

Risk Based Inspection
in the
Process and Refining Industry

by

Maarten-Jan Kallen

Faculty of Information Technology and Systems
Technical University of Delft
Delft, The Netherlands

December 2002

Preface

This text forms the thesis for the Master of Science degree in technical mathematics at the Technical University of Delft. The research was performed under the supervision of ing. van den Berg at Det Norske Veritas B.V. in Rotterdam and prof. dr. ir. van Noortwijk at TUDelft. The graduation committee is composed of:

prof. dr. ir. Jan van Noortwijk	(TUDelft, HKV consultants)
ing. Chris van den Berg	(Det Norske Veritas B.V.)
prof. dr. Tom Mazzuchi	(TUDelft, George Washington U.)
dr. Hans van der Weide	(TUDelft)

Acknowledgements

I would like to thank DNV for introducing the subject to me and giving me the opportunity to work within their organization. In particular I would like to thank my supervisor Chris van den Berg, Gert Koppen and Panos Topalis for their generous support. Many thanks go out to Jan van Noortwijk for his excellent supervision of my work for this thesis and to Roger Cooke and colleagues for an inspiring two years in the *Risk and Environmental Modelling* program.

Det Norske Veritas

Det Norske Veritas (DNV) is an independent, autonomous foundation working to safeguard life, property and the environment. DNV comprises 300 offices in 100 countries, with 5,500 employees. The headquarters are located in Høvik (Oslo) Norway.

As one of the world's leading maritime classification societies, DNV establishes rules for the construction of ships and mobile offshore platforms; some 20 percent of all ships currently on order will be build to DNV class.

DNV provides certification of management systems, products and personnel to a wide range of industries. DNV is also a leading provider of consulting services within safety-, health-, environment- and business risk.

Contents

List of figures	8
Introduction	9
1 Risk Based Inspection	11
1.1 Historical development	11
1.2 Dutch and European regulations	13
2 ORBIT™ Onshore	17
2.1 ORBIT™ overview	17
2.1.1 Thinning module	19
2.1.2 Stress corrosion cracking module	20
2.1.3 External corrosion	22
2.1.4 Brittle fracture	23
2.1.5 Fatigue	23
2.2 Inspection updating	23
2.3 Inspection planning	25
2.4 Sensitivity and uncertainty analysis	26
2.4.1 Example data	27
2.4.2 Sensitivity measures	27
2.4.3 Corrosion model	28
2.4.4 Cracking model	31
3 Reliability index methods	35
3.1 Reliability index concept	35
3.2 FORM	36
3.2.1 Rackwitz–Fiessler	38
3.2.2 Lagrange multipliers	39
3.2.3 FORM performance comparison	40
3.3 MVFOSM	42
3.4 Directional sampling	46

4	A new Gamma process inspection model	55
4.1	Model concept	56
4.2	Gamma stochastic process	58
4.3	Inspection updating	60
4.4	Optimal inspection and maintenance planning	64
4.5	Case study: inspecting a hydrogen dryer	66
4.5.1	Without uncertainty	67
4.5.2	With uncertainty	69
4.5.3	Influence of the replacement level	72
4.6	Model discussion	73
	Final conclusions	75
A	Mathematical details	77
A.1	FORM	77
A.2	Directional sampling	79
A.3	Gamma process inspection model	81
A.3.1	Gamma process	81
A.3.2	Perfect inspection(s)	82
A.3.3	Imperfect inspection(s)	85
A.4	Expected inspection and maintenance costs	88
B	Gamma process model implementation	93
B.1	modelinput: input data	93
B.2	corrgamma: main program	94
B.3	invgamfit: fitted inverted gamma prior	98
B.4	simulprep: simulation data preparation	98
B.5	eacptu: exp. ave. costs per time unit	99
B.6	calcfinish: simulation timer	101
	References	103
	Symbols, notations and definitions	105
	Index	108

List of Figures

1.1	The Terneuzen plant operated by Dow Benelux B.V. (Courtesy of Siebe Swart Fotografie Amsterdam).	12
2.1	ORBIT™ risk matrix example.	26
2.2	Cobweb plot for the thinning model at $t = 9$ years.	30
2.3	Cobweb plot for the thinning model at $t = 24$ years.	31
2.4	Cobweb plot for the thinning model at $t = 24$ years conditionalized on cr and p	32
2.5	Cobweb plot for the scc model at $t = 31$ years conditionalized on a and p	33
3.1	Graphical representation of the First Order Reliability Method.	36
3.2	Comparison of different FORM implementations and Monte Carlo simulation for the thinning model.	42
3.3	Results using MVFOSM for the thinning (a) and SCC (b) modules.	44
3.4	Monte Carlo simulation histogram (10000 samples) overlayed with the normal density determined using the MVFOSM technique for thinning and SCC at $t = 25yr$: (a) and (b) and the same for $t = 250yr$: (c) and (D) resp.	45
3.5	Graphical representation of the directional sampling technique.	47
3.6	3D rendering of the directional sampling technique.	49
3.7	Top-down view onto the plane defined by the material strength (S) and the corrosion rate (C) including the rendering of the regular Monte Carlo simulation.	50
3.8	Decision tree for choosing the radius in the directional sampling technique and calculating the correct failure probability.	51
3.9	Convergence of the directional sampling technique compared to the regular Monte Carlo technique for the thinning model at 22 years (a) and 26 years (b).	52
3.10	Comparison of the results for the cracking model failure probability using both reliability index and simulation techniques.	53

4.1	A graphical representation of the Gamma process deterioration process in a resistance minus stress model.	58
4.2	The inverted gamma density fitted to the ORBIT™ discrete prior.	61
4.3	Comparison between the posterior of a confirming perfect inspection and a perfect inspection resulting in a corrosion 3 times faster than the estimate.	62
4.4	Posterior densities for 1 imperfect inspection at 5 years and for 2 imperfect inspections at 5 and 15 years, each confirming a corrosion rate of 1mm/yr and $\sigma_\epsilon = 0.5\mu_C$	64
4.5	The expected cycle length (a), the expected costs per cycle (b) and the expected average costs per year (c) for the case with no uncertainty.	68
4.6	The prior and posterior density for μ given the inspection data of the hydrogen dryer.	69
4.7	Expected costs per cycle (a) and the expected cycle length (b) calculated using 199 simulations.	70
4.8	The expected average costs for the hydrogen dryer using two approaches with $N = 199$ simulations.	71
4.9	Expected cycle costs for simulation 36, where $\rho = 0.6078$	72
A.1	A comparison between a continuous density function and its discrete equivalent for the purpose of numerical integration.	86

Introduction

The original goal of my research at DNV was to perform a technical analysis of the ORBIT™¹ methodology and to have the methodology accepted by a Dutch committee as a valid model to estimate the remaining life of a component or pipeline. The analysis was necessary to get a clear understanding of the method itself, as well as its behavior. The latter has been covered by an uncertainty and sensitivity analysis, which included the use of the UNICORN package which is developed by the decision theory group at the faculty of Information Technology and Systems at the Technical University in Delft. The second chapter contains a summary of all modules in ORBIT™ followed by the results of the uncertainty and sensitivity analysis.

The analysis of ORBIT™ has resulted in a number of suggestions which have been presented to DNV's software development group in London. The specific details of these suggestions are not included in this thesis as they are for internal use only.

A historical overview of Risk Based Inspection (RBI) techniques and their development is given in the first chapter. Also in this chapter is a summary of the requirements which have been laid out by a technical committee in the Netherlands for the approval of risk-based methods for inspection planning. The third chapter deals with the most commonly used methods to approximate the probability of failure in the process and refining industries, namely MVFOSM and FORM. ORBIT™ uses the MVFOSM method, which works well for the degradation mechanisms under consideration. FORM is accepted as a more accurate method, but it is far less efficient. This chapter considers a number of different approaches to implementing the FORM approximation, including the most common Rackwitz–Fiessler method, and compares their performance on the ORBIT™ modules.

The fourth and last chapter came about in the second half of the project, when it came to our attention that the use of a stochastic process to model the material degradation seemed very fitting. The most challenging aspect of this model was to incorporate the possibility to deal with results from imperfect inspection in the continuous Bayesian updating model. Together with this approach, the chapter also includes a case study where the cost-

¹ORBIT™ is a trademark owned by Det Norske Veritas.

criterion of expected average costs per unit time is applied to a hydrogen dryer. Finally, in the conclusions, we take a look back at the stochastic process model and discuss its applicability in the process industry. The method is not new, but it has never been applied to this industry in particular.

In order to keep the thesis legible for engineers with a non-mathematical background, all mathematical details have been placed in appendix A. The implementation of the case study was done in MATLAB (version 6.1) from The Mathworks, Inc. and the code listing supplemented with various comments can be found in appendix B.

Chapter 1

Risk Based Inspection

This chapter gives a short overview of the historical framework in which the risk based inspection (RBI) methodology in the process industry has developed to its current status.

Besides the historical aspect of the methodology, there is also the aspect of national regulations which govern the safety of the production and use of pressurized equipment. The main objective of my research for DNV was to have DNV's risk based inspection methodology accepted by the Dutch authorities as a valid tool to determine safe inspection intervals for pressurized components. The second part of this chapter gives an outline of these regulations and the requirements for validation of the methodology.

1.1 Historical development

What is risk based inspection? DNV puts it as follows: "Risk Based Inspection is a combination of technologies providing industries with a risk based method for evaluating and developing inspection plans. RBI works by calculating both the consequences of possible failures and the likelihood of those failures. The combination of consequences and likelihood identifies which equipment warrants the most attention for managing risk." RBI is therefore a tool to first select which items require attention and then plan when and how to inspect the component or system. As a rule of thumb, 80% of the risk is caused by only 20% of the equipment. It is therefore essential that these high risk components are identified beforehand.

The origin of the RBI methodology for the process industry lies in the cooperation between DNV and the American Petroleum Institute (API) in the early 90's. This cooperation was initiated by DNV and also involved a large industry sponsor group consisting of some of the largest process companies in the United States and abroad. The research led to a large publication which is commonly known as the Base Resource Document, but was later published as the API581 [1]. There was a great need for this

methodology since most of the plants (especially those build in the post–2nd World War Europe) were reaching a critical age which required some form of inspection or maintenance planning. These plants were quickly built in order to get the European industry back on track, but they had a design life of only 15 years. Because the plants were still functional and profitable at this age, it was decided that they would not be closed or dismantled. Most countries have had regulations in some form or another which prescribed fixed intervals for the inspection of certain types of pressurized equipment. However, large parts of the plants did not fall under these regulations and the owners of the plants were eventually left with hundreds or even thousands of uninspected components with which they did not know what to do.

After laying the ground work for the implementation of RBI, DNV eventually parted with the research group and the API to work on their own implementation called ORBIT™ which stands for Onshore Risk Based Inspection Technology. Compared to the API methodology, ORBIT™ is a more fine tuned package which further exploits the extensive knowledge and experience within the DNV organization. Features were added which are not covered by the API program.

The software is easy to use due to the availability of several levels of detail for the calculations and the presence of ORBIT™ Launcher. The latter tool provides the user with the option to perform an initial screening of all items such that the detailed analysis can be concentrated on the high risk items only. A project like the one DNV Benelux performed for one of its largest customers, Dow Benelux B.V. (part of the Dow Chemical Company), shows the benefit of such a screening tool. The project involved the analysis of around 4500 items at their Terneuzen location (see figure 1.1) in the Netherlands. This location is Europe's largest ethylene production



Figure 1.1: The Terneuzen plant operated by Dow Benelux B.V. (Courtesy of Siebe Swart Fotografie Amsterdam).

site. By 1995, this plant was 30 years old and the maintenance cost formed a substantial part of the operating budget. This is due to the periodical inspections prescribed by the regulations. The DOW Chemical Company adopted

DNV's RBI methodology for its inspection planning worldwide, which results in addressing the maintenance spending to the equipment where we get the most benefit in terms of increased safety and reliability (up time of the plants). The risk methodology as used by DOW includes the aspects of safety, environment and business interruption. In the longer run, DOW is convinced that savings are created by switching from the old periodic inspections to a full risk based approach. Concentrating on the high risk components increases the level of safety, environmental care and reliability.

In Europe there is currently a large effort to develop a industry independent risk based maintenance and inspection methodology. It is known as the *RIMAP project*¹ (**R**isk **B**ased **M**aintenance **P**rocedures for European Industry), which is funded by the European Commission and joins 16 organizations from 8 countries in the effort. The goal is to create a methodology for maintenance and inspection planning for all types of equipment (not only pressurized equipment), which accounts for both the technical and managerial aspects. The application of this methodology will not apply to the nuclear power industry and will be limited to the in-service phase of the equipment.

There are a number of other industries which have already or are in the process of applying similar methodologies. See [24] for an overview and examples of the application of risk based inspection in the nuclear power industry.

1.2 Dutch and European regulations

As mentioned before, most countries have (limited) regulations regarding the maximum intervals between inspections for different types of equipment (e.g. drums, pipelines, columns etc.). The Netherlands is the first European country to extend these regulations to allow for so-called flexible inspection periods or risk based inspection periods. These changes followed the introduction of new European regulations regarding the production phase of pressurized equipment. These have been accepted by the European Union in 1997 as the Pressurized Equipment Directive (PED: [6]) and were implemented in the Netherlands as of November 1999.

The PED only regulates the certification of new equipment built inside or imported from outside the EU for use under a pressure of more than 0.5 bar. All equipment produced under this directive will receive a CE marking, which allows them to be freely traded throughout the member states of the European Union. Another aspect of the PED is the fact that the certification is no longer in the hands of a single (governmental) body. Any body which is appointed by the government will be able to perform the certification. This means that national institutes which had a monopoly over these

¹see <http://www.mpa-lifetech.de/rimap> for more information.

certifications until now (this was Stoomwezen B.V. in the Netherlands), will receive competition from other companies like DNV and AIB–Vinçotte.

One of the remaining aspects of the rules for pressurized equipment, is the regulation of periodical inspections during the operational phase of the equipment. A team of people representing governmental institutes and companies in the process industry have worked to update the Rules for Pressure Vessels [21] in The Netherlands to include the regulation of flexible inspection periods. Other aspects which will be included in the near future are regulations concerning taking the component into service and periodic repair and maintenance.

A key element of these regulations is that a production facility can apply for extended or flexible inspection periods if the time until the next inspection is obtained using an accepted risk based inspection planning methodology. Any company will therefore need to use their own method or that of a third party, which has been accepted by a committee and listed in the Rules for Pressure Vessels. A summary² of the requirements for the application is given by the following list:

1. The component for which the application is done, is categorized according to sheet G0402 and has already undergone one or two (depending on the category) inspections. There are some categories for which no flexible periodical inspections can be applied for.
2. The operator/owner of the equipment must comply with a number of organizational requirements according to sheet G0205. These include requirements concerning certification of the safety management system in place at the facility, the presence of a qualified inspection engineer etc.
3. The methodology used to determine the inspection period has to be listed in the rules as a valid method for risk based inspection planning.

