

Inspection and maintenance decisions based on imperfect inspections

M.J.Kallen

Delft University of Technology, Delft, Netherlands

J.M. van Noortwijk

HKV Consultants, Lelystad, and Delft University of Technology, Delft, Netherlands

ABSTRACT: The process industry is increasingly making use of Risk Based Inspection (RBI) techniques to develop cost and/or safety optimal inspection plans. This paper makes use of the gamma stochastic deterioration process to model the corrosion damage mechanism. This model is successfully extended to update prior knowledge over the corrosion rate with imperfect wall thickness measurements. This is very important in the process industry as current non-destructive inspection techniques are not capable of measuring the exact material thickness, nor can these inspections cover the total surface area of the component. The model is illustrated by examples using actual plant data.

1 INTRODUCTION

In order to illustrate the requirements, which are necessary for the development of a suitable model, we start with a general introduction on current practices in the process industry concerning the use of decision models for inspection planning.

1.1 Risk Based Inspection

Since the late 1980's, numerous companies and organizations have developed several qualitative and quantitative models to aid plant engineers with the prioritization of component inspections. The average chemical process plant or refining installation will have thousands of components like pipelines, columns, heat exchangers, steam vessels etc. which operate under significant pressure. This pressure, combined with the corrosive nature of the chemicals inside the systems and the exposure to the weather on the outside, will degrade the quality of the construction material of the components. Among the most common degradation mechanisms are internal and external thinning (e.g. corrosion), cracking, brittle fracture and fatigue.

The highest uncertainties in the decision model are associated with the rate at which these mechanisms reduce the resistance of the construction material. Inspections are used to reduce this uncertainty, but since current wall thickness or crack length measurement techniques are not perfect, the measurements contain (small) errors. Some inspection techniques, e.g. ultra-

sonic wall thickness measurements, are highly accurate, but there is always the spatial variability in the measurements. This means that the quality of the material is not uniform for the whole component. Inspections have to be carried out such that the most critical spots are covered and such that the average measurement is representative for the complete item. Decision models should take these measurement errors into account.

1.2 Current methodologies

A common approach in tackling the inspection problem, is to start with a prioritization. This means that usually a more qualitative model is used to determine the components which constitute the highest risk. The assumption is that 80% of the risk is generated by only 20% of the components. This prioritization is then followed by a detailed quantitative analysis of these high-risk items, in which the remaining lifetimes are estimated and the consequences of failures are modelled. The American Petroleum Institute (API) has published a large document (American Petroleum Institute 2000), which describes the methodology developed in cooperation with an industry sponsor group. This document has become a methodology in itself and is used by many companies as a basis for further development.

Typically, structural reliability methods are used to estimate the failure probabilities of the components. The condition of the construction material is

described by a state function g , which is represented by a resistance (R) minus stress (L) model:

$$g = R - L. \quad (1)$$

The uncertain variables in this model have normal densities with suitable parameters assigned to them. The failure probability is then approximated by using a reliability index method (e.g. FORM: First Order Reliability Method), which makes use of the limit state $g = 0$ or $R = L$. Many of the methods currently applied in the process industry, also make use of a simple discrete version of Bayes' theorem to update prior knowledge, or an estimate for the rate of degradation, with inspection results. The likelihood in Bayes' formula represents the likelihood of an inspection correctly identifying the state of the component.

In the following sections we will take a different approach to this problem. We will use the gamma stochastic process to model the degradation and we will update the parameters of this process using inspection data. This is inspired by the successful application of this stochastic process to the problem of optimally inspection large civil structures like dikes and storm-surge barriers, e.g. see (van Noortwijk 1996). It is not claimed that this is the best approach to this problem. The purpose of this paper is to illustrate the versatility of the gamma process and to show that there is great potential for its use in the process industry. Characteristically, in the development of most current models, many assumptions and simplifications are applied in order to keep the models easy to use. We will use the same approach in order to avoid large amounts of input and to keep the necessary input as simple as possible.

2 MODELLING THE DETERIORATION

In this paper we will only consider degradation due to corrosion, but the application of this model is not restricted to this damage mechanism. For this purpose we will first define the corrosion state function, after which the gamma process with suitable parameters is introduced.

