TPK4120 - Lecture summary

Jørn Vatn
eMail: jorn.vatn@ntnu.no

Updated 2022-01-22

Chapter 12 - Preventive Maintenance

This memo is based on the book: System Reliability Theory - Models, Statistical Methods, and Applications by Rausand, Barros and Hoyland (2021). John Wiley & Sons, and in particular on Chapter 12.

The objective of this chapter is to demonstrate aspects of maintenance as part of a reliability analysis. In addition to passively treat maintenance as part of the reliability, we will also investigate some models for maintenance optimization.

Note that elements from Chapter 9 is included in the following presentation.

Maintenance

Definition: The combination of all technical and management actions during the life cycle of an item intended to retain the item in, or restore it to, a state in which it can perform as required.

Maintenance is important to achieve a high availability. Generally availability depends on the following factors:

1. Inherent reliability (e.g., quality, type of material used and design principles)
2. Maintainability (how easy it is to perform maintenance)
3. Maintenance support (resources, spare parts etc)

Maintenance Categories

The maintenance is often categorized into:

1. Corrective maintenance (CM), i.e., tasks performed as a result of a detected item failure or fault, to restore the item to a specific condition. CM tasks may be carried out *immediately* or be *deferred*.

2. Preventive maintenance (PM), i.e., planned maintenance tasks performed prior to failures. The activities are carried out in order to reduce the probability of failure, or increase the mean time to failure (MTTF). There are several types of PM tasks:

   (a) Age-based

   (b) Clock-based (calendar based)

   (c) Condition-based

   (d) Opportunity-based

   (e) Overhaul, e.g., as part of a turnaround

3. Predictive maintenance, i.e., maintenance based on prognoses for the degradation of the item.

Note that the categorization varies from standard to standard, e.g., some standards include predictive maintenance as part of condition-based maintenance.

Preventive Maintenance Policies

A preventive maintenance policy is a strategy that aims at minimizing the long run cost. A policy both deals with qualitative issues like replace an item periodically at a given age, and quantitative issues like what age that should be. The classical maintenance policies were basically considering age or calendar time as the decision variable to use in the optimization. In light of "predictive maintenance" the condition of an item and future operational loads are becoming more important in order to minimize long run cost. Examples of both types of models will be investigated.

Preventive Maintenance

Definition: Maintenance carried out at predetermined intervals or according to prescribed criteria and intended to reduce the probability of failure or functional degradation of an item.

Terminology and Cost Function

- *Maintenance task:* A specific task to maintain an item determined by "what, where, how and when". A task is part of the task space, $\mathscr{A}$, i.e., $\mathscr{A} = a_1, a_2, a_3, \dots$

- *Maintenance decision:* A process δ to select a specific maintenance task $a_i \in \mathscr{A}$. δ depends on available data $\mathscr{D}$, cost, operating conditions etc.

- *Maintenance strategy:* An overall framework describing how the maintenance decision problem shall be approached. A strategy embraces an objective function, often denoted the cost function:

- *Cost Function:* $C = C(a, \delta, t, \mathscr{D}, \mathscr{D}_{\mathrm{OC}}, t_{\mathrm{cal}}, \ldots)$. In addition to the maintenance task and the data the cost function depends on the time t of executing the maintenance, the operational context $\mathscr{D}_{\mathrm{OC}}$, the calendar time $t_{\mathrm{cal}}, \ldots$ (e.g., inside / outside working hours) and so on. A specific note is made regarding the notation used for the time. In the textbook the default notation is to use t or t_0 for the time axis, but in many other presentations we use τ to denote time, for example the length of a maintenance interval. Also note that time may be multi-dimensional, for example if we carry out both a failure-finding-task and a replacement-task.

To optimize maintenance we would like to minimize the cost per unit time. In many situations this will be to minimize the expected maintenance cost in a renewal period divided by the expected length of the renewal period:

$$C = \frac{\mathrm{E}[C(T_R)]}{\mathrm{E}(T_R)}$$

Age and calendar based policies

In this presentation the maintenance interval or replacement period is denoted τ whereas the textbook use the notation t_0. However the arguments are similar.