In order to have the methodology accepted by the technical committee, the applicant has to submit a model description in order to demonstrate the compliance with the model specifications. This description should include the following aspects (as described in sheet T0260 appendix 1):

1. Concerning how the remaining life is determined, the following considerations have to be included:
 - an inventory of the most sensitive locations on the component together with the reasoning which supports the choice of these locations

²this summary is an unofficial translation and interpretation of the rules for pressure vessels and it should therefore not be considered as anything but such.

- determination of the possible interactions between degradation mechanisms
 - determination of the course of degradation under the given operational conditions and an extrapolation/trend analysis if the degradation can be modelled or approximated
2. The following aspects should be treated to determine the quality and the reliability of the calculated (expected) remaining life for the component:
 - indication of the applicability of the trend analysis or degradation model or in other words: is the chosen model a good choice to predict the active degradation mechanism?
 - the number of inspection measurements which are available and the time period over which this data is taken
 - the quality and applicability of the chosen inspection methods for the active degradation mechanism
 - how representative the locations at which the inspection results were obtained are for the overall condition of the component
 - the stability of the process of which the component is a part. Are the operating conditions stable or variable?
 3. the requirements for the calculations of the consequence of failure are quite limited. The rules only prescribe that at least the following parameters should be considered in the analysis:
 - flammability
 - toxicity
 - physical explosiveness
 - hazard probability with regard to the surrounding residential area
 4. the determination of the inspection interval should include the information of all three previous points. The maximum length of the interval is 4 times the regular fixed inspection period.

Together with the above model specification, the applicant has to supply the results of at least three test cases and compare them to the result obtained using the reference model which is specified in appendix 2 of sheet T0260.

Chapter 2

ORBIT™ Onshore

This chapter is a short description of some of the main components in the ORBIT™ methodology. In essence, it is a resumé of the ORBIT Technical Manual [16]. This chapter forms the basis for the topics discussed in later chapters. Although a software program is not a methodology in itself, it embodies DNV's methodology for risk based inspection planning. To explain how the methodology works, I've taken the outline of ORBIT™ as a basis.

The last section of this chapter includes the results from an uncertainty and sensitivity analysis, which includes the use of the UNICORN software package. The behavior of the corrosion and cracking models have been analyzed using cobweb plots and the sensitivity of the models to their input variables has been studied with the use of numerous sensitivity measures. This helps the user of the ORBIT™ methodology and software to better understand the influence which the input exercises on the result. A better understanding of the mathematical model behavior also helps DNV to clearly explain to its customers what ORBIT™ does with their data.

2.1 ORBIT™ overview

The methodology is built up around the concept of *risk*. Risk is defined as the product of the failure probability of an item with the consequence value due to failure of the item. The ORBIT™ software package therefore consists of two distinctive parts:

1. calculation of the failure probability based on the active degradation mechanism(s), material properties and operating conditions.
2. calculation of the possible consequence when failure occurs based on the type of failure, inventory characteristics and mass and information concerning the surrounding equipment and personnel.

Each part has two calculation levels: a simple qualitative method and a detailed quantitative analysis. The qualitative method is mainly used to

get an indication of which items have the highest risk associated with them. With this information the user can reduce the amount of work needed for the detailed analysis by filtering out the unimportant or low-risk items.

The calculation of the failure probability for a component is done in a number of steps. First the program uses the information on the type of inspection and the measurement obtained from it, to calculate a posterior density over the degradation rate using Bayesian updating. If no inspection data is available the program assumes a (conservative) default density for the degradation rate. Next the program approximates the failure probability using a reliability index method which is commonly used in structural reliability. This part of the program is performed by so-called *technical modules*, where each module corresponds to a degradation mechanism. The following degradation mechanisms are currently coded into ORBIT™ :

- internal thinning
- stress corrosion cracking (SCC)
- external thinning due to corrosion
- brittle fracture
- fatigue due to vibration

The two main modules are the thinning and SCC modules. They are also used by the external thinning module and the brittle fracture module is very similar to the cracking module. In the following text I will treat each module separately.

Since the concept of probability is not always easy to grasp for (inspection) engineers, the program uses so-called Technical Module SubFactors (*TMSF*). This factor is obtained by dividing the calculated failure probability with a generic cumulative failure probability. This generic value is an industry average value for a specific type of component, which is taken from a source like the OREDA database [17]. The TMSF therefore represents how much more or less likely the component is to fail compared to the industry average. Also, the program outputs a failure frequency (events per year) which it names the likelihood of failure (LoF). This name is somewhat confusing, especially since the industry cannot agree on the difference between PoF (probability of failure) and LoF, but we will not go into this discussion here.

ORBIT™ will add each individual TMSF obtained from the active technical modules and multiply this overall TMSF with the component specific generic failure probability to obtain an overall failure probability.

The second part of the ORBIT™ methodology is the calculation of the consequence of failure. Since this aspect of RBI is not heavily regulated (see section 1.2) and because the methodologies used for the calculations are

widely accepted as some of the best available, I will only briefly discuss this matter. The detailed consequence of failure is calculated using lookup tables generated by *PHAST*. This program uses the information on the release type (instantaneous or continuous) and the fluid properties, which are given by the user, to calculate the area equation coefficients for a number of weather scenarios. The area equations are based on the power law and are in the following form:

$$A = ax^b,$$

where A is the affected area, x is the release size and a and b are coefficients. ORBIT™ uses these scenario tables to calculate the size of the affected area. Combined with data on equipment costs and the amount of possible employees in this affected area, ORBIT™ calculates the consequence of failure.

This is however by far not a complete picture of the consequence modelling side. Besides calculating the affected area, the program also determines the toxicity and the flammability of the released fluid. All these aspects are incorporated using event trees and probability calculations in order to obtain the results. For a more complete description of the consequence calculations, I refer the reader to the ORBIT™ Technical Manual [16].

2.1.1 Thinning module

Thinning is the most common damage mechanism causing leaks in process piping and equipment. Thinning may take form of general wall loss or localized wall loss such as pitting. Carbon steel or copper are usually more susceptible to general thinning (although localized corrosion and pitting cannot be excluded), while stainless steels and higher alloy materials are usually more prone to localized thinning and pitting.

In ORBIT™ the user inputs either an estimated value for the corrosion rate obtained using expert opinion or he inputs the measured rate which is calculated from thickness measurements and the component's age. The program then uses a state function based on the material resistance (R) minus the applied stress (S):

$$g = \underbrace{S \left(1 - \frac{C \times \Delta t}{th} \right)}_R - \underbrace{\left(\frac{P \times d}{2th} \right)}_S, \quad (2.1)$$

where the variables and parameters with their description and dimensions are listed in table 2.1. The dimensions given in table 2.1 are used throughout this report. The tensile and yield strengths are determined by the material grade. The most common system for the classification of material grades is one developed by the American Society of Mechanical Engineers (ASME) and ASTM International (formerly known as the American Society for Testing and Materials). If you know the material grade, you can look up all the

<i>variable</i>	<i>dimension</i>	<i>description</i>
S	[MPa]	Flow stress = $\min \{1.1(YS + TS)/2, TS\}$
TS	[MPa]	Material tensile strength
YS	[MPa]	Material yield strength
C	[mm/yr]	Corrosion rate
Δt	[yr]	Time since service start
th	[mm]	Material thickness
P	[bar]	(Maximum) Operating pressure
d	[mm]	Component (internal) diameter

1 MPa = 10^6 Pa and 1 MPa = 10 bar.

Table 2.1: Description of the variables for the thinning module.

material properties, including the tensile and yield strength, in the ASME Materials Handbook. ORBIT™ has a database with the full material catalogue from which it obtains the necessary values relating to the material properties.

The *state function* is a commonly used concept in structural reliability. It determines when the system, structure or component is functional and when it is considered to have failed. As long as the resistance is larger than the applied stress, the item is considered to function safely. In this case $g > 0$. As soon as the stress becomes larger than the resistance the state function becomes negative: $g < 0$. The point at which $g = 0$ is called the *limit state* and represents the level at which the item will fail.

In the state function g in (2.1), there are three variables with some variability/uncertainty, namely S , C and P . The user's input is used as the mean value for each variable and the standard deviation is then determined using the following coefficients of variation:

$$\begin{aligned}\sigma_S &= 0.20\mu_S, \\ \sigma_C &= 0.10\mu_C, \\ \sigma_P &= 0.05\mu_P.\end{aligned}$$

These coefficients have been determined using expert judgment and are fixed in the methodology.

2.1.2 Stress corrosion cracking module

Stress corrosion cracking (SCC) is environmentally induced crack propagation, resulting from a combination of mechanical stress and corrosion reactions. It is a delayed failure process where cracks initiate and propagate until failure occurs. There are many specific material/environmental combinations which can result in a SCC problem. Some common chemicals associ-

ated with this are chlorides, sulfides and hydrogen. However, there are many other material/environmental combinations within which SCC can occur.

ORBIT™ assesses the failure probability for different types of SCC mechanisms for both carbon steels and stainless steels. It includes models for caustic cracking, amine cracking, carbonate cracking, sulfide stress cracking, hydrogen induced cracking, stress oriented hydrogen induced cracking for carbon steels and polythionic acid cracking, and chloride cracking for stainless steels.

As with the thinning module, the SCC module uses a resistance minus stress model based on Paris' crack growth law used in Linear Elastic Fracture Mechanics (LEFM):

$$g = \underbrace{K_{IC}}_R - Y \underbrace{\left(\frac{P \times d}{2th} + S \right)}_S \sqrt{A\pi}, \quad (2.2)$$

where the variables are described in table 2.2. The crack depth is calculated

<i>variable</i>	<i>dimension</i>	<i>description</i>
K_{IC}	[MPa√mm]	Material fracture toughness
Y	n/a	Geometrical factor
P	[bar]	(Maximum) Operating pressure
d	[mm]	Component (internal) diameter
S	[MPa]	Residual stress
th	[mm]	Material thickness
A	[mm]	Crack depth

Table 2.2: Description of the variables for the SCC module.

using information on the type of cracking, the susceptibility of the material to the type of cracking and whether or not the component has been post weld heat treated (PWHT). *PWHT* is performed on a component to relieve it of stresses which were created during welding on a part of the component. The residual stress depends heavily on this information. If there has been no PWHT, then $S = 0.2 \times YS$, whereas a PWHT'd component will have $S = 0.75 \times YS$. This treatment is therefore able to retain 75% of the material's original yield stress, compared to only 20% which is left after welding has been done on the material. The crack growth rate is determined by

$$C = \frac{l_{char}}{t_{grow}^n},$$

where the characteristic crack length l_{char} and the growth time t_{grow} are taken from a database included in the ORBIT™ program. The values in this database have been determined by expert opinion, as is the crack growth

exponent n . For PWHT'd components $n = 0.25$ and for non PWHT'd components $n = 0.50$. The crack depth is then obtained by dividing the crack length l with the crack length-to-depth ratio $R_{l/a} = 6$, which is based on API579 and is commonly used in LEFM. The crack length is given by $l = C \times \Delta t^n$, where Δt is the time since the service start of the component.

For stainless steels $K_{IC} = 300 \text{ ksi}\sqrt{\text{in}}^1$ and for carbon and low alloy steel

$$K_{IC} = \min\{33.2 + 2.806e^{(0.02(T+100))}, 200\} \text{ ksi}\sqrt{\text{in}},$$

where T is the operating temperature. For these types of steel the fracture toughness increases as the temperature increases until it levels off at about 110°F, therefore the failure probability decreases with increasing temperature.

In the state function g in (2.2), there are five variables with some variability/uncertainty, namely K_{IC} , Y , P , S and A . The input given by the user is used as a mean value for these variables and the standard deviation is determined using the following coefficients of variation:

$$\begin{aligned}\sigma_K &= 0.20\mu_K, \\ \sigma_Y &= 0.50\mu_Y, \\ \sigma_P &= 0.05\mu_P, \\ \sigma_S &= 0.25\mu_S, \\ \sigma_A &= 0.10\mu_A.\end{aligned}$$

These coefficients are based on expert opinion and cannot be changed by the user.

2.1.3 External corrosion

External corrosion can occur on materials with unprotected or insulated surfaces. The most likely form of external corrosion on carbon steels is general or localized thinning, while the most likely form on stainless steels is localized thinning or stress corrosion. ORBIT™ assesses whether a material is susceptible to external stress corrosion cracking (ESCC) based on information regarding material of construction, operating conditions, coating, insulation and climatic conditions. A combination of relatively high temperatures, austenitic stainless steels and poor quality insulation and coating usually leads to high susceptibility. If the material is susceptible to ESCC, ORBIT™ passes derived information to the stress corrosion cracking module discussed in the previous section. If the material is susceptible to external thinning, it passes an estimated external corrosion rate to the thinning module described in section 2.1.1. This corrosion rate is dependent on the operating temperature and the amount of rainfall in the area. There are three possible weather conditions: marine, temperate and arid.

¹1 ksi√in = 6.895√25.4 MPa√mm

2.1.4 Brittle fracture

At low temperatures carbon steels may suffer brittle fracture at loads significantly below the intended design loads. This is because carbon steels change from a ductile behaviour to a brittle behaviour at low temperatures. The temperature where this change occurs is known as the *transition temperature*. While most materials are designed to operate well above the transition temperature, some materials in low temperature service at older plants may operate near or below the transition temperature and as a result, be susceptible to brittle fracture. ORBIT™ directly estimates the TMSF based on information on the lowest possible temperature exposure, the material specification and grade, whether the material was impact tested and whether the component has been stress relieved.

2.1.5 Fatigue

Vibration may cause process equipment and piping to fail prematurely. In particular, improperly supported piping near vibration sources is prone to fatigue. In contrast to the other degradation modules that use quantitative methods to calculate a TMSF, the fatigue module is based on a simple indexing methodology to determine the TMSF.

2.2 Inspection updating

ORBIT™ uses Bayes' theorem to update the prior knowledge of the degradation rate with the information gained from an inspection. There are a number of different types of inspections ranging from simple visual inspection to ultrasonic thickness measurements. All inspections are of course non-destructive inspections and in ORBIT™ they are categorized in 5 levels of efficiency: highly effective, usually effective, fairly effective, poorly effective and ineffective.

