by a random event. In order to be performed frequently, any PM task should be non-intrusive and relatively inexpensive to perform. Condition-monitoring tasks typically possess these characteristics. Trending the results of some condition-monitoring tasks may permit prediction of the failure time.

However, the primary goal of condition-monitoring tasks is to monitor the condition of the equipment; predicting the time of failure is an additional capability possessed by some condition-monitoring tasks which could therefore be called predictive maintenance tasks. In view of the need, described above, to plan restorative actions to avoid a functional failure but also, if possible, to avoid onerous operational costs, it is clear that such predictive capabilities can be very beneficial.

Despite these distinctions, common usage makes little differentiation between condition-monitoring tasks and predictive maintenance tasks, and the two terms are used more or less interchangeably.

Even though condition-monitoring tasks are relatively non-intrusive, they may require equipment isolation (for example, MOVATS testing). They have to be non-intrusive if they are to carry less risk of introducing early life failures when performed as frequently as required to address random failures. This could be as frequently as every month. These tasks may be done on a

sampling basis to fit the schedule and to make most efficient use of special equipment that is needed.

Condition monitoring is practically the only way to defend against failures caused by random events (such as improper tightening of cable terminations) and short-term wear-out mechanisms which are driven by a random event (such as rapid wear caused by contamination of lubricant).

Examples of condition-monitoring tasks are:

- Thermographic IR scan
- Oil sampling and analysis
- Ultrasonic minimum wall thickness test
- Differential pressure measurement
- Vibration analysis
- Observation of packing leaks – operator rounds
- External visual inspection

The ultrasonic minimum wall thickness test in the above list is an example of a PM task which is difficult to classify. Although it is non-intrusive and reveals the condition of the equipment, it involves special equipment and access requirements, which usually means it is not performed very frequently. Nevertheless, to be effective it needs to be done on a time scale which is suited to the target failure mechanism, which is usually erosion or corrosion on interior surfaces. Even if such processes are initiated by a random event (such as sudden influx of particulates into a fluid stream), in many cases the resulting wear-out process may take several years to significantly erode a pressure boundary. So despite its purely diagnostic and non-intrusive characteristics this task might be better listed as a time-directed task.

2.2.3 Failure Finding

Failure finding tasks are scheduled tests to determine if a failure has already occurred. In a nuclear power plant they include essentially all the surveillance tests for standby equipment. However, it is a mistake to assume that these comprise all the necessary failure finding tasks. Any attempt to systematically design the maintenance program will undoubtedly discover the need for additional failure finding tasks. As discussed in a previous section, their primary goal is to control unavailability by preventing hidden failures from remaining in the system.

2.2.4 On-Condition

On-condition PM tasks are *unscheduled* tasks done *as a result of* poor as-found condition, discovered, for example, during vibration analysis or internal inspection. The on-condition task

is likely to be an intrusive activity. An example of an on-condition PM task is "replace motor bearings as a result of oil sample and vibration analysis."

In principle, these tasks are not periodic or scheduled in any sense, but in practice some of them may be former periodic tasks whose interval has been extended by knowledge of equipment condition. For example, mechanical refurbishment may have been deleted as a scheduled task for medium voltage motors, by building up condition-monitoring capabilities, but it would be performed if these tasks indicated it were necessary at a particular time.

In addition to such obvious examples, there are a large number of other activities, which include adjusting the packing on valves, adjusting drive belt tension, and repairing small leaks of air, oil, steam, or water whenever the need is observed during operator rounds. The list of maintenance work orders at most plants contains a majority of this type. Despite what the work orders are called at the plants, the work is not truly corrective maintenance provided no important equipment functions were lost.

Examples of on-condition PM tasks include:

- Replacing packing when leak rate is >1gal/hour
- Retorquing bonnet and flange bolts when process fluid is observed leaking
- Lubricating breaker operating mechanism when manual operation reveals binding
- Cleaning heat exchanger when ΔP reaches 20 psig
- Deferring overhaul to next outage because detailed inspection showed only slight degradation

On-condition tasks are the natural sequel to condition monitoring. However, they may result from many other ways of observing equipment condition—for example from internal inspection. Condition monitoring and on-condition tasks are sometimes collectively referred to as condition-directed PM activities.

2.3 Preventive Maintenance Objectives

Preventive maintenance is expensive because it has to be performed over and over again on thousands of components. In order to control the scope and costs of this activity and to maximize its effectiveness, maintenance managers need to apply more comprehensive PM tasks where they are really needed, and to minimize PM application for unimportant components. To get a clear view of how to do this, it is instructive to ask what the objectives should be in performing PM for different categories of equipment.

2.3.1 Prevent All Failures: The "Critical" Category

Critical equipment is equipment which must not fail at all. In other words, individual cases of equipment failure are intolerable. By this is meant equipment whose function is sufficiently important that plant managers are willing to expend significant PM resources to prevent these

functions from failing, even once. Of course, such failures will occur from time to time because no PM program can be perfect, but each failure carries a very significant penalty in safety performance, personnel injuries, or large economic losses. It is apparent that maintenance managers should do everything possible to avoid such failures. The PM program should be as comprehensive as possible. The PM objective is simply to go the extra mile to prevent all such failures.

In a nuclear plant environment, losses of high-level safety functions are clearly of extreme gravity. However, because designers have employed extensive functional and equipment redundancy there should be little danger that a single equipment failure could fail a high-level safety function. Despite continuing to be watchful for such single point failures, at the current mature stage of the industry there are very few safety SSCs which are truly critical in the sense described above. In contrast, the ability of single trains of *multiple train* safety systems to experience a failure with little impact on overall safety is enshrined in many Maintenance Rule performance criteria at most plants, which explicitly allow a single failure, or even more, in a relatively short period of one operating cycle.

This is not at all true of equipment vital to the generation of electricity. Loss of plant generation, even for a short time, has a very high immediate cost, but plant designs incorporate much less functional and equipment redundancy in the balance of plant than for safety systems. Single point failures of equipment important to generation are not uncommon, and are almost certain to be classed as critical.

Despite the fact that loss of one train of redundant safety equipment is not critical by the above definition, experience has taught that failure of even one train of a redundant safety system engenders significant penalties. Likewise, partial loss of generation which does not extend to a plant trip is nevertheless very costly. It is the prerogative of management at each plant to decide which consequences of failure are sufficiently burdensome to be labeled critical and to require comprehensive PM coverage. Relatively uncontroversial choices are that a critical failure should: 1) lead to loss of a high-level safety function; 2) cause partial (>50%) or complete plant outage; or 3) create a personnel hazard causing serious injury or death.

When redundant trains provide a high-level safety function, each train must be highly reliable so that if a failure on one train occurs, one of the other trains has a very high probability of providing the required functions. Clearly a high level of PM is likely to be needed to ensure that the reliability target is met. However, it has already been made clear that a single failure of any such train is indeed usually tolerable without severe consequences, providing high reliability of all the trains is achieved. For this purpose it is useful to define a category called "*significant*," rather than "critical." The "significant" category will be described in detail below.

Given the possibility of recognizing the "significant" category, the following list contains suggestions for criteria to be used to establish a "critical" PM category SSC. To be critical, the failure of the SSC must meet one or more of the following criteria:

1. Causing possible death or serious injury to plant personnel

2. Causing plant trip

3. Causing reduction in plant power generation or heat rate > 50% of maximum
4. Causing loss of the single train of a single-train safety function
5. Causing radiation release in excess of ODCM or other post-accident release limit
6. Resulting in greater than a 1-day delay in the outage or startup schedule
7. Causing entry into a technical specification 72-hour action statement
8. Being risk significant by the RAW criterion (see Section 2.3.5)
9. Being assigned as critical within an important plant program [such as GL 89-10 MOV, RG 1.97 (Accident Instrumentation), RG 1.155 (Station Blackout), EQ, ASME code requirement, Environment Effluent Monitoring, HELB/MELB, Risk-Based ISI/IST, Check Valve, Rosemount Oil Loss Monitoring, Appendix R (Fire Protection), Appendix J (Containment Leak Testing)] or being deemed critical to a management commitment

It is worth noting that some failure modes of a component may be critical and others noncritical. In that case, the component should be classed as critical, because PM activities are directed at the whole equipment, not at specific failure modes. It is also necessary to emphasize that the critical characteristics cover the domain of safety as well as production.

In particular, note that the following characteristics are *not* usually adequate for deciding if a component is critical for PM application: being safety-related, risk significant in general, being within the Maintenance Rule scope, having Maintenance Rule performance criteria or being standby equipment referenced in emergency operating procedures.

2.3.2 Prevent Most Failures: The "Significant" Category

As stated in the last section, there exists a category of equipment for which reliability must be maintained at a high level, but for which an individual failure can be tolerated, within expectations set by the reliability target. In other words, there must not be more than a single failure or so, in a certain time period, but a single failure is not itself remarkably destructive. In order to achieve a satisfactory level of reliability, most failure mechanisms of the SSC need to be addressed by PM tasks. The PM objective is not to prevent all failures, but it is certainly to prevent most of them.

In a nuclear power plant context, the PM program for the "significant" category will often be almost indistinguishable from the comprehensive program for a critical component. However, important differences are likely to emerge, such as the willingness to employ predictive maintenance techniques which have not been extensively used in the industry, or the ability to extend PM task intervals to find a more cost-effective PM program. For this category, achieving an appropriate balance between reliability and availability may also be less important than for critical components. Depending on the application, other industries may prefer to employ a different PM program for significant rather than critical components.

The following list contains suggestions for criteria to be used to establish a "significant" PM category SSC. To be significant, the failure of the SSC must meet one or more of the following criteria:

1. Causing reduction in plant power generation or heat rate > 10% but < 50% of maximum
2. Causing loss of one train of a multiple train high level safety function
3. Causing radiation release in excess of the plant administrative limit
4. Causing personnel radiation exposure in excess of the plant administrative limit
5. Causing delay in the outage or startup schedule greater than 1 shift but less than 1 day
6. Causing entry into a technical specification 7-day action statement
7. Causing major impact to the off-site environment
8. Causing unacceptable repair or replacement costs
9. Causing major impact on plant personnel resources
10. Causing failure of another critical or significant component
11. Being risk significant by the FV criterion or by designation by an expert panel (see Section 2.3.5)
12. Being deemed to be significant within an important plant program [such as GL 89-10 MOV, RG 1.97 (Accident Instrumentation), RG 1.155 (Station Blackout), EQ, ASME code requirement, Environment Effluent Monitoring, HELB/MELB, Risk-Based ISI/IST, Check Valve, Rosemount Oil Loss Monitoring, Appendix R (Fire Protection), Appendix J (Containment Leak Testing)] or being deemed significant to a management commitment

2.3.3 Prevent Some Failures: The "Minor" Category

Some SSCs create no large consequences upon failure, but nevertheless have some consequences which are worth avoiding if cost-effective PM tasks can be found to serve that purpose. Traditionally these have been called noncritical SSCs, and this terminology has been used in the EPRI PM Basis database [6]. To be cost-effective, it is usually necessary to narrowly restrict the PM tasks which are applied for this category of equipment, and to direct them at particular failure mechanisms or groups of failure mechanisms which seem to cause the more significant consequences, or which are the most likely to occur. The PM objective is thus to prevent only some of the failures.

The following list contains suggestions for criteria to be used to establish a "minor" PM category SSC. This terminology essentially coincides with existing use of the term "noncritical," but makes it clear that it does not include the "significant" category. "Minor" is also more

descriptive than "noncritical." To be minor, the failure of the SSC must meet one or more of the following criteria:

1. Causing possible injury to plant personnel requiring first aid
2. Causing reduction in plant power generation or heat rate <10% of maximum
3. Causing minor impact to the off-site environment
4. Causing entry into a technical specification 30-day action statement
5. Causing large repair or replacement costs
6. Causing minor impact on plant personnel resources
7. Having excessive corrective maintenance history
8. Having CM costs in excess of expected PM costs
9. Causing impairment of routine operational or maintenance activities
10. Causing reduced capability to perform equipment performance monitoring activities
11. Causing high personnel radiation exposure during repair, but less than the plant administrative limit

2.3.4 Prevent No Failures: The "Run-to-Failure" Category

Run-to-failure means that no PM tasks at all are performed. This option is chosen for equipment that does not fall into the "critical," "significant," or "minor" categories. Because no PM activities are being performed, there is no expectation that any failures will be prevented by PM. Consequently, the PM objective here is to prevent no failures. At a minimum, this category of equipment must be cheaper to repair when it fails than to perform PM on it. A typical example in this category would be replacing light bulbs only when they fail.

A question sometimes arises regarding equipment which is included in operator rounds, but which is not PM critical or significant. If the only reason for including it in operator rounds is, for example, that leaks and loose or missing parts can be spotted easily, this qualifies as satisfying element 8 (CM costs exceed expected PM costs) of the checklist for the "minor" category. Such equipment should be categorized as minor. Operator rounds may be the only PM activity, which is sufficiently cost-effective for such equipment. The employment of operator rounds legitimately prevents the equipment from being classed as run-to-failure. It is good to carefully reserve the "run-to-failure" category for SSCs which are truly run to failure. Equipment for which the only PM activity is to change a filter, to observe it during operator rounds, or to occasionally lubricate it is by definition not in the "run-to-failure" category.

2.3.5 PM Objectives and Risk Significance

The above checklists for the "critical" and "significant" PM categories contain items referring to the risk significance of SSCs. This provides a way to use the risk significance parameters, which are calculated in a PSA analysis when the SSCs in question are a part of the model, to assist in deciding on the PM category.

The Risk Achievement Worth (RAW) parameter measures the additional risk to core damage when the component is unavailable. It is defined to be the factor by which the core damage frequency (CDF) is increased during the period that the SSC's functions are unavailable. A threshold of RAW = 2 has become a de facto industry standard, above which the SSC is said to be risk-significant by the RAW measure.

If an SSC has a RAW value of 2, it means that the CDF is twice as large as the baseline value whenever the SSC is unavailable. This means that the failure or unavailability creates as much additional risk during the time until the SSC is returned to service as the entire risk of operating the baseline plant during the same period. This situation is interpreted as meaning that such failures are individually highly undesirable each time they occur, because they single-handedly double the plant risk for that period of time. Such an SSC merits the critical PM category.

In contrast, the Fussel-Vesely (FV) risk significance parameter indicates the *average* effect which failures or unavailabilities of the SSC have on the CDF over a long period of time. FV does not refer solely to what happens in an individual failure, but weights the increase in the CDF while an SSC is failed with the chance that such failures are going to occur, and it weights the consequences of other sources of unavailability with the frequency and duration of the unavailability events. The FV value is therefore sensitive to the failure rate, whereas the RAW value is completely independent of it. A threshold value of 0.005 for FV has also become a de facto industry standard, with values exceeding this threshold labeling the SSC as risk-significant by the FV measure.

In light of these definitions and interpretations it can be claimed that the SSC should be PM critical if the RAW value is 2 or greater, regardless of the value of FV, because the effect on the CDF of individual failures is very high. If an SSC is risk-significant by FV but not by RAW it means that there is a need to control reliability, because reliability may be quite poor (basic event probability must be >0.005 or whatever the FV threshold is), and it drives the FV value, even though there is no need for a comprehensive level of PM to prevent individual failures.

These considerations have led to the quadrant chart shown in Figure 2-2, which provides guidance on using risk significance to determine PM objectives and category. Because PM is performed on whole components, the FV value used for a component should be that for the component as a whole, obtained by adding the FV values for all the reliability basic events for the component. Select a value for RAW by selecting the largest RAW from all the basic events for the SSC.

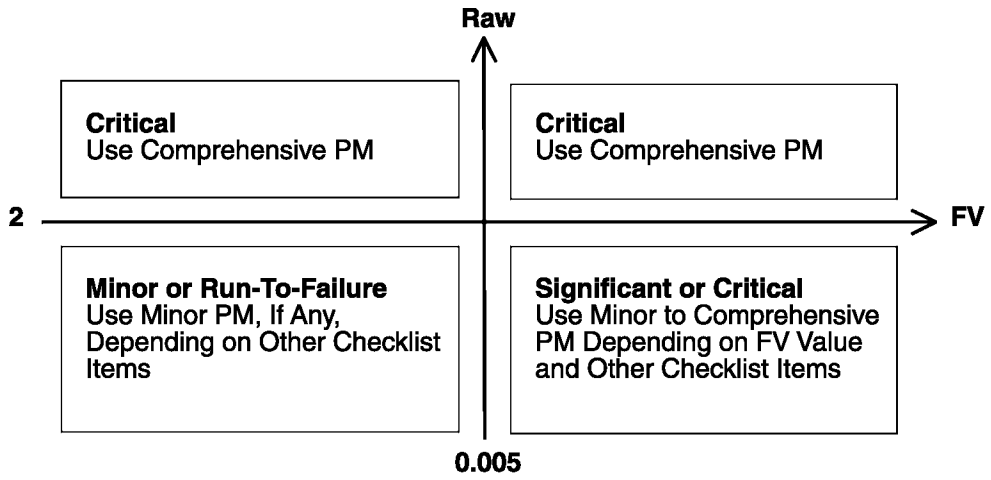

Figure 2-2
PM Category Depending on FV and RAW Values for the Whole Component (All Failure Modes)

2.4 Classical Reliability Centered Maintenance

The intent of plant maintenance optimization is to allocate limited plant resources to where they will provide the greatest return. There are several methods that have been employed to achieve this objective. The most comprehensive is the performance of a reliability centered maintenance (RCM) analysis. This approach was developed initially for application to commercial aviation [7]. This process currently is used extensively in military (DoD) and space (NASA) applications. The methodology is specified in various documents, including MIL-STD 2173(AS) [8] for applications to defense systems, and SAE Standard JA1011 [9] for commercial applications. These standards prescribe that for any process to be called RCM, the following seven steps must all be completed in the prescribed order:

1. Definition of the functions and acceptable performance standards for the asset within its current operating context

2. Definition of the functional failures for the asset

3. Identification of the appropriate failure modes

4. Identification of the corresponding failure effects

5. Identification of the resulting failure consequences

6. Specification of appropriate proactive tasks and intervals to detect or prevent failure

7. Specification of appropriate “default actions” (that is, alternative mitigation/elimination strategies) if no appropriate proactive task is possible

Notably absent from the above list is any mention of working through the list of components in the plant to figure out how they should be maintained. There are many examples of how this can lead to maintaining the wrong equipment, and over-maintaining the rest. Instead, RCM focuses on the functions that the plant must provide, and the ways in which these functions may fail to be provided. In practice, the analyst must find his or her own way to address the plant functional failures by considering intermediate levels of functions and functional failures. *The functional levels chosen are theoretically arbitrary and are not prescribed by RCM*, but are constrained by the level at which information is available, and at which designers, operators, and maintainers are accustomed to think about the plant and how it achieves its functions. Therefore, the plant assets, systems, and subsystems are usually the levels, which mediate attempts to characterize functions and functional failures. In principle and in practice, RCM requires the equipment to be encountered only when answering questions about the modes of failure of the functions of the assets, systems, and subsystems. Failure modes of equipment are introduced in the course of RCM analysis only in a *deductive* process, which inquires how each higher level functional failure can be caused to occur.

The SAE standard has been submitted to the American National Standards Institute (ANSI) for recognition as an American National Standard; however, compliance with it is voluntary. Its use is widespread throughout the automotive industry. Other process industries, including petrochemicals and metals production, have also employed it to a significant extent. However, this standard only pertains if the optimization process used is to be considered an RCM process. Other maintenance improvement/optimization approaches (to achieve particular business objectives) are outside the scope of the standard as long as they do not claim to be RCM processes.

This RCM approach, commonly referred to as classical RCM, uses a failure modes, effects, and criticality analysis (FMECA) to determine the plant equipment and failure causes which are important to each system functional failure. In many applications of RCM throughout the electric utility industry (and other process industries such as petrochemicals, gas transmission, and manufacturing), a more simplified failure modes and effects analysis (FMEA) is used. As mentioned above, in principle the FMECA deductively identifies an equipment failure mechanism because it is an important way in which the system function may be caused to fail. The deductive question is always, “How else could the failure of this function be caused?”

This is made very clear by Mowbray [10], in a process he calls RCM II. In practice, the deductive approach is compromised because in the context of hundreds of potentially contributing components, some contributions are likely to be missed unless the equipment list is systematically incorporated into the process. Mowbray provides many insights and suggestions but is ultimately silent on the procedural resolution of this problem, leaving it up to the experience of the analyst. However, a practical solution is offered by Smith [11], albeit one which departs from the deductive line of reasoning. In Smith’s process, the equipment list is explicitly introduced in a screening step, which systematically eliminates each item of equipment from the FMECA if there is no chance (conservatively judged) that its failure could lead to any of the system functional failures. From a practical point of view this is a good step to take, but it involves working through the equipment list asking *inductively* of each component, “If this fails what happens?” In Smith’s procedure, only equipment which passes the screen is subjected to

the FMECA. The two main references on classical forms of RCM discussed here (Smith and Mowbray) are both excellent sources of information on classical RCM and should be studied in detail by anyone embarking on an RCM project.

This deductive/inductive issue is highlighted here, not because it represents a particular weakness in the RCM method but because it is the source of an important procedural arbitrariness not resolved by the RCM standards. Furthermore, RCM purports to hold the line on a strictly deductive approach driven from system functional failures, and eschews interrogation of the equipment list as a throwback to older, less-intelligent methods of maintenance program development. Despite the philosophical posture of RCM, the author is inclined to believe that the equipment list must be used to provide closure—that is, some kind of assurance that the deductive process has considered all the potential contributors. Smith does this explicitly; Mowbray seems to do it informally.

In RCM, the "effects and criticality" part of the FMECA is conventional, asking for local, system, and plant level effects to be documented, but RCM does not close the loop to formally make use of the system functional failures which initiated the FMECA and which, in fact, control the organization of the work. The final step of criticality assignment takes place in a decision tree or logic tree analysis which considers whether each failure mode can be noticed by the operators, and whether it is of safety or significant production concern.

Once the functional importance of a particular piece of equipment is determined, classical RCM provides rough guidelines (more detailed in [10]) leading to the kinds of applicable maintenance activities described in Section 2.2. However, RCM itself does not contain any information on the specific PM tasks and intervals to be employed. In other approaches in process industry applications, component maintenance templates have been employed to specify these activities, based on the component's functional importance, service environment, and duty cycle. To the degree that the templates (as in [5, 6]) provide *candidate* tasks, they can be used in classical RCM without violating RCM principles. The templates, however, should not be used uncritically.

Finally, the classical RCM standards require specification of a living RCM program. This program requires a periodic review for both the information used to support decisions made and the decisions themselves. This task is one which is very important, to achieve continuous improvement and permit the organization to remain competitive in a changing business climate. However, it also is the one which is most often neglected. There are many instances in which organizations undertake maintenance improvement efforts, only to have the benefits erode over time due to lack of follow-up. Thus, regardless of the method chosen to achieve the business objectives of the organization (classical or streamlined RCM), it is of critical importance to develop and implement processes to implement this living program, and these processes should be made functional at an early stage of the implementation process.

The major drawback to use of a classical RCM approach is that experience has found it is very labor-intensive and expensive to perform. This is an acceptable limitation when the results can be applied to many identical assets. As an example, in commercial aviation, the maintenance program can be specified for an entire aircraft model (for example, Boeing 777) and be applied

to all aircraft of that type which are produced, thus spreading the costs over a large number of applications. Another example is its use for a production line asset such as a machine tool of which there may exist many essentially identical copies among a manufacturing company's production facilities.

However, since process facilities have many features which make their design and operation unique, these economies of scale are often not present. Thus, for application to many industrial facilities, particularly those in competitive industries, application of classical RCM has been found to be too expensive to justify its use across the board. In the electric power industry, this experience was confirmed in initial applications of classical RCM to operating nuclear power plants. One solution to the demands on analysis resources is to apply RCM only to the minority of systems or assets which present the majority of maintenance costs and problems—the so-called 80/20 approach, described in [11] and to some degree in [10]. This may be made to work rather well, but by definition it leaves maintenance on the majority of the plant's equipment to be optimized by some other method. There is also a risk that significant resources could be diverted to identify the minority of systems that would benefit, which may not be a simple task.

In response to these limitations, so-called "streamlined RCM" (SCRM) processes have been developed. A number of these approaches were developed specifically as a consequence of the lessons learned from application of classical RCM to nuclear power plants, several of which are described in [12]. Since development of these streamlined approaches, they have found widespread acceptance throughout the electric utility industry, including application to nuclear, coal, gas, and hydroelectric generating stations. Additionally, they have achieved success in applications to other process industries including petrochemicals, pharmaceuticals, gas transmission, rail transportation, and manufacturing. In this section, the two most widely applied alternatives to classical RCM will be discussed and their advantages and potential disadvantages outlined. Despite the use of RCM in their names, these alternative approaches should not claim to be RCM, especially now that that epithet is being enshrined in industry standards. They do, however, lay claim to many of the benefits of RCM, and avoid some of RCM's disadvantages, although they encounter some difficulties of their own. The rest of this section is intended to throw some light on these comparative advantages and disadvantages.

2.5 Critique of RCM and Streamlined Methods

The first approach which has found widespread application throughout the electric utility industry can be classified as *streamlined RCM* (one version of which is described in [13]). The streamlined RCM approach attempts to follow the intent of classical RCM, but combines or eliminates steps perceived to provide limited added value. Depending on the shortcuts used, the method can be close to or far from the classical approach. For this reason the streamlined, or abbreviated, forms of RCM will be discussed not by referencing a particular method, but by focusing on the issues raised by a range of possible shortcuts. Because of the close relation between the steps of classical and streamlined RCM, these two methods will be compared first.

The second approach used as an alternative to RCM corresponds to what has been called a *criticality checklist approach*. In a simple checklist method, the functional analysis and the FMEA are replaced by use of a single list of questions designed to address the connection

between equipment failures and plant functional failures expressed at the highest level. Obviously, in such an approach there is little methodology which resembles RCM, but it has been found to be quite effective in specific situations.

2.5.1 Functions and Functional Failures

The first step in the streamlining process is applied to the development of system functions. In classical RCM, all system functions are identified and analyzed. Additionally, the referenced classical RCM standards ([8, 9]) require that all function statements be *quantified* in every case where possible. As an example, for the reactor core isolation cooling (RCIC) system of a boiling water reactor, one function would be "provide process flow from the condensate storage tank to the reactor pressure vessel (RPV) at a flow rate of 600 gpm at 1050 psig pressure to provide RPV makeup during RPV isolation conditions." In addition, there would be a second function to provide this makeup capability from the suppression pool. In classical RCM, each of these functions are then analyzed separately.

In streamlined forms of RCM, the functions are not necessarily quantified, and qualitative modifiers may be used in the function description. Additionally, similar functions could be grouped together. A streamlined version of the previous example could be "provide RPV makeup at sufficient pressure and flow during isolation conditions." The classical RCM approach is much more explicit. One potential drawback to being less explicit about the functions is that the functions may need to be modified (that is, expanded upon) if they are to be used to support other applications. For example, use in Maintenance Rule applications typically requires a greater degree of specification than is provided in some streamlined methods. As a particular example, quantitative criteria are typically required when evaluating whether a Maintenance Rule functional failure has occurred [14].

Additionally, because many of the components are identical between the two systems (the only difference in the example provided is the suction isolation valves and their associated control instrumentation and logic), in the streamlined approach, they may only be analyzed once, versus once in each system in classical RCM. This provides significant savings in cost and labor over the classical approach. This is particularly true for instrumentation and control devices which typically support multiple system functions and can constitute a large fraction of the total plant equipment which requires analysis.

We find these issues discussed in Mowbray's book, but not resolved. In his equipment screening step (mentioned above), Smith provides an analogue to combining functions, by referencing one functional failure for a specific component whenever it is judged that the FMECA analysis to be performed for the referenced functional failure would envelop that for the functional failure from which the reference is made. In other words, if pump A can contribute to functional failures 2, 4, and 5, but the FMECAs for 2 and 5 would be copies or subsets of the FMECA for 4, then his procedure carries out the FMECA only for functional failure 4, and references it under 2 and 5. Of course, this does not combine functions or functional failures, because they have already been explicitly defined, but it does combine what would otherwise be the resulting FMECAs for different functional failures and greatly reduces the amount of analysis. However, if the RCIC and Suppression Pool systems were analyzed as separate systems, even Smith's procedure would

analyze these functional failures twice. One could claim that analysts need to keep aware of the potential to avoid duplicating work already performed when it is appropriate to do so. However, in this case the work has not already been performed, because the inductive use of the component list and referencing across functional failures takes place before any FMECA is attempted.

Finally, if the suppression pool and RCIC systems were analyzed together as one system (for example, RPV makeup), the combined generalized function could be the correct one to use even in classical RCM, providing it is explicitly quantified. Since the system level is an arbitrary choice (even within existing systems some are much larger and more complex than others) and is not prescribed by RCM, the question of the right way to handle functions dissolves into matters of feasibility, economy, style, and convenience rather than revealing fundamental differences in methods. The essential point is that at whatever level they are addressed, the functions and functional failures must be defined and used to focus attention on the right equipment failures. Mowbray suggests this should be done at the highest level at which it is technically feasible.

2.5.2 Component Failure Modes

The next step to which the streamlining process is applied is the FMECA/FMEA. In classical RCM, all functions, regardless of their importance, are analyzed to specify their functional failures and the individual failure modes which lead to them. In streamlined processes, functions identified as unimportant may not be analyzed in depth, although equipment which supports only unimportant functions is reviewed for economic impact in a run-to-failure determination. A screen on the importance of functions or functional failures seems to be an obvious way to focus analysis resources. Components which support only the unimportant functions must still be subjected to a run-to-failure screen to identify those which should not be run to failure.

Furthermore, in classical RCM, each failure mode for components which support the function are analyzed, regardless of the relative likelihood of expected occurrence, with plant level, system level, and local level effects specified for each failure mode. This detailed level of analysis is one of the primary contributors to the large costs associated with performing classical RCM. In SRCM processes, only those failure mechanisms identified as being dominant are analyzed. This focus on only the most likely failure modes has resulted in significant reduction in analysis time and cost. This focus also has resulted in a much greater degree of acceptance of the recommendations by plant staff, because they are limited to the set of failures for which they have reasonable expectations they can experience.

Both RCM and SRCM rely heavily on analyst experience and judgment to identify those component failure modes which are worth entering into the FMEA. A streamlined method such as [17] limits the number to those which are truly dominant. RCM is somewhat vague on this topic but expects the analyst to include all plausible failure modes irrespective of their perceived probability of occurrence. The operating context usually provides good clues as to what is plausible, and the level of consequences can act as a control on the level of detail required. A procedural difficulty of using the level of consequences in this way, as suggested by Mowbray for classical RCM, is that at this point in the FMEA the level of consequences has not formally been addressed (see next paragraph). In any case, limiting the FMEA failure modes will result in

the analysis being less complete and more vulnerable to the omission of a significant failure mode. Additionally, most streamlined forms of RCM combine all of the dominant failure modes and their effects in a single record in the implementation software. This also results in a reduction in the specificity of the analysis and, again, limits its utility to support other applications. The level of detail in the FMEA appears to be the most significant difference between streamlined and classical RCM identified to this point.

One further note is required, on the analyst's ability to control the level of detail in the FMEA. In RCM the FMEA is carried out to the level of component failure modes and causes before subjecting them to the questions about the effects of the failure on the system and plant. If these effects are minor or insignificant, the effort expended in developing the causes will have been largely wasted. A significant improvement in RCM would be to insert the effects part of the process before developing the failure causes. The same comment is probably true of streamlined methods.

An additional difference in the conduct of the functional importance evaluation is that at least the classical RCM *standards* specifically require consideration of events and processes that would likely result in a functional failure, including causes due to human error, unless human error is actively addressed by a separate process. It is not clear how far RCM or the RCM standards intend human errors to be pursued, or what constitutes a separate process.

Smith requires human causes of component failures to be entered into the FMEA but further analysis of them to be curtailed, on the grounds that PM tasks cannot address them. He recommends further examination of any significant human error issues among "Items Of Interest." This is a list of related issues compiled during the analysis, often focusing on design modifications—items which may be outside the strict RCM analysis boundary. Mowbray also requires human error causes to be entered into the FMEA but suggests they be considered under categories of anthropometric, human sensory, physiological, and psychological factors. He admits that these topics are vast in themselves and require expert consideration outside the boundary of PM task assignments. RCM therefore "considers" human error, but not at all on an equal footing with its treatment of other failure causes. In streamlined methods applied to nuclear plants, human error is almost universally excluded from the analysis. In part this is due to the fact that all nuclear power plants (and most plants using fossil fuel sources) do have extensive programs and processes which separately address human error, as referenced in the RCM standard. Even so, these programs are neither referenced by nor integrated into the streamlined effort.

A further point concerns the level at which equipment is analyzed in the FMEA. Plant instrumentation and control devices can constitute a large proportion of plant equipment. These devices typically also support multiple functions, many of which cross system boundaries. Thus, significant effort has been expended in reducing the analysis required for these components. The most often employed strategy (and the one used in most streamlined methods) to address this issue is to analyze plant instrumentation in loops. In this approach, the entire group of instruments (for example, primary process sensor, electronic transmitter, signal conditioning devices, bistable trip units, and local and remote indication and recorders) is analyzed as a single "supercomponent," with all maintenance activities referenced to the primary designated

component (typically the transmitter). This approach greatly reduces the amount of analysis required and maintenance work orders generated. It also eases the burden of coordinating work in the planning and scheduling process. The reduction in cost is achieved at the expense of reducing the specificity of the analysis. As an example, failure of a recorder in an instrument loop may have no significant effect on plant operations; however, other devices in the loop may provide critical functions. Since the functional importance is assigned to the loop as a single entity, the recorder is assigned the functional importance classification of the loop, and no distinction is made in the importance of the devices within the loop.

Treating a group of instruments as a loop does not appear to violate any fundamental requirements of RCM, providing all the ways in which the group can contribute to functional failures are recorded. One could argue that treating the loop as an entity focuses more attention on the integrated effect of the instruments on such functional failures, or alternatively, that the loop treatment encourages the omission of such contributions. The trade-off seems to be more a matter of style and experience of the analyst. However, treating a group of components as a block for the FMEA analysis is the norm in RCM, because otherwise a pump or motor with lubricating and cooling components would have to be broken down to its hardware, making any RCM analysis intractable.