2.1 Corrosion state function

The state function for corrosion is taken from (American Petroleum Institute 2000) and is defined as:

$$g(t) = \underbrace{S \left(1 - C \frac{t}{x_0} \right)}_{\text{Resistance}} - \underbrace{P \frac{d}{2x_0}}_{\text{Load}}, \quad (2)$$

where S [MPa = 10 bar], C [mm/yr] and P [bar] are random variables representing the material strength, corrosion rate and operating pressure respectively. The component diameter d [mm] and the initial material thickness x_0 [mm] are assumed to be known.

Next to the variables are the dimensions which are used throughout this paper. Note that the amount of wall loss at time t is given by $w(t) = C \times t$ [mm].

The component is assumed to have failed at time t when $g(t) < 0$. If t' is the time at which failure occurs, i.e. $g(t') = 0$, then $w(t') = C \times t'$ is the maximum amount of wall thickness which can be lost until the component fails. If we call this the safety margin m , then we can calculate this amount from the state function (2):

$$m = w(t') = x_0 - P \frac{d}{2S}. \quad (3)$$

Each component will usually have a so-called corrosion allowance c_{max} associated with it. This is a value given by the manufacturer of the component and represents the maximum amount of wall loss up to which the component is assumed to be able to function safely. It should hold that $0 \leq c_{max} \leq m$, which means that we can write the corrosion allowance as a percentage of the safety margin:

$$c_{max} = \rho m, \quad 0 \leq \rho \leq 1. \quad (4)$$

We will use this value as the replacement level, i.e. we will preventively replace the component if the wall loss is more than the corrosion allowance.

2.2 Gamma deterioration process

Instead of using a reliability index method, we will model the cumulative wall loss with a gamma process. We will use the following definition for the gamma density with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$:

$$\text{Ga}(x|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp\{-\beta x\} \text{ for } x \geq 0 \quad (5)$$

The gamma process with shape function $at > 0$, $t \geq 0$ and scale parameter $b > 0$ is a continuous-time process $\{X(t) : t \geq 0\}$ with the following properties:

1. $X(0) = 0$ with probability one,
2. $X(\tau) - X(t) \sim \text{Ga}(a(\tau - t), b)$ for $\tau > t \geq 0$,
3. $X(t)$ has independent increments.

Let $X(t)$ denote the amount of deterioration at time t , then the probability density function of $X(t)$ is given by

$$f_{X(T)}(x) = \text{Ga}(x|at, b) \quad (6)$$

In essence, we replace the amount of wall loss $w(t)$ with the process $X(t)$. There are a number of advantages to using the gamma process. Most interestingly, the increments are always positive, therefore

increases in wall thickness are not allowed for. The fact that increments are independent and therefore exchangeable fits well with the physics of corrosion and deterioration in general.

Using moment generating functions, it can be proven that the expectation and the variance of the process $X(t)$ are given by:

$$\mathbb{E}(X(t)) = \frac{a}{b}t \quad \text{and} \quad \text{Var}(X(t)) = \frac{a}{b^2}t. \quad (7)$$

Assuming that the expectation and variance are linear in time, i.e.

$$\mathbb{E}(X(t)) = \mu t \quad \text{and} \quad \text{Var}(X(t)) = \sigma^2 t,$$

we find that the parameters of the process $X(t)$ are defined as:

$$a = \mu^2/\sigma^2 \quad \text{and} \quad b = \mu/\sigma^2, \quad (8)$$

where μ is the expectation for the average corrosion rate and σ^2 is the variance of the process. These two parameters are uncertain and assessing both variables for each individual component is too much work. In order to keep the method practical and easy to use for the plant engineer, we will fix the standard deviation σ relative to the mean μ through the use of a coefficient of variation ν (this coefficient is often referred to as the COV):

$$\sigma = \nu \times \mu \implies \nu = \sigma/\mu. \quad (9)$$

This approach is used in many of today's models: the variances of the uncertain variables are predetermined by expert judgment and subsequently fixed in the model. Using (9), the density for $X(t)$ in (6) reduces to

$$f_{X(t)}(x) = \text{Ga}\left(x \left| \frac{t}{\nu^2}, \frac{1}{\mu\nu^2} \right.\right). \quad (10)$$

Now we are only left with the uncertain variable μ for the average corrosion rate. In the following section we will define a suitable prior density for this variable and we will discuss the consequence of fixing σ relative to μ .

