CHAPTER

4

# Failure, Its Nature and Characteristics

In the last chapter we looked at aspects of reliability engineering that can be of use to the maintenance practitioner. We discussed some of the underlying principles that can help us identify reliability parameters from historical maintenance records. In order to apply this knowledge, it is useful to understand the nature of failure. In this chapter, we will examine the following.

- Failure in relation to the required performance standards; critical, degraded, and incipient failures;
- Significance of the operating context;
- Use of failures as a method of control of the process;
- Role of maintenance in restoration of desired performance;
- Incipiency and its use in condition-based maintenance;
- Age-related failures;
- System-level failures;
- Human errors and the effect of stress, sleep cycles, and shift patterns;
- Role of feelings and emotions and how these affect our reactions to situations.

## 4.1 FAILURE

### 4.1.1 Failure - a systems approach

Failure is the inability of an item of equipment, a sub-system, or system to meet a set of predetermined performance standards. This means that we have some expectations, which we can express quantitatively. For example, we can expect the discharge pressure of a centrifugal pump to be 10 bar gauge at 1000 liters per minute. In some cases, we can define our expectations within a band of acceptable performance. For example, the discharge flow of this pump should be 950-1000 liters per minute at 10 bar gauge. The performance standard may be for the system, sub-system, equipment, or component in question. These standards relate to what we need to achieve and our evaluation of the item's design capability and intrinsic reliability.

### 4.1.2 Critical and degraded failures

As a result of a failure, the system may be totally incapacitated such that there is a complete loss of function. For example, if a fire pump fails to start, it will result in the unavailability of water to fight fires. If there had been a real fire and only one fire pump installed, this failure could result in the destruction of the facility. In this case, the failure-to-start of the pump results in complete loss of function. As a second example, let us say that we have a set of three smoke detectors in an enclosed equipment-housing. The logic is such that an alarm will come on in the control room if any one of the three detectors senses smoke. If any two detectors sense smoke, the logic will activate the deluge system. It is possible that one, two, or all three detectors are defective, and are unable to detect smoke. When there is smoke, there is no effect if only one detector is defective, as the other two will activate the deluge. If two of them are in a failed state at the same time, the initiation of the deluge system will not take place when there is smoke in the housing. Lastly, with the loss of all three, even the alarm will not initiate. The loss of all three units will result in total loss of function, so this is a critical failure. If two of the three fail, the third can still initiate the alarm on demand. The operator then has the ability to respond to the alarm and initiate the deluge system manually. The system can still be of use in raising the alarm, so it has partial or degraded functionality.

### 4.1.3 Evident failures

When the impeller of a pump wears out, the operators can see the change in flow or pressure and hence knows about the deterioration in its performance. We call it an evident failure as the operator knows its condition. Similarly, an increase in the differential pressure across a filter or exchanger indicates an increase in fouling. When we take bearing vibration readings and plot the changes, it is possible to predict when it needs replacement. In each case, the operator knows the condition of the equipment, using their own senses or instruments. The operator, in this context, is the person who is responsible for starting, running, and stopping of the equipment. For example, the driver of an automobile is its operator.

### 4.1.4 Hidden failures

These failures, by contrast, are unknown to the operators during normal operation. Do you know if your automobile brake lights work? Similarly one does not know whether a smoke detector or a pressure relief valve is in a working condition at any point in time. A second event, such as a fire (causing smoke) will initiate the smoke detector, if it is in a working condition. If the vessel pressure exceeds the relief valve's set-pressure, it should lift. The standby power generator must start when there is a power failure. Will the pressure relief-valve lift or the standby generator start?

Hidden failures are also observed with protective instruments. Once equipment complexity increases, the designer provides various protective devices to warn the operator, using alarms, or bring it to a safe condition, using trips. These protective devices are rarely called upon to work and the operator will not know if they are working. These are subject to hidden failures.

If the operator is not physically present when the event takes place, is it an evident or hidden failure? For example, a pump seal may leak in a normally unattended unit. There will normally be some evidence of the leak, such as a pool of process liquid on the pump-bed. Merely because the operator was not present and did not see it does not change the event from an evident to a hidden failure. If the operator had been present, the leak would have been obvious, and a second event is not necessary. The question is not whether a witness was present, but whether the consequence occurred at the same time as the failure. To identify a hidden failure a second event must take place, and unless this condition occurs, it is an evident failure. Thus the time the operator sees the failure is not an issue.

To revert to the earlier question of the brake lights, you know that at the time you inspected the vehicle it was road-worthy, and the lights were working. If you ask a friend to stand behind the automobile while you press the brake pedal, you will soon know the answer. This is an example of a test on an item subject to hidden failures.

### 4.1.5 Incipient failures

If the deterioration process is gradual, and takes place over a period of time, there is a point where we can just notice the start of deterioration. Incipiency is the point at which the onset of failure becomes detectable. As the deterioration progresses, there is a point when the performance is no longer acceptable. This is the point of functional failure. The incipiency interval is the time from onset of incipiency to functional failure. When the failures are evident and exhibit incipiency, it is possible to predict the timing of functional failures.

