

RELIABILITY-CENTERED MAINTENANCE

F.S. Nowlan, et al

United Airlines  
San Francisco, California

December 1978

REPRODUCED BY  
U.S. DEPARTMENT OF COMMERCE  
NATIONAL TECHNICAL  
INFORMATION SERVICE  
SPRINGFIELD, VA 22161



Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

| REPORT DOCUMENTATION PAGE                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                       | READ INSTRUCTIONS<br>BEFORE COMPLETING FORM                    |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------|----------------------------------------------------------------|
| 1. REPORT NUMBER<br><b>AD-A066-579</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | 2. GOVT ACCESSION NO. | 3. RECIPIENT'S CATALOG NUMBER                                  |
| 4. TITLE (and Subtitle)<br><b>Reliability-Centered Maintenance</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |                       | 5. TYPE OF REPORT & PERIOD COVERED<br><b>Final</b>             |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                       | 6. PERFORMING ORG. REPORT NUMBER                               |
| 7. AUTHOR(s)<br><b>F. Stanley Nowlan<br/>Howard F. Heap</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                       | 8. CONTRACT OR GRANT NUMBER(s)<br><b>MDA 903-75-C-0349</b>     |
| 9. PERFORMING ORGANIZATION NAME AND ADDRESS<br><b>United Airlines<br/>San Francisco International Airport<br/>San Francisco, Ca 94128</b>                                                                                                                                                                                                                                                                                                                                                                                          |                       | 10. PROGRAM ELEMENT, PROJECT, TASK<br>AREA & WORK UNIT NUMBERS |
| 11. CONTROLLING OFFICE NAME AND ADDRESS<br><b>Office of Assistant Secretary of Defense<br/>(Manpower, Reserve Affairs and Logistics)<br/>Washington, DC 20301</b>                                                                                                                                                                                                                                                                                                                                                                  |                       | 12. REPORT DATE<br><b>December 29, 1978</b>                    |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                       | 13. NUMBER OF PAGES<br><b>476</b>                              |
| 14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)                                                                                                                                                                                                                                                                                                                                                                                                                                                        |                       | 15. SECURITY CLASS. (of this report)<br><b>Unclassified</b>    |
|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |                       | 15a. DECLASSIFICATION/DOWNGRADING<br>SCHEDULE                  |
| 16. DISTRIBUTION STATEMENT (of this Report)<br><b>Approved for public release; distribution unlimited. Copies may be obtained from<br/>the National Technical Information Service or the Defense Documentation Center.</b>                                                                                                                                                                                                                                                                                                         |                       |                                                                |
| 17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)                                                                                                                                                                                                                                                                                                                                                                                                                                         |                       |                                                                |
| 18. SUPPLEMENTARY NOTES<br><b>This work was performed by United Airlines under the sponsorship of the<br/>Office of Assistant Secretary of Defense (Manpower, Reserve Affairs and Logistics)</b>                                                                                                                                                                                                                                                                                                                                   |                       |                                                                |
| 19. KEY WORDS (Continue on reverse side if necessary and identify by block number)<br><b>Actuarial analysis Aircraft maintenance Cost effectiveness Decision theory<br/>Evaluation Failure effects Flight safety Logistics Maintenance program<br/>Mechanical safety Preventive maintenance Product improvement<br/>Reliability Scheduled maintenance System effectiveness</b>                                                                                                                                                     |                       |                                                                |
| 20. ABSTRACT (Continue on reverse side if necessary and identify by block number)<br><b>This book explains basic concepts, principles, definitions, and applications of a<br/>logical discipline for development of efficient scheduled (preventive) maintenance<br/>programs for complex equipment, and the on-going management of such programs.<br/>Such programs are called reliability-centered maintenance (RCM) programs because<br/>they are centered on achieving the inherent safety and reliability capabilities of</b> |                       |                                                                |

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

**Unclassified**

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. ABSTRACT (Continued)

equipment at a minimum cost. A U.S. Department of Defense objective in sponsoring preparation of this document was that it serve as a guide for application to a wide range of different types of military equipment.

There are essentially only four types of tasks in a scheduled maintenance program. Mechanics can be asked to:

- Inspect an item to detect a potential failure
- Rework an item before a maximum permissible age is exceeded
- Discard an item before a maximum permissible age is exceeded
- Inspect an item to find failures that have already occurred but were not evident to the equipment operating crew

A central problem addressed in this book is how to determine which types of scheduled maintenance tasks, if any, should be applied to an item and how frequently assigned tasks should be accomplished. The use of a decision diagram as an aid in this analysis is illustrated. The net result is a structured, systematic blend of experience, judgment, and operational data/information to identify and analyze which type of maintenance task is both applicable and effective for each significant item as it relates to a particular type of equipment. A concluding chapter emphasizes the key importance of having a mutually supportive partnership between the personnel responsible for equipment design and the personnel responsible for equipment maintenance if maximum RCM results are to be achieved.

Appendices are included as follows:

- Procedures for auditing the development and implementation of an RCM program
- A historical review of equipment maintenance evolution
- Techniques of performing actuarial analyses
- An annotated bibliography

**Unclassified**

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

AD/A066 579

## **RELIABILITY-CENTERED MAINTENANCE**

**reliability-centered**

# **maintenance**

**F. STANLEY NOWLAN**

Director, Maintenance Analysis  
United Airlines

**HOWARD F. HEAP**

Manager, Maintenance Program Planning  
United Airlines

**Produced by Dolby Access Press**

**Designer Nancy Clark, illustrators David A. Strassman  
and Evanell Towne, compositor Jonathan D. Peck,  
printer Braun-Brumfield, Inc.**

# **preface**

**THIS VOLUME** provides the first full discussion of reliability-centered maintenance as a logical discipline for the development of scheduled-maintenance programs. The objective of such programs is to realize the inherent reliability capabilities of the equipment for which they are designed, and to do so at minimum cost. Each scheduled-maintenance task in an RCM program is generated for an identifiable and explicit reason. The consequences of each failure possibility are evaluated, and the failures are then classified according to the severity of their consequences. Then for all significant items—those whose failure involves operating safety or has major economic consequences—proposed tasks are evaluated according to specific criteria of applicability and effectiveness. The resulting scheduled-maintenance program thus includes all the tasks necessary to protect safety and operating reliability, and only the tasks that will accomplish this objective.

Up to this point the only document describing the use of decision diagrams for developing maintenance programs has been MSG-2, the predecessor of RCM analysis. MSG-2 was concerned primarily with the development of prior-to-service programs and did not cover the use of operating information to modify the maintenance program after the equipment enters service or the role of product improvement in equipment development. The chief focus was on the identification of a set of tasks that would eliminate the cost of unnecessary maintenance without compromising safety or operating capability. There was no mention of the problem of establishing task intervals, of consolidating the tasks into work packages, or of making decisions where the necessary information is unavailable. The treatment of structure programs was sketchy, and zonal and other general inspection programs were not discussed at all.

The difficulty that many people experienced in attempting to apply the concepts of MSG-2 indicated the need for changes and additions simply to clarify many of the points. It was also abundantly clear, however, that the scope of the material should be expanded to cover the topics that had not been discussed in that document. This volume includes a major expansion of the discussion on the problem of identifying functionally and structurally significant items. The RCM decision diagram itself is quite different from the one used for MSG-2. Instead of beginning with the evaluation of proposed maintenance tasks, the decision logic begins with the factor that determines the maintenance requirements of each item—the consequences of a functional failure—and then an evaluation of the failure modes that cause it. This new diagram also recognizes the four basic maintenance tasks that mechanics can perform (instead of three maintenance processes), thereby clarifying the treatment of items with hidden functions. The role of a hidden-function failure in a sequence of multiple independent failures is stressed, and it is also shown that the consequences of a possible multiple failure are explicitly recognized in the definition of the consequences of the first failure.

Another important aspect of the RCM decision logic is that it includes a default strategy for making initial maintenance decisions in the absence of full information. There is a full discussion of the problem of assigning task intervals, particularly those for first and repeat on-condition inspections. The role of age exploration and the use of information derived from operating experience, both to modify the initial maintenance program and to initiate product improvement, is discussed at length. The content of scheduled-maintenance programs developed by experienced practitioners of MSG-2 techniques may be quite similar to the programs resulting from RCM analysis, but the RCM approach is more rigorous, and there should be much more confidence in its outcome. The RCM technique can also be learned more quickly and is more readily applicable to complex equipment other than transport aircraft.

Part One of this volume presents a full explanation of the theory and principles of reliability-centered maintenance, including a discussion of the failure process, the criteria for each of the four basic tasks, the use of the decision logic to develop an initial program, and the age-exploration activities that result in a continuing evolution of this program after the equipment enters service. Part Two describes the application of these principles to the analysis of typical items in the systems, powerplant, and structure division of an airplane; the considerations in packaging the RCM tasks, along with other scheduled tasks, for actual implementation; and the information systems necessary for management of the ongoing maintenance program. The concluding chapter discusses the relationship of scheduled maintenance to operating safety, the design-maintenance partnership, and the application of

RCM analysis both to in-service fleets and to other types of complex equipment.

The text is followed by four appendices. Appendix A outlines the principles of auditing a program-development project and discusses some of the common problems that arise during analysis. This material provides an excellent check list for the analyst as well as the auditor and should be especially useful as a teaching aid for those conducting training groups in RCM methods. Appendix B is a historical review of the changes in maintenance thinking in the airline industry. Appendix C is a discussion of the engineering procedures and techniques used in actuarial analysis of reliability data. Appendix D, written by Dr. James L. Dolby, is a discussion of the literature in reliability theory, information science, decision analysis, and other areas related to RCM analysis and provides an annotated guide to this literature as well as to the specific literature on reliability-centered maintenance. Dr. Howard L. Resnikoff has written an accompanying mathematical treatment of the subject, titled *Mathematical Aspects of Reliability-Centered Maintenance*.

A book of this nature is the result of many efforts, only a few of which can be acknowledged here. First of all, we wish to express our gratitude to the late W. C. Mentzer, who directed the pioneering studies of maintenance policy at United Airlines, and to the Federal Aviation Administration for creating the environment in which this work was developed over the last twenty years. We also thank Charles S. Smith and Joseph C. Saia of the Department of Defense, who defined the content of the present text and counseled us throughout its preparation. James L. Dolby of San Jose State University, in addition to preparing the bibliography, contributed his expertise to the text. In particular, he helped to develop the concept of partitioning to identify significant items and the concept of default answers as part of the decision logic, as well as advising us on the actuarial appendix. Nancy Clark edited our efforts and organized them for clear exposition. Her logical thought processes resulted in numerous major improvements throughout and made possible the successful translation of our manuscript to textbook form.

Much help on specific areas of the text has come from friends and coworkers in the industry. We especially wish to thank Mel Stone of Douglas Aircraft for his extensive help with the structure chapter, John F. McDonald of the Flying Tiger Line for his comments on the theoretical chapters, and John F. Pirtle of General Electric for his comments on the powerplant chapter. Of the many others whose contributions influenced the text in some important respect, we give particular thanks to Thomas M. Edwards of United Airlines, Thomas D. Matteson of United Airlines, Ernest Boyer of the Federal Aviation Administration, Captain L. Ebbert of the U.S. Navy, Edward L. Thomas of the Air Transport Association, and Robert Gard of the University of Missouri.

We are also grateful to the many people at United Airlines who provided us with specific help and assistance. The manuscript itself would not have materialized without the efforts of Marie Tilson, who cheerfully typed and retyped the material through many drafts. We also thank Claudia Tracy, whose artwork made the draft manuscript more readable, and J. Douglas Burch, whose efforts throughout the project helped bring it to completion. Finally, we would like to thank the management of United Airlines for its patience and our wives for their encouragement over the many long months of authorship and publication.

*F. Stanley Nowlan*  
*Howard F. Heap*

# contents

PREFACE *vii*

A MAINTENANCE PHILOSOPHY *xvi*

## CHAPTER ONE **rcm: a maintenance discipline** 2

- 1-1 THE EVOLUTION OF RCM ANALYSIS 4
- 1-2 THE BASIS OF RCM DECISION LOGIC 6
- 1-3 RELIABILITY PROBLEMS IN COMPLEX EQUIPMENT 9
- 1-4 AN OVERVIEW OF MAINTENANCE ACTIVITY 11

## PART ONE THEORY AND PRINCIPLES 15

## CHAPTER TWO **the nature of failure** 16

- 2-1 THE DEFINITION OF FAILURE 17
- 2-2 THE DETECTION OF FAILURES 20
- 2-3 THE CONSEQUENCES OF FAILURE 25
- 2-4 MULTIPLE FAILURES 28
- 2-5 THE FAILURE PROCESS 31
- 2-6 FAILURE IN COMPLEX ITEMS 37

2-7 QUANTITATIVE DESCRIPTIONS OF FAILURE 39

2-8 AGE-RELIABILITY CHARACTERISTICS 45

### **CHAPTER THREE the four basic maintenance tasks 50**

3-1 SCHEDULED ON-CONDITION TASKS 51

3-2 SCHEDULED REWORK TASKS 56

3-3 SCHEDULED DISCARD TASKS 58

3-4 SCHEDULED FAILURE-FINDING TASKS 61

3-5 CHARACTERISTICS OF THE BASIC TASKS 65

3-6 THE DIMENSIONS OF A SCHEDULED-MAINTENANCE PROGRAM 71

3-7 PRODUCT IMPROVEMENT AS PREVENTIVE MAINTENANCE 75

### **CHAPTER FOUR developing the initial program 78**

4-1 THE NATURE OF SIGNIFICANT ITEMS 80

4-2 THE RCM DECISION PROCESS 86

4-3 USE OF THE RCM DECISION DIAGRAM 91

4-4 DETERMINING COST EFFECTIVENESS 100

4-5 AGE EXPLORATION 106

4-6 PACKAGING THE MAINTENANCE TASKS 109

### **CHAPTER FIVE evolution of the rcm program 112**

5-1 THE USES OF OPERATING DATA 113

5-2 REACTING TO SERIOUS FAILURES 116

5-3 REFINING THE MAINTENANCE PROGRAM 121

5-4 REVISIONS IN MAINTENANCE REQUIREMENTS 126

5-5 THE PRODUCT-IMPROVEMENT PROCESS 128

5-6 RCM PROGRAMS FOR IN-SERVICE EQUIPMENT 137

## **PART TWO APPLICATIONS 139**

### **CHAPTER SIX applying rcm theory to aircraft 140**

- 6-1 A SUMMARY OF RCM PRINCIPLES 141**
- 6-2 ORGANIZATION OF THE PROGRAM-DEVELOPMENT TEAM 145**
- 6-3 BEGINNING THE DECISION PROCESS 147**
- 6-4 THE INFORMATION FLOW IN DECISION MAKING 153**

### **CHAPTER SEVEN rcm analysis of systems 158**

- 7-1 CHARACTERISTICS OF SYSTEMS ITEMS 159**
- 7-2 ASSEMBLING THE REQUIRED INFORMATION 161**
- 7-3 ANALYSIS OF TYPICAL SYSTEMS ITEMS 166**
- 7-4 ESTABLISHING TASK INTERVALS 192**

### **CHAPTER EIGHT rcm analysis of powerplants 194**

- 8-1 CHARACTERISTICS OF POWERPLANT ITEMS 195**
- 8-2 ASSEMBLING THE REQUIRED INFORMATION 199**
- 8-3 FAILURES OF THE BASIC ENGINE FUNCTION 205**
- 8-4 FAILURES OF SECONDARY ENGINE FUNCTIONS 217**
- 8-5 THE ROLE OF AGE EXPLORATION 224**

### **CHAPTER NINE rcm analysis of structures 228**

- 9-1 CHARACTERISTICS OF STRUCTURAL ITEMS 229**
- 9-2 THE STRUCTURAL INSPECTION PLAN 238**
- 9-3 ASSEMBLING THE REQUIRED INFORMATION 247**
- 9-4 RCM ANALYSIS OF STRUCTURAL ITEMS 252**
- 9-5 ESTABLISHING INITIAL INSPECTION INTERVALS 256**
- 9-6 STRUCTURAL AGE EXPLORATION 273**

## **CHAPTER TEN completing the maintenance program 276**

- 10.1 OTHER SCHEDULED-MAINTENANCE TASKS 277**
- 10.2 PACKAGING THE MAINTENANCE WORKLOAD 284**

## **CHAPTER ELEVEN the use of operating information 292**

- 11.1 TYPICAL INFORMATION SYSTEMS 293**
- 11.2 TYPICAL TYPES OF ROUTINE ANALYSIS 301**
- 11.3 MODIFYING THE MAINTENANCE PROGRAM 307**
- 11.4 INTERVALS: AN INFORMATION PROGRAM 324**
- 11.5 RESOLVING DIFFERENCES OF OPINION 325**
- 11.6 PURGING THE PROGRAM 328**

## **CHAPTER TWELVE the role of scheduled maintenance 330**

- 12.1 SAFETY, RELIABILITY, AND SCHEDULED MAINTENANCE 331**
- 12.2 AIR-TRANSPORT SAFETY LEVELS 337**
- 12.3 THE DESIGN-MAINTENANCE PARTNERSHIP 341**
- 12.4 RCM PROGRAMS FOR IN-SERVICE FLEETS 343**
- 12.5 EXPANSION OF RCM APPLICATIONS 347**

## **PART THREE APPENDICES 349**

### **APPENDIX A auditing rcm program development 350**

- A.1 AUDITING THE PROGRAM-DEVELOPMENT PROJECT 351**
- A.2 AUDITING THE DECISION PROCESS 354**
- A.3 AUDITING ANALYSIS OF THE EQUIPMENT 362**
- A.4 AUDITING THE ONGOING PROGRAM 367**
- A.5 AUDITING NEW PROGRAMS FOR IN-SERVICE FLEETS 368**

## **APPENDIX B the history of rcm programs 370**

- B-1 THE HARD-TIME PARADOX 371**
- B-2 CHANGING PERCEPTIONS OF THE HARD-TIME POLICY 376**
- B-3 THE INTRODUCTION OF ON-CONDITION MAINTENANCE 383**
- B-4 THE AIR TRANSPORT ASSOCIATION MSG-1 AND MSG-2 PROGRAMS 385**
- B-5 THE RELATIONSHIP OF SCHEDULED MAINTENANCE TO OPERATING SAFETY 387**

## **APPENDIX C actuarial analysis 390**

- C-1 ANALYSIS OF LIFE-TEST DATA 391**
- C-2 ANALYSIS OF DATA FROM A DEFINED CALENDAR PERIOD 395**
- C-3 THE SMOOTHING PROBLEM 402**
- C-4 ANALYSIS OF A MIXED POPULATION 408**
- C-5 USEFUL PROBABILITY DISTRIBUTIONS 411**
- C-6 A SPECIAL USE OF THE EXPONENTIAL DISTRIBUTION 417**

## **APPENDIX D bibliography 420**

- D-1 HISTORICAL DEVELOPMENT 422**
- D-2 RELIABILITY THEORY AND ANALYSIS 425**
- D-3 INFORMATION SCIENCE AND DECISION ANALYSIS 427**
- D-4 MAINTENANCE THEORY AND PHILOSOPHY 430**
- D-5 MAINTENANCE APPLICATIONS 436**
- D-6 A GUIDE TO OTHER SOURCES 438**

**GLOSSARY 453**

**INDEX 467**

XV

# a maintenance philosophy

An operator's maintenance program has four objectives:

- ▶ To ensure realization of the inherent safety and reliability levels of the equipment
- ▶ To restore safety and reliability to their inherent levels when deterioration has occurred
- ▶ To obtain the information necessary for design improvement of those items whose inherent reliability proves inadequate
- ▶ To accomplish these goals at a minimum total cost, including maintenance costs and the costs of residual failures

Reliability-centered maintenance is based on the following precepts:

- ▶ A failure is an unsatisfactory condition. There are two types of failures: *functional failures*, usually reported by operating crews, and *potential failures*, usually discovered by maintenance crews.
- ▶ The consequences of a functional failure determine the priority of maintenance effort. These consequences fall into four categories:
  - ▶ Safety consequences, involving possible loss of the equipment and its occupants
  - ▶ Operational consequences, which involve an indirect economic loss as well as the direct cost of repair
  - ▶ Nonoperational consequences, which involve only the direct cost of repair
  - ▶ Hidden-failure consequences, which involve exposure to a possible multiple failure as a result of the undetected failure of a hidden function

- ▶ Scheduled maintenance is required for any item whose loss of function or mode of failure could have safety consequences. If preventive tasks cannot reduce the risk of such failures to an acceptable level, the item must be redesigned to alter its failure consequences.
- ▶ Scheduled maintenance is required for any item whose functional failure will not be evident to the operating crew, and therefore reported for corrective action.
- ▶ In all other cases the consequences of failure are economic, and maintenance tasks directed at preventing such failures must be justified on economic grounds.
- ▶ All failure consequences, including economic consequences, are established by the design characteristics of the equipment and can be altered only by basic changes in the design:
  - ▶ Safety consequences can in nearly all cases be reduced to economic consequences by the use of redundancy.
  - ▶ Hidden functions can usually be made evident by instrumentation or other design features.
  - ▶ The feasibility and cost effectiveness of scheduled maintenance depend on the inspectability of the item, and the cost of corrective maintenance depends on its failure modes and inherent reliability.
- ▶ The inherent reliability of the equipment is the level of reliability achieved with an effective maintenance program. This level is established by the design of each item and the manufacturing processes that produced it. Scheduled maintenance can ensure that the inherent reliability of each item is achieved, but no form of maintenance

nance can yield a level of reliability beyond that inherent in the design.

A reliability-centered maintenance program includes only those tasks which satisfy the criteria for both applicability and effectiveness. The *applicability* of a task is determined by the characteristics of the item, and its *effectiveness* is defined in terms of the consequences the task is designed to prevent.

- ▶ There are four basic types of tasks that mechanics can perform, each of which is applicable under a unique set of conditions. The first three tasks are directed at preventing functional failures of the items to which they are assigned and the fourth is directed at preventing a multiple failure involving that item:
  - ▶ On-condition inspections of an item to find and correct any potential failures
  - ▶ Rework (overhaul) of an item at or before some specified age limit
  - ▶ Discard of an item (or one of its parts) at or before some specified life limit
  - ▶ Failure-finding inspections of a hidden-function item to find and correct functional failures that have already occurred but were not evident to the operating crew
- ▶ A *simple item*, one that is subject to only one or a very few failure modes, frequently shows a decrease in reliability with increasing operating age. An age limit may be useful in reducing the overall failure rate of such items, and safe-life limits imposed on a single part play a crucial role in controlling critical failures.
- ▶ A *complex item*, one whose functional failure may result from many different failure modes, shows little or no decrease in overall reliability with increasing age unless there is a dominant failure mode. Age limits imposed on complex components and systems (including the equipment itself) therefore have little or no effect on their overall failure rates.

The RCM decision diagram provides a logical tool for determining which scheduled tasks are either necessary or desirable to protect the safety and operating capability of the equipment.

- ▶ The resulting set of RCM tasks is based on the following considerations:
  - ▶ The consequences of each type of functional failure

- ▶ The visibility of a functional failure to the operating crew (evidence that a failure has occurred)
- ▶ The visibility of reduced resistance to failure (evidence that a failure is imminent)
- ▶ The age-reliability characteristics of each item
- ▶ The economic tradeoff between the cost of scheduled maintenance and the benefits to be derived from it
- ▶ A multiple failure, resulting from a sequence of independent failures, may have consequences that would not be caused by any one of the individual failures alone. These consequences are taken into account in the definition of the failure consequences for the first failure.
- ▶ A default strategy governs decision making in the absence of full information or agreement. This strategy provides for conservative initial decisions, to be revised on the basis of information derived from operating experience.

A scheduled-maintenance program must be dynamic. Any prior-to-service program is based on limited information, and the operating organization must be prepared to collect and respond to real data throughout the operating life of the equipment.

- ▶ Management of the ongoing maintenance program requires an organized information system for surveillance and analysis of the performance of each item under actual operating conditions. This information is needed for two purposes:
  - ▶ To determine the refinements and modifications to be made in the initial maintenance program (including the adjustment of task intervals)
  - ▶ To determine the needs for product improvement
- ▶ The information derived from operating experience has the following hierarchy of importance:
  - ▶ Failures that could affect operating safety
  - ▶ Failures that have operational consequences
  - ▶ The failure modes of units removed as a result of failures
  - ▶ The general condition of unfailed parts in units that have failed
  - ▶ The general condition of serviceable units inspected as samples

- ▶ At the time an initial program is developed information is available to determine the tasks necessary to protect safety and operating capability. However, the information required to determine optimum task intervals and the applicability of age limits can be obtained only from age exploration after the equipment enters service.
- ▶ With any new equipment there is always the possibility of unanticipated failure modes. The first occurrence of any serious unanticipated failure immediately sets in motion the following product-improvement cycle:
  - ▶ An on-condition task is developed to prevent recurrences while the item is being redesigned.
  - ▶ The operating fleet is modified to incorporate the redesigned part.
  - ▶ After the modification has proved successful, the special task is eliminated from the maintenance program.
- ▶ Product improvement, based on identification of the actual reliability characteristics of each item through age exploration, is part of the normal development cycle of all complex equipment.

## **RELIABILITY-CENTERED MAINTENANCE**

# reliability-centered maintenance

**THE TERM** *reliability-centered maintenance* refers to a scheduled-maintenance program designed to realize the inherent reliability capabilities of equipment. For years maintenance was a craft learned through experience and rarely examined analytically. As new performance requirements led to increasingly complex equipment, however, maintenance costs grew accordingly. By the late 1950s the volume of these costs in the airline industry had reached a level that warranted a new look at the entire concept of preventive maintenance. By that time studies of actual operating data had also begun to contradict certain basic assumptions of traditional maintenance practice.

One of the underlying assumptions of maintenance theory has always been that there is a fundamental cause-and-effect relationship between scheduled maintenance and operating reliability. This assumption was based on the intuitive belief that because mechanical parts wear out, the reliability of any equipment is directly related to operating age. It therefore followed that the more frequently equipment was overhauled, the better protected it was against the likelihood of failure. The only problem was in determining what age limit was necessary to assure reliable operation.

In the case of aircraft it was also commonly assumed that all reliability problems were directly related to operating safety. Over the years, however, it was found that many types of failures could not be prevented no matter how intensive the maintenance activities. Moreover, in a field subject to rapidly expanding technology it was becoming increasingly difficult to eliminate uncertainty. Equipment designers were able to cope with this problem, not by preventing failures, but by

preventing such failures from affecting safety. In most aircraft all essential functions are protected by redundancy features which ensure that, in the event of a failure, the necessary function will still be available from some other source. Although fail-safe and "failure-tolerant" design practices have not entirely eliminated the relationship between safety and reliability, they have dissociated the two issues sufficiently that their implications for maintenance have become quite different.

A major question still remained, however, concerning the relationship between scheduled maintenance and reliability. Despite the time-honored belief that reliability was directly related to the intervals between scheduled overhauls, searching studies based on actuarial analysis of failure data suggested that the traditional hard-time policies were, apart from their expense, ineffective in controlling failure rates. This was not because the intervals were not short enough, and surely not because the teardown inspections were not sufficiently thorough. Rather, it was because, contrary to expectations, for many items the likelihood of failure did not in fact increase with increasing operating age. Consequently a maintenance policy based exclusively on some maximum operating age would, no matter what the age limit, have little or no effect on the failure rate.

At the same time the FAA, which is responsible for regulating airline maintenance practices, was frustrated by experiences showing that it was not possible for airlines to control the failure rate of certain types of engines by any feasible changes in scheduled-overhaul policy. As a result, in 1960 a task force was formed, consisting of representatives from both the FAA and the airlines, to investigate the capabilities of

scheduled maintenance. The work of this group led to an FAA/Industry Reliability Program, issued in November 1961. The introduction to that program stated:\*

The development of this program is towards the control of reliability through an analysis of the factors that affect reliability and provide a system of actions to improve low reliability levels when they exist. . . . In the past, a great deal of emphasis has been placed on the control of overhaul periods to provide a satisfactory level of reliability. After careful study, the Committee is convinced that reliability and overhaul time control are not necessarily directly associated topics; therefore, these subjects are dealt with separately. Because the propulsion system has been the area of greatest concern in the recent past, and due to powerplant data being more readily available for study, programs are being developed for the propulsion system first as only one system at a time can be successfully worked out.

This approach was a direct challenge to the traditional concept that the length of the interval between successive overhauls of an item was an important factor in its failure rate. The task force developed a propulsion-system reliability program, and each airline involved in the task force was then authorized to develop and implement reliability programs in the area of maintenance in which it was most interested. During this process a great deal was learned about the conditions that must obtain for scheduled maintenance to be effective.† It was also found that in many cases there was no effective form of scheduled maintenance.

## 1.1 THE EVOLUTION OF RCM ANALYSIS

development of decision-  
diagram techniques  
predecessor programs:  
MSG-1 and MSG-2

At United Airlines an effort was made to coordinate what had been learned from these various activities and define a generally applicable approach to the design of maintenance programs. A rudimentary decision-diagram technique was devised in 1965 and was refined over the next few years.‡ This technique was eventually embodied in a docu-

\*FAA/Industry Reliability Program, Federal Aviation Agency, November 7, 1961, p. 1.

†Handbook for Maintenance Control by Reliability Methods, FAA Advisory Circular 120-17, December 31, 1964.

‡H. N. Taylor and F. S. Nowlan, Turbine Engine Reliability Program, FAA Maintenance Symposium on Continued Reliability of Transport-type Aircraft Propulsion Systems, Washington, D.C., November 17-18, 1965. T. D. Matteson and F. S. Nowlan, Current Trends in Airline Maintenance Programs, AIAA Commercial Aircraft Design and Operations Meeting, Los Angeles, June 12-14, 1967. F. S. Nowlan, The Use of Decision Diagrams for Logical Analysis of Maintenance Programs, United Airlines internal document, August 2, 1967.

ment published under the title *Handbook: Maintenance Evaluation and Program Development*, generally known as MSG-1.\* MSG-1 was used by special teams of industry and FAA personnel to develop the initial program issued by the FAA Maintenance Review Board for the Boeing 747. As described by the FAA, these teamst

... sorted out the potential maintenance tasks and then evaluated them to determine which must be done for operating safety or essential hidden function protection. The remaining potential tasks were evaluated to determine whether they were economically useful. These procedures provide a systematic review of the aircraft design so that, in the absence of real experience, the best [maintenance] process can be utilized for each component and system.