3 INSPECTION UPDATING

Due to the fact that most components are usually at most inspected every 2 years, with the average inspection interval being about 6 to 8 years, there is not enough data available for a statistical method like regression analysis. Using Bayesian updating we can efficiently incorporate the measurements and the engineer's prior knowledge or estimate of μ .

3.1 Choosing the prior

Due to the large number of components and the usually limited experience of the plant engineer with probability theory, it is not feasible to ask this engineer to assess a suitable density for μ . In line with current practices in the process industry, we will only ask for an average corrosion rate over which a default density will be placed. The API (American Petroleum Institute 2000) uses the simple discrete prior density, which is shown in Figure 1. The idea behind this den-




Figure 1: The API discrete prior and the related (continuous) inverted gamma density.

sity is that the model has 50% confidence in a corrosion rate which is less than or equal to the rate assessed by the plant engineer. The other 50% is divided into 30% between 1 and 2 times the assessed rate and 20% between 2 and 4 times the engineer's estimate.

For our gamma process model we will not use this discrete prior, but a continuous inverted gamma density. The definition of this density is given by

$$\text{Ig}(x|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \left(\frac{1}{x}\right)^{\alpha+1} \exp\left\{-\frac{\beta}{x}\right\} \quad (11)$$

for $x \geq 0$. Notice that a random variable X is inverted gamma distributed if $Y = X^{-1} \sim \text{Ga}(\alpha, \beta)$ with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$. With suitable choices for the parameters, we can again define a default prior density for μ as is shown in Figure 1. Most confidence is placed between the engineer's assessment and 2 to 3 times this value. The prior density in Figure 1 is arbitrarily chosen for the purpose of demonstration. In practice, the parameters of this density should be determined using expert judgment. One could also define a number of different default priors from which the practitioner can select the one which best represents his own confidence/uncertainty in the degradation rate.

In the following sections we will use Bayes' theorem to update the prior density with the likelihood of the inspection measurement to obtain a posterior density. The continuous version of Bayes' theorem is given by

$$\pi(\mu|x) = \frac{l(x|\mu)\pi(\mu)}{\int_{\mu=0}^{\infty} l(x|\mu)\pi(\mu)d\mu}, \quad (12)$$

where $\pi(\mu)$ is the prior density for μ and $l(x|\mu)$ is the likelihood of measurement x given μ .

3.2 Case 1: perfect inspections

The choice for the inverted gamma prior becomes clear when we only consider perfect inspections. In other words, we first assume that we can measure the exact wall thickness. We can then precisely determine the amount of wall loss x at time t and it can be proven that the posterior $\pi(\mu|x)$ is given by

$$\pi(\mu|x) = \text{Ig}\left(\mu \left| \alpha + \frac{t}{\nu^2}, \beta + \frac{x}{\nu^2} \right.\right), \quad (13)$$

where α and β are the parameters from the prior density. The fact that the product of the inverted gamma density and gamma distributed likelihood is again proportional to a inverted gamma density is also used in (van Noortwijk et al. 1995).

The equivalent of (13) for $n \geq 1$ inspections is

$$\begin{aligned} \pi(\mu|x_1, \dots, x_n) &= \\ &= \text{Ig}\left(\mu \left| \alpha + \sum_{i=1}^n \frac{t_i - t_{i-1}}{\nu^2}, \beta + \sum_{i=1}^n \frac{x_i - x_{i-1}}{\nu^2} \right.\right). \end{aligned}$$

which reduces to

$$\pi(\mu|x_n) = \text{Ig}\left(\mu \left| \alpha + \frac{t_n}{\nu^2}, \beta + \frac{x_n}{\nu^2} \right.\right), \quad (14)$$

if we assume that $x_0 = 0$ at $t_0 = 0$. In other words, because we fixed the standard deviation σ relative to the mean μ , only the last inspection is needed to calculate the posterior when using perfect measurements. Besides the fact that perfect measurements do not exist, this result will be hard to sell to any plant engineer or regulator.