## 4.2 THE OPERATING CONTEXT

The operating context describes the physical environment in which the equipment operates, demands made on it and the details of how it is used. The way in which we operate equipment has a bearing on how it performs, and affects its rate of deterioration. How close to the duty point does it operate? What is the external environment in which it operates? Does the internal environment affect its performance? What is the loading roughness? Does it have an installed spare unit that can come on stream if it fails? If the net positive suction head (NPSH) available to a pump is just acceptable, is the suction piping alignment such that the spare pump has as much NPSH as the duty pump? The answers to this type of questions will help define the operating context.

To illustrate this concept, let us take the example of an automobile or bus, and examine how we use it. For the purpose of this discussion, consider the following two contrasting requirements:

- We use it for long distance travel, mainly on freeways (highways, autobahns, or motor-ways);
- We use it for city travel only.

In the first case, the vehicle operates in a steady state, generally at cruising speed for much of its operating life. So the vehicle is operating predominantly at constant loads, well below duty point and with a smooth loading. In the second case, there will be frequent starts and stops, and driving speeds will be changing most of the time. The load on the engine will be variable due to the rapid changes in speed. The fluctuating power requirement from the engine means there will be more wear on the main elements of the power transmission, such as the clutch and gearbox. One should expect that brakes, tires, and indicator lights will need more frequent replacement.

We now add the driver profile, and the situation becomes more complex, for example,

- The driver has many years of experience, and has a 'clean' license, or
- The driver received the (first) license three weeks ago, and has already had one accident.

Turning to driving styles, we know that some drivers like to accelerate rapidly and use brakes frequently. Some are fond of taking corners at high speeds. Others prefer to cruise at a steady pace most of the time, use brakes infrequently, and take corners on all four wheels! Assume that you are buying a used car, and have the following options. One car belongs to a person who drives at a steady 40 mph, accelerates gently, and uses brakes sparingly. The other car, identical in make, model, vintage, and miles on the clock, belongs to a person who comes in screeching round the corner and slams on the brakes. If the price of the two cars is about the same, which one do you choose? It is an easy call, and you will decide quite quickly. The example highlights the significance of loading roughness, which contributes greatly to wear and tear.

External factors are next on our list of variables. These include dust or sand in the air, road surface, and weather conditions. One can see that the differences in performance as a result of these factors can be quite important.

Each of the changes in operating context will affect different sub-systems or components differently. For example, demanding driving habits will result in accelerated wear and tear on brakes, clutches, and tires, while dusty conditions will clog up air and lubrication filters more frequently. In an industrial context, the situation is quite similar. People wonder why identical pieces of equipment in the same process service perform differently. They believe that a pump is a pump is a pump! When we examine the differences in operating context, the reasons for performance variations become evident. As in the case of the vehicle, the operating context is one of the most significant contributors to performance.

## 4.3 THE FEEDBACK CONTROL MODEL

Let us examine how the driver of a vehicle controls it. The driver's eyes measure the position and attitude of the vehicle. These measurements are with respect to the edge of the road, other vehicles on the road, as well as any pedestrians who may be trying to cross the road. The change in position and attitude is being measured all the time. This information reaches the driver's brain, which compares these measurements with acceptance standards. The brain calculates the rate of change in position and attitude, and checks them against the norms. The driver's knowledge of the traffic regulations and past experience determine these acceptance criteria. The brain computes deviations from the norms, generating error signals. These signals initiate control actions, which are similar to those in section 1.4. The driver's brain instructs the hands to move the steering wheel, or the foot to press the brake or accelerator pedals, so that the car remains in control.

Other control systems follow a similar process, whether the unit in question is a battle-tank gun control or a chemical-process control system. Figure 4.1 shows a model illustrating the control mechanism.

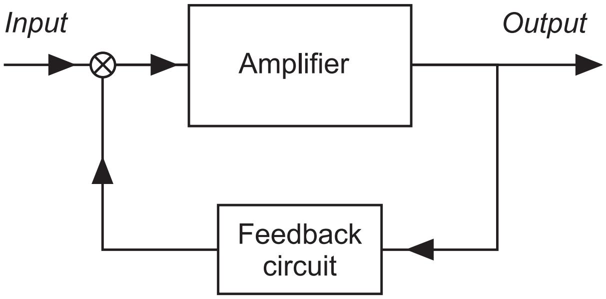

**Figure 4.1** Input signal, amplifier, output signal, and feedback loop.

## 4.4 LIFE WITHOUT FAILURE

Would it not be wonderful to have life without failure? The fewer the failures, the higher the intrinsic reliability that we can enjoy. A good designer tries hard to make the product or service as reliable as possible, within given economic and technical constraints.

