# Chapter 3: Literature Survey

## 3.1. Introduction

There is a scarcity of publications on the development and use of RCM. The only textbooks worth noting are those of Anderson and Neri (1990), Moubray (1991) and Smith (1993). The only other textbook giving proper coverage to RCM is Coëtzee (1997), while Jones (1995) and Campbell (1995) has overview chapters on the subject. No academic research of importance is being done on RCM as a technique. Academics that do write in the area of the design of maintenance plans do so from a mathematical modelling approach (there are a plethora of them) or has created own techniques - examples are Gits (1992) and Kelly (1997). Even non-academics are active in the creation of own methodologies for maintenance plan design - an example is Jones (1995). The only research of any consequence done to develop the RCM technique, is done by practitioners of the technique - examples are Smith (1993) and Moubray (1991).

At least two of the major textbooks on the RCM methodology Moubray (1991) and Smith (1993) have followed the broad RCM structure proposed by Nowlan and Heap (1978). This same structure will be used for both this and the next chapters to facilitate comparison with the RCM standard works. The main headings that will be used for this purpose are:

- The principle behind RCM: preserve function or preserve equipment?
- The selection of application areas for RCM
- Information Assembly
- Identification of Failure Modes
- Prioritisation of Failure Modes
- Classification of Failure Modes
- Task Selection
- Task Frequencies
- Task Packaging
- Critical assessment of resulting program

## 3.2. A survey of publications and trends in the development of the RCM technique and related techniques

Most of the groundwork for the following discussions will be laid by first exploring and discussing the original RCM work [Nowlan and Heap (1978)]. The work of later authors of works on RCM and related techniques will then be evaluated against these original principles, with a view to identifying any real improvements to the original methodology.

Because RCM is sometimes thought to be complex, a simple overall structure is introduced, which serves to communicate to the reader of this chapter the position of the specific topic in the overall RCM structure. This framework is introduced in figure 3-1.

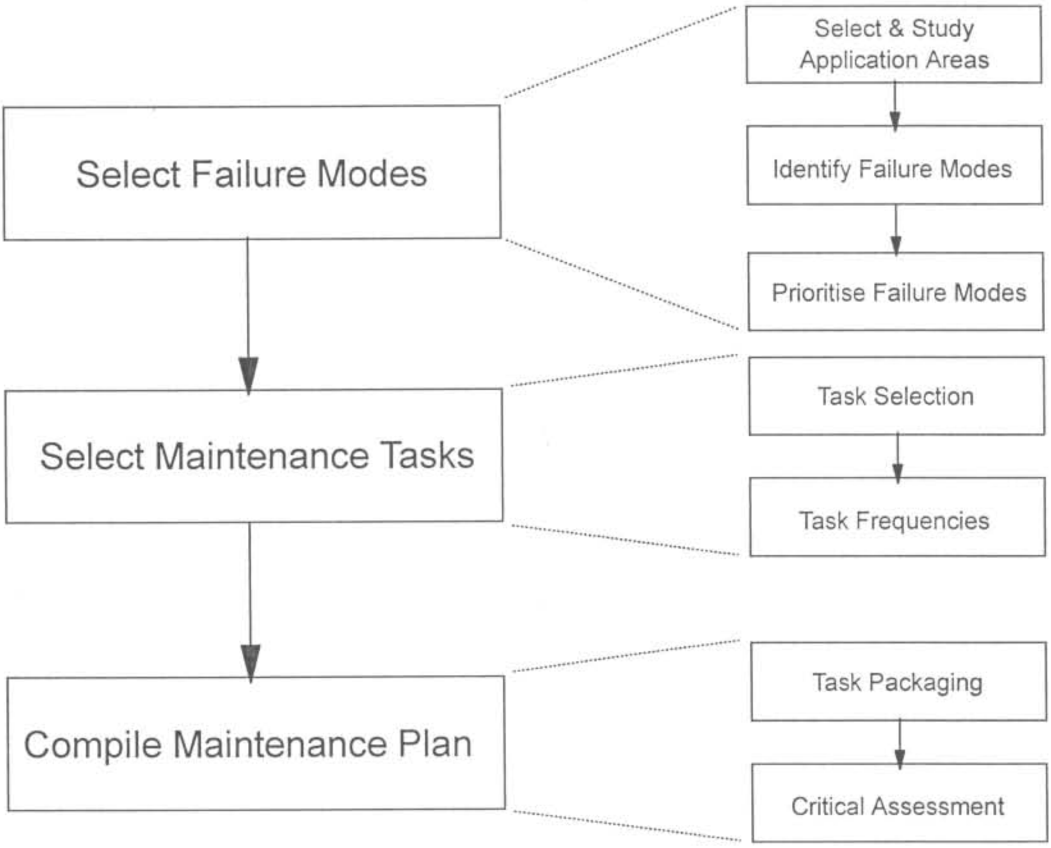

**Figure 3-1: Outline of RCM process**

This framework shows that the total RCM process consists of three sub-processes:

- Select failure modes, which has as objective to identify the failure modes, which are most detrimental to achieving the objectives of the organisation.
- Select maintenance tasks, consisting of finding the best maintenance tasks and task frequencies to deal effectively with the failure modes selected in the first step.
- Compile maintenance plan, where the maintenance tasks found in the previous step are packaged into a practical maintenance programme and the resultant programme critically evaluated for correctness and practicality.

The figure above will be used in the form of the tracking diagram shown in figure 3-2 to assist the reader in following which part of the RCM methodology is being addressed in any particular paragraph. The smaller blocks on the right of the main RCM process blocks in figure 3-2 represents the sub-processes depicted in figure 3-1, and will be filled to indicate that a particular sub-process is being addressed in a particular paragraph.

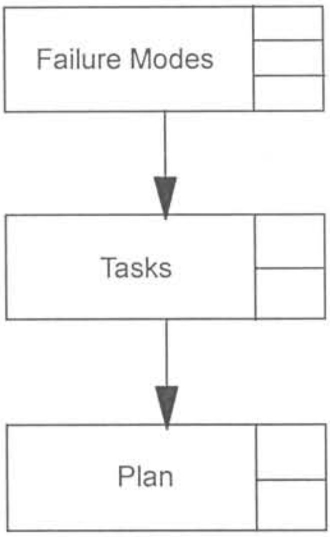

Figure 3-2: Tracking Diagram

### 3.2.1. The principle behind RCM: preserve function or preserve equipment?

Although the designation 'Reliability Centred Maintenance' was coined by the U.S. Department of Defence in 1975 [Matteson (1989), p. 6], the task of writing the first textbook (really a book-length report) [Nowlan and Heap (1978)] on RCM was given to United Airlines under contract. The content of this textbook was based on and refined the principles by now already firmly established in MSG-2, as successor to the successful Boeing 747 maintenance program design methodology, MSG-1.

One of the principal drives behind the development and use of MSG-1 was to move away from the principle of preserving equipment, to a principle of preserving function (especially functions essential to operating safety, but also those affecting economy of operation). Under the pre-RCM principle of preserving equipment, all equipment and all functions of such equipment had to be maintained (preserved), regardless of whether the preservation mattered or not. Under the RCM principle of preserving function, functions for which preservation does not matter are disregarded and the total force of maintenance actions is focussed on preserving essential equipment functions. Nowlan and Heap (1978), p.7 state this as follows: "*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.*" If a specific failure does not have significant consequences at the equipment level, it is totally disregarded in the design of the maintenance plan for the equipment. That is not to say that maintenance resources will not be spent on it – but it will be handled in a corrective way and not preventively.

This change in attitude is one of the major contributions of RCM to improving the quality of maintenance plans. Any text on the subject should thus emphasise both the fact of the change and its importance. The user of the methodology should realise that RCM does not merely assist him in designing a universal 'good' maintenance plan. It rather creates a maintenance plan that will, because it preserves function, result in the organisation reaching its goals.

The principle of preserving function is neglected in the RCM literature. Because of this, the majority of the users of the technique do not appreciate the enormity of the move from preserving equipment to that of preserving function. The responsibility for this can certainly be laid at the door of the various authors, who fail to make RCM analysts and users aware of the subtle, but critical difference. Moubray (1991), when defining RCM, states (p.7) that 'a fuller definition of RCM could be "a process used to determine *what must be done to ensure that any physical asset continues to fulfil its intended functions* in its present operating context".' He then refers to this in a later chapter (p.37) regarding asset maintenance as being "the state we wish to preserve is one in which it continues to fulfil its intended *functions*." Nevertheless, not once is the difference between the two concepts accentuated. This is not only true of Moubray, so that this critique is not aimed only against him, but against all the other authors that fail to get this concept across.

Smith (1993) is the only author that gets this concept across properly. When listing the principal features of RCM, he states that '*Traditional methods for determining PM tasks start with the issue of preserving equipment operability, and such methods tend to focus the entire task selection process on what can be done to the equipment. As a rule, why it should be done is never clearly addressed (or documented, if such consideration was, in fact, ever investigated). RCM is a major departure from this traditional practice! Its basic premise is "**preserve function**" - not "**preserve equipment**." This approach forces the analyst to systematically understand (and document) the system functions that must be preserved without any specific regard initially as to the equipment that may be involved. It then requires the analyst to think carefully about how functions are lost - in functional failure terms, not equipment failure terms. The purpose of this approach is to develop a credible rationale for why one might eventually desire to perform an appropriate PM task rather than just arbitrarily deciding to do something because "it sounds right."* (italics and accentuation added)

Jones (1995) mentions in his introduction to RCM that "The shift of emphasis, however, is from maintenance of *equipment operation* to the RCM theme of maintenance of *system function*." This states the point exactly, but like Moubray, the difference between the two concepts is not explained in the remainder of his text, nor is it accentuated.

Business Centred Maintenance (BCM) [Kelly (1997)] is an extremely practical approach to the design of a maintenance plan, as is evident from the short description in chapter 2 (paragraph 2.6.4). It is typical of the way in which maintenance plans were developed in the pre-RCM era. The two opposing principles of preserving *equipment function* versus that of preserving *equipment operation* are not mentioned and the method uses the outdated principle of preservation of *equipment operation*.

MSG-3 (1993), Anderson and Neri (1990), Coetzee (1997/2) and Campbell (1995) do not even mention the principle of *preserving function* at all.

### 3.2.2. The selection of application areas for RCM

When one anticipates the maintenance of a complex system, you are confronted by the decision regarding which of the myriad of components and their failure modes should be taken into account to base the (successful) maintenance plan on. In the case of a single system, such as an aircraft, the answer to this question is less complex than in a typical industrial system (plant, mine, or the like). In the case of an aircraft, although the system is very complex, it does not have the complexity of as many different and uncoordinated manufacturers of different parts of the system.

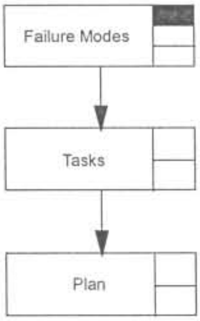

Nowlan and Heap addressed this problem through the introduction of the concept of a *Maintenance Significant Item (MSI)*, which they defined as follows:

> *A (maintenance) significant item is one whose failure could affect operating safety or have major economic consequences* [Nowlan and Heap (1978), p.80].

They also suggested that a 'partitioning' process be used to break the equipment down, first into major divisions (systems, powerplant and structure in their case), then further into systems, subsystems, assemblies and lastly parts (see figure 3-3).

The resultant tree in figure 3-3 must then be 'pruned' upwards to remove the obviously non-significant items. These are those items for which it can be demonstrated that they are non-significant because:

- their functions are unrelated to operating capability *or*
- the functions are replicated (this will be the case where redundancy has been designed into the system, either by parallel sub-systems or through standby sub-systems) *or*
- the item can be restored quickly following failure *or*
- the item cannot benefit from scheduled maintenance

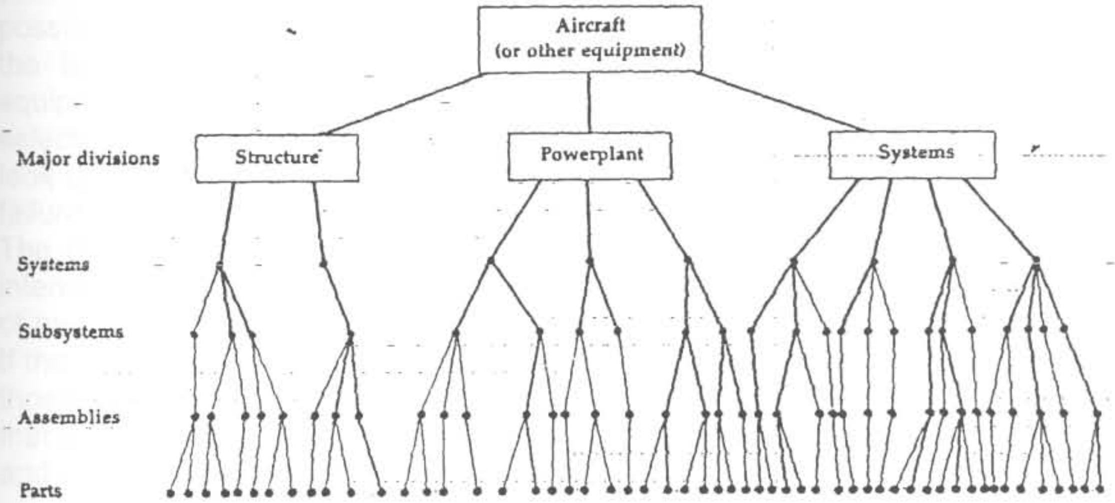

Figure 3-3: Tree breakdown structure (partitioning process)

The only exception in this process of classifying items non-significant is in the case of items with hidden functions. Such items might be classified as non-significant using the reasoning process described above, but will still be listed as items having to receive scheduled maintenance to reduce the occurrence of hidden failures to an acceptable level (so to reduce the probability the resultant multiple failure[1] to an acceptable level). The tree of significant items is shown in figure 3-4.

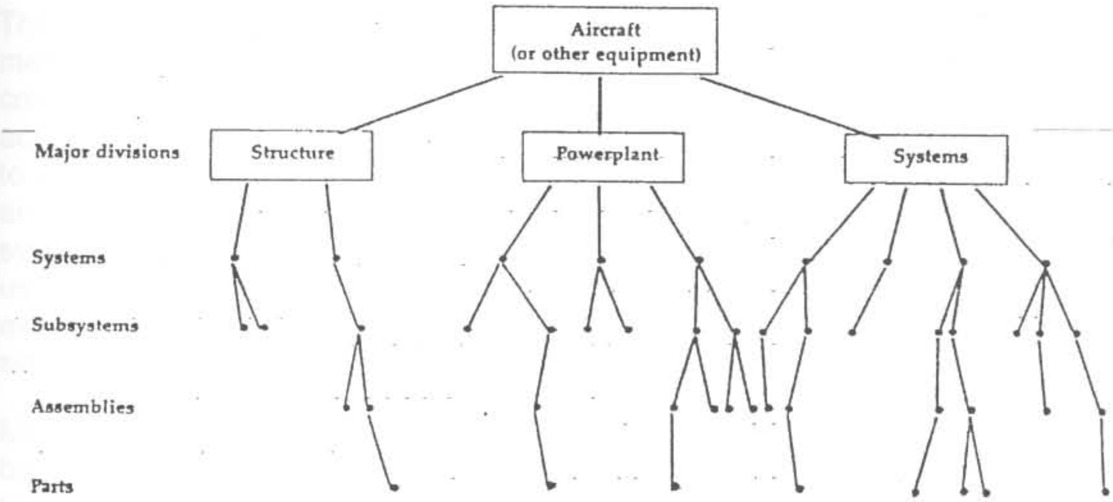

Figure 3-4: Tree of Maintenance Significant Items

The objective of this process, according to Nowlan and Heap (1978), is to find the most convenient level of each system or assembly to classify as significant. This resultant level must be low enough to ensure that no failure possibilities are overlooked, but high enough for the loss of function to have

---

[1] A typical example of a hidden failure consequence item comprises a non-failsafe safety device. If the device has failed there are no direct adverse consequences, but when the safety function of the device is needed and it is not available a catastrophic failure of the protected system mostly occurs. This failure plus the original (hidden) failure of the safety device are called a 'multiple failure'.

an impact on the equipment itself. The higher the level chosen, the more unwieldy the number of failure possibilities becomes, and the higher the possibility of missing important failure modes. On the other hand, the lower the level chosen, the higher the probability of losing contact with the equipment itself. Nowlan and Heap states (p.83) that "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 *comprising a sub-system* instead of many thousands *in a major division*." (italicised portions added for clarity). If the process is followed correctly, the parts selected as significant are usually those in which a critical failure mode originates. The authors caution (p.354) that there is a tendency amongst analysts to identify items based on their cost and complexity instead of based on their failure consequences.

As can be anticipated, the original Nowlan and Heap technique works best for the situation for which it was developed. It is perfect for a situation where many of a specific design of equipment are sold to a large market. In such a situation, one can afford the expense of a meticulous dissection of the equipment with the objective of creating the best possible maintenance plan. To that, one can add the high profile of operating safety in the airline business as an additional reason why this approach works perfectly well for the airline situation. On the other hand, in the general industrial application of RCM, the method tends to be overly demanding and users then seek for shortcuts.

The RCM method that has been popularised in South Africa is the adapted method as suggested by Moubray (1991). Moubray has a reasonably comprehensive chapter on the plant register, which seems to suggest that he advocates the use of the plant register as a means to decide which machines to include in the RCM analysis. He then proceeds to the analysis of functions and failures without explaining how he selected certain systems/sub-systems/components for analysis. Appendix 2 of his book also proposes the use of the equipment hierarchy to guide the analysis process without explaining how the decision regarding the specific systems/sub-systems/components to be analysed should be taken.

Moubray warns (p.302) that "an analysis carried out at too high a level can become too superficial, while one done at too low a level can become unintelligible." He then follows up by giving guidelines regarding the dangers of analysing at too low or too high a level. The dangers of doing the analysis at too low a level includes:

- that it becomes difficult to conceptualise and define performance standards
- it becomes difficult to visualise failure consequences
- control and protective loops can become difficult to deal with
- the process generates vast amounts of paperwork, which can become quite intimidating.

On the other hand, the disadvantage of doing the analysis at too high a level is that one could overlook several failure modes. He then suggests that the correct level of analysis is that one which supports the function(s) of the asset being analysed best. One cannot find fault with this approach, which really expands and supports the approach of Nowlan and Heap, apart from the fact that it is not specific at all regarding the use of the plant register as guideline.

Smith (1993) suggests that, based on experience, the system level is the most appropriate choice for the design of a maintenance plan. In evaluating his approach, it must be taken into account that he and Tom Matteson [Matteson (1989)] was involved with the Electric Power Research Institute (EPRI) to transfer the knowledge gained in the airline industry to the power industry. Most of his experience was therefore gained in a process type industry. He, like Moubray, suggests that to carry the analysis out at too low a level causes many problems, such as:

- that it becomes difficult to define the significance of functions and functional failures
- it is difficult to perform priority rankings between failure modes that are competing for limited PM resources

At the other end of the spectrum, he suggests that even combining two systems in one analysis proved to be extremely cumbersome and difficult. He thus suggests that the analysis process start with a list of all the systems in the plant/business asset/major system to be analysed. This is very specific and it seems to be a very practical guideline, at least for use in process industries.

Neither Nowlan and Heap, nor Moubray suggested any way of concentrating the analysis effort on systems that are more important. Smith is the first author to suggest a choice of only certain systems for analysis. The choice is based on selection by using one of the following 'schemes':

i. systems with a present high preventive task and/or cost content
ii. systems with a history of a high number of corrective maintenance tasks being done
iii. a combination of the previous two schemes
iv. systems with high corrective maintenance cost
v. systems which contributes significantly towards partial and total plant/business asset shutdowns
vi. systems which cause concern regarding safety and the environment

He states that from experience it seems as if all of the above-mentioned schemes, apart from scheme (vi), have the same result. Scheme (vi) has, according to Smith, proven not to be a good indicator of where the maintenance improvement effort should be concentrated as it often tends to create a bias towards systems that, maintenance-wise, are low-cost/low-problem systems. This solution, excluding scheme (vi), looks like a good, solid practical approach.

The RCM descriptions given by Jones (1995) and Campbell (1995) roughly agree with that of the previous authors, although not one of them has sufficient detail to guide a user in using the RCM methodology. Their description of the method is meant more as an introduction to RCM than to expounding its detailed use. Jones' description of 'functional decomposition' follows a combination of the ideas of Nowlan and Heap and Smith. Campbell, on the other hand, follows the notions of Moubray regarding the use of the plant register as basis. Jones does not suggest any way of selection of only certain parts of the plant/business asset for analysis. Campbell suggests the use of the following parameters for this purpose: availability, process capability, quality, cost as well as safety and environmental risk. He does not give detail, though.

Gits (1984) first describes an elaborate approach of functional decomposition, leading to the identification of PFC's (Process Failure Combinations), but then settles for what he calls 'technical decomposition' in his 'satisficing approach'. This is none other than Nowlan and Heap's partitioning process. In fact, Gits defines technical decomposition as being the 'hierarchical partitioning of the TS (*TS ≡ Technical System*) into its parts.' He then goes on to explain that the level of analysis should be determined through 'pragmatic considerations'. These include:

- organisational aspects, such as:
  - spare part policy
  - outsourcing policy
  - the availability of maintenance resources
  - statutory requirements
- hardware structural aspects, such as:
  - replaceability of parts
  - accessibility of parts
- information that are available regarding failure behaviour

Through these 'pragmatic considerations' he leaves the field open for interpretation but gives good guidelines on how to do such an analysis. Le Clercq and Van den Broek (1999) is a case in point. Figure 3-5 shows an example of a resultant decomposition using the Gits method.

Kelly (1997) also uses a hierarchical breakdown of the plant/business asset in the order: plant -> units -> assemblies -> sub-assemblies -> components. The purpose of this decomposition is to identify the 'maintenance causing items'. These are very similar to the MSI's of Nowlan and Heap (1978). His analysis approach looks at the maintenance requirements of each maintenance-causing item on an item by item basis. He does not specifically address the level of analysis, but ranks units in order of criticality according to the impact of their failure on production and safety.

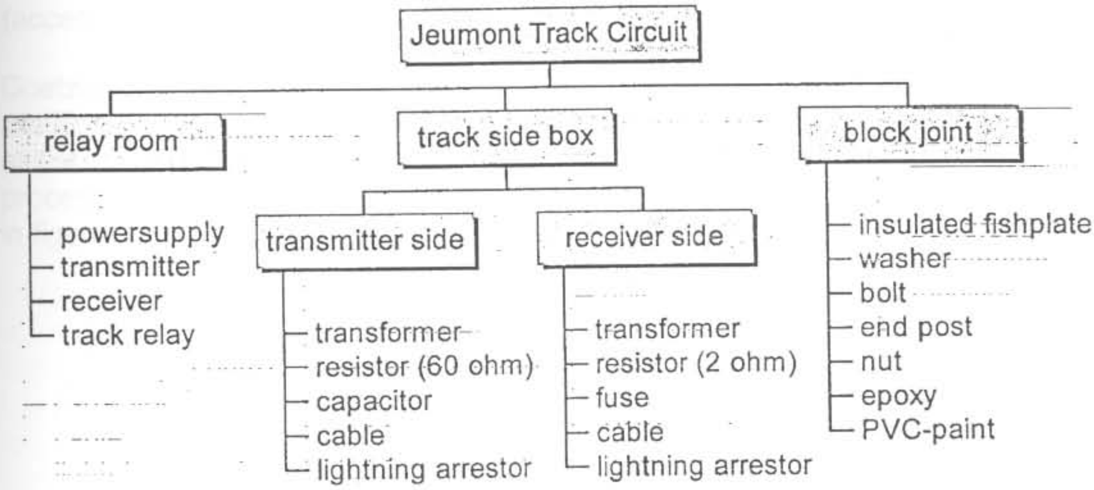

**Figure 3-5: Technical Decomposition a la Gits**

MSG-3 (1993) uses a top-down approach similar to that of Nowlan and Heap (1978) to identify the significant items on an aircraft. *Firstly* the aircraft is partitioned into major functional areas (or ATA Systems and Sub-Systems) and *then* further until sub-components, which are not replaced on-aircraft are identified. A candidate MSI is *then* usually a system or sub-system that is, in most cases, one level above the lowest (on-aircraft) level identified in the previous step. This level is considered the highest manageable level, i.e., one which is high enough to avoid unnecessary analysis, but low enough to be properly analysed to ensure that all functions, failures, and causes are adequately covered. They describe the process of identifying Maintenance Significant Items (MSI's) as a conservative process (using engineering judgement) based on the anticipated consequences of failure. The resultant MSI's are those items identified by the manufacturer whose failure could have any one or more of the following effects:

i. It affects safety (on ground or in flight).

ii. It could be undetectable or are not likely to be detected during operations.

iii. It could have significant operational impact.

iv. It could have significant economic impact.

The RCM handbook of the Naval Sea Systems Command (1983) prescribes the identification of all the ship's systems and then partitioning them in a logical way. It cautions that this partitioning process should rarely go below sub-system level. It states that "The need for analysis below the subsystem level depends on the complexity of the system and your knowledge and expertise. ... Nevertheless, understanding all of the functions of a complex system ... may require that you go, selectively, to the equipment level or below." It cautions that the tendency to want to analyse at too low a level is the direct result of misunderstanding the intention of an RCM analysis: "RCM is a methodology intended for use in developing ... the **preventive maintenance program for a ship**. If you are involved ... in the application of RCM, you should understand that intention. Otherwise, you may focus on some lesser level of assembly and its function rather than **on how it, in concert with other hardware, provides all the functions of the ship**." (accentuation added)

Coetzee (1997/2) makes use of the Nowlan and Heap method of partitioning/decomposition. He adds a simple decision diagram (p.86) that assists the user of the technique to decide when to stop the decomposition process down a specific leg of the decomposition tree. This diagram is shown in figure 3-6.