3.3 Case 2: imperfect inspections

This is where we present the main feature of this paper, namely the extension of the gamma process updating model with uncertain measurement data. Similar to (Newby and Dagg 2002), we consider a new process $Y(t)$, which includes the original process $X(t)$ together with a normally distributed error ϵ :

$$Y(t) = X(t) + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma_\epsilon). \quad (15)$$

Here we assume the error has a mean 0 and a standard deviation σ_ϵ . Taking a mean different from zero would mean that the inspection tends to over- or underestimate the actual wall loss. The likelihood of the measurement y given the corrosion rate μ is now determined by the convolution:

$$l(y|\mu) = f_{Y(t)}(y) = \int_{-\infty}^{\infty} f_{X(t)}(y - \epsilon) f_\epsilon(\epsilon) d\epsilon, \quad (16)$$

where $f_{X(t)}(y - \epsilon) = \text{Ga}(y - \epsilon|at, b)$ is the likelihood of the gamma increment $X(t)$ as given by (10). We immediately go over to the likelihood for more than one inspection:

$$l(\mathbf{y}|\mu) = \prod_k l_{Y(t_k)-Y(t_{k-1})}(y_k - y_{k-1}|\mu), k > 1, \quad (17)$$

where $\mathbf{y} = \{y_1, \dots, y_k\}$ are the wall loss measurements. In (17), we have used the fact that the increments are independent. Now we introduce some notation: the increment of $X(t)$ between two inspections is defined as $D_k = X(t_k) - X(t_{k-1})$, the difference between two measurements is $d_k = y_k - y_{k-1}$. Using this notation and the integral convolution as in (16), the likelihood (17) can be rewritten as

$$\begin{aligned} l(\mathbf{y}|\mu) &= \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \dots \\ &\dots \prod_k f_{D_k}(d_k - \delta_k) f(\delta_1, \dots, \delta_k) d\delta_1 \dots d\delta_k, \end{aligned} \quad (18)$$

where $\delta_k = \epsilon_k - \epsilon_{k-1}$. Clearly the δ 's are not independent since every δ_k depends on δ_{k-1} . We are left with two options: we calculate the covariances between the δ 's and analytically solve the likelihood using the joint distribution of the δ 's or we simulate the ϵ_k 's and approximate the likelihood. Since the first option will complicate matters considerably, we will use the simulation approach. The likelihood (18) can be formulated as an expectation, which in turn can be approximated by the average of the product:

$$\begin{aligned} l(\mathbf{y}|\mu) &= \mathbb{E} \left\{ \prod_k f_{D_k}(d_k - \delta_k) \right\} \approx \\ &\approx \frac{1}{N} \sum_{j=1}^N \prod_k f_{D_k}(d_k - \delta_k^{(j)}) \text{ as } N \rightarrow \infty. \end{aligned} \quad (19)$$

Here we use the law of large numbers to perform a so-called Monte Carlo integration. For each inspection k we

1. sample $\epsilon_k^{(j)}$ for $j = 1, 2, \dots, N$ and
2. calculate $\delta_k^{(j)} = \epsilon_k^{(j)} - \epsilon_{k-1}^{(j)}$, then finally
3. calculate (19).

Since the gamma distributed $f_{D_k}(x)$ is not defined for $x < 0$, we need to make sure that the argument is non-negative. We have worked our way around this problem by using the minimum function as follows:

$$l(\mathbf{y}|\mu) \approx \frac{1}{N} \sum_{j=1}^N \prod_k f_{D_k} \left(d_k - \min \left\{ \delta_k^{(j)}, d_k \right\} \right) \quad (20)$$

as $N \rightarrow \infty$, where

$$\begin{aligned} f_{D_k} \left(d_k - \min \left\{ \delta_k^{(j)}, d_k \right\} \right) &= \\ &= \text{Ga} \left(d_k - \min \left\{ \delta_k^{(j)}, d_k \right\} \left| \frac{t_k - t_{k-1}}{\nu^2}, \frac{1}{\mu\nu^2} \right. \right). \end{aligned}$$

Equation (20) can now be substituted in Bayes' formula (12), which can then be solved by discretization (i.e. simple numerical integration) to obtain the posterior $\pi(\mu|y_1, \dots, y_n)$.