2.5.3 Logic or Decision Tree Analysis

Classical RCM requires the explicit identification of hidden failure modes. Additionally, it requires the consequence categorization to clearly distinguish those failures which result in safety or environmental consequences from those which result in only economic consequences. This is accomplished by adherence to a simple series of questions, which despite its name does not constitute a logic tree. Nevertheless the RCM process does properly label each failure mode as hidden or not, a matter of safety or operational significance, and if operational, establishes a 2-bin scale of significance. The latter embodies the criticality assignment as critical or noncritical; safety issues are deemed critical. There is little that is wrong with this important—indeed, essential—process, but it is worth noting that its structure is trivial (meaning it is easy to remember and apply routinely) and it is inefficient to employ in the logic tree format.

First, it has already been pointed out that the functional failures themselves could have been binned into those which have critical or noncritical impact (perhaps not a final designation but at least as a screening tool). Second, the last part of the FMEA has already *explicitly asked and documented* the local, system, and *plant level effects* of each failure mode. The plant effects cannot avoid stipulating a safety or operational type of consequence. Third, at least in a nuclear power environment, there is a sharp distinction drawn between safety and non-safety issues, almost always highly apparent at many earlier points in the analysis. For nuclear plants, hidden failures have already been researched extensively to establish surveillance tasks prominent in the plant's operating license. Other hidden failure modes may indeed be added by RCM analysis, but they will be few in relation to the vast burden of surveillance tasks in a nuclear plant. As a consequence, the logic tree analysis has always appeared to be an impediment to rapid analysis, and redundant to at least one prior step.

Streamlined RCM as practiced in nuclear power plants typically addresses the issues of hidden failures, safety versus operational consequences, and a 2-bin criticality scale, but establishes these characterizations with less formality than the logic tree structure. This is greatly aided by the design, licensing, and culture of a nuclear power plant, and may not be quite as valid in an industrial facility. Even then, it should be possible to assign and codify the appropriate criticality characteristics during the documentation of the plant level effects, thus subsuming the inefficient logic tree analysis into the final stage of the FMECA. This would not violate any vital RCM fundamentals.

One additional point about critical failures in RCM is that the use of a 2-bin scale is very crude. Section 2.3 has described a slightly more effective criticality scheme which distinguishes an additional level of criticality. The "critical" level and the new "significant" level are aligned with the differing objectives of the PM program for single point critical and critical-in-redundant-combinations type of failures. They may be especially useful in non-nuclear plant applications.

2.5.4 Task Selection

The next significant difference between the classical and streamlined forms of RCM is in the method of task selection. Classical RCM uses a further series of formal questions to guide the selection of appropriate maintenance tasks. Once again, the series of questions is important but trivial in nature, and it can be claimed that its formal documentation is an exercise in pedantry.

RCM does not clearly link failure cause to the timing and age characteristics of the failure mechanism. In Section 2.2, time-directed PM tasks were described as very ineffective against random failures. Only at the level of cause can a failure mechanism be distinguished as random or wear-out. For example, an electric motor may have anti-friction bearings which wear out over a period of 2 to 6 years when the oil is allowed to degrade because it is not replaced often enough. However, if the bearing wears as a result of misalignment, its failure could be rapid and random. Thus "worn bearing" is not sufficient to establish the timing characteristics, and hence not adequate for the purpose of task selection. The best control on the level of cause detail to include in the FMEA is to continue to the point where the timing behavior can be stated. At a minimum this must include identification of random or wear-out failure mechanisms.

RCM is not very explicit on this question. In the task selection step, Smith asks if the age relation is known, instead of whether there is an age relation at all, and prefers to select a time-directed task if one is applicable and effective. Mowbray asks correct questions about the suitability of different kinds of tasks but assumes that an applicable and effective condition monitoring task should be preferred. In fact, it was the lack of clear direction on task selection from classical RCM which was the primary motivation for developing the EPRI PM Basis database. Classical RCM is therefore not clear on which kinds of tasks are preferred, and provides no information on selecting the actual tasks and intervals. Mowbray has begun to address this issue by including a useful appendix on condition-monitoring techniques.

Especially in light of its lack of specific technical information, the formal RCM "task selection" approach has been found to be time-consuming and labor-intensive out of proportion to its

benefits. It could be claimed that classical RCM does not include a task selection step at all, but does what it can to prepare the analyst for these essential decisions.

To address this issue, standard maintenance plans (that is, sets of PM tasks and recommended intervals, more commonly known as PM templates) have been developed for the major component types commonly found in process applications. Some, but not all, streamlined RCM approaches use them. See [5, 6] for examples developed by EPRI for application to components found in nuclear power plants. These templates suggest standard maintenance plans for different component types based on the equipment's functional importance, operating environment, and duty cycle. A major premise of the templates is that for each component type, there is only a finite number of tasks which are applicable to detect equipment degradation and restore performance. The tasks and intervals in the templates are recommendations for *candidate* tasks and intervals in the task selection step. The selection of tasks is further tailored by the analyst in the best streamlined approaches.

Additionally, there are operational constraints which determine the periodicity at which these activities can be performed (for example, fuel cycle length, or technical specification specified testing intervals). Thus, as a practical matter, much of the RCM task selection step provides little tangible benefit. For application to nuclear power plants, the templates (and supporting basis documentation provided in [5, 6]) also provide the following advantages, some of which are not present in a classical RCM.

First, the templates were constructed using the hierarchical questions contained within the RCM task selection step; thus this hierarchy is embedded within them. Second, the templates developed by EPRI were constructed via a structured process with experts from many different nuclear power plants. Additionally, representatives of various equipment manufacturers participated in the development of the templates. This diversified level of experience has now been fortified by incorporating reviews by expert industry groups to incorporate the majority of PM experience available in the industry. Typically, this level of operational experience is not available when applying classical RCM at a single site.

Third, each template is supported by an FMEA developed by these experts to a level of cause that supports the random/wear-out characterization of failure mechanisms, and to a completeness in failure mechanisms which is more than an order of magnitude more comprehensive than encountered in any RCM analysis (for example, ~100 mechanisms rather than 3 to 10). Computer tools which examine the level of protection afforded by each PM task, and the combined effect of any subsets of PM tasks, including the effect of changing task intervals, provide a much enhanced capability to judge the technical value of each task against the full range of failure mechanisms.

However, there are several minor drawbacks to the use of maintenance templates for specification of the maintenance program. First, the templates are still a generic tool whose implementation relies on the analyst's knowledge of the RCM approach to task selection to ensure that the dominant failure mechanisms are addressed in a cost-effective manner for the less critical components. Second, although the task (and implementation frequency) recommendations are derived from considerations of the various failure mechanisms, including

their time propagation characteristics and the effectiveness of the tasks to address them, the templates, *when used alone*, do not provide a direct, explicit link between the equipment failure mechanisms and the maintenance tasks selected. Third, because the templates are intended to provide guidance—that is, candidate tasks and intervals—it is imperative that the analyst document the basis for any deviations from the recommendations. However, other EPRI PM database tools can be used in conjunction with the templates to address these issues.

In classical RCM, each failure mechanism is analyzed for each plant component, and the hierarchy of task types is reviewed in order to determine the appropriate maintenance strategy in each case. The inclusion of only a subset (typically 3 to 10) of failure mechanisms for a component in the FMEA, instead of the complete set (which would make RCM utterly intractable) is not as serious a deficiency as it might seem, because the majority of the sensored failure mechanisms may nevertheless be protected by the PM tasks which are selected to address the 3 to 10 dominant failure mechanisms. Although this may often be the case, it cannot be assured. In the future it is planned to examine this completeness issue for major component types, by selecting only the most common failure mechanisms in the EPRI PM Basis FMEA tables and determining the degree to which the PM tasks and intervals that would have been selected to address only these subsets actually provide adequate PM protection for the complete set of failure mechanisms.

If classical RCM is performed in conjunction with a recognized engineering standard, it provides a defensible legal basis for the prudence of the plant maintenance program if failure of plant equipment should occur and result in significant consequences. Carefully executed streamlined methods of RCM may also be capable of providing this basis; however, the standard of proof will be more difficult and would probably require additional supporting evidence from the plant operator.

2.5.5 The Criticality Checklist Method

The second approach often utilized to specify an applicable maintenance strategy is to replace the functional analysis and the FMEA with a checklist of conditions which characterize the important high-level functional failures, and by inductive inference, the equipment's functional importance. Such checklist items were described in detail in Section 2.3. The premise behind this approach is that functional importance is primarily determined by conditions which plant management determines are unacceptable. Examples include failures which result in a plant trip, significant power reduction, environmental release, or significant personnel safety hazard. The approach is to define this categorization and then classify plant equipment against the criteria developed.

This approach permits a very rapid assessment of equipment functional importance. Also, because the criteria for the various component classifications are uniquely specified and approved by plant management, this approach can result in more uniform decision making then either classical RCM or its streamlined forms previously discussed. Because the approach uses simple specific questions which can be answered by knowledgeable plant personnel, training requirements are minimal (that is, there is no need to understand how to perform or interpret an FMEA). Finally, because the functional importance criteria are explicit, it is a very simple matter

to query the database in which the analysis is stored, to obtain specific information when desired. However, because the checklist process eliminates the FMEA (which is a generic process) and substitutes the checklist questions, there are some tradeoffs involved with its application.

The first tradeoff is that, in the checklist process, there no longer are any explicit ties between the criticality classification and system or plant level functions. The connections clearly exist for the plant level functional failures, except that the unimportant ones are simply not delineated. This is not a catastrophic limitation—Sections 2.4 and 2.5.1 have already established that the level at which functions are defined is arbitrary, and the checklist method explicitly addresses the most important plant level functional failures.

Second, because no FMEA is performed, the only effects analyzed are typically these very plant level effects which have been identified on the checklist. In some instances of application of this technique, there is provision made in the analysis software to provide textual remarks that offer additional information on intermediate, and possibly additional, effects—for example, at the system level. However, this is completely a function of the analyst's experience and knowledge. Because the checklist approach often is chosen specifically to reduce analysis time, there also are strong pressures and incentives not to provide this level of analysis depth.

Third, because the checklist process can be performed very rapidly, the quality of results obtained depends very heavily on the experience and knowledge of the analyst. Most checklist approaches accomplish task specification via use of maintenance templates, which at first sight appears to compound the shortcomings of not performing an FMEA. However, the task recommendations in the templates have been designed to address a comprehensive set of failure mechanisms, so the failure of the checklist method to document an FMEA is significantly mitigated by the use of templates.

Finally, failures which result in safety or environmental consequences for which no applicable maintenance activity is possible require special treatment in RCM. When this condition occurs, the classical RCM standards require that a one-time change (that is, design change or operating strategy modification) must be performed to reduce the level of failure probability to a level considered tolerable by the asset owner. However, in application of the streamlined and checklist approaches discussed here, these decisions are made on an ad hoc basis with no such requirement imposed. Because of the potential consequences which attend issues of this type, management should provide guidance (as a policy statement or some other form of control) on how these issues should be resolved when they occur.

2.5.6 Conclusions on RCM Methods

Classical RCM and other approximate methods are management decision processes that optimize use of plant resources in the performance of preventive maintenance, while assuring appropriate levels of equipment reliability and availability. In the critique sections above it has been established that although there are fewer areas than might appear at first sight where the streamlined RCM methods violate basic precepts of RCM, the streamlined results will be less comprehensive than classical RCM. This is simply an expression of the maxim that what you get is what you pay for. Despite its pedigree, there is some justification for the view that with

classical RCM it is possible to pay more than you should for what you get and that it takes more time than you have. The checklist method can lay no claims to being RCM of any kind, but it is nevertheless a very useful tool for nuclear power plants which need to make rapid progress in improving their PM programs. Ultimately all the approximate methods depend heavily on the skill and experience of the analysts. This dependence is particularly strong when the approximate methods are used.

It has often been found that the best way to proceed at a facility requiring a large PM optimization program is to first analyze a few of the most important systems using classical RCM. After that, a streamlined approach, carefully optimized considering all the caveats above in Sections 2.4 and 2.5, should be selected for the remaining important systems. Less-important systems could be analyzed using the checklist approach. This approach gets everyone in the program to appreciate the rigor of the classical RCM approach, but solves the problem of performing a credible analysis of the facility within reasonable resource limitations. In essence, such an approach roughly corresponds to the use of the 80/20 rule at every level in the analysis, not just to identify the systems to be subjected to classical RCM.

3
THE DEPENDENCE OF FAILURE RATE ON TIME AND PREVENTIVE MAINTENANCE

3.1 Failures as a Function of Time

In Section 1 it was pointed out that although equipment failures appear to be very random in their timing there are good reasons why some individual failure mechanisms are anything but random in their overall pattern of occurrence, there being a period of time when no failures are expected from a wear-out mechanism. After that time has passed, wear-out failures become more and more likely. This section shows how such a failure time distribution translates into a failure rate as a function of time. In some circumstances, such failure time distributions can be revealed by Weibull analysis of the times to failure. In principle, failure time distributions can be combined mathematically with the performance of PM tasks whose effect is to reset the time axis for the equipment, so that wear-out mechanisms essentially "start again" from the beginning of their expected failure-free periods. The opportunities for the development of complex statistical models to describe this behavior are extensive, and the reliability literature contains many examples of so-called alternating renewal processes.

In this section we will avoid all such treatments, preferring instead to point out many practical reasons why a much more simple approach to modeling the effect of PM on reliability is likely to be more successful. The approach described will correspond to that used in the EPRI PM Basis database [5]. In concise terms, 1) the simple approach can use most of the practical knowledge the industry has accumulated on what makes PM effective, and 2) the most significant omissions from the simple model, which are the failure time distributions, can be shown to play usually only a minor role in determining the failure rate as a function of time. This apparently rather extravagant claim will be substantiated in the following sections, in which first the dependence of the failure rate on time is discussed, then its dependence on preventive maintenance.

3.1.1 The Time to Failure Distribution and the Failure Rate

The failure time distribution, f(t), is a probability density which describes the occurrence of a particular failure mechanism over time, among a group of N components when the failed components are not replaced. For a wear-out type of failure mechanism there will be few failures initially, increasing numbers of failures with time, and fewer failures eventually, because most of the group will have already failed. The distribution f(t) will resemble Figure 3-1.

Figure 3-1
The Failure Time Distribution

The quantity f(t) is defined with the understanding that at a time t some of the components have failed and others have not. This means that f(t) is not conditional on any assumption about the failed or unfailed state of the components.

t

The cumulative function, $F(t) = \int f(t')dt'$, is the total fraction of the N components that

0

have failed up to time t, which equals the *probability* that at least one component will have failed by time t.

The failure rate as a function of time, $\lambda(t)$, with which we are familiar from previous sections, is defined so that the probability of a failure in t to t + dt is $\lambda(t)dt$, but *this assumes that the component is in the working condition at t*. This makes the failure probability, λdt, a conditional probability—that is, conditional on there existing an operational component at time t, which is capable of failing.

Because the probability of an event equals the probability of the event given a condition, multiplied by the probability of the condition:

Unconditional probability of failing in t to t + dt = Probability of failure in t to t + dt, given the component is operational at t x Probability that the component is operational at t

Or:

$$f(t)dt = \lambda(t)dt \; x \; (1 - F(t) \qquad \textbf{Eq. 3-1}$$

Or:

$$\lambda(t) = f(t) / [1 - F(t)] \qquad \textbf{Eq. 3-2}$$

$\lambda(t)$ is the analogue of f(t) in a process where the failed components are replaced as soon as they fail.

Clearly, a given wear-out failure mechanism always contributes a time-dependent component of the failure rate. We might expect that failure rates would therefore always be time dependent, because they consist of an aggregate of such time-dependent functions. The success of reliability modeling and probabilistic safety assessment depends largely on the fact that, in practice, this assertion is generally false. Failure rates are assumed to be constant in time for two main reasons: 1) they are, in fact, more or less constant in time; and 2) it is very difficult to determine any residual deviations from constancy.

3.1.2 Nine Reasons Why Failure Rates Are Constant in Time

The main reason why failure rates are indeed constant in time is that a complex piece of repairable equipment will be made up of many subcomponents, each of which has a number of failure mechanisms, many of which can be influenced by a number of different stressors which determine the rate of progression to failure. Many will be wear-out mechanisms with widely different times to the "rising part of the bathtub curve" of failure rate versus age. (A bathtub curve was shown in Figure 1-2A.) In most equipment the time scales for wear-out range from the very short (a few months) to the very long (40 years or more), and cover the whole range in between. As a result, the complex piece of equipment never may never completely wear out, and exhibits an approximately constant failure rate, as displayed by the flat parts of almost all the curves in Figure 1-2.

Weibull analysis to determine time dependence from a sample of times to failure is therefore best applied to a single failure mechanism that can display a clear wear-out effect, rather than to the whole failure rate of the equipment. Most texts on Weibull analysis (for example, [15]) point out that as soon as five or more different wear-out mechanisms contribute to the data, the failure rate tends to take on the appearance of a more or less constant failure rate. In practice, there can be 20, 30, or even more different wear-out mechanisms contributing to the overall failure rate of a complex SSC.

The second reason why the failure rate is constant is that most of the competing wear-out mechanisms are conditional on a random initiating event, as described in Section 1.7. This randomizes the start of the competing wear-out mechanisms in a way which continuously varies throughout the life of the component. These randomizing influences are not only plant-specific but also equipment-specific (not equipment type-specific, but actually dependent on individual ID tags).

The third reason failure rates are constant is that we usually need the failure rates for critical equipment. Critical equipment is subject to high levels of preventive maintenance, which systematically and with great success prevents the failure mechanisms from entering the bulk of the failure time distribution. To be effective, a time-directed PM task must be performed before or at least close to the start of the failure time distribution of a failure mechanism which it targets—that is, close to the beginning of the rising part of the bathtub curve. The shape of the rest of the distribution is simply irrelevant. To be sure, this is not the case for run-to-failure equipment, but we rarely seek failure rates of any kind for such functionally unimportant equipment.

Figure 1-2 resulted from the analysis of a large quantity of data [2] from the airline industry in the 1960s, and amply demonstrates this point—that is, that preventive maintenance removes the long-term time dependence in all but a small proportion (6% to 11%; see Figure 1-2A, B, and C) of components, which happen to include aircraft reciprocating engines (B) and turbine engines (C). In Figure 1-2, time is plotted along the horizontal axis, failure rate along the vertical axis. Only in the subsets A, B, and C would the long-term time dependence conceivably be of interest.

A fourth reason failure rates are constant is that the wide range of times to wear-out, the randomizing effects of wear-out initiators, and the performance of PM tasks interact to scramble the time dependence to an even greater degree. In fact, some subcomponents will wear out and may even be replaced several times before the wear-out characteristics of other items come into play. Also, if the degree of degradation is not sufficiently advanced to be observable, even when a PM task is appropriately timed, the task may "miss" the observable indications and the mechanism might proceed to the point of failure. This could occur when a wear-out mechanism is randomly initiated part way between successive performances of a PM task which is supposed to address it. The statistical effect is to smear out the failure time distributions by randomizing the time origin. No specific part of the failure time distribution can then be uniquely correlated with the cycle of PM tasks.

Similarly, no PM task is completely effective at detecting targeted degraded conditions, even when the conditions are present and observable at the performance of the task. Such degradation mechanisms proceed to cause failure in accordance with the failure time distribution, but the occurrence rate is attenuated by the effectiveness of the task. This is a random influence itself, and permits such a failure to "leak through" the PM defenses in only a small percentage of cases.

The fifth reason failure rates are constant is that a large fraction (typically 25% to 75%) of the many failure mechanisms are random in nature and therefore contribute directly to a constant failure rate. This is true for all equipment, some to a much greater degree than others.

It is clear that we should not expect to observe any striking time dependence in failure rates of complex equipment. There are also several practical or logistical reasons why the limited residual time dependence fails to be detected.

The first such logistical reason is that active equipment is likely to be partly or completely refurbished or replaced at certain intervals, involving the insertion of new subcomponents at different calendar times. This brings into question the meaning of "the" age of the component as a whole, and makes tracking of the age of the subcomponents an onerous task—one that has not yet been done systematically anywhere in the nuclear power industry. Further, most failure rate quantification in the industry has been done to provide failure rates for probabilistic safety assessment models, and these have not required time dependent data. Further evidence of this effect is that it is often stated in textbooks describing time dependent methods that wear-out applies most successfully to small subcomponents such as springs and elastomers. These are usually the only level at which a unique age could be ascribed to the hardware, and for which only a single, or at most a few, wear-out mechanisms might apply.

Second, most nuclear plant equipment is very reliable, and is not present in the very large populations typical of fleets of airplanes, motor vehicles, or consumer items. The consequent lack of failures makes failure rate estimation a very uncertain affair (see Section 5). The pressure is always to increase the sample size to increase the number of failures experienced in order to improve the accuracy of the estimate. This generally leads to the pooling of data from several plants, or even across the whole industry. At the least, this tends to mix data from different PM programs, duty cycles, and service conditions, as well as from different manufacturers and model lines. Influences on the failure rate from these effects tend to obscure trends with the age at failure.

Typical uncertainties associated with constant failure rates are about a factor of 2 each way at best, and more usually a factor of 3 to 5 or more. Paradoxically, the effort to reduce the numerical uncertainty by increasing the sample size in failure rate analysis leads to increased uncertainty over whether the result obtained from a wide range of different conditions actually applies to the particular component of interest. This is true regardless of whether the time dependence of the failure rate is in question. Against this background of uncertainty, the weak age trends exemplified by Figures 1-2A, B, and C are almost impossible to identify, and ultimately resist interpretation.

Third, subdividing the limited failure experience into subgroups of different age to determine a trend of failure rate over time would significantly increase the uncertainty in the failure rate for each subgroup. Although standard regression (trend) techniques could be applied to this kind of age data to determine its time dependence, the large uncertainties on individual data points, the small number of such data points, and the fact that confidence intervals on regression parameters widen toward either end of the range of data all combine to make this approach viable in only rare cases.

Fourth, Weibull analysis, the often-quoted method to determine time dependent failure rates, uses a sample of the ages at failure. Even if such data were available, it would need to apply to non-repairable or non-PM'ed items. Even more important, the method plots the times to failure against the cumulative failure fraction. It requires a large fraction of the population to fail in order to provide reasonable estimates of the Weibull parameters. This is not an impediment in a manufacturing environment where a number of items may be put on test and the test is run until most have failed. But this situation almost never arises in the nuclear power industry, where corrective and proactive actions must be taken as soon as the first few failures of a specific type occur on the same set of critical components in a plant, or even across the industry.

Constant failure rates for complex equipment are therefore a fact in both principle and practice.

3.1.3 Good as New and Bad as Old

Many of the above reasons for constant failure rates can assist in formulating a simple model for the effect of PM on reliability. Some of the potentially most useful features in this regard are the several ways in which combining many failure mechanisms among themselves, and with the performance of PM tasks, randomizes the times to failure and diminishes the importance of the individual time to failure distributions, almost all of which are likely to remain unknown.

However, it is a mistake to imagine that this means that a complex SSC can be treated as a homogeneous block, because different PM tasks and their task intervals must be represented explicitly in any useful model. Each PM task can have a very different capability to detect the different kinds of degradation. Moreover, the industry has amassed a wealth of experience to describe the effectiveness of different tasks for protecting against different degradation mechanisms. A model which includes this information could be equipment-specific and PM task-specific and encompass the accumulated and ongoing industry experience. This is the approach taken in the EPRI PM Basis database.

This approach also solves a ubiquitous problem in prior mathematical models of reliability as a function of preventive maintenance. In mathematical alternating renewal models, a decision has to be made about whether the performance of a PM task restores the equipment to an essentially as-new condition, or whether the task scope is so narrowly focused that most of the equipment's piece parts remain in the same condition after the task is performed as they were in before task execution. The two assumptions are referred to as Good-As-New (GN) and Bad-As-Old (BO), respectively. When such models are applied at the level of the whole equipment, the GN or BO assumption is such a gross approximation as to be almost meaningless.

The PM Basis database approach is to specify the effectiveness of each PM task in addressing each and every failure mechanism listed for the equipment. If its effectiveness is zero, it means that the task does not address the failure mechanism at all. The equipment condition after the task for that specific mechanism is thus equivalent to the BO assumption. If the task is assumed to be 100% effective for a specific failure mechanism, it means that after the task is performed the relevant subcomponent has been restored to an as-new condition with regard to that specific degraded condition, and is thus equivalent to the GN assumption. In practice, the 100% GN assumption does not seem realistic under any circumstances, because of opportunities for maintenance error and because some of the time a wear-out may be initiated too late for the condition to have degraded sufficiently to be detected with confidence during the relevant PM task. Consequently, task effectiveness between 0% and 100% is used (actually binned into Low, Medium, and High). A task effectiveness in this range corresponds to an assumption between the extremes of BO and GN. Consequently, the EPRI PM Basis model contains a more sophisticated treatment of the BO/GN issue at the level of each PM task and failure mechanism than is normally encountered in alternating renewal models.

This is very important to the validity of the model, because the task effectiveness information contains the essence of how PM tasks interact with degraded conditions, and it can easily be extracted from experienced maintenance personnel who have absorbed industry experience over the course of their careers.

3.1.4 Intrinsic and Overall Task Effectiveness

The task effectiveness is not quite as straightforward as it appears from the foregoing discussion, because it is obviously also a function of the timing of the task in relation to the failure time distribution. The discussion above refers to the overall task effectiveness, which combines the effect of the task's intrinsic capability to detect a degraded condition when it is performed at the right time, and the task interval, which determines whether it is performed at the right time.

Two quantities are thus introduced, the intrinsic task effectiveness and the overall task effectiveness. The intrinsic task effectiveness answers the question, "If the task is performed when the degraded condition is present, what is the probability (High, Medium, Low) that some relevant anomaly will be detected?" The reference to a relevant anomaly rather than the degraded condition itself is to admit that diagnosis of the actual condition may not be perfect, but some anomaly is discovered which leads to restoration of the condition.

For example, an overhaul which results in direct examination of piece parts might be expected to be intrinsically highly (almost 100%) effective in detecting wear, broken or missing parts, burned or defective insulation, and so on. On the other hand, operator rounds might only be of low intrinsic effectiveness at detecting a worn bearing, depending on many factors such as the amount of wear and noise produced, the noise level in the neighborhood of the equipment, and the experience of the operator. Operator rounds might be assessed to have a low intrinsic effectiveness for such a degraded condition, but of high intrinsic effectiveness in detecting oil leaking onto the floor or steam escaping from a faulty gasket.

Consider a wear-out mechanism which leads to some expectation of failures beginning at around 6 years and extending beyond, for which an overhaul has a high intrinsic effectiveness. If the overhaul were performed with an interval around 6 years, its overall effectiveness would remain high, and equal to its intrinsic effectiveness. Alternatively, if the overhaul is performed at 10 years its overall effectiveness cannot be expected to remain at the level of its intrinsic effectiveness. The overall effectiveness would be expected to be downgraded to perhaps medium or even low.

The algorithm used to modify the intrinsic effectiveness as a function of task interval and type of degradation mechanism is in accordance with engineering judgment, and is described more fully in the Application Guideline within the EPRI PM Basis database, and also in the User Manual. It is subject to modification according to the results of future validation activities.

3.1.5 Why the Failure Time Distributions Ultimately Are Not Important

Figure 3-2 shows the effect on the predicted failure rate of successively increasing the Overhaul task interval for a critical centrifugal compressor using the EPRI PM Basis database. The function g(t) is a form factor representing the changes relative to the current failure rate (that is, the rate when the overhaul task is performed at a 5-year interval). The time points are arbitrary.

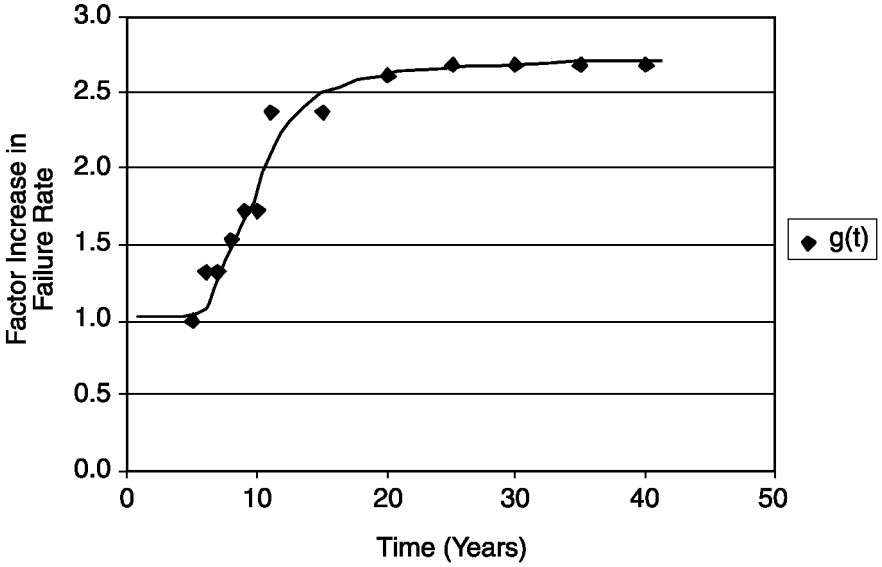

Figure 3-2
Effect on Failure Rate of Changing the Compressor Overhaul Task Interval

The line is drawn by eye to be a smooth curve, although local plateaus and jumps mark the actual predicted course of the results. Smoothing seems to be justified by the impossibility of verifying the calculated behavior over times as short as a year or two (see preceding discussion of time dependence in Section 3.1.2) and by the conceptual difficulty of even defining the failure rate over short periods of time (see Section 5). The jumps are caused by the "uncovering" of additional wear-out failure mechanisms as the increasing task interval successively moves to times longer than the leading edge of the successive time to failure distributions. In practice it is the failure rate itself which is crudely represented in the algorithm by a step function for each wear-out mechanism, rather than the failure time distributions.

The point is that many wear-out mechanisms contribute as the task is performed at successively later times. Although the detailed form of the failure time distributions can never be known, the addition of so many of them over a period of 10 to 40 years produces an effect mainly influenced by the inclusion of 100% of the failures for the shorter term mechanisms (that is, the whole distribution is encompassed as the task interval increases), and the more or less sudden experiencing of new mechanisms. The detailed form of the approach to the asymptote at long intervals will not be correct, but for practical purposes a straight line could be drawn to represent the curve between 5 years and 15 years with little loss of useful information for maintenance decision making.

It was stated earlier that for critical equipment, no one will propose moving important task intervals into the region where many wear-out modes are exposed, as in the graphical experiment above. However, it can be seen that the results could still be useful for decision making even in such an extreme case. Consider, for example, the decision to perform the overhaul at a longer interval for the "significant" category of equipment than for the "critical" category, or not at all for the "minor" category. Normal use of the model will be to detect the range of task interval at which significant impacts on equipment reliability will occur.

Even if the shape of the bulk of the failure time distributions could be determined, the shape of the leading edge would remain extremely uncertain (even though the mathematical model would undoubtedly contain an exact, but idealized and misleading formulation). This kind of information is essentially unknowable. In a decision-theoretic framework it is therefore appropriate to treat the leading edges of the failure rate functions as a set of uniform step functions, as in the PM Basis algorithm.

The leading edge of the time dependence of each failure mechanism can thus be conservatively represented by a step function rather than by more gradual increases. The later parts of the distributions have little impact, because their effect is either removed by the performance of PM tasks, or they add up to a result which is not sensitive to the detailed shape of any one of them, or they are smeared out to resemble uniform distributions by the various randomizing effects. For each wear-out mechanism the time of the step, the useful life, is a quantity which can be elicited directly from experienced maintenance personnel in workshop sessions.

The PM Basis model is a maintenance model of equipment reliability created to capture the essential nature of industry PM experience expressed through the ability of each task to address each failure mechanism, while justifiably avoiding relatively unimportant mathematical distractions.

3.2 Linear and Nonlinear Interactions Between PM Tasks

3.2.1 Gaps in Protection Between Tasks

Different PM tasks will exhibit different scope and effectiveness over the set of failure mechanisms for a complex SSC. Where one task may be very effective against a failure mechanism, another task may have no effectiveness at all. The total effect on SSC reliability will consist of the total effect of all the tasks on all the failure mechanisms, taking account of the nature of each failure mechanism (random or wear-out) and the task intervals and intrinsic effectiveness values.

It may happen that a subset of the failure mechanisms is not addressed by any of the PM tasks which are performed. These will exert their full effect on the failure rate, unattenuated by the protective effects of PM. All the other failure mechanisms will make some contribution to the failure rate, but those mechanisms will be heavily attenuated by highly effective PM tasks, or by combinations of partially effective tasks. When a comprehensive level of PM is applied, as for a critical SSC, the attenuated contributions of the majority of failure mechanisms which are protected against add up to a fairly low residual failure rate. The existence of one or two failure mechanisms which have no protection at all may be considerable, hence the significant impact on the industry of new failure mechanisms which are not adequately addressed by existing PM tasks.

Experience using the PM Basis database has shown that even with the full set of recommended PM tasks performed at recommended intervals it is common to find some failure mechanisms (perhaps 10% of the total) which are either not addressed at all by any task, or which are not

addressed by any task with an overall effectiveness better than "Low." These are the origins of the residual failures which occur despite careful attention to the PM program. In some cases, the proportion of such poorly protected failure mechanisms rises much higher. This usually occurs when applicable PM tasks exist to improve the level of protection, but where the tasks cannot be performed at the appropriate times because of inaccessibility of the equipment during power operations, or because the cost of performing the tasks at appropriate intervals would be prohibitively expensive in terms of resources or unavailability.