A marble rolling along a smooth glass surface may roll on for a long time. However, controlling its movement can be difficult. Similarly, an astronaut doing a space-walk faces a handicap. In the absence of friction or gravity, it is very difficult to navigate, because the only way to do so is to use reaction forces, applying the principle of conservation of momentum. Thus a lack of resistance or opposition may make the process energy-efficient, but control is more difficult. One could extend this approach to explain why democracies are superior to dictatorships, or why market forces are better than price controls. Seen in this context, failures can be useful, as they help identify deviations from expected performance and hence the scope for improvement.

Failures are deviations that we can measure, and provide the means to control a process. Resnikoff[1] identified the significance of failures when he presented his well-known conundrum. This is the fact that we require information about critical failures to identify the correct maintenance work, the purpose of which is to avoid the same failures. Hence with perfect maintenance, such critical failures will never take place, so we can never collect the relevant data! The inability to collect the data required for this purpose can stymie any organization attempting to go along the path of continuous improvement.

## 4.5 CAPABILITY AND EXPECTATION

Every component, equipment, or system has an intrinsic design capability. The bold line in Figure 4.2 shows this graphically.

The demand or expected value may be below this level, shown by the dashed line in Figure 4.2. In this case there should be no problem meeting the

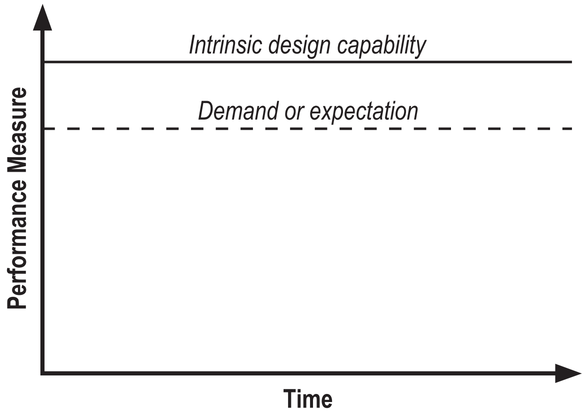

**Figure 4.2** Normal relationship of demand to capability.

demand. However the expectation may be higher than the design capability, as shown by the dotted line in Figure 4.3. In this case, we cannot achieve the expected values on a long-term basis. No amount of maintenance can increase the capability of the equipment to produce continuously above the intrinsic design levels.

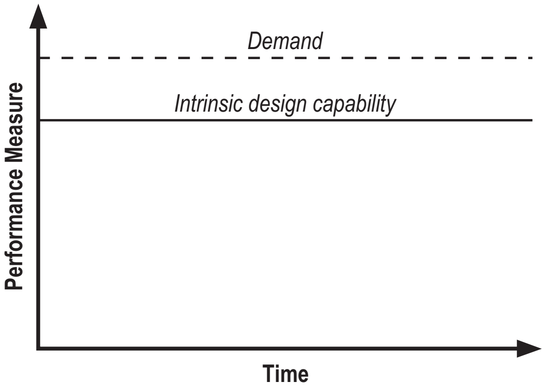

**Figure 4.3** Demand exceeds capability.

Designers tend to build in some 'fat,' stating a level of capability lower than the real value. This is partly due to the use of standard components, some of which are stronger than required, and partly due to built-in safety factors. When we exploit this 'fat,' there is a temptation to think that we are able to exceed the design values continuously. The reality is that this capability was always there, but the designers informed us differently.

Over time, the capability line will droop, due to fouling, wear, fatigue, or chemical attack.When this happens, some maintenance has to be done, to bring the capability up to the design level, as shown in Figure 4.4.

The demand profile may be flat, or as is more common, fluctuating, with peaks and troughs. We cannot meet the expected demand when the two lines intersect, so we need to do some maintenance at this time. Alternatively, we can do the maintenance in anticipation of this situation as illustrated in Fig.4.5.

The capability line will also exhibit some roughness. Thus there will be a spread or distribution of values, in the case both of the capability line and the demand line. These can be shown as bands of values as shown in Figure 4.6 and its inset. Normally, with smooth demand and capability lines, there is a single point of failure, shown by point B in the inset. With both curves having a band of values the earliest point of intersection is point A and the latest point C. There is, therefore, a range of points of functional failure. This leads to uncertainty in determining it and the lowest value will normally be chosen, so that we are on the 'safe side.'