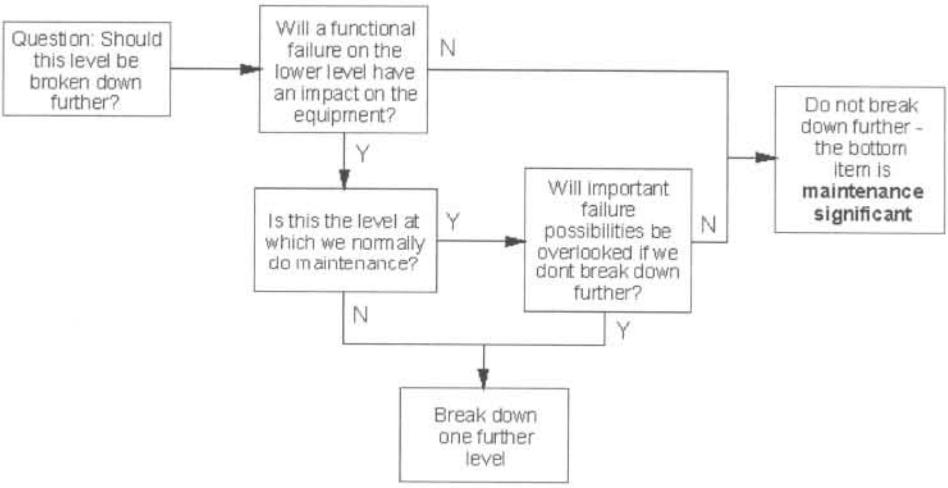

**Figure 3-6: Item breakdown decision diagram**

This diagram is based on the principle stated by Nowlan and Heap that the resultant level (of MSI's) must be low enough to ensure that no failure possibilities are overlooked, but high enough for the loss of function to have an impact on the equipment itself. This surely results in a more detailed RCM analysis than is the case with many of the later authors. The principle as stated above by the Naval Sea Command (1983) should be added to this diagram to make sure that the analysis remains function driven.

Coetzee also adds two prioritisation steps to ensure that the RCM analysis is concentrating on the high priority plant items or business assets. Before doing the partitioning as described above, he suggests that the plant/business asset be divided into logical groups of plant items/major systems/equipment. A Pareto analysis based on the profit contribution per group is then used to select those groups that should be analysed further using the RCM methodology. Following the partitioning process for these selected groups, he then suggests a further Pareto analysis of the resultant MSI's, based on each MSI's contribution to the downtime of the particular system/equipment to identify the MSI's that contribute most to the downtime of these, the highest profit contributing items. These two prioritisation processes constitute an effective 'funnelling' action to reduce the size of the total RCM analysis effort and to focus the effort on the most important items. The problem with the first prioritisation process is that it only focuses on profit and leaves out safety and quality as specific instances of results. Coetzee does however add in a footnote, that "Care should be taken not to disregard those plant items/business assets that do not contribute directly to profit, but which are of cardinal importance in ensuring long term profit." Care should be taken at all times while selecting the items for analysis that all important MSI's, based on the overall business objectives, be selected for further analysis.

### 3.2.3. Information Assembly

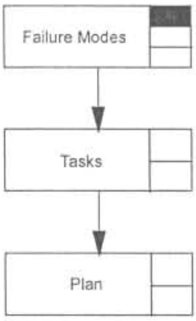

Probably because of the very specific business area for which it was written, the book of Nowlan and Heap (1978) does not specifically address the issue of the study of application areas. They do mention the necessity of providing analysts with schematics, full descriptions of the hardware and its relationship to other aspects of the aeroplane, though, in the appendix on auditing the RCM process (p.354). They apparently assume that the relevant information is normally available at the fingertips of the analysts performing RCM analyses for airlines and military aircraft. They even assume a close relationship with the designer of the system, which in their case might be a reasonable assumption. Moreover, they put a very large emphasis on the management of the process through auditors (see paragraph 3.2.10 for a discussion of the auditing process). However, in the industrial application of RCM this (information assembly - collecting all the relevant information that are necessary for the RCM analysis) is a very important aspect of the analysis process, and as such gets attention from most of the authors writing application texts.

Although Moubray (1991) stresses the importance of information, he depends on the knowledge of people (mostly operating and maintenance staff) to provide the information inputs to the process. The knowledge and experience of the people involved in the operation and maintenance of equipment are often one of the best sources of information regarding, especially, the short term history and behaviour of the plant/business system. In his chapter on implementing RCM (p.229) Moubray asks the question 'Who knows?' and then resolves the question with the comment 'More often than not, "somewhere" actually turns out to be "someone" – someone who has intimate knowledge and experience of the asset under consideration.' He then leaves it at that. His approach is centred on the knowledge of people operating and maintaining the plant/business system. He does make provision for technical experts to sit in on the RCM sessions, but even then he relies on the expertise of a person instead of verifying it with information from reputable sources, plant maintenance and operational manuals, plant maintenance and operational history and design documentation.

Smith (1993) has a high regard for the inputs from people, in line with that of Moubray. However, he does not see that as being sufficient and adds the use of plant-specific and generic data files (p.56). Under the heading 'Information collection' he presents the following list of documents and information that are required for the RCM analysis process:

- system piping and instrumentation diagram
- system schematic and/or block diagram
- individual vendor manuals for the various pieces of equipment in the system
- equipment history files
- system operation manuals
- system design specification and descriptive data

It is again apparent, that most of his later experience was gained in the process industries. The above-mentioned list of information items will not apply exactly as is to all industries, but the principle is sound. One should get as much help as is possible from documented sources. Smith adds that one will often have to create the system piping and instrumentation diagram via a 'system walkdown' and visual reconstruction of the as-built configuration. He mentions that the one obvious omission from the above list is the present (non-RCM) maintenance plan. He recommends that this should only be made available near the point of completion of the RCM plan. This principle is sound, taking into account the nature of mankind and is very specifically borne out by Nowlan and Heap in their appendix on RCM plan auditing (p. 361). They state that some analysts have the tendency to work backwards through the decision logic from either the present task or their own favourite task, in order to justify it and to save some analysis effort.

Gits (1984, 1992) does not specifically mention the aggregation of information as a specific step in his process. Nevertheless, the project of Le Clercq and Van den Broek (1999) had a very definite information gathering phase, which included the experience of maintenance and operations staff, manuals, schematics as well as material (such as annual reports and promotional brochures) to understand the business of Metrorail.

Kelly (1997) makes use of process flow diagrams (separate ones at system and sub-system levels) and a plant/business asset inventory, which are drawn up as part of the first step of his BCM process. The flow diagrams have to specify plant logistical information, such as plant production rates and buffer storage capacities, as well as item-redundancies. He specifies that the analyst should develop a proper understanding of the plant operating characteristics and the production management policies followed. The user requirements regarding maintenance results (i.e. availability) should be obtained as well and translated down to unit level as input to the plan design process. Additionally to this, the scheduling characteristics of the plant should be studied and documented regarding *firstly* maintenance opportunity windows arising from production schedules and operational characteristics (e.g. batch operation). *Secondly*, plant item/major system interdependencies should be determined (this can involve flexible downtime opportunities flowing from the buffer capacity created through in-process storage, as well as limitations concerning series system availabilities). He moreover makes use of manufacturer's manuals, drawings and schematics in the process of plan design.

MSG-3 (1993) suggests that one or more working groups, consisting of specialist representatives from the participating operators (airlines), the prime manufacturer, and the regulatory authority, be constituted. Their task is to ensure that the highest amount of operating and maintenance experience, data and other source material are available for the ensuing maintenance plan analyses. The steering committee, which steers the maintenance plan analysis, may alternatively arrange some other means for obtaining the detailed technical information necessary to develop recommendations for maintenance programs in each area. The point is that, irrespective of the organisation of the working activity, written technical data, which is not specified in detail in the MSG methodology, must be provided to support the recommendation to the steering committee. After approval by this committee, these analyses and recommendations are consolidated into a final report for presentation to the regulatory authority.

Naval Sea Systems Command (1983) has very specific requirements regarding the information requirements. They classify the information under the headings 'descriptive information' and 'operating information':

- Descriptive Information
  - Narrative descriptions
  - Design specifications
  - System schematics
  - Assembly drawings
  - Field and engineering changes
- Operating Information
  - Operating and maintenance instructions
  - Condition and performance standards
  - Failure data
  - Existing maintenance specification

They state explicitly that the last item should only be used for reference purposes after the RCM analysis has been completed.

### 3.2.4. Identification of Failure Modes

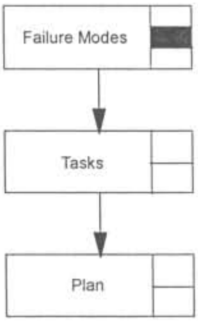

The identification of failure modes was not addressed in the case of Nowlan and Heap (1978). Their book has very few references to the process of identifying failure modes, as they use a Failure Modes and Effects Analysis (FMEA) which is supplied by the manufacturer (p. 80). In addition to this, they accentuate that airlines have knowledge and operating experience with similar items, as well as knowledge of the failure consequences in the particular operating context in which the equipment is used. The important implication is that RCM was developed around the technique of FMEA as a basis to provide the majority of the knowledge regarding failure modes for the RCM analysis process. Different to the typical industrial situation, the FMEA knowledge that is available up front gave them an edge in the choice of MSI's (paragraph 3.2.2).

In industrial applications the FMEA analysis is not a given and needs to be done as part of the analysis process. As the industrial version of the technique developed, it was in most cases enhanced by practical experience. In many cases, the FMEA was extended to a form of FMECA (Failure Mode, Effects and Criticality Analysis) to provide a way of classifying the failures as being more or less critical. The criticality aspect is addressed in paragraph 3.2.5. The format of an FMEA worksheet as given in MIL-STD-1629A will be regarded as the standard format and is shown in figure 3-7.

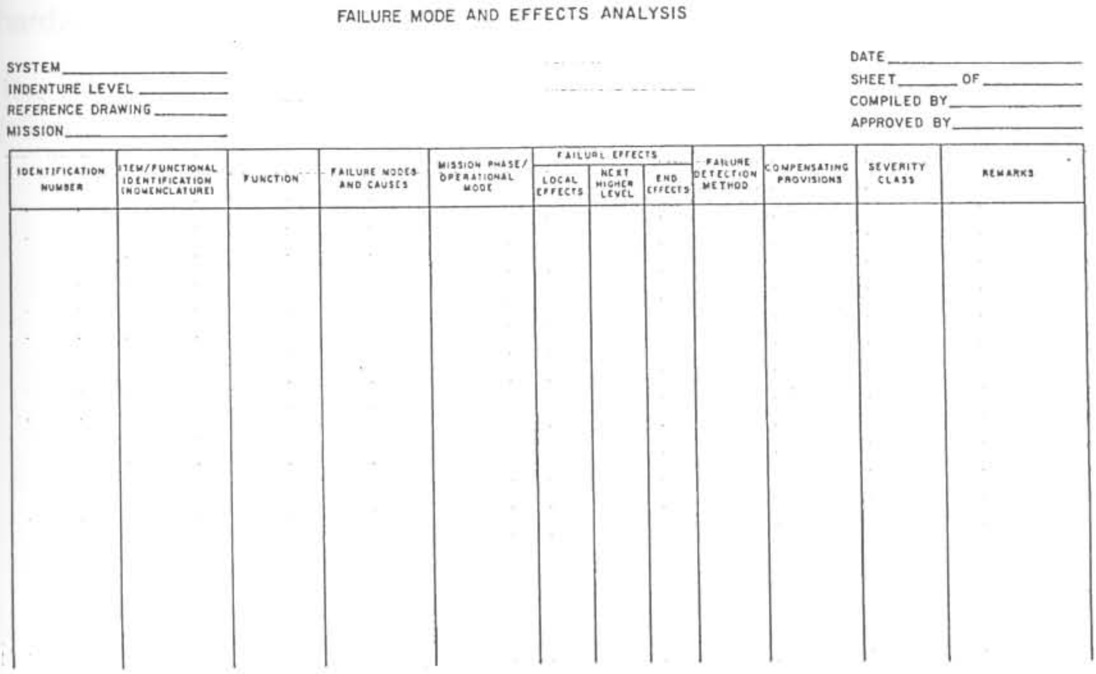

FAILURE MODE AND EFFECTS ANALYSIS

SYSTEM__________
INDENTURE LEVEL __________
REFERENCE DRAWING __________
MISSION__________

DATE__________
SHEET_____ OF _____
COMPILED BY__________
APPROVED BY__________

| IDENTIFICATION NUMBER | ITEM/FUNCTIONAL IDENTIFICATION (NOMENCLATURE) | FUNCTION | FAILURE MODES AND CAUSES | MISSION PHASE/ OPERATIONAL MODE | FAILURE EFFECTS | | | FAILURE DETECTION METHOD | COMPENSATING PROVISIONS | SEVERITY CLASS | REMARKS |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | LOCAL EFFECTS | NEXT HIGHER LEVEL | END EFFECTS | | | | |
| | | | | | | | | | | | |

Figure 3-7: Standard FMEA worksheet

The standard content format for the different columns of the FMEA worksheet (figure 3-7), as found in MIL-STD-1629A is given in Table 3.1. The FMEA is really a design document, which is used to analyse a technical system during the systems design process with the objective of designing out critical failure modes, to improve the inherent design reliability of the system. MIL-STD-1629A defines the purpose of the FMEA as follows:

> *The purpose of the FMEA is to study the results or effects of item failure on system operation and to classify each potential failure according to its severity.*

The idea is to eliminate at least Category I (catastrophic) and Category II (critical) failures (see Table 3.1 for the definition of severity classes) during the design and system development phases. The same technique can (and should) of course be used by production organisations in the process of purchasing systems 'off the shelf'. This could ensure that design flaws are corrected prior to the purchase, which will lead to less downtime, less maintenance cost and less logistical headaches. Unfortunately, this is done in very few cases.

There are two approaches that can be followed in the analysis process, a *hardware approach* and a *functional approach*. The *hardware approach* is usually followed in less complex situations where the hardware items can be uniquely identified from schematics and drawings and is typically performed in a bottom-up way (from the parts level upward). It will thus develop the FMEA analysis, progressing from the parts level upward until the systems level is reached and the analysis being complete. On the other hand, the *functional approach* is used in complex situations and where the hardware cannot be uniquely identified. Such analysis is typically conducted from the systems level downward. For complex systems, a combination of the functional and hardware approaches may also be considered.

The procedure for conducting an FMEA, according to MIL-STD-1629A, is as follows:

i. The first step is to define the system to be analysed. This includes the overall system function(s) in its intended operating situation. As part of this definition, the intended use of the system is analysed to reveal multiple functional use profiles and multiple modes of operation per functional use profile.

ii. Use is made of functional and reliability block diagrams to support the analysis process. The functional block diagram is required to show the functional flow sequence, while the reliability block diagram is used to show the series dependence/independence of functions and operations. More than one block diagram will be needed in the case of alternative operational modes. Both types of block diagrams is needed if the FMEA is conducted using the functional approach, while only the functional block diagram is used for the hardware approach, as the reliability interrelationships are already implied in the hardware method.

iii. The FMEA analysis is performed using Table 3.1 as guideline. Each single item failure, as its effects are analysed, is considered to be the only failure in the system (i.e. the remainder of the system's functions and components are in a non-failed state). There are two exceptions to this rule. The *first* of these is where the single item failure, that is being analysed, is hidden – in such a case the analysis will investigate whether the failure of a second component or more components, given the hidden failure, will lead to a catastrophic or critical multiple failure condition. The *second* exception occurs in the analysis of safety, redundant or back-up systems. In such a case, the analysis is widened to include the failure condition(s) that led to the incorporation of the safety, redundant or back-up items.

**Table 3.1: Standard FMEA process**

| Column Heading (figure 3-7) | Column Content |
|---|---|
| Identification number | A serial number or other reference number – used for traceability purposes – typically the block numbers in a Reliability Block Diagram – can also be used to refer to MSI numbers |
| Item/Functional identification | The name or number which identifies the item or system function being analysed |
| Function | A concise statement of the function performed by the item. This includes both the inherent function of the item and its relationship to interfacing items |
| Failure modes and causes | All predictable failure modes for each level analysed are identified and described. Failure modes are identified by studying the item outputs and functional outputs identified in applicable block diagrams and schematics<br>Each failure mode typically has more than one probable cause – all of these are listed for each failure mode identified. It is often helpful to also identify failure causes as seen from one level above and one level below the present level as well<br>To ensure that a complete analysis results, each failure mode and output should be examined in relation to the following failure possibilities:<br>▪ Premature operation<br>▪ Failure to operate at the prescribed time<br>▪ Intermittent operation<br>▪ Failure to cease operation at the prescribed time<br>▪ Loss of output or failure during operation<br>▪ Degraded output/operational capability |

**Table 3.1: Standard FMEA process (continued)**

| Column Heading (figure 3-7) | Column Content |
|---|---|
| Mission phase/operational mode | A concise statement of the mission phase and operational mode during which the failure occurs |
| Failure Effect | Evaluate and record the consequences of the failure on item operation, function or status. In this process, also consider mission objectives, maintenance requirements and personnel/system safety.<br>Because the failure may impact levels other than the present one, the effect is also evaluated under the following three sub-headings:<br>*Local effects*: impact of the failure at the present level of investigation. This forms a basis for evaluating compensating provisions (column 10). It is possible for the local 'effect' to be the failure mode itself.<br>*Next higher level effect*: the impact of the failure at the level of which the present one forms part.<br>*End effects*: the impact of the failure on the operation, function or status of the system itself. This end effect may be the consequence of more than one failure. |
| Failure detection means | The failure detection means, through which the operator detects the occurrence of the failure mode, is identified. These include:<br>▪ Identified warning devices and instruments<br>▪ Other indications<br>The most direct method to isolate the failure should also be investigated and reported |

**Table 3.1: Standard FMEA process (continued)**

| Column Heading (figure 3-7) | Column Content |
|---|---|
| Compensating provisions | This includes design provisions or operator actions, which circumvent or mitigate the effect of the failure.<br>Design provisions: include redundancy, safety/relief devices and alternative modes of operation.<br>Operator actions: The action that could best be taken by the operator should be identified and listed. Any incorrect action(s) by the operator should also be listed with the anticipated effect(s). |
| Severity classification | Each failure mode and each item is assigned a severity classification, according to the failure effect, from the following list:<br>▪ Category I – Catastrophic – a failure that may cause death or system loss.<br>▪ Category II – Critical – a failure that may cause severe injury, major property damage, major system damage, which will result in mission loss.<br>▪ Category III – Marginal – a failure that may cause minor injury, minor property damage, minor system damage, which will result in delay, mission degradation.<br>▪ Category IV – Minor – a failure that will not cause injury, property damage, system damage, bur which will result in unscheduled maintenance or repair. |
| Remarks | Any remarks which could clarify why certain conclusions were made in the preceding columns on this row.<br>Notes regarding recommendations for design improvements. |

A simpler approach to an FMEA analysis is proposed by McDermott et al (1996), which will be attractive to many maintenance users. They define (p.3) FMEA as "*a systematic method of identifying and preventing product and process problems before they occur*". They suggest a less rigorous approach, using teams of people to develop FMEA analyses, instead of the analyst in the corner approach of MIL-STD-1629A.

Their approach to FMEA's is a more generalised one, advocating the use of the FMEA process to solve any problem at hand – examples listed include safety hazards, financial strategies, software quality, marketing strategy and purchasing of major (capital) equipment.

The teams in this approach are formed whenever needed and disbanded when their task is complete. Each area of expertise involved should have at least one representative on the team, with a team size of four to six people working best, according to McDermott et al (1996). They also suggest that using some people on the team, who are relative outsiders regarding the problem at hand, may be helpful in achieving a balanced result. Such person(s), while not having the required expertise, often have valuable insights which may be missed by people working near to the problem from day to day. They also do not have an emotional investment in the specific situation and thus tend to be objective.

Team members must know the basics of working in a team. Training, if necessary, should include knowledge of consensus building techniques, team project documentation and idea-generating techniques such as brainstorming. They should also be comfortable in the use of flowcharting, data analysis and graphing techniques. The team leader is the only person that must have a good knowledge of the use of the FMEA process.

The leader of the team should be appointed by management or selected by the team members immediately following the formation of the team. Such an individual, who takes part in the process as facilitator and not decision-maker, is responsible for co-ordinating the FMEA process. Such co-ordination includes (1) setting up meetings, (2) facilitating meetings, (3) securing the necessary team resources, and (4) ensuring progress in the right direction. Other roles in the FMEA process includes that of scribe (a role which is typically rotated amongst the team members, excluding the leader) and the process expert. The process expert is someone who has extensive knowledge/experience of the process being investigated and typically has quite an investment in it. The purpose of the FMEA is to find fault with the process expert's work, which makes it difficult for the person. Whereas such a person can help speed up the process, the emotional aspects involved can also slow it down, if not handled correctly by the leader.

Figure 3-8 shows a form suggested by the authors [McDermott et al (1996)] to clarify the boundaries of freedom, within which the team performs its duties, as well as the scope of the FMEA. The former relieves the team of any unnecessary pressures and thus allows them to progress purposefully, while the latter ensures focus on the correct problem.

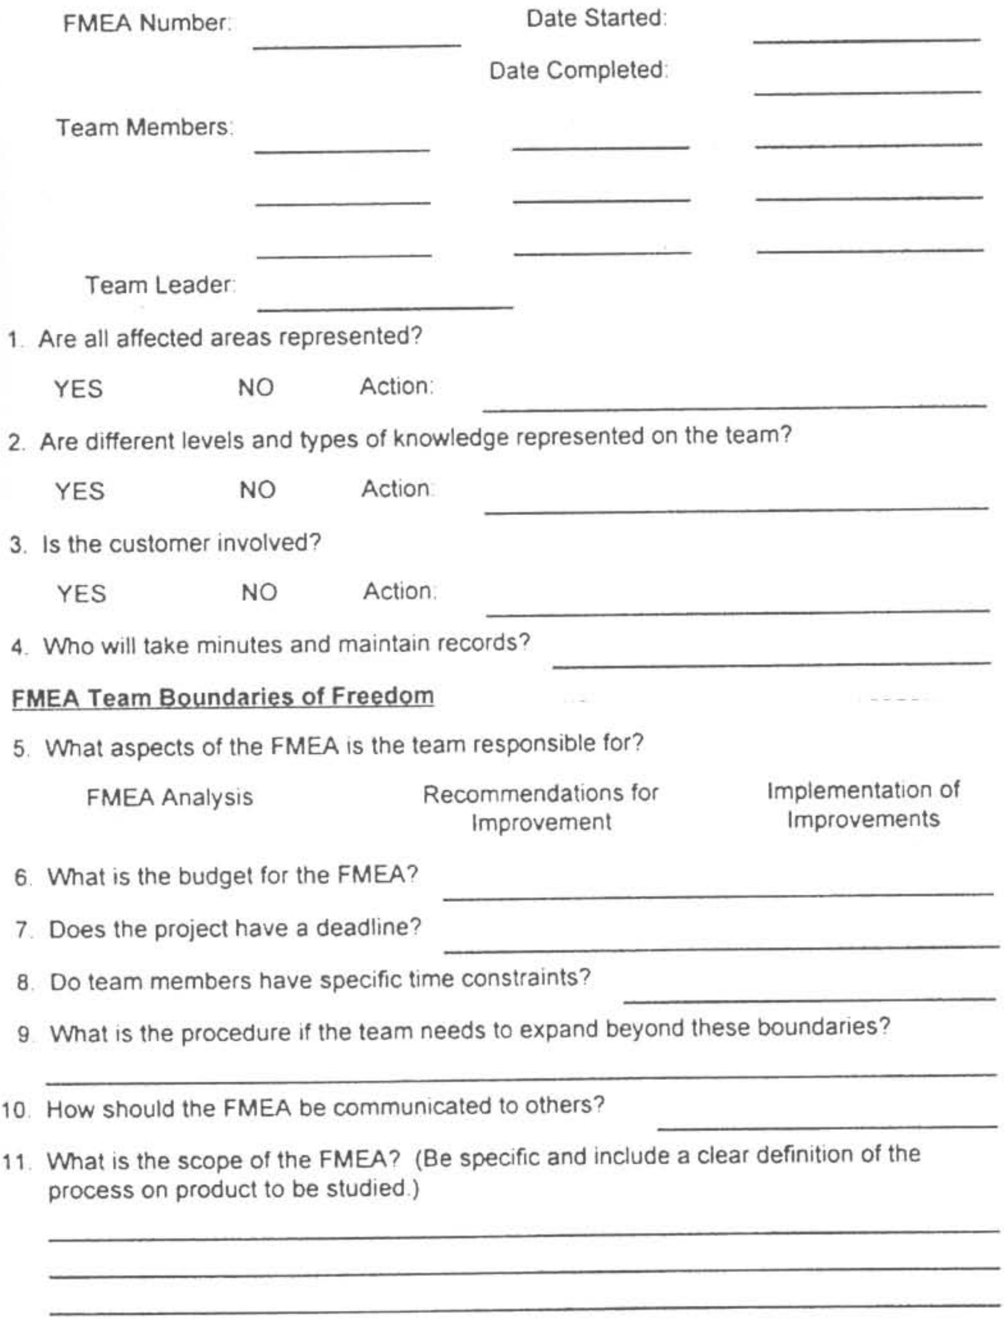

FMEA Number: ____________ Date Started: ____________

Date Completed: ____________

Team Members: ____________ ____________ ____________

____________ ____________ ____________

____________ ____________ ____________

Team Leader: ____________

1. Are all affected areas represented?

YES NO Action: ____________

2. Are different levels and types of knowledge represented on the team?

YES NO Action: ____________

3. Is the customer involved?

YES NO Action: ____________

4. Who will take minutes and maintain records? ____________

**FMEA Team Boundaries of Freedom**

5. What aspects of the FMEA is the team responsible for?

FMEA Analysis Recommendations for Improvement Implementation of Improvements

6. What is the budget for the FMEA? ____________

7. Does the project have a deadline? ____________

8. Do team members have specific time constraints? ____________

9. What is the procedure if the team needs to expand beyond these boundaries?

____________

10. How should the FMEA be communicated to others? ____________

11. What is the scope of the FMEA? (Be specific and include a clear definition of the process on product to be studied.)

____________

____________

____________

Figure 3-8: FMEA Team Start-up Worksheet [McDermott et al (1996)]

Figure 3-9 shows the FMEA worksheet used by the authors, utilising a 10-step process explained in Table 3.2. Only the first 8 steps are shown, as the last two are not relevant to the discussion.