It is important for utility management and regulators to appreciate that even a comprehensive PM program does not provide good protection for every failure mechanism which is known to occur. The poorly protected mechanisms are always preponderantly random in nature, underlining the difficulty of addressing random failure mechanisms, and the importance of developing a strong set of predictive PM capabilities (the only tasks which can be effective against random failure mechanisms). Plant personnel should ideally adopt a graded approach in their response to failures which are experienced, with more attention paid to the more preventable failures. The exception is obviously for failure mechanisms which are new to industry experience, which might be addressed by the addition of new PM activities.

The effect of gaps in the protection afforded by a PM program is essentially a linear phenomenon; the more gaps there are, the more they add directly to the failure rate.

3.2.2 Overlaps in Protection Between Tasks

When the scope and effectiveness of two PM tasks have a significant overlap, in that they address the same subsets of failure mechanisms to a significant degree, the effects on the failure rate of adding or deleting the tasks can be decidedly nonlinear. Consider an idealized case where the two tasks overlap precisely in the failure mechanisms they address and the effectiveness with which they address them.

From an initial PM program which includes neither of the tasks, suppose the addition of either one of the tasks significantly decreases the failure rate. We would claim that the task is obviously an important task. But if the second task is subsequently added to the program, the further decrease in the failure rate will be much smaller than the initial decrease. For example, if each task acts independently, and the first removes 90% of the failures attributed to one failure mechanism, the second will only remove 90% of the remaining 10%—a further 9% effect, ten times less than the effect of the first task. In practice the two tasks are unlikely to be completely independent, in that if the nature of the degraded condition leads the first task to "miss" it, the second task may be more likely to miss it also. The effectiveness of the second task may then only be 70% or 80%, and the incremental benefit of the second task only 7% or 8%. Either way, we might think that the second task (whichever one it is, because the same conclusion applies to both tasks) is not an important task. The conclusions about which task is important would be exactly reversed if the tasks were added in the reverse order.

The same effect means that if both of the tasks are performed in the initial PM program, and one is deleted, there will be practically no increase in the failure rate. We might conclude that the task is not worth doing. If, subsequently, the second task is also deleted there would be a

significant increase in the failure rate and the second task would appear to be important. The conclusions would again depend on the order in which the tasks are deleted.

The lesson here is that the effects of a task are often strongly dependent on which other tasks are being performed. Because the task intervals affect the overall task effectiveness parameters, the same kind of effects can be observed by changing task intervals rather than by deleting tasks altogether. Two conclusions are apparent: 1) the value of a PM task cannot be judged in isolation from the other PM tasks which are performed, and 2) attempts to rank the relative importance and benefit of different tasks in a PM program must take the initial program of tasks into account.

The effects of overlaps between the scope of different tasks can be highly nonlinear, with the potential to spring surprises on the unwary.

3.3 Task Deferral

Deferring a PM task is usually a temporary state of affairs limited to a single instance of delaying the task performance to a later time. This is to be distinguished from a permanent change in a task interval, which means having the intention to repeatedly perform the task at the new interval. In the remainder of this section guidelines are first provided for making permanent changes in PM task intervals, and then guidelines are provided for one-time task deferrals. Finally, the statistical treatment of deferrals, as manifested in the late performance of a significant number of PM tasks, is discussed. The Appendix describes the generic numerical model to assess the impact of PM task intervals on reliability, which was used to guide the numerical recommendations in the following sections.

3.3.1 Permanent Changes in Task Interval

In all cases below, where an interval change is discussed, it is assumed that the correct PM tasks are being performed. It is important to be sure that the set of PM tasks is appropriate, that they all have the appropriate scope, and that the required skills, procedures, and vendor and industry information are employed in their execution, before assuming that a task interval might need adjusting.

3.3.1.1 Decreasing the Interval

The prime indicator of a need to decrease the interval is poor equipment condition. The interval should be decreased (that is, the task performed more frequently) if the equipment condition has deteriorated to the point where you lack confidence that, even after it has been restored by the PM task, the equipment will remain unfailed through the following (unchanged) task interval. Conceivably, in some cases equipment condition could be extrapolated to show the likely condition by the end of a modified interval.

If you have already experienced one or more failures, or have experienced a severely degraded condition, you need to establish the cause of failure or degradation to be sure that it is a

degradation mechanism that the task in question is supposed to address—that is, to be able to detect before failure occurs. In other words, do not adjust the interval for the wrong task. The right task to have its interval adjusted is not necessarily the same task as that in which the condition was discovered. In addition, the degradation mechanism must usually be of the wear-out kind for a time-directed task to have any significant chance of improving the situation. The EPRI PM Basis provides clear indications when a task could benefit from a reduced interval over the whole range of failure mechanisms it addresses. Before proceeding to consider a task deferral, be sure that the task was actually performed the last time it should have been. Also, be aware that a certain amount of degradation is expected to occur between PM tasks. A PM task is normally performed to detect such degradation if it is present, and the task often contains restorative actions regardless of whether degradation is detected. Consider decreasing the interval if 1) experience shows that the degradation is so advanced by the time the task is performed that you judge there may be a significant chance of a failure in the future if the interval remains the same, or 2) you judge for any reason that maintenance action should have been taken sooner.

The EPRI PM Basis provides immediate indications of the impact of the proposed change on the reliability and availability of the equipment, and advises on the suitability of the change from the point of view of the balance between reliability and availability (see Section 4).

If the poor condition results from a poorly protected failure mechanism, check in the PM Basis database to verify that decreasing the interval does indeed result in a higher level of protection to the failure mechanism.

3.3.1.2 Increasing the Interval

Two conditions suggest consideration of an increase in interval:

1. There is convincing evidence that the equipment condition at the existing interval is invariably good enough to enable an extension of the interval by the proposed amount, usually by at least 25% of the existing interval, but not normally more than 2 years at one time; and

2. No relevant failures have been experienced using the existing interval.

Note that the equipment condition requirement is necessary for a confident increase in interval, and that it requires a judgment that the condition is not merely good, but good enough to regularly last to the extended interval. There will be a more confident condition assessment if the condition of a number of similar components is observed, if the observer knows what kind of degradation to look for, and if some measured parameter can be trended. When a group of similar components is available for interval extension, it is beneficial to stagger the initiation of the interval increase among components so that some components can deliver condition information at the extended interval before the others reach the extended interval.

The EPRI PM Basis database provides explicit guidance on whether the recommended interval for critical components with a high duty cycle and operated in severe service conditions appears to have scope for interval extension.

Often, residual failure rates when PM is effective (that is, the interval is less than or equal to the shortest failure-free wear-out interval) are low enough to give MTBFs in the region of 10 to 25 years, even when randomly occurring failure modes are factored in. Therefore, the fact that zero failures have been observed in many (~MTBF/I) task intervals cannot be taken by itself as a justification to increase the interval. *Therefore, it is not valid to increase intervals by referring to the absence of failures over moderate periods of time, without also considering the condition of the equipment.* If equipment condition information is gathered in the normal course of performing PM tasks, the timely information available on which to base interval extension decisions will greatly increase in quantity and quality.

3.3.1.3 25% Increase

If equipment condition is satisfactory, the interval could be increased by about 25% without further analysis. This is because in the worst case, where the original interval is optimal (that is, the interval is already set at the shortest failure-free interval of all the relevant wear-out failure mechanisms), and some failure modes become unprotected by the change, the failure rate is unlikely to increase by more than 15% to 30%. This happens because the failure-free period only indicates the gradual beginning of the failure time distribution, which usually extends over a long time, and because only one or a minority of failure modes will be so affected. The importance of good equipment condition at the existing interval is that it adds a significant measure of conservatism to this estimate by making it much less likely that the worst case applies.

3.3.1.4 Larger Increases (>25%)

When increases are constrained for practical reasons to be larger than 25%, for example when 1.5 years would be changed to 3 years because access would be impossible at power, this 100% increase in interval may increase the failure rate by 125% to 250% (that is, the new rate could be 3.5 times the old rate) in the above worst case (that is, when you start at the optimum and failure modes become unprotected). However, this result is not only the worst case, but also supposes that you do not have any a priori knowledge of whether failure modes will become unprotected by the increase in interval. Observing consistently good equipment condition at the existing interval is one way to be sure that you are not in the very worst condition of uncovering failure mechanisms for even modest increases in interval. Judging that the condition will remain good for the duration of a large proposed increase can be more demanding.

To add additional confidence that a large increase will not result in unacceptable failures, you should prospectively try to assure that one or more of the relevant failure modes do not lose their PM protection when the interval is increased. Again, the EPRI PM Basis database has been designed to display this information explicitly.

Consider, however, that not every failure mechanism is relevant, that is, needs to be addressed, as the following list illustrates:

1. For the "minor" criticality category of components you may only need to defend against failure mechanisms which have occurred before at this plant, and against the most common failure causes experienced in the industry, and maybe not all of these, depending on economic factors.

2. For any category of components, random failure mechanisms cannot be effectively defended against by tasks with intervals longer than about 1 year and should be ignored for these tasks. In other words, if the task whose interval is being modified has an existing interval longer than 1 year, it will already be providing only feeble protection against random failure mechanisms, so concern about coverage of the random mechanisms should not unduly influence your decision. In contrast, the effect of random failure mechanisms could dominate the interval extension decision if the existing interval is less than 1 year.

3.3.2 One-Time Task Deferral

Because deferring a task one time is limited to a single occasion, the level of risk is generally less than that which would accompany a permanent change. However, deferrals are often sought for purely logistical reasons, and historically have not always been supported by equipment condition information.

The following procedure assumes that you have no historical *adverse* equipment condition information to suggest that the deferred task will lead to an unacceptable equipment condition or failure.

It will also be assumed that there is no specific reason to suspect that deferring the task will leave known failure mechanisms undefended. Nevertheless, since the possibility exists, two things will happen. Generically, there will be an increase in the failure rate, and an increase in the probability of experiencing a failure during the period of deferral, T. When making the decision to defer a task, the probability of having a failure during the period of deferral takes on a special significance, independently of what might be happening to the failure rate. The failure rate increases with the length of the deferral. The probability of a failure (λT) increases with the failure rate, λ, and also with the length of deferral, T. The quantity which may be of most concern is thus more strongly dependent on the deferral period than is the failure rate itself. This suggests different treatment for critical (both "critical" and "significant" categories) and noncritical equipment, because the failure rate will not be of much concern for the "minor" category if the probability of a failure is controlled.

For both critical and noncritical equipment the following deferral recommendations assumed (arbitrarily) that the probability of a failure should not be more than 0.1 or 10% in absolute terms as a result of the deferral. For critical equipment it is also assumed that the failure rate should not be permitted to increase by more than a factor of two. This is because an increase of a basic event probability in the plant PSA by a factor of two will cause a relative increase in the Core

Damage Frequency (CDF) by the FV fraction (that is, ΔCDF/CDF = FV). This should be acceptable on a one-time basis for a limited period except for the most risk-significant equipment. It is assumed that PM tasks for very risk-significant equipment (for example, FV>5%) would not be deferred without additional evaluation. The failure rate criterion is not applied in the case of noncritical equipment.

The following rules were derived from detailed but generic failure rate calculations which took into account the generic number and distribution of failure mechanisms, random failures from non-wear-out failure mechanisms, task effectiveness, and the proportion of wear-out failures addressed by a PM task. The rules assumed the latter proportion was conservatively 100%. The calculations also assumed there is no conservatism in the existing PM program, so as to maximize the effects of deferrals (that is, the existing task interval equals the shortest failure-free wear-out period).

3.3.2.1 Noncritical Equipment

Defer one time without further evaluation, up to the following limits. If the deferral is longer than these limits, it requires evaluation.

Interval (Years)	Defer by (Years)	A One-Time Deferred Interval May Thus Become (Years)
1	1	2
1.5	1.5	3
2	2	4
>5	3	>8

3.3.2.2 Critical Equipment

Defer one time without further evaluation, up to the following limits. If the deferral is longer than these limits, it requires evaluation. Tasks with intervals 1.5 years or less should not be deferred without evaluation.

Interval (Years)	Defer by (Years)	A One-Time Deferred Interval May Thus Become (Years)
≤1.5	Requires Additional Evaluation	_____
2	1	3
3	1	4
>4.5	1.5	>6

It must be stressed that the above recommendations are generic, and are based on arbitrary, but reasonable, rules. Their value lies in the fact that, given the rules and the generic assumptions, the results are in the range of typical deferral decisions. Going outside the range of these recommendations may lead to either a greater than 10% chance of a failure in the deferral period, a doubling of the failure rate, or both.

The above rules should cover almost all cases of interest without the need for further evaluation. For more component-specific and task-specific guidance, use the EPRI PM Basis database to estimate the effect on failure rate. Calculate the effect on the probability of a failure manually using the following expression:

$$Probability\ of\ a\ failure\ in\ the\ deferral\ period\ =\ \lambda T \qquad \textbf{Eq. 3-3}$$

For this you need access to a value for the current failure rate. In addition, any one of the following three methods may be used to justify task deferral outside the above limits:

1. Discover if the equipment condition has consistently been good enough so you can judge that it is able to reach the deferred task execution time without failure. This requires data from plant specific experience.

2. Discover from experience at other plants if the proposed deferral is likely to lead to a failure.

3. If the current interval is sufficiently less than the interval recommended in the EPRI PM Basis database, so that even with the deferral, the combined period does not exceed 125% of the recommended interval, it should be safe to defer the task unless there is plant specific experience to the contrary. This requires a very conservative initial task interval.

3.3.3 Guidance on the Use of Grace Periods

The same kind of generic failure rate modeling has been used to estimate the impact of employing different policies to control the number or proportion of plant PM tasks which are permitted to exceed their planned PM intervals. The Appendix describes the generic numerical model to assess the impact of PM task intervals on reliability.

Practical constraints result in some PM tasks at nuclear power plants being performed later than scheduled. This is unavoidable even in good maintenance programs where the PM intervals are optimal or conservative. To limit the risk of additional failures, most plants adopt a “grace period” for performing a PM task, limited to (for example) 25% beyond the scheduled time. PM tasks delayed longer than the grace period are reported as delinquent.

Some plants schedule the tasks at intervals which are 20% shorter than the technically optimal intervals so that a grace period 25% beyond the scheduled interval still meets the intent of the optimal interval. Consequently, most PM tasks get scheduled and performed considerably sooner than their optimal intervals in order to reduce to almost zero the number that become delinquent. This trend adds to PM costs and may harm reliability by introducing unnecessary maintenance error.

The objective of the model results discussed in this section is to generate a reasonable strategy from a reliability perspective, which plants can adopt as policy, regarding 1) how long the grace period should be; 2) when tasks should be performed within the grace period; 3) how to track the plant performance in meeting this goal; and 4) the degree to which the overdue date might be exceeded without undue risk.

3.3.3.1 The Length of the Grace Period

Intrusive PM tasks performed before their technically optimal task intervals are likely to increase the failure rate, because maintenance error and material defects are introduced more often. The effect (infant mortality) is commonplace for a wide range of equipment (for example, switchgear, AOVs, check valves, relays), and is evidenced by significant levels of rework soon after a maintenance outage [16]. Therefore, scheduling intrusive tasks too soon is detrimental. Nevertheless, to prevent many tasks from becoming delinquent it is a practical necessity to perform a significant proportion of all tasks before their optimal intervals.

If an intrusive task is performed 20% earlier than its optimal interval, the infant mortality part of the failure rate, which is already roughly equal to the best failure rate that good PM can produce (see Appendix), will increase by a commensurate 20%, regardless of the fact that no "naturally occurring" failure modes are expected when the task is performed at this early time.

We distinguish two cases, in both of which the technically optimal interval is that beyond which wear-out failure modes can be expected to occur.

In case A, conservatism is built into the scheduled intervals:

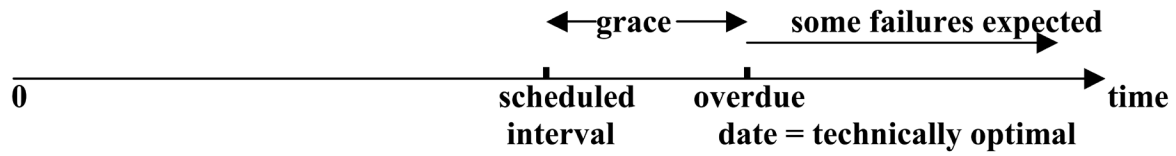

In case B, no conservatism is built into the scheduled intervals:

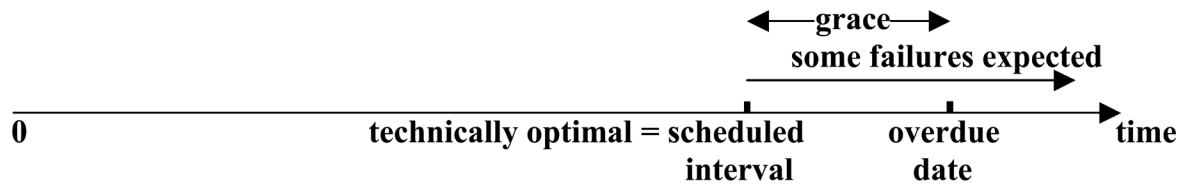

Because infant mortality erodes the benefits of good PM, in case A the degree of conservatism (and hence the grace period) should not exceed an amount which, following current industry practice, we will initially consider to be 20% of the technically optimal interval (that is, 25% of the scheduled interval). Less is better if it is also practical.

In case B, we also initially consider the grace period to be 25% of the scheduled interval, but in this case no infant mortality considerations arise. Instead, there is a concern that reliability may be worsened, because wear-out failure modes could in principle occur during the grace period.

It will be assumed that despite this residual concern, no wear-out failure modes are actually *known* to occur with high probability within the grace period. This is a good assumption for grace periods in nuclear power plants. Both cases are considered because they are common in the industry.

3.3.3.2 Task Performance Within the Grace Period

Case A: Within the grace period it is advantageous from a reliability perspective to perform intrusive tasks as close as possible to the end of the grace period (that is, to the overdue date), so they are not performed too frequently. Condition monitoring tasks and other non-intrusive tasks may be performed sooner with little detrimental impact on reliability. Since some intrusive tasks must still be performed before others, it is better that these be the tasks with the longer intervals, because these will represent a smaller proportionate increase in failure rate. For example, 90 days before the overdue date is a 17% shortening for an 18-month interval, but only a 2.5% shortening for a 10-year interval.

Case B: Within the grace period it is advantageous from a reliability perspective to perform all tasks as close as possible to the start of the grace period (that is, to the scheduled date). Since some tasks must be performed before others, it is better that these be the tasks with the longer intervals, because these will represent a smaller proportionate increase in failure rate. For example, 90 days before the overdue date is an 8.3% extension for an 18-month interval, but is a 22.5% extension for a 10-year interval.

3.3.3.3 Tracking Task Performance Within the Grace Period

In case A, performing PM tasks before their overdue date confers no reliability benefit. The sole benefit is the practical matter of avoiding too many overdue tasks. Consequently, plants need track only those tasks approaching the overdue date, and only to the degree that it facilitates task implementation to avoid delinquency. For example, tracking tasks within 90 days of the overdue date could be a solution.

The number of tasks permitted to be within 90 days of their overdue date could be limited to somewhere in the range 50 to 200, depending on plant experience with getting tasks completed. There does not seem to be a useful purpose in limiting the overall number of tasks in the whole grace period, since there is no reliability penalty for being “in grace.”

In case B, performing PM tasks before their overdue date does reduce the reliability disbenefit of exceeding the scheduled date. The following section, Exceeding the Due Date, puts this concern into quantitative perspective, and demonstrates that performing tasks up to 25% beyond their due date does not lead to a significant reliability increase.

Tracking tasks that are within 90 days of the overdue date could be a practical solution.

The number of tasks permitted to be within 90 days of their overdue date should also be limited, depending on plant experience with getting tasks completed. In this case, there is also a useful purpose in limiting the overall number of tasks in the whole grace period, since there is a reliability penalty for being "in grace."

3.3.3.4 Exceeding the Due Date

An optimal PM interval for a time-directed PM task corresponds to the onset of a non-zero probability of failure after an expected failure-free interval. This probability distribution rises slowly. It will generally be many years before the chance of a failure has become a certainty. Consequently, to exceed the optimal interval does not necessarily result in immediate failures, but only in an increase in failure rate. The accompanying model of failure rates (see Appendix) shows that a population of components, with different task performance times in relation to their optimal intervals, can be permitted to extend past the technically optimal date to a considerable degree without causing a sudden large increase in failure rate. As a broad generalization, the increase in failure rate caused by exceeding the optimal intervals by a given percentage (<~50%) is roughly similar in magnitude to that caused by infant mortality when shortening the intervals by the same percentage.

In case A, if the population of actual task performance times is centered anywhere between the scheduled intervals and the optimal task intervals (with standard deviation 12.5%), the average failure rate increases by about 6% at the most. In case B, if most of the population is positioned between the scheduled date and the overdue date but with 15% of the components in the grace period *past the overdue date*, the overall failure rate would be increased by only 20%.

Moreover, a specific component which does not get its task performed until 25% (this is 2 standard deviations if the mean is at the optimal interval) beyond the overdue date in case A, or 25% beyond the scheduled date in case B, experiences an increase in failure rate of no more than about 30%. In fact, it is the relatively slow response of failure rate to increasing interval which permits the possibility of finding the right interval by trial and error without excessive danger from overshooting.

The results indicate that it should be possible to permit a certain proportion of tasks to be performed beyond the optimal date without significant harmful effect. It is suggested that the strategy to be followed should avoid designating tasks as delinquent unless their performance times exceed some limit *beyond* the overdue date in case A. The model results show that even if 15% of the components in the grace period go past *125% of the overdue date* in case A, the overall failure rate would be increased by only 20%. In case B, a delinquent component should be one that has not received its PM task by the overdue date. Even then, if 15% of the components in the grace period become delinquent, the overall failure rate would be increased by only 20%.

3.3.3.5 Proposed Strategy

1. The proposed strategy would focus on completing, by the technically optimal date (that is, by the overdue date in case A, and by the scheduled date in case B), PM tasks which:
   - Are for risk-significant components because their Fussel-Vesely (FV) parameter >0.5%.
   - Are technical specifications, surveillance tests, or code requirements.
   - Are for components in 10CFR50.65 (a)(1) (Maintenance Rule).
   - Are known to be needed to prevent a known high risk of failures—for example, replacing head valves in reciprocating compressors.

2. In case A, the due date (date scheduled) would be programmed at no more than 20% less than the overdue date. The grace period would be the time between these dates. There would be no negative connotation attached to being in the grace period. The grace period exists only to focus on completing tasks by the overdue date. This case is more expensive to implement than case B, and does not contain as much conservatism as might be expected from the adoption of task intervals which are shorter than the technically optimal intervals. The negative impact on reliability of infant mortality is likely to cancel out the benefits of conservative task intervals.

   In case B, the due date would be the technically optimal interval. The grace period would extend an additional 25% of this interval. There is a disbenefit to being in the grace period, but this is moderate and controlled by other steps. This case is less expensive to implement than case A.

3. In both cases, the number of tasks that are within 90 days of the overdue date could be tracked and limited to a number in the range 50 to 200, depending on plant experience. In case B only, the number of tasks that are in grace should also be limited to an overall maximum.

4. In case A, control workflow so that intrusive tasks with short intervals (for example, 2 years or less) are preferentially completed during this 90-day window, and not before—that is, as close as possible to the overdue date. Condition-monitoring and non-intrusive tasks could be performed earlier rather than later in the grace period to assist in workflow management.

   In case B, control workflow so that tasks with longer intervals (for example, 3 years or more) are preferentially completed before this 90-day window.

5. Screen tasks during the 90-day period before the overdue date so as to prevent tasks which are of the following type from going over the overdue date. The tasks involved:
   - Are for risk significant components because their Fussel-Vesely (FV) parameter >0.5%
   - Are technical specifications, surveillance tests, or code requirements
   - Are for components in 10CFR50.65 (a)(1) (Maintenance Rule)

- Are known to be needed to prevent a high risk of failures—for example, replacing head valves in reciprocating compressors at some plants

6. In both cases, permit some of the other tasks normally in the grace period to go over the overdue date if necessary for practical reasons (for example, when spare parts are not available), without being declared delinquent. Limit the total number of tasks to go beyond their overdue date to be no more than 15% of the total in the grace period.

7. Establish an upper time limit equal to the overdue date plus 25% in case A, and the overdue date plus 15% in case B (note that this is a proportion of the scheduled interval, not an absolute number of days), beyond which any task would be declared delinquent.

4
A BALANCE BETWEEN RELIABILITY AND AVAILABILITY

4.1 Introduction

The Maintenance Rule for U.S. nuclear power plants (10CFR50.65) requires licensees to ensure that a balance exists between the availability and the reliability of SSCs that are risk-significant [14]. NRC staff is concerned that in some instances licensees may take equipment out of service for preventive maintenance, making it unavailable to perform its safety function, for periods which are long compared to the safety benefit derived from the maintenance that is being performed.

Some of the unavailable hours which an SSC may accumulate during preventive maintenance may be avoided by adopting more efficient procedures. There may also be opportunities to shift some of the PM performed with the plant on-line to times when the plant is shut down, or to perform some of the PM when other equipment is out of service, effectively "shadowing" the unavailability in question. This analysis will assume that these opportunities have already been taken advantage of, to the maximum tolerable extent. What remains is the main objective of this section—a hard core of situations which do not yield to easy solutions.

It is not the objective of this section to propose a method for deciding if performance criteria on availability and reliability are balanced. These performance criteria have been set up as part of a monitoring mechanism to provide assurance (with considerable uncertainty) that the availability and reliability have values not too far from what they are supposed to be. The question of balance should be applied to the underlying assumptions and measurements concerning reliability and availability.

There is currently no universally accepted method for judging whether or not such a balance has been attained, and there has been relatively little discussion of the issues involved and their relationship to preventive maintenance and risk. The purpose of this section is to present a technically valid criterion for balance in the Maintenance Rule, which is based on general aspects of maintenance, reliability, and safety management.

The technical basis for the criterion may be found in sections 4.2 through 4.6. Section 4.2 provides a formal background in which the unavailability caused by on-line and off-line preventive maintenance can be related to reliability in general. Section 4.3 uses these relationships to derive a balance criterion for an SSC's availability and reliability. Section 4.4 extends this criterion to include hidden unavailability ($\lambda\tau/2$), repair unavailability, and multiple reliability failure modes (for example, fails to start, fails to run) of an SSC. Section 4.5 extends it

further to explicitly include risk significance—that is, to balance the effect of availability and reliability on the CDF, rather than just the bare availability and reliability of the SSC. This takes account of the situation where one or more of an SSC's basic events have a different level of risk significance compared to its other basic events. Section 4.6 addresses balance for supercomponents.

4.2 Formal Framework

We start by considering the restricted case of an SSC which has a single failure mode—for example, fails to start—with an *un*reliability represented by P_r. Also assume that it has *un*availability caused solely by preventive maintenance, designated as P_a. P_r and P_a are the basic event probabilities familiar from PSA, and represent the probability of the SSC being failed and unavailable, respectively, at a random moment when its functions are demanded.

The reliability basic event probability, P_r, depends on preventive maintenance being performed on the SSC. The whole issue of balance is a consequence of this and the fact that some of the necessary PM is performed on-line. It is reasonable to suppose that the failure probability will be higher if insufficient maintenance is performed, that performing the right kind of maintenance at the right time improves reliability, and that there exists a maximum reliability level (minimum unreliability) which can be reached with an appropriate maintenance program. This is an assumption which cannot be relaxed, because it means that the continued improvement of reliability with increasing amounts of appropriate PM has to show diminishing returns and flatten out at some point, an essential ingredient to the balance criterion.

Additional maintenance beyond that which is needed to achieve this maximum reliability may or may not introduce failure causes which would not otherwise have occurred, leading to an increase in the failure probability. This might be brought about by an excessive degree of intrusive PM involving tasks such as internal inspections, component replacements, and overhauls. For completeness, the following treatment will assume that reliability deteriorates as excessive PM is applied past the optimum point, but this is not a necessary assumption.

P_a is defined in the normal way as

$$P_a = H_{PMon} / H_{Req} \qquad \textbf{Eq. 4-1}$$

$$= H_{PMon} \cdot \gamma \qquad \textbf{Eq. 4-2}$$

where H_{PMon} is the number of hours spent performing preventive maintenance on the SSC, during which the SSC's functions are not available, at a time when the SSC's functions are normally required. For most SSCs which provide safety functions, which are the main target of the balance issue, this will be whenever the plant is on-line. Analogous definitions apply to systems required during shutdown. H_{Req} is the number of hours the SSC's function is normally required. H_{Req} is about 7,000 hours per year at an 80% capacity factor. The reciprocal $1/H_{Req}$ then has the value 1.43×10^{-4} (= γ), and is an important parameter in what follows, as it sets the scale for the

changes we will be considering. Because it appears in many places, this parameter has been given the name γ (gamma) to make it easy to refer to.

The numbers of hours are further related by

$$H_{PMtot} = H_{PMon} + H_{PMoff} \qquad \textbf{Eq. 4-3}$$

where H_{PMoff} is the number of hours spent performing PM on the SSC when the plant is off-line, and H_{PMtot} is the total number of hours spent performing PM on the SSC.

We know that the unreliability depends on the total number of PM hours deployed. The simplest way to represent the dependence of P_r in accordance with the preceding discussion is the curve shown in Figure 4-1, using H_{PMtot} as the independent variable (that is, the x-axis).

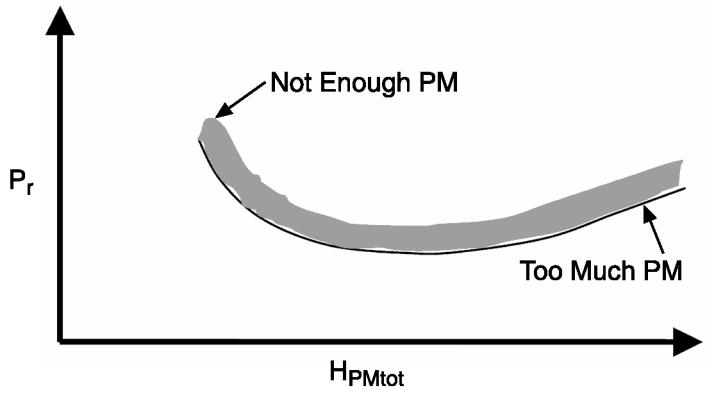

Figure 4-1
Unreliability as a Function of Total PM Hours

The development of a balance criterion will not depend on the detailed functional form of the curve, aside from the fact that it should possess a minimum somewhere. The minimum represents the best reliability that can be achieved through the exercise of PM, without considering a design change.

A given PM program may be operating above the optimal curve, in the region shaded in Figure 4-1, simply because the maintenance is being performed ineffectively, or tasks with little value or providing overlapping protection are being performed. Nevertheless, to assume that the PM tasks being performed are more or less appropriate ones, and that they are being performed in accordance with good industry practice and appropriate craft skills, would in fact usually be a reasonably good assumption for risk significant standby safety SSCs at nuclear power plants.

If the balance criterion depended on always being on the optimum curve, its utility might be limited, but we will not need such a strict constraint. It is best to consider that most PM programs operate in the zone shaded in Figure 4-1, which is bounded on the lower side by the optimum curve. Furthermore, implementation of the balance criterion will always involve practical, hence finite rather than infinitesimal, steps along the x- and y-axes, so that connections between points in the shaded region or on the curve will only enter through practical changes that can actually be

carried out. To develop the criterion we will proceed as if we are on the optimum curve. The criterion will subsequently be seen to also apply for points in the shaded region.

Figure 4-2 shows P_a as a function of H_{PMon}, as given by Equation 4-2. In terms of H_{PMon}, P_a is a straight line of slope γ starting at the origin, regardless of whether additional PM is done on-line or off-line. However, if P_a were represented as a function of H_{PMtot}, as in Figure 4-3, the same line is shifted to the right by an amount equal to H_{PMoff}.

Expressed in this way as a function of H_{PMtot}, P_a depends also on H_{PMoff}. If H_{PMon} is used as the independent variable, as will be preferred, we must obviously assume that H_{PMoff} is kept constant when we proceed to consider the effect on reliability and safety.

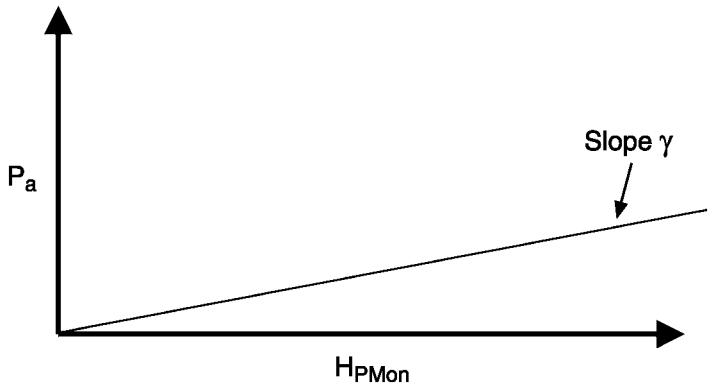

Figure 4-2
Unavailability as a Function of On-Line PM Hours

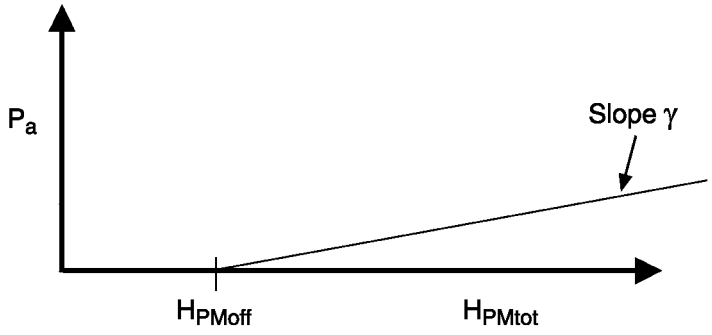

Figure 4-3
Unavailability as a Function of Total PM Hours

4.3 The Balance Criterion

Balancing is not intended to trade the *economic* impact of on-line PM against the *safety* impact of on-line PM. But it does trade the *safety* impact of the reliability achieved against the *safety* impact of the unavailability incurred to achieve it. However, different component types require dramatically different amounts of total PM to achieve their characteristic levels of reliability, quite apart from the question of how much of the PM could be done on-line. Clearly, a single

rule cannot be formulated for how much of the PM could be performed on-line, without involving parameters which depend on component types.