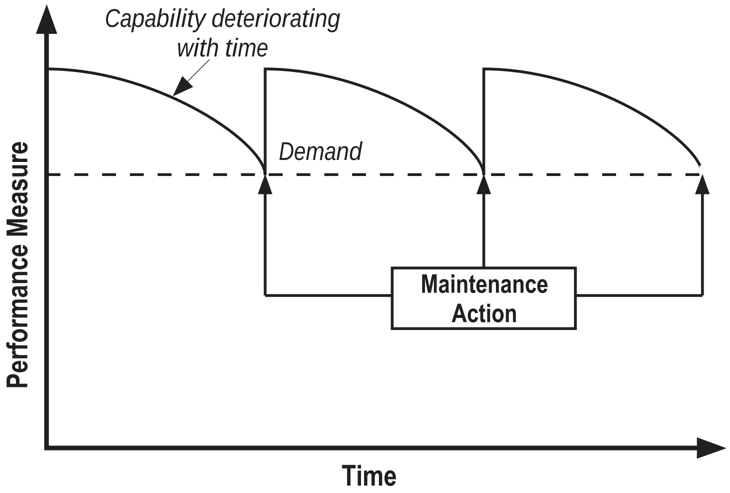

**Figure 4.4** Maintenance to restore capability

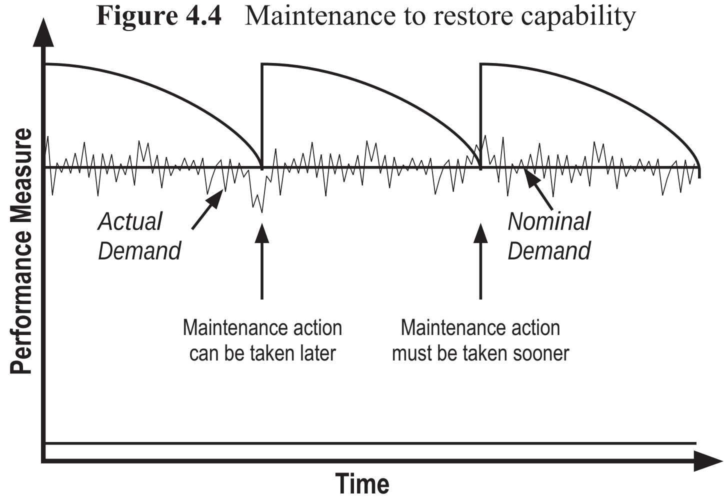

**Figure 4.5** Effect of demand fluctuations on maintenance timing.

## 4.6 INCIPIENCY

We mentioned incipiency briefly in section 4.1.5. Here we will examine the physical process in greater detail.

At the level of the smallest replaceable component, we will deal with items such as light bulbs, ball bearings, or structural welds. *Failure initiation is usually by fatigue or deformation caused by thermal or mechanical stress, or by chemical attack*. The rate of progression of the failure mechanism is variable, in some cases rapid, in others quite slow. Let us examine one or two common situations where we can observe the progress of the failure.