Since it can be very difficult to find a representative continuous prior density for the degradation rate such that the posterior is easily calculated, ORBIT™ uses a discrete version of Bayes' theorem:

$$P(A_i|B_k) = \frac{P(B_k|A_i)P(A_i)}{\sum_{j \in S} P(B_k|A_j)P(A_j)}, \quad (2.3)$$

where $P(B_k|A_i)$ is called the *likelihood* and A and B are events. The initial belief about the event A is given by the *prior* $P(A_i)$. The division by the sum over all states (S is the set of all states and in our case $S = \{1, 2, 3\}$) normalises the result. If we interpret event A_i as the corrosion rate being in state i and B_k as the inspection measuring a rate in state k , then $P(A_i|B_k)$ can be interpreted as the probability that the corrosion rate is in state i given that our inspection told us it's in state k . The likelihood is therefore

interpreted as the probability of an inspection giving us state k given that it's state i . This conditional probability represents the effectiveness of the inspection in determining the true state of the corrosion rate.

Currently ORBIT™ uses only three so-called *degradation states* for the corrosion rate in the thinning module and the cracking rate in the SCC module. Table 2.3 gives the discrete prior density for both modules, which is assigned to the degradation rate in the case when there is no inspection data available. These prior densities therefore put a default uncertainty towards

<i>state</i>	<i>prob.</i>	<i>thinning</i>	<i>cracking</i>
1	50%	less or equal	to the estimated rate
2	30%	up to 2 times	up to 5 times the estimated rate
3	20%	up to 4 times	up to 10 times the estimated rate

Table 2.3: Discrete prior densities for the thinning and SCC modules.

higher degradation rates on the rate which has been given by the user. This is done to ensure enough conservatism and to save the user the tedious process of giving his own uncertainty distribution over the degradation rate for each item. On average, a well trained user will be able to enter the data of approximately 50 components into ORBIT™ in one day. For large plants with hundreds of items, this process can therefore take up a large amount of time. The program avoids complicated data entry, because otherwise the time required for the data entry becomes too long.

The likelihood for the inspection is given by a table similar to table 2.4. The likelihood of a highly effective inspection correctly identifying a rate in state 1 is therefore 90%, leaving only a 10% probability of an incorrect measurement. The ORBIT™ documentation includes a detailed list of inspection types and measurement conditions from which the user determines the correct classification of the inspection based on the above system.

Consider corrosion under insulation, which occurs when the heat insulation is damaged and water penetrates the insulation material down to the wall of the component. In this case a highly effective inspection is one where there is first a total visual inspection of the whole component, followed by the removal of the insulation at almost all suspect areas and an ultrasonic wall thickness measurement at these locations. These areas are usually around the support anchors which hold the component in its place. Vibrations cause the outer insulation to tear enabling water or other types of fluid to enter the insulation and corrode the construction material. An ineffective inspection would be a total visual inspection followed by the removal of less than 5% of the insulation at sensitive areas and no ultrasonic testing. The intermediate levels of inspection depend on the amount of sensitive locations which are uncovered and whether or not a ultrasonic thickness measurement has been performed.

Highly effective inspection:	$k = 1$	$k = 2$	$k = 3$
$i = 1$	0.90	0.05	0.01
$i = 2$	0.09	0.90	0.09
$i = 3$	0.01	0.05	0.90
Usually effective inspection:	$k = 1$	$k = 2$	$k = 3$
$i = 1$	0.70	0.15	0.10
$i = 2$	0.20	0.70	0.20
$i = 3$	0.10	0.15	0.70
Fairly effective inspection:	$k = 1$	$k = 2$	$k = 3$
$i = 1$	0.50	0.25	0.20
$i = 2$	0.30	0.50	0.30
$i = 3$	0.20	0.25	0.50
Poorly effective inspection:	$k = 1$	$k = 2$	$k = 3$
$i = 1$	0.40	0.30	0.27
$i = 2$	0.33	0.40	0.33
$i = 3$	0.27	0.30	0.40
Ineffective inspection:	$k = 1$	$k = 2$	$k = 3$
$i = 1$	0.33	0.33	0.33
$i = 2$	0.33	0.33	0.33
$i = 3$	0.33	0.33	0.33

Table 2.4: Likelihood of measured state (k) given actual state (i).

2.3 Inspection planning

The ORBIT™ methodology determines the optimal inspection period with safety as its only objective. The program evaluates the risk at a future evaluation date which is entered by the user. It uses risk targets which have been defined by the user to determine the optimal type of inspection. Starting with no inspection at the evaluation date, the program will calculate the ratio of expected risk over the risk target. If this ratio is larger than 1, then it will recalculate this ratio using a fairly effective inspection at the evaluation date. It will keep trying better inspections until the ratio is smaller than one or the highly effective inspection can not reduce the risk adequately.

This methodology will recommend the cheapest type of inspection at a fixed time in the future, but it will not optimize for both economics and time simultaneously. There is also the option of using the semi-quantitative screening method which will place the item in a 5×5 risk matrix as is shown in figure 2.1.

Qualitative Inspection Planning:
Years Until Next Inspection by Risk

LIKELIHOOD CATEGORY	5	6	4	3	1	1
	4	8	6	4	3	2
	3	9	9	7	5	2
	2	10	10	8	7	5
	1	12	12	11	11	11
		A	B	C	D	E
		CONSEQUENCE CATEGORY				

Figure 2.1: ORBIT™ risk matrix example.

2.4 Sensitivity and uncertainty analysis

One of the goals of this research is to obtain a better understanding of the sensitivities and uncertainties that play a role in the ORBIT™ methodology. The consultants who work at DNV are mostly inspection engineers with a background in mechanical, civil or chemical engineering. They have excellent knowledge of material properties, the matter inside the component and the damage mechanisms which affect the safe operation of the item. During the input of the plant data into the ORBIT™ software, the consultants discuss each item and its properties with the client before the data is accepted. Usually the client will supply a process engineer who is highly familiar with the components under analysis to do this job. This engineer might ask a question like for example: “Suppose the pressure becomes greater than the value we are currently using for the analysis. Would this influence the result in such a way that the state of the component is considerably more or less critical than our assumption?” A question like this will be given a physical explanation by the consultant. He might tell the engineer that increasing the pressure will not substantially influence the state of the component unless the temperature is higher, which is due to the properties of the construction material. The consultant will not be able to exactly explain to the engineer how ORBIT™ models this effect or how sensitive the model output is to the pressure. This is because the consultant is not familiar with the detailed technical functioning of the methodology. It is therefore very interesting to look at the model sensitivities from a mathematical point of view. Also, I will apply the UNICORN uncertainty analysis program, developed by the decision theory group at the TUDelft, to this model. This will enable us to get a better insight into the dependencies and the uncertainties of the

model.

The sensitivity and uncertainty analysis is preceded by the introduction of the data for the calculations and a short overview of the sensitivity measures which will be applied to the ORBIT™ models.

2.4.1 Example data

For all calculations in the rest of this document I will use the data in table 2.5. The strength of the material can be obtained from ASME part 2 (see [2]),

Component type:	4in. pipe
Material type:	austenitic stainless steel
Material grade:	TP312-316L
Min. tensile strength:	482.63 MPa
Min. Yield strength:	172.37 MPa
Operating pressure:	4.1 bar
Operating temperature:	93° C
Internal diameter:	102 mm
Initial wall thickness:	6 mm
Corrosion rate:	0.13 mm/yr
Cracking cause:	chloride
Cracking susceptibility:	medium
PWHT:	no
Insulation:	no
Coating:	no
Service start:	Dec. 1st, 1997
Inspection:	none

Table 2.5: Example data.

which prescribes the necessary properties for specific grades of construction materials. ORBIT™ has these properties included in its database and they can be consulted by the user at any time.

2.4.2 Sensitivity measures

In the next section we will use three different sensitivity measures: the product moment correlation, the correlation ratio and the FORM sensitivity index.

The most common and well known measure of concordance is the *product moment correlation*, also referred to as simply the correlation (see for example Ross [20]). For two variables X and Y , the correlation is defined as

$$\rho(X, Y) = \frac{E(XY) - E(X)E(Y)}{\sigma_X \sigma_Y},$$

where σ_X is the standard deviation of X and $\sigma_X^2 = \text{Var}(X)$ is the variance of X . The correlation ρ is a directional sensitivity measure for which $-1 \leq \rho \leq 1$ holds. If X and Y are independent, then $\rho = 0$, but this is not necessarily true the other way around. For our analysis we will be interested in the correlation between each input variable and the model output: $\rho(g(\mathbf{X}), X_i)$.

The *correlation ratio* is a non-directional global sensitivity measure which is defined as

$$CR(g(\mathbf{X}), X_i) = \frac{\sigma_{E(G|X)}^2}{\sigma_G^2}.$$

This measure can be taken as the general sensitivity of the model output $g(\mathbf{X})$ to the input variable X_i . It is clear that the correlation ratio is always positive, therefore it does not give any information about the direction of the influence. As can be seen from the definition, the correlation ratio is a measure of how much variance (relative to the total variance) is caused by the variable X in the model output G .

The third and last sensitivity measure which will be used is the *FORM sensitivity measure*. For each variable X_i this measure is given by the corresponding element of the unit directional vector:

$$\alpha_i = \frac{\partial g}{\partial x_i}(\mathbf{x}) \|\nabla g(\mathbf{x})\|^{-1},$$

where the partial derivative and the vector norm are defined in the list starting on page 105. This measure is a direct result of applying the Rackwitz–Fiessler method for the FORM reliability index method as will be discussed in section 3.2.1. The point \mathbf{x} at which this measure should be evaluated is called the design point and α_i is the slope of the tangent line on the failure surface in this point. This FORM sensitivity measure is therefore a measure of the contribution of each variable to this slope.

2.4.3 Corrosion model

We assume that all variables in the model (2.1) are independent. The thinning module is sensitive to the time at which we consider the output sensitivity. At first, the failure probability due to corrosion is low. In this stage the model is fully determined by the material strength S . As time progresses, the wall loss will start increasing the failure probability considerably. Between this point and 100% probability of failure, the model is most sensitive to the corrosion rate C . For the example defined in section 2.4.1, I will determine the value of the sensitivity measures at $t = 9$ years and $t = 24$ years. The failure probability is approximately 0 and 50% (obtained using Monte Carlo simulation) respectively. Table 2.6 summarizes the result of the analysis². It is clear from these results that the pressure has hardly

²with many thanks to Daniel Lewandowski for his spreadsheet implementation of the correlation ratio calculations.

	for $t = 9$:			for $t = 24$:		
	S	C	P	S	C	P
$\rho(g(\mathbf{X}), X_i)$	0.96	-0.29	-0.04	0	-0.98	-0.04
$CR(g(\mathbf{X}), X_i)$	0.902	0.086	0.004	0.003	0.947	0.007
α_{X_i}	0.99	-0.002	-0.003	0.04	-0.99	-0.009

Table 2.6: Sensitivity analysis results for the thinning model at $t = 9$ years and $t = 24$ years.

any influence on the model output. This is of course due to the (very) low pressure of 4.1bar which is used. Otherwise the three sensitivity measures confirm one another. For $t = 9$ the material strength is highly positively correlated with the model output g . At $t = 24$, when the component has an increasing probability of failure, higher values for the corrosion rate result in lower values for the model output. There is almost perfect negative correlation. The fact that this correlation is negative is no surprise when we look at the model formulation 2.1, because this variable has a minus sign in front of it. This is also why the pressure is (slightly) negatively correlated with g .

To be able to find more complex interactions, we use so-called *cobwebs*. These plots are generated by the program UNIGRAPH, which comes with the UNICORN package. They are a representation of the multidimensional model output. The model output is generated by sampling the input variables and calculating the value of g for each set of samples. Besides the graphical benefit of the cobweb plots, UNICORN has the option of pre-defining a correlation structure between the input variables such that the dependencies are incorporated in the Monte Carlo simulation. Figure 2.2 shows the cobweb for the thinning model. The vertical lines represent the percentiles of the input and output variables of the model. This scaling is most useful, but UNIGRAPH also allows you to use the actual sample value on a natural or logarithmic scale. The distributions in the figure are (from left to right) the corrosion rate cr , model output g , material strength s and operating pressure p . The horizontal lines represent the sample path, i.e. each sample set is connected with the corresponding output samples. In this case there is only one output sample for g ; the rest are input samples.

To interpret the cobweb plot we have to look at the shape of the web which is formed by the sample paths. Since we assumed that the input variables are independent, we can look only at the interaction between these input variables and the model output. The pattern of horizontal lines between s and g suggest a positive correlation: higher values of s are connected with higher values of g and visa versa. This confirms the almost perfect positive correlation $\rho(g, s) = 0.96$ in table 2.6. A crossed pattern like the one between cr and g suggest a negative correlation: high values of cr tend to




Figure 2.2: Cobweb plot for the thinning model at $t = 9$ years.

result in low values in the output g . This is not a perfect ‘X’-shaped cross due to the not so strong correlation of $\rho(g, cr) = -0.29$. Figure 2.3 shows the same result for $t = 24$ years. Here the strong negative correlation between cr and the output g results in a perfect ‘X’. We also see that now s is no longer correlated with the output g . Things get really interesting when we make use of the option to set conditions into the variables. Figure 2.4 has the pressure p conditionalized on the top 10% and the corrosion rate cr is conditionalized on the top and bottom 10%. Also, there is an extra output called g_high which is the same model as g , but using a very large diameter $d = 1m$. Using the indications given by the consultants, I expect to find an interesting combined relationship between p and g for a large and small diameter. Physically, a high operating pressure will be of more concern for items with a large diameter. This means that a high pressure together with a small diameter should result in higher values of g compared to when the diameter is large. We can see this clearly in figure 2.4: the condition on the top 10% of p puts the samples in the top 20% of g (small diameter) and top 50% of g_high (large diameter). A large diameter in combination with a high operating pressure will therefore result in a higher failure probability (i.e. lower value for the state function g). The conditioning on the top 10% of cr and p show that the model output for both diameters is extremely low.

Overall, the influence of the variables is largely proportional to the variance which is associated with them (see section 2.1.1). The choice for the coefficients of variation should therefore be done with great care. Based on the sensitivity and uncertainty analysis, we can conclude that the corrosion




Figure 2.3: Cobweb plot for the thinning model at $t = 24$ years.

rate is by far the most important variable in this corrosion model. The more we can reduce the uncertainty over this rate, the better our results will be.

2.4.4 Cracking model

The example which we are using is not very likely to fail due to cracking. This is because it is a stainless steel, which has the property of having a high toughness coefficient and a high residual strength even if it hasn't been PWHT'd. The results are therefore much less dependent on the time at which we evaluate the model. Only in the very first years is there a significant difference in the sensitivity of the model to the input variables. As table 2.7 shows, the toughness coefficient K plays a leading role in the first year(s), with the geometry factor Y as runner-up. Later on, the influence of the material toughness reduces a little and the geometry factor becomes more important. All three sensitivity measures result in similar conclusions, except that the correlation ratio seems small for Y and S at $t = 1$ and for S at $t = 31$. Otherwise there are no special interactions which can be found in the model. Once again the influence of the variables is proportional to the variance which they have been given (see section 2.1.2). The pressure is of little influence to the output, therefore a change in pressure will not result in dramatically different results.