The Boeing 747 maintenance program so developed was the first attempt to apply reliability-centered maintenance concepts. This program has been successful.

Subsequent improvements in the decision-diagram approach led in 1970 to a second document, *MSG-2: Airline/Manufacturer Maintenance Program Planning Document*, which was used to develop the scheduled-maintenance programs for the Lockheed 1011 and the Douglas DC-10.† These programs have been successful. MSG-2 has also been applied to tactical military aircraft such as the McDonnell F4J and the Lockheed P-3, and a similar document prepared in Europe was the basis of the initial scheduled-maintenance programs for such recent aircraft as the Airbus Industrie A-300 and the Concorde.

The objective of the techniques outlined by MSG-1 and MSG-2 was to develop a scheduled-maintenance program that assured the maximum safety and reliability of which the equipment was capable and would meet this requirement at the lowest cost. As an example of the economic benefits achieved with this type of program, under traditional maintenance policies the initial program for the Douglas DC-8 included scheduled overhaul for 339 items, whereas the initial program for the DC-10, based on MSG-2, assigned only seven items to overhaul. One of the items no longer subject to an overhaul limit in the later program was the turbine engine. Elimination of this scheduled task not only led to major reductions in labor and materials costs, but also reduced the spare-engine inventory required to cover shop activities by more than 50 percent. Since engines for larger airplanes now cost upwards of \$1 million each, this is a respectable saving.

\*747 Maintenance Steering Group, *Handbook: Maintenance Evaluation and Program Development (MSG-1)*, Air Transport Association, July 10, 1968.

†Federal Aviation Administration Certification Procedures, May 19, 1972, par. 3036.

‡*Airline/Manufacturer Maintenance Program Planning Document: MSG-2*, Air Transport Association, R & M Subcommittee, Marcy 25, 1970.

As another example, under the initial program developed for the Boeing 747 it took United Airlines only 66,000 manhours on major structural inspections to reach an inspection interval of 20,000 hours. In contrast, traditional maintenance policies led to an expenditure of over 4 million manhours before the same interval was attained for structural inspections on the smaller and less complex Douglas DC-8. Cost reductions on this scale are of obvious importance to any organization responsible for maintaining large fleets of complex equipment. More important, they are achieved with no decrease in the reliability of the equipment; in fact, a clearer understanding of the failure process has actually improved operating reliability by making it easier to pinpoint signs of an imminent failure.

The specific developments that led to RCM concepts as a fundamental approach to maintenance planning are described in detail in Appendix B. Although MSG-1 and MSG-2 were short working papers, intended for use by a small number of people with extensive backgrounds in aircraft maintenance, further clarification of the basic principles has resulted in a logical discipline that applies to maintenance programs for any complex equipment.

## 1.2 THE BASIS OF RCM DECISION LOGIC

the nature of failure  
identification of  
significant items  
evaluation of failure  
consequences  
selection of applicable and  
effective tasks  
the role of age exploration

The principles of reliability-centered maintenance stem from a rigorous examination of certain questions that are often taken for granted:

- ▶ How does a failure occur?
- ▶ What are its consequences?
- ▶ What good can preventive maintenance do?

One of the chief drawbacks of the old hard-time approach to scheduled maintenance is that the resulting teardown inspections provided no real basis for determining when serviceable parts were likely to fail—that is, there was no objective means of identifying reduced resistance to failure. More than any other single factor, recognition of the specific need to identify potential-failure conditions has been responsible for the change from scheduled overhauls to on-condition inspections for signs of imminent failure.

Unfortunately, not all items can be protected by this type of maintenance task. In some cases the failure mechanism is imperfectly understood, in others it is random, and in yet others the cost of such inspections exceeds the benefits they might provide. In fact, preventive maintenance is not possible for many items of modern complex equipment. Nor, in all cases, is it necessary. Failures which could jeopardize the safety of the equipment or its occupants must be prevented. Under

modern design practices, however, very few items fall into this category, either because an essential function is provided by more than one source or because operating safety is protected in some other way. Similarly, hidden functions must be protected by scheduled maintenance, both to ensure their availability and to prevent exposure to the risk of a multiple failure.

In all other cases the consequences of failure are economic, and the value of preventive maintenance must be measured in economic terms. In some cases these consequences are major, especially if a failure affects the operational capability of the equipment. Whenever equipment must be removed from service to correct a failure, the cost of failure includes that loss of service. Thus if the intended use of the equipment is of significant value, the delay or abandonment of that use will constitute a significant loss—a fact that must be taken into account in evaluating the benefit of preventive maintenance. Other failures will incur only the cost of correction or repair, and such failures may well be tolerable, in the sense that it is less expensive to correct them as they occur than to invest in the cost of preventing them.

In short, the driving element in all maintenance decisions is not the failure of a given item, but the consequences of that failure for the equipment as a whole. Within this context it is possible to develop an efficient scheduled-maintenance program, subject to the constraints of satisfying safety requirements and meeting operational-performance goals. However, the solution of such an optimization problem requires certain specific information which is nearly always unavailable at the time an initial program must be developed. Hence we also need a basic strategy for decision making which provides for optimum maintenance decisions, given the information available at the time. The process of developing an initial RCM program therefore consists of the following steps:

- ▶ Partitioning the equipment into object categories to identify those items that require intensive study
- ▶ Identifying significant items, those whose failure would have safety or major economic consequences for the equipment as a whole, and all hidden functions, which require scheduled maintenance regardless of their significance
- ▶ Evaluating the maintenance requirements for each significant item and hidden function in terms of the failure consequences and selecting only those tasks which will satisfy these requirements
- ▶ Identifying items for which no applicable and effective task can be found and either recommending design changes if safety is involved or assigning no scheduled-maintenance tasks to these items until further information becomes available

- ▶ Selecting conservative initial intervals for each of the included tasks and grouping the tasks in maintenance packages for application
- ▶ Establishing an age-exploration program to provide the factual information necessary to revise initial decisions

The first of these steps is intended, as a purely practical matter, to reduce the problem of analysis to manageable size and to focus it according to areas of engineering expertise. The next three steps are the crux of RCM analysis. They involve a specific sequence of decision questions, worded to indicate the information required for a yes/no answer in each case. Where this information is not available, a default answer specifies the action that will best protect the equipment until there is a basis for some other decision. This decision-diagram technique, described in full in Chapter 4, not only provides an orderly basis for making decisions with limited information, but also results in a clear audit trail for later review.

In the airline industry all scheduled-maintenance programs are, of course, subject to FAA review and approval. The initial program for each new type of equipment is promulgated by the FAA Maintenance Review Board. This document, developed in conference with the equipment manufacturers and the purchasing airlines, forms the basis of the initial program submitted by each airline for FAA approval. Organizations operating other equipment in the civilian and military spheres may define their initial maintenance programs differently, but some comparable review procedure is usually involved.

Because any initial scheduled-maintenance program must be developed and implemented in advance of actual operational data, an important element of RCM programs is age exploration, a procedure for systematic gathering of the information necessary to determine the applicability of some maintenance tasks and evaluate the effectiveness of others. As this information accumulates, the same decision diagram provides a means of revising and refining the initial program. Much of this information is already available, of course, for equipment that has been in service for some time. Although the specific data needed may have to be retrieved from several different information systems, and the remaining useful life of the equipment will be a factor in certain decisions, RCM analysis under these circumstances will result in fewer default decisions, and hence a near-optimum program at the outset. Such programs usually include a larger number of on-condition inspections than the programs arrived at under older policies, and fewer of the scheduled rework tasks which had been included simply because there was no evidence that they should not be done.

An effective scheduled-maintenance program will realize all the reliability of which the equipment is capable. However, no form of preventive maintenance can alter characteristics that are inherent in the

design. The residual failures that occur after all applicable and effective preventive tasks have been implemented reflect the inherent capability of the equipment, and if the resulting level of reliability is inadequate, the only recourse is engineering redesign. This effort may be directed at a single component to correct for a dominant failure mode or it may be directed at some characteristic that will make a particular preventive technique feasible. Product improvement of this kind takes place routinely during the early years of operation of any complex equipment. Thus, although reliability-centered maintenance is concerned in the short run with tasks based on the actual reliability characteristics of the equipment, it is also concerned with improvements that will ultimately increase delivered reliability.

### 1-3 RELIABILITY PROBLEMS IN COMPLEX EQUIPMENT

Failures are inevitable in any complex equipment, although their consequences can be controlled by careful design and effective maintenance. The reason for this failure incidence is apparent if we consider some basic differences between simple and complex equipment. Simple equipment is asked to perform very few different functions. Such equipment therefore consists of only a few systems and assemblies, and these in turn may be so simple that some are exposed to only one possible failure mode. In most cases this simplicity extends to the structural elements as well, and both the structure and the various items on the equipment are relatively accessible for inspection.

failure possibilities in  
complex equipment  
performance requirements  
and reliability  
the role of design in reducing  
failure consequences

As a result, simple equipment has certain distinct failure characteristics. Because it is exposed to relatively few failure possibilities, its overall reliability tends to be higher. For the same reason, these failures tend to be age-related; each type of failure tends to concentrate around some average age, and since only a few types of failure are involved, they govern the average age at failure. However, in the absence of redundancy and other protective features, such failures may have fairly serious consequences. Thus simple equipment is often protected by "overdesign"; components are heavier and bulkier than necessary, and familiar materials and processes are used to avoid the uncertainty associated with more complex high-performance equipment.

All in all, the traditional idea that failures are directly related to safety and that their likelihood varies directly with age is often true for simple equipment. In any case, it is fairly easy to make an exhaustive study of such equipment to determine its scheduled-maintenance requirements.

The situation is quite different with the complex equipment in use today. The general-aviation aircraft of the 1930s usually had a simple

SECTION 1-3 9

Reproduced from  
best available copy.



reciprocating engine, a fixed-pitch propeller, fixed landing gear, and no wing flaps. The modern airplane may have several turboprop or turbojet powerplants, retractable landing gear, movable high-lift devices, an airframe anti-icing system, pressure- and temperature-control systems for the cabin, extensive communications and navigation equipment, complex cockpit instrumentation, and complex ancillary systems to support all these additional items. This increased complexity has greatly expanded the safe operational capability of the aircraft. The simple airplane of the 1930s was restricted to trips of a few hundred miles under reasonably favorable weather conditions. The higher performance capability demanded of modern equipment, however, has greatly increased not only the number of items that can fail, but the types of failure that can occur.

Each new design of any high-performance equipment is essentially an attempt to make earlier designs technologically obsolete, with the usual measure of improvement being potential operating capability (including operating costs). In other words, this is the operating capability expected in the absence of any failures that might change the circumstances. The basis for evaluating new aircraft designs usually includes performance factors such as the following.

- ▶ The maximum payload (military or commercial) that can be carried over a given distance
- ▶ The maximum distance over which a given payload can be carried
- ▶ The minimum size of the vehicle that can carry a given payload over a given distance
- ▶ The highest speed that can be attained under defined payload/range conditions
- ▶ Special capabilities, such as the ability to traverse rough terrain, operate from short runways, or withstand punishment

In some cases these factors are weighed against the anticipated direct operating costs (including maintenance costs) associated with attaining such capabilities, since a major objective may be to achieve the minimum cost per unit of payload transported. In other cases performance takes precedence over cost. This is true not only of military equipment, but of certain types of civilian equipment, where there is an adequate market for specialized capability despite its cost.

Another aspect of performance demands, of course, is the trend toward increasing automation. Examples are everywhere—automatic flight-control systems in aircraft, including automatic approach and landing equipment; automatic transmissions in automobiles; automated traffic-control systems for rapid-transit trains; and automatic aperture-setting devices in cameras.

The design of complex equipment, therefore, is always a tradeoff between achieving the required performance capability and acceptable reliability. This tradeoff entails an intentional compromise between the lightness and compactness required for high performance and the weight and bulk required for durability. Thus it is neither economically nor technologically feasible to produce complex equipment that can sustain trouble-free operation for an indefinite period of time. Although the reliability of certain items that perform single functions may be improving, the number of such items has been vastly multiplied. It is therefore inevitable that failures will occur—that is, that certain parts of the equipment will lose the capability of performing their specified functions.

Our concern is not with the number of these failures, but with the consequences of a given failure for the equipment as a whole. Will the loss of a particular function endanger the equipment or its occupants? If not, is it necessary to abort the mission or take the equipment out of service until repairs can be made? Or can unrestricted operation continue and the repair be deferred to a convenient time and place? The ability to defer failure consequences depends largely on the design of the equipment. One strategy is the use of redundancy and fail-safe construction. Another strategy is failure substitution, the use of a minor failure to preempt a major one, as in the use of fuses and circuit breakers. This latter concept extends to maintenance activities in which potential failures are used to preempt functional failures. Thus the design may include various instrumentation to give some warning of an impending failure or other features which facilitate inspection for possible deterioration. All these features actually increase the number of failure possibilities in the sense that they add more items that could fail. However, they greatly reduce the consequences of any single failure.

## 1.4 AN OVERVIEW OF MAINTENANCE ACTIVITY

The activities of a maintenance organization include both the scheduled work that is performed to avoid failures and the corrective work that is performed after failures have occurred. Our present concern is with *preventive maintenance*, the program of scheduled tasks necessary to ensure safe and reliable operation of the equipment. The complete collection of these tasks, together with their assigned intervals, is termed the *scheduled-maintenance program*. This program includes only the tasks that are scheduled in advance—servicing and lubrication, inspection, and scheduled removal and replacement of items on the equipment. Exhibit 1.1 lists some typical tasks in such a program.

In order to accomplish the anticipated corrective and scheduled maintenance, an operating organization must establish an overall sup-

corrective and scheduled  
maintenance  
scheduled-maintenance  
program  
maintenance stations  
line maintenance and shop  
maintenance

**EXHIBIT 1-1** Typical scheduled-maintenance tasks for various items on aircraft. Some scheduled tasks are performed on the aircraft at line-maintenance stations and others are performed at the major maintenance base, either as part of a larger maintenance package or as part of the shop procedure whenever a failed unit is sent to the maintenance base for repair.

| nature of item                                         | scheduled-maintenance task                                                                                          | task interval                                                                                 |
|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| <b>SYSTEMS ITEMS</b>                                   |                                                                                                                     |                                                                                               |
| Fuel-pump assembly<br>(Douglas A4)                     | On-condition (on aircraft):<br>Inspect filter for contamination                                                     | 60 operating hours                                                                            |
|                                                        | On-condition (on aircraft):<br>Inspect drive shaft for spline wear                                                  | 1,000 operating hours                                                                         |
| Brake assembly, main landing gear (Douglas DC-10)      | On-condition (on aircraft):<br>Inspect brake wear indicators                                                        | During overnight stops and walkaround checks                                                  |
|                                                        | On-condition (in shop): Test automatic brake adjuster                                                               | Whenever brake assembly is in shop                                                            |
| <b>POWERPLANT ITEMS</b>                                |                                                                                                                     |                                                                                               |
| Compressor rear frame<br>(General Electric CF6-6)      | On-condition (on aircraft):<br>Inspect front flange for cracks emanating from bolt holes                            | 500 flight cycles or phase check (134 days), whichever is first                               |
| Nozzle guide vanes<br>(Pratt & Whitney JT8D-7)         | On-condition (on aircraft):<br>Perform borescope inspection for burning, cracking, or bowing of guide vanes         | 1,000 operating hours                                                                         |
| Tenth-stage compressor blades<br>(Pratt & Whitney JT4) | Scheduled rework:<br>Shot-peen blade dovetail and apply antigalling compound                                        | 6,000 operating hours                                                                         |
| Stage 3 turbine disk<br>(Pratt & Whitney JT9D)         | Scheduled discard: Replace turbine disk with new part                                                               | 15,000 flight cycles or 30,000 operating hours, whichever is first                            |
| <b>STRUCTURAL ITEMS</b>                                |                                                                                                                     |                                                                                               |
| Rear spar at bulkhead intersection (Douglas DC-10)     | On-condition (on aircraft):<br>Inspect specified intersections in zones 531, 631, 141, 142 for cracks and corrosion | Primary strength-indicator areas 5,000 operating hours, internal fuel-tank areas 20,000 hours |
| Shock strut, main landing gear (Boeing 737)            | On-condition (in shop): Strip cadmium plate and inspect for cracks and corrosion                                    | 19,500 operating hours                                                                        |

port plan which includes the designation of maintenance stations, staffing with trained mechanics, provision of specialized testing equipment and parts inventories, and so on. The overall maintenance plan of an airline is typical of that for any transportation system in which each piece of equipment operates through many stations but has no unique home station.

A large proportion of the failures that occur during operation are first observed and reported by the operating crew. Some of these must be corrected after the next landing, and a few are serious enough to require a change in flight plan. The correction of many other failures, however, can be deferred to a convenient time and location. Those line stations with a high exposure to the need for immediate corrective work are designated as *maintenance stations* and are equipped with trained mechanics, spare-parts inventory, and the facilities necessary to carry out such repairs. United Airlines serves 91 airline stations with 19 such maintenance stations.

The decision to designate a particular station as a maintenance station depends chiefly on the amount of traffic at that station and the reliability of the aircraft involved. A station at which the greatest volume of repairs is expected is the logical first choice. However, other considerations may be the frequency with which the operating schedule provides overnight layovers, the relative ease of routing other aircraft to that station, the availability of mechanics and parts to support other types of aircraft, the planned volume of scheduled-maintenance work, and so on.

Line-maintenance stations themselves vary in size and complexity. The facilities needed for immediate corrective work establish the minimum resources at any given maintenance station, but operating organizations generally consolidate the bulk of the deferrable work at a few of these stations for greater economy. To simplify the control of scheduled maintenance, individual tasks are grouped into a fairly small number of maintenance packages for execution. Like deferrable corrective work, these scheduled-maintenance packages can be assigned to any convenient maintenance station. Thus the more involved work is generally assigned to those line stations already equipped with the staff and inventories for extensive corrective work.

Not all scheduled-maintenance tasks can be carried out at line stations. Major structural inspections, scheduled rework, and inspections which entail extensive disassembly are best handled at a major maintenance base equipped with shop facilities. The major base also repairs failed units that are removed from aircraft at the line stations. Few such maintenance bases are needed, and reliability considerations generally determine their size and manpower requirements, rather than their location. Many large airlines operate efficiently with only one maintenance base. The work performed at a maintenance base is generally

termed *shop maintenance* to differentiate it from *line maintenance*, which consists primarily of replacing failed units rather than repairing them.

The entire process by which a detailed support plan is developed is beyond the scope of this volume. Suffice it to say that a detailed plan is necessary in order to implement a scheduled-maintenance program. Our concern here is with the development of such a program – or rather, with the principles underlying its development. In the following chapters we will examine the nature of failures, the basis on which their consequences are evaluated, and the specific criteria that determine the applicability and effectiveness of a given type of preventive task. With this framework established, we will consider the decision logic that results in a scheduled-maintenance program based on the actual reliability characteristics of the equipment. This reliability-centered approach ensures that the inherent safety and operating capability of the equipment will be realized at the minimum cost, given the information available at any time.

The chapters in Part Two illustrate the application of RCM decision logic to specific hardware examples and discuss some of the information processes involved in the continuing evolution of the maintenance program after the equipment enters service. All these illustrations are drawn from commercial-aircraft applications. However, it should be clear from the discussion in Part One that the basic principles of RCM programs extend not just to other operating contexts, but to maintenance programs for any complex equipment.

**PART ONE**

---

# **theory and principles**

## the nature of failure

**THE PARTS** of any mechanical equipment are subject to wear, corrosion, and fatigue which inevitably result in some deviation from the conditions that existed when the equipment was new. Ultimately the deviation will become great enough that the equipment, or some item on it, no longer meets the required performance standards—that is, it fails. The role of scheduled maintenance is to cope with the failure process. For years, however, the chief focus has been on anticipating the age at which things were likely to fail, rather than on how they fail and the consequences of such failures. As a result, there has been insufficient attention to the failure process itself, and even less attention to the question of precisely what constitutes a failure.

One reason for this lack of attention has been the common assumption that all equipment “wears out” and inevitably becomes less reliable with increasing operating age. This assumption led to the conclusion that the overall failure rate of an item will always be reduced by an age limit which precludes operation at ages where the likelihood of failure is greater. In accordance with this hard-time policy, all units were taken out of service when they reached a specified age and were sent to the major maintenance base for complete disassembly and overhaul, a procedure intended to restore each part to its original condition.

It is now known that the reliability of most complex items does not vary directly with operating age, at least not in such a way as to make hard-time overhaul a useful concept. Procedures directed at obtaining some precise evidence that a failure is imminent are frequently a far superior weapon against failure. However, to understand the specific nature of such procedures as they pertain to an RCM program, it is necessary to take a closer look at the entire concept of failure. Without a precise definition of what condition represents a failure, there is no

way either to assess its consequences or to define the physical evidence for which to inspect. The term *failure* must, in fact, be given a far more explicit meaning than "an inability to function" in order to clarify the basis of reliability-centered maintenance.

In this chapter we will examine the problem of defining failures and some of the implications this has for the analysis of failure data. We will also see how failure consequences are evaluated, both in terms of single failures and in terms of multiple failures. Finally, we will discuss the process of failure itself and see why complex items, unlike simple items, do not necessarily wear out.

## 2·1 THE DEFINITION OF FAILURE

Each of us has some intuitive notion of what constitutes a failure. We would all agree that an automobile engine, a fuel pump, or a tire has failed if it ceases to perform its intended function. But there are times when an item does continue to function, although not at its expected level. An automobile engine may run powerfully and smoothly, but its oil consumption is high; a fuel pump may pump fuel, but sluggishly; a tire may hold air and support the car, but its bald tread indicates that it will do neither much longer.

failure  
functional failure  
potential failure

Have these items failed? If not, how bad must their condition become before we would say a failure has occurred? Moreover, if any of these conditions is corrected, the time required for unanticipated repairs might force a change in other plans, such as the delay or cancellation of a trip. In this event could it still be argued that no failure had occurred?

To cover all these eventualities we can define a failure in broad terms as follows:

---

**A failure is an unsatisfactory condition.**

---

In other words, a failure is any identifiable deviation from the original condition which is unsatisfactory to a particular user. The determination that a condition is unsatisfactory, however, depends on the consequences of failure in a given operating context. For example, high oil consumption in an aircraft engine may pose no problem on short- or medium-range flights, whereas on long-range flights the same rate of consumption would exhaust the oil supply. Similarly, engine-instrument malfunctions that would not disrupt operations on multi-engine equipment would be clearly unsatisfactory on a single-engine plane, and performance that is acceptable in a land-based environment might not be good enough for carrier operation.

In short, the exact dividing line between satisfactory and unsatisfactory conditions will depend not only on the function of the item in question, but on the nature of the equipment in which it is installed and the operating context in which that equipment is used. The determination will therefore vary from one operating organization to another. Within a given organization, however, it is essential that the boundaries between satisfactory and unsatisfactory conditions be defined for each item in clear and unmistakable terms.

#### **FUNCTIONAL FAILURE**

The judgment that a condition is unsatisfactory implies that there must be some condition or performance standard on which this judgment can be based. As we have seen, however, an unsatisfactory condition can range from the complete inability of an item to perform its intended function to some physical evidence that it will soon be unable to do so. For maintenance purposes, therefore, we must classify failures further as either *functional failures* or *potential failures*:

---

**A functional failure is the inability of an item (or the equipment containing it) to meet a specified performance standard.**

---

A complete loss of function is clearly a functional failure. Note, however, that a functional failure also includes the inability of an item to function at the level of performance that has been specified as satisfactory. This definition thus provides us with an identifiable and measurable condition, a basis for identifying functional failures.

To define a functional failure for any item we must, of course, have a clear understanding of its functions. This is not a trivial consideration. For example, if we say that the function of the braking system on an airplane is to stop the plane, then only one functional failure is possible — inability to stop the plane. However, this system also has the functions of providing modulated stopping capability, providing differential braking for maneuvering on the ground, providing antiskid capability, and so on. With this expanded definition it becomes clear that the braking system is in fact subject to a number of different functional failures. It is extremely important to determine all the functions of an item that are significant in a given operating context, since it is only in these terms that its functional failures can be defined.

#### **POTENTIAL FAILURE**

Once a particular functional failure has been defined, some physical condition can often be identified which indicates that this failure is imminent. Under these circumstances it may be possible to remove the item from service before the point of functional failure. When such conditions can be identified, they are defined as *potential failures*:

---

**A potential failure is an identifiable physical condition which indicates a functional failure is imminent.**

---

The fact that potential failures can be identified is an important aspect of modern maintenance theory, because it permits maximum use of each item without the consequences associated with a functional failure. Units are removed or repaired at the potential-failure stage, so that potential failures preempt functional failures.

For some items the identifiable condition that indicates imminent failure is directly related to the performance criterion that defines the functional failure. For example, one of the functions of a tire tread is to provide a renewable surface that protects the carcass of the tire so that it can be retreaded. This function is not the most obvious one, and it might well be overlooked in a listing of tire functions; nevertheless, it is important from an economic standpoint. Repeated use of the tire wears away the tread, and if wear continues to the point at which the carcass cannot be retreaded, a functional failure has occurred. To prevent this particular functional failure, we must therefore define the potential failure as some wear level that does not endanger the carcass.

The ability to identify either a functional or a potential failure thus depends on three factors:

- Clear definitions of the functions of an item as they relate to the equipment or operating context in which the item is to be used

- ▶ A clear definition of the conditions that constitute a functional failure in each case
- ▶ A clear definition of the conditions that indicate the imminence of this failure

In other words, we must not only define the failure; we must also specify the precise evidence by which it can be recognized.

## 2·2 THE DETECTION OF FAILURES

the role of the operating crew  
evident and hidden functions  
verification of failures  
interpreting failure data

Both functional failures and potential failures can be defined in terms of identifiable conditions for a given operating context. In evaluating failure data, however, it is important to take into account the different frames of reference of several sets of failure observers—the operating crew, the line mechanic, the shop mechanic, and even passengers. Understanding how and when the observer sees a failure and how he interprets it is crucial both to operating reliability and to effective preventive maintenance.

The detection and reporting of failures depends on two principal elements:

- ▶ The observer must be in a position to detect the failure. This “right” position may be a physical location, a particular moment in time, or access to the inspection equipment that can reveal the condition.
- ▶ The observer must have standards that enable him to recognize the condition he sees as a failure, either functional or potential.

### THE ROLE OF THE OPERATING CREW

Members of the operating crew are the only people in a position to observe the dynamic operation of the equipment in its normal environment. Whereas an airplane in a maintenance facility is in a static environment, during flight its systems are activated and the whole machine is subjected to airloads and to both low atmospheric pressure and low outside temperatures. As a result, the operating crew will be the first to observe many functional failures. Such failures are often detected at the time a crew member calls on a function and finds that it is impaired.

In most complex equipment the crew’s ability to observe failures is further enhanced by extensive instrumentation, warning lights, or other monitoring devices. In some cases these indicators make failures evident at the moment they occur, when otherwise they might go undetected until the function was needed. Such early warning provides more time for changes in operating strategy to offset the consequences

of the failure. For example, certain engine malfunctions may require the shutdown of one engine and perhaps the selection of an alternate landing field, or an auxiliary hydraulic pump may have to be turned on after one of the main ones fails. Even when the flight can be continued without incident, the crew is required to record the failure as accurately as possible in the flight log so the condition can be corrected at the earliest opportunity.

This instrumentation also permits the crew to determine whether items that are still operative are functioning as well as they should. In some cases reduced performance is an indication of an imminent failure, and these conditions would also be examined later to see whether a potential failure exists.

Not surprisingly, the operating crew plays a major role in detecting failure conditions. This is illustrated by a study of the support costs on a fleet of Boeing 747's over the first ten months of 1975 (a total of 51,400 operating hours). In this case 66.1 percent of all failure reports while the plane was away from the maintenance base originated with the operating crew, and these failures accounted for 61.5 percent of the total man-hours for corrective line maintenance. The other 33.9 percent of the reported failures included potential failures detected by line mechanics, along with other failures not normally evident to the operating crew.

#### **HIDDEN-FUNCTION ITEMS**

Although most functional failures are first detected by the operating crew, many items are subject to failures that the crew is not in a position to observe. The crew duties often include special checks of certain hidden-function items, but most such failures must be found by inspections or tests performed by maintenance personnel. To ensure that we will know when a failure has occurred, we must know that the observer is in a position to detect it. Hence for maintenance purposes a basic distinction is made between *evident* and *hidden functions* from the vantage point of the operating crew:

---

**An evident function is one whose failure will be evident to the operating crew during the performance of normal duties.**

---

**A hidden function is one whose failure will not be evident to the operating crew during the performance of normal duties.**

---

An item may have several functions, any one of which can fail. If the loss of one of these functions would not be evident, the item must be classified from the maintenance standpoint as a hidden-function item.

Hidden functions may be of two kinds:

- ▶ A function that is normally active but gives no indication to the operating crew if it ceases
- ▶ A function that is normally inactive, so that the crew cannot know whether it will be available when it is needed (usually the demand follows some other failure)

The fire-detection system in an aircraft powerplant falls into the first category. This system is active whenever the engine is in use, but its sensing function is hidden unless it detects a fire; thus if it fails in some way, its failure is similarly hidden. The fire-extinguishing system that backs up this unit has the second kind of hidden function. It is not activated unless a fire is detected, and only when it is called upon to operate does the crew find out whether it works.

In addition to inspecting for potential failures, maintenance personnel also inspect most hidden-function items for functional failures. Thus the operating crew and the maintenance crew complement one another as failure observers.

#### **VERIFICATION OF FAILURES**

Operating crews occasionally report conditions which appear unsatisfactory to them, but which are actually satisfactory according to the defined standards for condition and performance. This is a basic principle of prevention. The operating crew cannot always know when a particular deviation represents a potential failure, and in the interests of safety the crew is required to report anything questionable. In most airlines the operating crew can communicate directly with a central group of maintenance specialists, or controllers, about any unusual conditions observed during flight. The controllers can determine the consequences of the condition described to them and advise the crew whether to land as soon as possible or continue the flight, with or without operating restrictions. The controllers are also in a position to determine whether the condition should be corrected before the plane is dispatched again. This advice is particularly important when a plane is operating into a station which is not a maintenance station.