The main deficiency of this method, when compared with the more rigorous standard FMEA analysis in MIL-STD-1629A, is that it does not lean as heavily on documentary proof of actual experience, but tend to be more subjective due to the team approach. On the other hand, it is much simpler to implement and will certainly lead to an FMEA, which can be used for the generation of a maintenance plan. Some failure modes may be missed, however, and non-existent failure modes may be listed due to the 'group talk' (subjective) effect.

Process/Product: ______
FMEA Team: ______
Team Leader: ______

FMEA Number: ______
FMEA Date: (Original) ______
(Revised) ______
Page: __ of __

| FMEA Process | | | | | | | | | | | Action Results | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Item and Function | Potential Failure Mode | Potential Effect(s) of Failure | Severity | Potential Cause(s) of Failure | Occurrence | Current Controls | Detection | RPN | Recommended Action | Responsibility and Target Completion Date | Action Taken | Severity | Occurrence | Detection | RPN |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | |
| Total Risk Priority Number | | | | | | | | | Resulting Risk Priority Number | | | | | | |

Figure 3-9: FMEA analysis worksheet [McDermott et al (1996)]

| Rating | Description | Definition |
|---|---|---|
| 10 | Dangerously high | Failure could injure the customer or an employee. |
| 9 | Extremely high | Failure would create noncompliance with federal regulations. |
| 8 | Very high | Failure renders the unit inoperable or unfit for use. |
| 7 | High | Failure causes a high degree of customer dissatisfaction. |
| 6 | Moderate | Failure results in a subsystem or partial malfunction of the product. |
| 5 | Low | Failure creates enough of a performance loss to cause the customer to complain. |
| 4 | Very Low | Failure can be overcome with modifications to the customer's process or product, but there is minor performance loss. |
| 3 | Minor | Failure would create a minor nuisance to the customer, but the customer can overcome it in the process or product without performance loss. |
| 2 | Very Minor | Failure may not be readily apparent to the customer, but would have minor effects on the customer's process or product. |
| 1 | None | Failure would not be noticeable to the customer and would not affect the customer's process or product. |

*Should be modified to fit the specific product or process.

Figure 3-10: Generic Severity rating scale [McDermott et al (1996)]

**Table 3.2: Simplified FMEA process [McDermott et al (1996)]**

| Column Heading (figure 3-9) | Column Content |
|---|---|
| Step 1: Review the process | Before starting the process the team should review the available documentation (see paragraph **Error! Reference source not found.**). If a detailed flowchart is not available, the team should create one (appendix 1 of the booklet). An expert is helpful here to answer any questions that the team may have. |
| Step 2: Brainstorm potential failure modes | To facilitate a thorough list of failure modes, the authors suggest that a series of brainstorming sessions be used, each focussed on a different element (e.g. people, methods, equipment, materials and the environment).<br>Once the brainstorming is complete, the team can organise the various ideas into like categories to facilitate the elimination of non-viable ideas and the combination of similar ideas. |
| Step 3: List potential failure effects of each failure mode | There may be several effects for each failure mode – each must be listed separately.<br>This step should be performed very thoroughly, as its result will form the basis of the work performed in the next steps.<br>The analysis is performed by repeatedly asking the question: *if* the failure occurs; *then* what are the consequences? |
| Step 4: Assign Severity Rating | Severity ≡ Consequence of failure<br>Rating done on a scale of 1 (not severe at all) to 10 (extremely severe).<br>Each effect should be given its own severity rating, even if there are several effects for a particular failure mode.<br>Authors give a generic rating scale – see figure 3-10. |

**Table 3.2: Simplified FMEA process [McDermott et al (1996)] (continued)**

| Column Heading (figure 3-9) | Column Content |
|---|---|
| Step 5: Assign Occurrence Rating | Occurrence ≡ Probability of failure occurring (or the frequency at which it occurs)<br>Rating done on a scale of 1 (probability virtually nil) to 10 (probability extremely high).<br>Authors give a generic rating scale – see figure 3-11. |
| Step 6: Assign Detection Rating | Detection ≡ Probability of failure being detected<br>Rating done on a scale of 1 (probability extremely high) to 10 (probability virtually nil).<br>Authors give a generic rating scale – see figure 3-12. |
| Step 7: Calculate the Risk Priority Number (RPN) for each failure mode | RPN = Severity x Occurrence x Detection<br>Those failure modes with highest RPN's are the ones with most potential for savings through preventive action. |
| Step 8: Prioritise the failure modes | The team must decide which of the failure modes deserve further attention (preventive action).<br>A cut-off RPN may be set by the organisation for this purpose.<br>Pareto-analysis may be used. |

On the micro level, the analysis worksheet (figure 3-9) has certain deficiencies when compared to the standard worksheet (figure 3-7). These include a lack of codification/cross referencing facilities and a lack of differentiation between different modes of operation. It also does not analyse the difference between failure effects between the own (local) level, the next higher level and the systems level. This is probably because the authors do not use any systems segmentation. This lack of sophistication may lead to an inferior analysis, with a resultant sub-optimal maintenance plan.

On the other hand, the RPN risk assessment is a simple, but effective, method to limit maintenance actions to the more risky failure modes.

| Rating | Description | Potential Failure Rate |
|---|---|---|
| 10 | Very High: Failure is almost inevitable | More than one occurrence per day or a probability of more than three occurrences in 10 events ($C_{pk} < 0.33$). |
| 9 | | One occurrence every three to four days or a probability of three occurrences in 10 events ($C_{pk} \approx 0.33$) |
| 8 | High: Repeated failures | One occurrence per week or a probability of 5 occurrences in 100 events ($C_{pk} \approx 0.67$). |
| 7 | | One occurrence every month or one occurrence in 100 events ($C_{pk} \approx 0.83$) |
| 6 | Moderate: Occasional failures | One occurrence every three months or three occurrences in 1,000 events ($C_{pk} \approx 1.00$) |
| 5 | | One occurrence every six months to one year or one occurrence in 10,000 events ($C_{pk} \approx 1.17$) |
| 4 | | One occurrence per year or six occurrences in 100,000 events ($C_{pk} \approx 1.33$). |
| 3 | Low: Relatively few failures | One occurrence every one to three years or six occurrences in ten million events ($C_{pk} \approx 1.67$). |
| 2 | | One occurrence every three to five years or 2 occurrences in one billion events ($C_{pk} \approx 2.00$). |
| 1 | Remote: Failure is unlikely. | One occurrence in greater than five years or less than two occurrences in one billion events ($C_{pk} > 2.00$) |

* Should be modified to fit the specific product or process.

**Figure 3-11: Generic Occurrence rating scale [McDermott et al (1996)]**

| Rating | Description | Definition |
|---|---|---|
| 10 | Absolute Uncertainty | The product is not inspected or the defect caused by failure is not detectable. |
| 9 | Very Remote | Product is sampled, inspected, and released based on Acceptable Quality Level (AQL) sampling plans. |
| 8 | Remote | Product is accepted based on no defectives in a sample. |
| 7 | Very Low | Product is 100% manually inspected in the process. |
| 6 | Low | Product is 100% manually inspected using go/no-go or other mistake-proofing gauges. |
| 5 | Moderate | Some Statistical Process Control (SPC) is used in process and product is final inspected off-line. |
| 4 | Moderately High | SPC is used and there is immediate reaction to out-of-control conditions. |
| 3 | High | An effective SPC program is in place with process capabilities ($C_{pk}$) greater than 1.33. |
| 2 | Very High | All product is 100% automatically inspected. |
| 1 | Almost Certain | The defect is obvious or there is 100% automatic inspection with regular calibration and preventive maintenance of the inspection equipment. |

* Should be modified to fit the specific product or process.

**Figure 3-12: Generic Detection rating scale [McDermott et al (1996)]**

Moubray (1991) uses a worksheet that is a simplified version of that of MIL-STD-1629A, but with certain added features. The worksheet is shown in figure 3-13. The worksheet has a comprehensive identification part in the heading, which corresponds to and improves upon that of MIL-STD-1629A (especially when taking into account that it has an industrial, non-military function). Moubray uses one (or more) sheet per item, so that the first two columns in MIL-STD-1629A become redundant, at the expense of more paperwork. He then adds a column for Functional Failures after Nowlan and Heap (1978), who, although they did not include the FMEA in their text, implied this structure in their further analysis of failure consequences (p.87). This column, which is also used by practically all of the application texts (including Coetzee (1988), Smith (1993), Jones (1995) and Coetzee (1997/2)), assists the analyst in identifying all the relevant failure modes through a process more logical to the human mind. The process followed here is item | function | functional failure | failure mode, which corresponds with MIL-STD-1629A's item | function | failure mode | cause (in other words, it is only a matter of definition as to whether a failure mode is the functional failure or the root cause of failure – see Moubray p.57). The latter definition, as used by Moubray and other RCM texts, is the preferred one in the RCM world.

Moubray (1991) inadvertently addressed at least part of the problem identified by Harris (1985) (see paragraph 3.2.6) by including functional standards in the function descriptions. This makes it less probable that functional failures of the 'process degradation' class will not be identified - these are often overlooked in standard RCM, according to Harris.

Another feature, that Moubray (1991) adds (see figure 3-13), is a codification that is used as cross reference back to the FMEA analysis sheet from his later 'RCM II Decision Worksheet'. He uses numeric characters to codify the Functions and Failure Modes, while he uses alphabetic characters to codify functional failures. Then, lastly, he prefers to write down the failure effect(s) in words instead of analysing them with the purpose to determine the size of the risk involved (MIL-STD-1629A's severity class and McDermott's RPN [McDermott (1996)]). The verbatim description of the failure effect is a large benefit during the further analysis and for archiving purposes, but is not sufficient to assist the user to limit the application of his preventive maintenance resources to critical/important failure modes only.

Smith (1993) divides the FMEA into three separate worksheets within two of his RCM steps. He first uses the worksheet depicted in figure 3-14 within his step 4: System Functions and Functional Failures. This worksheet is used *firstly* to list all the system functions – in this case, 'black box' functions provided by the present sub-system to other systems/sub-systems as well as passive functions required for system success (such as fluid boundary integrity). *Secondly*, the functional failures are listed per function, as is the case in the discussion above on Moubray's approach.

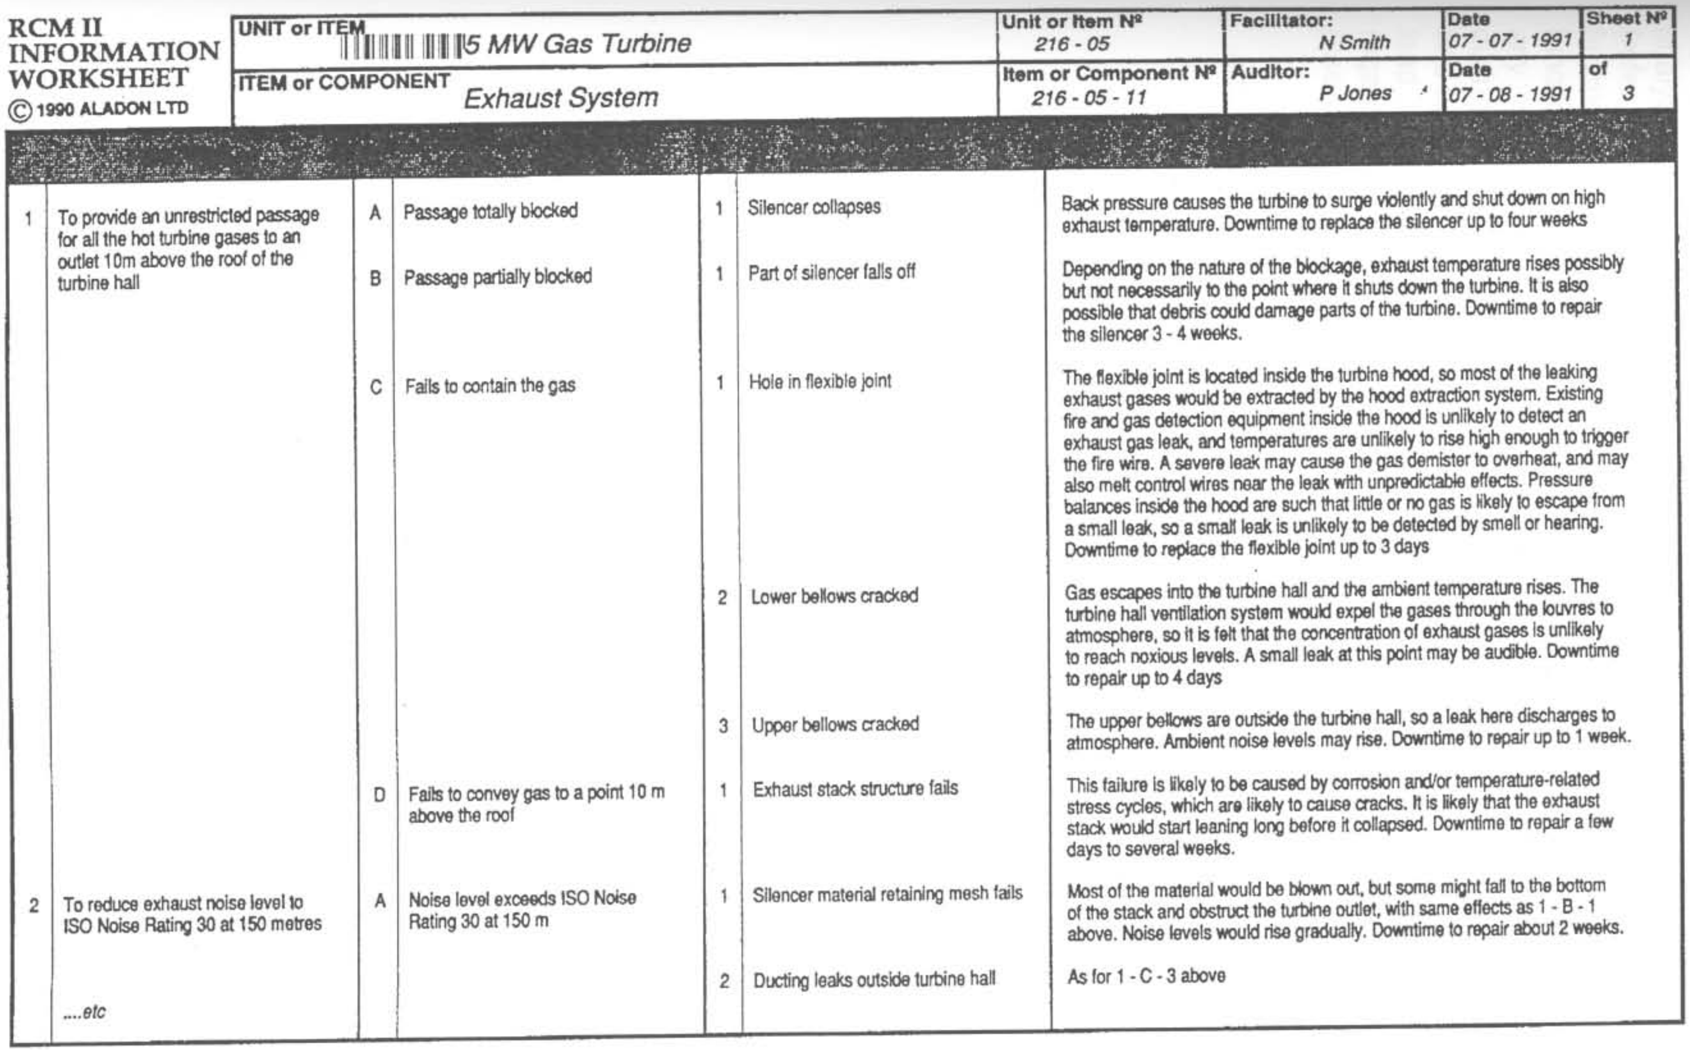

**RCM II INFORMATION WORKSHEET**
© 1990 ALADON LTD

| | | | | | | |
|---|---|---|---|---|---|---|
| **UNIT or ITEM** | [illegible]5 MW Gas Turbine | **Unit or Item Nº** 216 - 05 | **Facilitator:** N Smith | **Date** 07 - 07 - 1991 | **Sheet Nº** 1 |
| **ITEM or COMPONENT** | Exhaust System | **Item or Component Nº** 216 - 05 - 11 | **Auditor:** P Jones | **Date** 07 - 08 - 1991 | **of** 3 |

| | | | | | | |
|---|---|---|---|---|---|---|
| 1 | To provide an unrestricted passage for all the hot turbine gases to an outlet 10m above the roof of the turbine hall | A | Passage totally blocked | 1 | Silencer collapses | Back pressure causes the turbine to surge violently and shut down on high exhaust temperature. Downtime to replace the silencer up to four weeks |
| | | B | Passage partially blocked | 1 | Part of silencer falls off | Depending on the nature of the blockage, exhaust temperature rises possibly but not necessarily to the point where it shuts down the turbine. It is also possible that debris could damage parts of the turbine. Downtime to repair the silencer 3 - 4 weeks. |
| | | C | Fails to contain the gas | 1 | Hole in flexible joint | The flexible joint is located inside the turbine hood, so most of the leaking exhaust gases would be extracted by the hood extraction system. Existing fire and gas detection equipment inside the hood is unlikely to detect an exhaust gas leak, and temperatures are unlikely to rise high enough to trigger the fire wire. A severe leak may cause the gas demister to overheat, and may also melt control wires near the leak with unpredictable effects. Pressure balances inside the hood are such that little or no gas is likely to escape from a small leak, so a small leak is unlikely to be detected by smell or hearing. Downtime to replace the flexible joint up to 3 days |
| | | | | 2 | Lower bellows cracked | Gas escapes into the turbine hall and the ambient temperature rises. The turbine hall ventilation system would expel the gases through the louvres to atmosphere, so it is felt that the concentration of exhaust gases is unlikely to reach noxious levels. A small leak at this point may be audible. Downtime to repair up to 4 days |
| | | | | 3 | Upper bellows cracked | The upper bellows are outside the turbine hall, so a leak here discharges to atmosphere. Ambient noise levels may rise. Downtime to repair up to 1 week. |
| | | D | Fails to convey gas to a point 10 m above the roof | 1 | Exhaust stack structure fails | This failure is likely to be caused by corrosion and/or temperature-related stress cycles, which are likely to cause cracks. It is likely that the exhaust stack would start leaning long before it collapsed. Downtime to repair a few days to several weeks. |
| 2 | To reduce exhaust noise level to ISO Noise Rating 30 at 150 metres | A | Noise level exceeds ISO Noise Rating 30 at 150 m | 1 | Silencer material retaining mesh fails | Most of the material would be blown out, but some might fall to the bottom of the stack and obstruct the turbine outlet, with same effects as 1 - B - 1 above. Noise levels would rise gradually. Downtime to repair about 2 weeks. |
| | | | | 2 | Ducting leaks outside turbine hall | As for 1 - C - 3 above |
| | ....etc | | | | | |

Figure 3-13: FMEA worksheet in RCM II

He then uses the two worksheets depicted in figures 3-15 and figure 3-17 in his step 5: Failure Mode and Effects Analysis, for the remainder of the FMEA analysis. He now uses the segmentation of the system that he did previously (refer to paragraph 3.2.3) to show the relationship between the various functional failures and the equipment/components using the Equipment-functional failure matrix (figure 3-15). He specifically refrains from creating any connection between sub-assemblies/components and functions earlier during the process and only now identifies the connection between hardware and functional failures. This very important distinction should, in most cases, lead to a better RCM analysis - this follows from the previous discussion regarding the difference between *preserve function* and *preserve equipment* (see page 3-4).

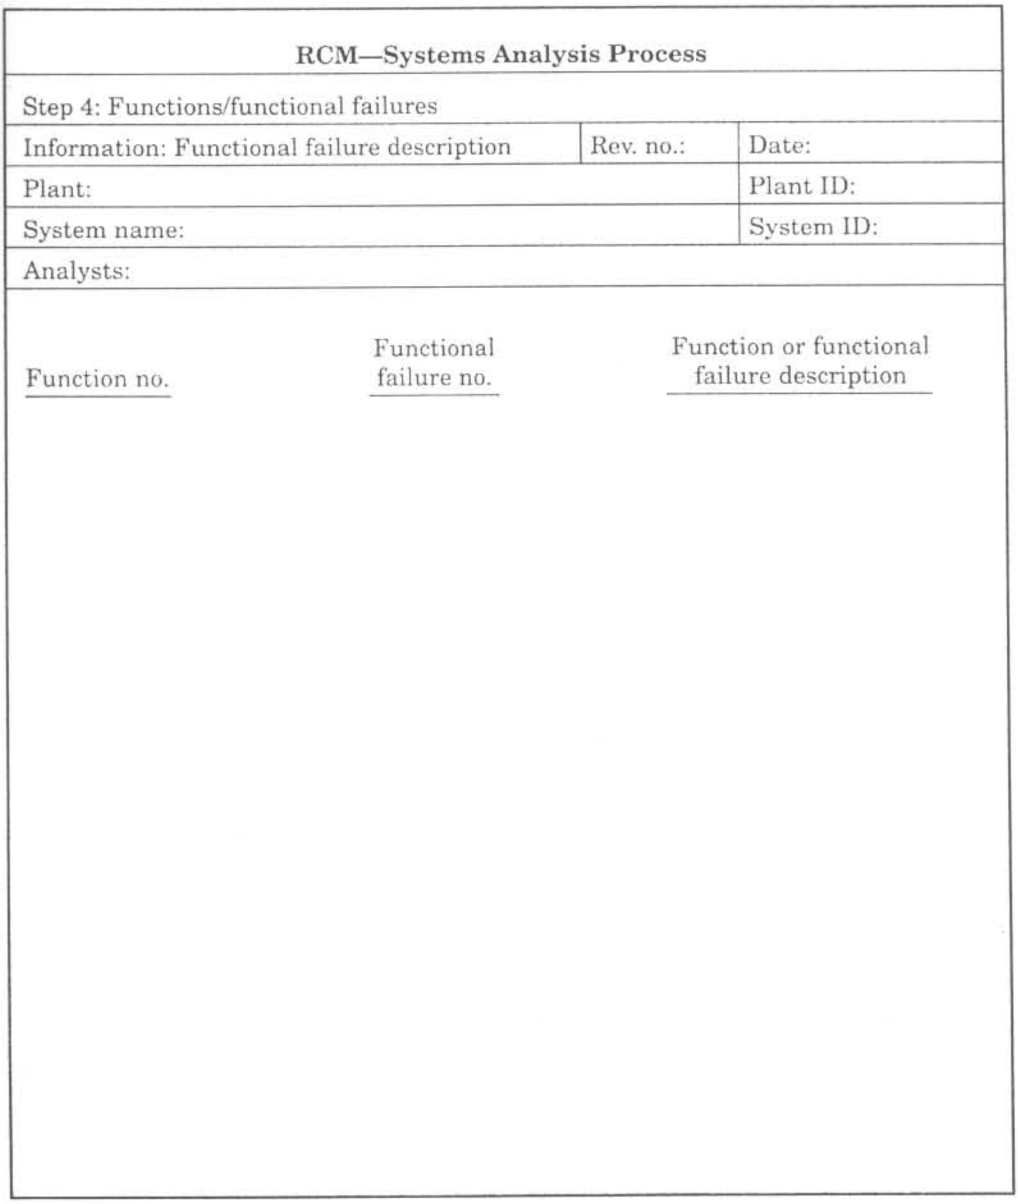

| | | |
|---|---|---|
| RCM—Systems Analysis Process | | |
| Step 4: Functions/functional failures | | |
| Information: Functional failure description | Rev. no.: | Date: |
| Plant: | | Plant ID: |
| System name: | | System ID: |
| Analysts: | | |
| Function no. | Functional failure no. | Function or functional failure description |

Figure 3-14: Functions/Functional Failure Worksheet [Smith (1993)]

The FMEA Worksheet used by Smith (figure 3-17) is to a large extent compatible with the standard FMEA Worksheet as depicted in figure 3-7. The *first* difference is that he uses one sheet per functional failure, on which he then lists the various Component | Failure Mode | Cause combinations. The *second* difference is that he now extends the <item | function | functional failure | failure mode> or <item | function | failure mode | cause> (page 3-26) to <item | function | functional failure | failure mode | cause>. However, the reason for this is fairly clear – he uses the term 'functional failure' for a collection of smaller functional failures and then uses 'failure mode' for the real functional failure and 'cause' for the failure mode. The way in which the various authors confuse these terms and their use must of necessity be very confusing to the average user of RCM. A part of this problem of course lies with the difference in use of these terms between MIL-STD-1629A and that of