Age Replacement Policy - ARP

In the age replacement policy an item is replaced or overhauled to an as-good-as-new condition when the item reaches a specified age. We usually consider a replacement rather than an overhaul, but the situation is the same if an overhauled item is as-good-as-new after an overhaul. This age is denoted τ and the challenge is to find the optimal τ, say τ^*. The situation is characterized by:

- The item is replaced when it reaches the age τ

- If the item fails before a periodic activity, the unit is replaced and the "maintenance clock" is set to 0

- In Figure 1 T_1 and T_2 are times-to-failure where the item is replaced

- The cost of a preventive replacement is c

- The cost of a corrective replacement, i.e., replacing a failed item is $c + k$

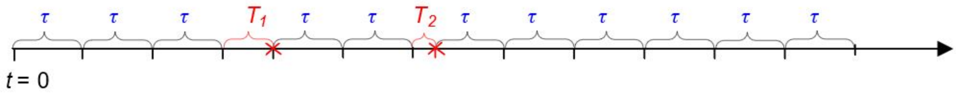

Figure 1: ARP

In the modelling we let $f(t)$ denote the life time distribution of the item, and we assume an item to be as-good-as-new after a replacement. The time between two consecutive replacements is called a replacement period. This period is stochastic, and the mean time between replacements is:

$$\mathrm{MTBR}(\tau) = \int_0^\tau t f(t) dt + \tau \Pr(T > \tau) = \ldots = \int_0^\tau (1 - F(t)) dt = \int_0^\tau R(t) dt$$

where T is the (potential) time to failure, and we have used partial integration.

For each replacement period we always have to pay the cost c. If a replacement period ends with a failure, we have to pay an extra cost k. The probability of paying the extra cost is $\Pr(T \le \tau) = F(\tau)$. The long run cost per unit time is then given by:

$$C_A(\tau) = \frac{\text{Cost in a cycle}}{\text{Expected length of a cycle}} = \frac{c + kF(\tau)}{\int_0^\tau [1 - F(t)] dt}$$

Numerical methods are required to minimize $C_A(\tau)$

Numerical methods for $\mathrm{MTBR}(\tau)$

The `ARPexample.xlsm` MS-Excel file available on Blackboard contains some Visual Basic code (VBA) for numerical integration required for calculating $\mathrm{MTBR}(\tau)$. The essential code is:

```
Function MTBR(Tau As Single, alpha As Single, lambda As Single)
MTBR = NumInt( iMTBR, 0, Tau, alpha, lambda, 0)
End Function
```

where `iMTBR()` is a function returning the integrand in $\int_0^\tau (1 - F(t)) dt$ where a Weibull distribution is assumed in the code. The `NumInt()` function performs the numerical integration.

Note the following:

- VBA is not very advanced, therefore a special trick is used to pass a function to another function, i.e., this is what we need if we have a general purpose integration function we would like to apply for any integrand function we are going to write. In more "math-oriented" languages such as MATLAB, the use of so-called "function pointers" is more easy to use and understand

- When we pass a function as an argument to the `NumInt()` function we need to pass the function name as a text string. The integrand is the first argument.

- In our integration function we may then use `Application.Run` function to "execute" the function specified by it's name. This is a rather slow approach, so a faster approach is to test the function name by a `Select Case` statement.

- The second and third argument of the `NumInt()` function are the lower and upper integration limit respectively

- We always need to pass exactly 3 additional arguments (parameters) to the `NumInt()` function. We only need 2 of them here, but we have to send the third one which here is set to 0 (how many parameters to pass can be changed, but it needs hard coding in the implementation).

- The integrand function, here `iMTBR()` is in fact a `SUBROUTINE`. Since a subroutine can not return a value, we have to use a special function `fReturn` to return the value, this is shown below:

```
Sub iMTBR(x As Single, alpha As Single, lambda As Single, dummy As Single)
fReturn Exp(-((lambda * x) ^ alpha))
End Sub
```

Block Replacement Policy - BRP

In a block replacement policy we would like to replace an item at a specific point of time. The argument for this could be that we have many identical components, and it is more convenient to perform the preventive maintenance at the same time (i.e., a block replacement). Another argument for a block replacement policy could be that this is much easier to manage by our computerized maintenance management system (CMMS). Figure 2 illustrates the situation. T_1 and T_2 are failure times, but they will not affect the time of a preventive activity. It seems a bit "waste" of useful life to replace at time 4τ but this may be defended by lower administrative cost.

