PK8207

Jørn Vatn
eMail: jorn.vatn@ntnu.no

Updated 2022-05-06

Lead-time model and marginal cost approach

In this memo we will investigate a predictive maintenance model where we assume that we can observe the health of a component in real-time, and we will order maintenance when the maintenance limit is reached. There is a lead-time from a maintenance order is placed until it is executed.

Theoretical background for the probability modelling is described elsewhere...

The Wiener process

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). In this memo only the Wiener process is considered.

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 $0u, W_{t+u} - W_t \sim \mathcal{N}(0,u)$.

4. W has continuous paths: W_t is continuous in t.

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 $\sigma^2 t$. Further X has Gaussian increments: $X_{t+u} - X_t$ is normally distributed with mean μ and variance $\sigma^2 u, X_{t+u} - X_t \sim \mathcal{N}(\mu u, \sigma^2 u)$.

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:

$$f_X(x;\alpha;\beta) = \sqrt{\frac{\beta}{2\pi x^3}} \exp\Big(-\frac{\beta(x-\alpha)^2}{2\alpha^2 x}\Big) \tag{1}$$

and

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

The expected value and variance are given by:

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

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:

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

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

The objective function, or cost equation to minimize is:

$$C(m) = \frac{c_\mathrm{R} + c_\mathrm{F} F(T_\mathrm{L}|m) + c_\mathrm{U} \int_0^{T_\mathrm{L}} f(t|m)(T_\mathrm{L} - t)dt}{\mathrm{MTBR}(m)} \tag{3}$$

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 equations (2) and (1) respectively, and the nominator of $C(m)$ may be obtained by numerical integration
- $\mathrm{MTBR}(m) = m/\mu + T_\mathrm{L}$

In Equation (3) the lead-time is assumed to be fixed. In some cases we can influence the lead-time, and it is reasonable that shorter lead-times means higher cost of renewal, i.e., $c_\mathrm{R} = c_\mathrm{R}(t_\mathrm{L})$, where t_L now is a decision variable. In this case we rewrite Equation (3):

$$C(m, t_\mathrm{L}) = \frac{c_\mathrm{R}(t_\mathrm{L}) + c_\mathrm{F} F(t_\mathrm{L}|m) + c_\mathrm{U} \int_0^{T_\mathrm{L}} f(t|m)(t_\mathrm{L} - t)dt}{\mathrm{MTBR}(m)} \tag{4}$$

Marginal approach

Minimizing Equation (3) wrt m represents a situation where we search the long run minimum under static conditions. In many situations there will be local variation in cost structures, and lack of opportunities for execute the maintenance. We will demonstrate the use of the marginal approach in this situation. For the marginal approach we start with the following assumptions:

- We minimize Equation (3) wrt m
- Denote the optimal maintenance limit by m^*, and the corresponding minimal cost is denoted C^*
- For some reason, we have lost the numerical value of m^*, but still we have kept the numerical value of C^*
- Assume we at some time, say t observe the degradation level $X(t) = x$

Essentially this means that we know the long run average cost per unit time, but we do not exactly now what is the best "here and now" decision to make. To solve the marginal approach problem We know that if $x < m^*$ it is beneficial to wait with placing the maintenance order. Thus if we calculate the expected cost of two strategies

1. Place an order at time t
2. Place an order at time $t + \Delta t$

then strategy 2 should be cheaper if $x < m^*$. The expected cost of strategy 1 from now on until $T_\mathrm{L} + \Delta t$ is:

$$C(S1|x) = c_\mathrm{F} F(T_\mathrm{L}|x) + c_\mathrm{U} \int_0^{T_\mathrm{L}} f(t|x)(T_\mathrm{L} - t)dt + C^* \Delta t \tag{5}$$

and the expected cost of strategy 2 from now on until $T_\mathrm{L} + \Delta t$ is:

$$C(S2|x) = c_\mathrm{F} F(T_\mathrm{L} + \Delta t|x) + c_\mathrm{U} \int_0^{T_\mathrm{L}+\Delta t} f(t|x)(T_\mathrm{L} + \Delta t - t)dt \tag{6}$$

Thus, given that the current state is x, the change in cost by postponing placing the order by Δt from now on is given by $C(S2|x) - C(S1|x)$. It is thus beneficial to postpone maintenance as long as $C(S2|x) - C(S1|x) < 0$.

Exercise 1

Assume the following values: $\mu = 1, \sigma = 2, l = 100, c_R = 1\,000, c_F = 10\,000, c_U = 5\,000$ and $T_L = 5$. Use Equation (3) to obtain the optimal maintenance limit, m^*

Exercise 2

Assume the same model parameters as in Exercise 1, and assume that the current condition is $x = m^* - 1$. Calculate the change in cost by slightly postponing placing the order. Repeat for $x = m^* + 1$. Does the example confirm the marginal cost idea?