The choice for use of simulation to determine (18) greatly reduces the efficiency of the model, but we also see that the choice for the prior distribution is then no longer restricted to the inverted gamma density. Also, we can easily introduce correlated measurement errors (ϵ_k) using a multivariate normal distribution in order to see the effect of inspection dependence on the end result. Due to the absence of the required data, we have not included this here.

4 DECISION CRITERION

We now have a posterior density over the average corrosion rate, which incorporates the plant engineer's prior knowledge and all the available inspection data. Using this density and the densities for the uncertain variables S (material strength) and P (operating pressure), we can calculate the probability of failure or a preventive replacement of the component as a function of time.

Inspections in a process plant are very expensive due to the fact that the process often has to be stopped during the inspection. Also, because many components contain highly corrosive and/or toxic chemicals, they have to be flushed and cleaned before an internal inspection can be done. The waste resulting from this rinsing needs to be treated, which brings added costs to the whole procedure. We therefore would like to make a cost optimal decision on when to inspect the components. This means that we want to maximize

the time interval between two subsequent inspections. On the other hand, we have to ensure that the component operates safely and therefore we also need to make sure that we do not inspect at too large intervals. For this purpose we use a cost based criterion, suggested by (Wagner 1975), called the expected average cost per time unit. This cost criterion is derived using renewal theory and uses the concept of the component life cycle. The length of this cycle is the time from the service start until a renewal, which is either a preventive replacement or a corrective replacement due to failure. The expected average costs per time unit is given by the ratio of the expected costs per cycle over the expected cycle length, see e.g. (van Noortwijk et al. 1997):

$$C(\theta, \Delta k) = \frac{\sum_{i=1}^{\infty} c_i(\theta, \Delta k) p_i(\theta, \Delta k)}{\sum_{i=1}^{\infty} i p_i(\theta, \Delta k)}, \quad (21)$$

where c_i and p_i are respectively the costs incurred during time unit i and the probability of renewal during this time unit. The ratio is a function of the replacement and failure levels, which is represented here by $\theta = \{m, \rho\}$, and we calculate this ratio for each inspection interval Δk . The failure probability also depends on the three uncertain variables S , P and C . For the sake of legibility, they are not included in the above equation. We refer to (Kallen 2002) for the details of the implementation of this particular model, which can be obtained by request from the authors.

In order to include the uncertainty over the material strength, operating pressure and the corrosion rate, we apply the same technique as in equation (19). We need to determine the expectation of (21), therefore we sample a large number of sets for the three uncertain variables: $\{s^{(j)}, p^{(j)}, c^{(j)}\}$, $j = 1, 2, \dots, N$. For a sufficiently large N , the average of the calculated (21) will approximate the expectation. The corrosion rate can be sampled from the posterior $\pi(\mu|y_1, \dots, y_n)$, which we have determined in section 3.3.

5 CASE STUDY: INSPECTING A HYDROGEN DRYER

We will illustrate the gamma process decision model with a case study on a hydrogen dryer. Table 1 summarizes the operational and measurement data, which is taken from an undisclosed plant in the Netherlands. The cost numbers are fictive. The material strength is determined using the tensile (TS) and yield (YS) strengths with the equation:

$$S = \min \{1.1(TS + YS)/2, TS\}.$$

The values for these strengths can be looked up in (ASME 2001). The uncertainty in S comes from the fact that these values represent the minimum requirements for this particular material type and therefore

Component type:	vertical drum	
Material type:	carbon steel	
Service start:	1977	
Tensile strength (TS):	413.69	MPa
Yield strength (YS):	206.84	MPa
Operating pressure:	32	bar(g)
Drum diameter:	1180	mm
Init. material thickness:	16.8	mm
Corrosion rate (est.):	0.1	mm/yr
Corrosion allowance:	4.5	mm
<i>Ultrasonic wall thickness measurements:</i>		
1982:	15.0	mm
1986:	15.6	mm
1990:	14.6	mm
1994:	14.2	mm
1998:	13.8	mm
<i>Costs for different actions:</i>		
Inspection:	10,000	\$
Preventive replacement:	50,000	\$
Failure + replacement:	1,000,000	\$

Table 1: Operational, material and inspection data for a hydrogen dryer.

the material could be stronger. On the other hand, the material loses strength when it is processed by the manufacturer of the component, therefore it could be weaker than indicated.