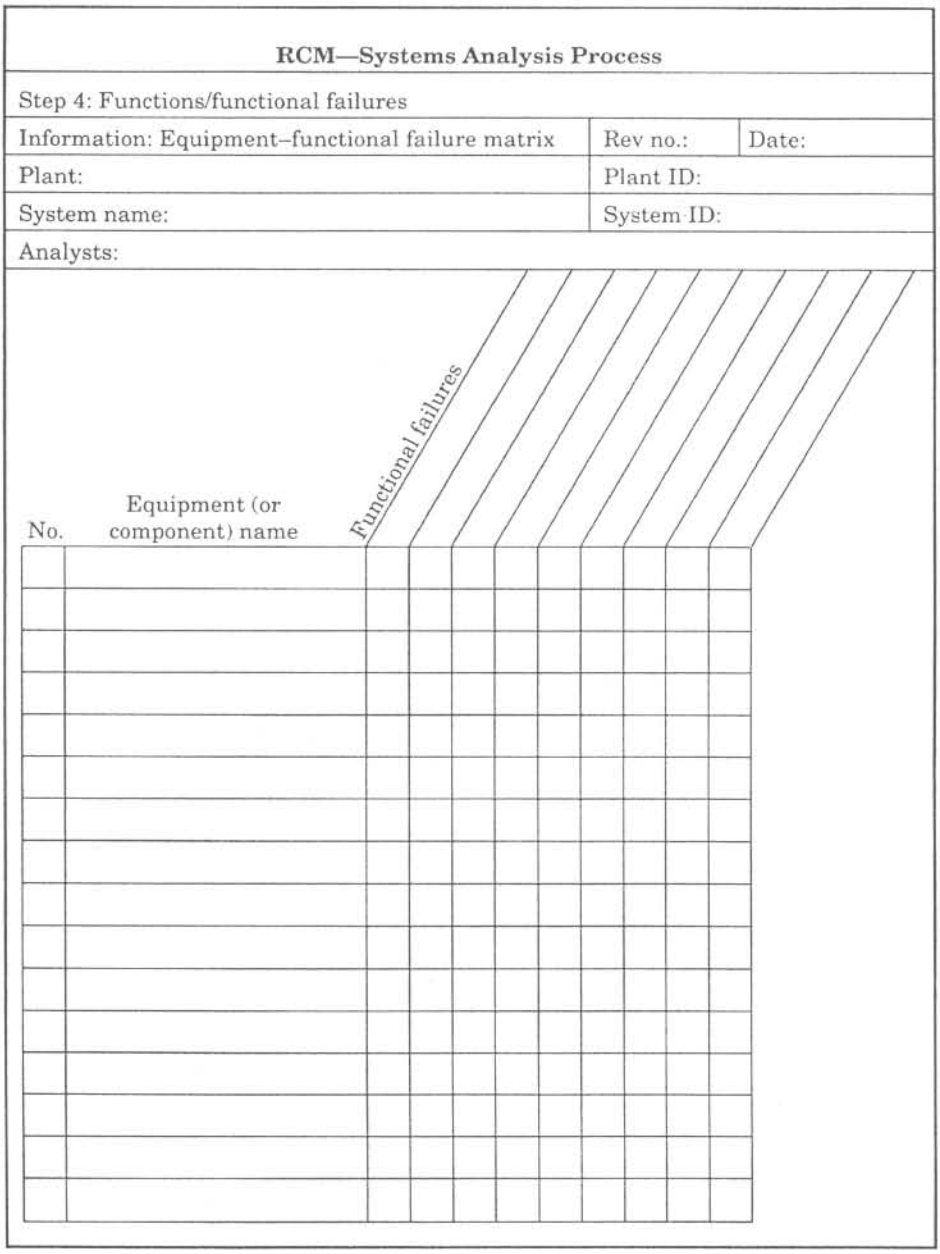

**RCM—Systems Analysis Process**

Step 4: Functions/functional failures

| | | |
|---|---|---|
| Information: Equipment–functional failure matrix | Rev no.: | Date: |
| Plant: | Plant ID: | |
| System name: | System ID: | |
| Analysts: | | |

| No. | Equipment (or component) name | Functional failures | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |
| | | | | | | | | | |

Figure 3-15: Equipment-functional failure matrix [Smith (1993)]

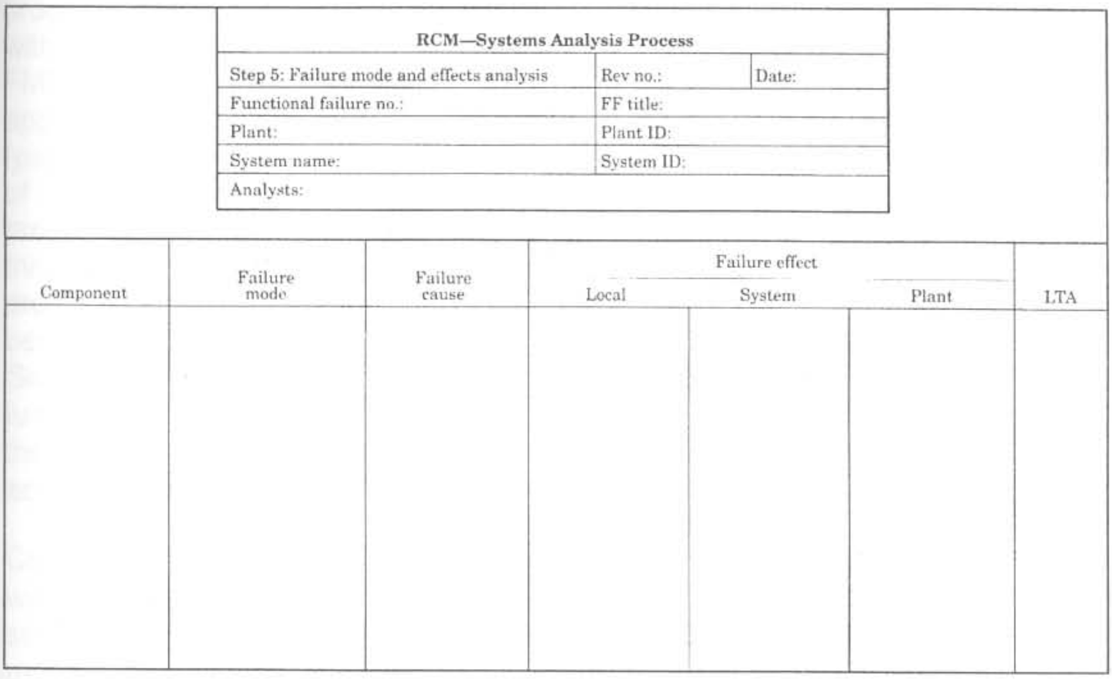

| | | |
|---|---|---|
| RCM—Systems Analysis Process | | |
| Step 5: Failure mode and effects analysis | Rev no.: | Date: |
| Functional failure no.: | FF title: | |
| Plant: | Plant ID: | |
| System name: | System ID: | |
| Analysts: | | |

| Component | Failure mode | Failure cause | Failure effect | | | LTA |
|---|---|---|---|---|---|---|
| | | | Local | System | Plant | |
| | | | | | | |

Figure 3-17: FMEA Worksheet [Smith (1993)]

standard RCM, so that the various authors are not fully to be blamed. The problem really lies in the definition (or lack thereof) of the different concepts. *Thirdly*, the FMEA Worksheet (figure 3-17) does not have columns for identifying different operational modes, failure detection methods and the severity class. The first of these (lack of columns) can be a severe restriction in military applications, while the others are more specific to the way in which Smith structures his version of RCM.

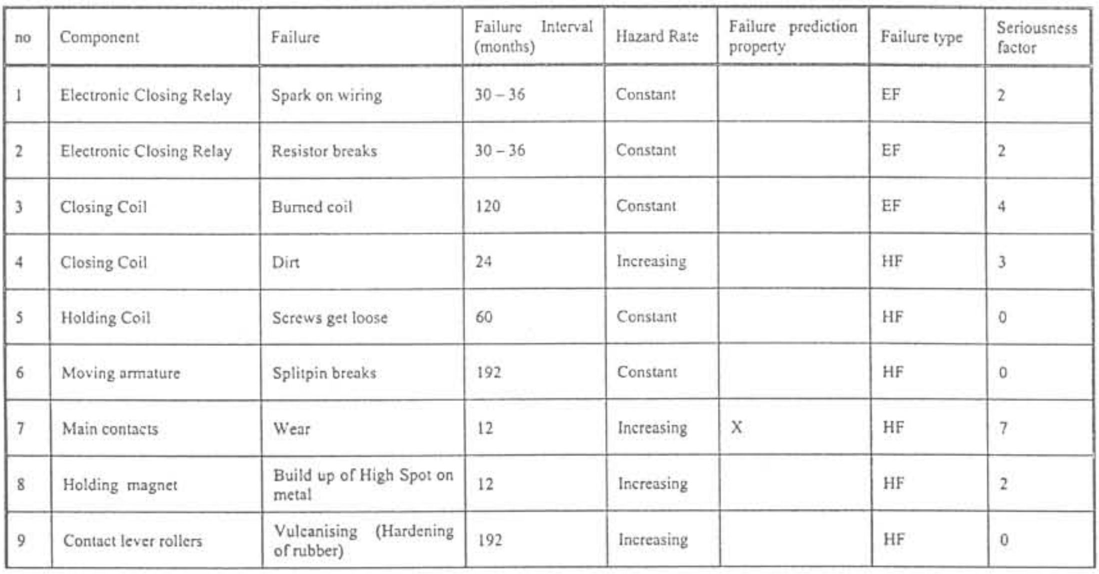

| no | Component | Failure | Failure Interval (months) | Hazard Rate | Failure prediction property | Failure type | Seriousness factor |
|---|---|---|---|---|---|---|---|
| 1 | Electronic Closing Relay | Spark on wiring | 30 – 36 | Constant | | EF | 2 |
| 2 | Electronic Closing Relay | Resistor breaks | 30 – 36 | Constant | | EF | 2 |
| 3 | Closing Coil | Burned coil | 120 | Constant | | EF | 4 |
| 4 | Closing Coil | Dirt | 24 | Increasing | | HF | 3 |
| 5 | Holding Coil | Screws get loose | 60 | Constant | | HF | 0 |
| 6 | Moving armature | Splitpin breaks | 192 | Constant | | HF | 0 |
| 7 | Main contacts | Wear | 12 | Increasing | X | HF | 7 |
| 8 | Holding magnet | Build up of High Spot on metal | 12 | Increasing | | HF | 2 |
| 9 | Contact lever rollers | Vulcanising (Hardening of rubber) | 192 | Increasing | | HF | 0 |

Figure 3-16: FMEA Worksheet [Le Clercq and Van den Broek (1999)]

Gits (1984) (p.99), when discussing his ‘satisficing approach’, comments that his way of ‘part classification’ has ‘far reaching analogies’ with FMEA. He never, in any of his referenced work, explains this part of his methodology properly, but in Gits (1988) he shows an example, which can be combined with the work of Le Clercq and Van den Broek (1999) to get an idea of the FMEA analysis used. These two examples are practically identical in approach and a small part of table 4.2 of Le Clercq and Van den Broek (1999) (page 35) is shown here (figure 3-16). This is very much a simplified version of the standard FMEA – when compared with the RCM standardised procedure of analysis of <item | function | functional failure | failure mode>, this only has two of the analysis steps, i.e. <item | failure mode>. The problem with this is, *firstly*, that the idea of *preserve function* (page 3-4) is lost because the analysis works only with hardware items and no functions. *Secondly*, this leads to a mixture of functional failures and failure modes in the functional failure column and *thirdly*, because of the extreme simplification, there is no certainty regarding whether all failure possibilities have been accounted for.

Columns 4, 5, 6 and 7 in the worksheet above are only used in later steps and will be referred to then. The seriousness column (column 8) is used in a way similar to the severity class in the standard FMEA worksheet (figure 3-7), but using a purpose-made classification system. One of the seriousness factor tables used by Le Clercq and Van den Broek (1999) (table 4.1, p. 35) is shown in figure 3-18. This table is used to allocate a seriousness factor (weight) to each failure based on its combination of failure effects on input/output variables, such as functionality, costs, manpower needed, quality impact, safety impact and secondary damage. This is a good, yet somewhat simplistic method of assessing risk factors that can be used with good effect to enhance the RCM methodology. However, this should be combined with the full functional analysis and failure effects estimation of the full FMEA methodology.

| Failure consequence | Weight factor | | | |
|---|---|---|---|---|
| | 0 | 1 | 2 | 3 |
| 1. Functional | None | Tripping | HSCB won't close | N/a |
| 2. Material costs | Low | Medium | High | N/a |
| 3. Workload needed | ≤1 hour | 1 ≤ 2 hours | ≥ 2 hours | N/a |
| 4. Safety | No | N/a | N/a | Yes |
| 5. Secondary Damage | No | N/a | Yes | N/a |

**Figure 3-18: Seriousness Factor Table**

The Business Centred Maintenance (BCM) approach of Kelly (1997) does not perform its analysis through a formal study of item functionality, and the failure modes that cause functional failures. It rather relies on either the manufacturer's unit life plan or, if that is not available, an engineering study of the item and its known failures.

The MSG-3 document of the Air Transport Association of America (ATA) [MSG-3 (1993)], differs from Nowlan and Heap (1978) in that it does not specify the use of the FMEA at all. What it does is to specify an 'FMEA like' approach where the following must be identified for each MSI: *firstly* functions (defined as the normal characteristic actions of the item), *secondly* functional failures (how the item fails to perform its function), *thirdly* failure effects (the result of the functional failure) and *fourthly* failure causes (why the functional failure occurs). This is similar to the item | function | functional failure | failure mode approach of Moubray (1991) and other authors, where *failure causes* ≡ *failure modes*. The problem with the MSG approach is that it does not provide any specific guidelines regarding the process of identifying failure modes. It probably means to suggest that a technique such as the FMEA or FMECA or equivalent be used. The premise presumably is not to prescribe the specific technique, but to allow freedom of choice. In the airline industry, this is possibly a valid approach, but in general industry, with its dearth of failure analysis expertise, one has to provide a much more detailed 'recipe'.

The version of RCM used for the maintenance programs of USA army aircraft is based on the original (1979) version of MSG-3 [Anderson and Neri (1990), p.19]. Nevertheless, it does specify the use of an FMECA as a basis of understanding the failure process [p.16 op. cit.].

### 3.2.5. Prioritisation of Failure Modes

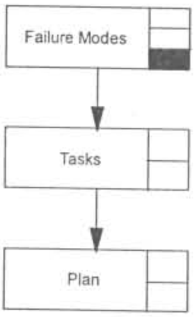

The original implementation of RCM [Nowlan and Heap (1978)] did not address the problem of the relative criticality of failure modes at all. That is, apart from the classification of failure modes (paragraph 3.2.6), which differentiates between different failure modes on the basis of their relative criticality, but never attempts to assign an absolute criticality value to individual failure modes. Such value is indispensable in applying limited preventive resources to a system (especially when combined with the prioritisation suggested by Coetzee (1997/2), described in paragraph 3.2.2, page 3-11). Both versions of the FMEA technique described above (paragraph 3.2.4) includes a form of criticality allocation:

- MIL-STD-1629A (1980) allocates a severity class to each failure mode, based on the severity scale (repeated from paragraph 3.2.4):
  - Category I – Catastrophic – a failure that may cause death or system loss.
  - Category II – Critical – a failure that may cause severe injury, major property damage, major system damage, which will result in mission loss.
  - Category III – Marginal – a failure that may cause minor injury, minor property damage, minor system damage, which will result in delay, mission degradation.
  - Category IV – Minor – a failure that will not cause injury, property damage, system damage, bur which will result in unscheduled maintenance or repair.

- McDermott et al (1996) uses a Risk Priority Number (RPN), which is the product (through multiplication) of three factors (repeated from paragraph 3.2.4):
  - Severity ≡ Consequence of failure. Rating is done on a scale of 1 (not severe at all) to 10 (extremely severe). Each effect of the failure mode is given an own severity rating, even if there are several effects for a particular failure mode. The authors suggest a generic rating scale – see figure 3-10
  - Occurrence ≡ Probability of failure occurring (or the frequency at which it occurs). Rating is done on a scale of 1 (probability virtually nil) to 10 (probability extremely high). The authors suggest a generic rating scale – see figure 3-11.
  - Detection ≡ Probability of failure being detected. Rating is done on a scale of 1 (probability extremely high) to 10 (probability virtually nil). The authors suggest a generic rating scale – see figure 3-12.

As discussed in paragraph 3.2.4 above, Gits (1988) uses a column (column 8) in his FMEA-type worksheet (figure 3-16) in a way similar to the severity class in the standard FMEA worksheet (figure 3-7), but using a purpose-made classification system. One of the seriousness factor tables used by Le Clercq and Van den Broek (1999) (table 4.1, p. 35) is shown in figure 3-18. This table is used to allocate a seriousness factor (weight) to each failure based on its combination of failure effects on input/output variables, such as functionality, costs, manpower needed, quality impact, safety impact and secondary damage.

Additional to the above-mentioned criticality measures, many practitioners of RCM uses a prioritisation of failure modes through the criticality analysis afforded by the extension of the FMEA to a full FMECA (Failure Mode, Effects and Criticality Analysis). One of the more important standard works on FMECA is again MIL-STD-1629A (1980), that describes the purpose of a criticality analysis as follows:

> "*The purpose of the criticality analysis (CA) is to rank each potential failure mode identified in the FMEA Task 101, according to the combined influence of severity classification and its probability of occurrence based upon the best available data.*"

The FMECA supports two ways of establishing the criticality of a failure mode. When no data regarding the failure rate of the item is available, it suggests a qualitative approach based on a standard probability of occurrence scale as follows:

- Level A – Frequent. A single failure mode probability larger than 20% of the total probability of failure of the item during the operating time interval.
- Level B – Reasonably probable. A single failure mode probability larger than 10% but less than 20% of the total probability of failure of the item during the operating time interval.

- Level C – Occasional. A single failure mode probability larger than 1% but less than 10% of the total probability of failure of the item during the operating time interval.
- Level D – Remote. A single failure mode probability larger than 0,1% but less than 1% of the total probability of failure of the item during the operating time interval.
- Level E – Extremely Unlikely. A single failure mode probability less than 0,1% of the total probability of failure of the item during the operating time interval.

On the other hand, if sufficient data regarding the failure rate of the item is available, it suggests a quantitative approach based the combination of several parameters, as follows:

- Failure effect probability $\beta$ - referring back to the FMEA, the $\beta$ value is the conditional probability that the failure effect will result in the allocated severity classification, given that the failure mode occurs. The $\beta$ value thus represents the judgement of the analyst regarding the conditional probability that the loss will actually occur. This judgement is based on the scale:

| | |
|---|---|
| Actual loss: | $\beta$=1,00 |
| Probable loss: | 0,1<$\beta$<1,00 |
| Possible loss: | 0<$\beta$<0,1 |
| No effect: | $\beta$=0 |

- Failure mode ratio $\alpha$ - the fraction of the item failure rate attributable to the particular failure mode (i.e. the probability that the item will fail in the specific failure mode, expressed as a decimal fraction).
- The item failure rate $\lambda_p$ (1/MTTF).
- Operating time t (in hours) or the number of operating cycles of the item per mission.

The failure mode criticality number $C_m$ is then calculated as the specific failure mode's contribution to the total criticality of the item, after which the total item criticality $C_r$ is calculated by adding the $C_m$ values for all the n failure modes of the item:

$$C_m = \beta\alpha\lambda_p t$$

$$C_r = \sum_{i=1}^{n}(C_m)_i$$

The standard format for the criticality analysis is presented in figure 3-19, which is self-explanatory, as it is based on a logical extension of the FMEA worksheet presented in figure 3-7.

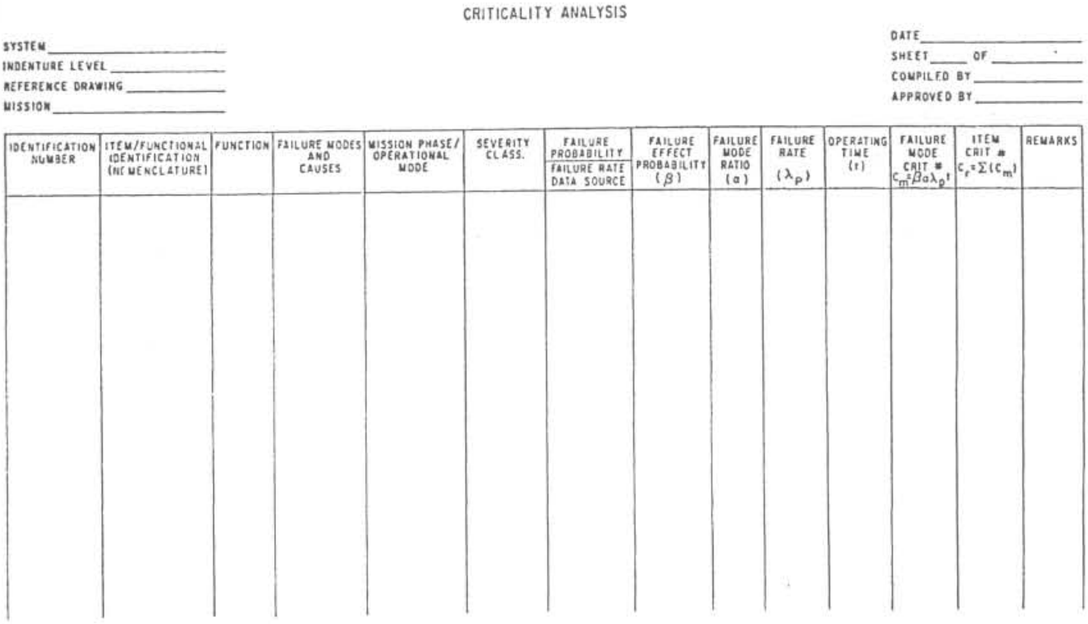

CRITICALITY ANALYSIS

SYSTEM ____________
INDENTURE LEVEL ____________
REFERENCE DRAWING ____________
MISSION ____________

DATE ____________
SHEET ____ OF ____
COMPILED BY ____________
APPROVED BY ____________

| IDENTIFICATION NUMBER | ITEM/FUNCTIONAL IDENTIFICATION (NOMENCLATURE) | FUNCTION | FAILURE MODES AND CAUSES | MISSION PHASE/ OPERATIONAL MODE | SEVERITY CLASS. | FAILURE PROBABILITY<br>FAILURE RATE DATA SOURCE | FAILURE EFFECT PROBABILITY $(\beta)$ | FAILURE MODE RATIO $(\alpha)$ | FAILURE RATE $(\lambda_p)$ | OPERATING TIME $(t)$ | FAILURE MODE CRIT # $C_m = \beta\alpha\lambda_p t$ | ITEM CRIT # $C_r = \Sigma(C_m)$ | REMARKS |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | | | | |

Figure 3-19: Criticality Analysis Worksheet

The criticality numbers ($C_m$ and $C_r$) as calculated above (or the probability of occurrence levels in the case of the qualitative approach) can be used directly to prioritise items/failure modes. However, the *criticality matrix* presented in figure 3-20 gives a better way of prioritisation. It provides the analyst with a template that can be used to prioritise by using both the criticality numbers (or probability of occurrence) and the *severity classification* allocated in the FMEA in a single prioritisation action. The matrix is constructed by inserting item or failure mode numbers in matrix locations representing the severity classification and either the probability of occurrence or the criticality number (whichever applies). The completed matrix then gives a visual representation of the relative priorities of the various items/failure modes.

Jones (1995) incorporates the concept of risk into RCM, which gives a criticality measure, but also gives an extra dimension, which can be used to reduce risk over the medium to long term. He defines risk as:

$$R_i = P_i \times C_i$$

$R_i$ = Risk of the $i^{th}$ risk consequence factor

$P_i$ = Probability of occurrence of the $i^{th}$ risk factor

$C_i$ = Consequence of occurrence of the $i^{th}$ risk factor

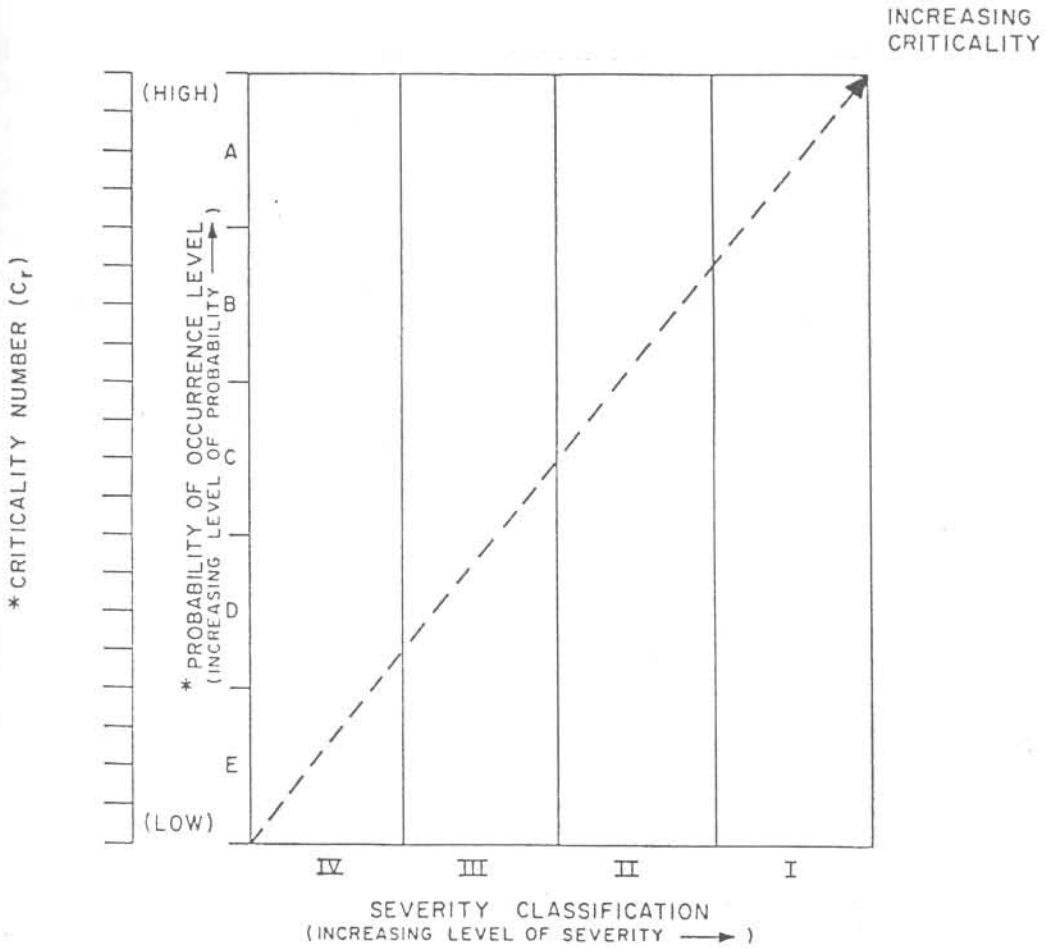

* NOTE: BOTH CRITICALITY NUMBER ($C_r$) AND PROBABILITY OF OCCURRENCE LEVEL ARE SHOWN FOR CONVENIENCE.

Figure 3-20: Criticality Matrix Template

The typical consequence factors considered are:

Risk consequence factor 1: **Safety**
Risk consequence factor 2: **Lost Production**
Risk consequence factor 3: **Lost Quality**
Risk consequence factor 4: **Environmental Effects**
Risk Consequence factor 5: **Maintenance**

The total risk for the combined consequence factors is given by:

$$R_t = \sum_{i=1}^{5} R_i = \sum_{i=1}^{5} \left(P_i \times C_i\right)$$

As is the case with the other criticality measures presented above, the total risk can be calculated per failure mode and totalled per item. This gives a quantitative measure which is based on the same premise as the McDermott et al (1996) Risk Priority Number (RPN)[2] and similar to that of FMECA [MIL-STD-1629A (1980)], but from a different perspective. In the case of FMECA, the probability of occurrence of a single failure effect ($\beta\alpha\lambda_p$) is multiplied with the consequence in terms of duration of mission or mission phase and then totalled over all the failure effects to give the total criticality (or risk) per item. In the present case the probability of occurrence of multiple failure effects are used to calculate a total risk per failure mode and per item in common consequence terms, which could be mission time in military applications and which will certainly be monetary terms in industrial applications.