Once the plane is available for maintenance inspection, the maintenance crew is in a better position to diagnose the problem and determine whether a failure condition actually does exist. Thus the suspect item may be replaced or repaired or marked "OK for continued operation." The fact that failure observers have different frames of reference for interpreting the conditions they see often makes it difficult to evaluate failure reports. For example, a broken seat recliner is recognizable to any observer as a failure. Frequently a passenger will notice the con-

dition first and complain about it to the flight attendant. The line mechanic at the next maintenance station will take corrective action, usually by replacing the mechanism and sending the failed unit to the maintenance base, where the shop mechanic will record the failure and make the repair. In this case all four types of observer would have no difficulty recognizing the failure.

The situation is somewhat different with an in-flight engine shutdown as a result of erratic instrument readings. Although the passengers would not be aware that a failure had occurred, the operating crew would report an engine failure. However, the line mechanic might discover that the failure was in the cockpit instruments, not the engine. He would then replace the faulty instrument and report an instrument failure. Thus the crew members are the only ones in a position to observe the failure, but they are not in a position to interpret it. Under other circumstances the situation may be reversed. For example, on certain engines actual separation of the turbine blades—a functional failure—is preceded by a perceptible looseness of one or more blades in their mounts. If the blades separate, both the operating crew and the passengers may become abruptly aware of the functional failure, but since the engine functions normally with loose blades, neither crew nor passengers have any reason to suspect a potential failure. In this case the crew members might be able to interpret the condition as a potential failure, but they are not in a position to observe it.

The line mechanic who inspects the engine as part of scheduled maintenance will check for loose blades by slowly rotating the turbine assembly and feeling the blades with a probe (typically a length of stiff rubber or plastic tubing). If he finds any loose blades, he will report a failure and remove the engine. The mechanics in the engine-repair shop are in an even better position for detailed observation, since they must go inside the engine case to get at the faulty blades. (On occasion they may be the first to observe loose blades in an engine removed for other reasons.) If they confirm the line mechanic's diagnosis, they will report the failure as verified.

Of course, the situation is not always this clear cut. Often there are no precise troubleshooting methods to determine exactly which component or part is responsible for a reported malfunction. Under these circumstances the line mechanic will remove several items, any one of which might have caused the problem. This practice is sometimes referred to as "shotgun" troubleshooting. Many of these suspect items will show normal performance characteristics when they are tested at the maintenance base. Thus, although they are reported as failures at the time they are removed from the equipment, from the shop mechanic's frame of reference they are unverified failures. By the same token, differences between the testing environment and the field environment

will sometimes result in unverified failures for items that are actually suffering functional failures in the field.

Units removed from equipment either as potential failures or because of malfunctions are termed *premature removals*. This term came into use when most equipment items had a fixed operating-age limit. A unit removed when it reached this limit was "time-expired," whereas one removed because it had failed (or was suspected of having failed) before this age limit was a "premature" removal.

#### **INTERPRETING FAILURE DATA**

The problem of interpreting failure data is further complicated by differences in reporting policy from one organization to another. For example, one airline might classify an engine removed because of loose turbine blades as a failure (this classification would be consistent with our definition of a potential failure). This removal and all others like it would then be counted as failures in all failure data. Another airline might classify such removals as "precautionary," or even as "scheduled" (having discovered a potential failure, they would then schedule the unit for removal at the earliest opportunity). In both these cases the removals would not be reported as failures.

Similar differences arise as a result of varying performance requirements. The inability of an item to meet some specified performance requirement is considered a functional failure. Thus functional failures (and also potential failures) are created or eliminated by differences in the specified limits; even in the same piece of equipment, what is a failure to one organization will not necessarily be a failure to another. These differences exist not only from one organization to another, but within a single organization over a long calendar period. Procedures change, or failure definitions are revised, and any of these changes will result in a change in the reported failure rate.

Another factor that must be taken into account is the difference in orientation between manufacturers and users. On one hand, the operating organization tends to view a failure for any reason as undesirable and expects the manufacturer to improve the product to eliminate all such occurrences. On the other hand, the manufacturer considers it his responsibility to deliver a product capable of performing at the warranted reliability level (if there is one) under the specific stress conditions for which it was designed. If it later develops that the equipment must frequently be operated beyond these conditions, he will not want to assume responsibility for any failures that may have been caused or accelerated by such operation. Thus manufacturers tend to "censor" the failure histories of operating organizations in light of their individual operating practices. The result is that equipment users, with some confusion among them, talk about what they actually saw, while the manufacturer talks about what they should have seen.

## 2·3 THE CONSEQUENCES OF FAILURE

While failure analysis may have some small intrinsic interest of its own, the reason for our concern with failure is its consequences. These may range from the modest cost of replacing a failed component to the possible destruction of a piece of equipment and the loss of lives. Thus all reliability-centered maintenance, including the need for redesign, is dictated, not by the frequency of a particular failure, but by the nature of its consequences. Any preventive-maintenance program is therefore based on the following precept:

safety consequences  
operational consequences  
nonoperational consequences  
hidden-failure consequences

---

**The consequences of a failure determine the priority of the maintenance activities or design improvement required to prevent its occurrence.**

---

The more complex any piece of equipment is, the more ways there are in which it can fail. All failure consequences, however, can be grouped in the following four categories:

- ▶ Safety consequences, involving possible loss of the equipment and its occupants
- ▶ Operational consequences, which involve an indirect economic loss as well as the direct cost of repair
- ▶ Nonoperational consequences, which involve only the direct cost of repair
- ▶ Hidden-failure consequences, which have no direct impact, but increase the likelihood of a multiple failure

### **SAFETY CONSEQUENCES**

The first consideration in evaluating any failure possibility is safety:

---

**Does the failure cause a loss of function or secondary damage that could have a direct adverse effect on operating safety?**

---

Suppose the failure in question is the separation of a number of turbine blades on an aircraft engine, causing the engine to vibrate heavily and lose much of its thrust. This functional failure could certainly affect the safety of a single-engine aircraft and its occupants, since the loss of thrust will force an immediate landing regardless of the terrain below. Furthermore, if the engine is one whose case cannot contain ejected blades, the blades may be thrown through the engine case and cause unpredictable, and perhaps serious, damage to the

plane itself. There is also danger from hot gases escaping from the torn engine case. In a multiengine plane the loss of thrust would have no direct effect on safety, since the aircraft can maintain altitude and complete its flight with one engine inoperative. Hence the loss of function is not in itself cause for alarm. However, both plane and passengers will still be endangered by the possible secondary damage caused by the ejected blades. In this case, therefore, the secondary effects are sufficient reason to classify the failure as critical.

A *critical failure* is any failure that could have a direct effect on safety. Note, however, that the term *direct* implies certain limitations. The impact of the failure must be immediate if it is to be considered direct; that is, the adverse effect must be one that will be felt before planned completion of the flight. In addition, these consequences must result from a single failure, not from some combination of this failure with one that has not yet occurred. An important fact follows from this:

- All critical failures will be evident to the operating crew. If a failure has no evident results, it cannot, by definition, have a direct effect on safety.

It may be necessary to remove a plane from service to correct certain failures before continuing operation, and in some cases it may even be advisable to discontinue the flight. However, as long as the failure itself has no immediate safety consequences, the need for these precautionary measures does not justify classifying this failure as critical.

Not every critical failure results in an accident; some such failures, in fact, have occurred fairly often with no serious consequences. However, the issue is not whether such consequences are inevitable, but whether they are possible. For example, the secondary effects associated with ejected turbine blades are unpredictable. Usually they do not injure passengers or damage a vital part of the plane—but they can. Therefore this failure is classified as critical. Similarly, any failure that causes an engine fire is critical. Despite the existence of fire-extinguishing systems, there is no guarantee that a fire can be controlled and extinguished. Safety consequences are always assessed at the most conservative level, and in the absence of proof that a failure *cannot* affect safety, it is classified by default as critical.

In the event of any critical failure, every attempt is made to prevent a recurrence. Often redesign of one or more vulnerable items is necessary. However, the design and manufacture of new parts and their subsequent incorporation in in-service equipment takes months, and sometimes years. Hence some other action is needed in the meantime. In the case of turbine-blade failure an identifiable physical condition—loose blades—has been found to occur well in advance of actual separation of the blades. Thus regular inspection for this condition as part of scheduled maintenance makes it possible to remove engines at the

potential-failure stage, thereby forestalling all critical functional failures. Note that this preventive-maintenance task does not prevent failures; rather, by substituting a potential failure for a functional failure, it precludes the consequences of a functional failure.

#### **OPERATIONAL CONSEQUENCES**

Once safety consequences have been ruled out, a second set of consequences must be considered:

---

**Does the failure have a direct adverse effect on operational capability?**

---

Whenever the need to correct a failure disrupts planned operations, the failure has operational consequences. Thus operational consequences include the need to abort an operation after a failure occurs, the delay or cancellation of other operations to make unanticipated repairs, or the need for operating restrictions until repairs can be made. (A critical failure can, of course, be viewed as a special case of a failure with operational consequences.) In this case the consequences are economic: they represent the imputed cost of lost operational capability.

A failure that requires immediate correction does not necessarily have operational consequences. For example, if a failed item on an aircraft can be replaced or repaired during the normal transit time at a line station, then it causes no delay or cancellation of subsequent flights, and the only economic consequence is the cost of corrective maintenance. In contrast, the plane may be operational, but its reduced capability will result in such costs as high fuel consumption. The definition of operational consequences will therefore vary from one operating context to another. In all cases, however, the total cost of an operational failure includes the economic loss resulting from the failure as well as the cost of repairing it. If a failure has no operational consequences, the cost of corrective maintenance is still incurred, but this is the only cost.

If a potential failure such as loose turbine blades were discovered while the plane was in service, the time required to remove this engine and install a new one would involve operational consequences. However, inspections for this potential failure can be performed while the plane is out of service for scheduled maintenance. In this case there is ample time to remove and replace any failed engines (potential failures) without disrupting planned operations.

#### **NONOPERATIONAL CONSEQUENCES**

There are many kinds of functional failures that have no direct adverse effect on operational capability. One common example is the failure of a navigation unit in a plane equipped with a highly redundant navigation system. Since other units ensure availability of the required func-

tion, the only consequence in this case is that the failed unit must be replaced at some convenient time. Thus the costs generated by such a failure are limited to the cost of corrective maintenance.

As we have seen, potential failures also fall in this category. The purpose of defining a potential failure that can be used to preempt a functional failure is to reduce the failure consequences in as many cases as possible to the level of direct cost of replacement and repair.

#### HIDDEN-FAILURE CONSEQUENCES

Another important class of failures that have no immediate consequences consists of failures of hidden-function items. By definition, hidden failures can have no direct adverse effects (if they did, the failure would not be hidden). However, the ultimate consequences can be major if a hidden failure is not detected and corrected. Certain elevator-control systems, for example, are designed with concentric inner and outer shafts so that the failure of one shaft will not result in any loss of elevator control. If the second shaft were to fail after an undetected failure of the first one, the result would be a critical failure. In other words, the consequence of any hidden-function failure is increased exposure to the consequences of a *multiple failure*.


## 2.4 MULTIPLE FAILURES

probability of a multiple failure  
evaluation of multiple-failure  
consequences

Failure consequences are often assessed in terms of a sequence of independent events leading to a multiple failure, since several successive failures may result in consequences that no one of the failures would produce individually. The probability of a multiple failure is simple to calculate. Suppose items *A* and *B* in Exhibit 2.1 both have a probability of 0.99 of surviving a given two-hour flight (this would correspond to one failure per 100 flights, which is in fact a very high failure rate). If items *A* and *B* are both functioning at takeoff time, there are only four possible outcomes:

- ▶ Item *A* survives and item *B* survives:  $P = 0.99 \times 0.99 = 0.9801$
- ▶ Item *A* survives and item *B* fails:  $P = 0.99 \times 0.01 = 0.0099$
- ▶ Item *A* fails and item *B* survives:  $P = 0.01 \times 0.99 = 0.0099$
- ▶ Item *A* fails and item *B* fails:  $P = 0.01 \times 0.01 = 0.0001$

In other words, the probability that *A* and *B* will both fail during the same flight is only 0.0001, or an average of once in 10,000 flights. If we were considering a multiple failure of three items, the average occurrence, even with the high failure rate we have assumed here, would be once every million flights.



**EXHIBIT 2·1** Tree diagram showing the probability of a multiple failure of two items during the same flight when both items are serviceable at takeoff.

Note the difference, however, if item A is in a failed state when the flight begins. The probability that B will fail is .01; thus the probability of a multiple failure of A and B depends only on the probability of the second failure—.01, or an average of one occurrence every 100 flights. This becomes a matter of concern if the combination has critical consequences. Because of the increased probability of a multiple failure, hidden-function items are placed in a special category, and all such items that are not subject to other maintenance tasks are scheduled for failure-finding tasks. Although this type of task is intended to discover, rather than to prevent, hidden failures, it can be viewed as preventive maintenance because one of its objectives is to reduce exposure to a possible multiple failure.

To illustrate how the consequences of a multiple failure might be evaluated, consider a sequence of failures all of which are evident. If the first failure has safety consequences, there is no need to assess the consequences of a second failure. This first critical failure is the sole concern, and every effort is made to prevent its occurrence. When the first loss of function is not critical, then the consequences of a second

| nature of failure consequences |                |               |                |  |                                                                                                                                                                                         |
|--------------------------------|----------------|---------------|----------------|--|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| first failure                  | second failure | third failure | fourth failure |  | effect on previous failures in sequence                                                                                                                                                 |
| Critical                       |                |               |                |  | The critical nature of the first failure supersedes the consequences of a possible second failure.                                                                                      |
| Operational                    | Critical       |               |                |  | A second failure would be critical; the first failure must be corrected before further dispatch and therefore has operational consequences.                                             |
| Nonoperational                 | Operational    | Critical      |                |  | A third failure would be critical; the second failure must be corrected before further dispatch, but correction of the first failure can be deferred to a convenient time and location. |
| Nonoperational                 | Nonoperational | Operational   | Critical       |  | A fourth failure would be critical; the third failure must be corrected before further dispatch, but correction of both the first and second failures can be deferred.                  |

**EXHIBIT 2-2** The consequences of a single failure as determined by the consequences of a possible multiple failure. A failure that does not in itself affect operating capability acquires operational consequences if a subsequent multiple failure would be critical.

loss of function must be investigated. If the combined effect of both failures would jeopardize safety, then this multiple failure must be prevented by correcting the first failure as soon as possible. This may entail an unscheduled landing and will at least require taking the equipment out of service until the condition has been repaired. In this case, therefore, the first failure has operational consequences.

Note in Exhibit 2.2 that multiple-failure consequences need be assessed only in terms of two successive failure events. If a third loss of function would be critical, the second failure has operational consequences. However, the first failure in such a sequence can be deferred to a convenient time and place; thus it has no operational consequences. Hidden-function failures are assessed on the same basis. If the first failure under consideration is a hidden one, scheduled maintenance is necessary to protect against a multiple failure. The intensity of this maintenance, however, is dictated by the consequences of the possible multiple failure. If the combination of this failure with a second failure

would be critical, every effort is made to ensure that the hidden function will be available.

What we are doing, in effect, is treating any single failure as the first in a succession of events that could lead to a critical multiple failure. It is this method of assessing failure consequences that permits us to base a maintenance program on the consequences of single failures.

## 2.5 THE FAILURE PROCESS


### FAILURE IN SIMPLE ITEMS


One reason for identifying unsatisfactory conditions at the potential-failure stage is to prevent the more serious consequences of a functional failure. Another reason, however, is that the removal of individual units on the basis of their condition makes it possible to realize most of the useful life of each unit. To see how this procedure works, consider a simple item such as the airplane tire in Exhibit 2.3. Although a tire has other functions, here we are concerned with its retread capability. Hence we have defined a functional failure as the point at which the

failure in simple items

a model of the failure process  
the age at failure

**EXHIBIT 2.3** Tire tread wear as an illustration of the failure process in a simple item. The potential-failure condition is defined in this case as the tread depth at point A. At point B, when the tire is smooth, it can still be removed as a potential failure, but if wear continues to point C the carcass will no longer be suitable for retreading, and the loss of this function will constitute a functional failure.





**EXHIBIT 2-4** The use of potential failures to prevent functional failures. When tread depth reaches the potential-failure stage, the tire is removed and retreaded (recapped). This process restores the original tread, and hence the original failure resistance, so that the tire never reaches the functional-failure stage.

carcass plies are exposed so that the carcass is no longer suitable for retreading. The remaining tread is thus the tire's *resistance to failure* at any given moment. The *stresses* to which the tire is subjected during each landing reduce this resistance by some predictable amount, and the number of landings is a measure of the total *exposure* to stress. With increasing exposure in service, the failure resistance is gradually reduced until eventually there is a functional failure—visible plies.

Because the reduction in failure resistance is visible and easily measured, it is usual maintenance practice to define a potential failure as some wear level just short of this failure point. The tires are inspected periodically, usually when the aircraft is out of service, and any tire worn beyond the specified level is replaced. To allow for periodic inspections, the condition we choose as the potential-failure stage must not be too close to the functional-failure condition; that is, there must be a reasonable interval in which to detect the potential failure and take action. Conversely, setting the potential-failure limit too high would mean replacing tires that still had substantial useful life.

Once the optimum potential-failure level has been defined, inspections can be scheduled at intervals based on the expected amount of


tread wear over a given number of landings. Exhibit 2.4 shows a smooth tread noticed at inspection 5. At this point the tire is replaced, and if its carcass is sound, it will be retreaded. Retreading restores the original tread, and hence the original resistance to failure, and a new service cycle begins.

Failure resistance, as we are using the concept here, is somewhat analogous to the structural engineering practice of determining the stresses imposed by an applied load and then adding a safety factor to determine the design strength of a structural member. The difference between the applied load and the design strength is then the resistance to failure. The same principle extends to servicing and lubrication requirements, for example, where a specified oil quantity or lubrication film represents a resistance to functional failure. Similarly, loose turbine blades are taken as a marked reduction in failure resistance. There is a subtle difference, however, between this latter situation and the tire example. In the case of the tire the decline in failure resistance is visible and the approximate unit of stress (average tread wear per landing) is known. In the case of turbine blades the unit of stress is unknown and the decline in failure resistance is not apparent until the resistance has become quite low.

#### **A MODEL OF THE FAILURE PROCESS**

So far we have discussed a reduction in failure resistance that is evidenced by some visible condition. The more general failure process involves a direct interaction between stress and resistance, as shown in Exhibit 2.5. The measure of exposure may be calendar time, total operating hours, or number of flight or landing cycles, depending on the item. Because the measurable events occur over time, it is common to refer to total exposure as the *age* of an item. Possible measures for the stress scale are even more varied. Stresses may include temperature and atmospheric conditions, vibration, abrasion, peak loads, or some combination of these factors. It is often impossible to separate all the stress factors that may affect an item; hence exposure to stress is usually generalized to include all the stresses to which the item is subjected in a given operating context.

The primary age measure for most aircraft equipment is operating hours, usually "off-to-on" (takeoff to landing) flying hours. Some failure modes, however, are related to the number of ground-air-ground stress cycles, and in these cases age is measured as number of landings or flight cycles. Flight cycles are important, for example, in determining the number of stress cycles experienced by the aircraft structure and landing gear during landing. They are also of concern for powerplants. Engines undergo much more stress during takeoff and climb than during cruise, and an engine that experiences more takeoffs in the same number of operating hours will deteriorate more rapidly.



**EXHIBIT 2-5** Generalized model of the failure process. Resistance to failure is assumed to decline steadily with exposure to stress, measured over time as operating age, flight cycles, and so on. A functional failure occurs when the amount of stress exceeds the remaining failure resistance. In reality both stress and resistance can fluctuate, so that there is no way to predict the exact age at which the failure point will be reached.

For this reason all aircraft equipment is monitored in terms of both operating hours and flight cycles, usually on the basis of total flying time and total flight cycles for the entire aircraft. Thus if an engine is installed in a plane that has accumulated 1,000 operating hours and is removed at 1,543 hours, the engine has aged 543 hours since installation. If that engine was 300 hours old when it was installed, its age at removal is 843 hours.

Some military aircraft are equipped with acceleration recorders which also monitor the number of times the structure is stressed beyond a certain number of G's during operation. The loads can be counted and converted to an equivalent number of flight hours at the plane's designed operating profile. Like operating hours or flight cycles, these "spectrum hours" provide a basis for estimating the reduction in resistance to a particular failure mode.


A functional failure occurs when the stress and resistance curves intersect—that is, when the stress exceeds the remaining resistance to failure. Either of these curves may take a variety of different shapes, and the point at which they intersect will vary accordingly (see Exhibit 2.6). Until they do intersect, however, no functional failure occurs. In practice this failure model can be applied only to *simple items*—those subject to only one or a very few failure modes—and to individual failure modes in complex items. The reason for such a limitation be-


comes apparent if we consider some of the variables in just a single failure mode.

#### THE AGE AT FAILURE

Our examples thus far imply that any given component, such as a tire, has a well-defined initial resistance to failure and that the rate of decline in this resistance is more or less known and predictable. It follows that the time of failure should be predictable. In reality, however, even nominally identical parts will vary both in their initial failure resistance and in the rate at which this resistance declines with age. Suppose we have two nominally identical units of a simple item, or perhaps two identical parts in a complex item. To simplify matters further, let us say they are exposed to only one type of stress and are subject to only one type of failure. On this basis we might expect their failure resistance to decline at the same rate and therefore expect both units to fail at approx-

**EXHIBIT 2-6** Variability of stress, failure resistance, and the age at failure. In example *A* the resistance remains constant over time, but a sudden peak in stress causes failure to occur. In *B* the stress and resistance curves do not intersect, but the peak in stress has permanently lowered the remaining failure resistance. In *C* the reduction in failure resistance caused by the peak stress is temporary. In *D* the peak stress has accelerated the rate at which the remaining resistance will decline with age.





**EXHIBIT 2-7** The difference in failure age of two nominally identical parts subjected to similar stress patterns. The two units begin their service lives with comparable initial resistance to failure, but unit *B* is exposed to greater stress peaks and reacts to them consistently. Unit *A* behaves less accountably; its resistance is unaffected by stress peaks at 600 and 1,120 hours but declines rapidly between 1,200 and 1,300 hours. As a result, one unit fails at 850 hours and the other at 1,300 hours.

imately the same age. However, all manufactured components are produced to specified tolerance limits, which results in a variation in initial resistance. These variations are insignificant from a performance standpoint, but the result is that the two units will begin their service lives with slightly different capacities to resist stress, and these capacities may decline at somewhat different rates.

Stress also varies from moment to moment during operation, sometimes quite abruptly. For example, the different loads exerted on an aircraft structure by atmospheric turbulence can vary markedly even in the course of a short flight. Moreover, the effect of these stresses will be further influenced by the condition of the item at the particular moment it is stressed. As a result, each component will encounter a different stress pattern even if both are operating as part of the same system. Although the variations in either stress or resistance may be slight, their interaction can make a substantial difference in the length of time a given component will operate before failing. Units *A* and *B* in Exhibit 2.7 are relatively alike in their initial resistance, and the stress placed on each does not vary much from the constant stress assumed in the generalized model. However, the time of failure is the point at which the stress and resistance curves intersect; thus unit *B* failed at an age of 850 hours, whereas unit *A* survived until 1,300 hours.

Despite the variation in the failure ages of individual units, if a large number of nominally identical units are considered, their failures will tend to concentrate about some average age. For purposes of reliability analysis, however, it is necessary to employ statistical techniques that describe the variation about this average age.

It is also important to recognize that the actual age at failure depends on the stress the unit experiences. The wing-to-fuselage joints of an aircraft will stand up to normal air turbulence for a very long time, but perhaps not to the loads encountered during a tornado. The fan blades of a turbine engine can withstand thousands of hours of normal stress, but they may not be able to tolerate the ingestion of a single goose. In nearly all cases random stress peaks markedly above the average level will lower the failure resistance. This reduction may be permanent, as when damage to several structural members lowers the failure resistance of a wing, or resistance may be affected only at the time the stress exceeds a certain level. In some cases resistance may change with each variation in stress, as with metal fatigue. From the standpoint of preventive maintenance, however, the important factor is not a prediction of *when* an item is likely to fail, but whether or not the reduction in failure resistance can be identified by some physical evidence that permits us to recognize an imminent failure.

Many functional failures are evident at the time they occur, and in these cases the exact age at failure is known. Unless a failure is evident to the operating crew, however, it is impossible to determine precisely when it occurred. A potential failure detected by mechanics is known only to have occurred some time between the last inspection and the inspection at which it is observed. Similarly, although there is some exact age at which a hidden function fails, the only age we can pinpoint is the time at which the failure is discovered. For this reason the *age at failure* is defined, by convention, as the age at which a failure is observed and reported.

## 2.6 FAILURE IN COMPLEX ITEMS

A *complex item* is one that is subject to many different failure modes. As a result, the failure processes may involve a dozen different stress and resistance considerations, and a correspondingly tangled graphic representation. However, each of these considerations pertains to a single *failure mode*—some particular type or manner of failure. For instance, a bearing in a generator may wear; this causes the unit to vibrate, and ultimately the bearing will seize. At this point the generator will suffer a functional failure, since it can no longer rotate and produce electric power. Generators can also fail for other reasons, but the failure mode in this case is bearing seizure.


Of course, the bearing itself is also subject to more than one failure

failure modes  
dominant failure modes  
failure age of a complex item

mode. It may wear as a result of abrasion or crack as a result of excessive heat. From the standpoint of the generator both conditions lead to the same failure, bearing seizure. However, the maintenance analyst must know the physical circumstances leading to a particular failure in order to define an identifiable potential-failure condition. The manufacturer also needs to know that the bearing is prone to failure and that a modification is needed to improve the reliability of the generator. Such a design modification is obviously desirable if one particular failure mode is responsible for a significant proportion of all the failures of the item. Such failure modes are called *dominant failure modes*.

As with failures in simple items, the failure ages for a single failure mode tend to concentrate about an average age for that mode. However,

**EXHIBIT 2-8** Experience with 50 newly installed Pratt & Whitney JT8D-7 engines over the first 2,000 operating hours. The 21 units that failed before 2,000 hours flew a total of 18,076 hours, so the total operating time for all 50 engines was 18,076 hours plus 58,000 hours for the surviving engines, or 76,076 hours. The mean time between failures was therefore  $76,076/21$ , or 3,622 hours. The average age of the failed engines, however, was only 861 hours. (United Airlines)



the average ages for all the different modes will be distributed along the exposure axis. Consequently, unless there is a dominant failure mode, the overall failure ages in complex items are usually widely dispersed and are unrelated to a specific operating age. This is a unique characteristic of complex items. A typical example is illustrated in Exhibit 2.8. In a sample of 50 newly installed Pratt & Whitney JT8D-7 engines, 29 survived beyond 2,000 operating hours. The disparate failure ages of the 21 units that failed, however, do not show any concentration about the average failure age of 861 hours.

Nevertheless, even in complex items, no matter how numerous the failure modes may be, the basic failure process reduces to the same factor—the interaction between stress and resistance to failure. Whether failures involve reduced resistance, random stress peaks, or any combination of the two, it is this interaction that brings an item to the failure point. This aspect of the failure process was summed up in a 1960 United Airlines report:<sup>\*</sup>

The airplane as a whole, its basic structure, its systems, and the various items in it are operated in an environment which causes stresses to be imposed upon them. The magnitudes, the durations and the frequencies with which specific stresses are imposed are all very variable. In many cases, the real spectrum of environmentally produced stresses is not known. The ability to withstand stress is also variable. It differs from piece to piece of new nominally identical equipment due to material differences, variations in the manufacturing processes, etc. The ability to withstand stress may also vary with the age of a piece of equipment.

It is implied that an instance of environmental stress that exceeds the failure resistance of an item at a particular time constitutes failure of that item at that time.

## 2·7 QUANTITATIVE DESCRIPTIONS OF FAILURE

Any unanticipated critical failure prompts an immediate response to prevent repetitions. In other cases, however, it is necessary to know how frequently an item is likely to fail in order to plan for reliable operation. There are several common reliability indexes based on the failure history of an item. Methods for deriving certain of these measures

failure rate  
mean time between failures  
probability of survival  
probability density of failure  
conditional probability  
of failure

<sup>\*</sup>F. S. Nowlan, A Comparison of the Potential Effectiveness of Numerical Regulatory Codes in the Fields of Overhaul Periodicity, Airplane Strength, and Airplane Performance, United Airlines Report POA-32, April 14, 1960. These remarks paraphrase a report prepared by D. J. Davis of the Rand Corporation in 1950, which offered intensive analysis of failure data. For an excellent detailed discussion of the physical processes present in the failure mechanism, see Robert P. Haviland, *Reliability and Long Life Design*, Van Nostrand Company, Inc., New York, 1964.

are discussed in detail in Appendix C, but it is helpful at this point to know what each measure actually represents.

#### **FAILURE RATE**

The failure rate is the total number of failures divided by some measure of operational exposure. In most cases the failure rate is expressed as failures per 1,000 operating hours. Thus if six failures have occurred over a period of 9,000 hours, the failure rate is ordinarily expressed as 0.667. Because measures other than operating hours are also used (flight cycles, calendar time, etc.), it is important to know the units of measure in comparing failure-rate data.

The failure rate is an especially valuable index for new equipment, since it shows whether the failure experience of an item is representative of the state of the art. It is also useful in assessing the economic desirability of product improvement. Early product-improvement decisions are based on the performance of units that have been exposed to fairly short individual periods of time in service, and this performance is adequately measured by the failure rate.

#### **MEAN TIME BETWEEN FAILURES**

The mean time between failures, another widely used reliability index, is the reciprocal of the failure rate. Thus with six failures in 9,000 operating hours, the mean time between failures would be  $9,000/6$ , or 1,500 hours. This measure has the same uses as the failure rate. Note that the mean time between failures is not necessarily the same as the average age at failure. In Exhibit 2.8, for example, the average age of the failed engines was 861 hours, whereas the mean time between failures was 3,622 hours.\*

#### **PROBABILITY OF SURVIVAL**


With more extended operating experience it becomes possible to determine the *age-reliability characteristics* of the item under study—the relationship between its operating age and its probability of failure. At this stage we can plot a survival curve, showing the probability of survival without failure as a function of operating age. This curve relates directly to the generally accepted definition of reliability:

---

Reliability is the probability that an item will survive to a specified operating age, under specified operating conditions, without failure.