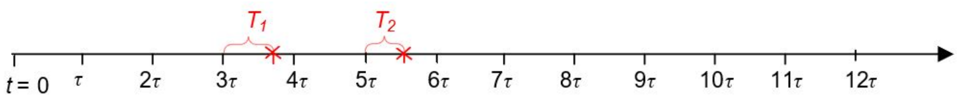

Figure 2: BRP

The situation is slightly different from the ARP, and also a different cost structure is used:

- The item is replaced every τ time unit

- An item may fail one ore more times between periodic replacements, if this happens it it assumed that the item is immediately repaired or replaced to a-good-as-new condition
- The cost of a preventive replacement is c
- The cost of a corrective replacement, i.e., replacing a failed item is k

In this situation the replacement period is always τ. In each period we always have to pay the cost c. In addition we pay the cost k for each failure. The expected number of failures is given by the *renewal function*, $W(\tau) = \mathrm{E}(N(\tau))$. Thus the average cost per time unit is:

$$C_B(\tau) = \frac{c + kW(\tau)}{\tau}$$

Note that $W(\tau)/\tau$ is the average expected number of failures per time unit when the item is replaced every τ time unit. This is often written:

$$\lambda_{\mathrm{E}}(\tau) = \frac{W(\tau)}{\tau}$$

For small values of τ compared to the MTTF, it is unlikely that we have more than one failure in an replacement period. This means that the expected number of failures in a replacement period is given by the average value of the failure rate function, $z(t)$. If we assume that failure times are Weibull distributed, we obtain:

$$\lambda_{\mathrm{E}}(\tau) = \left(\frac{\Gamma(1+1/\alpha)}{\mathrm{MTTF}}\right)^{\alpha} \tau^{\alpha-1}$$

where $\mathrm{MTTF} = \Gamma(1+1/\alpha)/\lambda$. In this situation it is straight forward to find an analytical solution for the optimal replacement period:

$$\tau^* = \frac{\mathrm{MTTF}}{\Gamma(1+1/\alpha)} \sqrt[\alpha]{\frac{c}{(\alpha-1)k}}$$

Note that the numerical approach is slightly different from Example 12.2 in the textbook. The analytical solution by τ^* deviates slightly from the example. For example if $k/c = 10$ and $\alpha = 2$ we obtain $\tau^* = 0.36\mathrm{MTTF}$, whereas the textbook example gave 0.39MTTF.

In Chapter 10 the fundamental renewal equation, $W(t) = F_T(t) + \int_0^t W(t-x) f_T(x)dx$ was given, and in case we have a reasonable initial approximation, for $W(t)$, say $W_0(t)$ we may use the following iteration scheme:

$$W_i(t) = F_T(t) + \int_0^t W_{i-1}(t-x) f_T(x)dx$$

to obtain better and better solutions for $W(t)$. We may use

$$W_0(t) = \lambda_{\mathrm{E}}(t)t = \left(\frac{\Gamma(1+1/\alpha)}{\mathrm{MTTF}}\right)^{\alpha} t^{\alpha-1} t = \left(\frac{\Gamma(1+1/\alpha)}{\mathrm{MTTF}}\right)^{\alpha} t^{\alpha}$$

Degradation Models

When maintenance is condition based we will utilize the understanding of degradation of the item to determine appropriate maintenance action and time for maintenance.

In the following we distinguish between:

- $X(t), t \geq 0$ = a stochastic process describing the actual degradation of the item at time t
- $Y(t), t \geq 0$ = a stochastic process describing the *measurements* of degradation of the item at time t

where the measurements typically contain noise. Degradation could be crack lengths, corrosion depths, vibration levels etc. In some situations it may be difficult to distinguish the variation in the degradation process from the measurement errors. In this presentation we will not explicitly consider imperfect measurements of the degradation in order to make simple presentations. We will therefore not make an explicit definition of what is the difference between $X(t), t \geq 0$ and $Y(t), t \geq 0$.

Remaining Useful Lifetime

In degradation modelling (prognostics) the term Remaining Useful Lifetime (RUL) is introduced. RUL(t_j) is a stochastic variable that measures the time from t_j until the item is not "useful" any more. "Useful" need to be defined, for example a failure, or some other bad performance. Since RUL is a stochastic variable, we often need the distribution function, i.e.,

$$\Pr(\mathrm{RUL}(t_j) \leq t) = F_{\mathrm{RUL}(t_j)}(t) \tag{1}$$

where t_j is the current time, and t is a future point of time, typically measured from t_j as the starting point.