Nevertheless, one might suspect that a truly universal balance criterion might be constructed which involves only *changes* in total and on-line PM, and consequent *changes* in reliability. The reason is that the rate at which unavailability is changing as H_{PMon} changes is always the universal, constant, and known rate, γ. The result will be a universal rule which is simple to apply. Its limitation, and perhaps its strength, is that balance cannot then be addressed in isolation from the practical steps which would be required to change the existing degree of reliability and on-line PM.

Proceeding in this direction, a tentative balance rule seems to have two characteristics: 1) it involves changes in PM and in reliability; and 2) if a change in PM affects reliability, the change in total PM performed must be reflected in on-line PM rather than off-line PM—that is, we will not trade on-line PM for off-line PM—so H_{PMoff} is indeed to be kept constant. The rule will decide if a change in on-line PM (and therefore total PM) seems appropriate from the point of view of balance alone.

Figure 4-4 restates the functional dependence of P_r in terms of H_{PMon}, rather than as a function of H_{PMtot}. Further development will retain H_{PMon} as the independent variable. Figure 4-4, by analogy with Figures 4-2 and 4-3, is just Figure 4-1 shifted to the left by H_{PMoff}.

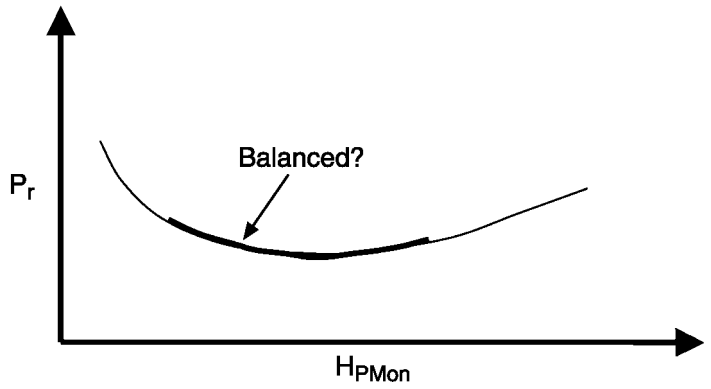

Figure 4-4
Unreliability as a Function of On-Line PM Hours

If some part of the P_r curve is ultimately to be designated as the "balanced" part, meaning that *the absolute value* of P_r appears to be balanced with the amount of on-line unavailability being caused by the PM, presumably this will correspond to some region around the minimum, shown in Figure 4-4 as a thicker line. The idea of designating a region around the minimum as being balanced is to permit an area of discretion in balancing availability and reliability. This is intended to underline the fact that being balanced is not necessarily the same as having the optimal reliability.

The problem is to decide how far the "balanced region" should extend around the minimum of the curve. Practical experience suggests that the minima of the above curves are often relatively

wide and shallow, because optimal reliability is usually approached slowly, in the sense of being achieved through increasing PM in the face of diminishing reliability returns.

To address the extent and location of the balanced region, we note that the probability that the SSC cannot provide its functions, given that it can be either failed or unavailable, is the sum of the probabilities of the two events, P_a and P_r. Represent this sum by P_{Sum}:

$$P_{Sum} = P_r + P_a \quad \textbf{Eq. 4-4}$$

The behavior of this quantity as a function of H_{PMon} is shown in Figure 4-5. The curve for P_{Sum} can be thought of as the curve for P_r sitting on top of the line representing P_a. At any value of H_{PMon}, the value of P_a can be visualized by the length of the vertical line from the x-axis to the straight line for P_a, and the value of P_r by the length of the vertical line from the straight line to the curve. But the amount of P_a which gets added to P_r increases as one goes to larger values of H_{PMon}. This shifts the minimum in P_{Sum} to a lower value of H_{PMon} than the minimum in P_r. Since the sloping line is completely determined, the uncertainty in the P_{Sum} curve is no larger than it was for P_r.

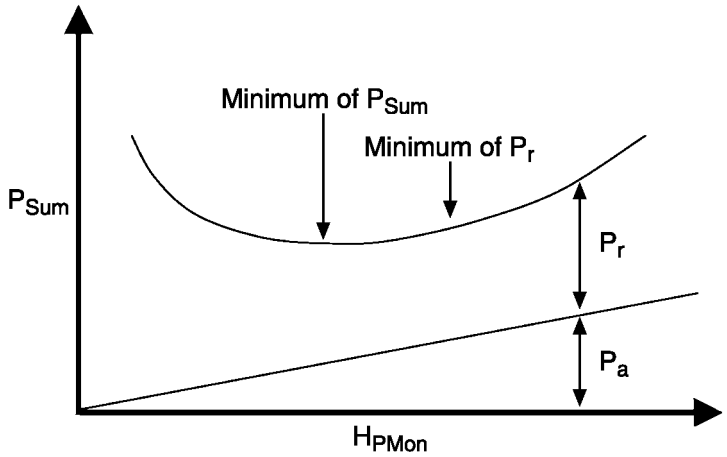

Figure 4-5
P_{Sum} as a Function of On-Line PM Hours

Algebraically:

$$dP_{Sum}/dH_{PMon} = dP_a/dH_{PMon} + dP_r/dH_{PMon} \quad \textbf{Eq. 4-5}$$

$$= \gamma + dP_r/dH_{PMon} \quad \textbf{Eq. 4-6}$$

Safety is optimized at the minimum of P_{Sum} where $dP_{Sum}/dH_{PMon} = 0$. At this point, Equation 4-6 gives (see Figure 4-6):

$$dP_r/dH_{PMon} \cdot = -\gamma \quad \textbf{Eq. 4-7}$$

The negative slope means that P_r is decreasing to the right. In other words:

> *Loss of SSC function (that is, P_{Sum}) is minimized (that is, safety is maximized) when there is still an opportunity to improve reliability by performing more PM.*

The "best" operating point, from the standpoint of the SSC providing its function, always comes at a worse value of reliability than the optimal reliability. This is because increasing on-line PM beyond the minimum in P_{Sum} can certainly improve the reliability, but at a rate which is less than γ per hour of unavailability, whereas the cost of the extra unavailability is always exactly equal to the rate γ.

Furthermore, at the point where P_r is a minimum, (that is, where $dP_r/dH_{PMon} = 0$), Equation 4-6 shows that the slope of P_{Sum} is equal to γ:

$$dP_{Sum}/dH_{PMon} = \gamma \qquad \textbf{Eq. 4-8}$$

This means that:

> *At the optimal reliability, protection against loss of function (that is, safety) is already deteriorating.*

Remarkably, even though we know next to nothing about the detailed shape of the curves, we have reached two significant conclusions, and we know quite a lot about the slope of the curves at key points:

When

$$dP_{Sum}/dH_{PMon} = +\gamma, dP_r/dH_{PMon} = 0 \qquad \textbf{Eq. 4-9}$$

when

$$dP_{Sum}/dH_{PMon} = 0, dP_r/dH_{PMon} = -\gamma \qquad \textbf{Eq. 4-10}$$

and when

$$dP_{Sum}/dH_{PMon} = -\gamma, dP_r/dH_{PMon} = -2\gamma \qquad \textbf{Eq. 4-11}$$

These points are marked in Figure 4-6, where the curve for P_r is displayed for comparison with P_{Sum}. The slopes of the curves are indicated in square brackets.

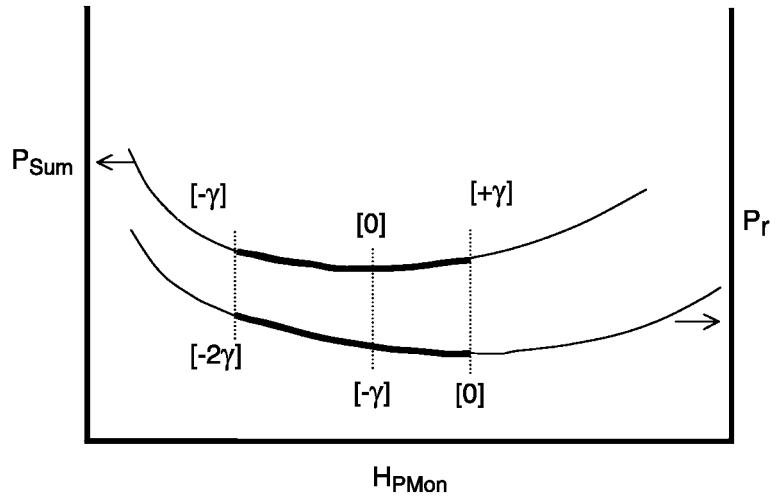

Figure 4-6
Corresponding Gradients on P_{Sum} and P_r

The significance of the point represented by Equation 4-9 (right end of heavy line) is that it is not reasonable to claim that points further to the right are balanced, because the reliability itself is beginning to deteriorate with increasing on-line maintenance. It is not in the economic interest of the industry to claim that cases of over-maintenance are properly balanced. Nor is this point balanced from the point of view of safety, because the value of P_{Sum} is already past its minimum and is further increasing to the right.

The significance of the point represented by Equation 4-11 (left end of heavy line) is that it does not seem reasonable to claim that points further to the left are balanced, because both P_r and P_{Sum} can be improved (that is, decreased) with more on-line maintenance. This point is not uniquely determined as the limit of the balanced region, but the rate at which P_{Sum} is changing at this point is equal to the rate of change in P_{Sum} at the point on the other side of the minimum. There is a sensible symmetry here, in the slope of P_{Sum}, although other choices could be made.

We will never know where we are on the curve, but a *change* of reliability with H_{PMon} can be quantified. If the heavy line is associated with reasonable balance, it should be noted that any change in reliability on this part of the lower curve has a negative slope, $\Delta P_r/\Delta H_{PMon}$, which is more shallow than -2γ.

If a proposed change has a negative slope which is steeper than -2γ, it means that either we are on the steeper part of the curve to the left of the balanced region, or we must be in the gray area somewhere above the curve. If we are on the curve, we are too far from the minimum, therefore not balanced, and the change is an effective way to improve reliability; if we are in the gray area we are also not balanced and the change is an effective way to get closer to the curve. In either case the initial situation is not well enough balanced, and the new situation will be better, through a change which is itself a balanced change in reliability and availability.

If the proposed change has a negative slope which is more shallow than -2γ, it means that either 1) we are already on the heavy line part of the curve, in which case balance is already

satisfactory and the proposed marginal improvement in reliability is not effective from the point of view of balance alone; or 2) we are in the gray area above the curve, in which case, although the situation may not be well balanced, there exist better options (that is, with steeper slopes) to improve balance than the change which is being proposed.

The balance criterion which emerges can be stated in several ways. The simplest approach tests for a steep negative slope, which is identified with the initial state being out of balance.

Statements of Balance: The test is formulated so that the test *succeeds* when the initial availability and reliability are *not* already sufficiently well balanced. The result of establishing that the initial situation is not well balanced is that the proposed change is justified on the basis of balance alone.

Passing the test means the initial situation is not sufficiently well balanced. The proposed change is then justified from a purely balance perspective.

Failing the test means that the proposed change is not an effective way to improve the balance, either because the initial situation is well enough balanced already, or because better options must exist to improve the balance. Failing the test means the proposed change is not justified from a balance perspective alone.

A. The Quantitative Balance Criterion

The initial availability and reliability are not balanced whenever reliability can still be improved by, 1) performing more on-line PM to improve reliability at a rate which exceeds 2γ per hour of additional on-line maintenance; or 2) performing less on-line PM.

Therefore, the existing situation is *not* balanced whenever:

$$\Delta P_r / \Delta H_{PMon} > 0 \quad \textit{(over - maintained)}$$

Or:

$$\Delta P_r / \Delta H_{PMon} < -2\gamma \quad \textit{(reliability is not good enough, given the unavailability)} \qquad \textbf{Eq. 4-12}$$

B. Statement of Maximum Allowed Unavailability

The maximum number of additional unavailable hours which, when resulting in an *improvement* in reliability, establish the initial condition as not sufficiently well-balanced (with $1/2\gamma = 3500$) is given by Equation 4-13.

The existing situation is not balanced when:

$$\textit{Additional unavailable hours} < \left|\textit{Change in reliability}\right| \times 3500 \qquad \textbf{Eq. 4-13}$$

C. Statement of Minimum Improvement in Reliability

The minimum improvement in reliability which, when accompanied by an *increase* in unavailable hours, establishes the initial condition as not sufficiently well balanced, is given by:

The existing situation is not balanced when:

$$\left|\textit{Change in reliability}\right| \times 3500 \; > \; \textit{Additional unavailable hours} \qquad \textbf{Eq. 4-14}$$

In each case, A, B, and C, the demonstration of initial imbalance justifies implementing the change.

None of the statements A to C is useful unless it is possible to estimate the rate at which reliability will change with a change in PM. From a practical point of view this may not be too difficult to accomplish, because in the context of an a(1) SSC in the Maintenance Rule, there is usually a small number of failures which, it is being claimed, would be prevented from occurring again, if additional PM were performed. This can be the basis for an estimate of reliability improvement.

The appearance of the factor of 2 along with γ is a direct consequence of incorporating the correct relationship between P_{Sum} and P_r. The factor of 2 requires a PM improvement to be achieved with fewer additional unavailable hours than if the factor were absent and a simple tradeoff were managed between ΔP_r and ΔH_{PMon}.

It should also be noted that the proposed criterion does not place any a priori restrictions on existing performance criteria. The criterion does not provide a test for such criteria—that is, an easy way to decide if the performance criteria are already balanced. In fact, it is dubious whether it makes sense to speak of balance in connection with performance criteria, because they are only the short-term means to provide uncertain assurance that the underlying availability and reliability are within bounds. Practical attempts to apply balance to performance criteria seem to end up not addressing the performance criteria at all, but addressing the underlying availability and reliability.

Again, there is no set of "allowable values" of availability and reliability embodied in the proposed balance rule. Instead, balance only arises in the context of a practical change which could be made in PM to affect reliability, for example a specific change in PM to prevent specific failures which have already been experienced.

4.4 Inclusion of Other Basic Events

For simplicity, the only basic events which have been considered so far are the unavailability stemming from the performance of PM, P_a, and a single unreliability basic event, P_r. In practice the following basic events could also be included: repair unavailability, hidden unavailability stemming from hidden failures before they are discovered, and a variety of other reliability basic events such as failure to run, failure to close, and so forth.

For a given component, the various reliability failure modes are mutually exclusive in the sense that a valve cannot fail open and fail closed at the same time, or a motor cannot fail to start and fail to run at the same time. The consequence is that the sum of the reliability basic event probabilities can be used to represent the total effect of reliability of the component on safety. This will be made more precise when we consider the inclusion of risk significance in the balance criterion, but the main point is that we can add the effects of different reliability failure modes.

Does this include the unavailability due to repair time? The repair time is usually included with other unavailable time in a probabilistic safety assessment to constitute a total maintenance unavailability, but is not a significant contribution to the total unavailability when equipment is reliable and significant PM is performed on-line. Nevertheless, the unavailability due to repair time is a simple consequence of failures, and cannot be adjusted to influence reliability. So it is clear that if it is to be included in a balance criterion, the repair unavailability must be counted in with the reliability basic events, and not with the PM unavailability.

What of the hidden unavailability? This is the $\lambda\tau/2$ term in many safety assessments which is used in place of a failure to start probability. The idea is that components can fail in standby with a standby failure rate, λ, during the time they are not being operated. The alternative is to assume that the failure occurs at the demand to function, with a constant probability of failure on demand. The complication here is that the time between surveillance tests, τ, can be adjusted to affect the reliability of the component through the $\lambda\tau/2$ term. Such an adjustment does nothing to alter the failure rate or the expected number of failures during surveillance tests, but it does change the size of this basic event in the PSA model, because it is a change in the standby mission time. The result is that for balance, even though this term is usually viewed as a contribution to unavailability, it is proportional to the standby unreliability, and must be added in with the other reliability basic events on the reliability side of the balance equation.

We are left with the same balance criterion as before, provided the reliability is construed as the sum of all the obvious reliability basic events, plus repair time unavailability, plus the hidden failure term. Represent this sum by P_R instead of P_r:

$$P_R = \Sigma_i P_{ri} + H_{repair}/H_{Req} + \lambda\tau/2 \qquad \textbf{Eq. 4-15}$$

P_a remains the unavailability due to on-line PM, as before.

The balance criterion should be applied using P_R instead of P_r, providing the related changes in reliability can be estimated. Improved PM would probably decrease (improve) λ and H_{repair}, as well as improve all the P_{ri}. Adding these extra terms will increase the numerical value of the slope of P_R over what it was formerly for a single failure mode, and will therefore make it somewhat more difficult to claim that the initial situation is already balanced.

Even with these additions, an important part of the overall balance picture still remains outside the current framework. The omission is that different failure modes (that is, basic event probabilities) can contribute to overall safety to very different degrees. Simply adding them together as in Equation 4-15 to create P_R misses their differing values of risk significance. A

more comprehensive balance criterion would include the Birnbaum factors which provide a weight for each basic event in the calculation of the core damage frequency.

4.5 Inclusion of Risk Significance

The appropriate measure of safety for a more complete balance criterion is the core damage frequency (CDF), as calculated in a probabilistic risk assessment (PRA). If attention is focused on a particular SSC, the CDF can always be written as:

$$CDF = \Sigma_i * a_{ri} P_{ri} + a_a P_a + b \qquad \textbf{Eq. 4-16}$$

This is an expansion of the simple CDF = aP + b form, where P is a basic event probability of interest, a is the remainder when P is factored out of the sum of accident sequence probabilities which contain the event of interest, and b is the sum of all other accident sequence probabilities. Equation 4-16 further expands the aP term using the complete set of reliability basic events $\{P_{ri}\}$ as defined by Equation 4-15, including the hidden failure and repair unavailability terms. P_a is a single unavailability basic event, the same as defined previously.

The parameters a_{ri}, a_a, and b are independent of H_{PMon}, provided H_{PMon} represents the maintenance hours applied only to the chosen component, (that is, it is not directly correlated with the unavailability of other components). In this case, the analog of Equation 4-5 is:

$$dCDF/dH_{PMon} = \Sigma_i (over\ all\ reliability\ modes)\, a_{ri}\, dP_{ri}/dH_{PMon} + a_a \gamma \qquad \textbf{Eq. 4-17}$$

The dependence of the CDF and P_{ri} on H_{PMon} is exactly analogous to what occurred for P_{Sum} and P_r. The result is displayed in Figure 4-7.

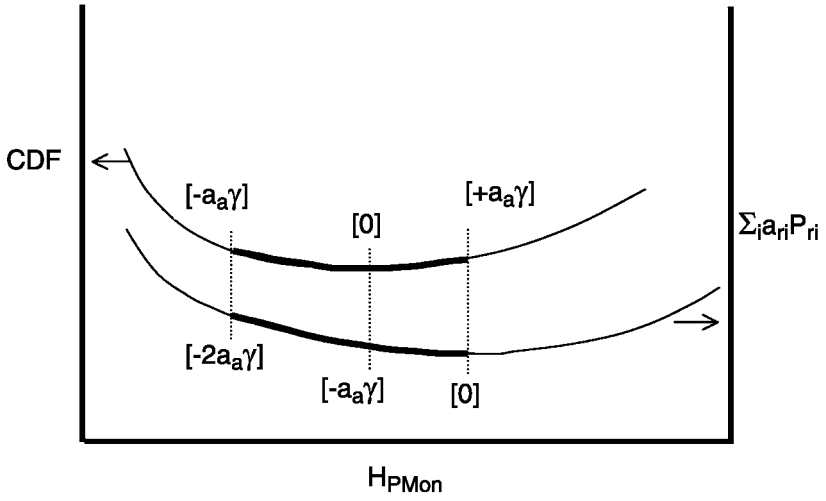

Figure 4-7
Corresponding Gradients on CDF and Σ_i $a_{ri}P_{ri}$

The slopes of the curves shown in square brackets in Figure 4-7 are derived from Equation 4-17, the same way they were before. The analog of the balance criterion, Equation 4-12, becomes:

Unbalanced when:

$$\Sigma_i \text{(over } \underline{\text{all}} \text{ reliability modes) } a_{ri}\, \Delta P_{ri} / \Delta H_{PMon} < -\, a_a\, 2\gamma; \qquad \text{or when } > 0 \qquad \textbf{Eq. 4-18}$$

The Birnbaum parameters a_a and a_{ir} can be replaced using the basic definitions:

$$CDF = a_x P_x + b' \qquad \textbf{Eq. 4-19}$$

Where b' equals the b of Equation 4-16, plus all the other $a_k P_k$ terms not explicitly involving x, and

$$RAW_x = (a_x + b) / CDF \qquad \textbf{Eq. 4-20}$$

whence

$$a_x = CDF\,(RAW_x - 1) / (1 - P_x) \qquad \textbf{Eq. 4-21}$$

or, to an extremely close approximation for hardware basic events, for which $P_x << 1$:

$$a_x = CDF\,(RAW_x - 1) \qquad \textbf{Eq. 4-22}$$

The balance criterion then becomes:

The existing situation is not balanced when:

$$\Sigma_i \text{(over } \underline{\text{all}} \text{ reliability modes)} \frac{(RAW_{ri} - 1)}{(RAWa\ 1)} \times \Delta P_{ri} / \Delta\Delta_{PMon} < -\, 2\gamma; \text{ or when } > 0$$

Eq. 4-23

This is the same as Equation 4-12 except that each failure mode is weighted with the factor $(RAW_{ri} - 1)/(RAW_a - 1)$.

To summarize:

Let us recall that Equation 4-12 expressed the balance test for a single reliability failure mode without regard to the relative risk significance of the failure mode and PM unavailability:

Unbalanced when:

$$\Delta P_r / \Delta H_{PMon} < -\, 2\gamma; \text{ or when } > 0 \qquad \textbf{Eq. 4-24}$$

The simplest way to take account of additional reliability failure modes simply adds them (including repair unavailability and the hidden failure term) without regard to their relative risk significance. The test then becomes:

Unbalanced when:

$$\Sigma_i (\text{over } \underline{\text{all}} \text{ reliability modes}) \quad \Delta P_{ri} / \Delta H_{PMon} < -2\gamma; \text{ or when } > 0 \qquad \textbf{Eq. 4-25}$$

The most comprehensive way to address balance is to account for all of the failure modes, and to include their relative risk significance, so that the test becomes:

Unbalanced when:

$$\Sigma_i (\text{over } \underline{\text{all}} \text{ reliability modes}) \frac{(RAW_{ri} - 1)}{(RAW_a - 1)} \cdot \Delta P_{ri} / \Delta H_{PMon} < -2\gamma; \quad \text{or when } > 0$$

Eq. 4-26

At whatever level the balance test is applied, it reduces to requiring a sufficiently large improvement in reliability compared to the additional unavailability which will be incurred.

4.6 Balance and Supercomponents

When components are included within the boundary of a supercomponent for the purpose of simplifying the modeling in a PSA, various important parameters for the individual components may need to be extracted from the parameters for the supercomponent before the balance test can be employed. These parameters include reliability, availability, and the RAW risk significance parameters.

In principle, the balance test could be applied to the supercomponent itself, provided separate reliability and availability basic events can be identified, which is not always the case, but the choice of potential PM improvements (potentially addressing PM for many components) is likely to be too wide for straightforward application. In addition, there is the risk that a single component becomes the focus, instead of adjustments to several components which might individually need better balance. It is better to apply the balance rule for all types of supercomponents at the level where PM is performed—at the level of individual major components.

The most common kind of supercomponent has a structure which permits the risk significance parameters for individual components to be extracted in a generic way from the parameters for the supercomponent. These supercomponents consist of a logical chain from the point of view of the success and failure of the combination. All the components must function properly to ensure success of this type of supercomponent, and the failure of any one of them causes the failure of the supercomponent. Such a chain has no internal equipment redundancy and could be an instrumentation loop or a typical “train” of components such as a suction valve, a pump, a driver and auxiliaries, a check valve, and a discharge valve. The essence of such an arrangement is that

the failure probability of the supercomponent is equal to the simple sum of failure probabilities of the constituent components, rather than a more complicated function of them. The immediate benefit is that a subgroup of the component basic events can be identified with the supercomponent unreliability basic event, and the rest with the supercomponent unavailability basic event. Denote the supercomponent unreliability basic event probability by P_R, and the constituent component unreliability basic event probabilities by P_{ri}. Then:

$$P_R = \Sigma_i P_{ri} \quad \textbf{Eq. 4-27}$$

The sum has to be taken over all relevant reliability basic events—that is, those which can logically contribute and, as before, includes hidden failure and repair unavailabilities.

The supercomponent PM unavailability event, P_A, is composed of only the PM unavailability contributions:

$$P_A = \Sigma_j P_{aj} \quad \textbf{Eq. 4-28}$$

Where j runs over all the constituent components. Notice that within the supercomponent, all these events are simply added without supplementary weighting.

Equations 4-27 and 4-28 define the unreliability and PM unavailability of the most common type of supercomponent in terms of the constituent unreliabilities and unavailabilities. More complex supercomponents will not permit separation of the reliability and availability basic events in this way, as there will be cross terms which mix the characteristics. *Clearly then, balance for the more complex types can only be addressed via balance for the individual components.*

Individual component basic event FV_{ri} factors can be written in terms of the supercomponent's FV_R factor as:

$$FV_{ri} = FV_R \, P_{ri} / P_R \quad \textbf{Eq. 4-29}$$

The same is true for the FV factors for PM unavailability:

$$FV_{aj} = FV_A \, P_{aj} / P_A \quad \textbf{Eq. 4-30}$$

The RAW factors can be extracted using:

$$RAW_{rk} = RAW_R + FV_R - FV_{rk} \quad \textbf{Eq. 4-31}$$

Or:

$$RAW_{rk} = RAW_R + FV_R (1 - P_{rk} / P_R) \quad \textbf{Eq. 4-32}$$

The same derivation applies for RAW_{ak}:

$$RAW_{ak} = RAW_A + FV_A - FV_{ak} \quad \textbf{Eq. 4-33}$$

$$RAW_{ak} = RAW_A + FV_A (1 - P_{ak} / P_A) \quad \textbf{Eq. 4-34}$$

The balance rule can now be applied for the individual components which comprise the supercomponent, because all the required quantities for the constituent components can be obtained from the input and output of the PSA for the supercomponent—that is, P_R, P_A, FV_R, FV_A, RAW_R, and RAW_A. Notice that all the P_{ri} and P_{aj} are also needed, and are known, because they were the means by which the values P_R and P_A were originally obtained for the supercomponent, Equations 4-27 and 4-28.

4.7 Summary

The balance criterion applies equally to any component type, regardless of the reliability, and regardless of whether the component requires more or less maintenance unavailability than other component types, or components of the same type. The criterion relies on the correct relationship between reliability, availability, and safety, which leads it to be two times more conservative, in its demand for reliability for a given amount of unavailability, than a simple trade-off between reliability and availability would indicate. The criterion also permits a range of different reliability/availability values to be classed as balanced, and does not focus relentlessly on achieving the optimum reliability. It easily encompasses any or all of the failure modes of a component, and has been extended to account for differing values of risk significance of all the failure modes.

On the negative side, there is some arbitrariness in the choice of the balanced region, and there may be difficulty in deciding what the reliability pay-off of a PM change is likely to be. This is very important, because there is nothing in the balance rule which defends against an incorrect presumption of the effect of a PM change. For example, if a component is over-maintained, and its reliability is thereby adversely affected, a user who is unaware of this could incorrectly assume that additional PM and unavailability will improve the reliability. The balance rule would then advocate the change, incorrectly, provided the change itself passed the test, based on the incorrect assumptions. The rule therefore does nothing to diminish the responsibility of nuclear power plant licensees to understand the strengths and weaknesses of their existing PM programs. In fact, the rule can only be used correctly when the user has established the technical basis of the PM program, and is confident that a proposed change in PM will, indeed, improve the reliability.

5
RELIABILITY AND AVAILABILITY MONITORING

5.1 The Reliability Context for Monitoring Failures

Reliability monitoring refers to the process of estimating the failure rate, or the probability of failure-on-demand as new failures occur. Reliability monitoring and improvement are connected by the fact that when the reliability of critical SSCs is seen to deteriorate, something may need to be done to improve it. Even though monitoring does not prevent the observed failures from occurring, it should, of course, lead to actions which prevent most of the failures from recurring. Programmatic measures, such as developing PM programs with a strong technical basis, and adopting a comprehensive component reliability improvement program as exemplified in the INPO AP-913 proposals [17], are the best proactive means to achieve high reliability and exert a strong influence on failure rates.

In practice, the effectiveness of reliability monitoring is constrained by the large uncertainties which attend estimates of failure rates based on small numbers of failures. When these uncertainties are at least a factor of two or three in each direction, it is difficult to discern any but the most egregious departures from the baseline. In spite of this, there would seem to be a clear value in attempting to monitor reliability if such large departures from expected values can be flagged automatically. Unfortunately, uncertainties as small as a factor of two are difficult to achieve within the constraints of monitoring a small number of components over a period as short as two years, which is typically the time duration implemented in the Maintenance Rule, 10CFR50.65.

It is worth questioning the above objectives of detecting significant changes in failure rates over short time periods, in relation to Section 3.1.2, Nine Reasons Why Failure Rates are Constant in Time, where it was stated that we do not expect to see large departures from a constant failure rate for critical components that are well maintained. Experience has also shown that in general, failure rates of key equipment in nuclear power plants have indeed remained more or less constant over several decades, except for improving trends which correspond to the elimination of early life design-related failures, and the implementation of improved PM and equipment reliability programs, as the industry has matured.

The answer to this question is that the achievement of constant or improved reliability over long time periods in fact conceals many short-term fluctuations. We expect such fluctuations to be of two types: 1) non-deterministic, in that failures are a statistical phenomenon so that measures of failure rate based on few failures will exhibit a large variance which is *purely without any causal basis*, and 2) deterministic, in that the quality of PM and other causal influences can change over the short term, and new failure mechanisms can be encountered which are not protected by

existing PM activities. Both types of fluctuations have serious implications for nuclear power operations, but only the non-deterministic type can be relied on to average to zero over the long term without deliberate corrective action by plant personnel.

The non-deterministic fluctuations are only a serious matter because they cannot easily be distinguished from the deterministic type, which require active intervention to prevent them from becoming permanent features. Active participation at all times is therefore required to detect such fluctuations, and to improve maintenance or take other measures to ensure that the failure rate is not permanently affected. It is imperative to distinguish the non-deterministic fluctuations from the deterministic ones to avoid reliability improvement efforts being driven by irrelevant issues. An even worse outcome might result if unnecessary PM actions are added to PM programs, making failure rates deteriorate further through the introduction of maintenance error. There is potential for magnification of this problem when conclusions are drawn erroneously from failures experienced on relatively few components, and applied to much larger populations of similar components.

These points underline the need to understand the technical basis for the PM program. When failures occur, it is useful to know whether the current PM program is supposed to have prevented them or not, and to estimate the potential for improved PM tasks to be more effective. If current PM tasks and intervals are adequate to have prevented failure mechanisms of the type observed, it most likely means that the PM tasks are not being performed properly. If the PM program appears to be technically inadequate to provide protection against those failure mechanisms, and cost-effective PM improvements *cannot* be made, then the failures must be seen as part of the fluctuating background of failures which have to be accepted on an ongoing basis, or prevented by a design modification. Ultimately, only this level of understanding, combined with good engineering judgment, can ensure that reliability monitoring over the short term does more good than harm.

On a philosophical level, the failure rate and probability of failure-on-demand are no more than descriptions and measures of an aggregate of phenomena that defy more exact analysis. For reliable equipment, in the short term there cannot be many failures to analyze, and short-term statistical uncertainties become overwhelming. There is therefore a practical limit on the shortness of monitoring period and the smallness of the population being monitored, below which it is not possible to determine the failure rate with any practical utility. Therefore, from a theoretical point of view as well as in a real practical sense, the failure rate has no meaning as an instantaneous quantity. We should not approach this topic believing that there is a “true” value of the failure rate at any one point in time. It is inherently a property representing average behavior over a fairly long period of time.

The failure rate depends upon the number of failures and the number of component years of operation (see Section 1.9). So providing the population being monitored does not change, and successive monitoring periods have the same duration, monitoring the number of failures is a surrogate for monitoring the failure rate. However, as discussed above, the number of failures will be variable, and trends impossible to interpret, unless as much or more attention is paid to the nature of the failure mechanisms as to the number of failures. In this light, exceeding a

performance criterion placed on the number of failures is an appropriate trigger for more in-depth engineering investigation.

If the probability of false alarms (that is, statistical fluctuations) was properly included when setting up a performance criterion, numbers of failures which do *not* exceed the criterion can be viewed as most likely to be statistical fluctuations. In practice, it is better if all failures of a critical or significant category SSC can be investigated for the possibility that something new is occurring, or something has deteriorated in the operating or maintenance environment of the equipment, with the understanding that the conclusion could be that nothing new is occurring, and that occasional failures of this kind are expected.

5.2 Limitations of Poor Statistics

Section 6 (particularly Sections 6.3 and 6.4) provides the means to calculate confidence intervals on both failure rates and probabilities of failure-on-demand. The conclusions are closely analogous in each case. They can be demonstrated using a simple expression (Equation 6-4) for the upper one-sided confidence limit at confidence level (1-α), when there have been no failures at all in T component years of exposure:

$$\lambda_{upper,\,one\ sided} = -(\log_e (1-\alpha)\times)/T \qquad \textbf{Eq. 5-1}$$

This "no failure" case is where we expect to find the worst manifestations of the uncertainties. A confidence level of (1-α) means that the probability is (1-α) that the true failure rate is between 0 and $\lambda_{upper,\ one\ sided}$. If the confidence limit is calculated for a 50% confidence level (α=0.5), the actual value has a 50% chance of being both above and below this level. The upper limit then becomes $\lambda_{upper,\ one\ sided} = 0.693/T$. It is equivalent to assuming that about 0.7 failures have occurred, even though the actual number of failures is zero.