---

For this reason the survival curve is commonly referred to as the *reliability function*.



**EXHIBIT 2.9** Survival curve for the Pratt & Whitney JT8D-7 engine of the Boeing 737, based on 58,432 total operating hours from May 1 to July 31, 1974. The average life is computed by partitioning along the vertical axis to form small incremental areas whose sum approximates the area under the curve. With an age limit of 1,000 hours, only the shaded area enters into this computation, since no engines can contribute to the survival curve beyond this limit, despite the fact that they would have survived had they been left in service. (United Airlines)

Exhibit 2.9 shows a typical survival curve for an aircraft turbine engine. The curve represents the percentage of installed engines that survived to the time shown on the horizontal axis, and this is usually our best estimate of the probability that any individual engine will survive to that time without failure.

A survival curve is more useful than a simple statement of the failure rate, since it can be used to predict the percentage of units that will survive to some given age. If the engines in Exhibit 2.9 were scheduled for removal at 1,000 hours, for example, 69 percent of them would survive to that age limit, whereas 31 percent could be expected to fail before then. The area under the survival curve can also be used to measure the average life of the item under consideration. If the probability scale is divided into small increments, each of which is projected to intersect the curve, the contribution of each of these incremental areas can be calculated and added to determine the average life. Thus, the triangle at the top is the contribution of the first 10 percent of the units that

fail (90 percent survive beyond this age):

$$[(\text{age at } P = 1) + (\text{age at } P = .90)] \times \frac{.10}{2}$$

The next incremental area represents the contribution to the average life of the next 10 percent of the units that fail:

$$[(\text{age at } P = .90) + (\text{age at } P = .80)] \times \frac{.10}{2}$$

and so on. Completion of this computation for the entire area under the curve would show that, with no age limit, the average life expected for each engine in service would be 1,811 hours.


Note, however, that an age limit of 1,000 hours removes all the surviving units from service at that age. In this case, therefore, the area under the curve represents only the area up to that age limit. The probability of survival to 1,000 hours is .692, so the contribution of any surviving unit to the average life is only 1,000 hours  $\times$  .692 = 692 hours. This contribution, added to the incremental contributions above it for the units that failed, yields an average realized life of 838 hours for failed and unfailed engines. Any engines that would have survived to ages higher than 1,000 hours, and thus have added to the average life, do not count. The average lives that would be realized with other age limits in this case are as follows:

| <i>age limit</i> | <i>average realized life</i> |
|------------------|------------------------------|
| 1,000 hours      | 838 hours                    |
| 2,000 hours      | 1,393 hours                  |
| 3,000 hours      | 1,685 hours                  |
| No limit         | 1,811 hours                  |

#### PROBABILITY DENSITY OF FAILURE

The probability that an engine in Exhibit 2.9 will survive to 1,000 hours is .692, and the probability that it will survive to 1,200 hours is .639. The difference between these probabilities, .053, is the probability of a failure during this 200-hour interval. In other words, an average of 5.3 out of every 100 engines that enter service can be expected to fail during this particular interval. Similarly, an average of 5.0 engines can be expected to fail during the interval from 1,200 to 1,400 hours. This measure is called the *probability density* of failure.

Exhibit 2.10 shows the probability densities for each 200-hour age interval, plotted from the probabilities of survival at each age. A decreasing percentage of the engines will fail in each successive age interval because a decreasing percentage of engines survives to enter that interval.




**EXHIBIT 2-10** Probability density of failure for the Pratt & Whitney JT8D-7 engine of the Boeing 737. Density values are plotted at the midpoint of each 200-hour interval and represent the probability that a failure will occur during this interval. (United Airlines)

#### CONDITIONAL PROBABILITY OF FAILURE

The most useful measure of the age-reliability relationship is the probability that an item entering a given age interval will fail during that interval. This measure is usually called the *conditional probability of failure*—the probability of failure, given the condition that the item enters that age interval. Sometimes it is also referred to as the hazard rate or the local failure rate.\* The conditional probability is related to both the probability of survival and the probability density. For example, an engine beginning at zero time has a probability of .692 of reaching the age of 1,000 hours; once it has reached this age, the probability density of failure in the next 200-hour interval is .053. Each engine that survives to 1,000 hours therefore has a conditional probability of failure between 1,000 and 1,200 hours of  $.053/.692$ , or .077. The complete conditional-probability curve for this engine is shown in Exhibit 2.11.

If the conditional probability of failure increases with age, we say that the item shows *wearout characteristics* and immediately wonder if an age limit would be effective in reducing the overall failure rate. (Note

\*In some literature these terms are defined in a narrower sense to mean the value obtained by computing the limit of the ratio as the age interval goes to zero.




**EXHIBIT 2-11** Conditional probability of failure for the Pratt & Whitney JT8D-7 engine of the Boeing 737. Probability values are plotted at the midpoint of each 200-hour interval and represent the average probability that an engine that survives to enter the interval will fail during this interval. (United Airlines)

that the term *wearout* in this context describes the adverse effect of age on reliability; it does not necessarily imply any evident physical change in individual units.) With an age limit of 1,000 hours the average realized life of the engine in question is 838 hours. The probability that an engine will survive to this age is .692, so the failure rate with this limit would be the probability of failure (.308) divided by the average life, or a rate of 0.37 failures per 1,000 hours.

Exhibit 2.12 shows this failure rate plotted as a function of various age limits. If the age limit is raised from 1,000 hours to 2,000 hours, the overall failure rate is 0.42, an increase of only 13.5 percent due to the second thousand hours of operation. However, the conditional probability of failure in the 200-hour interval just before each of these age limits goes up from .075 to .114, an increase of 52 percent. The *rate of increase* in the failure rate falls off with age because it depends on the conditional probability for each interval weighted by the probability of survival to that interval—and there is a continual reduction in the survival probability.

What this means is that the effectiveness of an age limit in controlling failure rates depends not only on large increases in conditional probability at higher ages, but also on a high probability of survival to



**EXHIBIT 2-12** Relationship between the failure rate and various age limits for the Pratt & Whitney JT8D-7 engine of the Boeing 737. (United Airlines)

those ages. It follows that the desirability of an age limit on any item cannot be investigated until there are sufficient operating data to construct survival and conditional-probability curves.


## 2-8 AGE-RELIABILITY CHARACTERISTICS

At one time it was believed that all equipment would show wearout characteristics, and during the years when equipment overhaul times were being rapidly extended, United Airlines developed numerous conditional-probability curves for aircraft components to ensure that the higher overhaul times were not reducing overall reliability. It was found that the conditional-probability curves fell into the six basic patterns shown in Exhibit 2.13. Pattern A is often referred to in reliability literature as the *bathtub curve*. This type of curve has three identifiable regions:

the bathtub curve  
age-reliability relationships  
of complex items  
the "life" of a complex item

- An *infant-mortality region*, the period immediately after manufacture or overhaul in which there is a relatively high probability of failure

**EXHIBIT 2-13** Age-reliability patterns. In each case the vertical axis represents the conditional probability of failure and the horizontal axis represents operating age since manufacture, overhaul, or repair. These six curves are derived from reliability analyses conducted over a number of years, during which all the items analyzed were found to be characterized by one or another of the age-reliability relationships shown. The percentages indicate the percentage of items studied that fell into each of the basic patterns (United Airlines)



46

- A region of constant and relatively low failure probability
- A *wearout region*, in which the probability of failure begins to increase rapidly with age

If the failure pattern of an item does in fact fit this curve, we are justified in concluding that the overall failure rate will be reduced if some action is taken just before this item enters the wearout zone. In these cases allowing the item to age well into the wearout region would cause an appreciable increase in the failure rate. Note, however, that such action will not have much effect on the overall rate unless there is a high probability that the item will survive to the age at which wearout appears.

The presence of a well-defined wearout region is far from universal; indeed, of the six curves in Exhibit 2.13, only *A* and *B* show wearout characteristics. It happens, however, that these two curves are associated with a great many single-celled or simple items—in the case of aircraft, such items as tires, reciprocating-engine cylinders, brake pads, turbine-engine compressor blades, and all parts of the airplane structure.

The relative frequency of each type of conditional-probability curve proved especially interesting. Some 89 percent of the items analyzed had no wearout zone; therefore their performance could not be improved by the imposition of an age limit. In fact, after a certain age the conditional probability of failure continued on at a constant rate (curves *D*, *E*, and *F*). Another 5 percent had no well-defined wearout zone (curve *C*) but did become steadily more likely to fail as age increased. For a very few of these items an age limit might prove useful, provided that it was cost-effective.

Only 6 percent of the items studied showed pronounced wearout characteristics (curves *A* and *B*). Although an age limit would be applicable to these items, as we have seen, its effectiveness depends on a high probability that the item will survive to that age. However, the conditional-probability curves make it possible to identify those items that might benefit from such a limit, and the question of effectiveness can then be investigated. Although it is often assumed that the bathtub curve is representative of most items, note that just 4 percent of the items fell into this pattern (curve *A*). Moreover, most complex items had conditional-probability curves represented by curves *C* to *F*—that is, they showed no concentration of failures directly related to operating age.

The basic difference between the failure patterns of complex and simple items has important implications for maintenance. Usually the conditional-probability curve for a complex item will show some infant mortality; often the probability of failure right after installation is fairly

high. Usually, also, the conditional-probability curve shows no marked point of increase with increasing age; the failure probability may increase gradually or remain constant, but there is no age that can be identified as the beginning of a wearout zone. For this reason, unless there is a dominant failure mode, an age limit does little or nothing to improve the overall reliability of a complex item. In fact, in many cases scheduled overhaul actually *increases* the overall failure rate by introducing a high infant-mortality rate in an otherwise stable system.

In contrast, single-celled and simple items frequently do show a direct relationship between reliability and increasing age. This is particularly true of parts subject to metal fatigue or mechanical wear and items designed as consumables. In this case an age limit based on some maximum operating age or number of stress cycles may be highly effective in improving the overall reliability of a complex item. Such limits in fact play a major role in controlling critical-failure modes, since they can be imposed on the part or component in which a given type of failure originates.

It is apparent from our discussion thus far that most statements about the "life" of equipment tell us little about its age-reliability characteristics. For example, the statement that an aircraft engine has a life of 2,000 operating hours might mean any of the following:

- ▶ No engines fail before reaching 2,000 hours.
- ▶ No critical engine failures occur before 2,000 hours.
- ▶ Half the engines fail before 2,000 hours.
- ▶ The average age of failed engines is 2,000 hours.
- ▶ The conditional probability of failure is constant below 2,000 hours.
- ▶ Some part in the engine has a life limit of 2,000 hours.

The definition of reliability is the probability that an item will survive a given operating period, under specified operating conditions, without failure. In discussions of reliability, therefore, it is insufficient to state an operating period alone as the "life" of an item. This statement has no meaning unless a *probability of survival* is associated with it.

It should also be apparent by now why the failure rate plays a relatively unimportant role in maintenance programs: it is too simple a measure. Although the frequency of failures is useful in making cost decisions and in establishing appropriate intervals for maintenance tasks, it tells us nothing about what tasks are appropriate or the consequences that dictate their objective. The effectiveness of a particular maintenance solution can be evaluated only in terms of the safety or economic consequences it is intended to prevent. By the same token, a

maintenance task must be applicable to the item in question in order to have any effect at all. Hence we must now consider the possible forms of preventive maintenance and see how an understanding of the failure process and the age-reliability characteristics of an item permit us to generate maintenance tasks on the basis of explicit criteria.

## the four basic maintenance tasks

**RCM PROGRAMS** consist of specific tasks selected on the basis of the actual reliability characteristics of the equipment they are designed to protect. All these tasks can be described in terms of four basic forms of preventive maintenance, each of which is applicable under a unique set of circumstances:

- ▶ Scheduled inspection of an item at regular intervals to find any potential failures
- ▶ Scheduled rework of an item at or before some specified age limit
- ▶ Scheduled discard of an item (or one of its parts) at or before some specified life limit
- ▶ Scheduled inspection of a hidden-function item to find any functional failures

The first three types of tasks are directed at preventing single failures and the fourth at preventing multiple failures. Inspection tasks can usually be performed without removing the item from its installed position, whereas rework and discard tasks generally require that the item be removed from the equipment and sent to a major maintenance base.

The development of a scheduled-maintenance program consists of determining which of these four tasks, if any, are both applicable and effective for a given item. *Applicability* depends on the failure characteristics of the item. Thus an inspection for potential failures can be applicable only if the item has characteristics that make it possible to define a potential-failure condition. Similarly, an age-limit task will be applicable only if the failures at which the task is directed are related to age. *Effectiveness* is a measure of the results of the task; the task objec-

tive, however, depends on the failure consequences involved. A proposed task might appear useful if it promises to reduce the overall failure rate, but it could not be considered effective if the purpose in applying it was to avoid functional failures altogether.

For inspection tasks the distinction between applicability and effectiveness is usually obvious: the item either does or does not have characteristics that make such a task applicable. For age-limit tasks, however, the distinction is sometimes blurred by the intuitive belief that the task is always applicable and therefore must also be effective. In reality imposing an age limit on an item does not in itself guarantee that its failure rate will be reduced. The issue in this case is not whether the task can be done, but whether doing it will in fact improve reliability.

### 3·1 SCHEDULED ON-CONDITION TASKS

Scheduled inspections to detect potential failures are commonly termed *on-condition tasks*, since they call for the removal or repair of individual units of an item "on the condition" that they do not meet the required standard. Such tasks are directed at specific failure modes and are based on the feasibility of defining some identifiable physical evidence of a reduced resistance to the type of failure in question. Each unit is inspected at regular intervals and remains in service until its failure resistance falls below a defined level—that is, until a potential failure is discovered. Since on-condition tasks discriminate between units that require corrective maintenance to forestall a functional failure and those units that will probably survive to the next inspection, they permit all units of the item to realize most of their useful lives.

detection of potential failures  
applicability criteria  
effectiveness criteria  
inspection intervals


This type of task is applicable to tires, brakes, many parts of an aircraft powerplant, and much of its structure. Many routine servicing tasks, such as checking oil quantity and tire pressure, are on-condition tasks. The applicability of an on-condition task depends to some extent on both maintenance technology and the design of the equipment. For example, borescope and radioisotope techniques have been developed for inspecting turbine engines, but these techniques are of value chiefly because the engines have been designed to facilitate their use.

If on-condition tasks were universally applicable, all failure possibilities could be dealt with in this way. Unfortunately there are many types of failures in which the failure mode is not clearly understood or is unpredictable or gives insufficient warning for preventive measures to be effective. There are three criteria that must be met for an on-condition task to be applicable:

- ▶ It must be possible to detect reduced failure resistance for a specific failure mode.
- ▶ It must be possible to define a potential-failure condition that can be detected by an explicit task.
- ▶ There must be a reasonably consistent age interval between the time of potential failure and the time of functional failure.

As an example, suppose a visible crack is used as a measure of metal fatigue, as shown in Exhibit 3.1. Such an item is most failure resistant when it is new (point A). The resistance drops steadily with increasing age and is already somewhat reduced by the time a crack appears (point B). Thereafter it is possible to monitor the growth of the crack and define a potential-failure point C far enough in advance to permit removal of the item before a functional failure occurs (point D). Once a crack has appeared, the failure resistance drops more rapidly; hence the rate of crack growth in this item must be known in order to establish an inspection interval  $\Delta T$  that will effectively control this failure mode.

The data for the entire population of this item would define a range of failure ages rather than one specific age. Hence both the defined potential failure and the frequency of inspections depend on the objective of the task. If a functional failure would have safety consequences, then the objective is to prevent all such failures. In this case an on-condition task may be applicable, but it would be considered effective only if it minimized the likelihood of a critical failure. If the failure does not involve safety, then effectiveness is measured in economic terms — that is, the task is effective only if it is *cost-effective*. In the case of operational consequences this means that the cost of finding and correcting potential failures must be less than the combined cost of the operational consequences plus the cost of repairing the failed units. It follows from



**EXHIBIT 3-1** Determining the interval for on-condition inspection of an item subject to metal fatigue. Once the rate of decline in failure resistance has been determined, an inspection interval  $\Delta T$  is established that provides ample opportunity to detect a potential failure before a functional failure can occur.

this that when an on-condition task is effective in reducing the failure rate, and hence the frequency of operational consequences, it is usually also cost-effective, since the cost of inspection is relatively low.

Exhibit 3.2 shows some typical on-condition tasks for an aircraft. The first example concerns a specific failure mode of an aircraft engine that has a set of 24 tie bolts between the fourth and fifth stages of its turbine to hold an air seal in position (and a similar set of tie bolts between the fifth and sixth stages). Failure of this set of tie bolts would result in a loose air seal and cause major damage to the engine. Lowered resistance to failure is evidenced by the failure of one or more individual bolts. (Note that although this would be a functional failure of the tie bolts, it is a potential failure from the standpoint of the engine.)

The second example concerns the nozzle guide vanes of the same engine. These vanes are subject to burning by the hot exhaust gases of the engine, and also to erosion by hard carbon particles from the combustor. The required borescope inspection is a visual inspection to determine how much damage has occurred on the airfoil and inner platform of the vane. The definition of potential-failure conditions in this case is quite complex; in practice the interval between inspections is reduced as the condition deteriorates, until a point is reached at which the engine must be removed from service.

**EXHIBIT 3•2** Examples of on-condition inspection tasks as specified for maintenance mechanics. (United Airlines)

**1 LOW-PRESSURE TURBINE SECTION**

Check for failed airseal tie bolts.

Note: Airseal tie bolts between fourth- and fifth-stage and sixth-stage rotors (last three stages) are failing. These broken bolts are trapped in the airseal between the rotors and cause a rattling sound as they roll when the turbine is slowly rotated.

- A Have fan rotated 180 degrees very slowly. Repeat 180-degree rotation as often as necessary.
- B Listen at tailcone for rattling sound caused by broken bolts rolling around (do not confuse with clanking sound of blades). Attempt to determine number of broken bolts by counting rattles.
- C Failed-bolt limits.  
Three or fewer broken bolts: engine may remain in service.  
Four or more broken bolts: engine must be borescoped within 75 hours.
- D Supply the following information:
  - (1) Plane number \_\_\_\_\_ Engine position \_\_\_\_\_  
Engine time since last shop visit \_\_\_\_\_
  - (2) Number of broken bolts estimated from  
"listening" check \_\_\_\_\_
- E Send DIS•P5106 message giving above information.

**2 FIRST-STAGE NOZZLE GUIDE VANES**

Borescope inspection (Boeing 747 JT9D powerplant).

- A Perform initial borescope inspection of first-stage nozzle guide vanes at 600 hours. Perform repeat inspections at 600, 200, 75, or 30 hours, depending on conditions found.
- B Distress limits as given in MM/OV 72-00-99:
  - (1) Trailing-edge cracks: maximum of 5 cracks per vane extending to window (slot) leading edge. If distress exceeds this limit, remove engine; otherwise, repeat inspection in 600 hours.
  - (2) Trailing-edge erosion: If burning-surface burn-through does not exceed 1/2 by 1/2 inch, repeat inspection in 600 hours; if burn-through does not exceed 3/4 by 3/4 inch, repeat inspection in 200 hours; if burn-through does not exceed 1 by 1 inch, repeat inspection in 75 hours. If surface burn-through is up to 5/8 inch from leading edge, repeat inspection in 30 hours.

Note: 30-hour limit is a maximum fly-back limit, to be used one time only.

### **3 FIRE-DETECTOR INSTALLATIONS**

Intensified inspection of installations, leads, and connections.

- A** Check for minimum clearance of 1/16 inch between sensing elements and engine, as well as between various engine components. Provide necessary clearance.
- B** Check for any signs of wear.
- C** Wear limits:
  - Acceptable: Flat spots not exceeding 0.035 inch in width; any length acceptable.
  - Not acceptable: Flat spots exceeding 0.035 inch in width or worn spot exposing inner conductor or composition material between inner conductor and outer sensing-element shell.

Note: Nominal diameter is 0.070 inch.

### **4 BRAKE ASSEMBLY, MAIN LANDING GEAR**

Check brake-lining wear at each assembly, using small scale.

- A** Set parking brakes.
- B** Measure wear-indicator pin extension at both indicator pins.
- C** Wear limits:
  - If either pin is less than 0.25 inch in length, replace brake assembly.

Note: Replacement may be deferred, with approval from SFOLM, provided wear-indicator pin measures longer than 13/64 inch. If wear-indicator pin length is 13/64 inch or less, immediate replacement is required.

### **5 PNEUMATIC DRIVE UNITS, LEADING EDGE FLAP**

Check oil level and service as required.

Note: Drive units are numbered from outboard to inboard, 1 to 4, left and right wing.

- A** Check oil level in proper sight glass. If oil level is visible in sight glass, no service is required.
- B** If oil is not visible, slowly add oil (OIL 2380) through fill port until sight glass is filled. Use 53769 oil dispenser.
- C** Allow excess oil to drain out before installing fill plug.

In the third example the potential failure may be either lack of adequate clearance or visible wear on fire-detector sensing elements and leads. The fourth and fifth examples involve less judgment in the inspection process. Exact limits are given for the brake wear-indicator pin in the first case and oil level in the pneumatic unit in the second case. Both require a clearcut response on the part of the inspecting mechanic.

Whenever an on-condition task is applicable, it is the most desirable type of preventive maintenance. Not only does it avoid the premature removal of units that are still in satisfactory condition, but the cost of correcting potential failures is often far less than the cost of correcting functional failures, especially those that cause extensive secondary damage. For this reason on-condition inspection tasks are steadily replacing older practices for the maintenance of airline equipment.

### 3·2 SCHEDULED REWORK TASKS

applicability criteria  
effectiveness criteria  
effect of an age limit on  
maintenance workload

Many single-celled and simple items display wearout characteristics—that is, the probability of their failure becomes significantly greater after a certain operating age. When an item does have an identifiable wearout age, its overall failure rate can sometimes be reduced by imposing a hard-time limit on all units to prevent operation at the ages of higher failure frequency. If the item is such that its original failure resistance can be restored by rework or remanufacture, the necessary rework task may be scheduled at appropriate intervals.\* For example, the airplane tire in Exhibit 2.4 could have been scheduled for rework after a specified number of landings, since retreading restores the original failure resistance. However, this would have resulted in the retreading of all tires at the specified age limit, whether they needed it or not, and would not have prevented functional failures in those tires that failed earlier than anticipated.

Where no potential-failure condition can be defined, on-condition inspection of individual units is not feasible. In such cases a rework task may be applicable, either for a simple item or to control a specific failure mode in a complex item. Although the age limit will be wasteful for some units and ineffective for others, the net effect on the entire population of that item will be favorable. This is not the case, however, for complete rework of a complex item. As we saw in Chapter 2, failures in complex items are the result of many different failure modes, each of

\*The term *overhaul* has the connotation that the unit is completely disassembled and remanufactured part by part to restore it as nearly as possible to a "like-new" physical condition. *Rework* refers to a set of maintenance operations considered sufficient to restore the unit's original resistance to failure. Thus rework for specific items may range from replacement of a single part to complete remanufacture.

which may occur at a different average age. Consequently the overall failure rate of such items remains relatively constant; in some cases reliability decreases gradually with age, but there is no particular age that can be identified as a wearout zone. Thus, unless there is a dominant failure mode which is eliminated in the course of rework, complete rework of a complex item will have little or no effect on the overall failure rate.

A rework task can be considered applicable to an item only if the following criteria are met:

- ▶ There must be an identifiable age at which the item shows a rapid increase in the conditional probability of failure.
- ▶ A large proportion of the units must survive to that age.
- ▶ It must be possible to restore the original failure resistance of the item by reworking it.

Because the information required to develop survival and conditional-probability curves for an item is not available when equipment first goes into service, scheduled rework tasks rarely appear in a prior-to-service maintenance program (only seven components were assigned to scheduled rework in the initial program developed for the Douglas DC-10). Often, however, those items subject to very expensive failures are put into an age-exploration program to find out as soon as possible whether they would benefit from scheduled rework.

Even when scheduled rework is applicable to an item, very often it does not meet the conditions for effectiveness. A reduction in the number of expected failures, for example, would not be sufficient in the case of safety consequences, and in the case of economic consequences the task must be cost-effective. Moreover, since an age limit lowers the average realized age of an item, it always increases the total number of units sent to the shop for rework.

As an example, consider the effect scheduled rework would have on the turbine engine discussed in Section 2.7. With no age limit, the failure rate of these engines is 0.552 failures per 1,000 hours. Thus over an operating period of 1 million hours an average of 552.2 failed units ( $1,000,000/1,811$ ) are sent to the shop for repair (see Exhibit 3.3). A rework age limit of 2,000 hours will reduce the failure rate to 0.416; however, it will also reduce the average realized age from 1,811 hours to 1,393 hours. Since 42 percent of the units survive to 2,000 hours, over the same operating period an average of 717.9 units would be sent to the shop—the 416.3 units that failed plus the additional 301.6 scheduled removals. In other words, there would be about 135 fewer failures, but 166 more engines that required rework. On this basis scheduled rework at 2,000-hour intervals would not be cost-effective unless the rework cost for scheduled removals were substantially lower than the cost of

| age limit<br>(hours) | failure rate<br>(per 1,000 hours) | percentage of units<br>surviving to age limit | averaged realized<br>engine age (hours) |
|----------------------|-----------------------------------|-----------------------------------------------|-----------------------------------------|
| 1,000                | 0.3681                            | 69.2                                          | 838                                     |
| 2,000                | 0.4163                            | 42.0                                          | 1,393                                   |
| 3,000                | 0.4871                            | 17.9                                          | 1,685                                   |
| None                 | 0.5522                            | 0                                             | 1,811                                   |

**EXHIBIT 3-3** Effect of several rework age limits on shop workload. The total number of engines sent to the shop is computed by dividing the total hours of engine operation by the average realized age for each age limit. The number of scheduled removals is then the percentage of this total that survives to the age limit in question.

repairing failures (in this case the rework cost would have to be less than 135.9/301.6, or 45.1 percent, of the repair cost).

Of course, the direct cost of rework is not the only economic factor to be taken into account. If the failure is one that has operational consequences, the reduction in the number of failures may more than offset the additional cost of rework. Determining the economic desirability of a proposed rework age limit will be discussed in greater detail in the next chapter. In general, however, the effect of at least four possible rework intervals must be analyzed before an optimum limit can be determined—if indeed one does exist. In most cases a rework task will not prove cost-effective unless the item has an unusually expensive failure mode or the cost of a functional failure includes economic losses other than the direct cost of repair.

### 3-3 SCHEDULED DISCARD TASKS

criteria for safe-life limits  
criteria for economic-life limits

The scheduled rework of items at a specified age limit is one type of hard-time task; the other is scheduled discard of items or certain of their parts at some specified operating age. Such tasks are frequently termed *life-limit tasks*. Life limits may be established to avoid critical failures, in which case they are called *safe-life limits*, or they may be established because they are cost-effective in preventing noncritical failures, in which case they are called *economic-life limits*.

#### SAFE-LIFE LIMITS

A safe-life limit is imposed on an item only when safety is involved and there is no observable condition that can be defined as a potential failure. In this case the item is removed at or before the specified maximum

| shop workload per 1,000,000 engine operating hours |                    |                |
|----------------------------------------------------|--------------------|----------------|
| failed engines                                     | scheduled removals | total workload |
| 368.1                                              | 825.2              | 1,193.3        |
| 416.3                                              | 301.6              | 717.9          |
| 487.1                                              | 106.4              | 593.5          |
| 552.2                                              | —                  | 552.2          |


age and is either discarded or disassembled for discard of a time-expired part. This practice is most useful for simple items or individual parts of complex items, such as pyrotechnic devices in ejection seats, which have a limited shelf life, and turbine-engine disks or nonredundant structural members, which are subject to metal fatigue.

The safe-life limit itself is usually established by the equipment manufacturer on the basis of developmental testing. A component whose failure would be critical is designed to begin with to have a very long life. It is then tested in a simulated operating environment to determine what average life has actually been achieved, and a conservatively safe fraction of this average life is used as the safe-life limit.

Safe-life items are nearly always single-celled parts, and their ages at failure are grouped fairly closely about the average. However, the correlation between a test environment and the actual operating environment is never perfect. Moreover, because testing a long-lived part to failure is both time-consuming and expensive, the volume of test data is often too small to permit us to draw a survival curve with much confidence. For this reason safe-life limits are usually established by dividing the average failure age by a large arbitrary factor—sometimes a factor as large as 3 or 4. The implication is that the conditional probability of failure at this limit is essentially zero; that is, a safe-life limit is based on a 100 percent probability of survival to that age. The difference between a safe-life limit and the average age at failure is illustrated in Exhibit 3.4.

A safe-life discard task is applicable only under the following circumstances:

- ▶ The item must be subject to a critical failure.
- ▶ Test data must show that no failures are expected to occur below the specified life limit.



**EXHIBIT 3-4** Comparison of the average age at failure (average life) determined from operating data and a safe-life limit determined on the basis of test data.

Since the function of a safe-life limit is to avoid the occurrence of a critical failure, the resulting discard task is effective only if it accomplishes this objective. Thus the only information for assessing effectiveness in this case will be the manufacturer's test data. Sometimes these tests have not been completed at the time the initial program is developed, but until a limit can be established, the available test data must show that the anticipated in-service aging of the item will be safe. An operating organization rarely has the facilities for further simulation testing that might justify increasing a safe-life limit, nor is there usually a reasonable basis for reducing it, unless failures occur.

#### **ECONOMIC-LIFE LIMITS**

In some instances extensive operating experience may indicate that scheduled discard of an item is desirable on purely economic grounds. An *economic-life limit*, however, is established in the same manner as an age limit for scheduled rework; that is, it is based on the actual age-reliability relationship of the item, rather than on some fraction of the average age at failure. Whereas the objective of a safe-life limit is

to avoid accumulating any failure data, the only justification for an economic-life limit is cost effectiveness. Thus the failure rate must be known in order to predict how the total number of scheduled removals at various age limits would affect the cost-benefit ratio.

In general, an economic-life task requires the following three conditions:

- ▶ The item must be subject to a failure that has major economic (but not safety) consequences.
- ▶ There must be an identifiable age at which the item shows a rapid increase in the conditional probability of failure.
- ▶ A large proportion of the units must survive to that age.