Unlike in the analysis of the thinning module, a larger diameter does not have the same effect in the SCC module. Conditioned on the top 10% of the operating pressure, the density over the model output g is essentially




Figure 2.4: Cobweb plot for the thinning model at $t = 24$ years conditionalized on cr and p .

the same for both $d = 102\text{mm}$ and $d = 1\text{m}$. The equivalent of figure 2.4 for the cracking model is given in figure 2.5. In this figure, the following variables are depicted from left to right: a_{t31} , k , y , s , g_{t31} , p and g_{t31h} . Here a_{t31} represents the crack depth at $t = 31$ years. This is the damage specific parameter equivalent to the corrosion rate in the thinning model. In the figure p is conditioned on the top 20% and a_{t31} is conditioned on the top 10% and the lower 10%. We can see that the crack depth has no correlation with the output g , but more importantly: the conditionalization on the operating pressure results in an equal distribution of the sample paths along g_{t31} (small diameter) and g_{t31h} (large diameter).

	<i>for t = 1:</i>				
	<i>K</i>	<i>Y</i>	<i>P</i>	<i>S</i>	<i>A</i>
$\rho(g(\mathbf{X}), X_i)$	0.94	-0.28	-0.01	-0.14	-0.01
$CR(g(\mathbf{X}), X_i)$	0.897	0.089	0.002	0.022	0.004
α_{X_i}	0.903	-0.353	-0.001	-0.239	-0.059
	<i>for t = 31:</i>				
$\rho(g(\mathbf{X}), X_i)$	0.77	-0.55	-0.01	-0.28	-0.07
$CR(g(\mathbf{X}), X_i)$	0.585	0.306	0.002	0.076	0.010
α_{X_i}	0.655	-0.617	-0.002	-0.424	-0.107

Table 2.7: Sensitivity analysis results for the SCC model at $t = 1$ years and $t = 31$ years.




Figure 2.5: Cobweb plot for the scc model at $t = 31$ years conditionalized on a and p .

Chapter 3

Reliability index methods

In this chapter I will discuss a number of different methods to approximate the failure probability using a state function like the thinning and SCC state functions. The most common method to approximate the failure probability in structural and civil engineering is the First Order Reliability Method (FORM). It is a relatively simple technique to approximate the solution of what is essentially a high-dimensional integral. Besides FORM there is also the more simple, but in most cases not correct, Mean Value First Order Second Moment (MVFOSM) method. Also, we can approximate the failure probability using simulation (e.g. Monte Carlo simulation, importance and directional sampling). I will consider all of these methods and also a number of different approaches to implementing the common FORM method.

3.1 Reliability index concept

The state functions (2.1) and (2.2) for the thinning and SCC models respectively are functions of time t , i.e. $g = g(t)$. However, we will evaluate the state of the material or structure at a fixed point in time, therefore the state function will be a function of the random variables in the model. These are the input variables which have uncertainty associated with them and this uncertainty/variability is represented by a suitable probability distribution. The most commonly used distribution in structural engineering is the normal distribution.

In the general case we consider the limit state function $g(X_1, \dots, X_n)$, where the X_i 's are n independent normally distributed random variables. The probability of failure $P(g(\mathbf{X}) < 0)$ is the total mass of the joint density for g which is in the failure region $\Omega = \{\mathbf{x} | g(\mathbf{x}) < 0\}$. This mass is given by the n -dimensional integral:

$$P(g(\mathbf{X}) < 0) = \int_{\mathbf{x} \in \Omega} f_g(x_1, \dots, x_n) dx_1 \dots dx_n, \quad (3.1)$$

where $f_g(\mathbf{x})$ is the joint density for g . This integral can usually not be solved analytically, therefore we would like to approximate the probability of failure. This is done using the concept of the *reliability index*, which is usually denoted by β . The FORM and MVFOSM are two different methods of determining the reliability index. When we know the value of β , the probability is found by calculating the value of the standard normal distribution at $-\beta$: $\Pr\{g < 0\} = \Phi(-\beta)$. The following discussion on the FORM method will clarify the geometrical interpretation of the reliability index. For more on this subject, I suggest the well known book by Madsen et al. [10] or the slightly older but equally useful book from Thoft–Cristensen and Baker [23]. A good overview of the following reliability index methods is given in the article of Haldar and Mahadevan [7].

3.2 FORM

The FORM method is a relatively simple method which approximates the limit state g in the so-called design point by a linear function (this is why it's called a first order method). In the case where g is a linear function of the X_i 's, the implementation is very straightforward, but in general the function g is non-linear. This requires us to pay good attention to what we are doing.

To be able to visualize how the method works, I will consider the two dimensional case: $g = g(X_1, X_2)$. Both X_1 and X_2 are normally distributed with mean μ_{X_1} , variance $\sigma_{X_1}^2$ and mean μ_{X_2} , variance $\sigma_{X_2}^2$ respectively. We can plot the contours of the 2-dimensional joint normal distribution together with the line $g(X_1, X_2) = 0$ as is done on the left in figure 3.1. Next we




Figure 3.1: Graphical representation of the First Order Reliability Method.

perform a transformation T , which transforms the variables to standard

normal (mean is 0 and variance is 1) variables:

$$T : Y_i = \frac{X_i - \mu_{X_i}}{\sigma_{X_i}}. \quad (3.2)$$

The contours are therefore moved to the origin of the X_1X_2 -plane and the limit state $g = 0$ is also moved accordingly. Now we define the reliability index β as the shortest distance to the limit state $g(Y_1, Y_2) = 0$. This index is also known as the Hasofer–Lind ([8]) reliability index. It can be interpreted as being the distance to the worst combination of variables resulting in failure of the system. Since the variables are standard normally distributed, the density along the dotted line in the right plot of the figure is also standard normal. We can therefore approximate the probability of failure by determining $P(g(\mathbf{y}) < 0) = \Phi(-\beta)$, where $\Phi(x)$ is the cumulative standard normal distribution.

The limit state function $g(\mathbf{y})$ is linearized in the design point y^* . This is done by expanding the function using *Taylor series* about $(Y_1, \dots, Y_n) = (y_1^*, \dots, y_n^*)$ as follows:

$$g(\mathbf{y}) \approx g(y_1^*, \dots, y_n^*) + \sum_{i=1}^n \frac{\partial g}{\partial Y_i} (Y_i - y_i^*) + H.O.T. \quad (3.3)$$

We disregard the higher order terms (H.O.T.) under the assumption that these are small compared to the first order terms.

The main part of this method is how to find the design point \mathbf{y}^* . The problem can be formulated as a nonlinear equality constrained optimization problem:

$$\begin{aligned} \text{minimize} \quad & d = \|\mathbf{y}\| = \sqrt{\mathbf{y}^T \mathbf{y}} \\ \text{such that} \quad & g(\mathbf{y}) = 0 \end{aligned} \quad (3.4)$$

In almost every textbook you will see a β in the place of the distance d in the above model. Unfortunately this is incorrect, because this distance measure is always positive (due to the norm) whereas the β index can also be negative. This is the case when the origin is in the failure region of the state function. Therefore the reliability index should be given by

$$\beta = \begin{cases} d & \text{if } g(\mathbf{0}) \geq 0, \\ -d & \text{if } g(\mathbf{0}) < 0. \end{cases}$$

This is very important, because otherwise the failure probability never increases above $\Phi(0) = 0.5$! Either the authors of these textbooks know this, but do not find it interesting enough to mention or they've never implemented this method themselves and simply copied the text from previous textbooks or papers.

Next, I will describe and compare the performance of three different methods to find the design point. Essential in these methods is that the state function $g(\mathbf{y})$ is in the standard normal space:

$$g(y_i) \equiv g(\sigma_i y_i + \mu_i). \quad (3.5)$$

3.2.1 Rackwitz–Fiessler

Almost any text you will find today uses the Rackwitz–Fiessler method ([19]), which is a sequential search for the design point. The design point \mathbf{y}^* is the limit of a sequence $\mathbf{y}^{(1)}, \mathbf{y}^{(2)}, \dots, \mathbf{y}^{(i)}, \dots$ which is calculated by iteratively solving

$$\mathbf{y}^{(i+1)} = \underbrace{\left[\left(\mathbf{y}^{(i)} \cdot \boldsymbol{\alpha}^{(i)} \right) + \frac{g(\mathbf{y}^{(i)})}{\|\nabla g(\mathbf{y}^{(i)})\|} \right]}_{\beta^{(i)}} \boldsymbol{\alpha}^{(i)}, \quad (3.6)$$

where $\nabla g(\mathbf{y})$ is the gradient vector given by

$$\nabla g(\mathbf{y}) = \left(\frac{\partial g}{\partial y_1}(\mathbf{y}), \dots, \frac{\partial g}{\partial y_n}(\mathbf{y}) \right)$$

and $\boldsymbol{\alpha}^{(i)}$ is a unit normal vector directed towards the failure set:

$$\boldsymbol{\alpha}^{(i)} = -\frac{\nabla g(\mathbf{y}^{(i)})}{\|\nabla g(\mathbf{y}^{(i)})\|}. \quad (3.7)$$

The *norm* of a vector \mathbf{x} is defined as $\|\mathbf{x}\| = \sqrt{x_1^2 + x_2^2 + \dots + x_n^2}$ and $\mathbf{y} \cdot \boldsymbol{\alpha}$ is the *inner product* of the vectors \mathbf{y} and $\boldsymbol{\alpha}$ (see also the list on page 105). The vector $\boldsymbol{\alpha}$ is a *unit normal vector*, which means that $\|\boldsymbol{\alpha}\| = 1$.

The above definition for the algorithm is taken from Madsen, Krenk & Lind [10]. The algorithm for the Rackwitz–Fiessler method is to first assume a starting point $\mathbf{y}^{(0)}$ and to iteratively solve equation (3.6) until the difference $\|\mathbf{y}^{(i+1)} - \mathbf{y}^{(i)}\| < \epsilon$, where ϵ is very small. When convergence is reached, the design point is given by

$$\mathbf{y}^* = \beta \boldsymbol{\alpha}^*, \quad g(\mathbf{y}^*) = 0.$$

There is of course a possibility that the algorithm diverges in which case a different starting vector should be chosen. Since the state functions are rarely linear, there will be more than one local optimum. It is advised to run the procedure a number of times with different starting points and to take the smallest β from these runs as the solution to the problem. In this way there is a bigger chance of finding the global optimum.

The elements of the unit vector (3.7) are often considered as sensitivity measures. For example, α_1^* would be a measure of how sensitive the model is to a change in the first variable Y_1 . This sensitivity measure has already been introduced in section 2.4.2 for the sensitivity analysis of the ORBIT™ corrosion and cracking models.

3.2.2 Lagrange multipliers

People familiar with mathematical optimization techniques will immediately think of using Lagrange multipliers when they see a problem of the form in (3.4). The general form of a nonlinear equality constrained optimization problem is given by

$$\begin{aligned} & \text{minimize} && f(x) \\ & \text{such that} && g_i(x) = b_i \quad \text{for } i = 1, 2, \dots, n \end{aligned}$$

where either f and g or both are nonlinear. The definition of the *Lagrangian function* $L(x, \boldsymbol{\lambda})$ is as follows

$$L(x, \boldsymbol{\lambda}) = f(x) + \sum_{i=1}^n \lambda_i (g_i(x) - b_i). \quad (3.8)$$

The λ_i 's are the *Lagrange multipliers*. Now it is easy to prove that the problem can be reduced to

$$\text{minimize } L(x, \boldsymbol{\lambda}).$$

The solution $(x^*, \boldsymbol{\lambda}^*)$ will give us the solution x^* to the original problem. By setting the derivatives to each variable of the Lagrangian function equal to zero and solving the resulting system of equations, we can find the optimal solution. In vector form we may write this as

$$\nabla_x L(\mathbf{x}, \boldsymbol{\lambda}) = \nabla f(\mathbf{x}) + \sum_{i=1}^n \lambda_i \nabla g_i(\mathbf{x}) = 0, \quad (3.9)$$

$$\nabla_\lambda L(\mathbf{x}, \boldsymbol{\lambda}) = \mathbf{g}(\mathbf{x}) - \mathbf{b} = 0. \quad (3.10)$$

If we apply this theory to our reliability index optimization problem (3.4), then we get

$$L(\mathbf{y}, \lambda) = \|\mathbf{y}\| + \lambda g(\mathbf{y})$$

and therefore we need to solve the following set of equations:

$$\begin{aligned} \frac{\partial L}{\partial y_1} &= \frac{\partial}{\partial y_1} \|\mathbf{y}\| + \lambda \frac{\partial g}{\partial y_1} = 0, \\ &\vdots \\ \frac{\partial L}{\partial y_n} &= \frac{\partial}{\partial y_n} \|\mathbf{y}\| + \lambda \frac{\partial g}{\partial y_n} = 0, \\ \frac{\partial L}{\partial \lambda} &= g(\mathbf{y}) = 0. \end{aligned}$$

This works fine, but we can simplify this problem even further. Since we are trying to find the closest point on the limit state function g to the origin, we

can also minimize $f = \|\mathbf{y}\|^2 = \mathbf{y}^T \mathbf{y}$, because the optimal point will be the same. To find the distance, we have to take the square root of the resulting value of the objective function f . In the example in the next section we will see that this makes solving the problem much more efficient. It's easy to see why this is when you compare the derivatives of the norm and the squared norm:

$$\begin{aligned} \text{regular: } \frac{\partial}{\partial y_i} \|\mathbf{y}\| &= \frac{1}{\|\mathbf{y}\|} y_i, \\ \text{simplified: } \frac{\partial}{\partial y_i} (\mathbf{y}^T \mathbf{y}) &= 2y_i. \end{aligned}$$

Doing things the regular way will result in a much more nonlinear equation compared to the simplified form. When the state function $g(\mathbf{y})$ is linear and we use the simplified objective function, the solution can be found by solving a simple set of linear equations. Unfortunately, this is usually not the case, but the simplified objective function will significantly reduce the required number of iterations to solve the system of nonlinear equations. I will refer to this method as the *simplified Lagrange multipliers*¹ method.

3.2.3 FORM performance comparison

Using the thinning model introduced in section 2.1.1 and the data from the example introduced in table 2.5, we apply the three methods (i.e. Rackwitz–Fiessler, regular Lagrange multipliers method and the simplified Lagrange multipliers method) which I've described in the previous section and compare their performance and results.

For the calculations and the implementation of the FORM methods for the thinning model, I refer to appendix A.1. Using a spreadsheet we can calculate the β and the failure probability $\Phi(-\beta)$ at increasing points in time. Also, to compare the performance of the methods, we keep track of the number of iterations n which are required for convergence and we evaluate the state function at the resulting design point (i.e. $g(\mathbf{y}^*)$) to check how close this is to the limit state). All three methods continue iterating until $\|\mathbf{y}^{(i+1)} - \mathbf{y}^{(i)}\| \leq 0.0001$. The Rackwitz–Fiessler method is limited to 10000 iterations, whereas the Lagrange methods are limited to 100 iterations. Table 3.1 summarizes these results at three year intervals.