Note that we will disregard the measurement taken in 1982, because it is too inaccurate to be considered for this analysis. This is not uncommon practice in the process industry, as older measurements (older than 10 to 15 years) are not considered to be reliable. This is due to the fact that the quality of the inspection techniques used in those days is not comparable to the accuracy of current techniques. Our gamma process model can incorporate this measurement only if the uncertainty distribution of the error term is chosen such that the probability of material growth is very small.

In the Netherlands, the inspection of pressurized vessels is regulated by the Dutch rules for pressure vessels (Stoomwezen 1997). All types of components are categorized and for each category there exists a fixed mandatory inspection interval. For the hydrogen dryer in our example the fixed interval is 4 years, which can be clearly observed in the measurement dates. If a plant engineer wants to extend this interval based on the results of a RBI analysis, then he will have to submit an application to the proper authorities for acceptance. The rules require all components to have had at least one regular inspection, according to the fixed intervals, before this application can be made. For items with a fixed interval of 2 years, this minimum is two inspections. There will thus always be at least one set of measurement data available for the model to be applied.

Using the Bayesian updating mode, which we have discussed in section 3, we can calculate the posterior density for the corrosion rate given the measurement data in Table 1. The results are shown in Figure 2. For




Figure 2: Posterior densities for 1 perfect and 4 imperfect inspections compared to the prior density.

comparison purposes, we have also included the posterior density, which we would get if we assumed that the measurements were exact. In this case only the last measurement is of importance, as we have determined in equation (14). The difference between the perfect and imperfect inspections is quite large. The assumption of perfect inspections clearly increases the confidence in the estimated corrosion rate considerably compared to the assumption of imperfect inspections.

The next step is to use the posterior density for the corrosion rate, the densities for S and P , and the cost criterion from section 4 to determine the optimal inspection period until the next inspection.

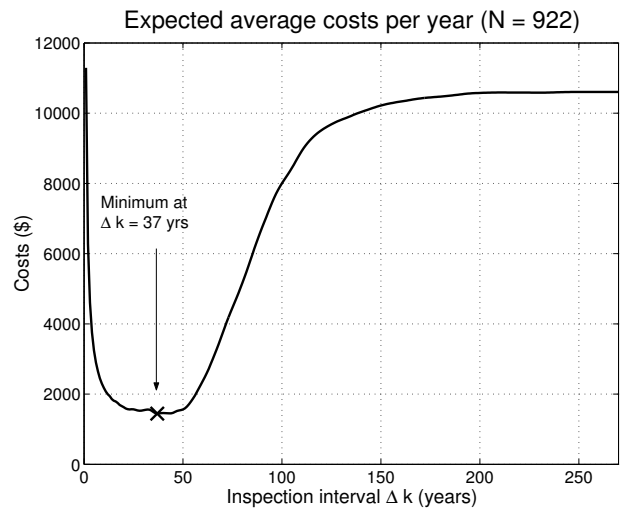


Figure 3: Expected average costs per year for the hydrogen dryer.

In Figure 3, we can see that the expected average costs per year will be high when we inspect too often and when we inspect too little. The result for the optimal inspection interval length is $\Delta k = 37$ yrs, but since the column was taken into service in 1977, we have to deduct its age of 25 years from this result. Therefore, the optimal period until the next inspection is 12 years. This is a very acceptable result, because it is less than the absolute maximum of 50 years and no more than 4 times the regular prescribed inspection interval of 4 years for this type of component. If this next inspection increases the confidence in the corrosion rate, then this model can be used again to determine the optimal inspection interval starting from the service start. A risk-averse decision maker can decide for a shorter inspection interval based on the result in Figure 3. A choice for Δk between 29 and 37 years will not significantly increase the expected costs per year. With current models, the decision maker will determine the failure probability as a function of time using the state function (2). Subsequently he will either choose to inspect before the fifth quantile of this distribution or he will define a risk criterion to determine the time before which an inspection should take place.