Although an item that meets the first criterion may be put into an age-exploration program to find out if a life limit is applicable, there are rarely sufficient grounds for including this type of discard task in an initial scheduled-maintenance program.

### 3 • 4 SCHEDULED FAILURE-FINDING TASKS

Whenever an item is subject to a functional failure that would not be evident to the operating crew, a scheduled task is necessary to protect the availability of that function. Although hidden-function failures, by definition, have no immediate consequences, failures that are undetected increase the exposure to a possible multiple failure. Hence, if no other type of maintenance task is applicable and effective, hidden-function items are assigned *failure-finding tasks*, scheduled inspections for hidden failures. Although such tasks are intended to locate functional failures rather than potential failures, they can be viewed as a type of on-condition maintenance, since the failure of a hidden-function item can also be viewed as a potential multiple failure. The chief difference is in the level of item considered; a functional failure of one item may be only a potential failure for the equipment as a whole.

Most items supported by failure-finding inspections remain in service until a functional failure is discovered. Some items, however, have several functions, of which only one or a few are hidden. Such items will be removed from service to correct evident failures, and if the removal rate is sufficient to ensure adequate availability of the hidden function, the shop specifications may include a failure-finding inspection at that time. Other items may not require scheduled failure-finding tasks because the operating crew is required to check them periodically. Many hidden functions, especially in systems, are made evident by the addition of instrumentation, so that a separate inspection for hidden failures is unnecessary.

applicability criteria  
determining average  
availability

A scheduled failure-finding task is applicable to an item under the following two conditions. Note that the second criterion is in fact a default condition:

- ▶ The item must be subject to a functional failure that is not evident to the operating crew during the performance of normal duties.
- ▶ The item must be one for which no other type of task is applicable and effective.

The objective of a failure-finding task is to ensure adequate availability of a hidden function. The level of availability that is needed, however, depends on the nature of the function and the consequences of a possible multiple failure. Some hidden functions, such as the fire-warning system in an aircraft powerplant, are sufficiently important that they are tested before every flight.

Appropriate intervals for failure-finding tasks cannot be determined as exactly as those for other types of tasks. In the case of emergency equipment hidden-function items which are replaced at specified intervals, such as pyrotechnic devices, are tested prior to rework or discard to see if they would have functioned had they been needed. The test results at any given interval provide a basis for increasing or decreasing the interval. In other cases the expected availability of a hidden function can be approximated by assuming that the age-reliability relationship is exponential,\* assigning a conservatively high failure rate, and then determining the probability of survival across a given inspection interval.


As an example, suppose some hidden function has an anticipated failure rate of 0.5 per 1,000 hours. The mean time between failures is then 2,000 hours. If the proposed inspection interval is 500 hours, a unit that is serviceable at one inspection will have aged 500 hours by the next inspection. The probability that it will survive this 500-hour interval (one-fourth of the mean time between failures) is .78 on an exponential curve (Exhibit 3.5). The average availability would thus be

$$\frac{1.00 + 0.78}{2} = 0.89$$

or a probability of .89 that the item will function if it is needed. If this degree of reliability is inadequate, the inspection interval must be reduced. Failure-finding tasks are always effective if the inspection interval is short enough.

To be considered effective a failure-finding task must ensure the required level of availability. However, this task must also be cost-

\*If the conditional probability of failure is nonincreasing, this is a conservative assumption.



**EXHIBIT 3-5** Establishing the interval for a failure-finding inspection. The age-reliability relationship of an item is assumed, in the absence of information, to be exponential over operating age. Thus at an inspection interval equal to one-fourth of the mean time between failures, the probability that the item will survive that interval is .78. This is true of the interval between any two inspections, regardless of the age of the item. On the basis of this inspection interval, the average availability of the unit would be 89 percent. An interval that represented a smaller fraction of the expected mean time between failures would yield a higher average availability.

effective with respect to the three other types of maintenance tasks—that is, it must be the least expensive means of ensuring the necessary level of availability. When a possible multiple failure is not related to safety, an availability goal of 95 percent is often used. Alternatively, the economic consequences of the multiple failure can be balanced against the costs of inspection to determine the most cost-effective interval and availability level.

Exhibit 3.6 shows some typical failure-finding tasks for a commercial aircraft. In each case the scheduled task is designed to identify a functional failure. In the second example the failure might or might not be evident to the operating crew, depending on whether a complaint was received from a passenger.

**EXHIBIT 3-6** Examples of failure-finding inspection tasks as specified for airline maintenance mechanics. In this case the mechanic is required only to replace the failed units. (United Airlines)

**1 SMOKE GOGGLES**

Replace missing or damaged goggles (not repairable) as required by the following conditions:

- A Plastic-foam face seal not adhering to goggle rim
- B Lens not retained within goggle groove
- C Dirt or scratches on lens
- D Any other detrimental condition

**2 READING LIGHTS, PASSENGER-SERVICE SYSTEM**

Test lights in zones A to E.

- A At positions 1, 2, 3, and 4 on right attendant's panel, position switches as follows:  
PES—OFF, PSS—OFF, CH—OFF, ATTND CALL—TEST (to illuminate blue)
- B For zone being checked, rotate reading-light switch to ON position:
  - (1) All reading lights in that zone should illuminate.
  - (2) Master call light should not blink.
- C Rotate reading-light switch to OFF position:
  - (1) All reading lights in that zone should not be illuminated.
  - (2) Master call light should not blink.
- D Rotate reading-light switch to SEAT position:  
All reading lights in that zone should return to individual seat CTL selector.

**3 EXTERIOR LIGHTS**

- A Turn on beacon, navigation, and wing-illumination lights, and at night turn on logo lights.
- B Walk around exterior of aircraft and check lights.
- C Turn off lights.

### 3.5 CHARACTERISTICS OF THE BASIC TASKS

The four types of scheduled-maintenance tasks employed in an RCM program differ both in terminology and in concept from traditional approaches to scheduled maintenance. In the airline industry, for example, it is customary to refer to three "primary maintenance processes": *on-condition*, *hard time*, and *condition monitoring*. All scheduled tasks are considered to be either on-condition or hard-time. On-condition tasks are defined by FAA regulations as:

... restricted to components on which a determination of continued airworthiness may be made by visual inspection, measurement, tests, or other means without a teardown inspection or overhaul. These "On-Condition" checks are to be performed within the time limitations prescribed for the inspection or check.

Although the term *hard time* is not specifically defined, it is implied by a number of FAA requirements. Airline maintenance specifications must include "time limitations, or standards for determining time limitations, for overhauls, inspections and checks of airframes, engines, propellers, appliances, and emergency equipment," and the basic principle for establishing these time limitations is:

... that the inspections, checks, maintenance or overhaul be performed at times well within the expected or proven service life of each component of the aircraft.

terminology differences  
the basis of task preference  
items that cannot benefit from  
scheduled maintenance

**EXHIBIT 3.7** Comparison of RCM task terminology and current regulatory usage.

| RCM terminology                                            | current regulatory usage                                           |
|------------------------------------------------------------|--------------------------------------------------------------------|
| <b>Inspection tasks:</b>                                   |                                                                    |
| On-condition tasks (to detect potential failures)          | On-condition process                                               |
| Failure-finding tasks (to detect hidden function failures) | Condition-monitoring process (inspection of hidden-function items) |
| <b>Removal tasks:</b>                                      |                                                                    |
| Scheduled rework                                           | Hard-time process                                                  |
| Scheduled discard                                          | Scheduled overhaul                                                 |
|                                                            | Life limit                                                         |
| <b>Servicing tasks</b>                                     | Servicing                                                          |
| <b>No scheduled maintenance</b>                            | Condition-monitoring process (no scheduled tasks)                  |

The process termed *condition monitoring* is one that is characterized by the absence of preventive-maintenance tasks. An item is said to be maintained by condition monitoring if it is permitted to remain in service without preventive maintenance until a functional failure occurs. However, since condition monitoring is oriented to after-the-fact detection of failures, this designation may refer in some instances to failure-finding tasks assigned to hidden-function items and in other instances to items assigned to no scheduled maintenance.

Despite the overlap in terminology, there are certain fundamental differences in concept between the tasks performed under traditional maintenance policies and the explicit task definitions required by an RCM program. The hard-time approach was based on the assumption that complex items do have an "expected or proven service life"—that is, that their overall reliability invariably decreases with age. On this premise overhaul specifications usually required that all units which had survived to the specified time limit be disassembled down to their smallest constituent parts and inspected in detail for signs of deterioration. Technical experts examined each part and formed opinions about whether a given component could have continued to operate satisfactorily to a projected new overhaul interval; in other words, they made judgments about the age at which the item was likely to fail.

These teardown inspections might at first appear to qualify as on-condition inspections. However, such inspections were rarely focused on the specific conditions required by an on-condition task. Unfortunately it is usually beyond human capability to look at a used part and determine what its likelihood of failure will be at some later age. As a result, the initial overhaul intervals for new equipment were short and were extended only by very small increments. At one point, in fact, the FAA limited extensions of the interval for engine overhauls to a maximum of 100 hours and required a period of at least three months between successive extensions.

Note that the traditional type of scheduled overhaul also fails to satisfy the criteria for a rework task. Shop specifications calling for the part-by-part remanufacture of complex items to restore them to "like-new" condition were intended to avoid operation in the age period at which failures were expected to be more likely. As we have seen, however, this expectation does not hold for most complex items. Consequently we cannot expect periodic overhaul at any operating age to make a noticeable difference in their reliability. Furthermore, even when a complex item does meet the applicability criteria for a rework task, it is difficult to satisfy the conditions for effectiveness. For this reason complete rework of items such as turbine engines is now relatively rare, and many organizations have abandoned rework of other rotating machinery, which was once considered a prime candidate for scheduled overhaul.

### THE BASIS OF TASK PREFERENCE

The applicability of any maintenance task depends on the failure characteristics of the item. However, the characteristics of the tasks themselves suggest a strong order of preference on the basis of their overall effectiveness as preventive measures. The first choice is always an on-condition inspection, particularly if it can be performed without removing the item from the equipment. This type of preventive maintenance has a number of advantages. Because on-condition tasks identify individual units at the potential-failure stage, they are particularly effective in preventing specific modes of failure. Hence they reduce the likelihood both of critical failures and of the operational consequences that would otherwise result from that failure mode. For the same reason, they also reduce the average cost of repair by avoiding the expensive secondary damage that might be caused by a functional failure.

The fact that on-condition tasks identify individual units at the point of potential failure means that each unit realizes almost all of its useful life. Since the number of removals for potential failures is only slightly larger than the number that would result from functional failures, both the repair costs and the number of spare units necessary to support the repair process are kept to a minimum. The scheduling of on-condition inspections at a time when the equipment is out of service concentrates the discovery of potential failures at the maintenance stations that perform the inspections. This fact, together with the lower probability of functional failures, further reduces the inventory of spare units that would otherwise have to be kept available at each line station.

If no applicable and effective on-condition task can be found, the next choice is a scheduled rework task. Scheduled rework of single parts or components leads to a marked reduction in the overall failure rate of items that have a dominant failure mode (the failures resulting from this mode would be concentrated about an average age). This type of task may be cost-effective if the failures have major economic consequences. As with on-condition inspections, the scheduled removals can be concentrated at a few maintenance stations, thus reducing the exposure of all line stations to the need to remove units after they have failed. A rework age limit usually includes no restriction on the remanufacture and reuse of time-expired units; hence material costs are lower than they would be if the entire unit had to be discarded.

Any scheduled rework task, however, has certain disadvantages. Because the age limit applies to all units of an item, many serviceable units will be removed that would otherwise have survived to higher ages. Moreover, as we saw in Section 3.2, the total number of removals will consist of failed units plus scheduled removals. Hence the total workload for this task is substantially greater than it would be with on-condition inspection, and a correspondingly larger number of spare units is needed to support the shop process.

Scheduled discard is economically the least desirable of the three directly preventive tasks, although it does have a few desirable features. A safe-life limit on simple components can prevent critical failures caused by certain failure modes. Similarly, an economic-life limit can reduce the frequency of functional failures that have major economic consequences. However, a discard task is in itself quite costly. The average life realized by an item subject to a safe-life limit is only a fraction of its potentially useful life, and the average life of an item subject to an economic-life limit is much less than the useful life of many indi-

**EXHIBIT 3-8** Comparison of various characteristics of the four basic scheduled-maintenance tasks.

| characteristic                              | on-condition task                                                                                                                              | scheduled rework task                                                                                                                                                                                 |
|---------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Applicability criteria                      | Reduced resistance to failure must be detectable; rate of reduction in failure resistance must be predictable.                                 | Conditional probability of failure must increase at an identifiable age; a large proportion of the units must survive to that age.                                                                    |
| Effectiveness criteria                      | For critical failures the task must reduce the risk of failure to an acceptable level; in all other cases the task must be cost-effective.     | For critical failures the task must reduce the risk of failure to an acceptable level (a rework task alone is unlikely to meet this requirement); in all other cases the task must be cost-effective. |
| Usual availability of required information  | Applicability prior to service; effectiveness after age exploration.                                                                           | Applicability after age exploration; effectiveness after age exploration.                                                                                                                             |
| Effect on occurrence of functional failures | Failures due to specific failure mode eliminated or greatly reduced in frequency.                                                              | Frequency of failures somewhat less than with no scheduled maintenance.                                                                                                                               |
| Distribution of removals                    | Removals for potential failures concentrated at few stations where inspections are performed; removals for functional failures at any station. | Scheduled removals concentrated at a very few stations; removals for functional failures at any station.                                                                                              |
| Effect on shop volume                       | Slightly greater than with no scheduled maintenance.                                                                                           | Much greater than with on-condition or no scheduled maintenance.                                                                                                                                      |

vidual units. In addition, a discard task involves the cost of replacement; new items or parts must be purchased to replace the time-expired units, since a life limit usually does not permit remanufacture and reuse.

Hidden-function failures have no immediate consequences; hence our interest is in the least expensive means of ensuring the necessary level of availability for the item. When none of the other three tasks is applicable, the default action for hidden-function items is a failure-finding task. Otherwise, the choice of task is determined by cost effectiveness.

| scheduled discard task                                                                                                                                                                                                                          | failure-finding task                                                                                                         |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| For safe-life items conditional probability of failure must be zero below life limit; for economic-life items conditional probability of failure must increase at an identifiable age and a large proportion of units must survive to that age. | The occurrence of a functional failure must not be evident to the operating crew.                                            |
| A safe-life limit must reduce the risk of failure to an acceptable level; an economic-life limit must be cost-effective.                                                                                                                        | The task must result in the level of availability necessary to reduce the risk of a multiple failure to an acceptable level. |
| Safe-life applicability and effectiveness prior to service; economic-life applicability and effectiveness after age exploration.                                                                                                                | Applicability prior to service; effectiveness after age exploration.                                                         |
| Failures due to specific failure mode eliminated (safe-life limit) or reduced in frequency (economic-life limit).                                                                                                                               | No effect on item inspected, but frequency of multiple failures greatly reduced.                                             |
| Scheduled removals concentrated at a very few stations; removals for functional failures (economic-life limit) at any station.                                                                                                                  | Removals concentrated at stations where inspections are performed; no removals at other stations.                            |
| Not applicable.                                                                                                                                                                                                                                 | Minimal.                                                                                                                     |

#### **ITEMS THAT CANNOT BENEFIT FROM SCHEDULED MAINTENANCE**

In the process of evaluating proposed maintenance tasks for an item there will be a number of instances in which no applicable task can be found—that is, items for which there is no evidence that a particular task will improve reliability. There will be far more instances, however, in which an applicable task does not satisfy the conditions for effectiveness. This may be because the failure has such minor consequences that the task is not cost-effective or because it has such major consequences that the task does not reduce the risk of failure to the required level. If safety consequences are involved, the objective of any task is to minimize the probability of a failure, and in this case all applicable tasks are assigned as preventive maintenance. Since most essential functions in well-designed equipment are protected by redundancy, the safety hazard is usually the possible secondary damage. However, the number of failure modes in which this is a factor is relatively small.

When an item cannot benefit from scheduled maintenance, in some cases product improvement may be necessary before the equipment goes into service. More often the chore of determining what preventive maintenance might accomplish for each item helps to clarify specific modifications that would improve reliability in subsequent designs.

Where safety consequences are not involved, any applicable task must be cost-effective, and this condition is usually difficult to satisfy unless the failure has operational consequences. Once again, the design often employs redundancy to limit the number of items subject to such failures. As a result, there are tens of thousands of items on complex equipment for which scheduled maintenance provides no advantage. Since such items cannot benefit from preventive maintenance, they are left in operation until a functional failure occurs. This strategy permits each unit to realize its maximum useful life.

Items that cannot benefit from scheduled maintenance are characterized by two properties:

- ▶ Such items have no hidden functions; hence a failure is evident to the operating crew and will therefore be reported and corrected.
- ▶ The failure is one that has no direct adverse effect on operating safety.

A further characteristic of such items is that many of them are complex. One reason for this is that when there is no evidence that a proposed task will actually improve the reliability of a complex item, there is always the possibility that it will introduce new problems, either by upsetting a stable state or, in some cases, by introducing workmanship problems. Thus where a complex system cannot be protected by on-condition inspections, from a purely practical standpoint the default action would be no scheduled maintenance. This is usually the case, for example, with electrical and electronic systems.

## 3·6 THE DIMENSIONS OF A SCHEDULED-MAINTENANCE PROGRAM

### THE ROLE OF THE BASIC TASKS

The maintenance activities required to support any type of complex equipment include routine servicing, periodic inspections, and the performance of any corrective maintenance necessary when a condition is found to be unsatisfactory. Scheduled tasks are selected, however, on the basis of the ways in which a particular item can fail. In considering all the known or anticipated failure modes of each item we find that many major components cannot benefit from any type of preventive maintenance, some will require a single task, and others will require several different tasks. The maintenance tasks assigned to a complex item such as an aircraft turbine engine, for example, are quite numerous. Following are just a few of the inspection tasks performed while the engine is installed:

- ▶ Oil-screen inspection to detect metal particles
- ▶ Borescope inspection of the combustor to detect signs of metal fatigue
- ▶ "Sniff test" of the fuel manifold to detect fuel odors
- ▶ "Broomstick check" to detect loose turbine blades
- ▶ Inspection of the fan blades and front compressor blades for possible damage
- ▶ Inspection for rattling noise to detect broken tie bolts
- ▶ Radioisotope inspection of nozzle guide vanes for deformation
- ▶ Spectrographic oil analysis to detect metallic indications of wear

Recognition of the criteria for applicability of scheduled rework has led to a great reduction in the number of items removed and sent to the shop for routine overhaul. Items are still removed from equipment and sent to the maintenance base, however, either because they have failed or because they contain parts that require scheduled rework or discard. In this case it is necessary to decide the extent of the work to be done before these items are returned to service. Within the frame of reference dictated by the applicability of rework tasks, there are only four circumstances under which rework would be specified:

- ▶ Single parts may require rework as the result of an inspection for potential failures that can be performed only when an item is disassembled in the shop. This applies to certain types of turbine blades.

the role of the basic tasks  
servicing and lubrication tasks  
zonal inspections and  
walkaround checks  
the total maintenance workload

- ▶ Single parts may require rework because their failure characteristics show that they will benefit from an age limit. This is the case with some fuel manifolds.
- ▶ Single parts may have to be discarded because they have reached a specified life limit. This applies to the safe-life limits imposed on most compressor and turbine disks.
- ▶ Single parts may have to be reworked or discarded because shop inspection discloses a functional failure that was not observable when the item was installed on the equipment.

The amount of work specified as part of shop maintenance depends, of course, on the nature of the item. With some the direct cause of a failure is corrected, and if the component can then meet its performance standards, it is returned to service. This practice is sometimes referred to as *conditional overhaul*. Other items, such as turbine engines, may have a great deal of additional work done on them while they are out of service. The work performed, however, is very much less than that done under hard-time overhaul policies. As a result, the RCM approach to rework tasks has substantially decreased engine maintenance costs, not only by reducing the volume of units flowing through the maintenance base, but also by reducing the amount of work required when they are there.

The propulsion system is not the only complex item on an aircraft; however, it is a system closely associated with operating safety, and the largest part of the maintenance costs for any aircraft stem from scheduled or unscheduled work on engines. Because of this, on-condition inspections play a major role in powerplant maintenance programs, and scheduled removals, when they are necessary, are set at the maximum interval that will allow satisfactory operation.

#### **SERVICING AND LUBRICATION**

Complex equipment requires numerous scheduled servicing and lubrication tasks to maintain satisfactory operation. There is usually no question about which tasks are required and whether they are applicable and effective. However, it is interesting to review this aspect of maintenance in light of our discussion thus far.

Lubrication, for example, really constitutes scheduled discard of a single-celled item (the old lubrication film). This task is applicable because the film does deteriorate with operating age and show wearout characteristics. Usually the condition of the film cannot be determined; hence conservatively short intervals are assigned for its replacement with new lubricant. Such tasks are also cost-effective. An item is lubricated whether it needs lubrication or not because the cost is minuscule in

comparison to the costs that would result from inadequate lubrication. In fact, the cost of this task is too low to justify studies to determine the most economical task interval. As a result, lubrication is rarely isolated for in-depth analysis in developing a maintenance program.

Whereas lubrication constitutes a discard task, the servicing tasks—checking tire pressure or fluid levels in oil and hydraulic systems—are on-condition tasks. In this case potential failures are represented by pressure or fluid levels below the replenishment level, and this condition is corrected in each unit as necessary.

#### **ZONAL INSPECTIONS AND WALKAROUND CHECKS**

In contrast to servicing and lubrication tasks, zonal inspections and walkaround checks of aircraft structures do not fall within the realm of RCM task definition. Walkaround checks are intended to spot accidental damage and fluid leaks and hence might be viewed as combination on-condition and failure-finding inspections. In fact, they do include a few specific on-condition tasks, such as a check of brake wear indicators. However, damage can occur at any time and is unrelated to any definable level of failure resistance. As a result, there is no basis for defining an explicit potential-failure stage or a predictable interval between a potential failure and a functional failure. Similarly, a check for leaks is not based on the failure characteristics of a particular item, but rather is intended to spot any unforeseen exceptions in failure behavior.

Zonal inspections are even less specific. They are not directed at any particular failure mode, but are merely a survey of the general conditions within a given zone, or area, of the equipment. Zonal inspections include a check of all the system assemblies and connecting lines in each zone for security (loose parts), obvious signs of damage or leaks, and normal wear and tear as a result of other maintenance activities. In the powerplant this inspection includes looking into the engine tailpipe and inlet, opening the cowlings and examining all the engine-mounted accessories, and so on. Such inspections play an important role in structural maintenance, since they also include a general inspection of the internal structural areas that can be seen with all installations in place. Thus they complement, but are not a substitute for, the program of detailed on-condition inspections developed for structurally significant items.

Although zonal-installation inspections do not meet the applicability criteria for any of the four basic tasks, their cost is such a small part of the total cost of scheduled maintenance that they are economically justified if they result in the discovery of even a few potential failures. For this reason any RCM program is supplemented by a separate program of scheduled zonal inspections.

**EXHIBIT 3-9** A breakdown of the total maintenance workload of 18.8 manhours per flight hour on the United Airlines fleet of Boeing 747's. Data are for January–November 1975 and do not include man-hours expended to accomplish modifications. (United Airlines)

| location of<br>work performed                  | scheduled<br>work  | corrective work        |                     | total manhours<br>per flight hour |
|------------------------------------------------|--------------------|------------------------|---------------------|-----------------------------------|
|                                                |                    | flight-crew<br>reports | mechanic<br>reports |                                   |
| ON THE AIRPLANE                                |                    |                        |                     |                                   |
| At stations                                    |                    |                        |                     |                                   |
| Below A-check level                            | — <sup>1</sup>     | 2.1                    | 0.2                 | 2.3                               |
| At A-check level                               | 0.2 <sup>2</sup>   | — <sup>3</sup>         | 0.2                 | 0.4                               |
|                                                | 0.2                | 2.1                    | 0.4                 | 2.7                               |
| At main maintenance base                       |                    |                        |                     |                                   |
| Phase check (combination of B<br>and C checks) | 0.7 <sup>2</sup>   | —                      | 0.7                 | 1.4                               |
| D check (heavy structural<br>inspection)       | 0.8 <sup>2,4</sup> | —                      | 0.8                 | 1.6                               |
|                                                | 1.5                | —                      | 1.5                 | 3.0                               |
| OFF THE AIRPLANE                               |                    |                        |                     |                                   |
| At main maintenance base                       |                    |                        |                     |                                   |
| Repair of failed engines                       | —                  | 2.3 <sup>5</sup>       | 6.9                 | 9.2                               |
| Repair of other failed items                   | —                  | 3.9                    | — <sup>1</sup>      | 3.9                               |
|                                                | —                  | 6.2                    | 6.9                 | 13.1                              |
|                                                | 1.7                | 8.3                    | 8.8                 | 18.8                              |

1 Workload was not significant.

2 Workload at checks was prorated, with one-half assigned to scheduled inspections and servicing and one-half assigned to corrective work.

3 A-check figures were adjusted to include only scheduled-maintenance work and the corrective work it generated. Corrective work resulting from flight-crew reports is aggregated with other below-A-check work.

4 The D check figure is not typical. During the reporting period there were two sample D checks for age-exploration purposes. A longer reporting period would lead to a smaller D check number.

5 The corrective engine work was prorated, with one-quarter assigned to pilot reports and the remainder assigned to mechanic findings.

#### **THE TOTAL MAINTENANCE WORKLOAD**

The total maintenance workload required to support complex equipment consists of all the work performed as scheduled maintenance, plus the corrective-maintenance work required to repair failed units. Exhibit 3.9 illustrates the ratio of these two aspects of maintenance for an aircraft supported by a scheduled-maintenance program that is essentially the same as an RCM program. The scheduled tasks comprised somewhat less than 10 percent of the total manhours spent on maintenance, yet these tasks ensured realization of all the reliability of which the equipment was capable. Additional scheduled work would have increased costs, but it would not have improved reliability.

Approximately 75 percent of the corrective work was done at the major maintenance base as a result of the line-maintenance practice of replacing failed units with serviceable ones. About half the corrective work was done on engines. The only way the corrective workload can be reduced is by design changes that improve the inherent reliability of the items that are failing. Such changes are usually directed at dominant failure modes in items whose failure has safety or major economic consequences. In this case the engine failures do have serious economic consequences, and this engine is still undergoing intensive development.

The absolute size of the scheduled workload for this aircraft will not change very much from its 1975 value, but the corrective workload will decrease substantially as product improvement overcomes those problems which require high manhour expenditures. Consequently the relative proportions of the workload components may change in the next several years. At some time in the future both components may increase again as a result of conditions that do not occur until much later ages.

### **3·7 PRODUCT IMPROVEMENT AS PREVENTIVE MAINTENANCE**

Over the years aircraft manufacturers have incorporated a number of design features that have increased the inherent capability of the equipment for reliable operation. In most cases these practices are intended not to prevent failure, but to reduce its consequences to the cost of corrective maintenance. Thus most systems items are designed with a high degree of redundancy to ensure that if one unit fails, the necessary function will still be available. On the same principle, structures are designed with multiple load paths so that they are damage-tolerant. Protective devices may also consist of entirely separate components, as in the case of emergency equipment—fire extinguishers, automatically released oxygen equipment in passenger aircraft, and ejection seats in single-engine military aircraft.

inherent reliability  
characteristics  
methods of coping with the  
failure process

Another common practice is failure substitution. This may be the substitution of a minor functional failure to preempt a major one, as in the use of automatic shutoff devices. Or it may be a feature included to permit easy identification of a potential failure; for example, the outer skin of an aircraft may be designed to crack before the structural member beneath it fails, so that there is evidence of an imminent failure that can be detected by visual inspection. Inspection features such as bore-scope ports in engines also facilitate the detection of potential failures that would otherwise be difficult to check for.

All these features are important from the standpoint of preventive maintenance, since they determine both the feasibility of certain tasks and the failure consequences by which task effectiveness is measured. On a short-term basis, however, any scheduled-maintenance program must be built around the reliability characteristics of the equipment as it exists. In the case of new equipment, therefore, it is important to bear in mind a basic conflict between certain design goals and reliability goals. This problem is nowhere more apparent than in modern aircraft, where the requirement for lightness and compactness is in direct opposition to the strength and bulk that is necessary for failure resistance. A further difficulty is posed by the rush to new technology, since this means that the designer is often working with new components and even new materials whose reliability has not been proved by experience.


There are several pitfalls here. Designing for lightness, for example, correspondingly reduces the initial margin between resistance and stress. Even with familiar materials, the actual strength of a material may be less than its nominal strength, or the rate at which its failure resistance declines may be greater than expected. With unfamiliar materials and processes the likelihood is increased in both these areas. The design goal of compactness may lead to the same results and to other problems as well. In a more compact area an item that functioned well in a different environment may be exposed to higher temperatures or to vibration from neighboring components. Such items are also likely to be more difficult to reach for inspection or replacement.

Where reliability problems are inherent in the design itself, there are three ways of coping with the failure process:

- ▶ Increasing the initial resistance to failure
- ▶ Reducing the rate at which failure resistance decreases
- ▶ Reducing the stress to which the item is exposed

All three of these effects are shown in Exhibit 3.10.

Reliability improvement in each of these areas can take any number of forms. In some instances the solution may be a modification in operating procedures. For example, the use of more reverse thrust and less braking to slow an airplane after it has landed will reduce the stress on



**EXHIBIT 3-10** Methods of coping with the failure process. An item may be redesigned to increase its initial failure resistance, to reduce the rate at which failure resistance decays, or both. At the same time, various strategies may be employed to reduce the stress to which the item is exposed. Any or all of these procedures will improve reliability by moving the point of functional failure farther into the future, and thus increasing the mean time between failures.

the brakes (although it increases the cumulative stress on the reverser). Since this procedure will also increase the life of the tires, it has several implications for maintenance. In general, however, when unsatisfactory reliability characteristics result in exposure to critical failures or excessive operational or maintenance costs, the only effective form of prevention is redesign—either to alleviate the problem or to mitigate its consequences.

When a critical-failure mode is involved, and no form of scheduled maintenance can be found that will effectively control it, product improvement is mandatory. Otherwise the desirability of redesign depends on an assessment of the costs involved on both sides. Since this information is ordinarily not available until after the equipment has been in service for some time, items that may ultimately be redesigned on the basis of actual operating costs are often assigned to no scheduled maintenance in a prior-to-service program.