Another benefit of the risk approach, which should be shared by FMECA, is that the risk can be plotted on a Consequence-Probability matrix. This affords the possibility of managing the risk over the medium to long term by either reducing the consequence severity or the probability of occurrence or both. The possibilities are shown in figure 3-21, which shows the different risk change options – the preference for risk change should be in the order direction III | direction II, direction IV, direction I (direction I is one of negative growth). Figure 3-22 shows iso-risk lines on the same set of axes, while figure 3-23 depicts a typical two-year failure mode risk reduction result.

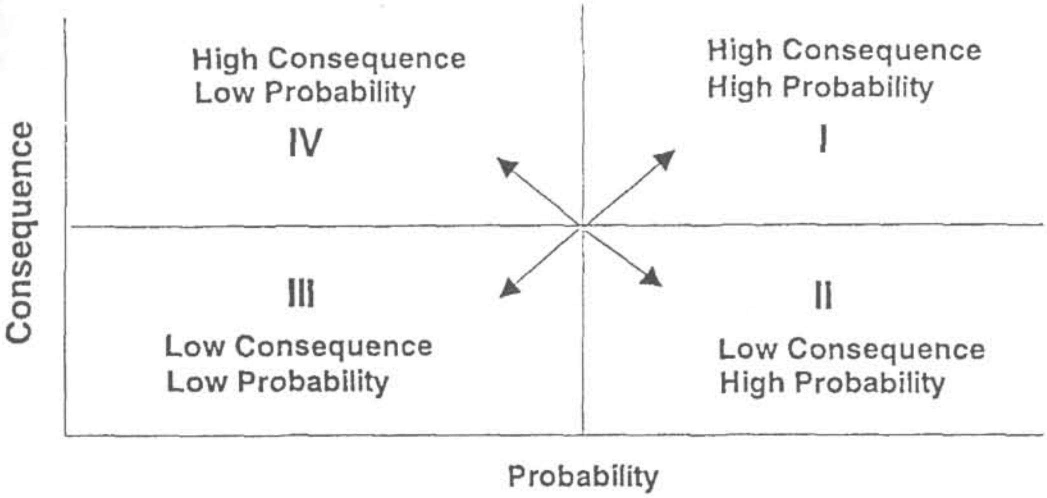

igure 1 Risk change possibilities

Kelly (1997) does not analyse at the level of failure modes (paragraph 3.2.4) and thus failure mode criticality is not relevant in the case of his Business Centred Maintenance.

This facet is not addressed at all by MSG-3 (1993). On the other hand Anderson and Neri (1990), p. 16, prescribes the use of FMECA, although their RCM approach for US army aircraft is based on MSG-3.

---

[2] In the case of McDermott et al (1996) the Risk Priority Number (RPN) also includes a third factor, the possibility of detection. On the other hand, it does not take into account the five consequence areas suggested by Jones (1995).995).I I

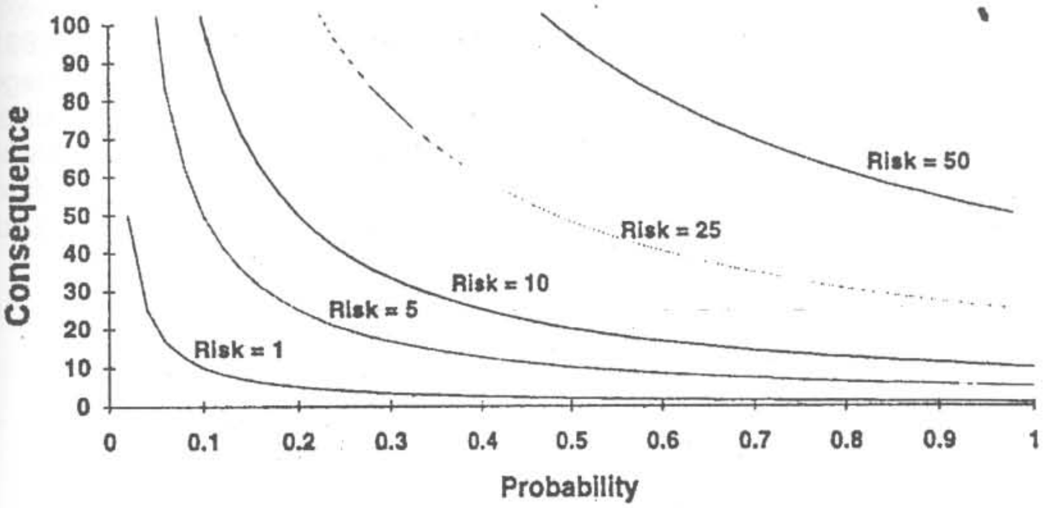

Figure 3-22: Constant risk contours

Although the methods presented by Gits (1988) and McDermott et al (1996) has a contribution to make and will be preferred by many users and practitioners of RCM, they are really too simplistic in nature and one will have to base any improved methods on FMECA [MIL-STD-1629A] and the risk calculations of Jones (1995).

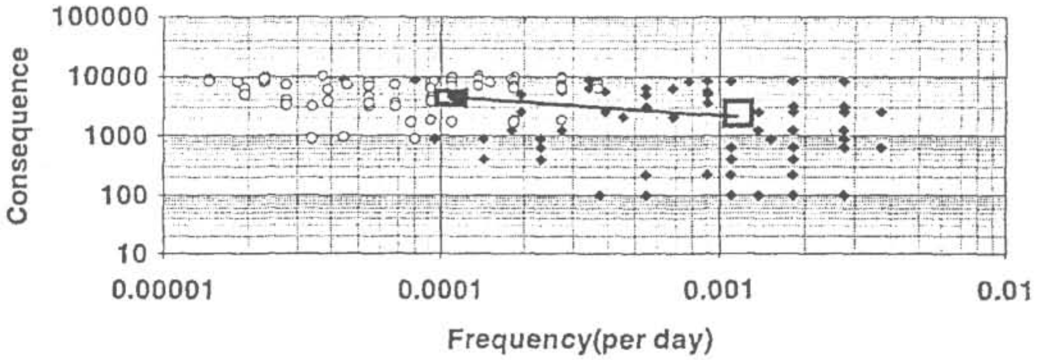

Figure 3-23: Two-year risk reduction result

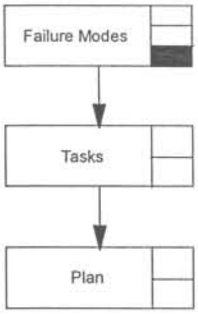

3.2.6. Classification of Failure Modes

The ‘new approach’ to scheduled maintenance that was developed during the late 60’s of the previous century had its foundation at United Airlines: “In 1966, an internal memorandum at United suggested that a decision tree approach, … might provide a more orderly and more objective means for designing preventive maintenance programs." [Matteson (1989)]. This basic decision tree was further developed to the one that, together with the segmentation process (paragraph 3.2.2) and task selection decision tree (paragraph 0) formed the basis of the original RCM approach of Nowlan and Heap (1978). This decision tree [Nowlan and Heap (1978) exhibit 4.2, p.88], which is shown in figure 3-24, forms the basis of the classification of failure modes in RCM.

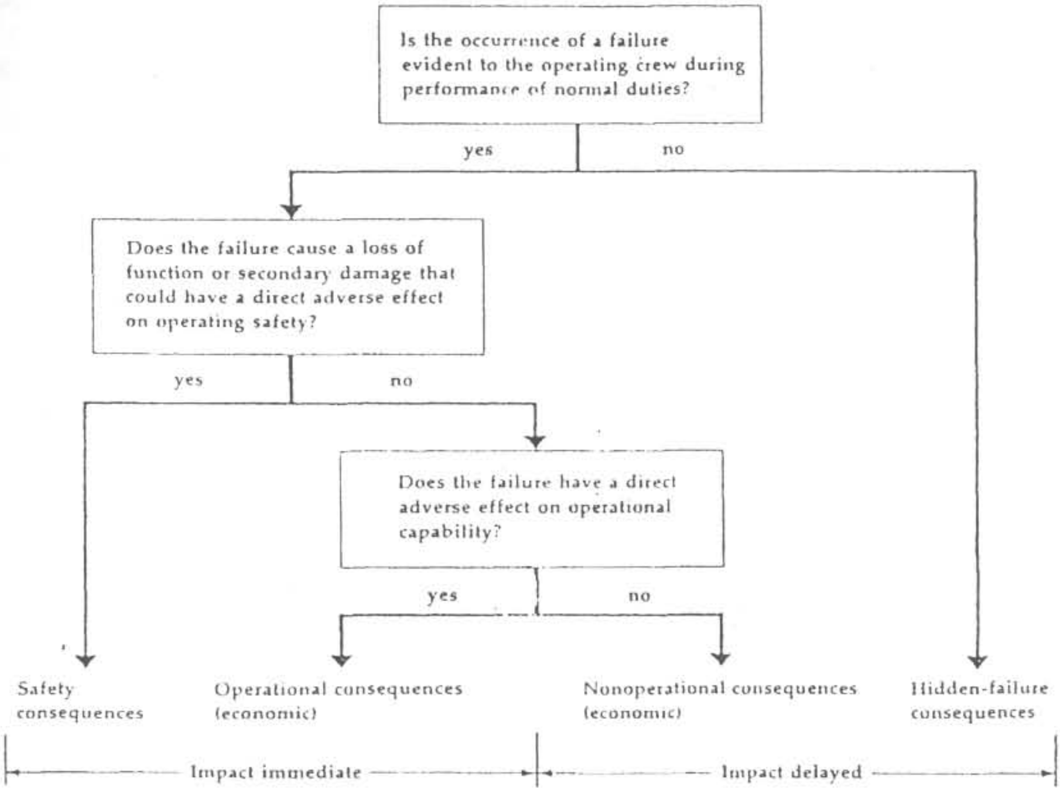

**Figure 3-24: Failure Consequence Decision Tree [Nowlan and Heap (1978)]**

The original version of RCM classified failure modes into four failure consequence classes:

- Hidden failures – these are failures that will not have an effect on the system performance at all, but which, if it remains undetected, can lead to a multiple failure. A typical example of a hidden failure is the failure of a non-failsafe safety device – if such a device fails, the failure will typically not be detected. When the function which is protected by the device misbehaves and the device is in the failed state, the protected function will fail, mostly catastrophically (a so-called multiple failure).
- Failures with safety and/or environmental consequences – this includes all failures that will cause danger to people, the environment and property.
- Failures with operational consequences – failures that will result in partial or total loss of operational capability, resulting in loss of revenue.

- Failures with non-operational consequences – failures that do not affect operational capability, but are costly due to the impact of repair cost.

The importance of this top structure classification of failure modes is that the task selection process differs depending on the failure consequences. The structure of the decision tree discussed above is also based on the relative criticality of the four classes of failure consequences. The most critical of the consequence categories is that of hidden failures. The relative criticality then progresses downward from hidden -> safety and environmental -> operational -> non-operational. The following short descriptions highlights this difference between the four failure consequence classes:

- Hidden failures, as previously defined, have a safety impact, mostly with a risk of multiple failure (often of catastrophic nature). For a preventive task to be the acceptable strategy in this category, it must reduce the risk of a multiple failure to an acceptably low level. If a preventive task is not found that reduces the risk to a low enough level, the default strategy is to specify a failure finding task or, if that is not effective, to redesign if a multiple failure can negatively affect safety or the environment. This (and the next) categories are those for which economics play a lesser role due to the possible loss of life, property and permanent environmental damage.
- Safety and Environmental Consequences includes all failures that will cause danger to people, the environment and property. For a preventive task to be the acceptable strategy in this category, it must reduce the risk of failure totally or to a very (acceptably) low level. If a preventive task is not found that reduces the risk to a low enough level, the default strategy is to redesign, as a compromise in this category is not acceptable. As was stated above this is the second category where economics play a lesser role due to the possible loss of life, property and permanent environmental damage.
- Operational Consequences - this and the next category (non-operational consequences) are the two categories for which the consequences are primarily of an economic nature. For a failure mode to have operational consequences, it must affect the production output of the business negatively. It thus causes a production loss, with an accompanying loss in sales and thus profit. Each time that a failure thus occurs, money is lost due to production being lost, as well as due to the cost of repairing the failure. For a preventive task to be an acceptable strategy, it must firstly be technically feasible in preventing a failure mode with operational consequences. That is, it must reduce the risk of failure to a low enough level so that the benefits of the preventive action (less production lost and lower breakdown costs) are worth the cost of implementing it. Secondly, it must also be economically feasible. The added cost due to the implementation of the task must be lower than the benefit in terms of reduction of lost production and cost of failure. If a preventive task is not found that is both technically and economically feasible, the default strategy is to repair the failure mode only after failure (corrective maintenance), with redesign as an option.

- Non-operational consequences - This category is very similar to the previous one (operational consequence) in that its consequences are also primarily of an economic nature. Its difference lies in the fact that it does not affect the production output of the business negatively. Its negative economic effect is thus limited to the cost incurred in repairing the failure each time that such failure occurs. As is the case for operational consequences, a preventive task will be an acceptable strategy if it is both technically feasible in preventing the failure mode and economically feasible (the added cost due to the implementation of the task lower than the benefit in terms of reduction of the cost of failure). If a preventive task is not found that is both technically and economically feasible, the default strategy is to repair the failure mode only after failure (corrective maintenance) with redesign as an option.

The original Nowlan and Heap approach suggested that the *first* question (whether failure is evident) in the tree be asked for each function of an item and the *second* (operational safety) for each functional failure and for each failure mode (pp.87/88). The reason for this differentiation is not clear and practically all of the later practitioners use each of the three questions in the decision tree per failure mode. Examples are Moubray (1991), pp.71and87, Smith (1993), p. 89 and Coetzee (1997/2), p.98 (also refer to table 3.3 below). This is a more conservative approach that will ensure that each failure mode (being one of the many failure possibilities of a typical item) is classified correctly according to its specific failure consequences.

The decision diagram of figure 3-24 has weathered the further development of the methodology well. Most of the later authors use it with only a few adaptations in terminology and detail changes in the three questions to promote clarity (apart from cosmetic changes). One of the more fundamental changes was the addition of the environmental consequence sub-class to what then became known as the safety *and environmental* consequence class by Moubray (1991). The detailed questions used by the various authors are shown in table 3.3. It is clear that, although no fundamental changes were made to the principle, there was an attempt towards producing more clarity in the questions.

MSG-3 (1993), on the other hand, has made a more fundamental change in introducing a fourth question to differentiate between hidden failures with a safety impact (the traditional stand) and those with an economic impact (bottom line in table 3.3). This is shown in figure 3-25, which is presented in a slightly different format than that of MSG-3, but with the same logic and wording as MSG-3. This change, which is a severity sub-classification of hidden failures, has, apart from Smith (1993), not been taken up in the commercial applications of RCM. MSG-3 also substituted the term 'consequence' by 'effect', such that the result of failure is called its 'failure effect' instead of the usual 'failure consequence'. This change could have been well accepted by the RCM user population, apart from the fact that it is in conflict with the terminology of FMEA and FMECA, where the 'E' represents a 'failure effect' with a different meaning. The old terminology should thus preferably be retained.

**Table 3.3: Comparison of Failure Consequence questions**

| Consequence class | Author | Question |
|---|---|---|
| Hidden | Nowlan and Heap (1978) | Is the occurrence of a failure evident to the operating crew during performance of normal duties? |
| | Moubray (1991) | Will the loss of function caused by this failure mode on its own become evident to the operating crew under normal circumstances? |
| | Smith (1993) | Under normal conditions, do the operators know that something has occurred? |
| | Coetzee (1997/2) | Is the loss of function due to the occurrence of this failure mode evident to the operations personnel? |
| | MSG-3 (1993) | Is the occurrence of a functional failure evident to the operating crew during the performance of normal duties? |
| Classification of hidden failures (safety/ economic) | MSG-3 (1993) | Does the combination of a hidden functional failure and one additional failure of a system related or back-up function have an adverse effect on operating safety? |

Harris (1985) divides the Operational Consequence category into two sub-categories, namely a 'mechanical performance' category and a 'process performance' category. His motivation for this is *firstly* that he defines the concept of 'chemical plant integrity', which he considers to be the sum of plant mechanical performance, plant process performance and plant financial performance. *Secondly*, he points out that the "the literature tends to largely ignore the relationship between maintenance and process performance" [Harris (1985), p. 47]. *Thirdly*, he considers the inclusion of a process performance category necessary "because the RCM technique does not explicitly account for the effect of functional failures on process performance." [Harris (1985), p. 255]. Examples of maintenance tasks following on process

**Table 3.3: Comparison of Failure Consequence questions (continued)**

| Consequence class | Author | Question |
|---|---|---|
| | Nowlan and Heap (1978) | Does the failure cause a loss of function or secondary damage that could have a direct adverse effect on operating safety? |
| | Moubray (1991) | Does the failure mode cause a loss of function or other damage, which could (i) hurt or kill someone or (ii) breach any known environmental standard or regulation? |
| Safety / Environment | Smith (1993) | Does this failure mode cause a safety problem? |
| | Coetzee (1997/2) | Does the failure mode cause a loss of function or secondary damage that could have a direct adverse effect on operating safety or the environment? |
| | MSG-3 (1993) | Does the functional failure or secondary damage resulting from the functional failure have a direct adverse effect on operating safety? |

performance degradation that he lists are the cleaning of boiler surfaces, the replacement of catalyst, overhaul of process pumps, cleaning of process piping and cleaning of heat exchangers. He adds that these are the more well-defined and quantifiable situations: "An additional dimension is added to the problem of relating maintenance to process performance when the less quantifiable effect of maintenance actions such as lubrication, servicing and minor adjustment is considered." [Harris (1985), p. 49]. Whether one should go to the extreme of adding an additional failure consequence category to the already well established four category consequence classification is an open question. It could further add to the confusion that the RCM user experiences. However, the concerns aired by Harris should possibly be addressed in the detailed definition of the operational consequence category.

**Table 3.3: Comparison of Failure Consequence questions (continued)**

| Consequence class | Author | Question |
|---|---|---|
| Operational | Nowlan and Heap (1978) | Does the failure have a direct adverse effect on operational capability? |
| | Moubray (1991) | Does the failure mode have a direct adverse effect on operational capability (output, quality, customer service or operating costs in addition to the direct cost of repair)? |
| | Smith (1993) | Does this failure mode result in a full or partial outage of the plant? |
| | Coetzee (1997/2) | Does the failure mode have a direct adverse effect on operational capability? |
| | MSG-3 (1993) | Does the functional failure have a direct adverse effect on operating capability? |

As was mentioned in the previous paragraph (paragraph 3.2.4), Kelly (1997) does not analyse at the level of failure modes and thus failure mode classification is not relevant in the case of his Business Centred Maintenance.

## 3.2.7. Task Selection

### *3.2.7.1. Introduction*

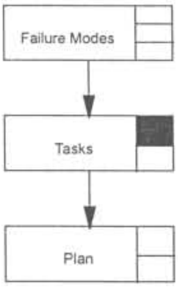

The task selection part of RCM is in a certain sense the heart of the methodology. Although all the previous are necessary to identify those failure modes that needs to be addressed, the task selection process is really where the constituent 'parts' of the maintenance plan are formulated. Again, the basis laid by Nowlan and Heap (1978) has withstood the test of time and is in most cases used in the same format to this day. MSG-3 (1993) has made certain changes in this area, but those are mostly in adding additional blocks of a minor nature (lubrication and pre-flight checks) to the front end of the original decision tree.

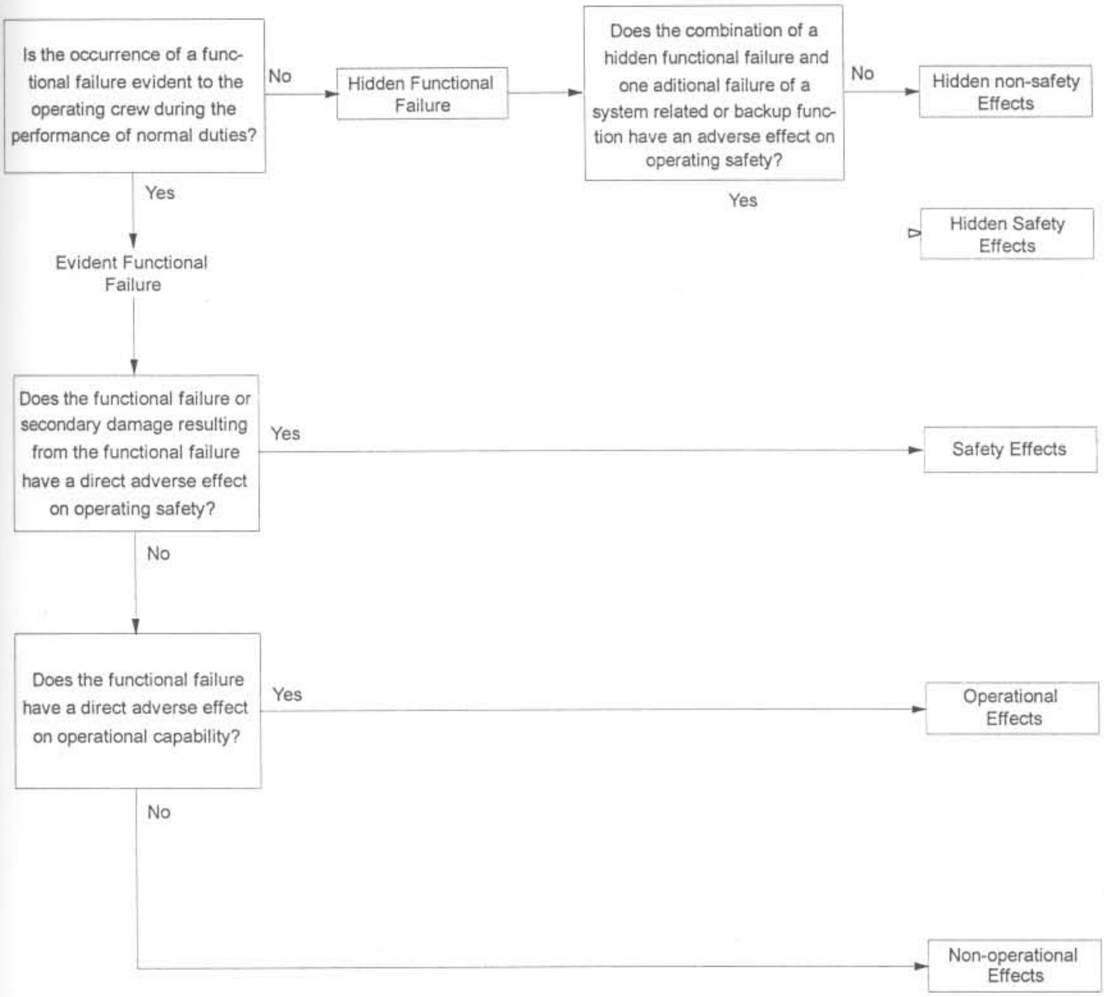

**Figure 3-25: Failure Effects Decision Tree [MSG-3 (1993)]**

### *3.2.7.2. Task types*

If one widens one's definition of what constitutes a maintenance task type, the original RCM text [Nowlan and Heap (1978)] suggests seven task types:

- On-condition task
- Rework task
- Discard task
- Failure Finding task
- Servicing and Lubrication task
- Redesign task
- Corrective Maintenance task

Apart from naming changes, none of the other authors on the subject made significant changes to this task structure:

(i) An On-condition task (inspection) is also called a Condition Monitoring task in general industry or a Predictive task. Coetzee (1997/2) retains the name On-condition task, Moubray (1991) calls it a Scheduled On-condition task, Smith (1993) a Condition-directed task and Kelly (1997) and Gits (1984) a Condition-based task. MSG-3 (1993) differentiates between two types of on-condition tasks, i.e. Inspection and Functional Checks. A third type, Operating Crew Monitoring, was included in the original (1979) version of MSG-3, but was removed as from MSG revision 1 (1987).

(ii) The name Rework task has not withstood the test of time. Moubray (1991) calls it a Scheduled Restoration task, Smith (1993) a Time-directed task (name shared with next category), Gits (1984) a Use-based task (name shared with next category), Kelly (1997) Fixed-time maintenance (name shared with next category), MSG-3 (1993) a Restoration task and Coetzee (1997/2) a Reconditioning task.

(iii) The name Discard task, on the other hand, did remarkably well. MSG-3 (1993) has kept it as is, while Moubray (1991) calls it a Scheduled Discard task. Smith (1993) calls it a Time-directed task (name shared with previous category), Gits (1984) a Use-based task (name shared with previous category), Kelly (1997) Fixed-time maintenance (name shared with previous category) and Coetzee (1997/2) a Replacement task.

(iv) The Failure-Finding task category is also retained by two of the major authors, i.e. Moubray (1991) and Smith (1993) using the exact name as proposed by Nowlan and Heap (1978) - the only major 'RCM' - publication that uses different terminology is MSG-3 (1993), which calls it a 'Check to verify operation'. Kelly (1997) and Gits (1988) (both non-RCM) also call it by different names. Kelly calls it Proof Testing and Gits State Inspection, but the idea is the same: an inspection must be done to find whether the function is still operative - this inspection can be visual or using instruments or through physically testing the function.

(v) Servicing / Lubrication task – this task type, although mentioned in most texts on the subject (see paragraph 3.2.7.1), never got formal attention in the strategy decision making process. MSG-3 (1993) is an exception and puts this decision as an extra one preceding even the Condition Based Maintenance decision. It seems as if authors of RCM texts agree that one will have something as mundane as servicing and lubrication, but that it needs not be part of the strategy-setting exercise.

(vi) Redesign task – most authors will immediately tell you that this is not a part of maintenance at all. This is of course true and is not to be argued with. However, in the process of the design of a maintenance plan one sometimes comes to a point where the correct strategy is that the particular failure mode should be designed out. The initiative then typically lies with the maintenance department to initiate (and sometimes manage) the design-out process. Thus, the name 'redesign task'. Most authors have kept the terminology as is, the exceptions being Smith (1993) with Design Modification and Coetzee (1997/2) with Design-out task.