As an example, if the failure rate is actually about 0.35 per year, which is high for most nuclear plant equipment, there is only about a 30% chance of a failure in one year ($1-e^{-0.35} = 0.295$), so we would expect to have no failures in 2 years out of every 3. If we monitored with a 1-year monitoring duration, such zero failure years would individually suggest a failure rate of 0.693 per year, twice the real value.

For more realistic (that is, lower) values of failure rate (certainly typical for safety system equipment) there will be a much larger proportion of 1-year monitoring periods in which there are no failures, and for each of these the ratio of estimate to real value will be higher than a factor of 2. For example, when the failure rate is really 0.05 per year, the zero failure estimate of 0.693 is nearly 14 times higher than the real value—and this is when the reliability during the monitoring period has been as good as it could have been. Consider also that there is as much chance that the real value is even higher than the estimated 0.693 as there is that it is lower! When the required confidence level is higher than 50%, the situation becomes much worse.

A different example, explained in detail in Section 6.3, shows the value of working with more failures, and uses two-sided confidence limits. Two failures are observed in a group of pumps

over a cumulative operating period of 3 pump years. The mean estimate for λ is 2/3 = 0.67 failures/year. The 5% and 95% confidence limits, explained in Section 6.3, are:

$$\lambda_{.95} = 2.1 \text{ per year} \quad \textbf{Eq. 5-2}$$

$$\lambda_{.05} = 0.12 \text{ per year} \quad \textbf{Eq. 5-3}$$

Notice that the ratio $\lambda_{.95} / \lambda_{.05} = 17.5$, which is a very wide range of uncertainty.

If the statistics are improved by making observations over a much longer time but at the same rate, so that 14 failures are observed in 21 pump years, the estimate for the failure rate remains the same at 14/21 = 0.67, but the confidence limits, explained in Section 6.3, are:

$$\lambda_{.95} = 43.8/42 = 1.04 \text{ per year} \quad \textbf{Eq. 5-4}$$

$$\lambda_{.05} = 16.9/42 = 0.40 \text{ per year} \quad \textbf{Eq. 5-5}$$

The ratio $\lambda_{.95} / \lambda_{.05}$ is now 2.6, so the uncertainty is much less than before.

The conclusion from these simple examples is that it is not possible to monitor the reliability of a single SSC, or even a small group of reliable SSCs, in a meaningful way over monitoring periods as short as 1 or 2 years. Section 5.5 gives a brief explanation of the best way to set performance criteria on failure rates and probabilities of failure-on-demand, which recognizes these difficulties.

The situation is not necessarily as difficult when monitoring unavailability. Unavailability caused by repair time following functional failures is often a small contribution to the total unavailability, compared to the contribution from preventive maintenance. The former should be expected to be even more variable than the number of failures, because the uncertainty derives from the statistics of failures and also from the distribution of repair times. However, the small size of this contribution diminishes the effect of its variability. Fortunately, the generally larger PM contribution derives mostly from scheduled events involving the repetition of standard activities, so is expected to be much less variable.

5.3 Benefits of Maintenance Rule Monitoring

If it is basically impossible to make sense out of the number of failures from individual components monitored over short time periods, we should ask if Maintenance Rule monitoring can serve a useful purpose. Most monitoring in the Maintenance Rule is directed at an SSC which represents an aggregate of equipment, such as a redundant cooling train. Broadly speaking, the failure rate of such a train will be the sum of the failure rates of the items of equipment it is made of.

If a typical train consists of 10 or more major components such as suction and discharge valves, check valves, a pump driver, coupling, pump, lubrication subsystem, and controls, it will have a

failure rate much higher than normally associated with a single reliable component. However, the train must still be very reliable overall, so the failure rate cannot usually be higher than the range 0.1 to 0.5 failures per year. Such a train therefore does not escape the uncertainty limitations discussed in the previous section. Therefore, monitoring can at best serve to flag the possibility of an underlying problem, as described in Section 5.1.

Probably the most valuable aspect of Maintenance Rule monitoring to performance criteria has been the increased focus it has brought to the identification of functional failures, the sharing of information about them, and the emphasis on preventing the occurrence of repetitive failures. Real benefits will accrue from the exercise of good engineering judgment in cause evaluations performed when performance criteria are exceeded. A utility cannot afford to rely on the monitoring results and spend few resources on cause evaluation, because the trade-off is far from linear. Paying too little attention to cause evaluation and the effectiveness of the PM program will not merely forego the benefits which might have been obtained, it will result in the frequent mis-identification of appropriate corrective actions driven by the widely fluctuating numerical results of monitoring.

5.4 Performance Criteria, PSA, and the False Alarm Rate

5.4.1 Unavailability

The motivation to monitor unavailability has much to do with the desire to eliminate procedural and logistical causes of unnecessary unavailability, and to be sure that PM activities continue to be carried out efficiently. A possible cause of unnecessary unavailability occurs when, for no good reason, PM work on equipment is not initiated until a day or more after it has been tagged out of service. Unavailability of safety SSCs is vitally important to the safety of the plant, because it influences the core damage frequency on an equal footing with unreliability. Because it is largely composed of PM unavailability, in a very real sense it is the price that must be paid to achieve high reliability. The relationship between availability and reliability is the subject of Section 4.

The value of unavailability being monitored will fluctuate above and below its average value. There is a need to prevent unwanted exceedances of the unavailability performance criterion, which can arise because of the random contribution to these fluctuations. It is therefore important to set the performance criterion for unavailability at some point above the average value.

Consider three key unavailability quantities: 1) the rolling average over ~ 1 cycle; 2) the PSA value that in principle is an average over very many cycles; and 3) the performance criterion. To prevent inadvertent exceedances, the performance criterion should obviously be larger than most 1-cycle rolling averages—that is, close to the upper limit of the expected range of *normal* performance. This upper limit of the SSC's range of performance has to be judged in light of an anticipated uncertainty, which is probably equal to or greater than the average unavailability itself. We should therefore expect the performance criterion to be around two times the average unavailability, or even higher. If the performance criterion is set much lower than this, there will

be too-frequent exceedances due to random effects, which do not reveal anything untoward happening in the PM program.

But the PSA value is supposed to represent the average SSC performance, and should approximate the average taken over many monitoring cycles. It should therefore lie within the SSC's range of *normal* performance. Therefore, the performance criterion should be set significantly higher than the PSA data value. It is equally obvious that the PSA value should not be a standard to which SSC monitoring performance is compared, because a significant number of the monitored values are bound to lie higher than the average, even for a quantity which may have a significant "tail" on the high side of the distribution, which makes the average much larger than the median.

As a practical matter, the performance criterion for unavailability could be the historical mean at the plant, or the PSA data value, plus one or two standard deviations of the historical data. If you have no indication of the historical standard deviation, it is reasonable to double the average value. In particular, do not use the historical average value as a performance criterion, at least unless you are sure the performance will improve as a result of recent programmatic activities. If historical values of unavailability include periods when there were abuses which caused large amounts of unnecessary unavailability, it will be necessary to estimate a reasonable correction factor, as the future will not resemble the past.

From the point of view of added risk, when the performance criterion is set above the PSA mean value, the risk increase of operating that SSC *permanently at, or higher than,* the limit of the performance criterion is:

$$\frac{The\ increase\ in\ CDF}{Baseline\ CDF} \geq \frac{FV \times (PC - P_{PSA})}{P_{PSA}} \qquad \textbf{Eq. 5-6}$$

where CDF means the core damage frequency, FV means the Fussel-Vesely risk significance value, PC represents the performance criterion, and P_{PSA} is the PSA data value. The whole idea of setting a performance criterion is that the unavailability will not be permitted to remain so high on a permanent basis.

5.4.2 Reliability

As with unavailability, a performance criterion on unreliability or failure rate has to contend with expected random fluctuations which are part of normal behavior, and which therefore do not portend ill health of the PM program. Once again, the performance criterion should be set higher than the long-term average value, which should be equal to the PSA value. However, unlike the unavailability, where the scale of uncertainty is best represented by the historical variations, the statistical distribution of the number of failures should follow a Poisson law for a failure rate model, and a binomial distribution for a constant probability of failure-on-demand model. This enables the false alarm rate associated with a performance criterion on the number of failures to be found quite easily. It also establishes the connection between the PSA expected value and the performance criterion, without irrationally setting the performance criterion to be equal to the

PSA value. In the following discussion, be aware that over 100 performance criteria may be set for purposes of monitoring reliability in an average plant's Maintenance Rule program. A false alarm rate larger than 10% to 15% may be damaging to PM program development.

The number of failures can be monitored as a surrogate for monitoring failure rate, provided the population of components being monitored is held constant between successive monitoring periods of constant duration. For the probability of failure-on-demand, it follows that there is no benefit in tracking the number of demands as well as the number of failures, unless the number of demands varies dramatically from one monitoring period to another. If the number of demands is not tracked for a given SSC, it is advisable to reexamine this assumption if there is a sudden increase in the number of failures.

In a monitoring period of duration T, or in n test demands, the expected (that is, average) number of failures is λT or $P_f n$, respectively, when average values are used for λ and P_f. Assuming the PSA quantification is reasonably up to date and valid, these values can be taken from the plant PSA. The Poisson and binomial distributions, respectively, give simple estimates of the probability that 0, 1, 2, 3, … failures will actually occur over the period T (or in n demands), rather than the expected number. For active equipment, the chance of one failure is usually at least a few percent in a 2-year period, and may be significantly higher for less reliable SSCs.

First, estimate the probabilities of different numbers of failures, using the equations below (also described in Section 6.7, Equation 6-15, and Section 6.8, Equation 6-17).

Poisson Distribution: The probability, P(r), of observing exactly r failures in time T when λ is the failure rate—that is, when you expect λT failures on average—is:

$$P(r) = (\lambda T)^r \times e^{-\lambda T} / r! \qquad \textbf{Eq. 5-7}$$

r! is the factorial function: r! = r(r-1)(r-2) …..3 x 2 x 1 (note: 0! = 1! = 1)

For example, if $\lambda T = 0.1$ failures are expected, the probability of actually getting 1 failure (r=1) is:

$$(0.1^1 \times e^{-0.1})/1 = 0.1 \, x \, e^{-0.1} = 0.1 \, x \, 0.905 = 0.0905, or \, 9\% \qquad \textbf{Eq. 5-8}$$

If you had 100 SSCs like this one, on average every monitoring period there would be 9 of them which experienced exactly one failure—and even one failure is ten times more than the expected value of 0.1.

The probability of getting 2 failures when $\lambda T = 0.1$ failures are expected is:

$$(0.1^2 \cdot e^{-0.1})/2 = 0.01 \, x \, e^{-0.1}/2 = 0.005 \, x \, 0.905 = 0.0045 = 0.45\% \qquad \textbf{Eq. 5-9}$$

In this case, you might choose a performance criterion (PC) of 0 failures because the chance of exceeding 0 failures by chance is only $1 - e^{-0.1} = 1 - 0.905 = 0.095 = 9.5\%$. But having 9 or 10

"accidental" exceedances for every 100 PCs of this kind each monitoring period seems to be at best a marginally acceptable false alarm rate. Given the fact that regulators seem to prefer not to see zero failures as a performance criterion, it is probably wiser to set the PC=1, and benefit from a very small false alarm rate of about half a percent.

Binomial Distribution: The probability, $P_n(r)$, of observing exactly r failures in n trials, where P_f = the probability of failure-on-demand, is:

$$P_n(r) = n!\ P_f^{\,r}\ (1 - P_f)^{(n-r)} \cdot / (r!(n - r)!) \qquad \textbf{Eq. 5-10}$$

Inserting specific values for r:

$$Pn(0) = (1 - P_f)^n$$

$$Pn(1) = nP_f \bullet (1 - P_f)^{n-1}$$

$$Pn(2) = n(n - 1) \bullet P_f^{\,2} \bullet (1 - P_f)^{n-2} / 2$$

$$Pn(3) = n(n - 1)(n - 2) \bullet P_f^{\,3} \bullet (1 - P_f)^{n-3} / 6 \qquad \textbf{Eq. 5-11}$$

For example, if a Terry turbine failure to start probability is 0.03 f/demand, and it is tested every 3 months, what is the chance of experiencing exactly 0, 1, and 2 failures in a 24-month monitoring cycle?

Assume there are 8 tests in 24 months, so n = 8, P_f = 0.03. Then:

$$P_8(0) = (1 - 0.03)^8 = 0.97^8 = 0.78 = 78\%$$

$$P_8(1) = 8x0.03x0.97^7 = 0.19 = 19\%$$

$$P_8(2) = 8x7x0.03^2 \bullet 0.97^6 / 2 = 0.02 = 2\% \qquad \textbf{Eq. 5-12}$$

To set a performance criterion (PC) on the number of failures in this case, you need to note that there is a 78% + 19% = 97% probability of getting either 0 or 1 failure. Therefore there is just a 3% chance of getting 2 or more failures by random fluctuations. If you set PC=2, and actually observed 2 failures, you need to consider whether the 2% chance of getting this result by random fluctuations makes you believe this is what has occurred, and that there is nothing wrong with the PM program, or whether it is much more likely that the failure rate is really much larger than the value assumed. It seems more rational to choose the latter, so 2 failures could represent an exceedance. We therefore would not want to set the PC=2. If you set the PC=1 and actually observed 1 failure, there is a 19% chance that this result occurred because of random fluctuations, through no fault of the PM program. The chance of getting two failures is about an order of magnitude less than the chance of getting one.

This kind of reasoning can provide a probability context for the choice of a rational performance criterion. To add the false alarm rate to this decision process, just note that in this example, when the PC=1, only 3 out of 100 equivalent PCs (the above 3% chance of exceeding 1 failure) will exceed the PC by chance—that is, the false alarm rate is only 3%, which is very satisfactory. If you had selected zero failures as the PC (regardless of other objections to a choice of zero), the probability of exceeding the PC by chance is 22% (=100-78). This means that your PM program development could be continually subject to misdirection caused by frequent inadvertent exceedances, which might delay or defeat convergence to a superior program over time.

It is not possible to provide a detailed guideline to implement this approach, because many factors need to be integrated to come to a final decision. Even within the narrow boundaries of the probability considerations, two factors must be weighed together:

1. The probability of getting more than PC=N failures. It must be small enough to make you believe that, given this result, the most likely reason is that the failure rate has increased. It must also be small enough to limit the false alarm rate to an acceptable value.

2. The *relative* values of the probabilities of getting PC and PC+1 failures. You would like to have a sharp decrease in the probability of getting an exceedance—that is, N=PC+1 failures, compared to N=PC.

For additional discussion of these options see [3], including the case of alternating equipment between running and standby using a single performance criterion, and averaging over trains to set criteria that are more restrictive than using one failure per train—for example, setting a performance criterion of two, or three failures between four trains.

6
RELIABILITY DATA

6.1 The Practical Context of Reliability Prediction

Quantitative reliability data includes information which summarizes a group of failures, such as the following: 1) the failure rate, or the probability of failure-on-demand; 2) the useful or minimum life—that is, the age at which the first wear-out failures occur; 3) the characteristic life—that is, the age at which 63% of the wear-out failures have occurred; 4) the age at each failure—that is, the failure times; and 5) the fraction of failures which have specific failure modes or causes. It is good to keep in mind that the failure rate and probability of failure-on-demand are no more than descriptions and measures of an aggregate of phenomena that defy more exact analysis. This description may be more or less adequate, but we should not approach this topic believing that there is a "true" value of the failure rate at any particular point in time. The failure rate in time and the probability of failure-on-demand are the most common of the quantitative parameters.

6.1.1 Sources of Data

The Maintenance Rule (10CFR50.65) and probabilistic safety analysis both provide opportunities to obtain reliability data. The value of the Maintenance Rule is that it provides a fairly uniform process for putting the spotlight on repetitive failures from the same cause, increases the knowledge and use of industry operating experience, and promotes finding early solutions to new failure mechanisms on an industrywide basis. The focus on preventing repetitive failures is especially important to achieving the goals of preventive maintenance and to keeping failure rates constant over the long term. Although the Maintenance Rule will provide information to assist in quantifying failure rates (see below), the Maintenance Rule program itself has no direct capability to supply quantitative failure rates, because it mostly monitors small numbers of like components over short periods of time (~2 years). Maintenance Rule programs therefore do not normally calculate or trend failure rates.

On the other hand, PSA is a practical source of failure rates for many active components that can be found in standby safety systems. Most plants improve their failure rate quantification by incorporating plant-specific information into the PSA every few years. Consequently, PSA personnel at a plant are usually the best source of knowledge on the status of plant-specific failure rates and of expertise in the methods used to update them. The Maintenance Rule has improved the recognition and timely reporting of functional failures, which should improve the quality of data included in plant-specific updates for PSAs, and will help to speed the process of performing the updates. Almost all critical components, both active and passive, are included in the Maintenance Rule failure reporting process.

Currently, apart from individual PSA studies and some NUREG reports on specific components, the only consistently high-quality databases for failure rates of nuclear plant equipment are three foreign databases. These are the European Industry Reliability Data System, EIReDa (formerly EuReData) [18], the TUD system from Scandinavia (formerly ATV) [19], and the ZEDB database from Germany [20]. Each of these data systems publishes data collected from a significant number of nuclear power plants in a format which is readily usable, using the statistical methods described in this section. The NERC-GADS (North American Electric Reliability Council-Generating Availability Data System), and EPIX (Equipment Performance Information Experience, formerly NPRDS), are failure event databases and do not currently provide preanalyzed failure rate information. The result is that it requires a considerable investment of time to extract the numerator of a failure rate from these event databases, and population information for the denominator may have to be found from other sources.

6.1.2 Improving Failure Rate Estimates Using Additional Data

Sections 6.2 to 6.16 describe a number of techniques used to analyze and manipulate failure rate data. Several of these techniques are used to improve a given sample of data when additional data becomes available. If both samples of data are available in the form of raw information on the number of failures experienced in a certain time, or in a given number of demands, then the numbers of failures and so forth can obviously simply be added together to create new estimates of the failure rate. What usually happens, however, is that the existing knowledge of the failure rate is in the form of a statistical distribution over a range of possible values of the failure rate, and the original numbers of failures and other raw data are not available. The new raw data (that is, in terms of the number of failures experienced in a certain time, or in a given number of demands) then has to be combined with the previous ("prior") failure rate distribution.

The most common application of these techniques is to improve generic data on the failure rate (obtained from nuclear industry sources, or from outside the nuclear power industry), with a local sample of recent failures more specifically representing the SSC in question. The need to do this is fairly evident—the generic data probably contains more statistical evidence, but its applicability to your SSC may be questionable. The local data is especially relevant, but it is unlikely to contain enough failures to be statistically meaningful by itself. Hence the need to combine the two sources. This process is accomplished by the Bayesian Updating procedure.

Because the plant-specific data (that is, the "local" data or the "new" data) is limited statistically, it is often not possible to know whether it represents a run of good luck or bad luck, or the emergence of a potential problem, and whether the reliability it suggests will turn out to be applicable over the long term. The Bayesian Updating procedure automatically takes care of the "strength" of the influence of the new data over the prior generic data, depending on the displacement of the medians between the two samples, and especially on the variance of the samples. Of course, the procedure can also be applied to improve any sample of data with the addition of any other data, as might be done when updating earlier estimates of the plant-specific failure rate with new plant-specific information.

6.1.3 Use of Generic Data

The following sections provide fairly detailed technical methods with which to address the generation or modification of failure rates or probabilities of failure-on-demand in the light of old and new data of various kinds. It remains only to indicate which of the methods would normally be used in certain situations. In practice, the choices are almost always very limited. Using plant-specific information has the sole advantage that you probably have detailed knowledge about the key parameters discussed in Section 1—that is, the PM program, the duty cycle, and service conditions, and that these factors are appropriate for your application. Generic data will most likely be superior in a strict statistical sense, but using it may require you to forego knowledge of these key parameters. The methods described here can be applied regardless of whether the data sources are generic or plant-specific. Therefore, the most important issue in seeking and combining data sources is the above concern over homogeneity of the application.

For many nuclear plant components, generic nuclear plant sources may be reasonably homogeneous in PM program, duty cycle, and service conditions because of the restricted application of the component types. For example, large, complex equipment such as pumps, motors, and medium- or high-voltage breakers will most probably receive a reasonable level of PM simply because of the high cost of repairing the equipment when it fails. Many components may also have reasonably similar service conditions. For example, charging pumps and instrument air compressors are likely to be positioned inside clean, air-conditioned buildings. Likewise, for many components, the duty cycle category will require only deciding if equipment is normally in standby, or is normally operating.

Precise statements about what constitutes a high or low duty cycle, and severe or mild service conditions, are provided for 60 major component types in the EPRI PM Basis database in the Definitions form, also accessible by clicking on headers in the Source form. It can be assumed that critical components will usually have a comprehensive PM program because of their functional importance.

The likely needs of the data analyst are 1) to combine different generic sources of data, each in the form of a given distribution over the parameter of interest; and 2) to update a generic or plant-specific distribution on the failure rate, or on the probability of failure-on-demand, with plant-specific data on numbers of failures, or failure times (failure rate only). The presumption should be that if the available data sources include PSA data, you need to carefully consider using the PSA data, because it may well be the most applicable of all the sources, especially if it has already been updated with plant-specific data.

Where no data is available for the equipment in question, you need to identify another component that shares design features which make it likely that its failure rate could be used as a surrogate. You can proceed to use the surrogate component data, but it may be necessary to modify the results using factors that account for remaining design differences. In this case, such factors could be derived from the EPRI PM Basis database [5, 6].

The remainder of this section and Section 6.2 give a guide to the use of the equations which follow in Sections 6.3 through 6.16. Some of the terminology used may only become clear after the sections containing the relevant equations have been read.

Combine generic data sources, providing they meet the applicability requirements discussed above, before updating the result with more recent plant-specific data which may also be available. Use Equation 6-19 to combine separate homogeneous generic distributions. The symmetry of the right-hand side of the equation shows that it does not matter which distribution you consider to be the prior and which the likelihood. Use Equation 6-19 sequentially to combine more than two distributions—that is, use the posterior obtained from combining the first two sources as the prior for combining the next. This procedure necessarily involves numerical analysis.

Do not assume that it is reasonable to combine sources of generic data just because the above procedure makes it possible. If one source has a much smaller variance (is much narrower) than another, it may be better to use the one with the narrower distribution on its own, because it will almost certainly include more failure experience derived from more homogeneous plant conditions. However, this is by no means a golden rule, because the narrower distribution may represent a set of conditions that is not a good match to the conditions appropriate for your application. If you do not know the application conditions for either distribution, you may benefit from using the wider distribution alone, in order to avoid too much specificity in the generic data. If the generic sources are not markedly different in this way, it is probably best to combine them all, providing you are sure they are truly independent sources.

6.2 Updating a Prior Distribution With New Failure Data

In general, to include new data along with prior information on failure rates, use Equations 6-11 and 6-12 or their self-conjugate equivalents with the following procedure:

1. Decide whether the new data is generated from an exponential, Poisson, or binomial statistical process. Calculate the likelihood of getting the new data with this process, using Equation 6-14, 6-16, or 6-18.

2. If the new data is generated from a Poisson or binomial statistical process you may use a gamma or beta self-conjugate prior, respectively. In that case, determine the parameters of the prior, by matching the mean and variance with those of the given prior distribution, as described in Section 6.13. Modify the parameters of the gamma or beta priors to obtain the posterior distributions using Equation 6-27 or 6-29. This procedure requires only a little algebra.

3. If the new data is not from a Poisson or binomial statistical process, or if you do not wish to use self-conjugate priors, use Equations 6-11 and 6-12 directly to obtain the posterior distribution. This procedure necessarily involves numerical analysis.

4. Choose a representative point estimate from the posterior distribution, such as the mean. You may need to calculate the mean if the posterior is not a standard distribution. Calculate

confidence bounds using Equations 6-30 and 6-31, or by using Equations 6-33 and 6-34 for a gamma distribution.

The following sections contain sufficient technical information and examples for the user to make the best use of all the quantitative failure data likely to be encountered, with a view to obtaining the best possible values of the current failure rate. These sections provide a fairly comprehensive set of tools usable by those who are not expert in the field of reliability. For the first-time reader, there is a benefit in reading through these sections in numerical order before attempting to use them individually.

Although the mathematics may appear quite complex in some parts, the equations provide all the essentials necessary to manipulate failure data. The accompanying text explains the equations to the extent that non-mathematicians and non-statisticians should have little difficulty in using the methods without further guidance. Wherever further development may be necessary to take advantage of more advanced methods, the text provides a note to that effect.

Statistical tables needed to evaluate any of the equations are presented in Section 7, in notation consistent with that used in the equations. When other compilations of statistical tables might differ in definition, a note is provided in the text.

6.3 Constant Failure Rate Over Time

A failure rate over time summarizes the number of failures experienced over a period of time. The actual number of failures can obviously vary from one occasion to another even when the same time period is involved, and is distributed according to a Poisson distribution, to be described in Section 6.7. The usual measure, λ, for the failure rate, does not require knowledge of the individual times to failure, but only the total number of failures, N_f, in the cumulative number of component years, T, in the operating environment:

$$\lambda = N_f / T \qquad \textbf{Eq. 6-1}$$

The upper (two-sided) confidence bound on λ, λ_{upper}, at a confidence level of $(1-\alpha)$, is:

$$\lambda_{upper} = \chi^2_{1-\alpha/2} (2N_f + 2) / 2T \qquad \textbf{Eq. 6-2}$$

The lower (two-sided) confidence bound on λ, λ_{lower}, at a confidence level of $(1-\alpha)$, is:

$$\lambda_{lower} = \chi^2_{\alpha/2} (2N_f) / 2T \qquad \textbf{Eq. 6-3}$$

A confidence level of $(1-\alpha)$ means that the probability is $(1-\alpha)$ that the true failure rate is between λ_{upper} to λ_{lower}. So, if the confidence level is 90%, $\alpha/2 = 0.05$. The notation used here corresponds to that used in most tabulations of the χ^2 (chi-squared) distribution, in which α is the area under the distribution of χ^2 from 0 to χ^2. The χ^2 distribution can be found in Table 7-1, and in most statistical texts or compilations, but be aware that some statistical tables tabulate the

complement of this quantity—that is, the area from χ^2 to 1, so check the definition. In Table 7-1, select the column with value ε equal to $\alpha/2$ or $1-\alpha/2$, and read the confidence limit using the row labeled with $\nu = 2N_f$ or $2N_f + 2$.

Example: Two failures are observed in a group of pumps over a cumulative operating period of 3 pump years (suppose 2 pumps over 1.5 calendar years). The estimate for λ is 2/3 = 0.67 failures/year. For the confidence bounds, look up the value of $\chi^2_{0.95}$ (6) (=12.6) for the upper bound, and $\chi^2_{0.05}$ (4) (=0.711) for the lower bound. Then:

$$\lambda_{upper} = 12.6/6 = 2.1 \text{ per year}$$

$$\lambda_{lower} = 0.711/6 = 0.12 \text{ per year}$$

Notice that the ratio $\lambda_{upper} / \lambda_{lower} = 17.5$, which is a very wide range of uncertainty.

If the statistics were improved by making observations over a much longer time, so that 14 failures were observed in 21 pump years, $\chi^2_{0.95}$ (2x14+2) becomes 43.8 for the upper bound, and $\chi^2_{0.05}$ (28) becomes 16.9 for the lower bound. Then the estimate for the failure rate remains the same at 0.67, but:

$$\lambda_{upper} = 43.8/42 = 1.04 \text{ per year}$$

$$\lambda_{lower} = 16.9/42 = 0.40 \text{ per year}$$

Notice that now the ratio $\lambda_{upper} / \lambda_{lower} = 2.6$, so the uncertainty is much less than before. Equation 6-2 for the upper confidence limit can be used even when there have been no failures at all ($N_f = 0$). The lower limit is then zero.

A different prescription for the upper limit is often used when there have been no failures, by quoting the upper *one-sided* confidence limit which has the value:

$$\lambda_{upper,\, one\ sided} = -(\log_e (1-\alpha)\cdot)/T \qquad \textbf{Eq. 6-4}$$

In nuclear plant practice, as discussed in [21], this estimate is usually calculated for a 50% confidence level, so that the actual value has a 50% chance of being both above and below this level. Equation 6-4 then becomes $\lambda_{upper,\, one\ sided} = 0.693/T$. It is clearly equivalent to assuming that about 0.7 failures have occurred, even though the actual value is zero. See also [22].

6.4 Constant Failure Probability on Demand

The number of failures on demand follows a binomial distribution, described in Section 6.8. The estimate for the probability of failure-on-demand, P_f, is simply:

$$P_f = n / N_d \qquad \textbf{Eq. 6-5}$$

where n is the number of failures on demand, and N_d is the cumulative number of demands on the group of components.

The confidence limits for this distribution are given by solutions to the following equations:

$$P_{f\,lower} \text{ is the } p \text{ value which satisfies : } \sum_{i=n}^{N_d} C^{i}_{Nd}\, p^{i} (1-p)^{Nd-i} = \alpha / 2 \qquad \textbf{Eq. 6-6}$$

$$P_{f\,upper} \text{ is the } p \text{ value which satisfies : } \sum_{i=0}^{n} C^{i}_{Nd}\, p^{i} (1-p)^{Nd-i} = \alpha / 2 \qquad \textbf{Eq. 6-7}$$

where $$C^{i}_{Nd} = N_d! / [i!(N_d - i)!] \qquad \textbf{Eq. 6-8}$$

with $x! = x(x-1)(x-2)\ldots.3.2.1$, and $0! = 1$.

You do not need to use Equations 6-6 and 6-7 directly, because solutions can be found in Tables 7-2 through 7-4 (and in statistical tabulations covering the binomial distribution). Select the table for the confidence level required, and read the confidence limits from the body of the table. The tables, unfortunately, have a range of application restricted to 48 demands or less. To work with larger numbers of demands, use the fact that at these larger numbers the number of demands can be treated as a continuum, like time, and there is a close analogy between the failure rate of the Poisson model and the probability of failure-on-demand. The third example below explains the procedure.

Caution: in some texts, the binomial distribution is described for p equal to the probability of success, rather than the probability of failure. In that case, the above statements are still all true, except that n becomes the number of successes.

Once again, a confidence level of $(1-\alpha)$ means that the probability is $(1-\alpha)$ that the true failure rate lies between λ_{upper} to λ_{lower}. So, if the confidence level is 90%, $\alpha/2 = 0.05$. The ratio between the upper and lower confidence bounds is just as sensitive to the number of failures as were the bounds for the Poisson distribution.

When no failures occur in N_d demands, the two-sided confidence bounds on p become:

$$P_{f\,lower} = 0 \qquad and \quad P_{f\,upper} = 1 - (\alpha/2)^{1/n} \qquad \textbf{Eq. 6-9}$$

Nuclear plant practice favors the use of a *one-sided* upper confidence bound in this situation, given simply by $1 - \alpha^{1/n}$, stated at a 50% confidence level. Thus:

$$P_{f\,upper} = 1 - 0.5^{1/n} \qquad \textbf{Eq. 6-10}$$

Example 1: Find the 90% confidence limits on the probability of failure-on-demand if there are 2 failures in 15 demands. The estimated probability of failure-on-demand is 2/15 = 0.13. In Table 7-3 (Two-Sided Confidence Limits for Binomial Distribution, Confidence Level: 1-α = 0.9), use the column headed 13 because 15 – 2 = 13. The confidence limits for 2 failures can be read as 0.024 to 0.363.

Example 2: State the one-sided upper 50% confidence limit when there are no failures in 15 demands. Equation 6-10 gives $1 - 0.5^{1/15} = 1 - 0.5^{0.0666} = 1 - 0.955 = 0.045$.

Example 3: 2 failures have occurred in 100 demands. What are the two-sided 95% confidence limits on the probability of failure-on-demand? The estimate of probability of failure-on-demand is 2/100 = 0.02. 100-2 = 98 is outside the range of the column headings of Table 7-4. However, when the number of demands exceeds the range of the tables, Equations 6-2 and 6-3 can be used as a rather accurate analogue, equating N_d with T, and n with N_f, and using Table 7-1:

$$P_{upper} = \chi^2\,1-\alpha/2\,(2x2+2) \;/\; 2x100 = \chi^2\,0.975\,(6) \;/\; 200 = 14.449 / 200 = 0.072$$

$$P_{lower} = \chi^2\,\alpha/2\,(2x2) \;/\; 2x100 = \chi^2\,0.025\,(4) \;/\; 200 = 0.484 / 200 = 0.0024$$

6.5 Updating Knowledge of Failure Rates With New Data

Sections 6.5 to 6.14 deal with the common situation where new data has come to hand. To make the best use of it requires some kind of combination of the new data with the estimates available before the new data was obtained. A trivial case is where values already exist for the failure rate or probability of failure-on-demand, *and the numbers of failures and so forth which gave rise to these values is known*. This might arise when plant-specific data is being updated with additional plant-specific data, and the details of the earlier calculations are still available. Combined values can then be computed from Equations 6-1 to 6-10, after the new experience (number of failures, number of component years of exposure, number of demands) is simply added to the old.

More often, the situation is not as trivial. The more normal situation has the following characteristics: 1) you have some kind of knowledge of the failure rate (the "prior" knowledge); 2) the prior failure history is not known in terms of the numbers of failures and component years of exposure; and 3) new estimates must be made for the failure rate and its confidence bounds.

In this case, the prior knowledge is represented by a probability distribution for the failure rate. The failure rate is thus treated as a random variable, with an uncertainty expressed by the prior distribution. The prior may be a very crude discrete representation, such as "there is a 75% chance that the failure rate is 0.1/year, and a 25% chance that it is 0.01/year," or it may be a completely specified probability distribution such as a lognormal. The prior distribution specifies what you know of the possible values of the failure rate, without any reference to an underlying statistical model of the failure processes, and with no access to the raw statistics which gave rise to it. Bayes' Formula provides the link between the prior distribution, the new data (for example, 1 additional failure in 28 demands), and the final distribution, which is called the posterior distribution:

Bayes' Formula:

$$Posterior\ (\lambda) = K \times Prior(\lambda) \times Likelihood\ of\ the\ Data\ Given\ \lambda \qquad \textbf{Eq. 6-11}$$

K is a constant explained below. The prior distribution of λ is given. It is the distribution you obtain from an industry database expressing prior knowledge about λ. The likelihood expresses the probability of getting exactly the results which were obtained for the new data, if the failure rate had had the value λ. This value is treated as a variable, so you need a statistical model of the underlying failure process, such as the Poisson or the binomial, to find the likelihood for a general value of λ.