## developing the initial program

**AN INITIAL** scheduled-maintenance program must be developed for new equipment long before it enters service. While it might be possible to obtain a small mountain of test data on every part, assembly, and subsystem, the information about their actual reliability comes only from operating experience. Thus the problem in basing a maintenance program on reliability characteristics might appear to be a lack of the very information that is needed. In reality the problem is not the lack of information; rather, it is knowing what information is necessary in order to make decisions.

The RCM solution to this problem is a structured decision process based, not on an attempt to estimate the reliability of each part, but on the consequences of functional failures for the equipment itself. The decision process thus proceeds from the top down, first to identify those items whose failure is significant at the equipment level and then to determine what scheduled maintenance can do for each of these items. At each step of the analysis the decision is governed by the nature of the failure consequences. This focus establishes the priority of maintenance activity and also permits us to define the effectiveness of proposed maintenance tasks in terms of the results they must accomplish. Once this determination has been made, we are in a position to examine each of the four possible forms of preventive maintenance to see which tasks, if any, are both applicable and effective for the item under consideration.

tasks is facilitated by a decision-diagram technique which employs an ordered set of priorities—in the case of both failure consequences and task selection—with the questions at each level worded to define the information required for that decision. In many cases the answer will be obvious from engineering expertise, the manufacturer's test data, and previous experience with similar items. However, in developing a prior-to-service maintenance program a strategy is required for decision making when the appropriate information is not available. Thus the decision logic also provides for default answers to meet this situation. For an item subject to critical failures, the default path leads ultimately to redesign. Where the consequences of failure are economic, the default decision may be to do nothing (no scheduled maintenance) until operating experience provides the information to justify some other choice.

The result of RCM analysis is a scheduled-maintenance program that includes all scheduled tasks necessary to ensure safety and operating economy, but only those tasks that will do so. Where there is no basis for determining whether a particular task will prove applicable and effective, the default strategy provides the most conservative answer, and as the maintenance program evolves, these initial decisions are systematically modified on the basis of actual operating data. This process continues throughout the service life of the equipment, so that the decision structure provides for an optimal program in terms of the information available at any time. In this chapter we will examine the

decision process as it relates to commercial aircraft. However, the decision logic itself is general and applies to any complex equipment that requires a maintenance support program designed to realize maximum operating reliability at the lowest cost.

#### 4.1 THE NATURE OF SIGNIFICANT ITEMS

identifying significant items  
structurally significant items  
functionally significant items

A transport plane consists of a vast number of parts and components, all of which have specific functions. All these items can be expected to fail at one time or another, but some of the failures have more serious consequences than others. Certain kinds of failures are a threat to safety, and others have a direct effect on operating capability. However, there are tens of thousands of items whose failure has no immediate impact on the equipment as a whole. The failures are simply corrected soon after they occur, and the only consequence is the cost of repair. These items have no significance from the standpoint of preventive maintenance in the sense that their consequences are tolerable. It is less expensive to leave them in service until they fail than it would be to prevent the failures. Thus the initial decision for these tens of thousands of items is no scheduled maintenance.

The information on which to base this decision ordinarily comes from the manufacturer, who has had to face the problem of failures during the design and development of the equipment. In order to qualify the aircraft for airworthiness, the manufacturer will have conducted a failure modes and effects analysis (FMEA) for all the major assemblies, subsystems, and systems to demonstrate how the equipment will perform when various items fail. In addition, the purchasing airlines will have knowledge of operating experience with similar items in the past, as well as knowledge of the failure consequences in the particular operating context in which the equipment is to be used.

The failures that are of concern are those which have serious consequences. Thus an RCM program directs tasks at a relatively small number of items—those systems, subsystems, and assemblies whose functional failure would be significant at the equipment level, either immediately or downstream in the event of a hidden failure.

##### IDENTIFYING SIGNIFICANT ITEMS

The first step in the development of a scheduled-maintenance program is a quick, approximate, but conservative identification of a set of *significant items*:

---

A significant item is one whose failure could affect operating safety or have major economic consequences.

---

The definition of "major economic consequences" will vary from one operating organization to another, but in most cases it includes any functional failure that has a direct effect on operational capability or involves a failure mode with unusually high repair costs.


So far we have used the term *item* in a very general sense to refer to some component of the equipment. An item can, in fact, be of any size; the entire aircraft might be viewed as an item, as might any one of its parts. However, the larger and more complex the item, the more unwieldy the set of failure possibilities that must be taken into account. To reduce the problem of analysis to manageable size, it is customary to partition the equipment into three major divisions—systems, powerplant, and structure—each of which involves different areas of engineering expertise. Each division is then partitioned in descending order of complexity, with successively fewer failure possibilities at each level.

The chore now is to sort through the functions and failure possibilities of the various components and eliminate all the obviously non-significant items. To ensure that borderline cases and items for which information is lacking will always receive further study, any items eliminated at this stage must be demonstrated to be nonsignificant. Items may be classified as nonsignificant because their functions are unrelated to operating capability or because they are replicated, so that a functional failure would not affect operating capability. Many items can be eliminated because their failures can be repaired quickly and therefore involve no operational consequences. Other items may be ruled out later because they are not candidates for on-condition or safe-life tasks and hence cannot benefit from scheduled maintenance (there is usually no information on the applicability of rework tasks at this time). At this stage, however, all the items that might benefit from scheduled maintenance must be listed for further study.

During the process of classifying items as significant or nonsignificant certain items will be identified that have hidden functions. All these items will require scheduled maintenance regardless of their significance. Although the loss of a hidden function has no direct effect on safety or operating capability, an undetected failure exposes the equipment to the risk of a multiple failure which might have serious consequences. Hence hidden-function items are subjected to the same intensive analysis as significant items.

Note that all items will in fact be included by this procedure, since the partitioning process itself has the following properties:


- ▶ Any item containing a significant item is itself significant.
- ▶ Any nonsignificant item is contained in a higher-level significant item.
- ▶ Any lower-level item contained in a nonsignificant item is itself nonsignificant.



**EXHIBIT 4-1** Partitioning an aircraft for preliminary identification of significant items. The equipment is first partitioned to show all items in descending order of complexity. Those items whose failure clearly has no significant consequences at the equipment level are then pruned from the tree, leaving the set of items on which maintenance studies must be conducted. Each significant item will include as failure modes all the failure possibilities it contains.

The objective, however, is to find the most convenient *level* of each system or assembly to classify as significant. The level must be low enough to ensure that no important failure possibilities are overlooked, but high enough for the loss of function to have an impact on the equipment itself, since the consequences of a functional failure are significant only at the equipment level—that is, for the aircraft as a whole.

Once the optimum level of item has been selected for study in each case, we can prune the "tree" back to a set of several hundred potentially significant items with the assurance that any failure possibilities they include at lower levels will be taken into account as failure modes. As an example, consider the engine described in Section 3.1, in which failure of one or more individual tie bolts in a set of 24 was defined as a potential failure. Although this might be viewed as a functional failure of the tie bolt, the failure of a single bolt does not affect engine performance enough to be evident to the operating crew; consequently the tie bolt is not a significant item. It does, however, have a hidden function, and if enough tie bolts failed, the resulting multiple failure



would indeed become evident. The inspection task selected to avoid such a multiple failure would still be the one described in Exhibit 3.2 — a check for broken tie bolts. However, viewed from the engine level this is an on-condition task, whereas at the parts level it would be considered a failure-finding task.

In other words, the level of item selected as significant is important only as a frame of reference. Whether we look up at a multiple failure or down at a failure mode, an analysis of all the failure possibilities will ultimately lead to exactly the same preventive task. The chief advantage of the partitioning process is that it allows us to focus intensive study on just a few hundred items instead of many thousands. In an aircraft these items will include some of the parts and assemblies, some subsystems, each of the systems, and each of the major divisions themselves. The parts selected as significant are usually those in which a critical failure mode originates. The structure division represents a special case, since the significant items are specific regions that require scheduled maintenance, rather than whole structural assemblies.

### STRUCTURALLY SIGNIFICANT ITEMS

The significant items in each of the major divisions of an aircraft have certain common characteristics which relate to their maintenance requirements. For example, the aircraft structure is a relatively static assemblage of single-celled elements, and except for items such as control surfaces, landing gear, or doors, the only structural movement is deflection under applied loads. However, the structure is subjected to a great many such loads in the course of its operating life. As we saw in Chapter 2, single-celled parts of a mechanism frequently exhibit wear-out characteristics. This is true of metallic structural elements, which are subject to metal fatigue—that is, to a reduction in failure resistance with increasing age.

Another physical process that can lead to the age-related failure of structural elements is corrosion, although the effects of corrosion are much less predictable than those of fatigue. Even minor pitting seriously reduces both static strength and fatigue life, since the loss of load-carrying material correspondingly increases the stress on the rest of the element. Accidental damage has a similar effect in preventing structural components from realizing their inherent fatigue resistance. Thus, although the aircraft structure is designed for a very long fatigue life, it is subject not only to age-related failure in general, but to physical processes that compound the decline in failure resistance with age.

The failure of a major structural assembly which causes the loss of some basic structural function—such as enabling aerodynamic lifting forces to balance the weight of the airplane or providing flight-control surfaces for maneuvering capability—clearly has safety consequences. Moreover, any failures short of a critical failure—failures that do not result in a loss of function to the aircraft—will usually not be evident to the operating crew. The primary consideration in identifying significant structural members, therefore, is the effect that failure of a member has on the residual strength of the remaining assembly, although consideration is also given to susceptibility to corrosion and accidental damage.

The generic term *structurally significant item* (SSI) is used to denote each specific structural region that requires scheduled maintenance to guard against the fracture of a significant member. This region may be defined as a site that includes a number of structural elements, it may be defined as the significant member itself, or it may be a particular region on the member that is the best indicator of its condition. Often such items are the points at which different structural elements are joined; for example, the wing-to-fuselage joint is always listed as a structurally significant item. Most aircraft structure is maintained by on-condition inspections of the regions identified as structurally significant items. These inspections are designed to identify and repair corrosion, fatigue, and other damage at the earliest possible stage, since the replacement of a failed structural element is both difficult and expensive.

#### FUNCTIONALLY SIGNIFICANT ITEMS

Unlike structural items, most systems are equipped with instrumentation to monitor the performance both of the system as a whole and of individual assemblies within it. As a result, the occurrence of any functional failure in a system is usually evident to the operating crew. Moreover, most systems are designed to be highly redundant, so that the failure of one unit often has no effect on operational capability. Unless a second unit fails, the aircraft is dispatched as usual, and corrective maintenance is simply deferred to a convenient time and location. Thus, although the system as a whole is a *functionally significant item* (FSI), the units that comprise it would be classified as nonsignificant, since their individual failures have no consequences at the equipment level.

Systems items differ in two other ways from structural items. Most systems components are themselves multicelled, or complex; hence their overall reliability shows little or no deterioration with age. Certain metal parts in mechanical systems are subject to fatigue and corrosion, but these are rarely responsible for a dominant failure mode. To meet space and weight requirements, systems components are usually designed with a narrow margin between initial failure resistance and stress. Since they are therefore subject to more frequent failure, the system is usually also designed to facilitate the replacement of failed units. A further distinction between systems and structural items is that certain systems items, such as electrical and electronic components, are characteristically unable to benefit from scheduled maintenance.

Although the powerplant is itself a system, it warrants a category of its own because of its complexity, its high cost, and the critical nature of some of its failure modes. The shutdown of one engine in a multi-engine aircraft has operational, but not safety, consequences. However, the failure of turbine or compressor disks—or any other failures that generate projectiles, cause fires, or leave the engine so that it cannot be shut down—can clearly affect safety. These failure modes are always given careful attention in a maintenance program.

The powerplant can be viewed as a functionally significant item in itself, but the failure characteristics of each of its modules, or major subassemblies, are often quite different from those of the engine as a whole. For example, the collective probabilities of all powerplant failures have little relation to operating age, whereas single important parts may be subject to directly age-related failures. Thus scheduled-maintenance tasks in the powerplant program may include safe-life limits for some items and scheduled rework for others. In as many instances as possible, however, on-condition inspections are employed, both to avoid the consequences of functional failures and to reduce the costs associated with scheduled removals. The powerplant is unique from a maintenance standpoint in that it is designed to permit extensive inspection capability on the aircraft, it can be replaced in a fairly

short time (although unscheduled replacements have operational consequences), and it is subject to extensive shop inspections as well.

In the case of new engines there may be some failure modes that cannot be effectively controlled except by redesign. The occurrence of an unanticipated type of failure in any engine prompts an immediate response on the part of maintenance. The failure consequences are quickly assessed and the engine is examined to determine the cause of the failure. Next, some method is usually devised for inspecting the rest of the engines in service (or the suspect group of engines) for early signs of the same kind of failure. These inspections forestall further failures while the part is being redesigned. The alternative, if the failure is critical and no preventive task can be found, is grounding the fleet until the problem can be solved.

Because items within the powerplant are exposed to many different forms of deterioration, including all those that affect the structure and the various systems, they have no common failure characteristic. Unlike systems items, however, all engine failures have operational consequences and some failure modes have safety consequences. For this reason significant items in the powerplant are identified primarily on the basis of their failure effects. The very complexity of the powerplant results in one further characteristic. Engines are subject to so many failure possibilities that operating data accumulate rapidly, especially with use on multiengine commercial aircraft. This rapid feedback, along with the high cost of corrective maintenance on engines, favors the initial selection of intensive on-condition inspections for powerplant items, since the applicability of age-limit tasks can be investigated before the point at which age-related failures would have any major economic impact.

## 4.2 THE RCM DECISION PROCESS

evaluation of failure  
consequences  
evaluation of proposed tasks

The partitioning procedure gives us a conservative first approximation of the items that might benefit from scheduled maintenance. Each of these items must now be examined in detail to determine whether its failure consequences actually qualify it as significant — and if so, whether the item can in fact benefit from scheduled maintenance. Even when the significance of an item is confirmed, there may be no form of preventive maintenance that is applicable and effective. Such items cannot be eliminated from consideration, however, without a full analysis.

### EVALUATION OF FAILURE CONSEQUENCES

The consequences of a functional failure depend on both the nature of the function and the nature of the failure. Hence it is necessary to begin the analysis with an accurate list of all the functions demanded of an

item and a clear definition of the conditions that constitute a functional failure in each case. It is also necessary to know the failure modes involved in order to determine the possible effects of each failure. Once this information has been assembled for every item to be examined, we are in a position to evaluate the actual consequences of failure.

As a result of the partitioning process certain items will have been identified that have hidden functions—that is, their failure will not necessarily be evident to the operating crew. The first matter to be ascertained in all cases, however, is whether we will know when a failure has occurred. The following question is necessary, therefore, to ensure that all hidden functions are accounted for:

---

Is the occurrence of a failure evident to the operating crew during the performance of normal duties?

---

This question must be asked, not for each item, but for each *function* of the item. The loss of an item's basic function may be evident, but in many cases the item will have secondary or other characteristic functions whose failure will not be evident to the operating crew.

Recall from our discussion in Chapter 2 that any functional failure which has a direct effect on operational capability—including critical failures—will always be evident to the operating crew. If the effects of a failure are not observable, the loss of function has no immediate impact. But by the same token, there is no assurance that the failure will be reported and corrected. Thus if the answer to this first question is no for any function, scheduled maintenance is required for that item. The purpose of the task is not necessarily to prevent failures of the hidden function, but to prevent exposure of the equipment to a multiple failure involving that item.


In the case of a failure that is evident to the operating crew, the consequences might be immediate; we therefore need to know how serious they are likely to be:

---

Does the failure cause a loss of function or secondary damage that could have a direct adverse effect on operating safety?

---

This question must be asked for each functional failure and for each *failure mode*. Modern design practices ensure that transport aircraft are exposed to very few critical losses of function. However, certain failure modes, especially in engines, do cause secondary damage that poses a safety hazard. Therefore a yes answer to either aspect of this question means that preventive maintenance is mandatory and can be considered effective only if it prevents all occurrences of this type of failure.



**EXHIBIT 4-2** Decision diagram to identify significant items and hidden functions on the basis of failure consequences. Failures that affect safety or operating capability have an immediate impact, since the aircraft cannot be dispatched until they have been corrected. The impact of nonoperational failures and hidden failures is delayed in the sense that correction can be deferred to a convenient time and location.

If the answer to the safety question is no, our next concern is with economic consequences:

---

**Does the failure have a direct adverse effect on operational capability?**

---

The consequences in this case include an immediate interruption of operations, reduced capability if the airplane continues in service, or

the delay or cancellation of subsequent flights to make unscheduled repairs—all of which involve an economic loss beyond the cost of the repairs. In this case, although scheduled maintenance is not required for safety reasons, it may be desirable on economic grounds. Thus if the answer to this question is yes, any applicable preventive tasks must be investigated for cost effectiveness.

If the failure has no direct effect on operational capability, the economic consequences include only the cost of repair. However, certain functional failures may be far more expensive to repair than to prevent, especially in the case of a failure mode that causes extensive damage to surrounding items. Although scheduled maintenance is more likely to prove cost-effective when operational capability is a factor, there are certain failure modes for which it is often desirable to investigate the economic benefits of a preventive task.

The relationship of these three questions and the decision outcomes in each case are illustrated in Exhibit 4.2. This simple decision-diagram approach provides us with the following basic information about each failure possibility:

- ▶ We know whether the failure will be evident, and therefore reported for correction.
- ▶ We know whether its consequences include a possible safety hazard for the equipment or its occupants.
- ▶ We know whether its consequences have a direct effect on operational capability.
- ▶ We know the objective of preventive maintenance in each case, and hence the criterion for evaluating task effectiveness.

With this information we are now in a position to evaluate the maintenance possibilities for each item.

#### **EVALUATING THE PROPOSED MAINTENANCE TASKS**

The next phase of RCM analysis involves a systematic study of each failure mode to determine whether one of the four basic maintenance tasks will satisfy both the criteria for applicability and the specific conditions for effectiveness. Since there is a clear order of preference for the first three preventive tasks, we can again use a decision-diagram approach, as shown in Exhibit 4.3.

The first task to be considered for each anticipated failure mode of the item being studied is an on-condition inspection:

---


**Is an on-condition task to detect potential failures both applicable and effective?**

---

If the answer is yes, an on-condition inspection task is put into the program for that failure mode. If we obtain yes answers for all the failure modes of an item, the analysis of that item is complete.

The applicability of an on-condition task can be determined by engineering specialists who are familiar with the design characteristics of the item, the materials used in it, and the inspection technology available. Thus this information will be on hand before the equipment goes into service. At the time an initial maintenance program is developed, however, there may not be enough information to determine whether the task will be effective. In this case we assume that it will be effective and establish the initial inspection intervals according to the

**EXHIBIT 4-3** Decision diagram to evaluate proposed scheduled-maintenance tasks. If none of the three directly preventive tasks meets the criteria for applicability and effectiveness, an item whose failures are evident cannot be considered to benefit from scheduled maintenance. If the item has a hidden function, the default action is a scheduled failure-finding task.



seriousness of the failure consequences. Any applicable inspection task can be made effective in terms of failure prevention if the intervals are short enough, and if operating experience later shows that it is not cost-effective, the task will be deleted from the program at the next review.

If an on-condition task is not applicable for certain failure modes, the next choice is a scheduled rework task:

---

**Is a rework task to reduce the failure rate both applicable and effective?**

---

In this case the question of applicability as well as effectiveness requires an analysis of operating data. Thus, unless the age-reliability characteristics of the item are known from prior experience with a similar item exposed to a similar operating environment, the assumption in an initial program is that an item will *not* benefit from scheduled rework. In the absence of information, the answer to this question is no, and we wait for the necessary information to become available after the equipment goes into service.

A no answer to the rework question brings us to the question of a scheduled discard task:

---

**Is a discard task to avoid failures or reduce the failure rate both applicable and effective?**

---

In an initial maintenance program the only items scheduled for discard will be those for which the manufacturer has specified safe-life limits. The tasks associated with those items are put into the program, but in nearly all other cases the answer at this stage will be no.


### **4.3 USE OF THE RCM DECISION DIAGRAM**

The small decision diagram in Exhibit 4.3 provides the essential mechanism for deciding which, if any, of the preventive-maintenance tasks are both applicable and effective for a particular item. To use this diagram, however, it is necessary to know the failure consequences that determine effectiveness in each case and also dictate the default action to be taken at each decision level.


the full RCM decision diagram  
use of the four consequence  
branches  
the role of the default strategy

#### **THE COMBINED DECISION DIAGRAM**

Exhibit 4.4, which brings together the decision questions in Exhibits 4.2 and 4.3, can be used to develop an RCM program either for new equipment or for equipment which is already in service. As we will see in Chapter 5, it can also be used to modify the initial program as new



**EXHIBIT 4-4** The RCM decision diagram. These questions must be asked for each type of functional failure listed for the item. The first three questions determine the consequences of that failure, and hence the objective of preventive tasks. (F. S. Nowlan and H. F. Heap)



information becomes available. The chapters in Part Two discuss the application of RCM analysis to each of the three major divisions of the aircraft—systems, powerplant, and structures. For the time being, however, let us see how the failure consequences influence the process of task selection.

Consider an item which is subject to a critical failure. The answer to question 1 is yes, since any failure that has a direct adverse effect on operating safety will be evident to the operating crew. (This answer refers, of course, only to a loss of the particular function under consideration.) The answer to question 2 is also yes, since the failure has been stated as critical. All subsequent questions about this failure possibility therefore fall in the safety branch of the diagram. This has two important implications for scheduled maintenance:

- ▶ Scheduled maintenance is *required* if an applicable preventive task can be found.
- ▶ A task can be considered effective only if it reduces the risk of critical failure to an acceptable level.

In the case of transport aircraft the risk must be at a level of extreme improbability to be acceptable, but in the general case an acceptable level does exist. For example, single-engine aircraft are utilized for various civilian and military applications.

Each failure mode that might result in this failure is now examined to determine which of the proposed preventive tasks will accomplish the necessary objective. If an on-condition task is applicable for some failure mode, it can usually be made effective by assigning conservatively short inspection intervals (a yes answer to question 4). If there are failure modes for which on-condition inspection is not applicable, the question of scheduled rework is considered. However, in an initial program the failure data necessary to determine the applicability of such a task are rarely available, and no operating organization can afford the number of critical failures required to provide this information. Thus in the case of a critical-failure mode the answer to question 5 is no.

This brings us to the question of scheduled discard of the item or part in which the critical failure originates—that is, to a safe-life limit. In determining initial program requirements engineering advice may indicate that such a task is applicable. Its effectiveness cannot be evaluated, however, unless a safe-life limit has been established by developmental testing under simulated operating conditions. If a safe-life limit has been established, scheduled discard at this limit is required; if a life limit has not been established for this item, the answer to question 6 is no.

When some failure mode cannot be adequately controlled by any one of the preceding tasks, we have one further recourse:

---

**7 Is a combination of preventive tasks both applicable and effective?**

---

There are occasional circumstances in which a combination of two or more preventive tasks will reduce the risk of critical failure to an acceptable level. In a single-engine aircraft, for example, any and all applicable tasks might be employed to reduce the likelihood of engine failure to the lowest level possible. In most instances, however, this is a stop-gap measure, pending redesign of the vulnerable part. If no combination of tasks can be found that will effectively avoid critical failures in the interim, it may be necessary to restrict operation of the equipment or even to remove it from service.

To return to the top of the decision diagram, suppose the failure of an item has no safety consequences (a no answer to question 2), but it does have operational consequences (a yes answer to question 3). In this event we are concerned only with the *economic* consequences of a functional failure:

- ▶ Scheduled maintenance is *desirable* if its cost is less than the combined costs of operational consequences and repair for those failures it prevents.
- ▶ A task can be considered effective only if it is cost-effective.

In scheduled airlines operational consequences can usually be measured in terms of the inability to deliver service to passengers in a timely fashion. In other operating contexts the cost of lost operational capability might be measured differently. However, a cost can always be imputed to any operational failure in terms of the *opportunity cost* of being unable to use the equipment as planned.

To determine whether a proposed maintenance task is economically desirable, it is necessary to know the imputed cost assigned to the expected operational consequences. In initial programs this will usually be an arbitrary figure based on the benefits anticipated at the time the equipment was purchased. In addition, it is necessary to have some idea of the likelihood of failure, the cost of the proposed task, and the cost of corrective maintenance if the item is allowed to fail. Generally, if the expected failure rate is low and the operational consequences are not excessive, the decision will be to use no scheduled maintenance. As the total cost of failure increases, preventive maintenance becomes more attractive. In most cases it is possible to make a decision without a formal economic-tradeoff study. (Later in the chapter we will examine

a procedure for determining whether an economic-tradeoff study is likely to be worthwhile.)

Where no applicable and cost-effective maintenance task can be found, we must either accept the operational consequences (no scheduled maintenance) or redesign the item to reduce the frequency of failures. This decision ordinarily depends on the seriousness of the operational consequences. If they represent a major economic loss, the default decision is redesign.

If the failure of an item has no operational consequences, the question of task effectiveness is evaluated in direct economic terms:

- ▶ Scheduled maintenance is *desirable* if its cost is less than the cost of repair for those failures it prevents.
- ▶ A task can be considered effective only if it is cost-effective.

Task effectiveness in this case is a simple tradeoff between the cost of prevention and the cost of cure. If both costs are of the same order of magnitude, the decision goes to no scheduled maintenance. The reason for this is that any preventive-maintenance task may disturb the steady-state conditions of the mechanism, and this risk should not be introduced without good cause. Thus a preventive task will be scheduled only where the cost of correcting failed items far outweighs the cost of preventing failures.

Note that many of the items designated for no scheduled maintenance through this decision process might well have been identified at the outset as those which cannot benefit from scheduled maintenance. This branch of the decision diagram, however, permits us to evaluate borderline items which might have benefited from a scheduled task if an applicable one could be found.

In the case of hidden-function items task effectiveness involves two criteria:

- ▶ Scheduled maintenance is *required* to avoid exposure to a possible multiple failure.
- ▶ A task can be considered effective only if it ensures adequate availability of the hidden function.

Some hidden functions are sufficiently important that their availability is protected by periodic checks by the operating crew—that is, they are made evident by defining the normal duties of the crew to cover them. In all other cases, however, scheduled inspections are necessary. Since hidden failures can have no direct effect on safety or operational capability, we can allow such items to fail, but we cannot afford the possible consequences of *undetected* failures. Thus, in the absence of any directly preventive task that is applicable and effective, a specific failure-finding task must always be assigned.

#### **THE ROLE OF THE DEFAULT STRATEGY**

The information to be channeled into RCM decisions requires analysis under two different sets of conditions. One is the development of an initial maintenance program on the basis of limited information. The other is modification of these initial requirements as information becomes available from operating experience. As information accumulates, it becomes increasingly easier to make robust decisions. In developing a prior-to-service program, however, there are many areas in which there is insufficient information for a clearcut yes-or-no answer or the study group is unable to reach a consensus. To provide for decision making under these circumstances it is necessary to have a backup default strategy which dictates the course of action in such cases.

The default strategy summarized in Exhibit 4.5 shows for each of the decision questions which answer must be chosen in case of uncertainty. In each case the default answer is based on protection of the equipment against serious consequences. For example, in the process of identifying significant items, if it can be demonstrated that the failure of an item has no effect on safety or operating capability, the item can be classified as nonsignificant and does not warrant further study to see if it can benefit from scheduled maintenance. If there is any doubt, however, it must be classified as significant and cannot be dismissed without further analysis. Similarly, if it is not certain that a loss of function will be evident to the operating crew, it is treated as hidden unless a failure mode involves critical secondary damage.

This default approach can conceivably lead to more preventive maintenance than is necessary. Some tasks will be included as protection against hazards that do not exist, and others may be scheduled far too frequently. The means of eliminating such excessive costs is provided by the age-exploration studies which begin as soon as the equipment goes into service. Through this process the information needed to refine the initial program (and make major revisions where necessary) is gathered systematically for evaluation. We will examine the techniques of age exploration and the nature of the information it provides in the next chapter.

Since an analysis of age-reliability characteristics requires failure data that will not become available until some time after the equipment has been in service, the default strategy will result in a no answer to nearly all questions concerning the applicability and effectiveness of scheduled rework and discard tasks. Consequently, any initial RCM program will consist essentially of on-condition tasks, a few safe-life discard tasks, and failure-finding tasks for hidden-function items, in addition to the usual servicing and lubrication tasks. Scheduled rework or economic-life discard tasks may be added at some later stage, after their applicability and effectiveness can be evaluated, but they rarely appear in an initial program.