In the textbook various data-driven approaches are given in order to assess the RUL distribution. In the following a limited number of ideas are pursued. The definition in Equation (1) will not help us since there is no explicit link to the condition or degradation of the item. A more explicit definition of RUL is therefore:

$$\mathrm{RUL}(t_j) = \min\left\{h : X(t_j + h) \in \mathcal{X}_l\right\} \tag{2}$$

where $\mathcal{X}_l$ is the set of states where the item is considered not useful, and the distribution (CDF) is defined as:

$$\Pr(\mathrm{RUL}(t_j) \leq t) = \Pr(\min\left\{h : X(t_j + h) \in \mathcal{X}_l\right\} \leq t | T > t_j, Y(t)_{t \in \mathcal{T}_{t_j}})$$

where we condition on the fact that the item is still useful ($T > t_j$), and the knowledge of the measurements, i.e., the various observations (Y) at various points in time, i.e., the set $\mathcal{T}_{t_j}$.

The Wiener and Gamma processes

The Wiener and gamma processes are popular stochastic processes used to model degradation. Both processes assume that the change in degradation level in a small time interval can be described by a stochastic variable. In the Wiener process these changes can be both positive and negative, whereas in the gamma process the changes are always positive, i.e., positive increments. There are various pros and cons for these two processes. The gamma process is more intuitive, since increments (degradation) is always positive which is true for man failure mechanism, i.e., we can not improve unless some measures are taken. On the other side, measurements of the degradation often show that the change in degradation level from one point of time to the next may be negative. This could then be caused by measurement errors (noise).

Wiener Process with Linear Drift

Before we define a Wiener process with drift we define the Wiener process $\{W_t, t \geq 0\}$ by:

1. $W_0 = 0$
2. W has independent increments: for every $t > 0$, the future increments $W_{t+u} - W_t, u \geq 0$, are independent of the past values $W_s, s \leq t$.
3. W has Gaussian increments: $W_{t+u} - W_t$ is normally distributed with mean 0 and variance $u, W_{t+u} - W_t \sim \mathcal{N}(0, u)$.
4. W has continuous paths: W_t is continuous in t.

Note the slightly different notation from Chapter 10, we use the notation $\{W_t, t \geq 0\}$ rather than $\{W(t), t \geq 0\}$.

We now define the stochastic process:

$$X_t = \mu t + \sigma W_t$$

as a Wiener process with linear drift μ and infinitesimal variance σ_2.

It follows that $X_t = X(t)$ is normally distributed with mean μt and variance $\sigma^2 t$. Note that in the textbook $\mu = \alpha$. Further X has Gaussian increments: $X_{t+u} - X_t$ is normally distributed with mean 0 and variance $u, X_{t+u} - X_t \sim \mathcal{N}(\mu u, \sigma^2 u)$.

It is well known from the theory of stochastic processes that the time T when the process for the first time reach the level ℓ is inverse-Gauss distributed with parameters $\alpha = \ell/\mu$ and $\beta = (\ell/\sigma)^2$.

For the inverse-Gauss distribution, i.e., $X \sim \mathrm{IG}(\alpha, \beta)$ we have:

$$f_X(x;\alpha;\beta) = \sqrt{\frac{\beta}{2\pi x^3}}\exp\Big(-\frac{\beta(x-\alpha)^2}{2\alpha^2 x}\Big) \tag{3}$$

and

$$F_X(x;\alpha;\beta) = \Phi\left(\frac{\sqrt{\beta}}{\alpha}\sqrt{x} - \sqrt{\beta}\frac{1}{\sqrt{x}}\right) + \Phi\left(-\frac{\sqrt{\beta}}{\alpha}\sqrt{x} - \sqrt{\lambda}\frac{1}{\sqrt{x}}\right)e^{2\beta/\alpha} \tag{4}$$

The expected value and variance are given by:

$$\begin{aligned}\mathrm{E}(X) &= \alpha \\ \mathrm{Var}(X) &= \alpha^3/\beta\end{aligned}$$

In the Wiener process with parameters μ, σ the time, T to first passage of the threshold ℓ is then

$$T \sim \mathrm{IG}(\ell/\mu, (\ell/\sigma)^2)$$

and the expected value and variance are given by:

$$\begin{aligned}\mathrm{E}(T) &= \ell/\mu \\ \mathrm{Var}(T) &= \sigma^2\ell/\mu^3\end{aligned}$$