(vii) Corrective Maintenance task – the alternative name for Corrective Maintenance conceived by Nowlan and Heap (1978) and followed by most authors is No Scheduled Maintenance. Smith (1993) calls it Fly to Failure, Kelly (1997) calls it Operate to Failure, while Gits calls it Failure Based Maintenance (FBM), while it seems as if MSG-3 regards it as non-existent in its task selection logic (although it refers to non-scheduled tasks under its listing of task types).

The list given above from Nowlan and Heap (1978) is nearly exhaustive. The first five of these task types can broadly be classified as being Preventive Maintenance [Smith (1993), p. 51], leading to the principle of three main task types, Preventive, Corrective and Design-out [Coetzee (1997/2), p.48]. Nevertheless, valid maintenance task types other than the ones listed above are listed by various authors and should be taken note of in the RCM process. These include Opportunistic Maintenance [Kelly (1997) and Coetzee (1997/2)] and Adjustment [Kelly (1997)]. One of the fads of the early to middle 90's was the so-called Pro-active Maintenance[3] [Fitch (1992)] – this is not really a new class of maintenance task as Fitch would have it, but is a way of thought [Coetzee (1997/2), p.51]. Nevertheless, it is an important aspect, which should be dealt with in RCM. Another aspect of importance mentioned by Kelly (1997) regards the maintenance task category repair. Although repair can include restoration and discard, often repair is used in an exclusive way. In other words, repair is when one brings an item back into operation by means other than full restoration or replacement. Mostly repair by definition does not include prevention, but is rather one of the ways of effecting Corrective Maintenance, which can be through reconditioning, replacement or repair.

### *3.2.7.3. Task Selection*

The standard task selection decision tree proposed by Nowlan and Heap (1978) is shown in figure 3-26. This decision tree is based on the fact that, in most instances, we do not know the full failure characteristics of a certain system or component. It consequently proposes a basic framework for the evaluation of maintenance tasks that is overly conservative in its order of maintenance task selection [Coetzee (1997/2), p.101]. The tasks are evaluated and selected in the order *on-condition* (condition based) task, *scheduled reconditioning* (restoration or overhaul) task, *scheduled replacement* (discard) task.

---

[3] Pro-active Maintenance was the name that Fitch (1992) gave to a strategy of measurement (similar to Condition Based Maintenance) - where the objective of the measurements in the Condition Based case is to detect the failure in time for preventive action to be taken, its objective in the Pro-active case is to discover failure modes that should be designed out.

Hence, if any one of the three basic tasks is selected, the remainder of the selection process is truncated. If one of these are not selected, the answer to the task selection process is *corrective maintenance* (or *no scheduled maintenance* or *replace only on failure* (r.o.o.f.)) - if this is not acceptable *design-out maintenance* is the only alternative left. The basic premise of the RCM technique is therefore to propose a condition based task as the first strategy option. Only if this is not viable will the technique propose either a scheduled reconditioning or scheduled replacement task. It so ensures that the task with the least risk of doing unnecessary (costly) work is selected in each case.

Again, as is the case with the consequence classes, table 3.3, the various authors use different question formats in asking the task selection questions. These are compared in table 3.4. The differences are as follows:

- Differences in maintenance task terminology – see paragraph 3.2.7.2.
- Differences in task effectiveness measures – see paragraph 3.2.7.4.
- Differences in purpose for use based tasks – reduce failure rate / reduce hazard rate / avoid failures. These important differences will have to be resolved - see chapter 4.
- Language preference.

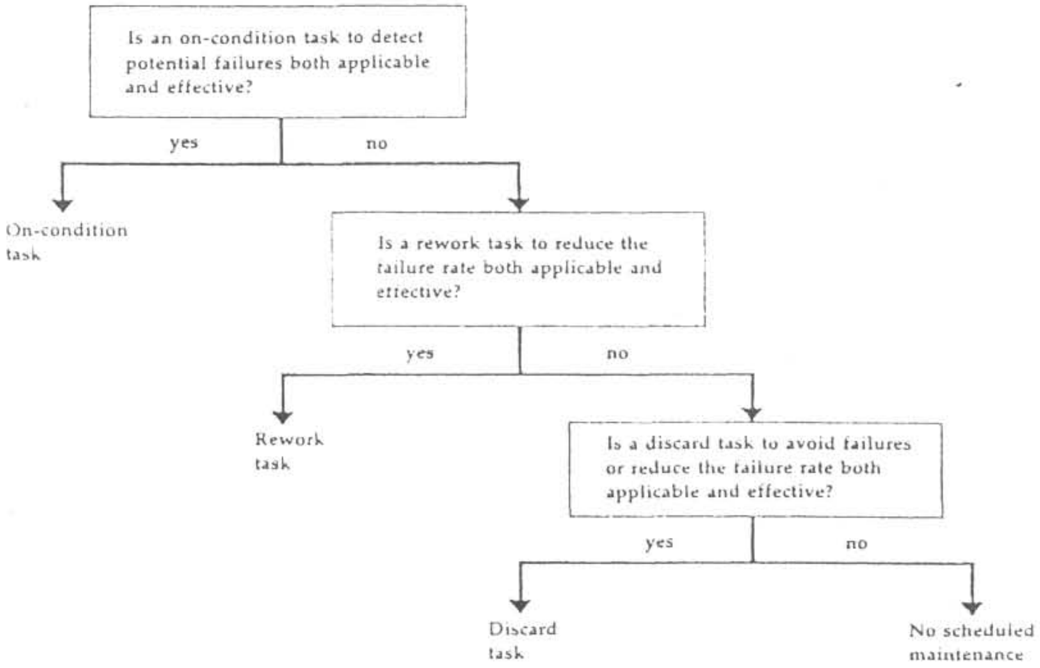

**Figure 3-26: Task selection decision tree [Nowlan and Heap (1978), p.90]**

Smith (1993) uses a similar but completely different looking concept in the choice of maintenance tasks, which will be considered in full when designing the new methodology in chapter 4. The method is shown in figure 3-27. The abbreviations used in the figure is as follows: CD ≡ Condition – directed; TD ≡

Time-directed; FF ≡ Failure Finding; Category C failure mode ≡ Minor to insignificant economic problem; Category D failure mode ≡ Hidden failure; Category D/C failure mode ≡ Hidden failure which only has minor to insignificant economic consequences. The differences between this and the standard decision tree are as follows:

- There is no formal distinction regarding the analysis of the four consequence classes, apart from the differentiation built into the decision tree itself. The standard tree affords a difference in handling the decisions for the four consequence classes – see paragraph 3.2.7.4. This can be a crucial difference in using / not using this approach.
- The decision tree is not truncated following task specification, as is the case with the standard tree. This typically results in multiple preventive tasks specified per failure mode of which one or more must be chosen which will achieve the required result. This is similar to the method of Gits (1988) – see discussion below.
- The order of questions has been drastically altered – because of the development of multiple task choices per failure mode. This is not necessarily negative, as the objective is to find the task with the best fit in terms of a combination of prevention and economy.
- It does not handle the compulsory redesign of the hidden and safety/environment consequence classes and the optional choice of redesign in the case of economic consequences well.

Gits (1988) uses a method similar to that of Smith above to develop a list of possible tasks based on its effectiveness. An example of such analysis is shown in figure 3-28 [Le Clercq and Van den Broek (1999), p.39]. The result of this method (compare with their FMEA – figure 3-16) is a list of task initiation possibilities (with the possibilities of task initiation being failure, age and condition). The result of such analysis is lists of the failures (represented by their relevant reference numbers) in each category. The following descriptions serve to shortly describe each category:

- Column 1: Hidden failures – task initiation when failure takes place – typical task specified Failure Finding.
- Column 2: Unimportant non-hidden failures (low seriousness factor) - task initiation when failure takes place – typical task specified Corrective Maintenance.
- Column 3: Important non-hidden failures (high seriousness factor) with increasing hazard rate and an identified failure indicating quantity – task initiation can take place based on any one of failure, use or condition – task choice based on efficiency (mostly cost) between these three possibilities.
- Column 4: Important non-hidden failures (high seriousness factor) with increasing hazard rate and no failure indicating quantity – task initiation can take place based on any one of failure and use – task choice based on efficiency (mostly cost) between these two possibilities.

Table 3.4: Comparison of task selection questions

| Task selection step | Author | Question |
|---|---|---|
| Lubrication/ Servicing | MSG-3 (1993) | Is a lubrication or servicing task applicable and effective? |
| Operational/ Visual Check | MSG-3 (1993) | Is a check to verify operation applicable and effective? |
| On-condition | Nowlan and Heap (1978) | Is an on-condition task to detect potential failures both applicable and effective? |
| | Moubray (1991) | Is a task to detect whether the failure is occurring or about to occur technically feasible and worth doing? |
| | Coetzee (1997/2) | Is an on-condition task to detect potential failures both technically and economically feasible? |
| | MSG-3 (1993) | Is an inspection or functional check to detect degradation of function applicable and effective? |
| Rework | Nowlan and Heap (1978) | Is a rework task to reduce the failure rate both applicable and effective? |
| | Moubray (1991) | Is a scheduled restoration task to reduce the failure rate[4] technically feasible and worth doing? |
| | Coetzee (1997/2) | Is a reconditioning task to reduce the hazard rate both technically and economically feasible? |
| | MSG-3 (1993) | Is a restoration task to reduce failure rate applicable and effective? |

---

[4] In the case of safety and environmental consequences, "reduce the failure rate" is replaced with "avoid failures".

**Table 3.4: Comparison of task selection questions (continued)**

| Task selection step | Author | Question |
|---|---|---|
| Discard | Nowlan and Heap (1978) | Is a discard task to avoid failures or reduce the failure rate both applicable and effective? |
| | Moubray (1991) | Is a scheduled discard task to reduce the failure rate[5] technically feasible and worth doing? |
| | Coetzee (1997/2) | Is a replacement task to reduce the hazard rate both technically and economically feasible? |
| | MSG-3 (1993) | Is a discard task to avoid failures or to reduce the failure rate applicable and effective? |

- Column 5: Important non-hidden failures (high seriousness factor) with constant or decreasing hazard rate and an identified failure indicating quantity – task initiation can take place based on any one of failure and condition – task choice based on efficiency (mostly cost) between these two possibilities.
- Column 6: Important non-hidden failures (high seriousness factor) with constant or decreasing hazard rate and no failure indicating quantity – task initiation when failure takes place – typical task choice corrective.

The method is simple and attractive if somewhat simplistic (for instance, a failure finding task is certainly not the only option in the case of a hidden failure - condition based and use based tasks should also be considered). The fact that this one diagram is used for all four failure consequence types (as in the case of Smith above) does not necessarily create a problem, but the loss of the reliability logic does. However, this type of diagram will have to be considered either on its own or possibly as an additional assist type of function in the task selection decision-making process. This diagram and the one of Smith highlight the fact that the standard decision tree, although theoretically sound, depends too much on the detailed knowledge of the analyst. It should be the objective to either replace or augment the standard tree by a diagram that builds all the relevant knowledge into the tree itself, as far as that is possible. It is a fact that most analysts in industry are not trained and capable to the level of expertise, which is necessary when using the standard decision tree. A 'decision diagram for dummies', if possible, will address this shortcoming.

---

[5] See footnote on previous page

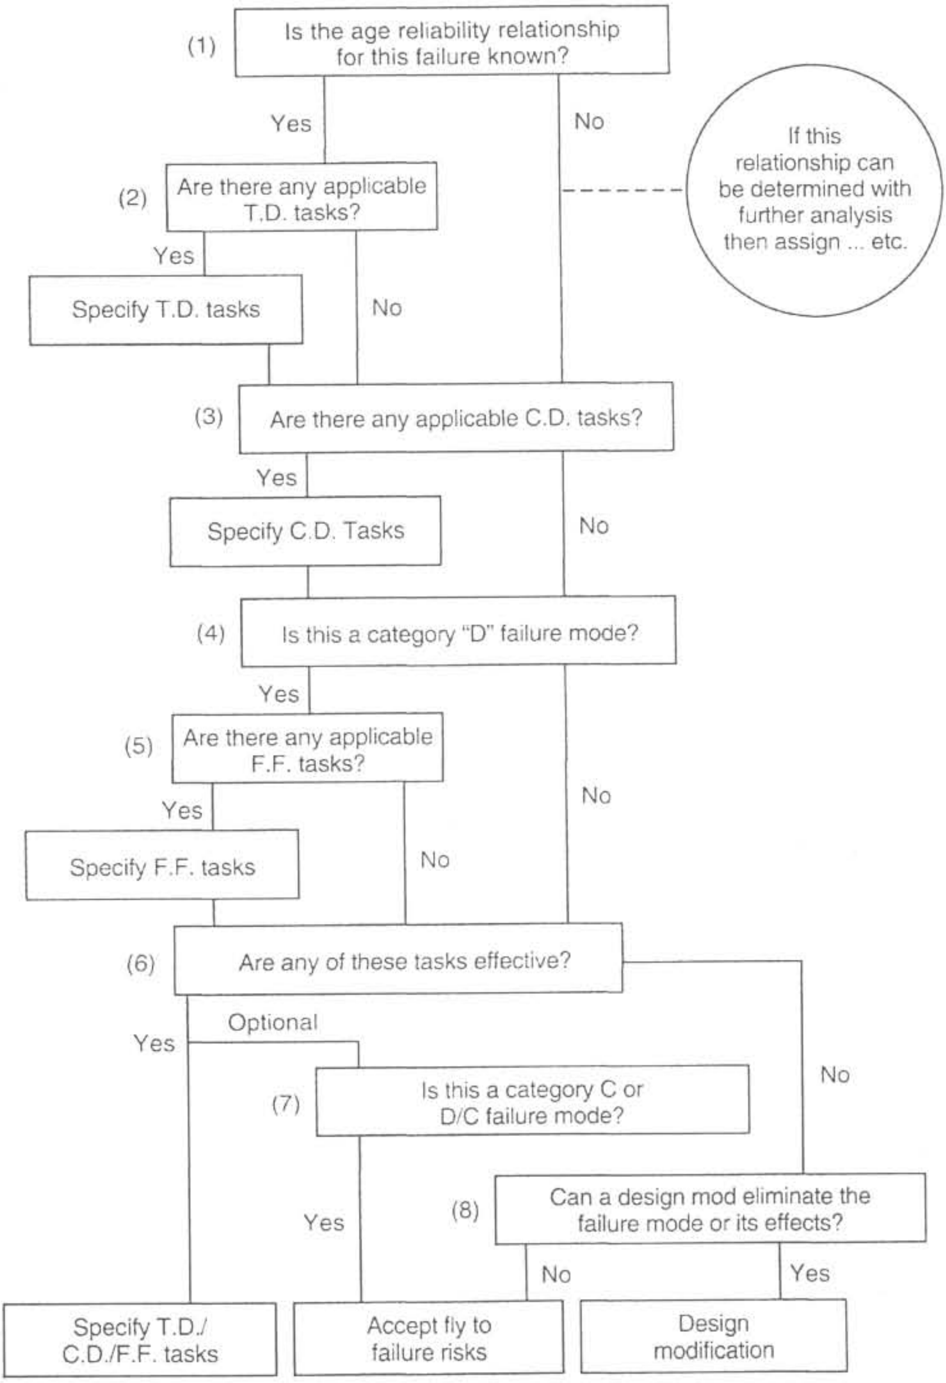

**Figure 3-27: Task selection road map [Smith (1993)]**

Task selection in the Business Centred Maintenance approach of Kelly [Kelly (1997)] is a very simplistic process of studying the design of the item and then using its known failures, as well as the analyst's prior knowledge regarding the operation to choose the best maintenance possibility.

MSG-3 (1993) modified the Nowlan and Heap task selection tree (figure 3-26) in a few ways. The standard MSG-3 tree (reproduced) is shown in figure 3-29, which is applicable for operational consequences and economic (non-operational) consequences. The *first* change is the introduction of the Lubrication/Servicing task at the top of the tree. This task is described as 'any act of lubrication or servicing for the purpose of maintaining inherent design capabilities'. As was stated earlier in this text, this was one of the important shortcomings of the original version of RCM. The *second* change is the fact that the decision-tree is not truncated if a lubrication/servicing task is specified.

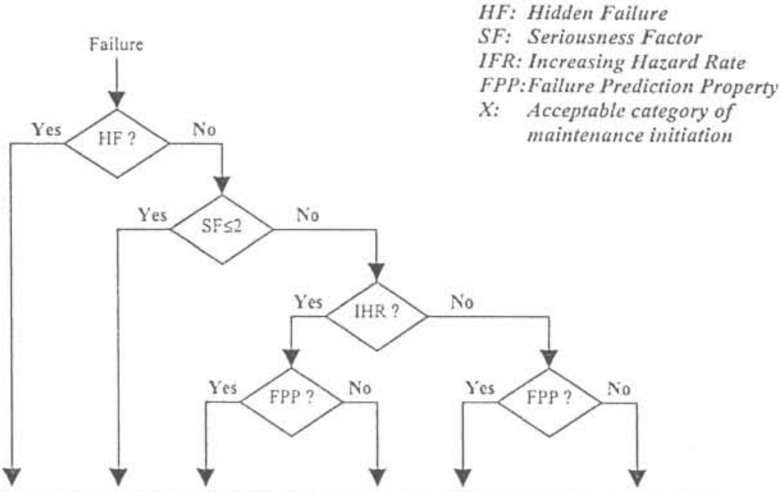

| | | | | | | |
|---|---|---|---|---|---|---|
| FBM | × | × | × | × | × | × |
| UBM | | | × | × | | |
| CBM | | | × | | × | |
| Failure | 4, 5, 6, 7, 8, 9,18, 20, 21, 24, 25, 36, 37 | 1, 2, 10, 11, 12, 13, 15, 16, 26, 27, 29, 30, 31, 33, 34, 35, 39, 40 | 14, 19, 22, | 17 | 23, 28, 38 | 3, 32 |

*FBM:* *Failure Based Maintenance*
*UBM:* *User Based Maintenance*
*CBM:* *Condition Based Maintenance*

**Figure 3-28: Qualification of maintenance initiation categories [Le Clercq and Van den Broek (1999)**

For safety tasks (both of the evident and hidden varieties), this last change is carried further in that no truncation takes place at any level of the task selection process (see figures 3-30 and 3-31). The *third* change is the fact that after each task level is tested for safety tasks (both evident and hidden), the last step is to choose the best combination of tasks selected in the previous steps (if such combination will be applicable and effective) – otherwise redesign is mandatory. The *last* change is the introduction of the 'verify operation' task (failure finding) between the lubrication/servicing task and the inspection/functional check task in the case of hidden consequences (both for safety and for economic varieties) – see figures 3-31 and 3-32. This task is described as 'a task to determine that an item is fulfilling its intended purpose - the check does not require quantitative tolerances'. The task can be an operational check or a visual check.

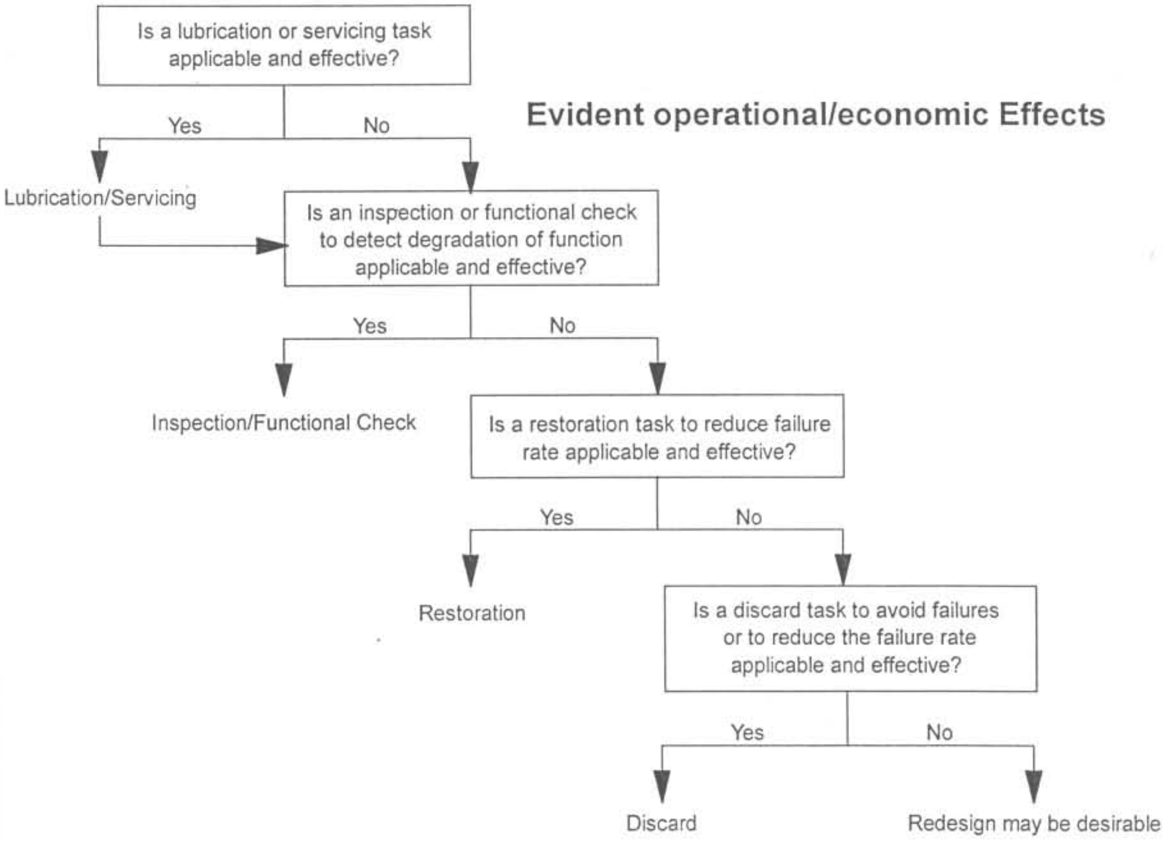

**Figure 3-29: Basic task selection tree [MSG-3 (1993)]**

Harris (1985) also adds an extra task at the top of the task selection ladder (before the on-condition task option). He calls this additional task category the "non-maintenance related improvement" (NMI) task category. He states that failure modes often occur, which cannot be eliminated by preventive maintenance, but which can be eliminated through different means. Examples of these are incorrect operational procedures, poor maintenance workmanship and faults in manufacture of equipment. The first of these can be solved through better operational procedures and training of operators, the second by improved maintenance procedures and the third through improved quality assurance. It is clear that all three examples relate to quality in some or other form. Thus a quality question may be the solution to this very important shortcoming in the present RCM methodology.

He also divides operational consequences into mechanical (technical) consequences and process consequences. This aspect was discussed on page 3-42. His proposed NMI task category will also address the problem of process consequences, without going to the extreme of introducing a new failure consequence category (page 3-43).

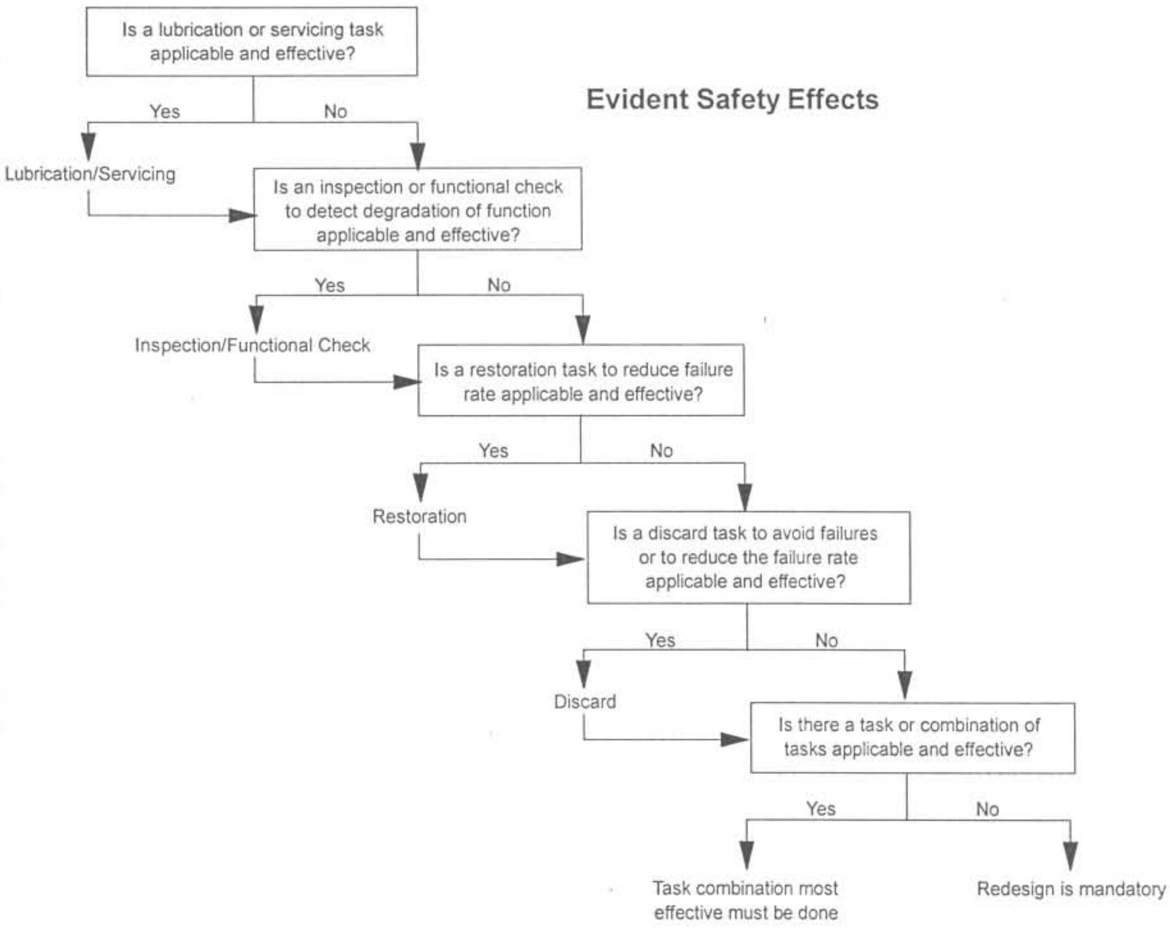

**Figure 3-30: Evident safety task selection tree [MSG-3 (1993)]**

### *3.2.7.4. Technical/Economic viability*

In the original version of RCM, Nowlan and Heap (1978) defined the terms 'applicable' and 'effective' as being two prerequisites for task selection (see figure 3-26). Applicability is a measure that ascertains whether the specific task is 'right' as preventive measure regarding the occurrence of the specific failure mode. If it is applicable, the issue of its effectiveness must be resolved – it must do the job at the right level functionally and in terms of cost, depending on the objective of the task. The two terms are defined [Nowlan and Heap (1978), p.50] as follows:

> *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 objective, 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.*

These terms (applicability and effectiveness) have been retained in MSG-3 (1993).