Future development of this model will include the calculation of the standard deviation of the expected costs per year. Interested readers can take a look at (van Noordwijk 2003), which discusses in great detail the required formulas for this purpose.

To finish, we note that the model will also be applicable when no inspection data is available. In other words: even when the component is new, we will be able to determine a responsible inspection interval with this decision model. The high uncertainty in the prior density for the degradation rate will initially make the optimal inspection interval quite short. As the number of measurements increase, so will the time between inspections increase.

6 CONCLUSIONS

There are a number of conclusions to be drawn from this research. First, we have determined that the gamma process is a suitable stochastic process to model the uncertain reduction of wall thickness due to corrosion. In line with what is currently done in the process industry, we have considerably simplified the parameters of this process by fixing the variance to the mean of the process. Together with the assumption of a fixed prior density for the average corrosion rate, this results in a model with minimum input requirements. This is a very desirable feature, as we are typically dealing with hundreds, maybe even thousands, of components in the average plant.

Next, we have created a simple extension to the Bayesian updating model, such that the model can in-

corporate the results from inaccurate measurements. In this step we have lost a lot of efficiency, because we have taken the path of simulation to solve the resulting equations. If we also consider other variables to be uncertain, e.g. the material strength or operating pressure, then the computational effort also becomes larger. The efficiency of the model is where the greatest improvements can be made in the future.

In order to make both cost optimal and safe inspection decisions, we have used the cost criterion of the expected average costs per year. Not only does this criterion fit well with our requirements, it also results in a graph which is easy to interpret by the plant engineers. This will ensure that the model will have some transparency and it will be less of a black box to the practitioners. Experience shows that the presentation of a single optimal value often leaves practitioners and regulators with more questions than they started out with.

The case study on a hydrogen dryer showed encouraging results of the whole model. We conclude that the use of the gamma Bayesian stochastic process is a viable alternative to the structural reliability methods which are currently used in the process industry.

7 ACKNOWLEDGMENTS

The research in this paper was performed by the author during an internship at Det Norske Veritas B.V. in Rotterdam, the Netherlands. We would like to thank Chris van den Berg and his colleagues at DNV for their support and for the plant data which is used in the example.

REFERENCES

- American Petroleum Institute (2000). *Publication API581* (2nd ed.). American Petroleum Institute.
- ASME (2001). Boiler and pressure vessel code. Technical report, American Society of Mechanical Engineers, New York, USA.
- Kallen, M. J. (2002). *Risk Based Inspection in the process and refining industry*. Master's thesis, Technical University of Delft, Faculty of Information Technology and Systems (ITS), Delft, the Netherlands.
- Newby, M. and R. Dagg (2002). Optimal inspection policies in the presence of covariates. In *Proceedings of ESREL 2002*, Lyon, France.
- Stoomwezen (1997). *Rules for pressure vessels*, Volume 1–3. The Hague, the Netherlands: Sdu Publishers.
- van Noordwijk, J. M. (1996). *Optimal maintenance decisions for hydraulic structures under isotropic deterioration*. Ph.d. thesis, Technical University of Delft.
- van Noordwijk, J. M. (2003). Explicit formulas for the variance of discounted life-cycle cost. *Reliability Engineering and System Safety*. Article in press.

- van Noortwijk, J. M., R. M. Cooke, and M. Kok (1995).
A bayesian failure model based on isotropic deterioration. *European Journal of Operational Research* 82(2), 270–282.
- van Noortwijk, J. M., M. Kok, and R. M. Cooke (1997).
Optimal maintenance decisions for the sea-bed protection of the eastern-scheldt barrier. In R. M. Cooke, M. Mendel, and H. Vrijling (Eds.), *Engineering Probabilistic Design and Maintenance for Flood Protection*, pp. 25–56. Kluwer Academic Publishers.
- Wagner, H. (1975). *Principles of operations research* (2nd ed.). Prentice–Hall.