You can't condition monitor random failures

RCM Notes series

Dr Mark Horton, Numeratis.com, January 2012

1 Introduction

Are you ever stumped for a good story at a reliability gathering? I thought not; but I'm sure that you can still appreciate a conundrum that can be kept safely tucked away until you meet someone who thinks he or she knows the very last wrinkle about reliability and risk.

Here's a question: why can't you apply condition-based maintenance to random failures?

2 What?

Yes, you read it correctly first time.

If you know anything about the development of modern reliability-centred methods, you are probably aware that the shift from age-based maintenance to condition-based methods happened for a number of reasons.

- The development of new, cheaper and more portable technologies which detect the signs of incipient failure
- The ability of condition-based methods to maximise the operational lifetime of assets, when compared with age-based overhaul and replacement
- The recognition that age-based methods do nothing to prevent failures that do not have a failure rate that increases with age

This short essay focuses on the third driver: that failure patterns that are not age-related cannot be managed by replacing or overhauling assets at fixed intervals. At best (if the failure rate is constant), age-based maintenance does nothing at all; at worst, when the failure rate decreases with age, replacing an asset or component increases the overall failure rate experienced.

3 Tell me something that I don't know!

Not just yet. That would spoil the surprise.

One of the key characteristics of modern equipment is that a large proportion of failure modes have either random (flat) failure rates, or they have failure rates that reduce with age. The story that most of us tell to the uninitiated is that the only proactive maintenance strategy that can manage these failures is condition-based maintenance. To demonstrate the principle, we draw a graph showing a flat hazard rate from time zero to infinity (well, as close to infinity as we can get on a slide). In Nowlan and Heap RCM terms, this is failure pattern E.

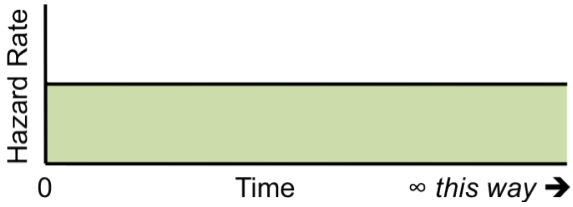

We imagine some out-of-date maintenance planner who is mad enough to replace the equipment after a fixed interval. We ask what the planner has achieved by replacing the equipment, and wait for the answer, "Nothing!"

Then we draw the diagram again, but this time showing an incipient failure at some time. It's a random failure pattern, so it doesn't matter when the failure starts to occur. If you imagine that this is a bearing failure, the "potential failure" might

be the point at which we can identify vibration due to deterioration.

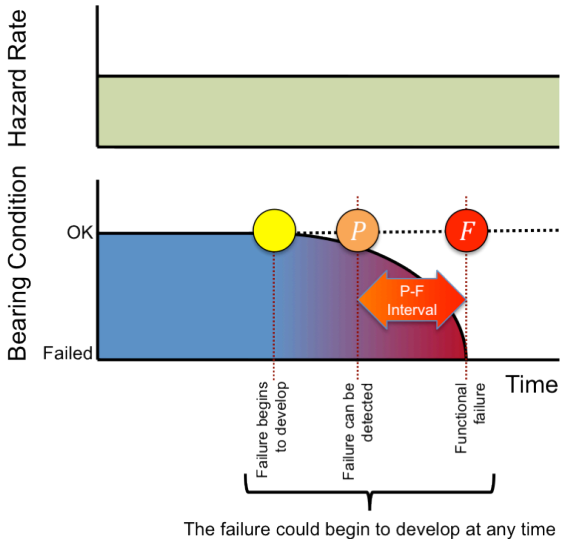

We state that there is a known minimum time between the incipient (or potential) failure and the full, functional failure that we are trying to avoid. This means that, if we look for potential failures at an interval less than P-F and have long enough to avoid the consequences of the functional failure, the failure can be managed *in spite of the fact* that the failure is random and we therefore don't have any idea when it might happen. We then congratulate ourselves for overcoming one of the curses of second generation maintenance [see John Moubray's book *Reliability-centred Maintenance*].

If you know all of this, it may come as a surprise to learn that it isn't possible to manage a *purely* random failure mode through condition monitoring.

The argument is so simple that, if you haven't worked out why by this point, you'd better get ready to kick yourself.

4 What's new?

Patience.

First, let's have a brief recap.

A random failure is one whose hazard rate is constant from time zero to infinity, like this.

One of the requirements that a condition monitoring technique must fulfil is that there is a clear potential-to-functional failure (P-F) interval. Specifically, there must be a *minimum* identifiable P-F interval which is used to determine the condition-based task interval. Otherwise failures can start to develop and go all the way through to functional failure before they can be detected.

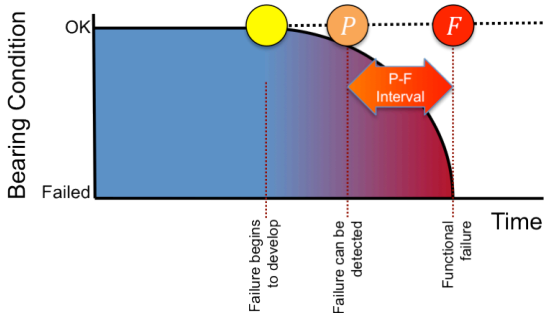

That's all the background we need: just two principles. Now imagine a new component that has just been installed. If we say that one of its failure modes can be subjected to condition monitoring, we must be able to define a minimum P-F interval (because otherwise condition-based maintenance would not be technically feasible). If we can define a minimum P-F interval T, and the new component is delivered without any incipient failures, it means that functional failure cannot occur before T, and therefore the hazard rate is zero until at least one P-F interval has passed.

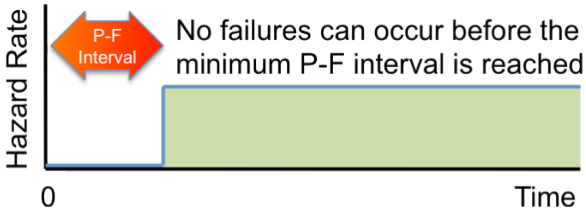

Therefore the hazard rate cannot be random, because it is zero for some time after t=0. In fact, the hazard rate must be more like Nowlan and Heap failure pattern D, where a period of low initial failure rate is followed by a flat hazard rate.

Alternatively, assume that a component shows a truly flat hazard rate from time zero onwards. If the hazard rate is flat, then it must be possible for a functional failure to occur within any stated P-F interval after the new component is installed. Therefore it must be impossible to define a minimum P-F interval, and condition-based maintenance is not technically feasible.

5 So what?

What practical application does this argument have? Almost none at all, except for bragging rights in understanding irritating, odd, dark corners of reliability theory. As long as the expected mean time between failures of a component is much longer than the failure mode's P-F interval (as it will be in almost all realistic cases), the effect on the choice of a maintenance policy is near zero.

6 References

Moubray, John (1997). *Reliability-centered Maintenance*. New York: Industrial Press.

Nowlan FS and Heap HF (1978). *Reliability-centered Maintenance*. San Francisco: Dolby Access Press.