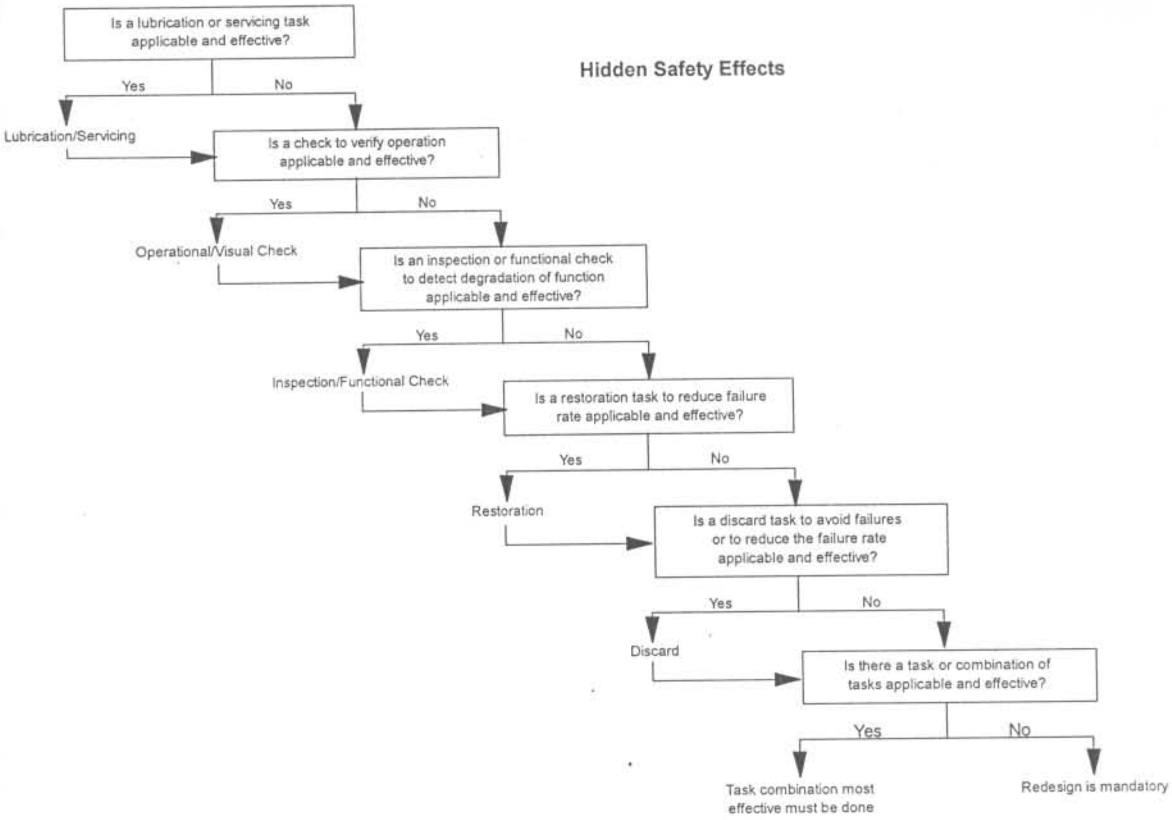

**Figure 3-31: Hidden safety task selection tree [MSG-3 (1993)]**

These terms are not completely satisfactory in that their definitions leave room for doubt and wrong interpretation. Additional problems regarding task 'effectiveness' in MSG-3 is that it freely talks of 'risk reduction' in the case of safety related tasks and 'cost effectiveness' in the case of economic related tasks. Nowhere does it however define what constitutes 'risk' or 'cost effectiveness'.

Moubray (1991) redefined and renamed the two measures for task selection as being *technical feasibility* and whether it is *worth doing*. He defined the two terms as (p.14):

> *Whether or not a task is technically feasible is governed by the technical characteristics of the task and of the failure, which it is meant to prevent.*
>
> *Whether it is worth doing is governed by how well it deals with the consequences of the failure*

These terms are more palatable and the definitions are given in much simpler and clearer terms. Notwithstanding, the basic idea are still the same. The applicability (technical feasibility) is dependent on the technical characteristics of the failure and that of the proposed task. Likewise, the effectiveness (worth doing) measures the success of the task, both in terms of *technical success* (hidden and safety/environmental consequences) and *economic success* (operational and non-operational consequences). It is especially this second category (which is a mixture of two different result types) that creates a problem in user's minds.

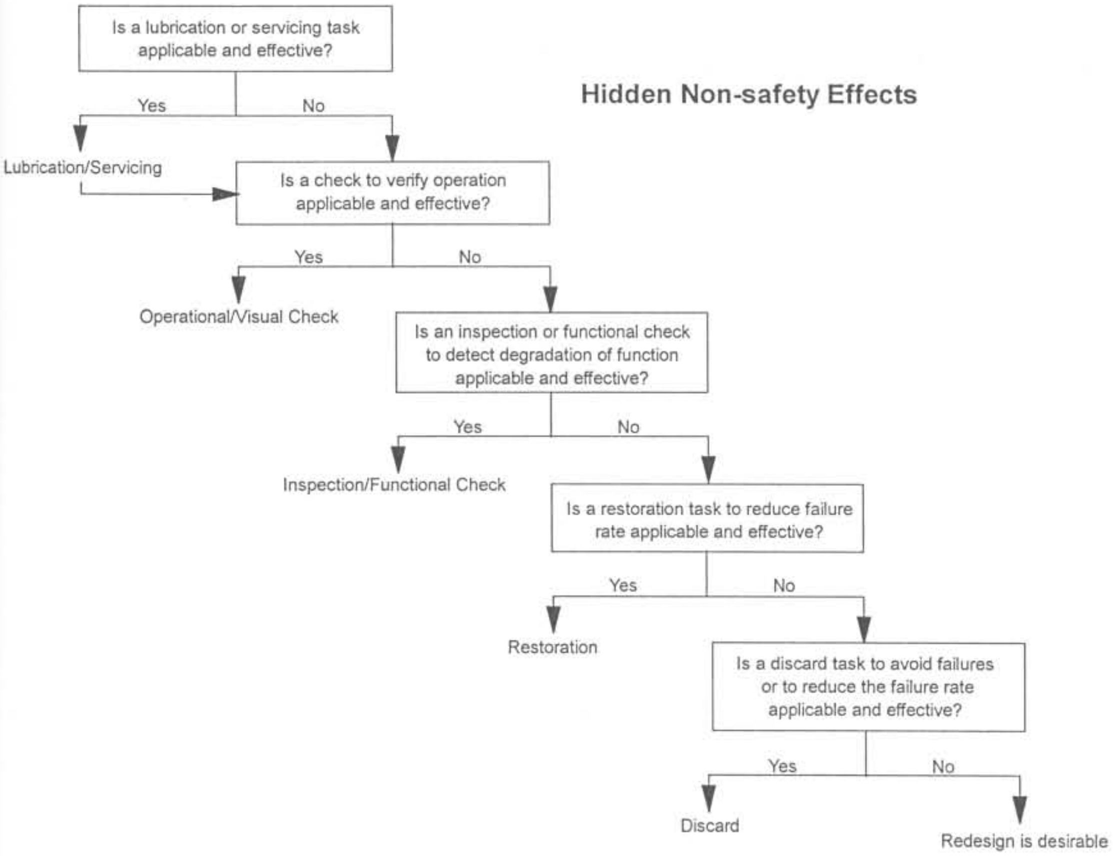

**Figure 3-32: Hidden economic task selection tree [MSG-3 (1993)]**

Coetzee (1997/2) contemplated these problems and used the terms *technically feasible* and *economically feasible*. The term *technical feasibility* thus includes the technical success of the task in dealing with the consequences of failure, while *economical feasibility* is only concerned with whether the task is an economical success. This makes the two primary decision criteria much simpler both to understand and work with. A full description of these and their meaning is extracted from Coetzee (1997/2).

*Technical Feasibility*

> Technical feasibility depends on the technical characteristics of both the failure mode and the task. The task must be able to prevent the failure mode to some specified level of certainty. This specification differs for the various failure consequences:
>
> ❖ *Hidden failure consequence*: the task must reduce the risk of a multiple failure to an acceptable level. In the case of a hidden failure, the failure itself does not have the negative effect, but if it is not found and repaired it will eventually lead to a multiple failure which may or may not be worth preventing depending on the risk involved. The fact

that the failure is hidden, does not necessarily imply high risk – the risk is high if the hidden failure can lead to a costly or unsafe multiple failure. On the other hand, if the hidden failure will not lead to a multiple failure or the resulting multiple failure is of no consequence then the risk involved is negligible (this is an example of a superfluous function - the specific function can in actual fact be removed). If a suitable task that reduces the risk to an acceptable level is not found then the initial default action is a scheduled failure finding task.

- *Safety or Environmental failure consequence*: the task should reduce the risk of the failure either totally or to a (very low) acceptable level. Where safety or environmental issues are at stake one cannot compromise. If a preventive task that reduces the risk of failure to a low enough level is not found, redesign is compulsory (the default action).
- *Operational and non-operational consequences*: the task should reduce the risk of failure to a low enough level so that the benefits of the preventive action are worth the cost of implementing it. If no suitable preventive task is found, the initial default action is corrective maintenance (no scheduled maintenance). Where-as in the case of hidden consequences and safety/environmental consequences, technical viability includes whether the prevention is successful, in this case it does not.

*Economical Feasibility*

Economical feasibility, as the name suggests, is only concerned with whether the task, which has already passed the sieve of technical feasibility, makes economical sense. The different failure consequence categories are dealt with as follows:

- *Hidden, safety and environmental consequences*: Here technical feasibility plays a much greater role than economical feasibility. The reason for this is that the consequence of failure in this area is in most cases detrimental to the health and life of persons and/or plant/business asset. Thus, one normally has to find a preventive task that works. However, one cannot disregard the cost of the proposed action altogether. In some or other way, one has to establish whether the proposed action is the only or least expensive way to solve the problem. This can entail comparing the cost of alternative preventive actions against one another or against the possibility of redesign to solve the problem.
- *Operational consequences*: In the case of operational consequences, the task is economically feasible if the total cost (cost of downtime + cost of failures incurred + cost of prevention) is at a minimum. That is, the cost of the preventive task per operational measuring unit is less than the cost of the operational consequences (in other words the value of production lost) plus the cost of repairing the failure per operational measuring unit.

- *Non-operational consequences*: In this case the answer to the question of economic feasibility is the same as for operational consequences, except that the cost of lost production does not feature (the failure does not affect the operation).

The technical feasibility criteria suggested by Coetzee (1997/2) and expounded above are not completely satisfactory for use in the operational and non-operational consequence categories. It still represents a mixture between technical risk and economic worth. This ambiguity should be removed to ensure that the term has a singular significance.

As a consequence of adding two additional task types to the task selection tree, MSG-3 (1993) has to address the 'applicability' and 'effectiveness' criteria of successfully selecting these tasks as part of the preventive strategy. The applicability criterion for *lubrication/servicing* tasks is stated as 'the replenishment of the consumable must reduce the rate of functional deterioration'. The effectiveness criteria differs for different consequence categories: for safety consequences 'the task must reduce the risk of failure', for operational consequences 'the task must reduce the risk of failure to an acceptable level', and for economic consequences 'the task must be cost-effective'. For the *verify operation* task (failure finding) the applicability criterion is that the 'identification of failure must be possible'. The effectiveness criteria again differs for different consequence categories: for safety consequences 'the task must ensure adequate availability of the hidden function to reduce the risk of a multiple failure', and for economic consequences 'the task must ensure adequate availability of the hidden function in order to avoid economic effects of multiple failures and must be cost-effective'.

As was mentioned previously (page 3-42), Harris (1985) shows that process degradation is not addressed adequately in the present versions of RCM. To address this problem, the cost of performance degradation should be added to the cost of lost production as described above.

Nowlan and Heap (1978) presented a 'quick and dirty' first order decision diagram for economic feasibility. This was improved by Coetzee (1997/2) and is shown in figure 3-33. As economic feasibility plays only a secondary role in the case of hidden failure consequences, as well as with safety and environmental consequences, this diagram is only applicable in the case of the operational and non-operational consequence categories.

### *3.2.7.5. Technical selection criteria*

Referring to Nowlan and Heap's (1978) task selection decision tree (figure 3-26) and table 3.4 above, the primary technical selection criteria for the different task types are as follows:

- On-condition task: detection
- Rework task: reduction of failure rate
- Discard task: failure avoidance *or* reduction of failure rate

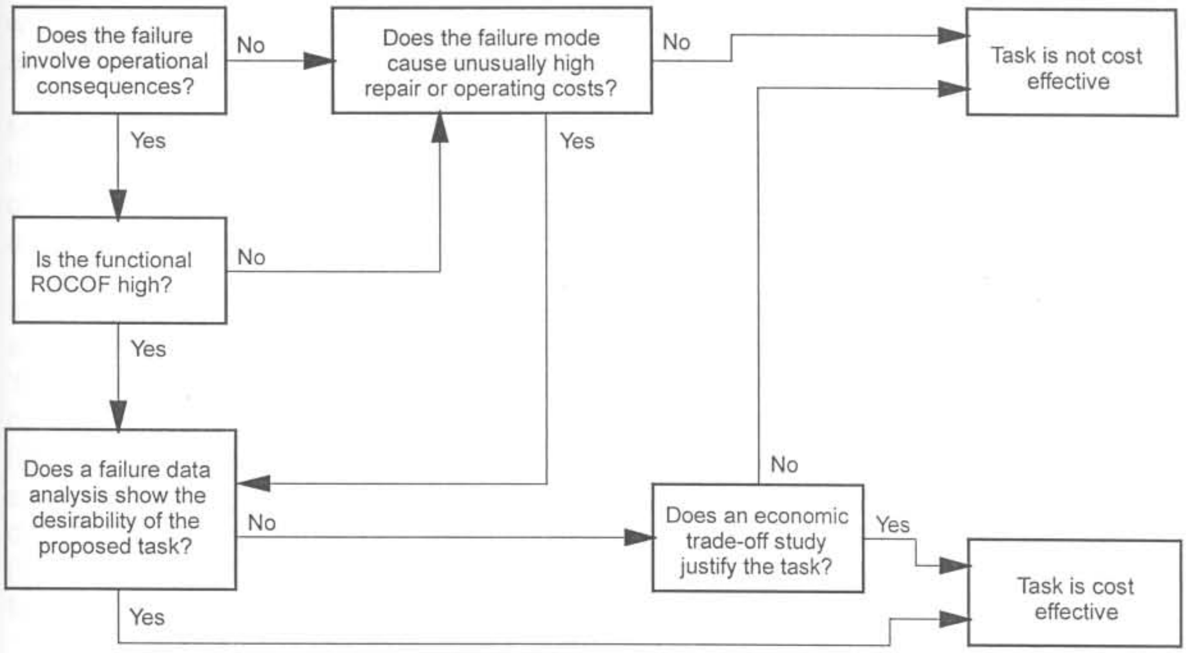

Figure 3-33: Cost Effectiveness Decision Diagram [Coetzee (1997/2)]

In the case of on condition tasks, all the authors concur that detection is the primary selection criterion. That includes the authors not specifically listed in table 3.4, because they mostly agree with the original Nowlan and Heap structure. The exception is Gits (1988), who does not agree with the Nowlan and Heap structure, having developed his own technique, but agrees that the primary selection criterion for condition-based tasks is detection (or the 'failure prediction property' (fpp) – see figure 3-28).

For rework/restoration/reconditioning tasks, all the authors agree that the primary selection criterion is the 'failure rate', although it is not totally clear that all of them appreciate the implication of that statement. The idea is that one can only practice Use Based Maintenance under conditions of IFOM[6] (Increasing Force of Mortality) or an increasing conditional failure rate (or Hazard Rate as Coetzee (1997/2) has it). The problem with the use of 'failure rate' here is that it can most probably be mistaken for an increasing ROCOF[7] (Rate of Occurrence of Failures) which will certainly increase in most cases of use based preventive action. A common problem with all of the authors (excluding Nowlan and Heap (1978) and Coetzee (1997/2)) is that they judge

---

[6] The Force of Mortality typically gives the risk that a component will fail at a certain age. Under an increasing Force of Mortality (IFOM), it could thus be justified to replace or recondition the component at some age, as the risk of failure of the renewed component will be much lower than that of the component before renewal. Cost factors will determine at what exact age renewal has to take place.

[7] An increasing Rate of Occurrence of Failures (ROCOF) is typical of a system which consists of many components, many of which will be IFOM, others displaying random shock failure patterns (constant FOM) and some even running-in failure patterns (decreasing FOM). As time passes, more and more components start failing for the first time, thus causing an increasing ROCOF. This is called reliability degradation.

whether the 'failure rate' is increasing through discussion with plant operations and maintenance personnel rather than through proper failure data analysis.

Moubray (1991) uses failure avoidance as the primary selection criterion in the case of safety/environmental consequences, which is probably more correct as the purpose of preventive action in this case should be total failure avoidance, which is a more conservative rule than FOM reduction.

In the case of discard/replacement tasks, Nowlan and Heap (1978) uses two alternative primary selection criteria, namely failure avoidance and increasing 'failure rate'. It is not clear why they include failure avoidance here as they never explain themselves regarding this issue. It probably has to do with safe life limits imposed through statutory or regulatory requirements. All the later authors use only the second criterion, as is the case with rework tasks. The discussion above for rework tasks further holds fully for this category as well (including the previous paragraph's remark regarding Moubray's position on failure avoidance).

MSG-3 (1993) has nothing new to contribute in this area, apart from the criteria for the two new task types, as described in paragraph 3.2.7.4 above. For on-condition tasks, it specifies detection and P-F interval[8] consistency as criteria as is the case with most other authors. In the case of restoration and discard tasks, MSG-3 has in fact not progressed at all compared to the original RCM [Nowlan and Heap (1978)], as it still sticks to the idea of the bathtub curve. It thus sees technical applicability for restoration and discard tasks as being dependent on 'the item must show functional degradation characteristics at an identifiable age and a large proportion of units must survive to that age'. For restoration tasks, it adds 'it must be possible to restore the item to a specific standard of failure resistance'.

### *3.2.7.6. Default tasks*

Although the task selection decision tree (figure 3-26) is principally the same for all consequence categories, the default action if no suitable preventive task was found differs considerably for each consequence category. Moubray (1991) coined the term 'default actions' for these consequence specific tasks. There is not much that warrants comparison in the different literature sources. The only author with a proper discussion of the default task area is Moubray (1991). A short discussion of the default tasks per consequence category from this source follows in the interest of completeness:

- Hidden failures – if no suitable preventive task is found in the case of hidden failures, the default action is a scheduled Failure Finding task (if it is possible to detect the failure through such means and if the task is technically and economically feasible). Should this possibility not be viable, and safety is involved, Redesign is compulsory – if safety is not

---

[8] The time between the point at which a Potential failure is identified using inspection or Condition Monitoring and the time at which the actual Functional failure occurs.

involved the default is Corrective Maintenance, with Redesign as an option.

- Safety/Environmental Consequences - if no suitable preventive task is found in the case of Safety/Environmental Consequences, the default is to consider a combination of preventive tasks (On-condition, Rework and Discard), as the failure mode with Safety/Environmental Consequences cannot be allowed to exist unchecked. Should this possibility not be viable, Redesign is compulsory.
- Operational and Non-operational Consequences – in the case of the so-called economic consequence categories, the lack of a suitable preventive action is not a crisis and the default action is Corrective Maintenance, with Redesign as an option.

## 3.2.8. Task Frequencies

### *3.2.8.1. Introduction*

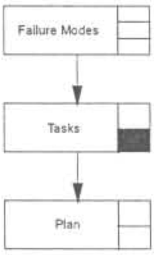

Maintenance task frequencies, especially concerning critical items, are a perpetual problem, as the task of maintenance is to prevent failures from occurring. The better the maintenance function fares regarding this objective, the less failure data will be available for analysis to determine scientifically calculated task frequencies. Nevertheless, one should do the best possible within the constraints of a scarcity of data. Fortunately, statistical methods allow for the use of suspended data in conjunction with the failure data that is available to enrich the data set. And, as the task decisions regarding rework and discard tasks require an increasing hazard rate, data analysis is essential if one does not want to choose your strategy from the On-condition category only, apart from a few statutory and regulatory life limiting tasks. The same analysis that leads to the hazard curves will also provide the task frequencies in the use based task categories. On-condition tasks and scheduled Failure Finding task frequencies, on the other hand, can be determined from past experience with similar equipment as well as manufacturer data, being very conservative in the setting of inspection intervals (typically only a small fraction of the assumed life) [Nowlan and Heap (1978), p. 324]. As the item's life gets nearer to the assumed failure point, the frequency of inspection is repeatedly increased. It is again important to note that accurate inspection frequencies also requires knowledge regarding the Mean Time to Failure and the variance of times to failure (i.e. the failure distribution must be known).

Nowlan and Heap (1978), (p.106) stresses the use of age exploration to develop task intervals as the equipment ages, especially in the case of On-condition tasks. For an On-condition task, three factors determine its technical feasibility. The first of these is the characteristics of the failure (especially whether the failure has a gradual onset). The second is the ability to measure reduced failure resistance and the third its ability to deal with the failure consequence. This third factor depends on the interval between inspections (the same is true for Failure Finding inspections). Thus, the intervals between inspections need to be conservatively short to ensure that the task is technically feasible, given the lack of experience/data. As the equipment gets older, this interval is even shortened further using the same reasoning. However, once in a while one of the items is left to fail without taking action, but monitoring it all the while, so to build up data/experience. The intent is to, over the longer term, increase the intervals based on operating experience (p.324 op.cit.).

For use based maintenance tasks Nowlan and Heap (1978), p.61 stresses the importance of the availability of failure data to determine task intervals. They argue that the only justification for an economic life limit is cost effectiveness, a fact that cannot be refuted. Cost effectiveness can only be proved if the failure distribution is known to predict how the age of preventive removal of an item will affect the cost-benefit ratio.

The task frequency selection approach of MSG-3 (1993) is very simplistic. It consists of the following three sources. The *first* concerns prior knowledge from other aircraft systems/powerplants and the *second* manufacturer's test data. *Thirdly*, if there is insufficient similarity between the previous and current systems, the task interval/frequency can only be established initially by an experienced working group and steering committee personnel using good judgement and operating experience in concert with accurate data.

Following this introductory passage, the various author's opinions regarding this important topic is explored.

#### *3.2.8.2. On-condition tasks*

Nowlan and Heap (1978) states that the task frequencies for On-condition tasks can be determined from past experience with similar equipment as well as manufacturer data. The approach is then to be very conservative in the setting of inspection intervals (typically only a small fraction of the assumed life) [Nowlan and Heap (1978), p. 324]. This high conservatism in the times between inspections is to ensure that the task is technically feasible, given the lack of experience/data. As the item's life gets nearer to the assumed failure point this interval of inspection is repeatedly even shortened further using the same reasoning. However, the intent is to increase the intervals – this is only possible based on operating experience. That is, accurate inspection frequencies requires knowledge regarding the Mean Time to Failure and the variance of times to failure (i.e. the failure distribution must be known). This is achieved through age exploration, i.e. once in a while leaving one of the items to fail without taking action, but monitoring it all the while, so to build up data/experience.

The original authors of RCM worked from a basis of developing a maintenance plan for new equipment, thus working from the premise of no experience/data. This is certainly one of the typical RCM scenarios (the one extreme). The other extreme is a situation where ample experience/data is available, and then one has all the possible situations between these two extremes, which possibly describes reality best.

Moubray (1991) following this practical experiential thought train states that On-condition task frequencies are governed by the P-F intervals (i.e. the time that expires between the point where a potential failure becomes measurable and the point of actual failure) (p.166). He suggests that a task interval of half the P-F interval should be conservative enough to ensure that the possibility of a failure is detected early enough.

Smith (1993) does not have anything substantial to add to the discussion and only echoes the thoughts of Nowlan and Heap given above in much shorter format. Gits (1988), on the other hand, just states that the inspection interval determines the cost of inspection, without even referring to the influence of the interval on task effectiveness.

Coetzee (1997/2) states that the effectiveness (and thus the technical feasibility) of an On-condition task to prevent failure from occurring depends on the following:

- It must be possible to unequivocally identify the potential failure through the condition monitoring or inspection method that was chosen.
- Normally when a component's condition deteriorates, it does so somewhat slowly in the beginning and then accelerates towards the point of functional failure. When we monitor such a deterioration process, it is important that the lag time between the point where our monitoring establishes that a potential failure exists and the point of actual functional failure is *firstly* reasonably consistent. It should, *secondly*, be more than enough to take preventive action (that is to prevent the failure completely or to limit the consequences of failure sufficiently).
- The condition monitoring or inspection frequency envisaged must be practically achievable. The time between inspections (condition monitoring) must be sufficiently less than the time interval between the point at which the potential failure becomes evident and the point of functional failure to ensure that the potential failure will be identified in time.
- The time between monitoring points should be set less than half the lead-time to failure (the time interval between the first identifiable indication of the potential failure and the functional failure). This will ensure that the potential failure will, with high probability, be identified in time for preventive action to be taken.
- In practice, it might sometimes be necessary to start monitoring at a higher frequency once the potential failure has been identified, to accurately predict when functional failure will occur. This is especially so in the case where the lead-time to failure spans a long time (sometimes weeks or even months).