It is clear that the Rackwitz–Fiessler search method does not perform very well: it requires a very large amount of iterations before convergence is reached. The Lagrange multipliers method with the regular objective function requires far less iterations, but the resulting failure probability jumps from 0 to 1 after 24 years. This can be seen in figure 3.2. The calculations performed by these two methods are of a very different type. The

¹I suggest to use the acronym “SLAM” as a reference to this method.

t	<i>Rackwitz–Fiessler</i>				<i>Regular Lagrange mult.</i>				<i>Simplified Lagrange mult.</i>			
	β	PoF	n	$g(\mathbf{y})$	β	PoF	n	$g(\mathbf{y})$	β	PoF	n	$g(\mathbf{y})$
3	46.94	0	786	0.060	136,24	0	19	8,39E-14	4,94	3,81E-07	5	3,35E-13
6	53.97	0	1032	0.060	61,52	0	23	9,02E-16	4,93	4,00E-07	6	-8,88E-15
9	63.42	0	1414	0.060	38,32	0	33	4,59E-13	4,92	4,26E-07	6	5,38E-12
12	76.71	0	2055	0.060	28,36	0	30	-1,05E-15	4,90	4,69E-07	6	5,46E-10
15	96.64	0	3292	0.060	11,44	0	11	-3,50E-14	4,87	5,50E-07	7	-2,17E-14
18	129.27	0	7046	0.061	6,78	0	8	-2,63E-12	4,81	7,52E-07	6	-6,70E-12
21	76.05	0	7335	0.061	20,83	0	39	1,48E-16	5,01	2,63E-07	100	-3,73
24	1.036	0.15	2138	0.061	20,61	0	8	7,24E-10	0,04	0,482	6	-1,67E-11
27	-49.04	1	4418	-0.061	-20,70	1	12	-3,64E-16	-2,07	0,980	7	3,67E-12
30	-86.72	1	3852	-0.060	-20,92	1	12	-1,14E-15	-3,76	0,999	8	6,75E-12
33	-118.47	1	4055	0.040	-21,21	1	11	-3,71E-16	-5,14	1	8	5,25E-10
36	-85.97	1	1567	-0.059	-21,51	1	16	-3,17E-16	-5,09	1	12	-2,27E-09
39	-68.85	1	1135	-0.060	-21,82	1	15	-4,71E-16	-7,27	1	11	2,87E-09
42	-57.58	1	854	-0.059	-22,11	1	11	-5,97E-16	-8,11	1	17	1,37E-14
45	-49.52	1	665	-0.059	-22,39	1	10	-1,10E-14	-22,39	1	21	-1,69E-13
48	-43.45	1	532	-0.059	-22,64	1	11	2,78E-16	-5,04	1	9	-3,91E-11
51	-38.73	1	435	-0.059	-22,88	1	13	5,68E-17	-5,04	1	9	3,28E-14
54	-34.94	1	363	-0.059	-23,11	1	15	-3,88E-15	-5,03	1	8	1,09E-10
57	-31.84	1	308	-0.058	-23,31	1	17	-6,84E-17	-5,03	1	9	-6,66E-15
60	-29.26	1	264	-0.058	-23,51	1	12	-1,35E-11	-5,03	1	9	4,88E-15
63	-27.08	1	230	-0.056	-23,68	1	11	3,33E-17	-5,03	1	9	-1,73E-14
66	-25.22	1	201	-0.058	-23,85	1	12	-5,54E-11	-5,02	1	9	-3,99E-14
69	-23.60	1	178	-0.057	-24,00	1	15	-6,56E-16	-5,02	1	9	8,74E-14
72	-22.20	1	159	-0.056	-24,15	1	24	1,78E-16	-5,02	1	8	7,49E-11
75	-20.96	1	143	-0.055	-24,28	1	21	-8,10E-17	-5,02	1	8	2,86E-11
78	-19.86	1	129	-0.055	-24,41	1	20	-2,91E-16	-5,02	1	8	8,71E-09
81	-18.88	1	117	-0.055	-24,53	1	13	-3,45E-13	-5,02	1	9	-6,66E-14
84	-18.00	1	107	-0.054	-24,64	1	14	-1,03E-11	-5,01	1	9	-3,15E-14
87	-17.21	1	98	-0.054	-24,74	1	14	3,10E-16	-5,01	1	9	1,29E-12
90	-16.49	1	90	-0.055	-24,84	1	21	-1,21E-16	-5,01	1	9	1,07E-10

Table 3.1: Results for three FORM implementations including failure probability and number of iterations used ($PoF = \Phi(-\beta)$).

Rackwitz–Fiessler method is a simple search method, therefore each iteration represents a step in the direction of the optimal point (this is a geometrical interpretation of the unit vector $\boldsymbol{\alpha}$). The iterations in the Lagrange multiplier method are those used by the Newton-Raphson iteration method to solve a system of nonlinear equations. This iteration scheme is generally the fastest method available, whereas the search method is very slow. The literature never mentions that the Rackwitz–Fiessler algorithm requires as many iterations as we see in table 3.1. Why does it perform so poorly in this case? Well, later on in this thesis we’ll see that the shape of the failure region is almost linear in the area of the design point. This means that the tangent hyperplane used by this algorithm will ‘jump’ around between points on either side of the design point. A more slanted plane will result in the approximation points to ‘move’ closer to the design point much faster.

The actual computing time required by both methods depends on the program which is used. The Newton-Raphson scheme requires a matrix inversion at every iteration and the speed at which this is done depends on

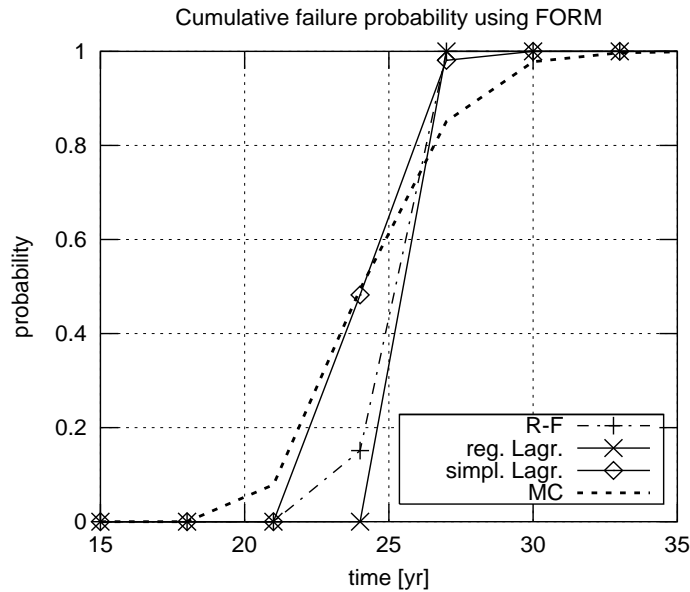


Figure 3.2: Comparison of different FORM implementations and Monte Carlo simulation for the thinning model.

the implementation in the program that is used. The method becomes less efficient as the matrix becomes larger.

The Lagrange multiplier method with the simplified objective function requires even less iterations than the regular Lagrange multiplier method. Also, this simplified version seems to follow the true solution much better. This conclusion is based on the fact that the result overlaps (more than the other two methods) the result obtained using simple Monte Carlo simulation. The accuracy of the simulation only depends on the number of samples which are used to evaluate the state of the component.

The last remark that can be made is that the two Lagrange multiplier methods result in design points which are closer to the limit state $g(\mathbf{y}) = 0$, than the Rackwitz–Fiessler method. The conclusion must be that the commonly used Rackwitz–Fiessler method is outperformed by the Lagrange multiplier method with the simplified objective function. This is a very useful result for anyone interested in applying an efficient method to perform a FORM.

3.3 MVFOSM

Before the FORM methodology was developed, there was the Mean Value First Order Second Moment (*MVFOSM*, also known as *FOSM*) method. It derives its name from the fact that it is based on a first-order Taylor series

approximation of the state function and uses only second-moment statistics (means and covariances) of the random variables. The original formulation by Cornell in 1969 [3] uses a simple two-variable approach. Assuming that R and S are independent normally distributed variables, a limit state function is defined as

$$g = R - S. \quad (3.11)$$

The mean of this function is then given by $\mu_g = \mu_R - \mu_S$ and variance is determined by $\sigma_g^2 = \sigma_R^2 + \sigma_S^2$. The probability of failure is then calculated as

$$F(0) = \Pr\{g < 0\} = \Phi\left(-\frac{\mu_g}{\sigma_g}\right).$$

This is based on the fact that the model output is normally distributed with mean μ_g and standard deviation σ_g :

$$\Pr\{g < 0\} = \Pr\left\{\frac{g - \mu_g}{\sigma_g} < \frac{0 - \mu_g}{\sigma_g}\right\} = \Phi\left(-\frac{\mu_g}{\sigma_g}\right)$$

Thus the probability of failure depends on the ratio of the mean value of g to its standard deviation. Cornell named this ratio the *safety index* and denoted it as $\beta = \mu_g/\sigma_g$.

In practice, the model will never be the simple linear model given by (3.11). Therefore the limit state function is linearized in the mean point $(\mu_{y_1}, \mu_{y_2}, \dots, \mu_{y_n})$ using Taylor series expansion as in (3.3). The mean of g is then approximated by

$$Eg(\mathbf{Y}) \approx g(\mu_{y_1}, \dots, \mu_{y_n}) + \sum_{i=1}^n \underbrace{(EY_i - \mu_{y_i})}_{=0} + H.O.T. \approx g(\mu_{y_1}, \dots, \mu_{y_n})$$

Again we assume that the variables Y_i are independent normally distributed random variables. The variance of g is approximated by

$$\sigma_g^2 = \sum_{i=1}^n \left(\frac{\partial g}{\partial Y_i}\right) \text{Var}(Y_i)$$

There are a number of arguments against using FOSM:

- the assumption that the model output is normally distributed with mean μ_g and standard deviation σ_g does not hold if the model is (highly) nonlinear.
- a nonlinear state function can also result in the approximation of the mean and variance by a Taylor series expansion having a high error associated with it.

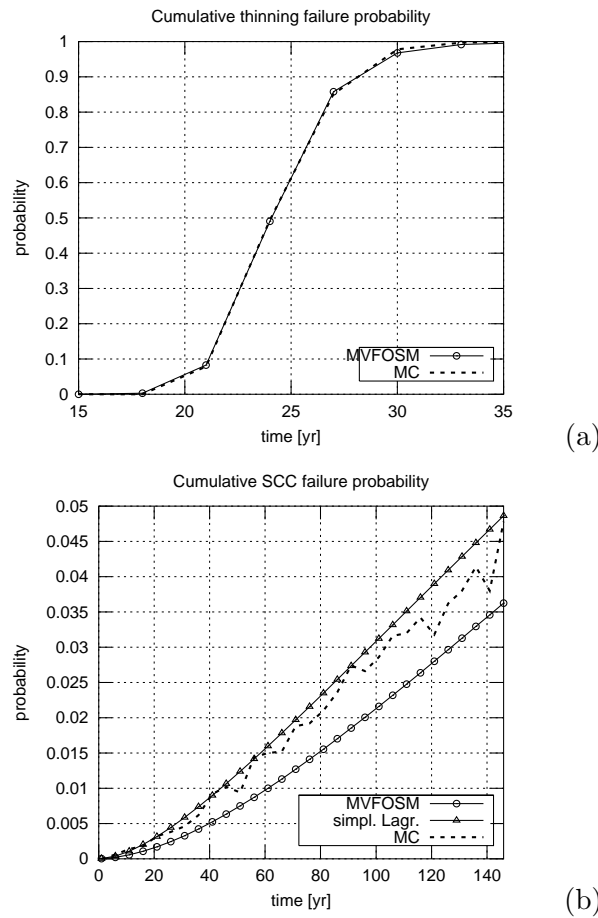


Figure 3.3: Results using MVFOSM for the thinning (a) and SCC (b) modules.

- the result of the FOSM method is not independent of the mathematical formulation of the problem (e.g. $R/S < 1$ is a different formulation of the same problem as $R - S < 0$, but the FOSM results can be different, see [22] for an example of this)

On the other hand, in the situation where the FOSM does give good results, no other method will be able to compete with the efficiency of the FOSM methodology. The evaluation of two simple equations is all that is needed to determine the reliability index β .

ORBIT™ uses MVFOSM to determine the reliability index. For the models (2.1) and (2.2) this works perfect, because the model output is normally distributed and the approximation of the mean and variance are very good. Figure 3.3 shows the results using MVFOSM for both models. The method works exceptionally well for the thinning module. For the SCC module the result is only slightly more progressive (i.e. lower failure probability)

than the simplified Lagrange multiplier method. The Monte Carlo simulation lacks enough samples to adequately approximate the failure probability. I've used 5000 samples per variable for both models, but the low probability of cracking for this type of steel means that the simulation will not be accurate enough. A different way of showing this would be to use the same sample set for each year in the figure, which would make the plot smooth. Unfortunately, the mean (and therefore also the standard deviation) of the crack depth variable A is a function of time, therefore this variable has to be sampled again every time. One advantage of this is that the amount of non-smoothness in the plot, reflects the lack of samples.