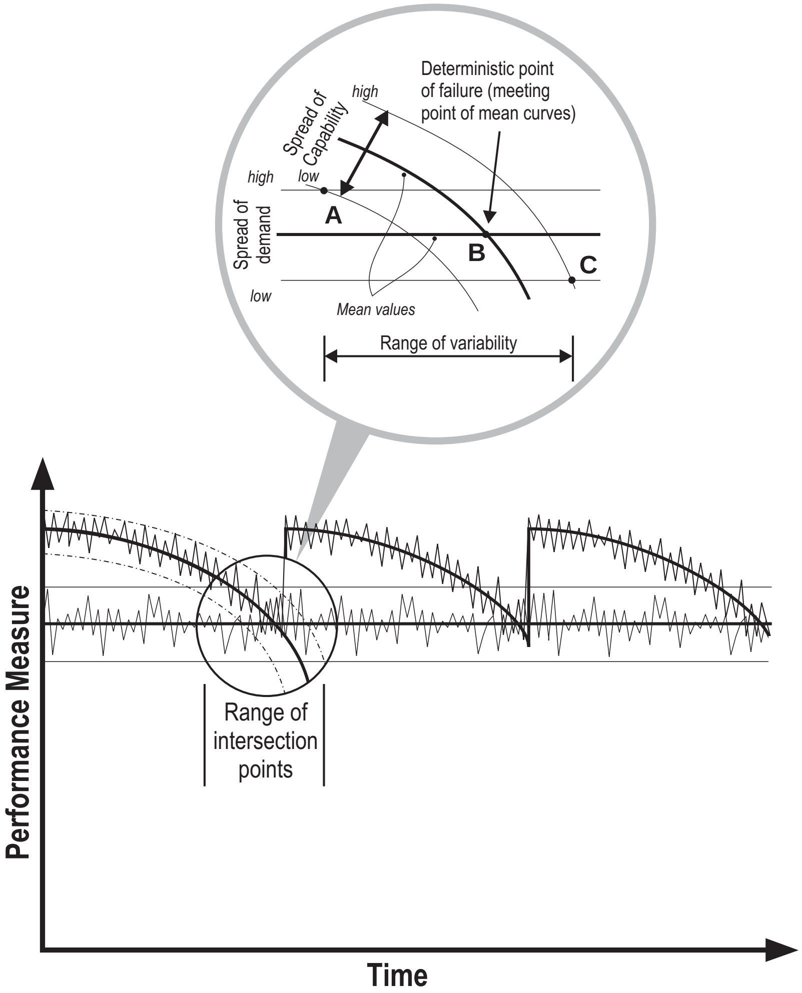

**Figure 4.6** Effect of fluctuations in demand and capability on the timing of maintenance.

The first example is of a road that has a small surface defect or unevenness caused by poor finishing. As vehicles pass over this unevenness, the tires enter the depression and then climb up to the original level. This causes an impact load on the road as well as on the vehicle suspension. The effect of this impact on the road is to damage it further, causing a deeper depression. The next truck gets a bigger bump, and causes even more damage to the road. If we do not carry out repairs, the depression eventually becomes a pothole, making it unsafe to drive on this section of the road. Figures 4.7 to 4.9 below illustrate the sequence of events.

The time when we notice the initial defect is the start of the incipient failure, denoted by point *x* at time $t_i$ in Fig.4.10 below. The droop of the curve shows the rate of growth of the pothole. At some point in time, this condition

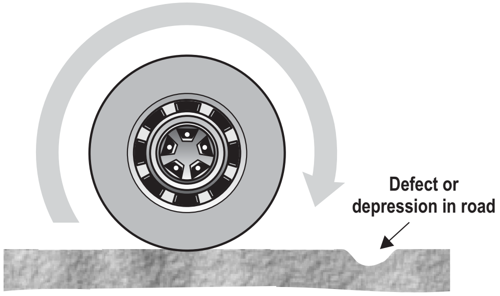

**Figure 4.7** How road surfaces get damaged.

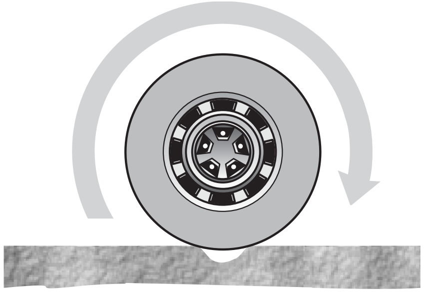

**Figure 4.8** Tires 'drop' into defect and climb.

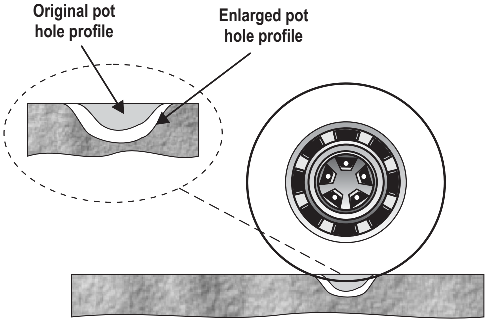

**Figure 4.9** The 'drop' energy damages the road further.

becomes unacceptable, as the road is no longer safe to use. This norm used to determine its acceptability is dependent on the operating context. The higher the speed of the vehicles and the greater their loading, the stricter are the acceptance standards. The dotted lines show the relative levels of acceptability, which are dependent on road speeds and loading. At the point of intersection with the curve, indicated by the point $y$ at time $t_f$, it is not safe to drive on the road any longer. In other words, it has failed. *The time taken for the condition to deteriorate from x to y, that is,* $t_f - t_i$, *is the incipiency interval*.

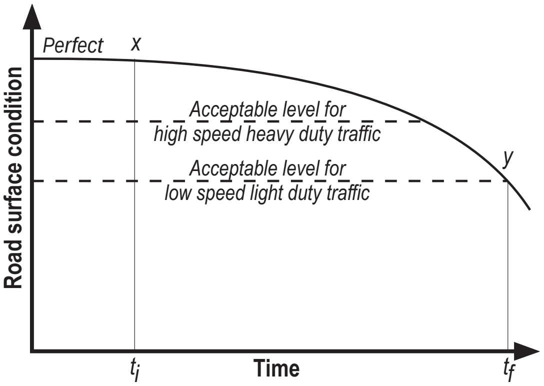

**Figure 4.10** Incipiency interval $(t_f - t_i)$

The second example is of a welded structure, such as a pressure vessel or steel frame of a building. When originally fabricated, some minor cracks would have remained in the welds. At the time of construction, these cracks either escaped detection or were not serious enough to trace and repair. After commissioning the structure, these welds experience loads, which can fluctuate in magnitude, direction or both. When there are cracks in the welds, the effective cross sectional area is smaller, resulting in higher stresses. At the tip of the crack (refer to figure 4.11) the material can become plastic due to stress concentration.The most stressed part of the weld will yield, resulting in the crack propagating further. This raises the stress just beyond this point, ensuring the continuous propagation of the crack. In due course, the crack can grow to such an extent that the weld as a whole is no longer able to perform its function.