Maintenance decision problem

We consider the following situation:

- Assume that we can observe the degradation process continuously without any uncertainty
- A failure occurs if $X(t) \geq \ell$ for some time t
- When degradation approaches the failure limit, ℓ we will place a request to replace the component with a new component
- We assume that there is a deterministic lead time, say T_L
- The objective is to determine the maintenance limit, $m < l$, i.e., how close to the failure limit we dear to go

A reasonable cost equation to minimize is:

$$C(m) = \frac{c_\mathrm{R} + c_\mathrm{F}F(T_\mathrm{L}|m) + c_\mathrm{U}\int_0^{T_\mathrm{L}} f(t|m)(T_\mathrm{L} - t)dt}{\mathrm{MTBR}(m)}$$

where

- c_R = cost of renewal/replacement
- c_F = cost of failure (additional cost for corrective maintenance and extra cost for the failure event)
- c_U = cost per hour down time
- $F()$ and $f()$ are CDF and PDF for remaining useful lifetime (RUL), given we are at the maintenance limit m at some point
- MTBR(m) = Mean Time Between Renewals, given the decision rule to request a maintenance at m

We now consider one maintenance cycle:

- Assume that we at time t in this cycle observe $Y(t) = m$
- Let RUL_m be the time from t until a failure occurs
- RUL_m is inverse-Gauss distributed with parameters $\alpha_m = (\ell - m)/\mu$ and $\beta_m = (\ell - m)^2/\sigma^2$, where μ and σ^2 are the parameters in the Wiener process, and ℓ is the failure threshold
- Thus, $F() = F(t; \alpha_m; \beta_m) =$ and $f() = f(t; \alpha_m; \beta_m)$ are given by is given by equations (4) and (3) respectively, and the nominator of $C(m)$ may be obtained by numerical integration
- $\text{MTBR}(m) = m/\mu + T_L$

Gamma process

A stationary *gamma* process $Y(t), t \geq 0$ is defined by:

1. $Y(0) = 0$
2. $Y(t), t \geq 0$ has independent and stationary increments
3. The increments in an interval $(s, t]$ is $Y(t) - Y(s)$ and are assumed to be gamma distributed with parameters $(t - s)\alpha$ and β

Since the increments are gamma distributed, the degradation in a time interval of length $s - t$ is $(t - s)\alpha/\beta$. Referring to the example given for the Weibull process, we could be tempted to assume that the expected time for a new item to reach a maintenance limit, m is $\text{MTBR}_m = m\beta/\alpha$. It can be shown (not in the textbook) that the mean time can be approximated by $m\beta/\alpha + 1/(2\alpha)$. The extra term $1/(2\alpha)$ is often denoted "overshooting" effect. The idea is that the gamma process is a jump process. This means that it will never exactly hit the value y_p but rather hit slightly above, and hence it takes some "extra" time compared to if it was an "exact hit".

Compared to the Wiener process, it is however easier to find the RUL_m distribution. Assume we order a maintenance when the process reaches the value m. We then have

$$\begin{aligned}F_{\mathrm{RUL}_m}(t) = \Pr(\mathrm{RUL}_m \leq t) = \Pr(X(T_m+t) \geq \ell | X(T_m) < \ell, \text{ history up to } T_m)\\ \approx \Pr(X(T_m+t) - X(T_m) \geq \ell - m) =\\ \int_{\ell-m}^{\infty} f_{\alpha t,\beta}(u)du = 1 - F_{\alpha t,\beta}(\ell - m)\end{aligned}$$

where T_m is the point of time when the process exceeds the maintenance limit. $f_{\alpha t,\beta}()$ and $F_{\alpha t,\beta}()$ are the PDF and CDF of the gamma distribution with parameters αt and β respectively.

Note the approximation which is due to the fact that we never exactly reach the maintenance limit m due to overshooting. If we pursue the maintenance model used in the Wiener process example, we should also take the "overshooting" into account for the expected time to reach the maintenance limit which could be approximated by MTBR $= m\beta/\alpha + 1/(2\alpha) + T_{\mathrm{L}}$. In the cost model we also need the PDF for the RUL in addition to the CDF derived above.

Note the slightly different presentation in the textbook where RUL is defined at a given point of time t_j rather than the first passage time of the maintenance limit.