**EXHIBIT 4-5** The default answers to be used in developing an initial scheduled-maintenance program in the absence of data from actual operating experience.

| decision question                                                                                                          | default answer to be used<br>in case of uncertainty                                      |
|----------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------|
| <b>IDENTIFICATION OF SIGNIFICANT ITEMS</b>                                                                                 |                                                                                          |
| Is the item clearly nonsignificant?                                                                                        | No: classify item as significant.                                                        |
| <b>EVALUATION OF FAILURE CONSEQUENCES</b>                                                                                  |                                                                                          |
| Is the occurrence of a failure evident to the operating crew during performance of normal duties?                          | No (except for critical secondary damage): classify function as hidden.                  |
| Does the failure cause a loss of function or secondary damage that could have a direct adverse effect on operating safety? | Yes: classify consequences as critical.                                                  |
| Does the failure have a direct adverse effect on operational capability?                                                   | Yes: classify consequences as operational.                                               |
| <b>EVALUATION OF PROPOSED TASKS</b>                                                                                        |                                                                                          |
| Is an on-condition task to detect potential failures applicable?                                                           | Yes: include on-condition task in program.                                               |
| If an on-condition task is applicable, is it effective?                                                                    | Yes: assign inspection intervals short enough to make task effective.                    |
| Is a rework task to reduce the failure rate applicable?                                                                    | No (unless there are real and applicable data): assign item to no scheduled maintenance. |
| If a rework task is applicable, is it effective?                                                                           | No (unless there are real and applicable data): assign item to no scheduled maintenance. |
| Is a discard task to avoid failures or reduce the failure rate applicable?                                                 | No (except for safe-life items): assign item to no scheduled maintenance.                |
| If a discard task is applicable, is it effective?                                                                          | No (except for safe-life items): assign item to no scheduled maintenance.                |

| stage at which question can be answered |                                     | possible adverse consequences<br>of default decision                                   | default consequences<br>eliminated with subsequent<br>operating information |
|-----------------------------------------|-------------------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| initial program<br>(with default)       | ongoing program<br>(operating data) |                                                                                        |                                                                             |
| X                                       | X                                   | Unnecessary analysis                                                                   | No                                                                          |
| X                                       | X                                   | Unnecessary inspections that are<br>not cost-effective                                 | Yes                                                                         |
| X                                       | X                                   | Unnecessary redesign or scheduled<br>maintenance that is not cost-<br>effective        | No for redesign;<br>yes for scheduled<br>maintenance                        |
| X                                       | X                                   | Scheduled maintenance that is not<br>cost-effective                                    | Yes                                                                         |
| X                                       | X                                   | Scheduled maintenance that is not<br>cost-effective                                    | Yes                                                                         |
| X                                       | X                                   | Scheduled maintenance that is not<br>cost-effective                                    | Yes                                                                         |
| —                                       | X                                   | Delay in exploiting opportunity<br>to reduce costs                                     | Yes                                                                         |
| —                                       | X                                   | Unnecessary redesign (safety) or<br>delay in exploiting opportunity<br>to reduce costs | No for redesign;<br>yes for scheduled<br>maintenance                        |
| X<br>(safe life only)                   | X<br>(economic life)                | Delay in exploiting opportunity<br>to reduce costs                                     | Yes                                                                         |
| X<br>(safe life only)                   | X<br>(economic life)                | Delay in exploiting opportunity<br>to reduce costs                                     | Yes                                                                         |

## 4.4 DETERMINING COST EFFECTIVENESS

criteria for cost effectiveness  
finding a cost-effective interval  
the impact of inherent  
reliability characteristics

Since a moderate amount of information gathering is necessary for calculations of cost effectiveness, it is helpful to know whether the effort is likely to be fruitful. The decision-diagram approach is also useful in this area. Exhibit 4.6 illustrates one method for deciding whether a detailed assessment of an applicable task might be worthwhile.

Up to this point we have not been concerned about failure rate, since it is not a primary measure of consequences. In the case of critical failures it has no bearing; in fact, the sole objective is to avoid any failures on which to base a rate. Where the consequences are economic, however, the total cost depends on the frequency with which these consequences are likely to occur. The first question in evaluating the cost effectiveness of prevention, therefore, concerns the frequency of functional failures:

---

Is the functional-failure rate high?

---

Since it is seldom worthwhile to deal with rare types of noncritical failures, this question rules out items that fail so seldom that the cost of scheduled maintenance would probably be greater than the benefits to be derived from it. The term *high*, of course, is open to interpretation. In airline practice a failure rate greater than 1 per 1,000 hours of flight time is usually considered high, whereas a rate of less than 0.1 per 1,000 hours is usually not considered important. This question is often easier to answer if the failure rate is described in terms of the number of failures per month.

If the failure rate is judged to be high, the next concern is the cost involved. Operational consequences are usually the major cost associated with a high failure rate:

---

Does the failure involve operational consequences?

---

Any failure that prevents continued dispatch of the equipment involves operational consequences. However, the extent of the economic loss depends largely on the intended use of the equipment. In a military context, for example, a much higher cost might be imputed to dispatch of an airplane with restrictions on its operating performance than would be the case in a commercial-airline context. If the failure does have operational consequences, the total cost of failure includes the combined cost of these consequences and the cost of repair.

Even when operational consequences are not involved, it may be advantageous to forestall a particularly expensive failure mode:


---

**Does any failure mode cause unusually high repair or operating costs?**

---

This question must be investigated separately, since such failure modes will usually be responsible for only a small fraction of the total number of failures.

**EXHIBIT 4-6** Decision diagram for evaluating the probable cost effectiveness of a proposed task when scheduled maintenance is not required to protect operating safety or the availability of hidden functions. The purpose of the decision technique is to reduce the number of formal economic-tradeoff studies that must be performed.



A yes answer to either of the preceding two questions means that we need further information:

---

Do real and applicable data show the desirability of the proposed task?

---

It is possible to arrive at a yes answer to this question if there is substantial evidence that this task was cost-effective in the past for this or a similar item. If so, the task can be scheduled without a formal study.

**EXHIBIT 4-7** A pro forma for analyzing the support costs associated with scheduled removals for rework. At least four proposed rework intervals must be examined to determine whether a cost-effective interval does exist.

|                                                                    |            |            |
|--------------------------------------------------------------------|------------|------------|
| item                                                               |            |            |
| annual volume of operation                                         |            |            |
| proposed interval                                                  |            |            |
| Number of failures per year <sup>1</sup>                           | X          |            |
| Average base cost of repairing a failed unit <sup>2</sup>          | <u>\$X</u> |            |
| Annual base cost of repairing failed units                         |            | \$X        |
| Number of failures that have operational consequences <sup>3</sup> | X          |            |
| Average cost of operational consequences after failure             | <u>\$X</u> |            |
| Annual cost of operational consequences                            |            | \$X        |
| Number of scheduled removals per year                              | X          |            |
| Average base cost for a time-expired unit <sup>2</sup>             | <u>\$X</u> |            |
| Annual base cost for time-expired units                            |            | \$X        |
| Number of spare units required to support workload                 | X          |            |
| Cost of unit                                                       | <u>\$X</u> |            |
| Annual cost of spare units required                                |            | <u>\$X</u> |
| Total annual support costs <sup>4</sup>                            |            | <u>\$X</u> |

1 It may be desirable to study a specific expensive failure mode separately.

2 Includes cost of removing and installing unit at line station and of transporting it to and from the maintenance base

3 The number of failures that have operational consequences may be different from the total number of failures, since not every failure will have such consequences.

4 If the change in volume of work at the maintenance base results in changes in facility requirements, the annual cost of such changes should be included in the support costs.

Otherwise the question of economic tradeoff must be evaluated for each of the applicable maintenance tasks:

---

**Does an economic-tradeoff study justify the task?**

---

An economic-tradeoff study involves several steps:

- ▶ An estimate of the incremental effect of the task on the failure rate of the item for several different task intervals
- ▶ A translation of the reduced failure rate into cost reductions
- ▶ An estimate of the cost of performing the proposed task for each of the intervals considered
- ▶ Determination of the interval, if one exists, at which the cost-benefit ratio is the most favorable

Exhibit 4.7 shows a pro forma for evaluating the cost effectiveness of a scheduled rework task. As we saw in Chapter 3, the cost factors for on-condition tasks and scheduled rework tasks are quite different. Scheduled removals increase both the total shop volume and the number of spare units required to replace the units that are undergoing rework. Consequently, unless the frequency of a very expensive failure is materially reduced by an age limit, the total cost of this task will usually outweigh its economic benefits.

In contrast, the total number of potential failures removed as a result of on-condition inspections is not appreciably greater than it would be if each unit were allowed to fail. Moreover, the cost of repairing potential failures is usually less than the cost of repair after a functional failure. As a result, on-condition inspection tasks, when they are applicable, are relatively easy to justify.

The important role of cost effectiveness in RCM decision making helps to clarify the nature of inherent reliability characteristics. The inherent reliability of an item is not the length of time it will survive with no failures; rather, it is the level of reliability the item will exhibit when it is protected by preventive maintenance and adequate servicing and lubrication. The degree of reliability that can be achieved, however, depends on certain characteristics that are a direct result of the design details of the equipment and the manufacturing processes that produced it. These characteristics determine both the need for preventive maintenance and the effectiveness with which it can be provided. Thus from a maintenance standpoint inherent reliability characteristics are decision factors such as those listed in Exhibit 4.8. Note that the answer to each of the questions in Exhibit 4.4 requires a knowledge of at least one of these characteristics.

| inherent reliability characteristic                           | impact on decision making                                                                                                                                                                |
|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Failure consequences                                          | Determines significance of items for scheduled maintenance; establishes definition of task effectiveness; determines default strategy when no applicable and effective task can be found |
| Visibility of functional failure to operating crew            | Determines need for failure-finding task to ensure that failure is detected                                                                                                              |
| Ability to measure reduced resistance to failure              | Determines applicability of on-condition tasks                                                                                                                                           |
| Rate at which failure resistance decreases with operating age | Determines interval for on-condition tasks                                                                                                                                               |
| Age-reliability relationship                                  | Determines applicability of rework and discard tasks                                                                                                                                     |
| Cost of corrective maintenance                                | Helps establish task effectiveness, except for critical failures                                                                                                                         |
| Cost of preventive maintenance                                | Helps establish task effectiveness, except for critical failures                                                                                                                         |
| Need for safe-life limits to prevent critical failures        | Determines applicability and interval of safe-life discard tasks                                                                                                                         |
| Need for servicing and lubrication                            | Determines applicability and interval of servicing and lubrication tasks                                                                                                                 |


**EXHIBIT 4-8** Examples of inherent reliability characteristics and their impact on decision making. Each decision question in Exhibit 4.4 requires a knowledge of at least one of these characteristics. In the absence of this knowledge, a default answer must be employed in developing an initial scheduled-maintenance program.

The test of cost effectiveness means that an RCM program will not include some tasks that might reduce the likelihood of noncritical failures. However, when a failure has economic consequences the inclusion of a task that is not cost-effective would merely transfer these consequences from one cost category to another; it would not reduce them. Thus the cost factors on both sides must be considered inherent reliability characteristics, since they dictate the level of reliability that is

feasible for an existing design. Within this framework, RCM analysis ensures all the operating reliability of which the equipment is capable. Moreover, it results in a selection of only those tasks which will accomplish this objective; hence it also provides the required maintenance protection at minimum cost.

Certain of the inherent reliability characteristics of new equipment are unknown at the time a prior-to-service maintenance program is developed. Consequently the initial program is somewhat more expensive than later refinements of it will be (although it is still a minimum-cost program in terms of the information available at the time). This situation is inevitable because of the default decisions necessary to protect the equipment in the absence of full information. It is not too serious a matter, however, because of the relatively slow rate at which fleets of new equipment grow. For example, the Boeing 727 fleet shown in Exhibit 4.9 took six years to reach its maximum size of 150 aircraft. Although the full fleet finally flew more than 400,000 total hours a year, the 20 planes in service by the end of the first year had flown a total of only 34,300 hours. Thus the maintenance costs stemming from these initial default decisions have little overall economic impact and will be materially reduced with the information available by the time the fleet reaches full size.

**EXHIBIT 4-9** Examples of fleet growth in a commercial airline. Each purchasing airline has a maximum rate at which it can accept new airplanes, determined by training and staffing requirements. The rate at which new equipment can enter service is highest for large airlines. (United Airlines)



## 4.5 AGE EXPLORATION


determination of  
potential-failure age  
opportunity sampling

One of the most important aspects of an initial RCM program is *age exploration* to determine the applicability of certain tasks and the most effective intervals for others. In the case of aircraft this process starts with the manufacturer's certification test flights, during which some of the most frequent types of failures will be identified. If some of these failures have major consequences, product improvement will be initiated before any equipment is delivered to the purchaser. The information obtained during the certification period, however, identifies only those items that have failed — presumably those with a high probability of failure. The entire certification program for a new commercial transport plane requires a total of only 1,500 to 2,000 flight hours accumulated on the five or six planes assigned to the program. The flying time for any one test plane is usually no more than 400 or 500 hours. In contrast, once a plane is put into service, it may fly 300 or more hours a month. At this point we can begin to acquire information on the additional reliability characteristics of the equipment.

As we saw in Section 3.1, the applicability of an on-condition task depends on the ability to measure reduced failure resistance. Its effectiveness, however, depends on the interval between inspections. The same holds true for failure-finding tasks assigned to hidden-function items. For this reason all such tasks are assigned conservatively short intervals in an initial program, and all items whose failure could have safety or major economic consequences are carefully monitored by frequent sample inspections to determine the exact effect of operating age on their condition. The simple metal part illustrated in Exhibit 3.1, for example, would initially be monitored at the intervals shown in Exhibit 4.10 to determine the exact point to be defined as a potential failure, the age at which inspections should start, and the most effective interval between inspections.

Because on-condition inspections play a large role in the maintenance programs for turbine engines, some interesting practices have evolved to reduce the cost of obtaining this information. When an initial program is being developed, experience with earlier types of engines will suggest many parts that might benefit from on-condition tasks, as well as some that might benefit from scheduled rework. Consequently the sample inspections required for age exploration make up a large part of the initial maintenance program for any powerplant.

Some of these inspections can be performed while the engine is installed, but others can be performed only at a major maintenance base after a certain amount of disassembly of the engine. The "on-the-wing"



**EXHIBIT 4-10** Initial sampling intervals assigned in an age-exploration program to determine the rate at which failure resistance declines. Reduced resistance is not detectable until a visible crack appears; thereafter the rate of crack propagation is monitored to determine the exact point to be defined as a potential failure, the point at which it is necessary to begin on-condition inspections, and the most effective inspection interval to ensure that all failing units will be identified at the potential-failure stage.

inspections are handled by an initial requirement for early inspection of the item on all engines. However, if inspection of the first few engines to reach this limit discloses no unsatisfactory conditions, the limit for the remaining engines is extended. Thus very few engines are actually inspected at any fixed time limit until the point at which it becomes desirable to stop extending the limit.

For those parts that require engine disassembly for inspection, the practice is to define an age limit at which inspection information is considered to be of value. The initial operating age of a part might be limited, for example, to 1,500 hours without inspection, and the threshold age for valid sampling information might be set at 500 hours. This was done for the General Electric CF6-6 engine in the Douglas DC-10. In that case the FAA required inspection of two sets of parts (equivalent to two engines) to justify an increase in the 1,500-hour limit. The initial maintenance program stated that sampling information could be obtained either from one part aged 500 to 1,000 hours and a second part aged 1,000 to 1,500 hours, or else from two parts that were both

aged 1,000 to 1,500 hours. The two sets of part-inspection reports could be based on the inspection of parts in any number of engines.

The reason for this flexibility in scheduling is to take advantage of *opportunity samples*, samples taken from engines that have failed and have been sent back to the main base for repair. Any undamaged parts from these engines can be used to meet the sampling requirements. This procedure makes it unnecessary to schedule engine removals for disassembly solely for the purpose of inspecting parts. Such *forced removals* are necessary only when the required volume of sampling cannot be obtained from opportunity samples. Because new types of engines usually have high failure rates that create abundant opportunity samples, it is possible to make a careful evaluation of the condition of each part before any engines on the aircraft actually age to the initial maximum limit.

On-condition inspections also play the primary role in the maintenance programs for structures. However, unlike powerplants, structure does not provide opportunity samples. The structure is designed as an integral unit, and corrective maintenance on any structural item removes the entire airplane from service. Moreover, because the failure of any major structural assembly is critical, all parts of the structure are designed to survive to very high ages. In the case of structure, therefore, the inspection program itself is the only vehicle for age exploration, and the inspection samples consist of individual airplanes, rather than samples of parts from different airplanes. The initial inspection interval for each structurally significant item is set at only a fraction of the age at which evidence of deterioration is expected to appear, not only to find and correct any conditions that may reduce the anticipated design life, but also to identify the age at which reduced failure resistance first becomes evident.

Whereas powerplant items are continually interchanged and replaced as part of the normal repair cycle, structural members are repaired, but are rarely replaced with new parts. Consequently the age of most parts of a given structure is the same as the total age of the airplane. This makes it possible to concentrate age-exploration activities on the highest total-time airplanes. The first few airplanes to reach the initial limit established for major structural inspections are designated as *inspection samples*. All inspection findings for these airplanes are carefully documented, so that any changes in their condition with age can be identified before younger airplanes reach this age. If there are no signs of deterioration, the starting intervals in the initial program will usually be increased for the remaining airplanes in the fleet.

Age exploration of systems items is conducted on still another basis. Systems items are generally characterized by low reliability; hence they provide abundant opportunity samples. However, because systems failures are rarely critical and so many systems items cannot benefit from

scheduled maintenance, extensive inspection of opportunity samples is usually not justified by the value of the information obtained. In this case the frequency of failures is likely to have greater economic impact than the consequences of individual failures. Thus for systems items age exploration is based primarily on the monitoring and analysis of failure data to determine the cost effectiveness of proposed tasks.

In the following chapter we will examine the many other aspects of the age-exploration process.

## 4•6 PACKAGING THE MAINTENANCE TASKS

Once each maintenance task in the prior-to-service program has been assigned an appropriate initial interval, either for the purpose of age exploration or on the basis of conservative judgment, the RCM tasks are combined with other scheduled tasks—the servicing and lubrication tasks specified by the manufacturer and the scheduled zonal-installation inspections. All the tasks with similar intervals are then grouped into a number of *maintenance packages*, each with its own interval. The principle is the same as that spelled out in new-car warranties, which specify a certain group of servicing and inspection tasks to be performed every 1,000 miles, another to be performed every 5,000 miles, and so on. For commercial aircraft these intervals range from between-flight checks at every station to major inspections at eight- to ten-year intervals at a maintenance base.

maintenance packages  
distribution of maintenance  
workload

This grouping results in slightly more frequent performance of some tasks than is strictly necessary, but the additional cost is justified by the increase in maintenance efficiency. Those tasks that are most expensive, both in actual cost and in terms of down time for out-of-service equipment, tend to shape the overall package. Thus if one task must be performed every 1,000 miles and another can be done easily at the same time, they will both be scheduled for that interval. If the second task is required, say, every 2,500 miles, it will be scheduled every other time the first task is done, and so on.

Airlines frequently give each of the major scheduled-maintenance packages an alphabetic designation; hence they are commonly known as *letter checks*. An A check might be performed every 125 hours of flight time, a B check every 900 hours, and so on. Exhibit 4.11 shows the sequence of letter checks as they would occur for an airplane over an operating period of 3,600 hours. The content of a given letter check will not necessarily be the same every time it is performed, since some tasks will come up only at every second or third occurrence of that check. However, the fact that the more extensive packages occur at longer intervals means that as the level of work increases, fewer stations need to be equipped to handle it.

| age<br>(flight hours) | work package            | age<br>(flight hours) | work package            |
|-----------------------|-------------------------|-----------------------|-------------------------|
| 125                   | #1 A Check              | 1,925                 | #17 A Check             |
| 250                   | #2 A Check              | 2,050                 | #18 A Check             |
| 375                   | #3 A Check              | 2,175                 | #19 A Check             |
| 500                   | #4 A Check              | 2,300                 | #20 A Check             |
| 625                   | #5 A Check              | 2,425                 | #21 A Check             |
| 750                   | #6 A Check              | 2,550                 | #22 A Check             |
| 875                   | #7 A Check              | 2,675                 | #23 A Check             |
| 900                   | #1 B Check <sup>1</sup> | 2,700                 | #3 B Check <sup>3</sup> |
| 1,025                 | #9 A Check              | 2,852                 | #25 A Check             |
| 1,150                 | #10 A Check             | 2,950                 | #26 A Check             |
| 1,275                 | #11 A Check             | 3,075                 | #27 A Check             |
| 1,400                 | #12 A Check             | 3,200                 | #28 A Check             |
| 1,525                 | #13 A Check             | 3,325                 | #29 A Check             |
| 1,650                 | #14 A Check             | 3,450                 | #30 A Check             |
| 1,775                 | #15 A Check             | 3,575                 | #31 A Check             |
| 1,800                 | #2 B Check <sup>2</sup> | 3,600                 | #1 C Check <sup>4</sup> |

1 Includes #8 A check.

2 Includes #16 A check.

3 Includes #24 A check.

4 Includes #4 B check and #32 A check.

**EXHIBIT 4-11** A sample schedule of maintenance packages. Each work package includes all the scheduled tasks to be performed at that interval. The A check includes all tasks scheduled at 125-hour intervals; the B check consists of all tasks scheduled at 900-hour intervals, as well as the A check that would otherwise be performed at that interval; and the C check, scheduled for 3,600-hour intervals, includes all the tasks scheduled for that interval, along with both the A and B checks that would ordinarily take place at that time. The A checks are performed at any of several line-maintenance stations. Planes are routed to a few large maintenance stations for B checks, and C checks are performed at the maintenance base.

for every stop at a line maintenance station, and a #2 service check might be scheduled for every stopover of more than five hours (unless a higher-level package is being performed), and so on.

In addition to the letter checks, which package the expensive or time-consuming tasks, there are a number of smaller service packages. For example, a #1 service check might include those tasks scheduled

The entire scheduled-maintenance program, packaged for actual implementation, must be completed and approved before any new air-

craft can enter service. Up to this point RCM analysis has provided us with a set of tasks based on those reliability characteristics that can be determined from a knowledge of the equipment and the operating context. Once the equipment enters service a whole new set of information will come to light, and from this point on the maintenance program will evolve on the basis of data from actual operating experience. This process will continue throughout the service life of the equipment, so that at every stage maintenance decisions are based, not on an estimate of what the reliability is likely to be, but on the specific reliability characteristics that can be determined at that time.

## evolution of the rcm program

**IN THE** preceding chapters we have examined the framework of RCM analysis and the decision process that leads to the selection of tasks for an initial maintenance program. After the equipment enters service information becomes available about its actual interaction with the operating environment. This information almost certainly contains some surprises—unanticipated types of failures, unexpected failure consequences, unusually high failure rates, or even an absence of anticipated failures. Because the volume of operation is small at first, information is gained at that time about the failures that are likely to occur soonest and with the greatest frequency. As operating time accumulates, the less frequent types of failures are discovered, as well as those that tend to occur at higher operating ages. All this information is used for continuing evolution of the ongoing maintenance program.

Any complex equipment is a failure generator, and failure events will occur throughout its whole operating life. The response to these events depends on the failure consequences. If an unanticipated failure has serious implications for safety, the first occurrence sets in motion an immediate cycle of maintenance and design changes. In other cases waiting until several failures have occurred allows a better assessment of their frequency to determine the economic benefits of preventive tasks, or possibly redesign. Very often waiting until enough failures have occurred to permit an evaluation of age-reliability relationships provides the information necessary to modify the initial maintenance decisions.

Evolution of the scheduled-maintenance program does not consist solely of reactions to unanticipated failures. The information that becomes available—including the absence of failures—is also used for systematic evaluation of all tasks in the initial program. On the basis of actual data, the initial conservative intervals for on-condition inspec-

tions can be adjusted and the applicability of scheduled rework and economic-life tasks can be investigated. Actual operations will frequently confirm the a priori assessments of failure consequences, but occasionally the consequences will be found to be more serious or less serious than anticipated, or a failure thought to be evident to the operating crew is not, and vice versa. The process by which all this information is obtained is called *age exploration*, both because the amount of information is a direct function of the age of the equipment in service and because some of this information relates to the ages of the items themselves.

## 5·1 THE USES OF OPERATING DATA

It is important to recognize, both in planning a prior-to-service program and at the age-exploration stage, that a fleet of equipment does not materialize overnight. In commercial aviation new planes are delivered to an airline at a rate of one to four a month, and as we saw in Exhibit 4.9, the number of aircraft in service and the associated volume of operations builds up slowly. This allows us to concentrate first on the most frequent failures (since those that occur early will continue to occur early after either delivery or repair) or on those failures with the most serious consequences. As the volume of operations increases, the less frequent failures come to light and can be dealt with later. In a military environment, where operating experience does not accumulate as rapidly, this latter information may be obtained by deliberate heavy use of the first few pieces of equipment—the fleet-leader concept—although the small size of the sample data presents a serious drawback.

the role of age exploration  
evolution of the initial program

The reliability information obtained from actual operating experience is quite varied. Although the failure rate plays a role early in operation in pinpointing design problems and evaluating task effectiveness, an age-exploration program is organized to provide the following kinds of information:

- ▶ The types of failures the equipment is actually exposed to, as well as their frequencies
- ▶ The consequences of each failure, ranging from direct safety hazards through serious operational consequences, high repair costs, long out-of-service times for repair, to a deferred need to correct inexpensive functional failures
- ▶ Confirmation that functional failures classified as evident to the operating crew are in fact evident during normal performance of duties
- ▶ Identification of the circumstances of failure to determine whether the failure occurred during normal operation or was due to some external factor, such as bird strike

**EXHIBIT 5-1** Summary of the uses of new information in the continuing evolution of the scheduled-maintenance program. After the equipment enters service age exploration and the evaluation of actual operating data continue throughout its entire service life.

| refinements of initial maintenance program                              |                                                                                                               |                                                                                                           |
|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| inspection tasks                                                        | results of age exploration                                                                                    |                                                                                                           |
|                                                                         | proposed age-limit tasks                                                                                      | items assigned to no scheduled maintenance                                                                |
| Confirm that reduction in failure resistance is visible.                | Determine age-reliability relationship to confirm that conditional probability of failure increases with age. | Monitor and evaluate operational data to see whether some applicable and effective task can be developed. |
| Determine rate of reduction in failure resistance.                      | If failures are age-related, determine whether a cost-effective age limit exists.                             |                                                                                                           |
| Confirm or modify defined potential-failure condition.                  |                                                                                                               |                                                                                                           |
| Adjust inspection interval and age for first inspection, if applicable. | If a cost-effective interval can be found, add task to program.                                               |                                                                                                           |

- ▶ Confirmation that on-condition inspections are really measuring the reduction in resistance to a particular failure mode
- ▶ The actual rates of reduction in failure resistance, to determine optimum inspection intervals
- ▶ The mechanism involved in certain failure modes, to identify new forms of on-condition inspection and parts that require design improvement
- ▶ Identification of tasks assigned as default actions in the initial program which do not prove applicable and effective
- ▶ Identification of maintenance packages that are generating few trouble reports
- ▶ Identification of items that are not generating trouble reports
- ▶ The ages at which failures occur, so that the applicability of scheduled rework and discard tasks can be determined by actuarial analysis

Exhibit 5.1 summarizes the uses of all this information in refining and

| major revisions to initial maintenance program                                                                                                           |                                                            |                                                                          |
|----------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------|--------------------------------------------------------------------------|
| results of technological change                                                                                                                          |                                                            |                                                                          |
| unanticipated failure modes or consequences                                                                                                              | new or redesigned item                                     | changes in inspection technology                                         |
| Develop on-condition tasks to prevent critical failures and to prevent or reduce frequency of expensive failures at low ages.                            | Conduct RCM analysis of item when it first enters service. | Evaluate applicability and effectiveness of new on-condition techniques. |
| Develop design changes necessary for permanent correction of problems.                                                                                   | Refine maintenance requirements through age exploration.   |                                                                          |
| Develop failure-finding tasks for hidden functions not identified in initial program.                                                                    |                                                            |                                                                          |
| Develop on-condition or other tasks to control critical or expensive failures at high ages, where product improvement may not be economically justified. |                                                            |                                                                          |

revising the initial maintenance program. The refinements are useful, but their overall economic impact is usually quite small. The major revisions are associated with unanticipated failures, design modifications, and the exploitation of new inspection technology; in this area far greater economies are realized.


## 5.2 REACTING TO SERIOUS FAILURES

the preventive-maintenance/  
redesign cycle  
the improvable failure rate  
prediction of reliability  
improvement

After new equipment enters service it may experience unanticipated types of failures and failure consequences. The most serious of these are usually in the powerplant and the basic structure. Although such failures can occur at any point in the life of the equipment, they are most likely to occur early in operation. The first failure may have such serious implications for operating safety or economics that all operating organizations and the manufacturer react immediately. Thus there is a structured pattern of events associated with unanticipated failures which results in a characteristic cycle of reliability improvement.

Suppose the unforeseen failure is a critical engine failure. As an immediate step, engineering investigations are undertaken to determine whether some on-condition inspection or other preventive task will be effective. This preventive measure may result in a substantial increase in maintenance costs. With a new engine a large number of engine removals, dictated either by the discovery of potential failures or by scheduled removal of all units, will also make it difficult to provide replacement engines. The next step is action to redesign the parts in which the failure mode originates. When the new parts are available, all the engines in service must then be modified to incorporate the change. Not all design changes are successful, and it may take several attempts over a period of two or three years to correct the problem. Once the problem has been eliminated, the scheduled-maintenance tasks instituted to control this type of failure are no longer necessary and can be discontinued.

Exhibit 5.2 illustrates this cycle. A year after this engine entered service two critical failures occurred during a three-month period. Both failures were found to be caused by notch wear in the third-stage turbine blades. Since this failure mode was also found to be detectable at the potential-failure stage, a line-maintenance on-condition inspection was specified to check for loose turbine blades. Frequent inspection intervals resulted in a large number of engine removals for this condition, but removal of these potential failures prevented any further functional failures. The turbine blade was redesigned, and halfway through the following year modification of the existing engines was started to incorporate the new "low-swirl" blades. The on-condition inspections were continued, but as more and more modified engines



**EXHIBIT 5-2** The pattern of events associated with an unanticipated critical failure mode in the Pratt & Whitney JT4 engine. The data represent all engine removals for this failure mode, the first two as functional failures and the rest as potential failures found by an on-condition task developed after the first failure events. These premature removals prevented all further functional failures, and as modified engines entered service, the number of potential failures also decreased. When no further potential failures were found, the on-condition task was deleted from the program. (United Airlines)

entered service, the number of premature removals (potential failures) dropped. Finally, about three years after the first two failures, the on-condition inspections were discontinued.

In new equipment the scheduled-maintenance tasks generated in response to early critical failures are nearly always on-condition inspections. Age-limit tasks are not likely to be feasible, since there are no data for actuarial analysis, and in the case of early failures, taking some fraction of the age at failure as a safe-life limit could easily be ineffective. Moreover, a short safe-life limit might effectively preclude continued operation of the equipment, since it would be difficult to provide the labor and spare parts needed for such intensive maintenance. The definition of an applicable on-condition task, however, may require great ingenuity. The failure mode must be determined, and a specific part that shows physical evidence of reduced failure resistance must be identified. Then some means of inspecting the part while it is still installed must be devised.

Under these circumstances both the potential-failure point and the inspection interval will be established on a very conservative basis. As soon as the on-condition task is implemented, all the equipment in

service is inspected. This first inspection of the fleet often leads to a large number of removals for the newly defined potential failure. The rate of removal after this first inspection will be much lower, of course. It may be low enough to justify increasing the initial conservative inspection interval, but the inspections themselves will be continued until experience has demonstrated that the problem no longer exists.

The cycle for early structural difficulties is similar. Once again, it is necessary to determine the failure mode and devise an on-condition inspection for potential failures. In this case the inspections may be scheduled as often as once every flight cycle or at intervals as long as 2,000 or 3,000 flight cycles. Again, even though the incidence of potential failures turns out to be relatively low after the first fleet inspection, the task itself is continued until the design can be modified.