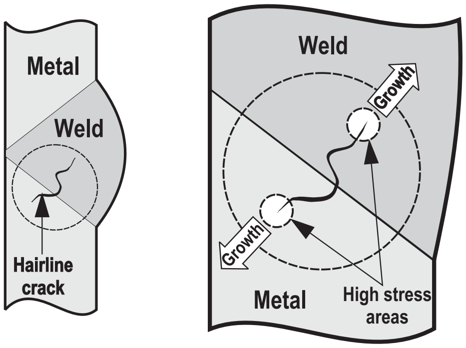

**Figure 4.11** Crack propagation in a weld.

The incipiency interval may be very short, as in the case of light bulbs, or very long, as in the case of weld crack propagation. A large number of failures have incipiency intervals ranging from weeks to several months or years. Bearing failures, general corrosion, and weld crack propagation are all examples of such failures. Nowlan and Heap[2] refer to the point x in Figure 4.10 as the point of potential failure, and the point y as the point of functional failure. Moubray[7] refers to it as the P-F curve, where points P and F correspond to points x and y in Figure 4.10. The range of variance in incipiency is shown in figure 4.12.

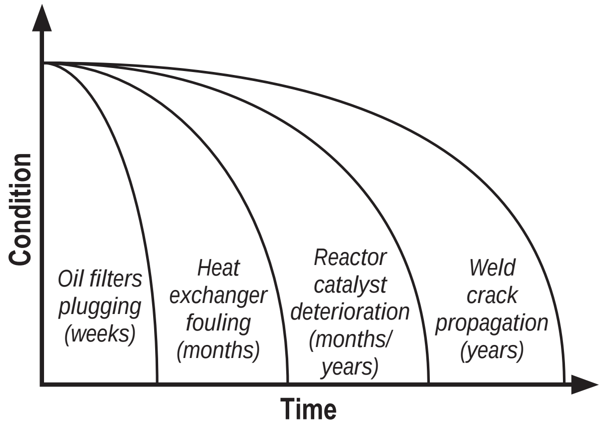

**Figure 4.12** Examples of incipiency intervals.

Even in the case of a single failure mode in a given operating context, the droop of the incipiency curve may vary. Thus, there is a range of incipiency intervals, as illustrated in Fig.4.13. This range introduces uncertainty in determining the incipiency interval.

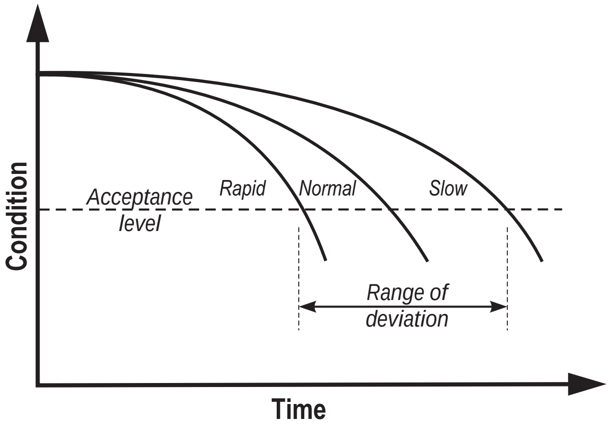

**Figure 4.13** Variations in incipiency intervals.

## 4.7 LIMITS TO THE APPLICATION OF CONDITION MONITORING

When the incipiency is very short, the time available to plan or execute maintenance action is also very small. In such cases, it is difficult to plan replacement before failure by monitoring the component's condition. When incipiency intervals are in weeks, months, or years, condition monitoring is often an effective way to plan component replacement. Condition monitoring is feasible when it is possible to measure the change in performance, using human senses or instruments. It follows that we cannot monitor hidden or unrevealed failures.

*Proponents of condition-based maintenance are correct when they highlight their ability to predict failures. Any predictive capability enhances the decision making process. However they sometimes give the impression that condition monitoring systems will solve all our problems.* We know that all failures do not lend themselves to condition monitoring. The failure must exhibit incipiency, it must be feasible to measure it, and the interval must be of reasonable duration. We must always ask the providers of condition monitoring services to demonstrate how they meet these requirements.

## 4.8 AGE RELATED FAILURE DISTRIBUTION

A system consists of many pieces of equipment, each of which has several components. Each component can fail in one or more ways. In Chapter 3, we looked at the six failure patterns identified by the Nowlan and Heap[2] team. You will recall that these failure patterns are plots of the hazard rates against time. The author obtained similar results in a study of failures in the offshore oil and gas industry. The experience of other industries is not available in the public domain, but we can expect broadly similar results.

Prior to the Nowlan and Heap study, the belief was that all failures followed the so-called bath-tub curve. Their results showed that this pattern was only applicable to 4% of all the failure modes.

Fourteen percent showed a constant failure pattern, and if we ignore the failures that took place early in life, a further 75% also followed this pattern. The remaining 11% (including 4% of the bath tub) of the failure modes exhibited a distinct relationship to age. Should we concern ourselves with this relatively small proportion of failures that exhibit an age-relationship?