### *3.2.8.3. Failure Finding tasks*

Moubray (1991) is the only author to have a proper discussion on the frequency of Failure Finding tasks. He states that these frequencies are governed by the *consequences of the multiple failure*, which dictate the necessary availability of the hidden function, and the *mean time between occurrences of the hidden failure.*

He presents two approaches to determining the Failure Finding task interval. The first of these is the 'rigorous' approach, where the required availability of the hidden function is calculated from:

$$Av_h = P_m \times P_p$$

$Av_h$ = Availabilty required of the hidden function
$P_m$ = Acceptable probability of occurrence of resultant multiple failure
$P_p$ = Probability of occurrence of failure of protected function during period under consideration

The required Failure Finding inspection interval can then be calculated from:

$$FF_i = f_i \times MTTF_h$$

$FF_i$ = Failure Finding inspection interval
$f_i$ = Failure Finding inspection interval fraction
$MTTF_h$ = Mean Time to Failure of hidden function

with the inspection interval fraction $f_I$ corresponding to the calculated value of $Av_h$ read from the following table (which is given without proof):

| **$Av_h$** | 0,995 | 0,975 | 0,95 | 0,93 | 0,91 | 0,89 |
|---|---|---|---|---|---|---|
| **$f_I$** | 0,01 | 0,05 | 0,10 | 0,15 | 0,20 | 0,25 |

His second approach is applicable to situations where the hidden function is not important enough to warrant the analysis effort. In such cases a decision regarding the hidden function required availability, $Av_h$ is made directly and a value for $f_I$ read from the table above, after which the Failure Finding inspection interval $FF_I$ is calculated.

### *3.2.8.4. Rework tasks*

As was stated in the introduction above Nowlan and Heap (1978), p.61 stresses the importance of the availability of failure data for determining the task intervals of use based maintenance tasks. They argue that the only justification for an economic life limit is cost effectiveness, a fact that cannot be refuted. Cost effectiveness can only be proved if the failure distribution is known to predict how the age of preventive removal of an item will affect the cost-benefit ratio.

Smith (1993) proposes a method whereby the task interval is based on the Failure Distribution function (an acceptable risk of failure is chosen and then the corresponding age at which that cumulative percentage of failures have occurred read off) in the case where data is available. In other cases (no failure data available) he suggests an initial conservative intelligent guess regarding the task interval. Age exploration is then used to improve this interval over time. Each time that the task occurs, the unit and all its parts are meticulously inspected, its condition recorded, and its task interval extended by a percentage based on the condition of the suspended unit. In this way, the correct task frequency is eventually found, without initial experience/failure data availability.

Moubray (1991) suggests that the interval, in most instances, be determined by the 'useful life'[9] through the 'consensus of people who have the most knowledge of the asset' (p.225). He states that analysis of failure data is not worth doing for more than 1 to 2% of failure modes and then for operational and non-operational failure modes only. This is a very simplistic view and fails to see the benefit of (i) understanding the failure mode well through a proper reliability analysis and (ii) analysing the failures of the units with the highest failure rates (the problem cases). Even if it then amounts to as low as 1 to 2%, the benefit will be disproportionately high (much higher than the 1 to 2% suggests). Furthermore, the 'consensus of people' is a very dangerous (and subjective) quantity to base this decision on. The method that Smith suggests above is much more balanced.

It is clear that this area needs a more fundamental investigation. It is true that there is a dearth of proper failure data in industry (for the reasons stated in the introductory paragraph above). Better ways of extracting information from the combined knowledge base of the organisation (formal data and the experience of people) will have to be investigated, perhaps using Bayesian statistics.

### *3.2.8.5. Discard tasks*

The argumentation for discard tasks broadly follows that for rework tasks above, with the only exception being safe life limit tasks. These are typically set by statutory or regulatory authorities and pose no problem in determining

---

[9] That life at which there is a rapid increase in the conditional probability of failure (p.206).

the task intervals. The only problem with such tasks is that the intervals are often unnecessarily restrictive and costly. In such cases, it is beneficial to enter into age exploration as described in paragraph 3.2.8.4 above in co-operation with the statutory or regulatory authority involved with the objective of extending the relevant safe life limits without jeopardising safety.

When the organisation itself imposes safe life limits on certain items, the proven minimum life is often divided by a safety factor of as high as 3 to 4 [Moubray (1991), p.114]. Such a limit would then also be subject to careful age exploration as described in paragraph 3.2.8.4 above.

### 3.2.9. Task Packaging

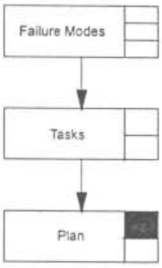

Following the analysis presented above in paragraphs 3.2.2 to 3.2.8, the result is an endless list of individual tasks, which need to be grouped into logical work packets that can be done at suitable intervals. While performing the analyses, no regard should be given to any present intervals, practical limitations and own preferences. These issues should be resolved now during the task packaging exercise.

The task packaging procedure suggested by Nowlan and Heap (1978) starts by adding known non-RCM tasks to the total list of tasks from the analysis – in their case this included servicing and lubrication tasks specified by the manufacturer and scheduled zonal walkabout checks. These are now grouped into standard 'number check' and 'letter check' maintenance packages, each done at an increased number of flight hours. The typical check frequencies used at that time was:

- #1 Check – every stop at a maintenance station
- #2 Check – every long layover if the aircraft has flown more than 20 flight hours since the previous #2 check
- A Check – every 125 flight hours
- B Check – every 900 flight hours
- C Check – every 3 600 flight hours
- D Check – every 20 000 flight hours

One principle involved in the execution of these checks was that each higher check, when performed, will automatically include the work of the nearest lower work packages. Thus, a number 2 check will include the task content of a number 1 check, a A-check will include the nearest #2 check's content (and through that that of the #1 check), a B-check will include the A-check at 875 hours, and so on.

The process of task packaging then follows roughly the following principles:

- Each task in the task list resulting from the RCM-analysis is included in a work package with the objective of having as few different check intervals as is possible.

- The work package intervals chosen, like the task intervals, are chosen to be highly conservative.
- The individual task intervals are adjusted, if possible upwards (if that would not jeopardise the expected result), to fit in with the closest work package interval.
- If the workload of the larger work packets (i.e. C and D in this case) will cause too large fluctuations in the maintenance workload, these can be spread evenly amongst lower level packages. For example, the C-check's workload can be distributed evenly between the 4 B-packets between and including the C-check. This will have the effect that the B-checks are differentiated, such that a 900 hour B-check will only be repeated at 4 500 hours, and so forth. This creates an additional scheduling, planning and control complexity, but evens out the workload. In this instance, there will be no real C-check, but only a phantom C-check.
- A group of tasks that are particularly time-consuming might be distributed between successive instances of the same letter check.
- Both the task intervals and the work package intervals in an initial program is subject to age exploration, although the individual task intervals are typically extended through the extension of the package intervals. Other possibilities include moving the task to another work package and freezing the work package interval if a task interval cannot be adjusted and this specific task controls the package.

Moubray (1991) has nothing significant to add to the above, but adds the possibility to incorporate some tasks in the operating procedures. This, and a suggested modification (redesign) list, is shown in figure 3-34 [Moubray (1991), p.251].

The method of Gits (1988) deals quite comprehensively with the task packaging issue. He proposes a three-step process of task packaging:

- Clustering of tasks – in this step the tasks are grouped per common set-up[10] type. The benefit of this step is that the number of times that a specific set-up has to be performed is limited with the resultant production income benefit, with the possible adverse effect of shortening the intervals of certain individual tasks. The practice of the clustering operation is to group the tasks per set-up type and per task interval. Thus two groups of tasks with the same task interval will constitute two separate clusters if they do not share a common set-up type.
- Structuring of maintenance intervals – here the task intervals are fitted into the constraints imposed by the user organisation. This is similar to the task package intervals chosen in Nowlan and Heap's method above (principle 2) and then fitting the task intervals into the available package intervals. The typical application of this rule is the decrease of task

---

[10] A set-up is an activity which must precede execution of the maintenance task. This includes production shutdown, opening up of the item, disassembly, administrative procedures.

intervals to fit in with the work package intervals. This agrees with the approach of Moubray (1991), but disagree with that of Nowlan and Heap (1978) above, who suggested a preference for increasing the intervals if possible. This is not necessarily conflicting as the position of Gits and Moubray reflects the more conservative stance and that of Nowlan and Heap the more economic, which should be pursued if possible but not at the expense of task effectiveness.

- Grouping of work packages – the objective of this step is to smooth the workload within the organisational constraints. This is similar to the actions in the 4th principle of Nowlan and Heap above.

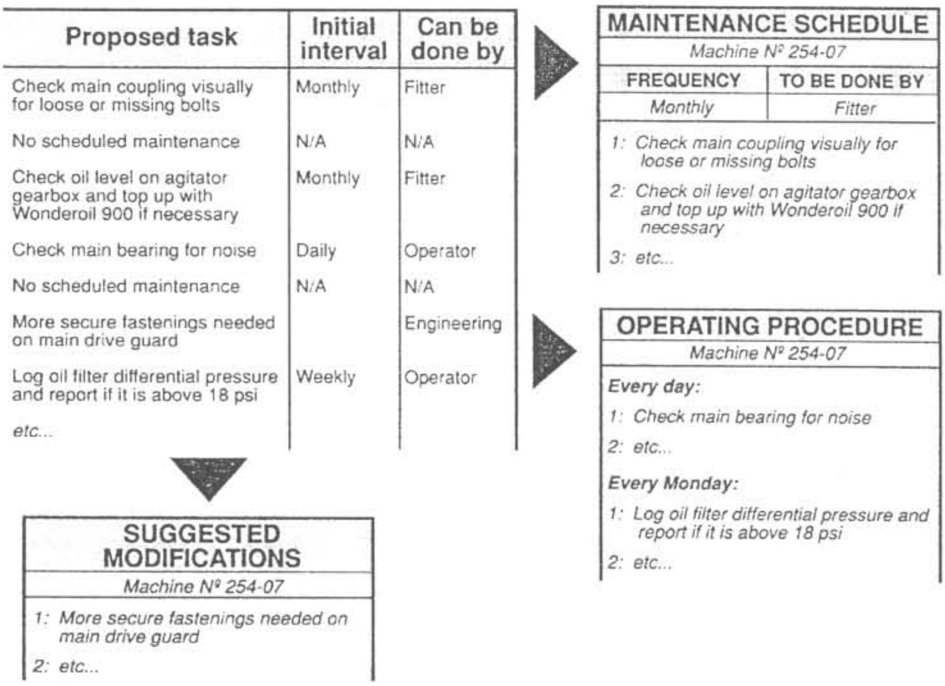

| Proposed task | Initial interval | Can be done by |
|---|---|---|
| Check main coupling visually for loose or missing bolts | Monthly | Fitter |
| No scheduled maintenance | N/A | N/A |
| Check oil level on agitator gearbox and top up with Wonderoil 900 if necessary | Monthly | Fitter |
| Check main bearing for noise | Daily | Operator |
| No scheduled maintenance | N/A | N/A |
| More secure fastenings needed on main drive guard | | Engineering |
| Log oil filter differential pressure and report if it is above 18 psi | Weekly | Operator |
| *etc...* | | |

**MAINTENANCE SCHEDULE**

*Machine Nº 254-07*

| FREQUENCY | TO BE DONE BY |
|---|---|
| *Monthly* | *Fitter* |

*1: Check main coupling visually for loose or missing bolts*
*2: Check oil level on agitator gearbox and top up with Wonderoil 900 if necessary*
*3: etc...*

**OPERATING PROCEDURE**

*Machine Nº 254-07*

***Every day:***
*1: Check main bearing for noise*
*2: etc...*

***Every Monday:***
*1: Log oil filter differential pressure and report if it is above 18 psi*
*2: etc...*

**SUGGESTED MODIFICATIONS**

*Machine Nº 254-07*

*1: More secure fastenings needed on main drive guard*
*2: etc...*

**Figure 3-34: Task packaging [Moubray (1991), p.251, figure 11.2**

Kelly (1997) divides the maintenance tasks of a unit (typically sub-plant) into three groups:

- Work that can be done during production.
- Work of a minor character, which can only be done during production stoppages.
- Work of a major character that can only be performed during production stoppages.

The first of these is bundled together and grouped into 'on line routines', the second scheduled for production window opportunities and the third added to long term shutdown definitions. As a part of this process a spare parts needs analysis is also done to serve as basis for spare parts provisioning.

MSG-3 (1993) does not address the issue of task packaging at all.

### 3.2.10. Critical assessment of resulting program

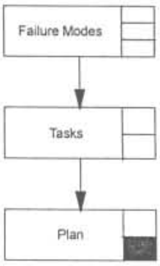

The only authors (including MSG-3 (1993)) that have a proper section on the evaluation of the resultant maintenance plan are Nowlan and Heap (1978). This is difficult to understand, given the importance of the decisions taken in the process of analysis and the implications of wrong analyses. Compare this to financial auditing where quite small transactions are audited to prevent loss and it is still more difficult to understand. The problem is not necessarily that the analysts will make wrong recommendations on purpose, but that they necessarily have some biases and blind spots, which should be checked by an independent auditor.

Such an auditor in the case of the aviation industry is typically a member of the RCM project steering committee. This has the benefit that the auditor is aware of the scope of the project and the interfaces between different working groups/analysts. Coetzee (1997/2), p.81 suggests that the maintenance manager for whom the RCM analysis is done should also act as the auditor and final approval of the program. This has the benefit that he is involved in the process. The result of the maintenance effort is certainly in his interest.

Nowlan and Heap (1978) discusses RCM auditing in quite some detail. Instead of a long discourse, it would be beneficial if this part of the analysis is carried out in an itemised manner, highlighting the audit concerns in a question and/or statement format, following the pattern of Nowlan and Heap:

- Project auditing
  - Scope of project
    - Does the team have clarity regarding the system/equipment boundaries?
    - Does the team know which maintenance actions should be included? Does it include lubrication/servicing and walkabout checks?
    - Is there clarity concerning who/which team is responsible for interfacing functions/items?
  - Final product definition
    - What is the expected final product? A list of tasks (with/without intervals) or a finalised set of work packages?
    - Are the task detail and the amount of descriptive material to be included in the final product clearly specified?
    - Will the procedure writers be able to translate the analysis results into clear task instructions that reflect the purpose of each task?
  - Project timetable
    - Is the timetable realistic when compared to the workload, the analysts assigned and their experience?

  - Do the project milestones constitute logical points for project control?
  - Is the auditor clear on his scheduled involvement? He typically has to do auditing chores at the following checkpoints:
    i. When the overall plan is drafted.
    ii. When the program development team has been organised and trained.
    iii. When each working group has agreed on a list of significant items.
    iv. When analysis of the first few items has been completed.
    v. At the completion of each major portion of the program.
    vi. When the final product has been assembled and is ready for approval.

- Program-development team
  - Are the team capable (in terms of experience, qualifications and organisation) to complete this analysis successfully?
  - Are the team managed by a capable manager with experience on similar projects?
  - Does the project team have access to the necessary resources, both within and outside the organisation?
  - Are all the necessary disciplines represented in the team or available for consultation?
  - Are there organisational obstacles that may impede communication?
  - Is each analyst responsible for a complete analysis, or are various aspects of the job assigned such that it makes the work difficult to integrate?
  - Is the designer available to answer questions about specific failure modes and effects?
  - Is there someone available to each working group who has extensive knowledge of RCM techniques?
- Standards and Procedures
  - Are all participants of the project trained in RCM procedures and the specific standards utilised for this project?
  - Are RCM reference materials available for easy reference?
  - Does each analyst have a copy of the cost trade-off models to be used, including the costs imputed by the organisation to various types of operational failures?
  - What level of failure rates or repair expenses are considered high enough to qualify an item for analysis?

    - Have all analysts been supplied with the necessary stationary, schematics and descriptions of part relationships/ operation as is necessary?
    - Are reliability data available, either from developmental testing or from service experience?
    - Is there access to an actual production model of the equipment if further questions arise?
- Auditing the decision process
  - The selection of items for analysis
    - Has the group arrived at a common definition of *significant item?* (there is often a tendency to identify items as significant on the basis or their cost and complexity, rather than on the basis of their failure consequences).
    - Is the group clear as to what constitutes *operational consequences*? (because the actual economic impact will vary from one operating context to another and even from organisation to organisation, it is necessary to have a clear definition of the circumstances that constitute operational consequences and the relative costs imputed to those consequences).
  - Reviewing the information worksheets
    - Does the design of worksheets used allow the full functionality of RCM to be performed?
    - Has the analysis been performed at the right level? (see paragraph 3.2.2)
    - Has each item's functions been identified correctly? What about secondary functions? And hidden functions?
    - Is there any confusion between functional failures and failure modes?
    - Do the worksheets list failure modes that have never actually occurred?
    - Are the failure modes reasonable in light of experience with similar equipment?
    - Have any important failure modes been overlooked?
    - Does the description of failure effects include all the information necessary to support the analyst's evaluation of the failure consequences?
    - Are the effects of secondary damage stated, as well as the effects of a loss of function, and is it clear from the description whether or not the secondary damage is critical? (in the case of hidden functions the ultimate effects will usually represent the combined effects of a possible multiple failure).
    - Do the failure effects represent overreaction by inexperienced analysts?

- Have all serious effects been included?

- Classification of failure consequences
  - Hidden functions:
    - Has the evident-failure question been asked, not for the item, but for each of its functions?
    - Has all instrumentation been included as a means of notifying the operating crew of malfunctions that would otherwise not be evident.
    - Have failures in replicated functions been identified as hidden?
    - Have the hidden functions of emergency items been overlooked?
    - Have hidden-function items with built-in test equipment been identified as being hidden regardless of the fact that failure-finding tasks are performed by the operating crew.
  - Safety Consequences:
    - Has a failure been identified as critical on the basis of multiple-failure consequences, rather than the consequences of a single failure?
    - Has the analyst taken into account redundancy and fail-safe protection that prevents a functional failure from being critical?
    - Was there a failure to identify secondary damage as critical when the system/equipment cannot be shown to be damage-tolerant in this respect.
  - Operational Consequences:
    - Is there a tendency to interpret failures that are expensive to repair as having operational consequences, or to ascribe operational consequences to failures that inconvenience the operating personnel?
    - A no answer to question 3 means that the failure in question has only non-operational consequences, and that function need not be protected by scheduled tasks in an initial program. If the item is subject to a particularly expensive failure mode, it will ordinarily be assigned to intensive age exploration to determine whether scheduled maintenance will be cost-effective. At this stage. however, any task analysis that falls in the third branch of the decision diagram is subject to challenge by the auditor and must be supported by a cost trade-off study based on operating data for the same or a similar item.
    - All answers to the first three task decision questions should be examined in detail, at least for the first few items completed by each analyst to ensure that the analyst fully understands the nature of the questions.
- Task selection: applicability criteria

- Does the analyst understand the relative resolving power of the four basic types of task and the specific conditions under which each type of task is applicable?
    - On-condition task: If the task is merely an inspection of the general condition of the item and is not directed at a specific failure mode, it does not constitute an on-condition task. The failure mode must also be one for which it is possible to define a potential-failure stage, with an adequate and fairly predictable interval for inspection. It is also important to evaluate proposed on-condition tasks in terms of their technical feasibility. The failure mode may be one for which on-condition inspection is applicable, but is the item accessible for inspection? Is the task one that is feasible within the maintenance framework of the organisation? Does each inspection task include the specific evidence the artisan is to look for - if not, the procedures writers may have difficulty converting the task to the proper job instruction.
    - Rework task: have the age-reliability characteristics of the item been established by actuarial analysis? Is the failure mode one for which rework will in fact restore the original resistance to failure? Is there a cost-effective interval for this task? Has the item been assigned to age exploration to obtain the necessary information?
    - Discard task: the only discard tasks that should appear in an initial program, is for items that have been assigned life limits by the manufacturer. Safe-life tasks are applicable only to items subject to critical failures; hence, they should appear only in the safety branch of the decision diagram. The life limits assigned to hidden-function emergency items (which are not in themselves subject to critical failures) are adjusted based on failure-finding tests and in the strict sense are not safe-life limits. The auditor should question any safe-life discard tasks that are not supported by on-condition inspections (where possible) to ensure that the safe-life age will be achieved.
    - Failure-finding task: is there a failure to recognise that these tasks are the default? Is there a failure to recognise that these tasks are limited to the detection of functional failures, not potential failures? The intervals for such tasks should also be examined tor mistaken assumptions concerning the required level of availability. Does the level of availability properly reflect the consequences of a possible multiple failure? Has the analyst overlooked the fact that the interval is based only on the required availability of the hidden function itself? Have failure-finding tasks covered by routine crew checks been accounted for on the decision worksheets?

❖ Task selection effectiveness criteria

- The applicability criteria for tasks pertain only to the type of task and are true for that task regardless of the nature of the failure consequences. The effectiveness criteria however, depend on the category of failure consequences.

  - Some practical problems often come up in interpreting the effectiveness criterion for the safety branch. Do the tasks and intervals selected have a reasonable chance of preventing all critical failures? If not, what is the basis for judging that the remaining risk level is acceptable?
  - On-condition tasks provide control of individual units and therefore have a good chance of preventing all functional failures if the inspection interval is short enough.
  - Scheduled removals merely control the overall failure rate for the item. The auditor should therefore question the decision outcome of scheduled rework in the safety branch, because a reduction in the failure rate is unlikely to reduce the risk of failure to an acceptable level.
  - In the case of Operational Consequences, does the analysis show the basis for determining that the task will be cost-effective? What costs are imputed to the operational consequences and what is the source of these costs?
  - Cost effectiveness is far more difficult to justify in the Non-operational Consequences branch. If a task has been assigned, what is the basis for the cost-tradeoff analysis?
  - In the hidden-function branch a proposed task must ensure the level of availability necessary to reduce the risk of a multiple failure to an acceptable level. Is there a policy concerning this risk level that can be used to interpret adequate availability?

- Use of the default strategy (the term 'default' is here used according to the Nowlan and Heap (1978) definition):
  - Have failures, which may or may not always be evident to the operating crew, been classified as hidden?
  - Where it cannot he demonstrated that any anticipated secondary damage will not be critical, has the failure been assigned to the safety branch?
  - Have any opportunities been overlooked to assign on-condition inspections that may be partially effective in pre-empting functional failures?
  - Have all items for which the necessary information was unavailable been assigned to age exploration?
- General use of the decision logic:
  - One major problem is the tendency to select a familiar maintenance task and then work back through the decision logic to justify it.

    - This handicaps the analysis in two ways: on one hand, more of the tasks tend to stay justified, and on the other, the possibilities of new tasks are not explored.
    - Some analysts may have a strong preference for rework tasks and will specify them whether they are applicable or not. Others will favour on-condition inspections under all circumstances.
    - The auditor should look for signs of individual bias during the progress-review meetings, and by actually counting the numbers of each type of task selected by the various analysts.
- Sundry auditing concerns
  - Analysis of systems items
    - The chief difficulty in analysing systems items is confusion about the appropriate level or analysis and the functions of the specific item under consideration.
    - In the case of aircraft, Nowlan and Heap (1978) states that if more than 500 systems items have been classified as significant at the aircraft level, the list is probably too long, and if there are fewer than 200, it may be too short. If any subsystem includes more than half a dozen functionally significant items, their classification should be re-examined.
    - Another problem is finding the dividing line between one system and another.
  - Non-RCM program elements
    - The zonal inspection program should be audited to ensure that all zones in the equipment are included.
    - Zonal inspections are general visual inspections; do the tasks clearly describe the elements in the zone to be inspected?
    - The servicing and lubrication tasks should be audited for completeness, and any deviations from the manufacturer's recommendations should be substantiated.
  - The completed program
    - Additional questions may arise when the program is examined as a whole.
      - Do the tasks for each portion of the equipment/system cover all levels of maintenance?
      - Do they still make sense when they are viewed together?
      - Are there any gaps or overlaps?
    - Packaging presents special auditing problems, since the standards to be applied depend on the organisation, and other factors such as the number and location of maintenance facilities.
      - Have these been taken into account?

        - Auditing the packaging of the tasks is primarily a matter of determining whether the tasks have been scheduled as efficiently as possible for a given set of circumstances.
        - The impact of the maintenance program on the intended use of the equipment should not be overlooked in the audit. Will the proposed maintenance schedule permit each equipment to have the longest possible uninterrupted production runs?
- Auditing the ongoing program
    - Certain information systems must be established before the equipment goes into service:
        - A system for reporting failures, their frequency, and their consequences.
        - An age-exploration system.
        - A system for controlling the addition of new scheduled tasks to ensure that they meet RCM criteria before they are accepted.
        - A system for periodic reevaluation of all tasks in the program to eliminate those which are no longer needed.
        - A system for reviewing the content of the work packages as the size of the fleet grows.
        - A system for evaluating unanticipated problems and determining the appropriate action.
    - Auditing an ongoing maintenance program may require different skills and experience from those needed to audit program development. At this stage the auditor may often find himself in an adversary situation and will have to be both inquisitive and objective.
- Auditing new RCM analyses for older plant/business systems/equipment:
    - Older equipment may not be as sophisticated or complex.
    - It often has fewer fail-safe or damage-tolerant features.
    - Much of the age-exploration information is already available.
    - It is especially important for the auditor to determine that the new RCM program is not being developed by an analysis of the existing tasks, but represents a completely independent analysis of the equipment.
    - There should be much data of actual operating experience. Often one has to delve to mine the wealth of data, as it mostly are not readily available. Thus one of the major differences in auditing the analysis itself is to determine that the data were in fact used and were used correctly. The auditor should make sure that rework tasks, for example, have not been selected without an actuarial analysis of the data for the specific item.

- The number of tasks in the program will ordinarily be somewhat greater for in-service equipment. These should be reviewed thoroughly to make sure they are necessary.
- Older equipment may require more rework tasks than new equipment:
  - First, the results of age exploration will show the economic desirability of some additional rework tasks.
  - Second, the older designs may actually have more assemblies that show a wear-out pattern.
- Older equipment may also have a larger number of scheduled tasks for hidden functions because of older design practices.
- The number of on-condition tasks for older equipment may be slightly higher because ways of exploiting these relatively inexpensive inspections will have been found by experience for a number of items.

This concludes the investigation of the RCM methodology as found in the literature. The next chapter will seek to build on this basis to suggest a proposed improved methodology that will serve the interest of industrial users better.