There is usually no difficulty at all in writing down the likelihood. For the normal case where the data contains multiple values (for example, a set of failure times), the likelihood will be the repeated product of the probability distributions for the type of data involved, because it expresses the probability of getting the first value, *and* the second, and so on.

The constant, K, is obtained by normalizing the right-hand side of Equation 6-11 to unity by integrating over the full range of possible λ values. The posterior is then a properly normalized probability density when K is calculated as:

$$1/K = \int_0^\infty Prior(\lambda') \bullet Likelihood\ Of\ The\ Data\ Given\ \lambda' \bullet d\lambda' \qquad \textbf{Eq. 6-12}$$

Although Equations 6-11 and 6-12 may require numerical methods to evaluate, confidence bounds on λ are conceptually easy to understand and to evaluate directly. In the general case, you have to integrate the posterior to find K from Equation 6-12, and to calculate confidence limits, but there are some ways to avoid the integration. These are described below.

If the new data is a large data set from a homogeneous population of components (in terms of PM, duty cycle, and so forth), combining it with the prior will have a dominating effect, with the posterior resembling the new data more than the prior. The more usual situation is for the prior to be a rather wide distribution, representing significant uncertainty about λ, as discussed before, and the new data to be of meager statistical weight, possibly differing markedly from the prior in

terms of mean value, and even in terms of its confidence limits. The Bayesian updating process of Equations 6-11 and 6-12 takes precise account of these disparities, and automatically results in a posterior distribution with appropriate weight given to the location and variance of both sources of the data.

There are many variations of the Bayesian approach. New data may take the form of 1) a set of n times to failure, t_1, t_2, t_3, …….t_n, or more simply, 2) n additional failures which occurred in a total operational time of T component years, or 3) additional estimates of λ, its confidence bounds, or its probability distribution from other sources. Similar considerations apply to the probability of failure-on-demand. In the remaining sections of Section 6, we consider each in turn. All the methods described fall in the general class known as Parametric Empirical Bayes (PEB).

6.6 Likelihood for New Times to Failure (Constant Failure Rate)

The most detailed level at which new data may become available is as a set of times to failure. In the power industry, it will be unusual to obtain data on failure times, but if you do obtain data in this form, use the following procedure to embody the assumption that the failure rate does *not* change in time. If we believe that the failure rate is constant in time with value λ, then the times to failure are distributed according to an *exponential* distribution, E(t, λ):

$$E(t,\lambda) = \lambda e^{-\lambda t} \qquad \textbf{Eq. 6-13}$$

The likelihood for a data sample of n new failure times is then the repeated product:

$$L(t_1,t_2,t_3,\ldots\ldots t_n,\lambda) = \lambda e^{-\lambda t_1} \bullet \lambda e^{-\lambda t_2} \bullet \lambda e^{-\lambda t_3} \ldots\ldots\ldots \lambda e^{-\lambda t_n}$$

$$L(t_1,t_2,t_3,\ldots\ldots t_n,\lambda) = \lambda^n \bullet exp\left(-\lambda \sum_{i=1}^{n} t_i\right) \qquad \textbf{Eq. 6-14}$$

This is because e^a times e^b times e^c times……. = $e^{(a+b+c+\ldots\ldots)}$.

Equation 6-14 gives one important part of Equation 6-11.

6.7 Likelihood for Number of New Failures (Failure Rate)

When new data is simply of the form that n failures have occurred in a time T, we use the Poisson distribution of the number of failures to create the likelihood function. The Poisson distribution also contains the assumption that the rate, λ, is a constant over time. P(x,λT) is the probability of observing exactly x failures when the expected (that is, the mean) value is λT failures:

$$P(x,\lambda T) = e^{-\lambda T} (\lambda T)^x / x! \quad x = 1,2,3... \qquad \textbf{Eq. 6-15}$$

The likelihood of observing exactly n failures is thus:

$$L(n,\lambda T) = P(n,\lambda T) \qquad \textbf{Eq. 6-16}$$

This likelihood does not involve a repeated product of the Poisson distribution, because there is only one result, n, in exposure T.

Notice that the Poisson distribution is actually a single parameter distribution which depends only on the product λT, rather than λ and T independently of each other. However, to use it to break out the failure rate obviously requires us to also know the value of T. The mean number of failures is λT for this distribution, with variance also equal to λT.

6.8 Likelihood for Number of New Failures (Failure-on-Demand)

When new data is of the form that n failures have occurred in a number of demands, N_d, we use the *binomial* distribution of the number of failures to create the likelihood function. The binomial distribution also contains the assumption that the probability of failure-on-demand, p, is a constant over time. $B(i,p,N_d)$ is the probability of observing exactly i failures when the expected (that is, the mean) value is pN_d failures:

$$B(i,p,N_d) = Nd!/[i!(N_d - i)!] \bullet p^i (1-p)^{Nd-i} \qquad i = 1,2,3 \ldots N_d \qquad \textbf{Eq. 6-17}$$

The likelihood of observing exactly n failures is thus:

$$L(n,p,N_d) = B(n,p,N_d) \qquad \textbf{Eq. 6-18}$$

The mean value of the number of failures for this distribution is pN_d, with variance $p(1-p)N_d$. Caution: in some texts, the method is described for p equal to the probability of success, rather than the probability of failure. In that case, the above statements are still all true, except that i and n become the number of successes.

6.9 Likelihood for a New Distribution

When new data is in the form of a probability density, g(λ), from a different source, the likelihood is simply the new data, since by definition g(λ) represents the appropriate relative probabilities—that is, the likelihood—of getting various values of λ. In this case, Equations 6-11 and 6-12 become the overlap probability, between the prior and the new distributions.

$$Posterior(\lambda) = \frac{Prior(\lambda) \bullet g(\lambda)}{\int_0^\infty Prior(\lambda') \bullet g(\lambda') d\lambda'} \qquad \textbf{Eq. 6-19}$$

This is therefore the approach to use to combine different generic sources of data on the same parameter.

6.10 Lognormal Prior Distribution of Failure Rate

This section simply introduces the lognormal distribution and the relations between its parameters. A common choice for the prior distribution of λ, or for the probability of failure-on-demand, is the lognormal. The reason is that lognormal distributions are favored by PSA practitioners, and so a large amount of existing failure rate knowledge, whether plant-specific or from generic industry sources, will be found in this form. The lognormal is perfectly serviceable but usually requires numerical methods to perform Bayesian calculations. The usual assumption is that the lognormal expresses uncertainty in λ, which is, however, a constant in time. The lognormal probability density, LN(λ), is given by:

$$LN(\lambda) = [1/\sqrt{(2\pi)}] \bullet [1/\lambda\sigma] \bullet \exp[-\{\ln(\lambda/\lambda_m)\}^2 / (2\sigma^2)] \quad \textbf{Eq. 6-20}$$

Where σ is the standard deviation of lnλ, and ln(x) is the logarithm of x to base e. The mean and variance are:

$$Mean = \lambda_m \bullet \exp(\sigma^2/2) \quad \textbf{Eq. 6-21}$$

$$Variance = \lambda_m^2 \bullet \exp(\sigma^2)\,(\exp(\sigma^2) - 1) \quad \textbf{Eq. 6-22}$$

Other useful values are:

$$Median = \lambda_m \quad \textbf{Eq. 6-23}$$

$$Mode = \lambda_m / \exp(\sigma^2) \quad \textbf{Eq. 6-24}$$

$$Error\ Factor = \exp(1.645\,\sigma) \quad \textbf{Eq. 6-25}$$

The error factor provides a way to calculate the symmetric lower and upper confidence bounds at a 90% confidence level. See Equation 6-32 of Section 6.14, and also Section 6.13, for examples of using the error factor.

6.11 Self-Conjugate Prior: Constant Failure Rate

It should be obvious that use of a general prior distribution, along with the likelihoods which stem from the statistical failure models described in Sections 6.6 to 6.9, will require numerical computation to evaluate the posterior from Equations 6-11 and 6-12. There are two extremely useful situations where use of likelihoods of the standard forms already described leads to the posterior distribution being of the same functional form as the prior, hence the term self-conjugate. All that needs to be done to perform a Bayesian update for these cases is to modify

the parameters of the prior distribution in trivial ways to immediately arrive at the posterior, without going through the rigors of solving Equations 6-11 and 6-12.

These two methods are of great utility, because 1) they are suited to likelihoods based on Poisson and binomial failure models, which as we have seen above are the most frequently needed cases; and 2) the prior distributions have very general forms which can be made to approximate almost any form of knowledge about the failure rate and the probability of failure-on-demand, including lognormal distributions.

There has been some criticism of self-conjugate priors on the grounds that they may be a little too resistant to modification by the new data, because they tend to underemphasize the uncertainty in the tail regions of the distributions. However, for the purposes of LCM applications they are a very convenient starting point. The tail region issue can be resolved by incorporating a noninformative prior into the parametric prior in a first Bayes updating step before the new data is introduced. This two-stage Bayes-Empirical Bayes approach (BEB) seems to be a robust improvement to the single stage Parametric Empirical Bayes procedures described here, but it is more complex, and users are advised to seek expert statistical input to select an appropriate non-informative prior.

The first case is presented for a constant failure rate, λ, which uses a Poisson-based likelihood as described in Section 6.7. The prior distribution is chosen to be the gamma distribution, $G(\lambda,b,c)$:

$$G(\lambda,b,c) = (\lambda/b)^{(c-1)} \bullet exp[-(\lambda/b)] / [b\Gamma(c)] \qquad \textbf{Eq. 6-26}$$

$\Gamma(c)$ is the gamma function which can be found in most statistical tabulations. The mean failure rate = bc, and the variance = b^2c. You can choose to restrict c to take only integer values, in which case $\Gamma(c) = (c-1)!$ Restricting c in this way has the justification that c is closely associated with the number of failures, at least when c is not too small, and it makes it easy to plot the distribution without using tabulations of the gamma function. However, it introduces extra error when using the method of matching moments (see Section 6.13) to determine equivalent distributions. Users should normally restrict the c parameter to integer values.

Suppose the parameters of the prior are b_0,c_0. When the new data consists of n failures in an exposure time of T years, the posterior distribution will still be of the gamma form, but with parameters, b_1,c_1, where:

$$b_1 = b_0/(1+T\,b_0)\ ;\ \ and\ \ \ c_1 = c_0 + n \qquad \textbf{Eq. 6-27}$$

The updated mean failure rate is b_1c_1 instead of b_0c_0 for the prior. This is a very easy way to avoid the complexities of Equations 6-11 and 6-12.

Example: Prior information for the failure rate is a gamma distribution, with parameters $b_0 = 1$, $c_0 = 0.05$, so that the mean value (bc) of the failure rate is 0.05/year, with a standard deviation (square root of the variance, b^2c) = 0.22. If more recent data consists of just 2 failures in 100 component years of operation, what are the new mean and standard deviation?
$b_1 = 1/(1+100x1) = 0.01$, and $c_1 = 2.05$. The new mean is thus 0.0205/year, and the new standard

deviation is 0.0143. If we had elected to restrict c_1 to an integer value ($c_1 = 2$) it would not have significantly influenced the result in this case, but there is usually no need to do this.

In this example, the new data dominates the mean because the prior distribution had a standard deviation about 4.4 times the mean, and the new data had more statistical weight. To see this, consider that the new data on its own would have given a mean value of $\lambda = 0.02$ with a 2-sided upper 90% confidence limit of $\chi^2_{0.95}(6)/200 = 0.053$. Although this is still 2.5 times the mean, the difference between the upper 90% limit and the mean is roughly 2 times the standard deviation, suggesting a standard deviation in the range 0.01 to 0.02. We can estimate the standard deviation exactly by stating that the standard deviation on the *number of failures* is the square root of the variance ($=\sqrt{(\lambda T)} = \sqrt{(N_f/T \times T)}$, and is thus $\sqrt{N_f}$, equal to 1.414, giving the standard deviation on the estimate of λ of 1.414/100 = 0.014. Therefore the new data alone would give $\lambda = 0.02$ with a standard deviation of ± 0.014. It is clearly more significant than the prior information which stated $\lambda = 0.05$ with a standard deviation of ± 0.22.

It is worth remembering that the standard deviation of the number of failures in a Poisson rate process is the square root of the number of failures, as this gives the analyst an immediate sense of the uncertainty in this number.

6.12 Self-Conjugate Prior: Constant Probability of Failure-on-Demand

This case uses a binomial-based likelihood, as described in Section 6.8. The prior distribution is chosen to be the beta distribution, BETA(p,V,W), with p the probability of failure-on-demand. For this application, the parameters, V and W, must be integers:

$$BETA(p,V,W) = \{(V+W-1)! / [(V-1)!(W-1)!]\} \bullet p^{(V-1)} \bullet (1-p)^{(W-1)} \qquad \textbf{Eq. 6-28}$$

The mean probability of failure-on-demand is given by $V/(V+W)$, and the variance is $VW / [(V+W)^2(V+W+1)]$.

Suppose the parameters of the prior are V_0, W_0. When the new data consists of n failures in N_d additional demands, the posterior distribution will still be of the beta form, with parameters V_1, W_1, where:

$$V_1 = V_0 + n \;\; ; \;\; and \;\; W_1 = W_0 + N_d - n \qquad \textbf{Eq. 6-29}$$

The parameter V_0 is thus modified by adding the number of additional failures, whereas W_0 is modified by adding the number of additional successes. The posterior mean probability of failure-on-demand is $V_1/(V_1+W_1)$ instead of $V_0/(V_0+W_0)$ for the prior. This method for the binomial distribution is just as straightforward as the previous use of the gamma prior for the Poisson distribution.

Caution: in some texts, the method is described for p equal to the probability of success, rather than the probability of failure. In that case, the above statements are still all true, except that n

becomes the number of new successes, V is associated with the number of successes rather than failures, and W is associated with the number of failures.

6.13 Parameters for the Prior: Method of Moments

When the prior distribution which is available to you is of the appropriate self-conjugate form, you use it directly with the parameters provided, following the procedures of Section 6.11.

However, it may happen that your prior is not in this form. For example, it will often be a lognormal prior distribution, and you then wish to convert it to an equivalent gamma or beta distribution so you can more conveniently use the conjugate prior methods of Section 6.11. A good way to match two distributions of any kind is to equate their means and variances. This is the method of matching moments for any one or two parameter distributions. Obviously, if there are more than two parameters to be specified, more than two moments must be matched, but we do not need to go beyond matching the mean and variance to address all the distributions mentioned in previous sections. The mean and variance of the lognormal, gamma, and beta distributions were given in Sections 6.10, 6.11, and 6.12, respectively.

For example, if a prior lognormal distribution had a mean of 0.01 failures per year and an error factor of 18, the standard deviation of $\ln\lambda$ must be $\sigma = \ln(18)/1.645 = 1.757$, by Equation 6-25. In that case, the variance is $\lambda_m^2 \bullet \exp(1.757^2)\,(\exp(1.757^2) - 1)$, by Equation 6-22. So:

$$Variance = \lambda_m^{\,2} \bullet 21.91 \bullet 20.91 = 458.19\,\lambda_m^{\,2}$$

$$Mean = \lambda_m \bullet exp(1.757^2/2) = 10.956\,\lambda_m \quad by\ Equation\ 6\text{-}21$$

$$= 0.01 \quad (given)$$

Therefore, $\lambda_m = 0.00091$, and the variance is 0.000382. Note that the median, λ_m, is 10 times smaller than the mean, not an unusual situation for failure rate distributions which tend to have long tails in the upper part of the range. If we need to match the lognormal prior to a gamma distribution, we put:

$$Mean: bc = 0.01$$

$$Variance: \quad b^2c = 0.000382, equivalent\ to\ a\ standard\ deviation\ of\ 0.02.$$

Whence: $b = 0.0382$, and $c = 0.262$. These two values would then be used as prior values, b_0 and c_0, before modifying them with new data. Suppose the new data were 1 failure (n) in 10 additional years (T) of component experience. Equation 6-27 gives

$$b_1 = b_0/(1 + b_0T) \quad i.e., \quad b_1 = 0.0382/(1 + 0.382)\ i.e.\ b_1 = 0.0276$$

$$c_1 = c_0 + n \quad i.e., \quad c_1 = 1.262$$

The new mean value of λ is thus 0.0276 x 1.262 = 0.034, and the new variance is 0.00096, whereas the new data alone would have given λ_{mean} = 0.1 failures per year, and the prior information had λ_{mean} = 0.01 failures per year. The new data does not completely dominate the prior, but it changes it significantly. This is because the standard deviation on just 1 failure is √1=1 failure, giving a standard deviation for λ based on the new data alone of 1/10—that is, λ_{mean} = 0.1 ± 0.1—whereas the prior had λ_{mean} = 0.01 ± 0.02.

The posterior distribution over λ is still of the gamma form:

$$Posterior(\lambda) = [(2/0.0276^2)(\lambda/0.0276)^{(1.262-1)} \bullet exp[-(\lambda/0.0276)] / \chi^2_{2.524}]$$

Of course, the example could be worked by numerically evaluating the posterior directly, using Equations 6-11 and 6-12, using the likelihood, $\lambda 10e^{-10\lambda}$, from Equation 6-13:

$$Posterior(\lambda) = \frac{[1/\lambda] \bullet exp[-\{ln(\lambda/0.00091)\}^2 / (2x1.757^2)] \bullet \lambda e^{-10\lambda}}{\int_0^\infty [1/\lambda'] \bullet exp[-\{ln(\lambda'/0.00091)\}^2 / (2x1.757^2)] \bullet \lambda' e^{-10\lambda'} \bullet d\lambda'}$$

The numerical constants which cancel out between the numerator and denominator have been omitted. This result would be somewhat more accurate than matching the moments, but it involves a lot more work, and the difference would only be seen by plotting the distributions.

When matching mean and variance for a beta distribution in the case of a probability of failure-on-demand, recall that the beta distribution parameters, V and W, are restricted to integer values. This means that you have to round off the values to the nearest integer. For example, if you find that the matching equations give you V = 24.31 and W = 1.66, then you select V = 24 and W = 2.

6.14 Point Estimates and Confidence Bounds

When you end up with a posterior distribution for the failure rate, but need to quote or use a single value for λ, point estimates of failure rate or probability of failure-on-demand can be chosen which correspond to the mean, median, or mode of the prior or posterior distributions.

In general, Bayesian confidence bounds are obtained by numerically integrating over the posterior distributions, although in special cases there exist closed forms for these integrals. If the distribution for the failure rate is P(λ), the confidence bounds at a (1-α) confidence level are the solutions of:

$$\alpha/2 = \int_0^{\lambda_{lower}} P(\lambda')\, d\lambda' \qquad \text{Eq. 6-30}$$

and

$$1 - \alpha/2 = \int_0^{\lambda_{upper}} P(\lambda')\, d\lambda' \qquad \textbf{Eq. 6-31}$$

In the case of a lognormal distribution, the 90% confidence bounds can be expressed very simply in terms of the error factor, EF, which was defined by Equation 6-25, such that:

$$\lambda_{lower} = \lambda_m / EF \qquad and \qquad \lambda_{upper} = \lambda_m \bullet EF \qquad \textbf{Eq. 6-32}$$

Bounds for other confidence levels for a lognormal distribution can be determined from tabulations of integrals of the normal distribution function ($\ln\lambda$ is normally distributed), as an alternative to using a numerical procedure to evaluate Equations 6-30 and 6-31. In general, bounds on a lognormal distribution are not very important in updating failure rate data, because even if you begin with a lognormal prior, the posterior distribution will not usually be lognormal.

In the case of a gamma posterior distribution, the chi-squared distribution gives the two-sided confidence bounds at the $(1-\alpha)$ level. Use Table 7-1, as shown previously, to evaluate:

$$\lambda_{lower} = (b_1/2) \bullet \chi^2_{\alpha/2}\,(2c_1) \qquad \textbf{Eq. 6-33}$$

and

$$\lambda_{upper} = (b_1/2) \bullet \chi^2_{(1-\alpha/2)}\,(2c_1) \qquad \textbf{Eq. 6-34}$$

For a beta distribution, use the tabulated values of percentage points of a beta distribution given in tables, or perform a numerical procedure based on Equations 6-30 and 6-31 in order to determine the confidence limits. Tables 7-5 through 7-7 give the lower confidence limit for $\alpha/2$ = 2.5%, 5%, and 10%—that is, for confidence levels of 95%, 90%, and 80%, respectively. To find the upper limits, interchange the values of V and W when using the tables, and then subtract the value obtained from the table from unity.

For example, if V=20, and W=10, Table 7-6 shows the lower 90% limit to be $P_{lower} = 0.52$. Use the table again with V=10 and W=20 to get $P_{upper} = 1 - 0.20 = 0.80$. You may need to interpolate for intermediate values of V and W.

6.15 Weibull Analysis of Times to Failure

In the case where a set of times to failure is available for an essentially non-repairable component, preferably for a single failure mechanism, the assumption of a Weibull distribution is the standard procedure. This is a general three-parameter power law model for the failure times, with a failure time distribution:

$$W(\eta, \gamma, \beta, t) = (\beta/\eta)\,[(t - \gamma)/\eta])^{(\beta-1)} \bullet exp - [(t - \gamma)/\eta]^{\beta} \qquad \textbf{Eq. 6-35}$$

which gives the time dependent failure rate as:

$$\lambda(\eta,\gamma,\beta,t) = (\beta/\eta)\ [(t-\gamma)/\eta]^{(\beta-1)} \qquad \textbf{Eq. 6-36}$$

This distribution is of wide generality, capable of accurately representing the exponential distribution when the shape parameter, β, is equal to 1, and even of approximating a normal distribution when β = 3.44.

The location parameter, γ, the useful life or minimum life, can be removed by making a shift of the time axis, because it only indicates that the time-dependent behavior begins at time, t= γ. Therefore, the standard analysis procedure assumes that you do not know the value of γ until you begin plotting the times to failure on Weibull paper. This is equivalent to initially assuming that the times to failure are distributed according to the two-parameter Weibull distribution (that is, Equation 6-35 with γ = 0). Unfortunately, the value of γ is not determined directly, even by later plots.

The parameter η is called the scale parameter or the characteristic life. This is the age at which 63.2% of the sample will have failed (when γ = 0; otherwise you need to add the value of γ to the characteristic life). Clearly, η provides some indication of the width of the distribution. The parameter, β, is a shape parameter, capable, as shown above, of making the distribution approximate the shape of many other distributions. When β>1, the failure rate is increasing with time.

The Weibull plot requires the times to failure to be ordered from the smallest to the largest. It also requires the total population of components subject to the sample conditions to be known. For example, if 20 components are to participate in the data sample, the failure of the first (shortest time to failure), represents a failure of 5% of the total. Failure of the second represents cumulative failure of 10% of the total, and so on. The Weibull plot consists of plotting the cumulative failure percentage on the y-axis, and the failure time on the x-axis (time is most common, but it could be cycles, revolutions, and so forth).

Draw a straight line through the points, usually by eye, but conceivably using linear regression. If this cannot be done because the line needs to curve, the points must be replotted using the set $(t_i - \gamma)$ rather than t_i. Estimate the value of γ as follows:

Draw a curved line through the data points and select an arbitrary point (y_2, t_2) approximately in the center of the line.

Choose two other points, one above and one below the center point, and both exactly equidistant from it in the vertical direction. Label the points 1, 2, and 3, with 1 for the shortest time.

Use $\gamma = t_2 - (t_3 - t_2)(t_2 - t_1) / [(t_3 - t_2) - (t_2 - t_1)]$ as an estimate for γ.

If the replotted points are still not linear, the data cannot be represented by a Weibull distribution. There may be two or more different wearout processes contributing to the data.

The estimate for η is found by reading the t value at which the straight line through the points intersects the dashed 'η estimate' line on the paper.

The estimate for β is found by drawing a line perpendicular to the plotted line through the estimation point marked on the top left corner of the graph paper. The estimate for β is read where the perpendicular crosses the β-scale along the top of the paper.

It is important that the group of components which provided the n times to failure must be defined before the failure times are observed. This means that you cannot allow a small number of failures which occur in a large population to define the sample, because you do not know beforehand which components will fail. Thus, in the normal power plant situation where there is a large population of components, N, and you find that n of them fail, you must use N as the sample size, not n. The cumulative failure percentage at the nth failure time is 100n/N. In power plants this will almost always be a small percentage, with the result that the Weibull plot will be confined close to the bottom of the Weibull chart. The result will be either that estimates of η and β are impossible to make, or that they will have uncertainties so large that they do not provide usable information.

This "no information" scenario is simply a statement that when only a small fraction of a population of components has failed, you cannot say anything about the time dependence of the failures of the other components at more distant times.

6.16 Linear Regression Applied to Estimates of Failure Rate

In the case where you may acquire multiple values for the failure rate which purport to address the same equipment at a variety of ages, it may be possible to determine an age dependence of the failure rate simply by analyzing the values using regression—that is, by drawing a line through them on a chart of failure rate versus age. The simplest lines would be straight lines, hence the name linear regression. The main problem with this approach is that such data samples will most likely be in the form of probability distributions which individually display a wide dispersion of possible values of λ, a dispersion which cannot be represented easily in the regression approach. The following sections outline the use of point values (mean, median, and mode) from such distributions, as well as upper and lower confidence bounds.

The assumption here is that a set of point estimates, λ_i, are correlated with the age of the equipment in each sample. Generally, only a linear time dependence is sought. The values of λ_i would be plotted as y values and the age of each sample, t_i, as x values. Linear regression can improve upon drawing a straight line through the points by eye, because it is a formal process capable of estimating the parameters, a and b, of the relation, $\lambda(t) = at + b$, and also of estimating confidence intervals for the parameters. However, the procedure will not take account of the prior uncertainty in each data point, other than by the degree of scatter displayed by all the points about the regression line.

To perform the analysis, use one of many standard software packages which perform a variety of least squares fitting procedures.

6.17 The Case of No Data

If you cannot find *any* data—generic, plant-specific, time-dependent, or time-independent—to represent the failure rate of a component type, you will have to use a different component type as a surrogate. This is not as unreasonable as it may seem. It is reasonable to expect that a DC motor will have a failure rate more similar to that of an AC motor than to that of a printed circuit board. Many different kinds of high-speed rotating machinery will share many failure mechanisms, but these could not be expected to resemble the failure mechanisms of a high-voltage breaker. It is clear that an appreciation of general design features will suggest which types of equipment may have failure rates which resemble each other. Of course, it is also common experience to find that two different models or manufacturers of the same basic component may nevertheless differ markedly in reliability. In these cases, experience usually points to specific design elements as the reasons for the differences.

Start with equipment in the same broad category. Eliminate any items which are known to have operating experience markedly different from the component of interest. Focus on one or two potential surrogate component types, and verify that failure rate information is available for them. Use the EPRI PM Basis database [5] to compare the expected unreliability of the component types of interest under equivalent assumptions about duty cycle, service conditions, and comprehensiveness of the PM program. The Statistics summary of the Default Vulnerability calculation will provide the necessary measure of unreliability. In the database, select the component type, then the Vulnerability button, then the Statistics button.

The Statistics form provides numerical results in terms of the numbers and percentages of opportunities for failures represented by subsets of the data. The number in the top right box is proportional to the number of failures which are *not* prevented by the PM program. These are the failures which are responsible for the residual unreliability which is experienced when using the PM program which has been analyzed. Compare this result between a potential surrogate component type and the component of interest. If they do not differ by more than a factor of 2 or 3 in this measure of unreliability, consider that the failure rate of the surrogate will form a satisfactory replacement.

In fact, it is quite reasonable to use the ratio of the two results as an adjustment factor on the failure rate. This comparison can be carried much further using the database, with consequent refinements of the results at each stage. For example, an examination of the “Red” records will reveal the reason why they are not well-protected by the PM tasks. In some cases the reason will lie in a paucity of tasks, in other cases in the fact that the tasks are not done frequently enough, in yet other cases in a larger proportion of randomly occurring failure mechanisms. If degradation mechanisms which are thought to not apply to the component of interest are removed (this requires administrative access to the data tables), and others added which are thought to be valid additions for the component of interest, the result can be made more realistic. Furthermore, adjustments can be made to the PM tasks, using the statistics results for Custom Vulnerability calculations, in which the user is free to make the PM coverage for the two components as comparable as possible.

The fact that these estimates can be made rather easily should not encourage any user to believe the results to be better than a factor of 3. However, since the normal margins of uncertainty on equipment failure rates are of this magnitude anyway, the method may have some utility. Its main disadvantage, of course, is that both the component of interest and potential surrogates must be present in the database.

7
STATISTICAL TABLES

Table 7-1
Percentage Points of Chi-Squared Distribution With ν Degrees of Freedom, $\chi^2_\varepsilon(\nu)$

	Values of ν									
ν	**0.005**	**0.025**	**0.05**	**0.10**	**0.20**	**0.80**	**0.90**	**0.95**	**0.975**	**0.995**
1	0.0000393	0.000982	0.00393	0.0158	0.0642	1.642	2.706	3.841	5.024	7.879
2	0.0100	0.0506	0.103	0.211	0.446	3.219	4.605	5.991	7.378	I0.597
3	0.0717	0.216	0.352	0.584	1.005	4.642	6.251	7.815	9.348	12.838
4	0.207	0.484	0.711	1.064	1.649	5.989	7.779	9.488	11.143	14.860
5	0.412	0.831	1.145	1.610	2.343	7.289	9.236	11.070	12.832	16.750
6	0.676	1.237	1.635	2.204	3.070	8.558	10.645	12.592	14.449	18.548
7	0.989	1.690	2.167	2.833	3.822	9.803	12.017	14.067	16.013	20.278
8	1.344	2.180	2.733	3.490	4.594	11.030	13.362	15.507	17.535	21.955
9	1.735	2.700	3.325	4.168	5.380	12.242	14.684	16.919	19.023	23.589
10	2.156	3.247	3.940	4.865	6.179	13.442	15.987	18.307	20.483	25.186
11	2.603	3.816	4.575	5.578	6.989	14.631	17.275	19.675	21.920	26.757
12	3.074	4.404	5.226	6.304	7.807	15.812	18.549	21.920	23.337	28.300
13	3.565	5.009	5.892	7.042	8.634	16.985	19.812	22.362	24.736	29.819
14	4.075	5.629	6.571	7.790	9.467	18.151	21.064	23.685	26.119	31.319
15	4.601	6.262	7.261	8.574	10.307	19.311	22.307	24.996	27.488	32.801
16	5.142	6.908	7.962	9.312	11.152	20.465	23.542	26.296	28.845	34.267
17	5.697	7.564	8.672	10.085	12.002	21.615	24.769	27.587	30.191	35.718
18	6.265	8.231	9.390	I0.865	12.857	22.760	25.989	28.869	31.526	37.156
19	6.844	8.907	10.117	11.651	13.716	23.900	27.204	30.144	32.852	38.582
20	7.434	9.591	I0.851	12.443	14.578	25.038	28.412	31.410	34.170	39.997
21	8.034	10.283	11.591	13.240	15.445	26.171	29.615	32.671	35.479	41.401
22	8.643	10.982	12.338	14.041	16.314	27.301	30.813	33.924	36.781	42.796
23	9.260	11.688	13.091	14.848	17.187	28.429	32.007	35.172	38.076	44.181
24	9.886	12.401	13.848	15.659	18.062	29.553	33.196	36.415	39.364	4S.558
25	10.520	13.120	14.611	16.473	18.940	30.675	34.382	37.652	40.646	46.928
26	11.160	13.844	15.379	17.292	19.820	31.795	35.563	3S.885	41.923	48.290
27	11.808	14.573	16.151	18.114	20.703	32.912	36.741	40.113	43.194	49.645
28	12.461	15.308	16.928	18.939	21.588	34.027	37.916	41.337	44.461	50.993
29	13.121	16.047	17.708	19.768	22.475	35.139	39.087	42.557	45.722	52.336
30	13.787	16.791	18.493	20.599	23.364	36.250	40.256	43.773	46.979	53.672
35	17.156	20.558	22.462	24.812	27.820	41.802	46.034	49.798	53.207	60.304
40	20.674	24.423	26.507	29.067	32.326	47.295	51.780	55.755	59.345	66.792
45	24.281	28.356	30.610	33.367	36.863	52.757	57.480	61.653	65.414	73.190
50	27.962	32.348	34.762	37.706	41.426	58.194	63.141	67.502	71.424	79.512
55	31.708	36.390	38.956	42.078	46.011	63.610	68.770	73.309	77.384	85.769
60	35.510	40.474	43.186	46.478	50.614	69.006	74.370	79.080	83.301	91.970
65	39.360	44.595	47.448	50.902	55.233	74.367	79.946	84.819	89.181	93.122
70	43.253	48.750	51.737	55.349	59.868	79.752	85.500	90.530	95.027	104.230
75	47.186	52.935	56.052	59.815	64.515	85.105	91.034	96.216	100.843	110.300
80	51.153	57.146	60.390	64.299	69.174	90.446	96.550	101.879	106.632	116.334
85	55.151	61.382	64.748	68.799	73.843	95.777	102.050	107.521	112.397	122.337
90	59.179	65.640	69.124	73.313	78.522	I0I.097	107.536	113.145	118.139	128.310
95	63.963	69.919	73.518	77.841	83.210	106.409	113.008	118.751	123.861	134.257
100	67.312	74.216	77.928	82.381	87.906	111.713	118.468	124.342	129.565	140.179
105	71.414	78.530	82.352	86.933	92.610	117.009	123.917	129.918	135.250	146.078
110	75.536	82.861	86.790	91.495	97.321	112.299	129.355	135.480	140.920	151.956
115	79.679	87.207	91.240	96.067	102.038	127.581	134.782	141.030	146.574	157.814
120	83.839	91.567	95.703	100.648	106.762	132.858	140.201	146.568	152.215	163.654

Table 7-2
Two-Sided Confidence Limits for Binomial Distribution, Confidence Level:1-α=0.8