To answer this question, we need to know whether any of these failure modes could result in serious consequences. If so, they acquire a new level of respect. With a skewed distribution, a strategy based on an assumed constant failure pattern will not be satisfactory. *Therefore, we cannot assume that all failures exhibit a constant hazard rate pattern, as long as any of the remaining 11% matter.*

## 4.9 SYSTEM LEVEL FAILURES

When we assemble components to build equipment, each component failure-mode affects the overall failure rate. These individual component failure-modes may have exhibited a distinct age-related failure pattern. When any failure takes place, we replace the affected part with a new one. In an ideal case, we do not replace any of the other components at this point. The latter are at different stages of deterioration in their own life cycles. One of these will fail some time thereafter because it has reached the end of its life. We replace it and start a new cycle, while other components continue from their partly worn-out state.The result is that at the assembly level, the failures tend to follow the exponential distribution.

The concept of Mean Time To Failures, or MTTF, is worth further consideration at this point. As discussed in Chapter 3, the mean does not tell us much about the distribution. With a given sample, many of the failures could have taken place early or late in terms of age. In such a case, the use of the mean distorts the picture, because one may wrongly infer that the failures take place uniformly over the life. Hence, the use of MTTF without a full understanding of the distribution may lead to inappropriate decisions.

*When the hazard rate is constant and the distribution is exponential, it is perfectly acceptable to use the MTTF.* At this point there is (approximately) a 63% probability that the component has failed, and only a 37% probability of survival. In cases where the consequences of failure are high, we must do whatever we can to reduce or eliminate them. If the failure is evident and exhibits incipiency, for example, as in a ball bearing, we can take vibration or other condition monitoring action. If the failure is hidden, as for example, in a gas detector, we carry out a test, or a failure finding task. We must plan preventive maintenance action well before t = MTTF, because we cannot accept a 37% probability of survival at the time of the test or repair. The lay person often thinks of the MTTF as the expected time of failure and, therefore, the maintenance interval, which is clearly not the case.

## 4.10 HUMAN FAILURES

Nearly three quarters of all accidents are due to the action (or inaction) of human beings. We cannot wish it away, as it is too large a contributor to ignore. Human beings are complex systems, with hundreds of failure modes. In the following discussion we will use the terms *human error* and *human failure* interchangeably.

The causes of human error are many and varied. Lorenzo[3] categorizes them as random, systematic, and sporadic. We can correct random errors by better training and supervision. A shift in performance in one direction indicates systematic variability. We can reduce these by providing a regular performance feedback. Sporadic errors are the most difficult ones to predict or control. In this case, the person's performance is fine for most of the time. A sudden distraction or loss of concentration results in sporadic error.

There is an optimum level of stress at which human beings perform well. *A certain level of stress is necessary to keep us alert, active, and expectant.* We call this facilitative stress. Too high a stress level can be as a result of physical or psychological pressures. This may result in tiredness and lack of concentration. Too low a stress can be due to the work being repetitive, intellectually undemanding, or otherwise boring. During World War II, the British Royal Navy noted that submarine lookouts became ineffective after about 30 minutes, as they could not remain alert, The lookouts knew that their own lives depended on their vigilance, so motivation was not an issue.

Swain and Guttmann[4] give the following examples of psychological stress:

- Suddenness of onset
- Duration of stress
- Task speed
- Task load
- High jeopardy risk
- Threats of failure, loss of job
- Monotonous, degrading, or meaningless work
- Conflicts of motives about job performance
- Reinforcement absent or negative
- Sensory deprivation
- Distractions such as noise, glare, flicker, color, or movement
- Inconsistent cueing

Each person is slightly different and thrives under different levels of stress. However, a number of the stress factors affect many people in similar ways.

In order to reduce human failures, we have to address the factors contributing to stress. By doing so, we can produce the right environment for each person. In most cases, we will not be able to influence stress caused by domestic matters, so we will focus on those at work. Job enrichment deals with the elimination of boredom and unacceptably low stress levels. We can attribute the remaining problems to high stress at work.

Control room operators perform critical functions. During plant upsets, startups, and shutdowns, their skills are in demand. We use alarms to catch their attention when things go wrong. Designers of control rooms have to take care to minimize the number of alarms they install. If too many alarms come on too quickly during a plant upset, operators can lose concentration and react incorrectly, thereby worsening the situation. In an article entitled 'How Alarming!,' Bransby and Jenkinson[5] report the results of a survey. They studied 96 control room operators in 13 different plants in the U.K. Their findings, listed below, indicate that we have to devote more attention to this issue at the design stage.

- In an average shift, during steady operations, operators receive an alarm about every two minutes;
- Many of these are repeats of ones that occurred in the previous five minutes;
- Operators stated that many of them were of little value to them, and that eliminating about 50% would have little or no effect;
- Following a plant disturbance, they estimated that there were about 90 alarms in the first minute and seventy in the next ten minutes;
- About half the operators said that they felt forced to accept alarms during plant upsets, without reading or understanding them;
- During the survey, they observed one such plant upset. The operator did not make a full check of the alarms for about half an hour. This behaviour was consistent with that reported by the others in the survey.