Another way of checking why the FOSM technique works well for these models is to overlay the simulation histogram with the normal density $f_X(x)$, where $X \sim \mathcal{N}(\mu_g, \sigma_g)$. This is shown in figure 3.4. As we look further into

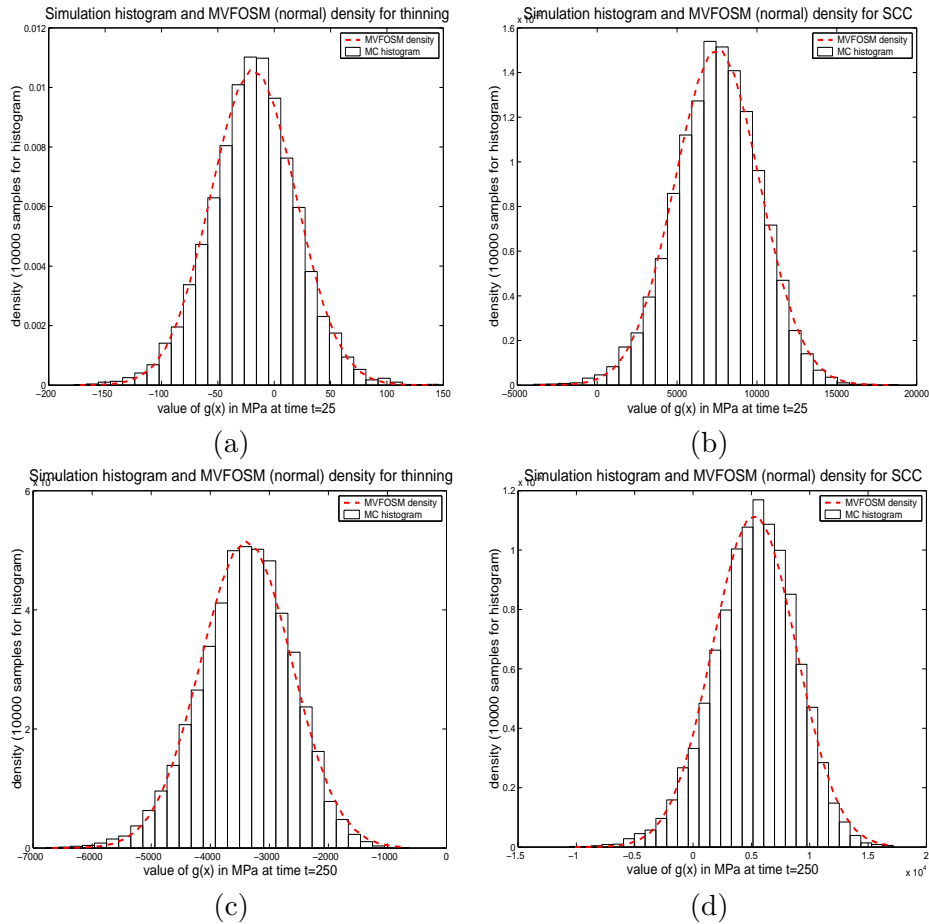


Figure 3.4: Monte Carlo simulation histogram (10000 samples) overlaid with the normal density determined using the MVFOSM technique for thinning and SCC at $t = 25yr$: (a) and (b) and the same for $t = 250yr$: (c) and (d) resp.

the future, the normal density seems to shift slightly to the left as can be seen in (c) and (d) of the figure. These are for $t = 250\text{yr}$, which is not realistic, but it helps us understand the behavior of the models. Taking this into account, we would expect the MVFOSM result in figure 3.3b to be more conservative than the simulation or FORM results. However, this figure suggests the opposite: the density should be shifting to the right. The explanation for this is the lower tail of the histogram compared to the lower tail of the MVFOSM density. There is clearly more mass in the lower tail of the histogram in figure 3.4c and 3.4d than in the equivalent tail of the MVFOSM density. This has a significant effect in the SCC module due to the low probabilities that we encounter here. This holds much less for the thinning model, but these conclusions are of course dependent on the example which we use. A different type of steel might be much more susceptible for cracking than the material used in this example.

3.4 Directional sampling

Simulation is technically not a reliability index method, but a sampling technique to approximate the failure probability. A regular *Monte Carlo simulation* will consist of taking many sets of samples from the input variables and to count the number of times this results in a failure (i.e. $\#\{i : g(\mathbf{x}^i) < 0\}$). If we take N samples of each variable, then the failure probability will be approximately the number of failures divided by N . This technique has the advantage that the quality of the approximation depends only on the size of N , therefore we can make the approximation as good as we like by taking $N \rightarrow \infty$. Unfortunately, this is also the disadvantage of the technique, because it makes it highly inefficient. For large N , the calculation is computationally very expensive. Especially for low probability events, a very large number of simulations is required to ensure that enough samples are located in the failure region. A nice example which shows this effect is figure 3.3b, where we can clearly see that there is a lack of samples.

Nonetheless, the use of simulation is still a very attractive method to approximate the failure probability. This is due to the fact that any distribution can be used for the input variables, which is not possible when using a regular FORM or MVFOSM technique. These methods require the user to transform the input variables to normally distributed variables. Also, the time which is required to perform a simulation has been dramatically reduced with the current processor speeds in personal computers.

A relatively new simulation technique which has gained much popularity over the last few years is called *directional simulation*. It's based on the idea of transforming the Cartesian space in which we usually look at the state function g , to polar coordinates. In short the technique involves sampling a direction on the unit hypersphere and integrating over the failure region

along this direction. When repeating this procedure many times, the failure probability will be approximated by taking the average of these probabilities.




Figure 3.5: Graphical representation of the directional sampling technique.

A single direction is shown in figure 3.5, where the directional integration starts at the point given by $r^*(\theta)$. This is the radius along the angle θ for which $g = 0$. In this case the variables in the state function have been transformed to standard normal variables Y_1, \dots, Y_n as in equation (3.2). Based on the chapter 9 of Monte Carlo Methods in Ditlevsen & Madsen [4], we can determine the failure probability along a direction by first sampling a unit direction vector \mathbf{A} , where

$$A_i = \frac{Z_i}{\sqrt{\sum_{i=1}^n Z_i^2}} \text{ with } Z_i \sim \mathcal{N}(0, 1) \text{ for } i = 1, \dots, n. \quad (3.12)$$

From this random unit vector, we sample a direction \mathbf{a} and calculate the radius for which $g(r, \mathbf{a}) = 0$. The probability that the directional radius R is larger than $r^*(\mathbf{a})$ is given by

$$\Pr\{R > r^*(\mathbf{a}) | \mathbf{A} = \mathbf{a}\} = \frac{\Gamma\left(\frac{n}{2}, \frac{[r^*(\mathbf{a})]^2}{2}\right)}{\Gamma\left(\frac{n}{2}\right)}, \quad (3.13)$$

where the *gamma function* and *incomplete gamma function* are defined as

$$\Gamma(a) = \int_0^{\infty} t^{a-1} e^{-t} dt \quad (3.14)$$

and

$$\Gamma(a, x) = \int_x^{\infty} t^{a-1} e^{-t} dt \quad (3.15)$$

respectively. The actual failure probability is then approximated by taking the average of the directional probabilities obtained by repeating the sampling and calculation procedure given by equations (3.12) and (3.13):

$$\frac{\sum_{j=1}^N \Pr\{R > r^*(\mathbf{a}^{(j)}) | \mathbf{A} = \mathbf{a}^{(j)}\}}{N} \longrightarrow \Pr\{g < 0\} \text{ as } N \longrightarrow \infty$$

The directional sampling method is more efficient than the regular Monte Carlo simulation method, because if there is a finite solution for the radius (i.e. $r^*(\mathbf{a}) < \infty$), then the integration will be done over all the mass in the failure region. This ensures that the average failure probability will converge to the actual probability much faster compared to the regular sampling technique. Most literature will also claim that the technique is more efficient due to the fact that we need to sample only $n-1$ angles instead of n variables. This does not hold in our example because sampling the angles such that the directions are uniformly distributed on the unit hypersphere is quite difficult (see appendix A.2 for more on this). This is the reason why we are using the unit vector defined by (3.12), which requires us to sample 3 standard normal variables Z_i . This increased efficiency therefore does not hold for the multivariate normal distribution, but will hold for other multivariate distributions.

Most authors underestimate the importance of determining the radius to the limit state function: $r^*(\mathbf{a})$. When considering complex models, this radius will usually have to be found numerically. This requires careful thought in order to keep the method from becoming inefficient. Also, there will always be multiple radii for which $g(r, \mathbf{a}) = 0$. For example, if we solve for r in the thinning module, then the result is a quadratic equation which means that there will (almost always) be 2 solutions for r . The question is now which radius to use in (3.13): should we take the smallest radius and what do we do with a negative radius? Melchers [12] acknowledges the problems that arise with multiple solutions for r , but only advises the reader to “carefully formulate the state function g .” Using the thinning model as an example, I will formulate some rules which can be applied to obtain a correct result for the radius. The implementation of the directional sampling technique for the thinning module can be found in appendix A.2.

The most appealing property of the thinning model which was introduced in (2.1) is that there are three input variables, which means that we can visualize the result in a 3-dimensional space. A pretty rendering of the directional sampling technique is given in figure 3.6. The dotted lines represent the sampled directions and the points are the radii on the failure surface. The thick cross indicates the center of the standard normal space. Unfortunately this figure does not help us any bit in determining which radii we have to use for calculating the failure probability. Figure 3.7 is a top-down view onto the XY-plane. These are the material strength (S) and corrosion rate (C) axes respectively, therefore we are looking down




Figure 3.6: 3D rendering of the directional sampling technique.

along the pressure (P) axis. The shape of the failure surface becomes very clear in this figure and it is not immediately what would be expected. Given the good results obtained using the MVFOSM technique, we concluded that the failure surface (close to) the center point is practically linear. The cloud formed by the sampled input variables is separated in two parts by the failure surface and it is indeed highly linear in this area. However, the directional sampling reveals much more of the failure surface. The reason for this is that this technique will also result in points $(x, y, z) = (rA_1, rA_2, rA_3)$ which are infeasible. Transforming these points back to the original variable space results in a negative value for one or more of the three input variables: material strength, corrosion rate and pressure. This is physically not possible, therefore these points should not be taken into account. These points are marked as an 'x' in figure 3.7 and the feasible points are marked with a solid dot. Also plotted are the feasibility bounds which are found by assuming

$$X_i \geq 0 \Leftrightarrow \sigma_i(rA_i) + \mu_i \geq 0 \Leftrightarrow rA_i \geq -\mu_i/\sigma_i.$$

To be able to obtain the correct failure probability using directional sampling, only the feasible radii should be used and if there are multiple feasible radii, then only the shortest (in absolute value) one should be used. This means that also the negative radii have to be used. The probability of failure (3.13) uses the square of the radius, therefore this probability does not depend on the direction of the radius. Equally important is to determine whether the center of the standard normal space is in the failure region or not, because we need to make sure that we are integrating over the failure region and not the safe region. The decision tree in figure 3.8 summarizes the steps which determine the probability of failure P_j for a single sampled direction j , where $j = 1, \dots, N$. In this figure $P(R > r)$ is short for (3.13).

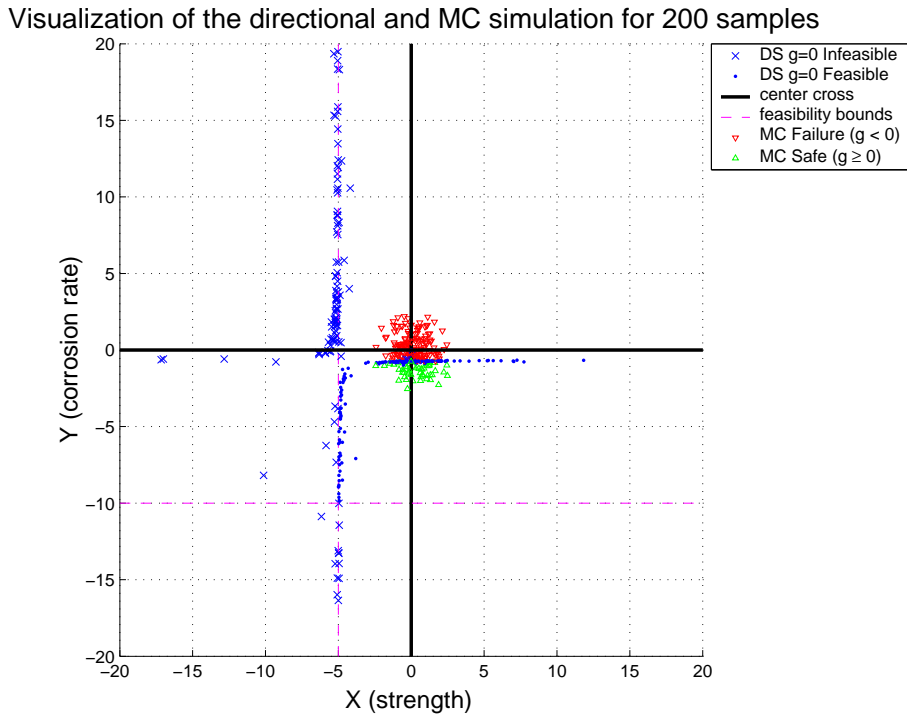


Figure 3.7: Top-down view onto the plane defined by the material strength (S) and the corrosion rate (C) including the rendering of the regular Monte Carlo simulation.

This probability is divided by 2, because the probability in only one direction is $\{R > 0 | \mathbf{A} = \mathbf{a}\} = 1$. Using this decision tree, we can compare the convergence of the directional sampling compared to the regular Monte Carlo sampling technique. This is shown in figure 3.9. The plot of the directional sampling failure probability stabilizes around the actual failure probability (obtained using the MVFOSM) faster than the regular sampling method. Also included are the 95% confidence bounds for both results, but these do not give any additional information about the difference in effectiveness of each method. We conclude that the directional sampling method indeed requires much less samples to effectively approximate the failure probability.