Example 1 - Opportunistic maintenance, first approach

Assume that maintenance opportunities are limited. Current time is t_0 and we are approaching the maintenance limit. Maintenance (renewal) has to be planned in advance. We can either carry out the maintenance at the first opportunity at time t_1 or at the next opportunity at time t_2. Which one to choose? For the calculation we need to realize that the two strategies will have slightly different "ending points", i.e., the renewal in strategy 2 will take place slightly later than in strategy 1. If we assume that both strategies continue to "pay" C^* from the time of renewal until some time $T > \max(t_1, t_2)$, the cost of the two strategies are respectively:

$$C(S1|x) = c_{\mathrm{F}} F(t_1 - t_0|x) + c_{\mathrm{U}} \int_0^{t_1 - t_0} f(t|x)(t_1 - t_0 - t)dt + C^*(T - t_1) \tag{7}$$

and

$$C(S2|x) = c_{\mathrm{F}} F(t_2 - t_0|x) + c_{\mathrm{U}} \int_0^{t_2 - t_0} f(t|x)(t_2 - t_0 - t)dt + C^*(T - t_2) \tag{8}$$

Note that C^* is calculated under the assumption that we always can maintain at the optimum value. Since opportunities are limited, we should use a slightly higher value than obtained by Equation (3).

Example 2 - Opportunistic maintenance, uncertainty regarding the future

In Example 1 we assumed that the two maintenance opportunities were always available. Both opportunities might be uncertain. Now, assume that we still have to plan for one of the two opportunities in advance, the first opportunity will be available with certainty, but there is a probability, say p, that opportunity 2 will not be available, and we have to stick to opportunity 3 available at time $t_3 > t_2$. In this case the expected cost of the second strategy is:

$$\begin{aligned} C(S2|x) = (1-p)&\left[c_{\mathrm{F}} F(t_2 - t_0|x) + c_{\mathrm{U}} \int_0^{t_2 - t_0} f(t|x)(t_2 - t_0 - t)dt + C^*(T - t_2)\right] \\ p&\left[c_{\mathrm{F}} F(t_3 - t_0|x) + c_{\mathrm{U}} \int_0^{t_3 - t_0} f(t|x)(t_3 - t_0 - t)dt + C^*(T - t_3)\right] \end{aligned} \tag{9}$$

Example 2 could be generalized to take into account that also opportunity 3 will not be available. This will typically be the problem when weather conditions will limit the opportunities for maintenance.

Example 3 - Opportunistic maintenance, portfolio example

Now, assume that we have opportunities at time $t_1, t_2, t_3, \ldots$. Further we have n components that are approaching their due date for maintenance. Assume we do not need to plan in advance, so depending on the condition we can utilize an opportunity immediately if this is the optimal choice. If we do not have any restriction in terms of how many components that could be maintained at each opportunity, a "greedy algorithm" seems reasonable, i.e., we have a rolling horizon where we at each point of time t_i consider what is best, i.e., to use the opportunity at time t_i or utilize the subsequent one at time t_{i+1}.

In many cases there will be a limitation on how many components that could be maintained at each opportunity. Assume the limit is K. The greedy algorithm is now likely to fail. The greedy algorithm will be to check which components would like to be maintained at time t_i and which ones at t_{i+1}. If more than K components prefer the first opportunity, we will prioritize the ones that would be most costly to postpone. Those not prioritized have to be postponed, and we repeat the process at time t_{i+1}. We see that the greedy algorithm may result in activities being stacked in a queue, where waiting time may be very costly. To avoid queuing, we could be more proactive, for example if less than K components prefer maintenance at time t_i, we should advance some of the components that prefer t_{i+1} if the opportunity at time t_{i+1} will be overbooked. There are two challenges here:

1. We need some dynamic programming approach, since also at t_{i+2} there might be an overbooking

2. At time t_i the future development of the degradation level for the various components are random. Therefore it is hard to say how many components that actually would require maintenance at the subsequent opportunities

An improved "greedy algorithm" would now be the following:

1. Calculate the number of components that prefer maintenance at time t_i

2. If the number is greater than K, prioritize the one that are most costly to postpone, carry out the maintenance, and wait for the next opportunity at time t_{i+1}

3. If the number equals K, carry out the maintenance for these, and wait for the next opportunity at time t_{i+1}

4. If the number is less than K, identify those components that prefer maintenance at time t_{i+1} or t_{i+2} according to the expected development in the degradation level

5. If the subsequent slots then will be overbooked, advance one or more components from the opportunities at t_{i+1} and/or t_{i+2} till the opportunity at t_i.

Note that looking beyond the next opportunity at time t_{i+1} is a proactive approach since components preferring the opportunity at t_{i+2} might deteriorate faster than expected, and hence contribute to extra overbooking at time t_{i+1}. A better solution is expected to be obtained by introducing stochastic programming, but this is far more demanding, see Arif.