Since the purpose of the alarm is to alert the operator, these results indicate that the designers have failed in their objectives. The authors state that improvements are possible, and that a variety of tools are available. Some of the simpler ones include tuning-up limit values and dead-bands, and adjusting priorities. The use of logic to suppress some non-essential post-trip alarms is also possible. As an example, they state that a review of the alarms resulted in a 30% reduction in the number of alarms.

One of the causes of human failures is tiredness, and this is often due to sleep deprivation. The human body operates with the help of a biological clock. Shift work can disturb normal (or circadian) sleep cycles. As a result, the reaction to stimuli can be slow. This can affect the ability of the operator to respond to a rapidly developing scenario. Night shift workers are more susceptible to this problem than the rest, because of the disturbance to their circadian rhythm. While there is no direct cause and effect relationship established, we note that some of the worst industrial disasters including Piper Alpha, Bhopal, Chernobyl, Three-mile Island, and Exxon Valdez occurred in the silent hours. This does not automatically mean that it is unsafe to work at night. Night-shift workers have completed many millions of hours of work without any incidents. It is the combination of circumstances that matter, so one must view this in context. Since we cannot eliminate night shift work, especially in continuous process plants, we have to try to understand the risks, so that we can take suitable steps to minimize them.

A factor affecting sleep cycles is the way we arrange shift patterns. Lardner and Miles[6] have explained why some shift patterns are superior to others from an ergonomic point of view. They propose a nine-day cycle, with 2 days each in the morning, afternoon, and night shifts, with a

3-day ‘weekend’ following the night shift. The ‘weekend’ may turn out to be in the middle of the week. They argue that this pattern is superior to the alternative 28-day cycle, which is quite common. The 28-day cycle consists of 7 night shifts and 7 evening shifts, followed by a 2 day ‘weekend’ after each block. This is followed by 7 morning shifts and a 3 day ‘weekend.’

Human errors occur due to a number of reasons, and lack of knowledge and experience are not necessarily the most common. Motivation and morale are often key issues to manage. Pride in work, a sense of being wanted, and being treated fairly are all important considerations. We all want ‘user friendly’ software; similarly, staff want managers who are ‘people friendly.’ When this is so, we are likely to experience lower absenteeism or sickness, better participation in team effort and suggestion schemes, lower accident rates, and higher productivity.

What makes human beings distinctly different from machines is their ability to think, often in a very creative manner. Feelings and emotions change the way a person responds to identical stimuli over time, and makes it hard to predict behavior. We have provided a brief introduction to the subject in this chapter and readers can refer to Lorenzo’s excellent guide for a more detailed discussion. A check-list of potential causes of human errors is available in Appendix 4-1.

## 4.11 CHAPTER SUMMARY

We began this chapter by defining failure in relation to the required performance standards. Failures can be critical, causing total loss of function, degraded where the loss is partial, or incipient where progressive deterioration has commenced, but will take some time before there is loss of function. We note the significance of the operating context, and how this explains why identical items of equipment perform differently. We saw how failures themselves provided a means of control on the process.

Our next topic is the role of maintenance in achieving the desired equipment performance, as long as it is lower than its capability. We discussed incipiency, and its use in condition-based maintenance, using some common examples to illustrate the concepts. Thereafter, we discussed age-related failures.

Finally, we looked at human errors, perhaps the most complex issue relating to failures. We noted that there is an optimum level of stress required to keep human errors as low as possible. The work done by experts on sleep cycles shows us how they can affect the body’s natural rhythm. The experts state that some shift patterns are superior to others when planning 24-hour coverage for continuous process plants.

Feelings and emotions play a major role in affecting the way people react to situations. Therefore, managers have to focus on motivation and morale, which are key issues in minimizing human failures.

## References

1) Resnikoff, H.L. 1978. Mathematical Aspects of Reliability Centered Maintenance. Dolby Access Press.
2) Nowlan, F.S., and H.F. Heap. 1978. *Reliability-Centered Maintenance.* U.S. Department of Defense. Unclassified, MDA 903-75-C-0349.
3) Lorenzo, D.K. 2001. *A Manager's Guide to Reducing Human Errors: Improving Human Performance in the Process Industries.* API Publication 770.
4) Swain, A.D., and H.E. Guttmann. 1983. *Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications.* NUREG/CR-1278-F SAND80-200. 3-34.
5) Bransby, M., and J. Jenkinson. 1998. "How Alarming!" The Chemical Engineer. January 15: 16.
6) Lardner, R., and R. Miles. 1997. "Better Shift Systems." The Chemical Engineer. September 11: 28.
7) Moubray, J. 2001. *Reliability-Centered Maintenance*. Industrial Press, Inc. ISBN: 0831131462.