A much more interesting application of the directional sampling technique is the stress corrosion cracking model. From figure 3.3b we concluded that the MVFOSM underestimated the failure probability. Also, the regular Monte Carlo simulation requires a very large number of samples due to the low probability of failure. We expect the directional sampling to have better results. The model, which is given by (2.2), is interesting because we now have to solve for r^* numerically. This is mainly due to the crack depth parameter A which is under a square root. Again the question arises how

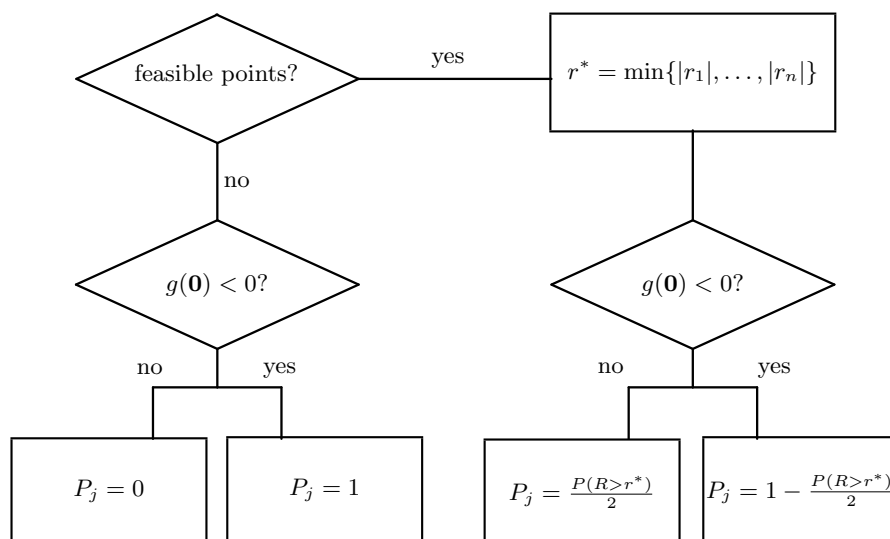


Figure 3.8: Decision tree for choosing the radius in the directional sampling technique and calculating the correct failure probability.

to choose the correct radius such that the result is also representative of the model. The details are given in appendix A.2. Here we use the Newton-Raphson iterative scheme to solve for the radius r^* for which $g(r^*, \mathbf{a}) = 0$. Assuming that we always want the closest radius, we take a zero radius ($r_0 = 0$) to start the iteration. The scheme will then proceed in the positive or negative direction until either the solution has been found or the scheme does not result in a feasible radius. The last case occurs when the scheme diverges or as soon as the radius becomes associated with an infeasible value for one of the variables. Besides there being only one feasible radius or none, the decision tree for this model is otherwise identical to the one in figure 3.8. Using 2000 samples for both simulation techniques, we get the result in figure 3.10. Again the regular Monte Carlo simulation suffers from a lack of samples, but the directional sampling technique performs slightly better. It's result is more stable and follows the FORM results more closely. Since the FORM is considered to be a good approximation of the actual failure probability, we conclude that the directional sampling still underestimates the failure probability albeit much less than the MVFOSM. This may be due to the fact that we reject an infeasible radius and subsequently assume $P_j = 0$ (or $P_j = 1$ depending on the value of $g(\mathbf{0})$), whereas there might be a feasible radius in the other direction. A smarter algorithm could result in slightly better results, but the directional sampling performs satisfactory in figure 3.10. Only the last point does not look good, but since this model has 4 directions to be sampled, we will need more samples compared to

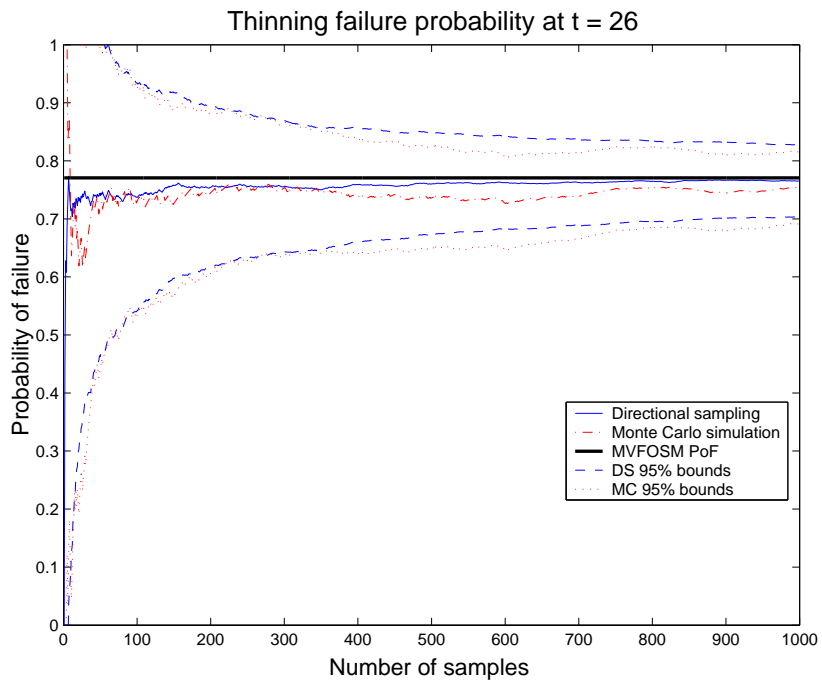
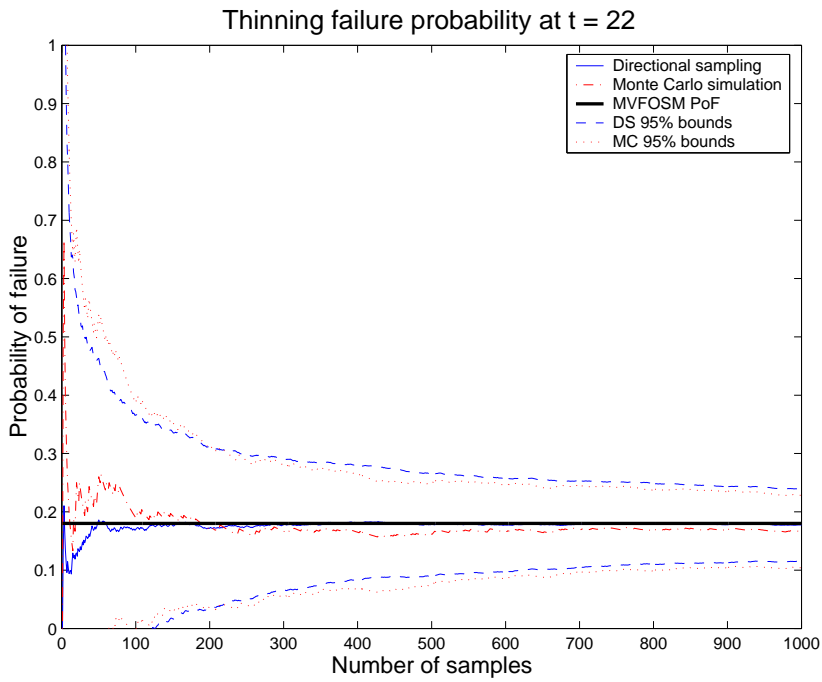


Figure 3.9: Convergence of the directional sampling technique compared to the regular Monte Carlo technique for the thinning model at 22 years (a) and 26 years (b).

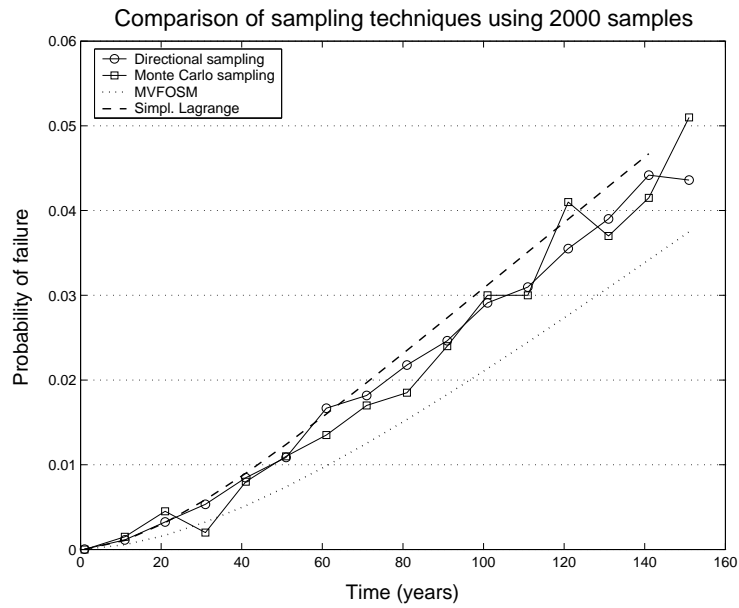


Figure 3.10: Comparison of the results for the cracking model failure probability using both reliability index and simulation techniques.

the thinning model to achieve similar accuracy. The more samples that are used, the smoother the plot will become. Also, the result is excellent until about 50 years into the service life of the component (which certainly can not be said about the regular sampling method). In the process and refining industry, components rarely stay in operation longer than such a length of time, therefore the performance during later years is less important.

Chapter 4

A new Gamma process inspection model

This chapter is entirely devoted to the development of a new decision model for optimal inspection planning in the process industry. The fundamental difference with the previous (deterministic) models is that the deterioration will now be modelled by a stochastic process. The idea is to use a Bayesian stochastic process, which is a more mathematically sound approach to the planning of maintenance and inspections. Unlike the deterministic model, we can not predict the exact behavior of the process in the future. With the deterministic model we determine a posterior density for the degradation rate and assume a linear degradation in time. This essentially fixes the point at which the component is assumed to fail, albeit with some uncertainty in the form of variance associated with it. A regression analysis, where a linear degradation fitted to measurement data is extrapolated into the future, has this characteristic as well. This statistical method has another problem: due to the lack of sufficient data, the quality of the regression is not good. Believe it or not, but a regression analysis methodology called KINT, has already been accepted by the technical committee as a valid risk based tool for inspection planning. This method is only valid for thinning due to corrosion, whereas ORBIT™ also considers cracking and fatigue.

The methods in this chapter are based on the techniques which have been developed and studied by J.M. van Noortwijk and have been successfully applied to hydraulic structures [15] (e.g. maintenance of dikes, sea-bed protection at a sea barrier, etc.). The principles on which the use of these methods are based, are such that the methods can be translated to the process industry.

First a concept is drawn up for the model, which includes some characteristics which we would like to see in a new model. Next, the gamma deterioration process is introduced, followed by a continuous Bayesian updating model which can handle both perfect as well as imperfect inspections.

In the last part of this chapter I will consider the expected average costs per unit time as a cost-based criterium for deciding on an optimal inspection and maintenance policy. The model as a whole will then be illustrated with a case study of a hydrogen dryer.

4.1 Model concept

Before we start working on a new model, we need to think about which characteristics this model should have such that it will be detailed, but also easy to use at the same time. We start by summing up some of the advantages of ORBIT™ which we certainly want to see again in the new model:

- **Easy and fast in use:** this is most important for the input of the model. The input should only consist of the standard operating conditions (e.g. average corrosion, pressure, temperature etc.) and material properties, since the model is used by practitioners in the process industry with a background in chemical engineering, civil engineering, mining etc. We should not ask these people for distribution parameters or uncertainty bounds, because this will make them feel uncomfortable due to their limited experience with probability theory and statistics. Also, this type of input requires a substantial amount of time in both preparation and actually typing it into the program. This is not an appealing prospect when dealing with thousands of components.
- **Model the degradation and update with available measurement data:** the advantage of ORBIT™ is that it combines the modelling of the degradation with measurements from inspections. It is therefore not only a statistical model, which only uses the inspection data to extrapolate the material condition into the future. The new model should also use state functions to model the condition of the construction material and update this model with the available inspection results using Bayesian updating. The use of Bayes' theorem is very common in structural reliability analysis and condition assessment. Pandey [18] applies it to update the prior with a probability of detection, Mahadevan & Zhang [11] incorporate the uncertainty over the model itself and Zhao *et al.* [26] apply it to the inspection of steel bridges.
- **Use the thinning and cracking state functions:** these models have been studied and applied in real-life applications for a very long time. They are accepted as valid models and they use typical operational conditions and material properties to model the degradation. Practitioners like to see these variables in the model as it gives them

confidence that all parameters which might influence the degradation are incorporated in the calculations.

Besides these desirable characteristics, we would also like to see the following aspects in the new model:

- **Use a stochastic model instead of a deterministic model:** we would like to use a stochastic process to model the cumulative damage to the material of the component. This in contrast to assuming an average degradation which is extrapolated to find the moment of failure supplemented with uncertainty bounds.
- **Choose a proper prior for the average degradation:** we would like to step away from the normal distribution which is commonly assigned to all variables in traditional structural reliability analysis. The normal distribution entails a (very small) probability of a negative value for the variable. Also, the normal density is symmetric, whereas a density with a longer tail towards higher values for the degradation rate would represent our intuition and uncertainty better. The goal is therefore to use a non-negative density with a long tail.
- **Incorporate the uncertainty created by imperfect inspections:** we define an *imperfect inspection* as an inspection technique which results in a measurement with a small error. We therefore assume that the measurement will give the actual state (e.g. wall thickness) of the material plus or minus a small error. We can model this error by assuming that the error is normally distributed with mean zero and a standard deviation which reflects the accuracy of the inspection method. The error does not account for other inaccuracies like the choice of the measurement location on the component, the (in)experience of the inspection engineer etc.

There has been much research done on the topic of inspections. In the process industry we come across a number of so-called non-destructive testing (NDT) techniques for evaluating the condition of the construction material. None of these methods are capable of perfectly identifying and measuring the damage or wall loss in the construction material of a component. Each inspection result will have an error associated with it, which represents the amount of uncertainty in the measurement. Following Ellingwood & Zheng [5], there are two types of uncertainty which are of primary concern for reliability based in-service inspection. These are the uncertainty in the flaw detection and the uncertainty in the flaw sizing. In the case of thinning, the probability of detection is of lesser concern, therefore we will not consider this in our model.

margin m , which occurs at time $t = T$ in figure 4.1. Later on, we will also consider the possibility of replacing the component before failure occurs. For this purpose we introduce $\rho \in (0, 1)$, which represents the percentage of the margin at which we decide to replace the component. Therefore we replace the item if the amount of thinning has reached the level ρm (this level is also known as the corrosion allowance). This is called a *preventive replacement*, whereas a replacement due to failure is referred to as a *corrective replacement*.

As is done in current practices, the deterioration is assumed to be linear in time. Both the mean value of the gamma process $X(t)$ as well as the variance are linear in time:

$$\mathbb{E}(X(t)) = \mu t \quad \text{and} \quad \text{Var}(X(t)) = \sigma^2 t. \quad (4.1)$$

This entails that the probability density function for $X(t)$ is given by

$$f_{X(t)}(x) = \text{Ga} \left(x \mid \frac{\mu^2 t}{\sigma^2}, \frac{\mu}{\sigma^2} \right). \quad (4.2)$$

A proof of the above properties, a more detailed definition of the gamma process and the definition of the gamma density (A.4) can be found in section A.3.1 of the appendix.

For the remainder of this chapter I will only consider the corrosion model. For this I define the wall loss as the corrosion rate times the time since service start of the component: $X(t) \equiv C \times t$. The margin m is the level of $X(t)$ at which $g = 0$. By rewriting the state function (2.1) we find that

$$m \equiv m(p, s) = th - \frac{p \times d}{2s}. \quad (4.3)$$

The margin therefore depends on the operating pressure p and the material strength s . We can take these as constant variables, but later on we will also assume that they can have some uncertainty or variability associated with them. The margin can therefore be uncertain, which we will have to incorporate into the model. In a plant, the operating pressure is usually controlled and most of the uncertainty will be in the material strength. This is because the material properties have been obtained under laboratory conditions and these are not always representative of the surroundings and conditions in which the component operates.

The deterioration process is determined by the average corrosion per year μ and the standard deviation in the corrosion rate σ . Since we only want to incorporate the uncertainty over the average corrosion rate, we fix the standard deviation with a coefficient of variation. In the chapter on ORBIT™ we've seen that the standard deviation is fixed in the model by a coefficient of variation. For example, the standard deviation for the

corrosion rate is 10% of the average corrosion rate per year: $\sigma_C = 0.1\mu_C$. This essentially fixes the ratio σ/μ and we define

$$\nu \equiv \frac{\sigma}{\mu},$$

which implicates that (4.2) becomes

$$f_{X(t)}(x) = \text{Ga}\left(x \mid \frac{t}{\nu^2}, \frac{1}{\mu\nu^2}\right). \quad (4.4)$$

Since the coefficient of variation ν is fixed, the only uncertain variable left in the model is μ (besides of course the pressure and material strength if we choose to make these random as well). This greatly reduces the complexity of the calculations and avoids the need for extra input compared to the ORBIT™ input. The choice for ν should be reconsidered as its role in the gamma deterioration process is quite different from its role in the MVFOSM approximation. Expert judgment can be used to determine suitable choices for this coefficient of variation. As different types of degradation mechanisms have different uncertainty associated with the (some deterioration is more predictable than others) it would make sense to agree on suitable ν 's for each of these mechanisms.