#Fail-	Number of Demands Minus the Number of Failures																	
ures	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
0	.900	.684	.536	.438	.369	.319	.280	.250	.226	.206	.189	.175	.162	.152	.142	.134	.127	.120
	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
1	.949	.804	.680	.584	.510	.453	.406	.368	.337	.310	.288	.268	.251	.236	.222	.210	.199	.190
	.051	035	.026	.021	.017	.015	.013	.012	.010	.010	.609	.006	.007	.007	.007	.006	.006	.006
2	.965	.857	.753	.667	.596	.538	.490	.450	.415	.386	.360	.337	.317	.300	.284	.269	.256	.245
	.196	.143	.112	.093	.079	.069	.061	.055	.049	.045	.042	.039	.036	.034	.032	.030	.028	.027
3	.974	.888	.799	.721	.655	.599	.552	.511	.475	.444	.417	.393	.371	.357	.334	.319	.304	.291
	.320	.247	.201	.170	.147	.130	.116	.105	.096	.088	.08l	.076	.071	.067	.063	.059	.056	.054
4	.979	.907	.830	.760	.699	.646	.599	.559	.523	.492	.464	.439	.416	.396	.378	.361	.345	.331
	.416	.333	.279	.240	.210	.188	.169	.154	.142	.131	.122	.114	.107	.101	.095	.090	.086	082
5	.983	.921	.653	.790	.733	.682	.638	.598	.563	.532	.503	.478	.455	.434	.415	.397	.381	.366
	.490	.404	.345	.301	.267	.240	.219	.201	.185	.172	.161	.151	.142	.134	.127	.121	.115	.110
6	.985	.931	.870	.812	.760	.712	.669	.631	.596	.565	.537	.512	.489	.467	.448	.430	.413	.398
	.547	.462	.401	.354	.318	.288	.264	.243	.226	.210	.197	.185	.175	.165	.158	.150	.143	.137
7	.987	.939	.884	.831	.781	.736	.695	.658	.625	.594	.567	.541	.518	.497	.477	.459	.442	.426
	.594	.510	.448	.401	.362	.331	.305	.282	.263	.246	.231	.218	.207	.196	.187	.178	.170	.163
8	.988	.945	.895	.846	.799	.757	.718	.682	.650	.620	.592	.567	.544	.523	.503	.464	.457	.451
	.632	.550	.489	.441	.402	.369	.342	.318	.297	.279	.263	.249	.236	.225	.214	.205	.196	.188
9	.990	.951	.904	.858	.815	.774	.737	.703	.671	.642	.615	.590	.568	.546	.526	.508	.491	.475
	.663	.585	.525	.477	.437	.404	.375	.350	.329	.310	.293	.278	.264	.252	.241	.230	.221	.212
10	.990	.955	.912	.869	.828	.790	.754	.721	.690	.662	.636	.611	.589	.567	.548	.529	.512	.496
	.690	.614	.556	.508	.468	.435	.406	.380	.356	.338	.321	.305	.290	.277	.265	.254	.244	.235
11	.991	.958	.919	.878	.839	.803	.769	.737	.707	.679	.654	.630	.608	.587	.567	.549	.532	.515
	.712	.640	.583	.536	.497	.463	.433	.408	.365	.364	.346	.330	.315	.301	.289	.277	.267	.257
12	.992	.961	.924	.886	.849	.815	.762	.751	.722	.695	.670	.647	.525	.604	.585	.567	.550	.533
	.732	.663	.607	.561	.522	.488	.459	.433	.410	.389	.370	.353	.336	.324	.311	.299	.288	.277
13	.993	.964	.929	.893	.858	.825	.793	.764	.736	.710	.685	.662	.641	.620	.601	.583	.556	.550
	.749	.683	.629	.584	.545	.511	.482	.456	.432	.411	.392	.375	.359	.345	.331	.319	.308	.297
14	.993	.966	.933	899	.866	834	.804	.775	.748	.723	.699	.576	.655	.635	.616	.599	.582	.566
	.764	.700	.648	.604	.566	.533	.503	.477	.454	.433	.413	.396	.380	.365	.351	.338	.327	.316
15	.993	.968	.937	.905	.873	.842	.813	.786	.759	.735	.711	.689	.669	.649	.630	.613	.596	.580
	.778	.716	.666	.622	.585	.552	.523	.497	.474	.452	.433	.415	.399	.384	.370	.357	.345	.333
16	.994	.970	.941	.910	.679	.850	.622	.795	.770	.746	.723	.701	.681	.662	.643	.626	.609	.594
	.790	.731	.681	.639	.603	.570	.541	.516	.492	.471	.451	.433	.417	.401	.387	.374	.362	.350
17	.994	.972	.944	.914	.885	.857	.830	.604	.779	.756	.733	.712	.692	.673	.655	.638	.622	.606
	.801	.744	.696	.655	.619	.587	.558	.533	.509	.488	.468	.450	.434	.418	.404	.391	.378	.366
18	.994	.973	.946	.918	.890	.863	.837	.612	.788	.765	.743	.723	.703	.684	.667	.650	.634	.618
	.810	.755	.709	.669	.634	.602	.574	.549	.525	.504	.485	.467	.450	.434	.420	.406	.394	.382
19	.995	.974	.949	.922	.895	.869	.843	.819	.796	.774	.752	.732	.713	.695	.677	.660	.645	.629
	.819	.766	.721	.682	.547	.617	.589	.564	.541	.519	.500	.482	.465	.449	.435	.421	.408	.396
20	.995	.976	.951	.925	.899	.874	.849	.826	.803	.782	.761	.741	.722	.704	.687	.671	.655	.640
	.827	.776	.732	.694	.660	.630	.603	.578	.555	.534	.514	.496	.480	.464	.449	.436	.423	.411
22	.995	.978	.955	.931	.907	.883	.660	.836	.817	.796	.776	.757	.739	.722	.705	.689	.674	.659
	.841	.793	.752	.716	.683	.654	.628	.603	.581	.560	.541	.523	.506	.491	.476	.462	.449	.437
24	.996	.979	.959	.936	.914	.891	.870	.849	.828	.809	.790	.772	.754	.737	.721	.706	.691	.677
	.853	.808	.769	.735	.703	.675	.650	.626	.604	.584	.565	.547	.530	.515	.500	.486	.473	.461
26	.996	.981	.961	.941	.919	.896	.678	.858	.838	.820	.802	.784	.767	.751	.736	.721	.706	.692
	.863	.821	.784	.751	.721	.694	.669	.646	.625	.605	.586	.569	.552	.537	.522	.508	.495	.483
28	.996	.982	.964	.944	.924	.905	.685	.866	.848	.830	.812	.796	.779	.764	.749	.734	.720	.706
	.872	.832	.797	.766	.737	.711	.687	.664	.643	.624	.606	.588	.572	.557	.543	.529	.516	.503
30	.997	.983	.966	.948	.929	.910	.691	.873	.856	.838	.822	.806	.790	.775	.760	.746	.733	.719
	.880	.842	.809	.778	.751	.726	.702	.681	.660	.641	.623	.606	.590	.575	.561	.548	.535	.522

Table 7-3
Two-Sided Confidence Limits for Binomial Distribution, Confidence Level:1-α=0.9

# Fail-	Number of Demands Minus the Number of Failures																	
ures	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
0	.950	.776	.632	.527	.451	.393	.348	.312	.283	.259	.238	.221	.206	.193	.181	.171	.162	.153
	.000	.000	.000	.000	000	.000	.000	.000	.000	.000	.000	.000	.000	000	.000	.000	.000	.000
1	.975	.865	.751	.657	.582	.521	.471	.429	.394	.364	.339	.316	.297	.279	.264	.250	.238	.226
	.025	.017	.013	.010	.009	.007	.006	.006	.005	.005	.004	.004	.004	.003	.003	.003	.003	.003
2	.983	.902	.811	.729	.659	.600	.550	.507	.470	.438	.410	.385	.363	.344	.326	.310	.296	.283
	.135	.098	.076	.063	.053	.046	.041	.037	.033	.030	.028	.026	.024	.023	.021	.020	.019	.0l8
3	.987	.924	.847	.775	.711	.655	.607	.564	.527	.495	.466	.440	.417	.396	.377	.359	.344	.329
	.249	.169	.153	.129	.111	.098	.087	.079	.072	.066	.061	.057	.053	.050	.047	.044	.042	.040
4	.990	.937	.871	.807	.749	.697	.650	.609	.573	.540	.511	.484	.461	.439	.419	.401	.384	.369
	.343	.271	.225	.193	.169	.150	.135	.123	.113	.104	.097	.090	.085	.080	.075	.071	.068	.065
5	.991	.947	.889	.831	.778	.729	.685	.645	.610	.577	.548	.522	.498	.476	.456	.437	.420	.404
	.418	.341	.289	.251	.222	.200	.18l	.166	.153	.142	.132	.124	.116	.110	.104	.099	.094	.090
6	.993	.954	.902	.050	.800	.755	.713	.675	.640	.609	.580	.554	.530	.508	.487	.469	.451	.435
	.479	.400	.345	.303	.271	.245	.224	.206	.191	.178	.166	.156	.148	.140	.132	.126	.120	.115
7	.994	.959	.913	.865	.819	.776	.736	.700	.667	.636	.608	.582	.558	.536	.516	.496	.479	.462
	.529	.450	.393	.350	.315	.287	.264	.244	.227	.212	.199	.188	.177	.168	.160	.152	.148	.139
8	.994	.963	.921	.877	.834	.794	.756	.721	.689	.659	.632	.606	.583	.561	.540	.521	.504	.487
	.571	.493	.436	.391	.355	.325	.300	.279	.260	.244	.230	.217	.206	.196	.185	.178	.170	.163
9	.995	.967	.928	.867	.847	.809	.773	.740	.709	.680	.653	.628	.605	.583	.563	.544	.526	.509
	.606	.530	.473	.427	.390	.360	.333	.311	.291	.274	.259	.245	.233	.222	.212	.202	.194	.186
10	.995	.970	.934	.896	.658	.822	.788	.756	.726	.698	.672	.647	.525	.603	.583	.564	.547	.530
	.636	.562	.505	.460	.423	.391	.364	.341	.320	.302	.286	.271	.258	.246	.236	.226	.217	.208
11	.996	.972	.939	.903	.868	.834	.801	.770	.741	.714	.689	.665	.642	.621	.602	.583	.565	.549
	.661	.590	.534	.489	.452	.420	.392	.368	.347	.328	.311	.296	.282	.270	.256	.246	.238	.229
12	.996	.974	.943	.910	.876	.844	.812	.783	.755	.729	.704	.681	.659	.638	.618	.600	.583	.566
	.684	.615	.560	.516	.478	.446	.418	.394	.372	.353	.335	.319	.305	.292	.280	.269	.259	.250
13	.996	.976	.947	.915	.884	.852	.823	.794	.767	.742	.718	.695	.673	.653	.634	.616	.598	.582
	.703	.637	.583	.539	.502	.470	.442	.417	.395	.375	.356	.341	.327	.313	.301	.289	.279	.269
14	.997	.977	.950	.920	.890	.860	.832	.804	.778	.754	.730	.708	.687	.667	.648	.630	.613	.597
	.721	.656	.604	.561	.524	.492	.464	.439	.417	.397	.379	.362	.347	.333	.320	.308	.297	.287
15	.997	.979	.953	.925	.896	.868	.840	.814	.788	.764	.742	.720	.699	.680	.661	.643	.627	.611
	.736	.674	.623	.581	.544	.513	.484	.460	.437	.417	.398	.382	.366	.352	.339	.327	.315	.305
16	.997	.980	.956	.929	.901	.874	.848	.822	.798	.774	.752	.731	.711	.692	.673	.656	.639	.623
	.750	.690	.641	.599	.553	.531	.504	.479	.456	.436	.417	.400	.384	.370	.357	.344	.333	.322
17	.997	.981	.958	.932	.906	.880	.854	.830	.806	.783	.762	.741	.721	.703	.685	.667	.651	.635
	.762	.704	.656	.616	.560	.549	.521	.496	.474	.453	.435	.417	.402	.387	.373	.351	.349	.338
18	.997	.982	.960	.935	.910	.885	.86l	.837	.814	.792	.771	.750	.731	.713	.695	.678	.662	.647
	.774	.717	.671	.631	.596	.565	.538	.513	.491	.470	.451	.434	418	.403	.389	.377	.365	.353
19	.997	.983	.962	.938	.914	.890	.866	.843	.821	.800	.779	.759	.740	.722	.705	.688	.672	.657
	.784	.729	.684	.645	.611	.580	.553	.529	.506	.486	.467	.450	.434	.419	.405	.392	.380	.368
20	.998	.984	.963	.941	.918	.894	.871	.849	.828	.807	.787	.767	.749	.731	.714	.698	.682	.667
	.793	.741	.696	.658	.625	.595	.568	.543	.521	.501	.482	.464	.448	.433	.419	.406	.394	.382
22	.998	.985	.967	.946	.924	.902	.881	.860	.839	.820	.801	.782	.764	.747	.731	.715	.700	.685
	.810	.760	.718	.682	.649	.620	.594	.570	.546	.528	.509	.492	.476	.461	.446	.433	.421	.409
24	.998	.986	.969	.950	.930	.909	.889	.869	.850	.831	.813	.795	.778	.762	.746	.731	.716	.702
	.824	.777	.737	.702	.671	.643	.617	.594	.572	.552	.534	.517	.500	.485	.471	.458	.445	.433
26	.998	.987	.971	.953	.934	.915	.896	.877	.859	.841	.823	.807	.790	.775	.759	.744	.730	.716
	.836	.792	.754	.720	.690	.663	.638	.615	.594	.575	.556	.539	.523	.508	.494	.481	.468	.456
28	.998	.988	.973	.956	.938	.920	.902	.884	.867	.850	.833	.817	.801	.785	.771	.757	.743	.730
	.847	.805	.768	.736	.707	.681	.657	.635	.614	.595	.577	.560	.544	.529	.515	.501	.489	.477
30	.998	.989	.975	.959	.942	.925	.908	.891	.874	.858	.842	.826	.811	.796	.782	.768	.755	.742
	.856	.816	.782	.751	.723	.697	.674	.652	.632	.613	.595	.579	.563	.548	.534	.521	.508	.496

Table 7-4
Two-Sided Confidence Limits for Binomial Distribution, Confidence Level:1-α=0.95

#Fail-	Number of Demands Minus the Number of Failures																	
ures	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
0	.975	.842	.708	.602	.522	.459	.410	.369	.336	.308	.285	.265	.247	.232	.218	.206	.195	.185
	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000	.000
1	.987	.906	.806	.716	.641	.579	.527	.483	.445	.413	.385	.360	.339	.319	.302	.287	.273	.260
	.013	.008	.006	.005	.004	.004	.003	.003	.003	.002	.002	.002	.002	.002	.002	.001	.001	.001
2	.992	.932	.853	.777	.710	.651	.600	.556	.518	.484	.454	.478	.405	.383	.364	.347	.331	.317
	.094	.068	.053	.043	.037	.032	.028	.025	.023	.021	.019	.018	.017	.016	.015	.014	.013	.012
3	.994	.947	.882	.816	.755	.701	.652	.610	.572	.538	.508	.481	.456	.434	.414	.396	.379	.363
	.194	.147	.118	.099	.085	.075	.067	.060	.055	.050	.047	.043	.040	.038	.036	.034	.032	.030
4	.995	.957	.901	.843	.788	.738	.692	.651	.614	.581	.551	.524	.499	.476	.456	.437	.419	.403
	.284	.223	.184	.157	.137	.122	.109	.099	.091	.084	.078	.073	.066	.064	.061	.057	.054	.052
5	.996	.963	.915	.863	.813	.766	.723	.684	.649	.616	.587	.560	.535	.512	.491	.471	.453	.436
	.359	.290	.245	.212	.187	.167	.151	.139	.128	.118	.110	.103	.097	.091	.087	.082	.078	.075
6	.996	.968	.925	.878	.833	.789	.749	.711	.677	.646	.617	.590	.565	.543	.522	.502	.484	.467
	.421	.349	.299	.262	.234	.211	.192	.177	.163	.152	.142	.133	.126	.119	.113	.107	.102	.098
7	.997	.972	.933	.891	.849	.808	.770	.734	.701	.671	.643	.616	.592	.570	.549	.529	.512	.494
	.473	.400	.348	.308	.277	.251	.230	.213	.198	.184	.173	.163	.154	.146	.139	.132	.126	.121
8	. 997	.975	.940	.901	.861	.823	.787	.753	.722	.692	.655	.639	.516	.593	.573	.553	.535	.518
	.517	.444	.390	.349	.315	.289	.266	.247	.230	.215	.203	.191	.181	.172	.164	.156	.149	.143
9	.997	.977	.945	.909	.872	.837	.802	.770	.740	.711	.685	.660	.636	.615	.594	.575	.557	.540
	.555	.482	.428	.386	.351	.323	.299	.278	.260	.244	.231	.218	.207	.197	.188	.180	.172	.165
10	.998	.979	.950	.916	.882	.848	.816	.785	.756	.728	.702	.678	.655	.634	.614	.595	.577	.560
	.587	.516	.462	.419	.384	.354	.329	.308	.289	.272	.257	.244	.232	.221	.211	.202	.194	.186
11	.998	.981	.953	.922	.890	.858	.827	.797	.769	.743	.718	.694	.672	.651	.631	.612	.594	.578
	.615	.546	.492	.449	.413	.363	.357	.335	.315	.298	.282	.268	.256	.244	.234	.224	.215	.207
12	.998	.982	.957	.927	.897	.867	.837	.809	.782	.756	.732	.709	.687	.666	.647	.628	.611	.594
	.640	.572	.519	.476	.440	.410	.384	.361	.340	.322	.306	.291	.278	.266	.255	.245	.235	.227
13	.998	.983	.960	.932	.903	.874	.846	.819	.793	.768	.744	.722	.701	.680	.661	.643	.626	.609
	.661	.595	.544	.501	.465	.435	.408	.384	.364	.345	.328	.313	.299	.287	.275	.264	.255	.245
14	.998	.984	.962	.936	.909	.881	.854	.828	.803	.779	.756	.734	.713	.694	.675	.657	.640	.624
	.681	.617	.566	.524	.488	.457	.430	.407	.385	.366	.349	.334	.320	.306	.295	.283	.273	.264
15	.998	.985	.964	.939	.913	.887	.861	.836	.812	.789	.766	.745	.725	.705	.687	.669	.653	.637
	.698	.636	.586	.544	.509	.478	.451	.427	.406	.386	.369	.353	.339	.325	.313	.302	.291	.281
16	.999	.986	.965	.943	.918	.893	.868	.844	.820	.798	.776	.755	.736	.717	.698	.681	.665	.649
	.713	.653	.604	.563	.529	.498	.471	.447	.425	.405	.388	.372	.357	.343	.331	.319	.308	.298
17	.999	.987	.968	.946	.922	.898	.874	.851	.828	.806	.785	.765	.745	.727	.709	.692	.676	.660
	.727	.669	.621	.58l	.547	.516	.488	.465	.443	.423	.406	.389	.374	.360	.347	.335	.324	.314
18	.999	.988	.970	.948	.925	.902	.879	.857	.835	.814	.793	.773	.755	.736	.719	.702	.686	.671
	.740	.683	.637	.597	.564	.533	.506	.482	.460	.440	.422	.406	.391	.376	.363	.351	.340	.329
19	.999	.988	.971	.950	.929	.906	.884	.862	.841	.821	.801	.782	.763	.745	.728	.712	.696	.681
	.751	.696	.651	.612	.579	.549	.522	.498	.476	.456	.439	.422	.406	.392	.379	.366	.355	.344
20	.999	.989	.977	.953	.932	.910	.889	.868	.847	.827	.808	.789	.771	.753	.737	.720	.705	.690
	.762	.708	.664	.626	.593	.564	.537	.513	.492	.472	.454	.437	.421	.407	.393	.381	.369	.356
22	.999	.990	.975	.956	.937	.917	.897	.877	.858	.839	.820	.803	.785	.768	.752	.737	.722	.707
	.781	.730	.688	.651	.619	.590	.565	.541	.519	.500	.481	.465	.449	.434	.421	.408	.396	.385
24	.999	.991	.976	.960	.942	.923	.904	.885	.867	.849	831	.814	.798	.782	.766	.751	.737	.723
	.797	.749	.708	.673	.642	.614	.589	.566	.545	.525	.507	.490	.475	.460	.446	.433	.421	.410
26	.999	.991	.976	.962	.945	.928	.910	.893	.875	.858	.841	.825	.809	.794	.779	.764	.750	.736
	.810	.765	.726	.693	.663	.636	.611	.588	.567	.548	.530	.513	.497	.483	.469	.456	.444	.432
28	.999	.992	.980	.965	.949	.932	.916	.899	.882	.866	.850	.834	.819	.804	.790	.776	.762	.749
	.822	.779	.743	.710	.681	.655	.631	.609	.588	.569	.551	.535	.519	.504	.491	.478	.465	.453
30	.999	.992	.981	.967	.952	.936	.920	.904	.889	.873	.858	.843	.826	.814	.800	.786	.773	.760
	.833	.792	.757	.725	.697	.672	.649	.627	.607	.588	.571	.554	.539	.524	.510	.498	.485	.473

Table 7-5
Values of P for Which the Cumulative Fraction of the Area Under the Beta Distribution Equals 2.5%—That Is, for a Confidence Level of 95%

Value of V	Value of W											
	1	2	3	4	5	6	10	12	15	20	30	60
1	0.02500	0.01258	0.00840	0.00631	0.00505	0.00421	0.00253	0.00211	0.00169	0.00126	0.00084	0.00042
2	0.15811	0.09429	0.06758	0.05274	0.04327	0.03669	0.02283	0.01921	0.0l551	0.01175	0.00791	0.00399
3	0.29240	0.19412	0.14663	0.11812	0.09899	0.08523	0.05486	0.04658	0.03798	0.02906	0.01977	0.01009
4	0.39764	0.28358	0.22278	0.18405	0.15701	0.13700	0.09092	0.07787	0.06409	0.04951	0.03403	0.01757
5	0.47818	0.35877	0.29042	0.24486	0.21201	0.18709	0.12760	0.11017	0.09147	0.07132	0.04953	0.02585
6	0.54074	0.42128	0.34914	0.29930	0.26238	0.23379	0.16336	0.14210	0.11893	0.09356	0.06562	0.03463
7	0.59038	0.47349	0.39991	0.34755	0.30790	0.27667	0.19753	0.17299	0.14588	0.11573	0.08194	0.04372
8	0.63058	0.51750	0.44390	0.39026	0.34888	0.31578	0.22983	0.20252	0.17198	0.13753	0.09827	0.05298
9	0.66373	0.55498	0.48224	0.42814	0.38574	0.35138	0.26019	0.23058	0.19708	0.15878	0.11444	0.06235
10	0.69150	0.58722	0.51586	0.46187	0.41896	0.38380	0.28864	0.25713	0.22110	0.17938	0.13038	0.07175
11	0.71509	0.61520	0.54553	0.49202	0.44900	0.41338	0.31528	0.28221	0.24402	0.19930	0.14601	0.08114
12	0.73535	0.63970	0.57187	0.51911	0.47623	0.44042	0.34021	0.30588	0.26587	0.21850	0.16130	0.09050
13	0.75295	0.66132	0.59540	0.54354	0.50101	0.46520	0.36355	0.32821	0.28667	0.23698	0.17622	0.09979
14	0.76836	0.68052	0.61652	0.56568	0.52363	0.48797	0.38542	0.34928	0.30647	0.25476	0.19076	0.10901
15	0.78198	0.69768	0.63559	0.58582	0.54435	0.50895	0.40594	0.36918	0.32532	0.27185	0.20492	0.11812
20	0.83157	0.76184	0.70839	0.66411	0.62616	0.59296	0.49168	0.45370	0.40697	0.34780	0.26997	0.16201
30	0.88430	0.83298	0.79193	0.75669	0.72550	0.69743	0.60674	0.57056	0.52422	0.46239	0.37498	0.24027
60	0.94037	0.91201	0.88828	0.86708	0.84764	0.82954	0.76678	0.73968	0.70299	0.65017	0.56658	0.41107
∞	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000

Table 7-6
Values of P for Which the Cumulative Fraction of the Area Under the Beta Distribution Equals 5.0%—That Is, for a Confidence Level of 90%

Value of V	Value of W											
	1	2	3	4	5	6	10	12	15	20	30	60
1	0.05000	0.02532	0.01695	0.01274	0.01021	0.00851	0.00512	0.00426	0.00341	0.00256	0.00170	0.00085
2	0.22361	0.13535	0.09761	0.07644	0.06285	0.05337	0.03332	0.02805	0.02268	0.01719	0.01158	0.00585
3	0.36840	0.24860	0.18926	0.15316	0.12876	0.11111	0.07187	0.06110	0.04990	0.03822	0.02604	0.01332
4	0.47287	0.34259	0.27134	0.22532	0.19290	0.16875	0.11267	0.09666	0.07969	0.06167	0.04248	0.02198
5	0.54928	0.41820	0.34126	0.28924	0.25137	0.22244	0.15272	0.13211	0.10991	0.08588	0.05978	0.03129
6	0.60696	0.47930	0.40031	0.34494	0.30354	0.27125	0.19086	0.16636	0.13955	0.11006	0.07739	0.04097
7	0.65184	0.52932	0.45036	0.39338	0.34981	0.31524	0.22669	0.19895	0.16818	0.13377	0.09499	0.05085
8	0.68766	0.57086	0.49310	0.43563	0.39086	0.35480	0.26011	0.22972	0.19556	0.15682	0.11240	0.06082
9	0.71687	0.60584	0.52991	0.47267	0.42738	0.39041	0.29120	0.25865	0.22164	0.17908	0.12950	0.07082
10	0.74113	0.63564	0.56189	0.50535	0.45999	0.42256	0.32009	0.28580	0.24639	0.20050	0.14622	0.08079
11	0.76160	0.66132	0.58990	0.53434	0.48925	0.45165	0.34693	0.31126	0.26985	0.22106	0.16252	0.09070
12	0.77908	0.68366	0.61461	0.56022	0.51560	0.47808	0.37190	0.33515	0.29208	0.24078	0.17838	0.10052
13	0.79418	0.70327	0.63656	0.58343	0.53945	0.50217	0.39516	0.35756	0.31314	0.25966	0.19379	0.11024
14	0.80736	0.72060	0.65617	0.60436	0.56112	0.52420	0.41685	0.37862	0.33309	0.27775	0.20875	0.11983
15	0.81896	0.73604	0.67381	0.62332	0.58088	0.54442	0.43711	0.39842	0.35200	0.29507	0.22326	0.12930
20	0.86089	0.79327	0.74053	0.69636	0.65819	0.62460	0.52099	0.48175	0.43321	0.37136	0.28936	0.17453
30	0.90497	0.85591	0.81606	0.78150	0.75070	0.72282	0.63185	0.59522	0.54807	0.48477	0.39458	0.25416
60	0.95130	0.92458	0.90192	0.88150	0.86266	0.84504	0.78342	0.75661	0.72016	0.66738	0.58326	0.42519
∞	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.0000	1.00000	1.00000

Table 7-7
Values of P for Which the Cumulative Fraction of the Area Under the Beta Distribution Equals 10.0%—That Is, for a Confidence Level of 80%

Value of V	Value of W											
	1	2	3	4	5	6	10	12	15	20	30	60
1	0.10000	0.05132	0.03451	0.02599	0.02085	0.01741	0.01048	0.00874	0.00700	0.00525	0.00351	0.00175
2	0.31623	0.19580	0.14256	0.11224	0.09259	0.07882	0.04945	0.04169	0.03375	0.02562	0.01729	0.00875
3	0.46416	0.32046	0.24664	0.20091	0.16964	0.14685	0.09565	0.08148	0.06667	0.05117	0.03494	0.01791
4	0.56234	0.41611	0.33319	0.27860	0.23966	0.21040	0.14161	0.12177	0.10064	0.07808	0.05393	0.02798
5	0.63096	0.48968	0.40382	0.34462	0.30097	0.26732	0.18513	0.16056	0.13394	0.10497	0.07330	0.03847
6	0.68129	0.54744	0.46178	0.40058	0.35422	0.31772	0.22559	0.19716	0.16587	0.13123	0.09260	0.04921
7	0.71969	0.59375	0.50992	0.44827	0.40053	0.36228	0.26292	0.23139	0.19619	0.15659	0.11161	0.05999
8	0.74989	0.63164	0.55040	0.48924	0.44100	0.40176	0.29726	0.26327	0.22483	0.18093	0.13019	0.07077
9	0.77426	0.66315	0.58484	0.52473	0.47657	0.43689	0.32885	0.29293	0.25182	0.20420	0.14828	0.08148
10	0.79433	0.68976	0.61448	0.55574	0.50803	0.46829	0.35793	0.32051	0.27721	0.22642	0.16583	0.09208
11	0.81110	0.71250	0.64022	0.58302	0.53603	0.49649	0.38475	0.34619	0.30111	0.24759	0.18283	0.10257
12	0.82540	0.73216	0.66279	0.60721	0.56108	0.52193	0.40954	0.37012	0.32361	0.26778	0.19928	0.11290
13	0.83768	0.74933	0.68271	0.62878	0.58361	0.54498	0.43248	0.39245	0.34481	0.28701	0.21518	0.12308
14	0.84834	0.76443	0.70044	0.64813	0.60398	0.56595	0.45378	0.41332	0.36479	0.30534	0.23054	0.13310
15	0.85770	0.77783	0.71630	0.66559	0.62247	0.58511	0.47359	0.43286	0.38366	0.32283	0.24539	0.14295
20	0.89125	0.82706	0.77578	0.73219	0.69412	0.66034	0.55476	0.51428	0.46386	0.39910	0.31243	0.18960
30	0.92612	0.88023	0.84212	0.80864	0.77851	0.75104	0.66029	0.62333	0.57545	0.51067	0.41750	0.27063
60	0.96235	0.93773	0.91643	0.89702	0.87897	0.86198	0.80192	0.77553	0.73946	0.68688	0.60235	0.44158
∞	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000	1.00000

8
REFERENCES

1. *Nuclear Power Plant Common Aging Terminology,* EPRI, Palo Alto, CA: 1992. TR-100844.

2. F. S. Nowlan, and H. F. Heap, "Reliability Centered Maintenance." National Technical Information Service Report No. AD/A066-579, December 1978.

3. *Reliability and Risk Significance: For Maintenance and Reliability Professionals at Nuclear Power Plants,* EPRI, Palo Alto, CA: 2002. 1007079.

4. Federal Register, Vol. 53, No. 56, March 23, 1988. "Rules and Regulations," page 9340.

5. The EPRI PM Basis Database, EPRI product 1003282, and The EPRI PM Basis Database: User's Manual, EPRI, Palo Alto, CA: 2001, EPRI product 1001448.

6. *Preventive Maintenance Basis Guidelines*. EPRI, Palo Alto, CA: 1999. TR-106857 Volumes 1-38.

7. F. S. Nowlan and H. F. Heep, "Reliability-Centered Maintenance." Department of Defense Report AD-A066579, 1978.

8. Reliability-Centered Maintenance Requirements for Naval Aircraft, Weapons Systems, and Support Equipment. MIL-STD 2173(AS), U.S. Naval Air Systems Command.

9. Evaluation Criteria for Reliability Centered Maintenance (RCM) Processes. SAE Standard JA1011, August 1999. Society of Automotive Engineers, Warrendale, PA.

10. J. Mowbray. RCM II, Reliability-centered Maintenance. Industrial Press, 1997.

11. A. M. Smith. Reliability-Centered Maintenance. McGraw-Hill, 1993.

12. *Comprehensive Low-Cost Reliability Centered Maintenance,* EPRI, Palo Alto, CA: 1995. TR-105365.

13. *Streamlined Reliability Centered Maintenance (SRCM) Implementation Guidelines,* EPRI, Palo Alto, CA: 1998. EPRIGEN Report TR-109795-V2.

14. "Industry Guideline for Monitoring the Effectiveness of Maintenance at Nuclear Power Plants," NUMARC 93-01. Nuclear Energy Institute, Washington, DC.

15. R. B. Abernethy, The New Weibull Handbook, Fourth Edition. Gulf Publishing, 2000.

16. M. R. Corio, and L. P. Costantini, "Frequency and Severity of Forced Outages Immediately Following Planned and Maintenance Outages," Generating Availability Trends Summary Report, NERC, 1989.

17. Equipment Reliability Process Description, INPO AP-913. Institute of Nuclear Power Operations, March 2000.

18. H. Procaccia, S. Arsenis, P. Aufort, and G. Volta. European Industry Reliability Data Bank. Crete University Press, 1998.

19. TUD Reliability, Maintenance, and Operation Database. TUD Office, SwedPower AB, P.O. Box 527, SE-162 16 Stockholm, Sweden.

20. Verlag Technisch-Wissenschaftlicher Schriften. VGB PowerTech Service GmbH, Postfach 10 39 32, D-45039 Essen, Germany.

21. A. Villemeur. Reliability, Availability, Maintainability, and Safety Assessment, Volume 1, Methods and Techniques. John Wiley and Sons, 1991.

22. Mann, Schafer, and Singpurwalla. Methods for Statistical Analysis and Life Data. John Wiley, 1974.

A
GENERIC MODELS FOR THE DEPENDENCE OF RELIABILITY ON PREVENTIVE MAINTENANCE INTERVALS

A.1 Basis for the Generic Approach

This appendix calculates the reliability of generic components as a function of preventive maintenance parameters. It includes the effects of changes in a PM task interval, the effectiveness of the PM task, and both random and wear-out failure mechanisms. It is applicable to either a single component with a single definite task interval, or to a group of components, which may have different task intervals but whose actual task execution times are distributed around and shifted from the designated intervals.

The method depends on a few observations that stemmed from the EPRI PM Basis database. A complex component (for example, a motor or valve) has a large number of failure mechanisms, divided between wear-out failure mechanisms (which have an expected period of failure-free operation before failures start to be observed), and random failure mechanisms (which can occur at any time). Further, the expected failure-free periods for the wear-out mechanisms seem to occupy all time scales available—that is, they range from less than one year to the design life of the equipment, say 40 years.