Serious unanticipated failures do not necessarily occur early in the life of new equipment. At later ages, however, such failures may not lead to design changes. The first response is still the same—the development of new scheduled-maintenance tasks. At this stage the imposition of safe-life limits may be both technically and economically feasible. On-condition tasks may also be applicable, but the inspections can be scheduled to begin at a relatively high age and may have longer intervals. Unless the failure mode is strongly related to age, in which case a life-limit task may be more appropriate, the number of potential failures found by on-condition inspections will be far lower than in relatively new equipment. Depending on the age of the equipment, the cost of redesign may not be warranted, since economic justification depends on the remaining technologically useful life of the equipment.


One further way of coping with failure is to restrict operating procedures to put less stress on a vulnerable component until it can be redesigned. Sometimes the opposite strategy is also useful. When no specific potential-failure condition can be identified, it may be possible to preempt a serious failure by inducing it under other circumstances. In one such case failures of a compressor disk on a tail-mounted turbine engine were occurring at very low ages, and no on-condition inspections were feasible. It was possible to keep the plane in service, however, by requiring the pilot to brake at the end of the runway and apply takeoff thrust with the aircraft stationary. The peak stress on the disk occurred when takeoff thrust was first applied and decreased as the disk warmed up. Thus if the disk did not fail during warmup, it was unlikely to do so during flight. This strategy resulted in several expensive failures, but they were not critical on the ground, whereas the secondary effects of disk failure would have been critical in flight.


A new piece of complex equipment often experiences a high failure rate. Often, too, the majority of these failures result from a small number of failure modes. In the case of aircraft engines the conditional proba-

bilities of such dominant failure modes will frequently increase rapidly with operating age. Exhibit 5.3 shows the results of successive analyses of an engine that entered service in 1964. At that time its initial reliability was poor, the conditional probability of failure was high, and this probability increased rapidly with age. However, the increase was linear and showed no identifiable wearout zone. Within a few months the reliability of this engine was substantially improved by design modifications directed at the dominant failure modes. The initial high failure rate brought the unmodified engines into the shop very frequently, which facilitated fairly rapid incorporation of the modified parts. Consequently the conditional probability of failure continued to drop, and ultimately the reliability of this engine showed no relationship to operating age.

Once the early dominant failure modes in an engine are disposed of, it becomes increasingly difficult to make further improvements. Because of its complexity, the engine will always be subject to many different failure modes, and some may even be dominant. However, the failure probability associated with any given mode is too low to justify further development of the engine. The difference between an item's initial and mature failure rate is its *improvable failure rate*—the

**EXHIBIT 5-3** Results of successive age-reliability analyses of the Pratt & Whitney JT8D engine of the Boeing 727. As engineering improvements gradually overcame dominant failure modes, the conditional-probability curve continued to flatten until it eventually showed no relationship of engine reliability to operating age. (United Airlines)





**EXHIBIT 5-4** Comparison of actual failure rates of the Pratt & Whitney JT8D engine with a forecast made in December 1965. During initial operation the failure rate based on small samples will show large variations in different calendar periods. However, since reliability improvement is characteristically exponential, it is possible to predict the expected reduction in failure rate over a longer calendar period. The temporary variation from the forecast level in this case was the result of a new dominant failure mode which took several years to resolve by redesign. (United Airlines)

portion that will be eliminated by product improvement. If a particular engine has a failure rate of 2 per 1,000 hours when it first enters service and we anticipate that its failure rate will ultimately drop to 0.3, then the improvable failure rate is 1.7.

In many cases the improvable failure rate declines exponentially over calendar time—that is, the percentage of reduction remains constant, although the amount of reduction becomes smaller as the failure rate is reduced. This percentage has been as much as 40 percent a year for engines in a commercial-airline environment. Such a high degree of improvement is possible only when a large number of engines are in service to generate the failure data required both to direct product improvement and to lower its unit cost. The fact that improvement is characteristically exponential enables us to plot reliability growth in new equipment with a fair degree of success. Exhibit 5.4 shows a comparison of actual failure experience with a forecast that was made in

1965. The forecast was reasonably good until 1968, when a new failure mode became dominant. This problem took nearly three years to resolve, after which the failure rate dropped back to the forecast level.

### 5.3 REFINING THE MAINTENANCE PROGRAM


The maintenance tasks added in response to unanticipated failures are only one aspect of the age-exploration process. At the time the initial program is developed certain reliability characteristics are unknown. For example, the ability to measure reduced failure resistance can be determined, but there is no information on the actual rate of reduction as various items age in service. Similarly, the information necessary to evaluate cost effectiveness and age-reliability relationships becomes available only after the equipment has been in service for some time. Once the maintenance program goes into effect, the results of the scheduled tasks provide the basis for adjusting the initial conservative task intervals, and as further operating data become available the default decisions made in the absence of information are gradually eliminated from the program.

adjusting task intervals  
uses of actuarial analysis

#### ADJUSTING TASK INTERVALS

As part of the initial program many items are scheduled for frequent sample inspections to monitor their condition and performance, and other tasks are assigned conservatively short initial intervals. All these tasks are then packaged for implementation. If the first few units to reach this check limit show no unsatisfactory conditions, it is safe to assume that the task interval for the remaining units can be extended. Any equipment that has aged to the present check limit is designated a *time-extension sample*.

In many cases, as we saw in Chapter 4, the required number of samples is provided by opportunity samples, units that are available for inspection because they have failed for some reason related to only one failure mode. In the case of engines, for example, the availability of samples of a particular part depends on the number of shop visits occasioned by failures in the section of the engine containing that part. Since a new type of engine is far more likely to experience failures of components in the hot section than in the cold section, the engine data in Exhibit 5.5 show far more opportunity samples for the exit guide-vane assembly than for the compressor assembly. In both cases, however, opportunity sampling provided a means of inspecting these parts as they aged in service. Since there was no great difference between the age of the highest-time installed part and the age of the highest-time sample inspected, it was possible to extend the check limits for both



**EXHIBIT 5-5** Effectiveness of opportunity sampling of the Pratt & Whitney JT8D engine. Opportunity samples of the exit guide-vane assembly (black) were more abundant than samples of the high-compressor assembly (red), but at every age the highest-time installed unit was only slightly older than the highest-time inspected sample. Thus any unsatisfactory condition detected in the sample would be found before the remaining installed units had reached this age. (United Airlines)

items until the age at which the sample units began to show signs of deterioration.

Task intervals for systems and structural items are ordinarily increased by increasing the interval of the letter-check package in which they have been included. However, if the inspection reports indicate that the interval for some particular task in this package should not be extended, the task must be moved to another package. A task originally assigned to the C-check package, for instance, might be reassigned to the package designated for every second B check. Conversely, there will be tasks whose original intervals now appear far too conservative. In this case the task interval might be increased, say, from C2 to C4 at the same time that the C-check interval itself is being revised upward. The same result can be achieved, of course, by leaving the intervals of all packages fixed and moving all tasks from one package to another.

The management of maintenance packages requires careful planning. First, a schedule is needed for conducting the analysis necessary to support each interval extension. This schedule must allow time for the first few units that have entered service to age to the existing check limit, and also time for the analysis necessary to assess the desirability

of extending the limit. The results of all inspections and corrective work performed on these sample units must be carefully analyzed so that the tasks for which intervals should not be extended can be moved to more compatible packages. Tasks producing marginal results may stay with the original package, but they should be noted for future attention. A hard-time directory is usually maintained to identify tasks for which a maximum interval appears likely. These tasks require closer study than the others, and maintenance planning is facilitated by advance knowledge that they may be moved to a different package in the near future.

#### **USES OF ACTUARIAL ANALYSIS IN AGE EXPLORATION**

Whereas serious unanticipated failures prompt an immediate response, action on infrequent failures or those with no major consequences is usually delayed until enough information has been gathered to make a full assessment of possible maintenance remedies. This is particularly true with regard to rework tasks, since these tasks are applicable only if the conditional-probability curve shows that an item has an identifiable wearout zone. Such curves are the result of an actuarial analysis in which the number of failures during various age intervals are measured in terms of the total exposure of the item (total operating time for all units) and the probability of survival to that age interval.

An actuarial analysis does not require hundreds of failure events. A survival curve can be constructed from the data on 20 functional failures, and if necessary, from a sample of 10. However, since it takes several thousand operating hours to accumulate this many occurrences of a given type of failure, there is sometimes concern about a surge of failures as a result of wearout after a certain age. If all the units in service were the same age this might be the case, but because of the slow buildup of a fleet of airplanes, the ages of the units in service are widely distributed. If the item is very reliable at lower ages, and the first failure does not occur until some time after the fleet has reached full strength, the age distribution of the in-service units at that time will be the same as that of the planes in the fleet. This means that there may be a difference of five years or more between the ages of the oldest unit and the newest one. If the item is not that reliable, there will be even fewer high-time units, since many of the units on the older airplanes will be replacements for units that have already failed.


It is this distribution in the ages of in-service units of an item that makes it feasible to use actuarial analysis as a tool for age exploration. If it is found that there is a sharp increase in the likelihood of failure at higher ages, there is ample time to take preventive steps, since very few units are actually approaching the "cliff" when it is discovered. It follows that attention is concentrated on the failure behavior of the oldest units, so that in the event that there is a wearout zone, a rework task


can be added to the maintenance program long before the other units reach this age.

Exhibit 5.6 shows the results of an actuarial analysis conducted to determine whether complete rework of a turbine engine would be an applicable task. The upper curve shows the total conditional probability for all units removed and sent to the shop for corrective work, and the lower curve shows the conditional probability of functional failures as reported by the operating crew. The distance between these two curves at any age represents the conditional probability of potential failures detected by on-condition inspections. It is functional failures that have safety or operational consequences, and the conditional probability of such failures in this case is constant. Since functional failures are independent of the time since engine installation (last shop visit), operating age is not a factor in the failure rate, and a rework task is therefore not applicable.

The conditional-probability curve that includes potential failures does show an increase with increasing age. However, we do not want to reduce the incidence of potential failures except by redesign, since these inspections for potential failures are clearly effective in reducing the number of functional failures. As it is, each engine can remain in operation until a potential failure is detected, and under these conditions

**EXHIBIT 5.6** Conditional-probability curves for the General Electric CF6-6 engine of the Douglas DC-10. The upper curve shows the total number of premature removals for both functional and potential failures, and the lower curve shows the number of these units removed as functional failures. Although the rate of potential failures increases with operating age, as a result of effective on-condition inspections the functional-failure rate is kept in check and shows no increase with age. (United Airlines)





**EXHIBIT 5-7** Partitioning of a conditional-probability curve to show the number of unverified failures and the number of verified failures resulting from each of three failure modes. Note that the only high infant mortality occurs from failure mode A; this results in an upturn of the curves above it in a layered representation.

there is no increase in the functional-failure rate with age. Thus the on-condition task itself prevents a wearout zone for functional failures and at the same time permits each engine to realize almost all of its useful life.

The relationship of verified and unverified failures can be examined in the same way to determine the effectiveness of troubleshooting methods. This information is of value to those concerned with stocking and allocating replacement units and spare parts, but it is also important in identifying the actual characteristics of verified failures, so that the failure mode can be pinpointed more exactly and a more accurate potential-failure condition can be defined.

Exhibit 5.7 shows the various age-reliability relationships that can be developed for an item subject to several different failure modes. The upper curve shows the conditional probability for all reported failures, and the curve below it shows the conditional probability of verified failures. The distance between these two curves represents the prob-

ability of unscheduled removals of units that are actually serviceable. Thus the first curve represents the apparent reliability of the item and the second curve represents its actual reliability.

To determine how we might improve the reliability of this item we must examine the contribution of each failure mode to the total of verified failures. For example, failure modes *A* and *B* show no increase with increasing age; hence any attempt to reduce the adverse age relationship must be directed at failure mode *C*. There is also a fairly high conditional probability of failure immediately after a shop visit as a result of high infant mortality from failure mode *A*. The high incidence of early failures from this failure mode could be due to a problem in shop procedures. If so, the difficulty might be overcome by changing shop specifications either to improve quality control or to break in a repaired unit before it is returned to service. In the case of aircraft engines, for example, shop procedures in commercial airlines include a test-cell run at the end of the shop process, during which some engines are rejected and sent back for further work. These test-cell rejects do not appear in the failure count, since this count begins only after the engine is installed on the aircraft.

An actuarial analysis such as that in Exhibit 5.7 can direct improvements toward a great many different areas by indicating which factors are actually involved in the failure behavior of the item. An analysis of the Boeing 727 generator, for example, showed that the conditional probability of generator failure did not increase with age until bearing failures started at an age of 2,000 hours. This failure mode usually results in destruction of the generator. Since a new generator costs about \$2,500, as opposed to \$50 for a bearing replacement, a generator rework task during which the bearing was discarded was both applicable and cost-effective at 4,000-hour intervals.

## 5.4 REVISIONS IN MAINTENANCE REQUIREMENTS

new diagnostic techniques  
design changes

The maintenance tasks instituted in response to serious unanticipated failures are usually interim measures, intended to control the problem until it can be resolved by redesign. Two kinds of technological change, however, may lead to revision of the requirements for scheduled maintenance: the development of new diagnostic techniques and modification of the present equipment.

### NEW DIAGNOSTIC TECHNIQUES

Most on-condition inspections are diagnostic techniques, since they measure resistance to failure to identify specific problems. The earliest and simplest technique used for aircraft was visual examination, perhaps aided by a magnifying glass. This visual inspection was extended

by development of the borescope. Numerous other techniques have been developed for detecting cracks in metallic items, such as eddy-current, magnaflux, and zygló inspections. Radiography is also widely employed, not only for detecting cracks, but also to check clearances and changes in configuration without the need to disassemble the item.

A useful diagnostic technique must be able to detect some specific condition that can confidently be defined as a potential failure. It should be sufficiently accurate to identify all units that have reached this condition without including a large number of units for which failure is remote. In other words, such techniques must provide a high power of discrimination. The demand for such discrimination depends in part on the consequences of failure. A technique with low resolving power might be of value for single-engine aircraft if it prevented even a small number of engine failures, despite the fact that it caused numerous unjustified removals. For a multiengine aircraft the same technique would be unnecessary as a safety precaution and undesirable in economic terms.

Certain diagnostic techniques appear to have great potential but will require further development before they can be universally adopted. For example, spectrographic analysis is sometimes used to detect wear in metal parts by measuring the concentration of metallic elements in lubricating oil. In many cases, however, it has been difficult to define a failure condition related to the metal concentrations. Parts have failed without the expected warning, and warnings have not necessarily been associated with imminent failure. Even a change in the brand of oil may necessitate new criteria for interpreting the analysis. Nevertheless, if the failure is one with major consequences, even a low incidence of successful interpretations (and prevented failures) may offset the cost of the inspections that produced no useful information.

Another recent technique is the use of computerized airborne integrated data systems (AIDS), which measure and record the performance characteristics of many items for later study. Some of these characteristics, especially in powerplants, are also monitored by the normal flight instrumentation, but the data are not automatically recorded and integrated with other data. This procedure opens up the possibility of correlating performance trends with the likelihood of failures, or "establishing a signature" for the failure mode. By revealing a previously overlooked indication of reduced resistance to failure, AIDS may make it possible to prevent certain functional failures by on-condition maintenance. The new data systems have in fact assisted in troubleshooting, and they have indicated engine conditions that increase the stress on certain internal parts. However, their success in performing a true (and continuous) on-condition surveillance has so far been limited. Once again, this system may be worthwhile for some organizations if analysis convinces them that the value of its contribution outweighs its costs.

As we have seen, scheduled rework tasks have limited applicability, and discard tasks apply only under rather special circumstances. Major improvements in maintenance effectiveness depend, therefore, on expanded use of diagnostic techniques. The search for additional techniques continues, and the economic desirability of such new developments must be reevaluated from time to time.

#### **DESIGN CHANGES**

The product-improvement process is also a factor in changing maintenance requirements, since design modifications may change the reliability characteristics of items either intentionally or otherwise. Hidden functions may be added or removed, critical-failure modes may be added or removed, dominant failure modes and/or age-reliability characteristics may be altered, and redesign may change the applicability of on-condition tasks.

Whenever an item is substantially modified, its maintenance requirements must be reviewed. It may also be necessary to repeat the age-exploration process for such items, both to find out whether the modifications have achieved their intended purpose and to determine how these modifications affect existing maintenance requirements for the item. Finally, entirely new items are added to most equipment during its service life. Initial requirements must be developed for each of these items, to be modified as necessary when operating data on them become available.

## **5.5 THE PRODUCT-IMPROVEMENT PROCESS**

determining the need for  
product improvement  
determining the desirability  
of product improvement  
information requirements  
the role of product improvement  
in equipment development

In the course of evaluating the maintenance requirements of complex equipment many items will be found that cannot benefit from scheduled maintenance, either because there is no applicable preventive task or because the available forms of prevention cannot provide the level of reliability necessary. Because of the inherent conflict between performance requirements and reliability requirements, the reliability problems identified and corrected during early operations are really a part of the normal development cycle of high-performance equipment.

The degree of reliability that can be achieved by preventive maintenance is limited by the equipment itself. Thus a product may be deemed unsatisfactory for any of the following reasons:

- ▶ Exposure to critical failures
- ▶ Exposure to failures that unduly reduce operational capability
- ▶ Unduly high maintenance costs
- ▶ A demonstrated need to make a hidden function visible

Failures may result from the stress and wear associated with the normal operation of the item, or they may be caused by external factors such as lightning strikes, bird ingestion, corrosive environments, and so on. Product improvement to increase resistance to these external factors may be just as necessary as modifications to withstand the effects of the normal operating environment.

#### **DETERMINING THE NEED FOR PRODUCT IMPROVEMENT**

Product improvement directed toward better reliability may take a number of forms. An item may be modified to prevent critical failures, to eliminate a particularly expensive failure mode, or to reduce its overall failure rate. The equipment, or an item on it, may be modified to facilitate replacement of a failed unit, to make a hidden function visible, to incorporate features that make on-condition inspections feasible, or to add redundant features which alter the consequences of failure.

Product improvement is expensive. It involves the cost of redesign and the manufacture of new parts or whole new items. The operating organization also incurs the direct cost of modifying the existing equipment and perhaps the indirect cost of taking it out of service while such modifications are being incorporated. Further risks are always introduced when the design of high-performance equipment is changed, and there is no assurance that the first attempt at improvement will eliminate or even alleviate the problem at which improvement is directed. For this reason it is important to distinguish between situations in which product improvement is necessary and those in which it is desirable.

The decision diagram in Exhibit 5.8 is helpful in evaluating the necessity or desirability of initiating design changes. In this case the answers to the decision questions are all based on operating experience. As always, the first consideration is safety:

---

**Does the failure cause a loss of function or secondary damage that could have a direct adverse effect on operating safety?**

---


If the answer to this question is yes, the next concern is whether such failures can be controlled at the maintenance level:

---

**Are present preventive measures effectively avoiding such failures?**

---

If the answer is no, then the safety hazard has not been resolved. In this case the only recourse is to remove the equipment from service until the problem can be solved by redesign. Clearly, product improvement is required.



**EXHIBIT 5-8** Decision diagram to determine whether product improvement is required or merely desirable if it is cost-effective. Unless product improvement is required for safety reasons, its cost effectiveness must be assessed (see Exhibit 5.9) to determine whether the improvement is in fact economically desirable.

If the present preventive measures are effectively controlling critical failures, then product improvement is not necessary for safety reasons. However, the problem may seriously restrict operating capability or result in unduly expensive maintenance requirements. It is therefore necessary to investigate the possibility of reducing these costs:

---

**Is product improvement cost-effective?**

---

Here we are concerned solely with economics. As long as the safety hazard has been removed, the only issue now is the cost of the preventive measures employed. By the same token, if the answer to the first question was no—that is, the failure has no direct effect on safety—it

may still have costly operational consequences. Thus a no answer to the safety question brings us directly to the question of cost effectiveness.

#### **DETERMINING THE DESIRABILITY OF PRODUCT IMPROVEMENT**

There is no hard-and-fast rule for determining when product improvement will be cost-effective. The major variables can be identified, but the monetary values assigned in each case depend not only on direct maintenance costs, but on a variety of other shop and operating costs, as well as on the plans for continuing use of the equipment. All these factors must be weighed against the costs of product improvement.

An operating organization is always faced with a larger number of apparently cost-effective improvement projects than are physically or economically feasible. The decision diagram in Exhibit 5.9 is helpful in ranking such projects and determining whether a proposed improvement is likely to produce discernible results within a reasonable length of time.

The first question in this case concerns the anticipated further use of the equipment:

---

**Is the remaining technologically useful life of the equipment high?**

---

Any equipment, no matter how reliable, will eventually be outmoded by new developments. Product improvement is not likely to result in major savings when the equipment is near the end of its technologically useful life, whereas the elimination of excess costs over a span of eight or ten years of continued service might represent a substantial saving.

Some organizations require for budget approval that the costs of product improvement be self-liquidating over a short period—say, two years. This is equivalent to setting the operational horizon of the equipment at two years. Such a policy reduces the number of projects initiated on the basis of projected cost benefits and ensures that only those projects with relatively high payback are approved. Thus if the answer to this first question is no, we can usually conclude that product improvement is not justified. If the economic consequences of failure are very large, it may be more economical to retire the equipment early than to attempt to modify it.


The case for product improvement is obviously strengthened if an item that will remain in service for some time is also experiencing frequent failures:

---

**Is the functional-failure rate high?**

---

**EXHIBIT 5-9** Decision diagram to assess the probable cost effectiveness of product improvement. If a particular improvement appears to be economically desirable, it must be supported by a formal economic-tradeoff study.



If the answer to this question is yes, we must consider the economic consequences of failure:

---

**Does the failure involve major operational consequences?**

---

Even when the failures have no operational consequences, there is another economic factor to be taken into account:

---

**Is the cost of scheduled and/or corrective maintenance high?**

---

Note that this last question may be reached by more than one path. With a no answer to the failure-rate question, scheduled maintenance may be effectively preventing functional failures, but only at great cost. With a no answer to the question of operational consequences, functional failures may not be affecting operating capability, but the failure mode may be one that results in exceedingly high repair costs. Thus a yes answer to either of the two preceding questions brings us to the question of product improvement:

---

**Are there specific costs which might be eliminated by product improvement?**

---

This question concerns both the imputed costs of reduced operational capability and the more tangible costs associated with maintenance activities. Unless these costs are related to a specific design characteristic, however, it is unlikely that the problem will be eliminated by product improvement. Hence a no answer to this question means the economic consequences of this failure will probably have to be borne.

If the problem can be pinned down to a specific cost element, then the economic potential of product improvement is high. But is this effort likely to produce the desired results?

---

**Is there a high probability, with existing technology, that an attempt at product improvement will be successful?**

---

Although a particular improvement might be very desirable economically, it may not be feasible. An improvement directed at one failure mode may unmask another failure mode, requiring several attempts before the problem is solved. If informed technical opinion indicates that the probability of success is low, the proposed improvement is unlikely to be economically worthwhile.

If the improvement under consideration has survived the screening process thus far, it warrants a formal economic-tradeoff study:

---

Does an economic-tradeoff study show an expected cost benefit?

---

The tradeoff study must compare the expected reduction in costs during the remaining useful life of the equipment with the costs of obtaining and incorporating the improved item. The expected benefit is then the projected saving if the first attempt at improvement is successful, multiplied by the probability of success at the first try. Alternatively, it might be considered that the improvement will always be successful, but only a portion of the potential savings will be realized.

There are some situations in which it may be necessary to proceed with an improvement even though it does not result in an actual cost benefit. In this case it is possible to work back through the set of decision questions and determine the values that would have to be ascribed for the project to break even. Also, improvements in the form of increased redundancy can often be justified when redesign of the offending item is not. This type of justification is not necessary, of course, when the in-service reliability characteristics of an item are specified by contractual warranties or when there is a need for improvement for reasons other than cost.

#### **INFORMATION REQUIREMENTS**

No manufacturer has unlimited resources for product improvement. He needs to know which modifications to his product are necessary and which are sufficiently desirable for him to risk the cost of developing them. This information must come from the operating organizations, who are in the best position to determine the consequences and costs of various types of failures measure their frequency, and define the specific conditions that they consider unsatisfactory.

Opinions will differ from one organization to another about the desirability of specific improvements, both because of differences in failure experience and because of differing definitions of a failure. A failure with safety consequences in one operating context may have only operational consequences in another, and operational consequences that are major for one organization may not be significant for another. Similarly, the costs of scheduled and corrective maintenance will vary and will also have different economic impacts, depending on the resources of each organization. Nevertheless, the manufacturer must assess the aggregate experience of the various users and decide which improvements will be of greatest value to the entire group.

With any new type of equipment, therefore, the operating organization must start with the following assumptions:

- ▶ Certain items on the equipment will need improvement.
- ▶ Requests for improvement must be supported by reliability and cost data.
- ▶ Specific information on the failure mode must be provided as a basis for redesign.



Critical failures must be reported by a safety-alert system so that all operating organizations can take immediate action against identified safety hazards. Failure with other operational consequences are reported at short intervals so that the cost effectiveness of product improvement can be assessed as soon as possible. The airline industry imputes high costs to delayed or cancelled flights, and these events are usually reported on a daily basis. In military applications it is important that operating data, especially peacetime exercise data, be examined carefully for its implications for operational readiness.

For items whose failure has no operational consequences, the only justification for product improvement is a substantial reduction in support costs. Many of these items will be ones for which there is no applicable and effective form of preventive maintenance. In this case statistical reliability reports at monthly or quarterly intervals are sufficient to permit an assessment of the desirability of product improvement. The economic benefits of redesign will usually not be as great under these circumstances. In general, the information requirements for product improvement are similar to those for management of the ongoing maintenance program. In one case the information is used to determine necessary or desirable design modifications and in the other it is used to determine necessary or desirable modifications in the maintenance program.

#### **THE ROLE OF PRODUCT IMPROVEMENT IN EQUIPMENT DEVELOPMENT**

The role of the product-improvement process in the development of new equipment is exemplified by the history of a fleet of Boeing 747's. The first planes in this fleet went into operation in 1970 and the last four planes were delivered in 1973. By April 1976 the airline had issued a total of 1,781 change-order authorizations. Of this total, 85 of the design changes were required by regulatory agencies, 801 were the result of altered mission requirements by the airline, and 895 were required by unsatisfactory reliability characteristics. The cumulative number of these change orders over the first six years of operation is shown in Exhibit 5.10. Most of the change orders to meet regulatory requirements were issued in compliance with FAA airworthiness directives. Such directives mandate specific design changes or maintenance requirements to prevent critical failures. The cumulative number of the 41 directives issued (some entailed more than one change) is shown by the second curve in Exhibit 5.10.

**EXHIBIT 5-10** History of change-order authorizations for design improvements in the Boeing 747 (top) and history of FAA airworthiness directives issued over the same time period (bottom). (United Airlines)



The 895 design changes required to improve reliability characteristics did not include those associated with critical failures. They consisted of the following types of product improvement:

- ▶ Those desirable to prevent or reduce the frequency of conditions causing delays, cancellations, or substitutions (495)
- ▶ Those desirable to improve structural fatigue life and reduce the need for frequent inspection and repairs (184)
- ▶ Those desirable to prevent or reduce the frequency of conditions considered to compromise ground or flight safety (214)

All these changes were based on information gathered from actual operations after the equipment went into service. Such information is an essential part of the development cycle in all complex equipment.

## **5·6 RCM PROGRAMS FOR IN-SERVICE EQUIPMENT**

The decision process outlined in Chapter 4 was discussed in terms of new equipment. However, this procedure also extends to the development of an RCM program for equipment that is already in service and is being supported by a scheduled-maintenance program developed on some other basis. In this case there will be much less need for default answers, since considerable information from operating experience is already available. For example, there will be at least some information about the total failure rate of each item, the actual economic consequences of various kinds of failures, what failure modes lead to loss of function, which cause major secondary damage, and which are dominant. Many hidden functions will have been identified, and there may be information on the age-reliability characteristics of many items.

Preparation for the program will still require a review of the design characteristics of the equipment to define a set of significant functions and functional failures. The usual result will be that items currently treated individually can be grouped as a system or subsystem to be considered as one significant item in the new program. A set of proposed maintenance tasks will have to be established which includes all those existing tasks that satisfy the applicability criteria; additional tasks may then be introduced if they also meet these requirements. The tasks would then be analyzed for effectiveness in terms of failure consequences, as with a prior-to-service program.

The new RCM program should be developed with minimal reference to the existing program, and the two programs should not be compared until the proposal for the new one is complete. This is essential

use of available information  
expected benefits

to avoid the influence of past biases and to allow for free exercise of the decision structure. When a comparison is finally made, the new RCM program will generally have the following features:

- ▶ Many systems and subsystems will be classified as significant items.
- ▶ There will be a smaller number of equipment items for which unique scheduled-maintenance tasks are specified.
- ▶ Most systems items will no longer be subject to scheduled rework.
- ▶ Turbine engines and other complex items will be subject to a few specific rework or discard tasks, rather than intensive scheduled overhaul.
- ▶ There will be age-exploration sampling of certain identified parts of the powerplant, which is continued until the parts reach very high ages.
- ▶ There will be increased use of on-condition tasks.
- ▶ There will be some new tasks that are justified by critical-failure modes, operational consequences, or hidden functions.
- ▶ The intervals of higher-level maintenance packages will be greatly increased, whereas intervals of lower-level packages, which consist primarily of servicing tasks and deferrable corrective work, will remain about the same.
- ▶ The overall scheduled-maintenance workload will be reduced.

If the existing program assigns a large number of items to scheduled rework, there may be some concern that eliminating these tasks will result in a substantial increase in the failure rate. This question can be resolved by conducting actuarial analyses of the failure data for these items under the new program, to confirm that the change in maintenance policy has not adversely affected their overall reliability. If these analyses show that rework tasks are both applicable and effective for some items, they can be reinstated.

The new RCM program will not be as labor-intensive as the program it replaces, and this fact will have to be taken into account in adjusting staff requirements at maintenance facilities. It may be necessary to estimate the volume of work that has been eliminated in each maintenance package and make these adjustments when the new program is first implemented. Otherwise the anticipated reductions in manhours and elapsed time for scheduled maintenance will often not be realized.