These failure mechanisms are actually a combination of a hardware (subcomponent) location of what fails (for example, a switch), a physical mechanism (for example, bent or damaged, misadjusted, worn, contaminated, or failed insulation), and the influences that may drive these occurrences, such as maintenance error, normal use, dirty environment, or heat. We will use the term *failure mechanism* for each combination of circumstances. A complex component may have hundreds of such failure mechanisms, although they can usually be grouped into about 20 major wear-out mechanisms and a similar number of random mechanisms. The details depend on the component.

At the end of the useful life (often referred to here as the failure-free period), a given wear-out mechanism has some probability each year of producing a failure. If you waited long enough (this could be 100 years or more), and did no PM, you could be pretty certain that each such mechanism would have produced a failure. Generally, mechanisms with short failure-free periods (for example, 1 year) will have higher subsequent annual failure probabilities over a shorter period of time (that is, the failure time distribution will be narrower) than mechanisms with longer failure-free periods (for example, 15 years).

If you want to calculate the expected number of failures per year when no preventive maintenance is performed, you would expect this failure time distribution to play a major role. Since no one normally possesses detailed information on the failure time distributions, such an RTF (run-to-failure) reliability prediction is not normally attempted. However, if a PM task is performed on a regular schedule, with an interval not too different from the failure-free period, it is clear that a reliability calculation will depend far less on the details of the failure time distribution. This is because an effective PM task will discover emerging degradation and correct it, thus restoring the component to something approaching an as-new condition. This will take place in the early part of the failure time distribution, so the bulk of its form and magnitude will scarcely be sampled in this situation.

Failure rates are also affected by whether the PM task is actually always performed, or performed on time, as well as by personnel errors which result in degradation not being recognized, repairs which are ineffective, or defects and faults being introduced during the tasks, as well as by task intervals which are longer than they should be, and by mechanisms which are not addressed by any task. Furthermore, although experienced maintenance personnel usually have a good idea of the most likely failure-free periods to expect, mechanisms can still be affected by many factors which change their failure-free intervals in ways which are hard to predict.

These observations suggest that a realistic maintenance decision model could be constructed using a uniform distribution of failure-free periods to represent the whole set of many wear-out failure mechanisms, and an overall effectiveness for each PM task. This effectiveness, E, would be the probability of diagnosing degradation and successfully correcting it, when such degradation exists and the task is performed. We would expect this parameter to be in the range of 75% to 95% for reasonably effective tasks (the PM Basis database uses 95% for highly effective and 75% for moderately effective tasks).

An individual wear-out failure mechanism with failure-free interval n years has a failure time distribution taken as uniform starting at n years and stretching out for another 2n years, so that normalization requires it to have an annual failure probability of 1/2n per year. All failure rates calculated in the model will be proportional to this probability, but the results are presented as ratios of failure rates so that the impact of this assumption is greatly reduced.

Assuming that N_w mechanisms are active for a component, and that these have failure-free periods uniformly distributed between some minimum, m years, and an upper limit of 40 years, there will be $N_w/(40\text{-}m)$ mechanisms “starting up” each year on average.

A.2 The Effective Maintenance Model—EM

Suppose a component is provided with a PM task at an interval of I years because it is known to begin experiencing failures from the shortest-term wear-out mechanism at m=I years. Such a component could be said to have the most effective PM possible, because the task is not being done too often, but on the other hand, it is always done just in time to intercept degradation from the earliest failure mechanism. Even then, it is possible that the task is not done well, or is ill-adapted to the failure mechanism. For this reason we assign a maintenance effectiveness, E, to

the task. This means that just after the task is performed, the degradation is still present with a probability (1-E), and can continue to cause failures.

In the interval between I and 2I, we expect to get (1-E)/2n failures per year from any mechanism with failure-free period n, and these will endure for (2I-n) years until the task is performed again. We expect this to happen every interval, so the failure rate for a single mechanism is thus:

$$\Lambda_{el} = (1-E)(2I-n)/2nI \qquad \textbf{Eq. A-1}$$

A chart of this relation against an accurate solution to the underlying alternating renewal process is shown in Figure A-1 for I = 5 years. Note that the renewal solution gives a larger rate because it includes contributions of order $(1\text{-}E)^2$ and higher.

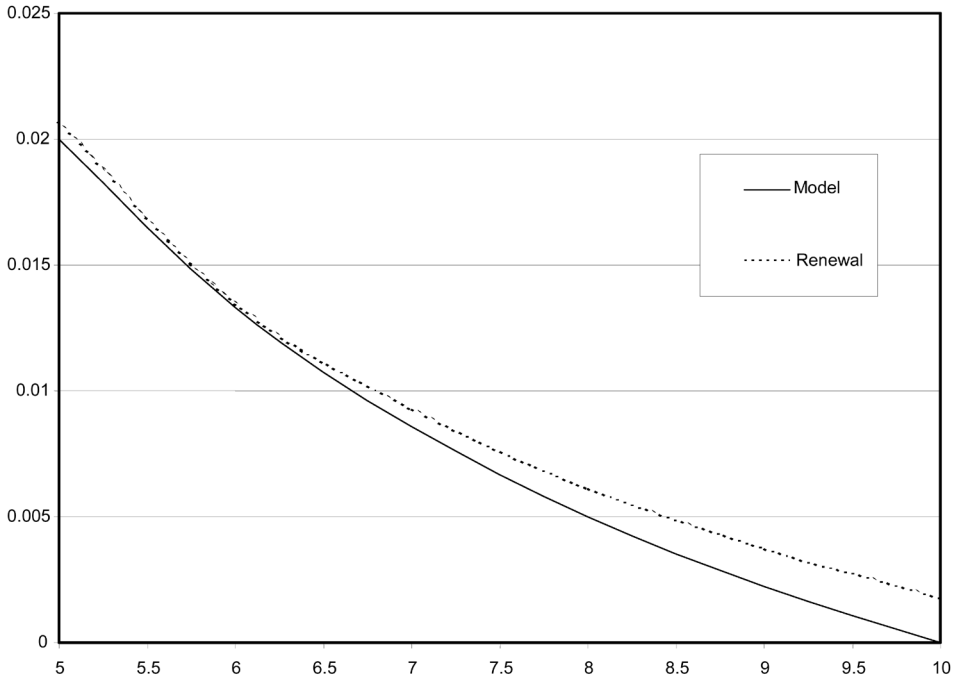

Figure A-1
Comparison of the Single Mode Model With an Accurate Renewal Solution

The chart shows that the model is only a few percent non-conservative (i.e. predicting low) for shorter-term mechanisms, and becomes more so for longer-term mechanisms. When this result is integrated across a spectrum of mechanisms, the shorter ones dominate, giving a result, below, that is a reasonable representation (that is, within a few percent) of the underlying renewal process.

Since N_w.dn/(40-I) mechanisms start up in dn years, the total contribution to the failure rate from all mechanisms that can contribute is:

$$\Lambda_e = \left[(1-E)\, N_w / 2I(40-I)\right] \cdot \int^{2I} I\,(2I-n)dn/n$$

$$\Lambda_e = (1-E)\cdot N_w \cdot (2\,1n2-1)/\,2(40-I) = 0.193(1-E)\cdot N_w /(40-I) \qquad \textbf{Eq. A-2}$$

Since I is usually much less than 40 years, the dependence on task interval is weak. Equation A-2 gives a failure rate of about 0.02 failures per year for effective PM with intervals up to 20 years, when E=80% and N_w = 20.

To this must be added the random rate, which cannot be protected with time-directed PM tasks. If we assume that it is not cost-effective to continue to reduce the wear-out failures below the level of the remaining random contributions, we would conclude that $\Lambda_r = B\Lambda_e$ with B ~ 1 or 2, so that the total effective maintenance failure rate must be close to $(1 + B)\Lambda_e \sim 0.04$ or 0.06 failures per year (that is, 17 to 25 years between failures when well maintained). This is within the range of operational experience.

A.3 The Risk of Performing PM

Intrusive PM tasks run a risk of introducing additional failures. A simple treatment enables the most important conclusion to be drawn. Consider that performing an intrusive task introduces an additional failure with a probability P_{im} (subscript for infant mortality). This applies each time the task is performed, so it increases the failure rate on average by $\Lambda_{im} = P_{im} / I$, where I is the task interval. The parameter P_{im} is perhaps in the range of 5% to 15% for a wide range of equipment.

A value of P_{im} = 0.1 with a 5-year interval adds Λ_{im} = 0.02 failures/year to $(1 + B)\Lambda_e$, above—an amount that equals the effective maintenance failure rate. If the interval is unnecessarily decreased from 5 years to 4 years, Λ_{im} will increase by 20%, a significant erosion of effective PM. Other values of P_{im} and I give a similar conclusion.

A.4 The Run-to-Failure Model—RTF

When there is no PM but failures are repaired in a time short compared to the mean time between failures, renewal theory provides an asymptotic solution for the above single mechanism that has a uniform time to failure distribution from n years to 3n years.

$$\Lambda_{RTF,1} = 2/(n+3n) = 1/2n \qquad \textbf{Eq. A-3}$$

If this also is integrated over N_w mechanisms:

$$\Lambda_{RTF} = \left[N_w / 2(40 - m)\right] \cdot \int_m^{40} dn / n$$

For m ~1 year we get:

$$\Lambda_{RTF} = 0.047 \cdot N_w \qquad \textbf{Eq. A-4}$$

This gives ~ 1 per year for 20 mechanisms and, with the result of Equation A-2, shows that effective PM reduces the failure rate by factors between 14 and 42 for E in the range 0.8 to 0.9 and B in the range 1 to 2, independent of the number of mechanisms. This result underlines the immense value of an effective PM program.

$$\textit{Failure Rate Reduction Factor By Effective PM} = 9.55 / (1 - E)(1 + B) \qquad \textbf{Eq. A-5}$$

These results are reasonable, but it is clear that we must account for the effectiveness of the task and also the level of random failures in the development which follows.

The general approach from this point is to develop a "Missed Modes Model" which will add the effect of failure mechanisms which cause failures because the PM interval is too long. Such mechanisms have nothing to prevent them from occurring and can greatly increase the failure rate. Armed with the Effective Maintenance and Missed Modes results we will later impose a statistical distribution of times at which the tasks actually get performed for application to the problem of grace periods.

A.5 The Missed Modes Model—MM

We envisage a single component with a set of failure mechanisms, as before. The shortest failure-free interval is at m years. The others are distributed uniformly between m and 40 years. Since the missed modes PM task interval is I>m, mechanisms with

$$m < \textit{the failure free interval} < I$$

are "missed" by the task and so are not attenuated by the factor (1-E). Mechanisms with failure-free interval > I would be treated as effective maintenance in the manner described for the EM model, above.

As before, failures from missed modes will accrue in the first interval at 1/2n per year per mechanism, for a total of (I-n) years. Integrating over all contributing mechanisms gives a failure rate of:

$$\Lambda_m = \left[N_w \cdot / 2I(40-m)\right] \cdot \int_m^I (I-n)dn/n$$

$$\Lambda_m = \left[N_w \cdot /2(40-m)\right] \cdot \left[m/I - (1+1n(m/I))\right] \qquad \textbf{Eq. A-6}$$

or

$$\Lambda_m = [N_w \cdot /2(40-\alpha I)].[\alpha - (1+1n\,\alpha)] \qquad \propto = m/I$$

The modified effective maintenance contribution to be added to Λ_m is given by:

$$\Lambda_e = [(1-E).N_w /2I(40-m)] \cdot \int_I^{2I} (2I-n)\,dn/n$$

so that

$$\Lambda_m = [N_w /2(40-\alpha I)].[\alpha - (1+1n\,\alpha)] + \Lambda_\varepsilon (40-I)/(40-\alpha I) \qquad \textbf{Eq. A-7}$$

We should also add the random contribution (=$B\Lambda_e$) as before:

$$\Lambda_{Tm} = [N_w /2(40-\alpha I)].[\alpha - (1+1n\,\alpha)] + \Lambda_e [(40-I)/(40-\alpha I) + B] \qquad \textbf{Eq. A-8}$$

The ratio of this total rate to the effective maintenance and random rate is:

$$Ratio = \{[N_w /2(40-\alpha I)].[\alpha - (1+1n\alpha)] + \Lambda_e [(40-I)/(40-\alpha I) + B]\} / \Lambda_e (1+B)$$

Eq. A-9

with Λ_e given by Equation A-2. The value of this ratio, minus 1, shows the fractional increase in the rate when the interval is such that some failure modes are left unprotected by the PM task. This ratio was used in Section 3.3 to guide the strategy for task deferral, and it is used in the EPRI PM Basis Application Guideline sections on adjusting task intervals and task deferral. Figures A-2, A-3, and A-4 show the ratio for different values of the parameters.

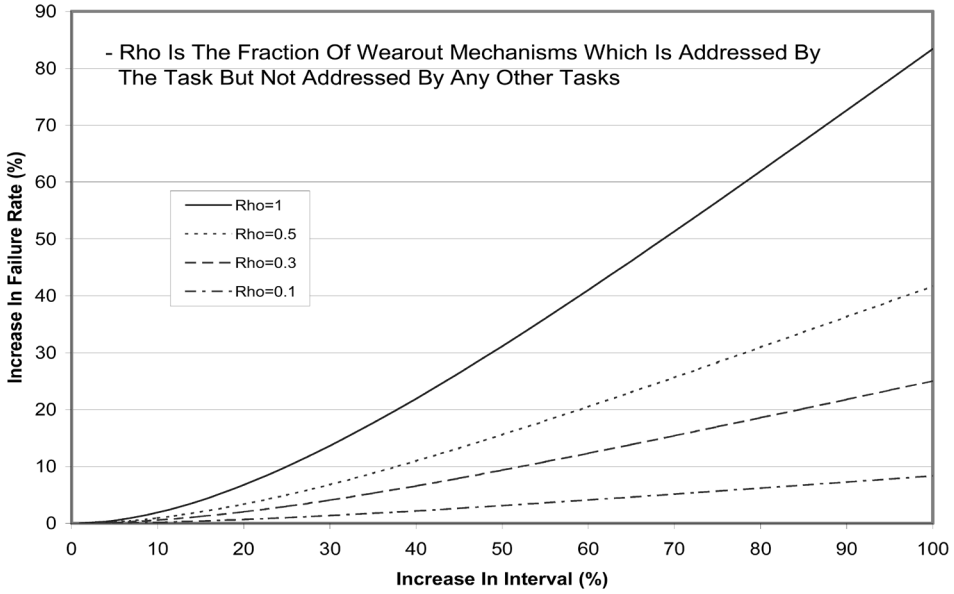

Figure A-2
Increase in Failure Rate Versus Increase in Task Interval for Task Effectiveness = 70%

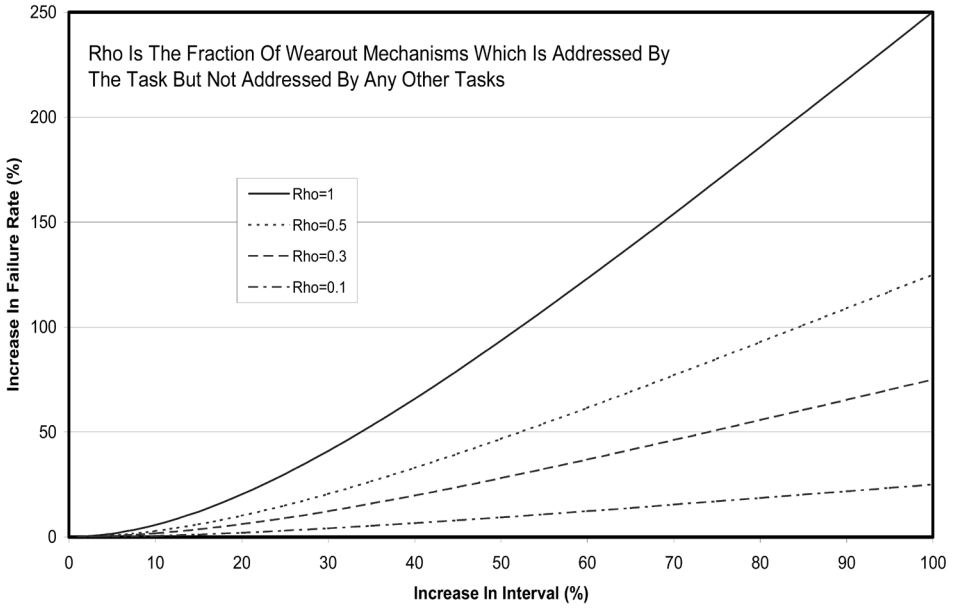

Figure A-3
Increase in Failure Rate Versus Increase in Task Interval for Task Effectiveness = 90%

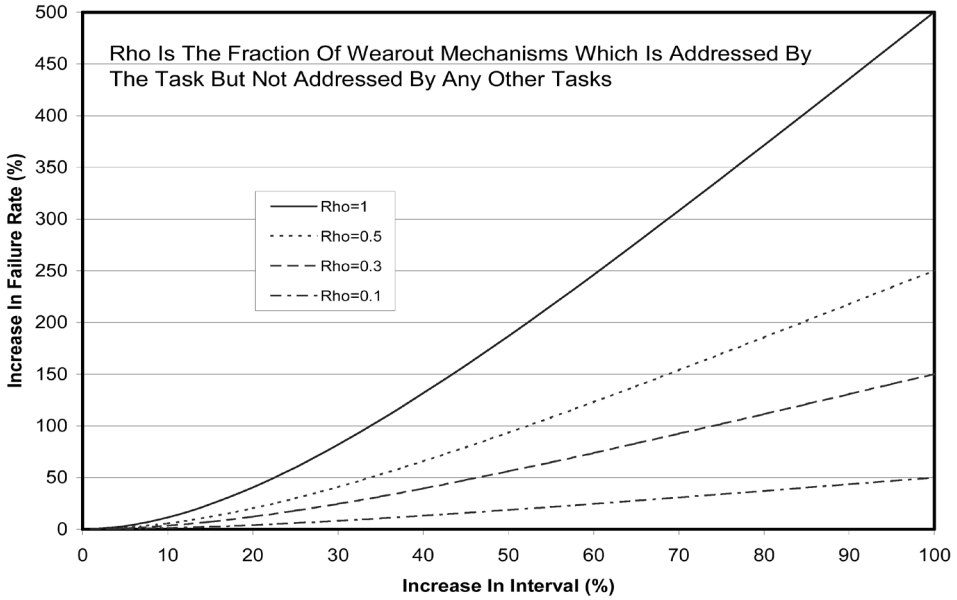

Figure A-4
Increase in Failure Rate Versus Increase in Task Interval for Task Effectiveness = 95%

A.6 The Statistical Model

A statistical version of the above failure rate ratio is needed when, for whatever reason, a large number of components have PM tasks actually performed at times different from their intended intervals, but distributed around the intended intervals. To address the general case where the intended intervals can have a range of values within the group considered, it will be necessary to cast the failure rate ratio in terms of the ratio of the actual performance time to the intended time.

The statistical model introduces a group of components which have individual task performance times, x_j, which do not necessarily equal their own task interval, I_j. Throughout, it will be convenient to use the dimensionless parameter $\gamma = (x_j - I_j)/I_j$ to represent the fraction by which any given task time exceeds its interval. Consequently, $\gamma = 0,1$ for all $x_j = I_j$, $2I_j$, and κ is the standard deviation in the space of an assumed normal distribution of γ. There is no further need to use the subscripts j. We assume that the task interval represents the effective maintenance case, that is, the shortest mode occurs at I; see Figure A-5. Components whose PM task is performed before this time are effectively being dealt with by effective maintenance because all mechanisms are longer than the task time. Components whose task time is later than the designated interval possess missed mechanisms, and so are treated with the missed mode model, plus a modified effective maintenance model for the mechanisms which arise after the task performance.

Complications which arise include 1) a normalization shift when the tails of the (infinite) normal distribution overlap the practical bounds of the problem at x = 0 and x = 2I; and 2) whether to correct the effective maintenance model to allow for the restricted range of mechanisms, since the shortest mode is at I > task time.

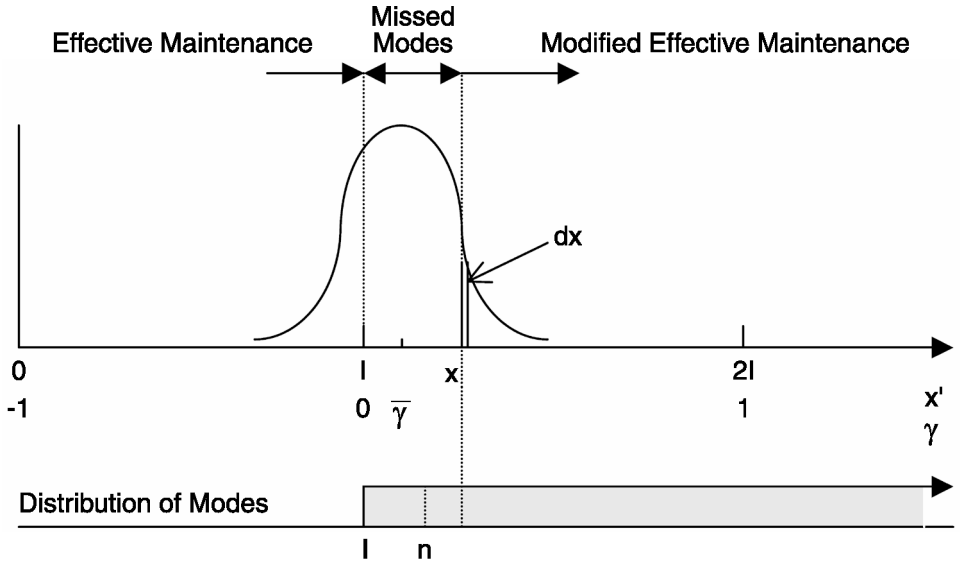

Figure A-5
Contributions to the Statistical Model

A.7 Missed Mode Contribution

When x > I, a missed mechanism at n provides an annual contribution of 1/2n failures for (x-n) years. Modes with n from I to x contribute. The rate is thus:

$$\Lambda_m(x)' = [N_w / (2I \cdot (40 - I)]. \int_I^x |(x-n)\, dn\ /n)$$

$$\Lambda_m(\gamma)' = [N_w / (2 \cdot (40 - I)]. [(1+\gamma)\, In\,(1+\gamma)\ -\ \gamma] \qquad \textbf{Eq. A-10}$$

The average of this quantity over the N components contributing to it is:

$$\Lambda_m(\bar{\gamma}, k) = [N/\varphi]. \int_0^I f(\gamma)\, \Lambda_m(\gamma)'\, d\gamma \qquad \textbf{Eq. A-11}$$

where N is the number of components in the population, ϕ is a normalization adjustment, and f(γ) is the population normal distribution. ϕ and f(γ) are given by:

$$f(\gamma) = [1/(k\sqrt{2\pi})] exp\,[-(\gamma - \bar{\gamma})^2 / k^2]\ d\gamma \qquad \textbf{Eq. A-12}$$

$$\varphi = \int_{-I}^{+I} f(\gamma)\, d\gamma \qquad \textbf{Eq. A-13}$$

A.8 Modified Effective Maintenance Contribution

The region to the right of x in Figure A-5 represents effective maintenance, but only for mechanisms arising at times greater than x. Contributing mechanisms each add failures for a time (x + I – n), and these contributions should be integrated from x to (x + I):

$$\Lambda_{e2}(x)' = [N_w .(1-E)/(2I.(40-I)]. \int_x^{x+I} (x+I-n)\, dn \;/n$$

$$\Lambda_{e2}(\gamma)' = [N_w .(1-E)/(2.(40-I)]. [(2+\gamma) ln\{(2+\gamma)/(1+\gamma)\} - I] \qquad \textbf{Eq. A-14}$$

The average of this quantity over the N components contributing to it is:

$$\Lambda_{e2}(\bar{\gamma}, k) = [N/\phi]. \int_0^{+I} f(\gamma)\, \Lambda_{e2}(\gamma)'\, d\gamma \qquad \textbf{Eq. A-15}$$

A.9 Effective Maintenance Contribution

Components which have the PM task performed earlier than the shortest mechanism, that is, with x < I, have effective PM, with two qualifications. The first is that performing the task too frequently can add a significant number of failures, increasing N_w and B in a way that cannot be modeled. The second is that when a mechanism adds failures, attenuated by the factor (1-E), the mechanisms that contribute are not the full spectrum from x to (x + I), as before, because the shortest mechanism is at I > x. These two effects oppose each other. The second can be calculated, but without the first the failure rate would be artificially reduced.

Consequently, the contribution from these components has been assumed to be the normal effective maintenance rate of Equation A-2 times the number of components in this part of the group:

$$\Lambda_{e1}(\bar{\gamma}, k) = [N\, \Lambda_e /\varphi]. \int_{-1}^{0} f(\gamma) d\gamma \qquad \textbf{Eq. A-16}$$

A.10 Total Rate and Excess Ratio

The new total failure rate, $\Lambda_T(\bar{\gamma}, \kappa)$, is obtained by adding the separate rates from Equations A-16, A-15, and A-11, plus the random contribution, $BN\Lambda_e$:

$$\Lambda_T(\bar{\gamma}, k) = \Lambda_{e1}(\bar{\gamma}, k) + \Lambda_{e2}(\gamma, k) + \Lambda_m(\bar{\gamma}, k) + BN\Lambda_e \qquad \textbf{Eq. A-17}$$

These rates are all functions of N_w and N.

A fractional excess failure rate can be examined to find the percentage change in the rate compared to having all components maintained at the effective maintenance rate:

$$Excess\ Ratio\ (\bar{\gamma}, k, B, E) = [\Lambda_T(\bar{\gamma}, k) - (1+B)N\Lambda_e] / (1+B)N\Lambda_e \qquad \textbf{Eq. A-18}$$

The Excess Ratio does not depend on N_w, or on N, because they cancel in the ratio. The excess ratio is the statistical model analogue of the failure rate ratio of Equation A-9.

A.11 Results

Figure A-6 shows the percentage Excess Ratio for four sets of E;B values as a function of standard deviation, when the population mean is in fact equal to the designated task intervals.

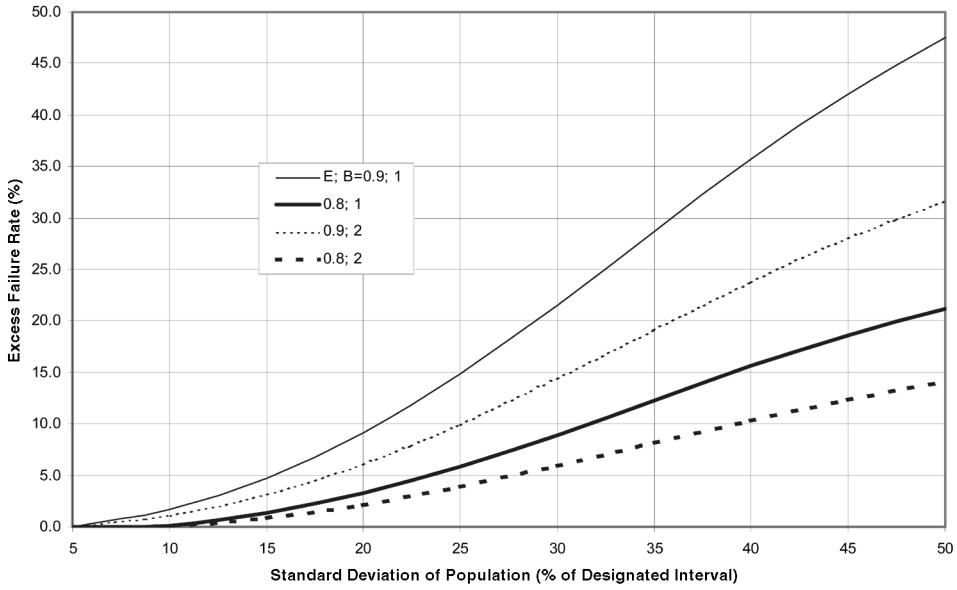

Figure A-6
Excess Failure Rate (%) Versus Standard Deviation When Mean of Deviations of Task Time From Designated Interval Equals Zero

When the population standard deviation is 25% of the designated intervals (25 on the x-axis), the worst case shown has an increase of 15% in the number of failures per year.

If the initial base case were less effective than an optimized PM program, this change would be smaller. For 80% effective tasks, and a random contribution that is twice the effective maintenance rate, the population would need to spread to a standard deviation of more than 50% of the intervals in order to increase the failure rate by 15%.

Figures A-7 and A-8 show the general behavior of the Excess Ratio. Using these results, and the fact that 16% of a normal distribution lies beyond one standard deviation, suggests the rule of thumb:

"Provided no more than 15% of PM tasks are executed beyond 125% of the optimal interval, the increase in failure rate will most likely be less than 20%."

Of course, these rules hold only as well as our assumptions about the values of E and B, a normal distribution of task execution times, and a uniform distribution of failure-free periods for wear-out failure modes. However, these were reasonable assumptions for generic PM programs and generic components. Furthermore, the baseline PM program was assumed to be well optimized (that is, it had the intended interval set to the shortest failure-free interval) to give the most sensitivity to the distribution of task times.

For PM programs which are not so well optimized, for example, where task intervals already have a conservatism built into them, or where PM tasks are less effectively performed, where the random background of failures is higher, or where the infant mortality effects of decreased intervals are present, the effects of delayed PM tasks will be smaller than estimated above.

Nevertheless, users of these models should be aware that they are very crude, generic representations, and individual cases may vary significantly from these results.

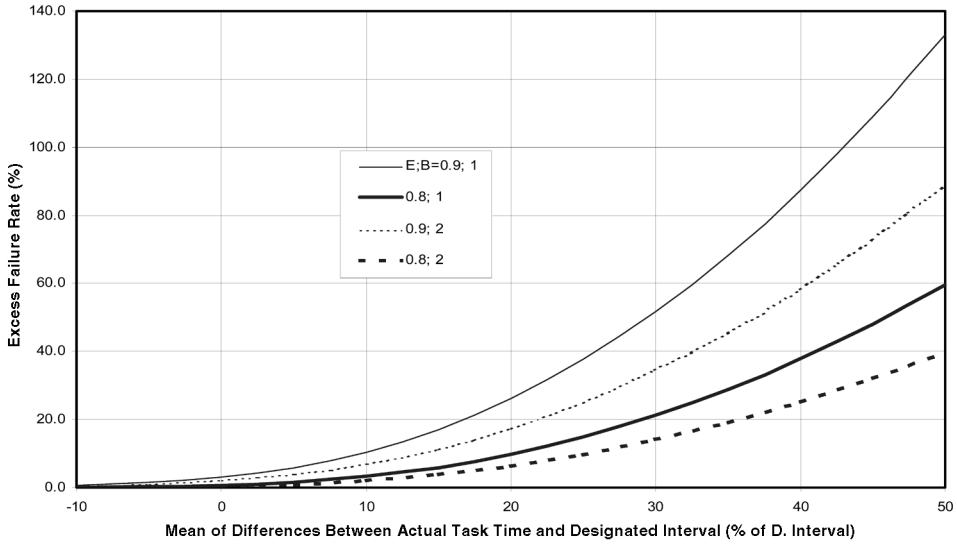

Figure A-7
Excess Failure Rate (%) Versus Deviation From Interval (%) With Standard Deviation = 12.5%

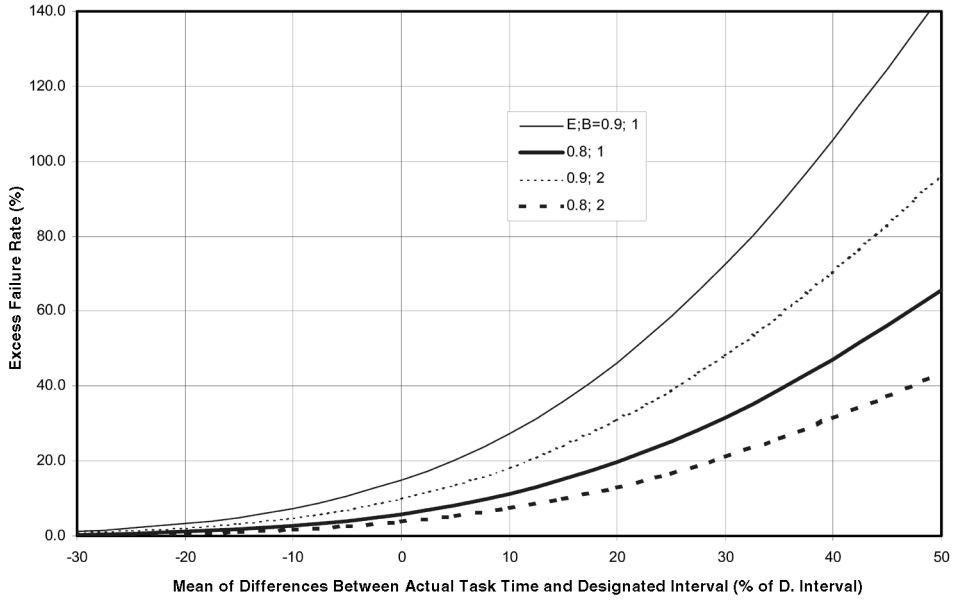

Figure A-8
Excess Failure Rate (%) Versus Deviation From Interval (%) With Standard Deviation = 25%

Target:
Nuclear Power

About EPRI

EPRI creates science and technology solutions for the global energy and energy services industry. U.S. electric utilities established the Electric Power Research Institute in 1973 as a nonprofit research consortium for the benefit of utility members, their customers, and society. Now known simply as EPRI, the company provides a wide range of innovative products and services to more than 1000 energy-related organizations in 40 countries. EPRI's multidisciplinary team of scientists and engineers draws on a worldwide network of technical and business expertise to help solve today's toughest energy and environmental problems.

EPRI. Electrify the World

1002936

EPRI • 3412 Hillview Avenue, Palo Alto, California 94304 • PO Box 10412, Palo Alto, California 94303 • USA
800.313.3774 • 650.855.2121 • askepri@epri.com • www.epri.com