4.3 Inspection updating

First we look at the case of perfect inspections. This means that we are able to exactly determine the wall thickness with the inspection. As with ORBIT™ we want to use a Bayesian updating model to incorporate the inspection data into our distribution over the degradation rate μ , only now we want to use a continuous version instead of the simple discrete variant of Bayes' theorem. The choice for a prior density is not easy, because we need to make sure that the posterior is not too hard to determine. Following the suggestion by van Noortwijk *et al.* [14], we will use the family of inverted gamma distributions. The inverted gamma density is given by

$$\text{Ig}(x|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} \left(\frac{1}{x}\right)^{\alpha+1} \exp\left\{-\frac{\beta}{x}\right\} \quad \text{for } x \geq 0. \quad (4.5)$$

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 $\beta > 0$. An inverted gamma density for the prior fits very well with the application to the corrosion rate. The density is non-negative and has a longer tail, which represents our uncertainty over higher corrosion rates very well. The inverted gamma density can be fitted to the ORBIT™ prior as is done in figure 4.2. This is not an accurate fit, but this is not really necessary either. The inverted gamma density has been fitted to the discrete prior given in

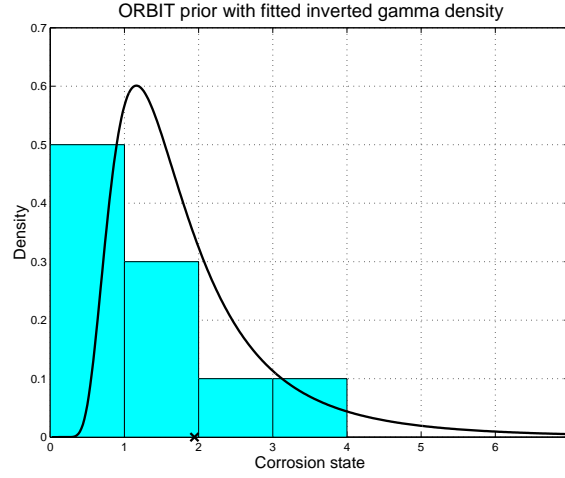


Figure 4.2: The inverted gamma density fitted to the ORBIT™ discrete prior.

table 2.3 through the use of simulation (see also the program code in section B.3 of the appendix). The main purpose is to have a prior which is well suited to our needs. At first you might think that an exponential density would be a better fit, but this would entail a high probability of a corrosion rate lower than the user's estimate (unless of course μ_C is the mean of the exponential distribution). The fact that most mass is concentrated around 1 time the estimated corrosion rate fits better with the real world. The expectation of both the discrete prior and continuous prior is around 1.9 times the estimated rate (indicated by a cross in the figure).

Say we measure the wall thickness at time t , then we know the amount of wall loss x which has occurred since the service start of the component. The likelihood of $X(t)$ given this measurement is given by the gamma density (4.2). The continuous posterior is found by calculating

$$\pi(\mu|x) = \frac{l(x|\mu)\pi(\mu)}{\int_{\mu=0}^{\infty} l(x|\mu)\pi(\mu)d\mu}, \quad (4.6)$$

where $\pi(\mu)$ is the prior density for μ and $l(x|\mu)$ is the likelihood of x given μ . The advantage of using the inverted gamma prior now comes in the fact that the posterior is given by

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

The proof of this can be found in appendix A.3.2.

As an example, I will show the effect of performing a perfect inspection after 5 years which reveals a wall loss of 5mm when the initial estimate for μ was 1mm/yr. This is therefore a confirming result for which the posterior

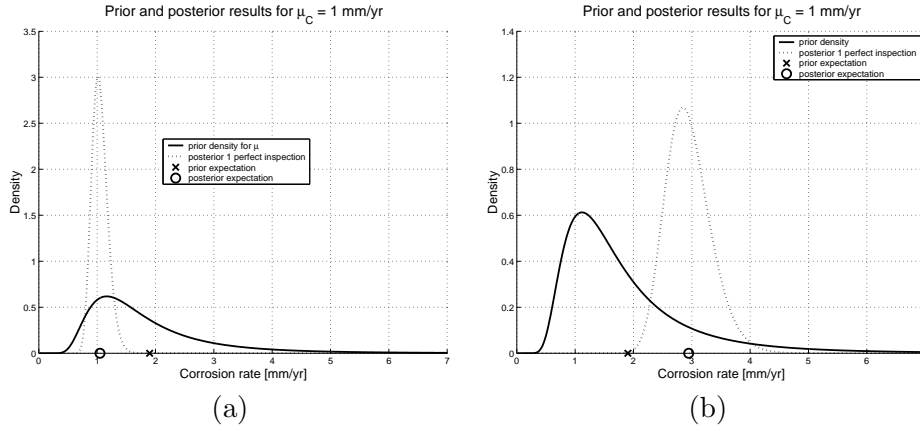


Figure 4.3: Comparison between the posterior of a confirming perfect inspection and a perfect inspection resulting in a corrosion 3 times faster than the estimate.

is shown in figure 4.3a. This inspection has clearly increased our confidence and decreased the uncertainty over the corrosion rate by concentrating all mass around the value of 1mm/yr. If, on the other hand, the inspection would have told us that the material has corroded three times faster than we had anticipated, the posterior would move towards this value. Also, because this is a contradicting result, the mass is more spread out like we can observe in figure 4.3b.

We can easily extend this model to incorporate the data from multiple perfect inspections. Say we measure the wall loss at n points in time, then we have a measurement x_1 at time t_1 , x_2 at t_2 etc. As the proof in section A.3.2 shows, the posterior for the degradation rate given n measurements is given by (A.10):

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

Again we have a simple analytical solution for the posterior density, but when we look closely at the above equation we see that the posterior actually only depends on the last inspection:

$$\pi(\mu|x_1, \dots, x_n) = \pi(\mu|x_n) = \text{Ig} \left(\mu \left| \frac{t_n}{\nu^2} + \alpha, \frac{x_n}{\nu^2} + \beta \right. \right),$$

if we assume that $x_0 = 0$ at $t_0 = 0$. This rather surprising result is a direct consequence of our choice to fix the ratio σ/μ . Besides the fact that perfect inspections do not exist, it will not be possible to convince a plant operator or safety technician that all but the last of his inspections are irrelevant in determining the failure probability of the component.

It is clear that for any model to be valuable in the process industry, it will have to include the possibility of modelling imperfect inspections. The first model which comes to mind is the stochastic process $Y(t)$ which encompasses the original gamma process $X(t)$ together with an extra factor which represents the error in the inspection measurement(s). This approach is also used by Newby & Dagg [13] and has the following form:

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

therefore we assume that an inspection will measure a wall loss y at time t , which is the actual wall loss x plus or minus a measurement error which is embodied by a standard deviation σ_ϵ . We assume that the mean error is zero, because any other choice would assume that the inspection is more likely to over- or underestimate the actual wall thickness.

The article from Newby & Dagg [13] is the only known attempt at implementing imperfect inspections together with a gamma deterioration process. Unfortunately, the paper only results in a complicated numerical integration and the output is presented in a non-intuitive manner. Sticking to our Bayesian model and avoiding complicated numerical integration schemes by applying a discretization, we can simplify and expand the model proposed in the previously mentioned article. The discretization is necessary, because (with the current model) we can not solve the problem analytically. Since the mathematical details and the required notation is relatively complex, I've put these in section A.3.3 in the appendix. I have successfully determined a (discrete) posterior density for one or more imperfect inspections, which are given by equations (A.14) and (A.17) respectively.

Figure 4.4 shows a comparison of the posterior densities for 1 imperfect inspection at 5 years and the posterior of two imperfect inspections at 5 and 15 years against the posterior for one perfect inspection at 5 years. All imperfect inspections have a standard deviation of 50% of the measured wall loss, therefore $\sigma_\epsilon = 0.5y$, where y is the measured wall loss. The error of an inspection technique is usually given as a percentage of the measured wall thickness, but if we use this in our implementation, then there is a large probability that the measurement error is (much) larger than the actual wall loss. It is therefore better to define the error as a percentage of the wall loss itself. An error of 50% of the measured wall loss seems very large and the choice for this value is purely demonstrational. As in the case of choosing a sensible value for ν , this standard deviation should be determined using some kind of expert opinion.

The figure confirms our expectations of the model. The imperfect inspection is much less effective at decreasing the uncertainty over the corrosion rate compared to the perfect inspection. An extra imperfect inspection ten years after the first one increases the confidence a little bit and shifts the expectation slightly towards the estimated rate of 1mm/yr.

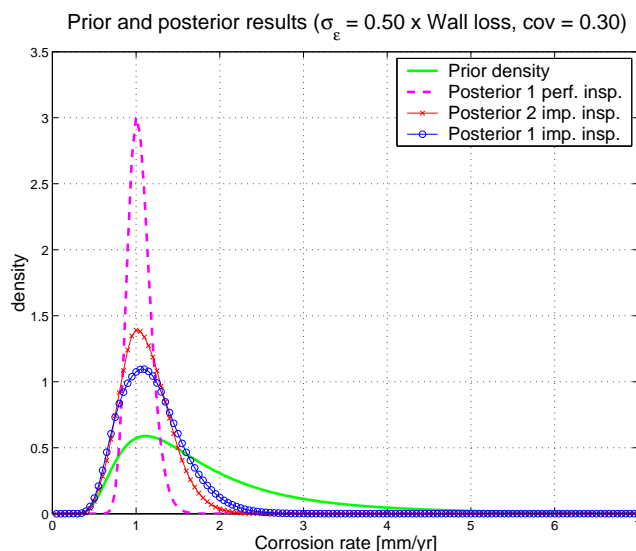


Figure 4.4: Posterior densities for 1 imperfect inspection at 5 years and for 2 imperfect inspections at 5 and 15 years, each confirming a corrosion rate of 1mm/yr and $\sigma_\epsilon = 0.5\mu_C$.

4.4 Optimal inspection and maintenance planning

DNV's ORBIT™ methodology suggests a type and a time of inspection such that a certain amount of risk is not surpassed. Theoretically this means that we could inspect the component at infinitely small time intervals, because this would minimize the risk of unexpected failure. This plan is however extremely expensive, certainly when the inspections are relatively expensive due to plant outage and other similar extra costs. We would therefore like to use a decision model which will give us an optimal inspection period such that safety is ensured and the costs are minimized.

The proper model for this purpose is the concept of the expected average costs per unit time. This is one of three cost-based criteria suggested by Wagner [25] and applied to a number of examples in the maintenance of hydraulic structures by van Noortwijk [15]. The expected average costs per unit time are obtained by simply averaging the costs over an unbounded horizon. This means that we will add all costs related to an inspection interval Δk and divide these costs by the expected duration of the life-cycle of the component. From a renewal theory point of view, the cycle of a component is the time until it is either preventively replaced or correctively replaced. The first is done when the condition of the construction material has degraded beyond an accepted level. In the process industry this level is commonly referred to as the *corrosion allowance*. A corrective replacement is performed when the component has failed.

As is also discussed in section A.4 in the appendix, the corrosion allowance is an amount of deterioration which is lower than the safety margin given by equation (4.3). In our model we will use a percentage ρ which represents the ratio c_{max}/m , where c_{max} is the corrosion allowance. We therefore have a corrosion allowance ρm , which is a percentage of the level m at which the component is assumed to fail. If the degradation is more than the corrosion allowance, then we replace the component, therefore we also call ρm the replacement level.

The expected average costs per unit time are given by the ratio of the expected costs per cycle over the expected cycle length:

$$C(\rho, \Delta k) = \frac{\sum_{i=1}^{\infty} c_i(\rho, \Delta k) p_i(\rho, \Delta k)}{\sum_{i=1}^{\infty} i p_i(\rho, \Delta k)}, \quad (4.7)$$

where c_i are all costs incurred during cycle i and p_i is the probability of incurring these costs during cycle i . This equation is a result of renewal theory, where a *cycle* is defined as the in-service time of a component. One cycle is therefore the duration between the service start and the time of replacement or failure, whichever occurs first. The use of this theory requires the assumption that the component will be renewed, i.e. it will be replaced by an identical component. For components with a short design life, this will be mostly the case and for the other components this assumption will not result in incorrect results.

The expected average costs per unit time is a function of the interval length Δk and the percentage ρ , which is determined by the corrosion allowance. For this model we will assume that the corrosion allowance is always given and that we will not use any other preventive replacement level, therefore the percentage ρ is fixed. In other cases we can consider ρ as another variable over which we can optimize the inspection plan. A certain choice for ρ might entail substantially lower inspection and maintenance costs. This option is interesting for the user who would also like to determine the optimal amount of degradation after which the component should be replaced.

The costs per cycle c_i include costs for inspections, preventive replacement and for failure (including corrective replacement). The probability of failure per cycle p_i depends on the degradation process, therefore it not only depends on the cycle length Δk , but also on the operating pressure, material strength and the corrosion rate.

To find the optimal inspection interval we simply compute (4.7) for different values of Δk . For example, if we use a time unit of one year, we can start with $\Delta k = 1\text{yr}$ and continue until the interval length is such that the component is sure to have failed during the interval. Besides the uncertainty over the corrosion rate, for which we determined the posterior density in the previous section, we would also like to consider the operating pressure and the material strength as uncertain variables. Since the calculations require

quite a bit of attention, they have been placed in section A.4 of the appendix. For the rest of this chapter I will restrict myself to the case study in the following section to illustrate the concept.

4.5 Case study: inspecting a hydrogen dryer

For this case study we use the data of a vertical drum which is used as a hydrogen dryer by one of the customers of DNV in the Netherlands. All required data¹ is given in table 4.1. From the inspection history, one can

Component type:	vertical drum	
Material type:	carbon steel	
Service start:	1977	
Tensile strength:	413.69	MPa
Yield strength:	206.84	MPa
Operating pressure:	32	bar(g)
Drum diameter:	1180	mm
Material thickness:	15	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 4.1: Operational, material and inspection data for a hydrogen dryer.

clearly observe that the drum has been inspected every 4 years according to the fixed regime prescribed by the Dutch Rules for Pressure Vessels [21].

The second observation concerning the inspection history is that the measurements are not consistent. In 1986, they measured a material thickness larger than the specified thickness given by the constructor of the component. It is actually very common that a thickness inspection results in a measurement which is larger than the previous measurement. This is jokingly referred to as ‘material growth’ by the inspection engineers. This phenomenon is due to the fast increase in accuracy of the inspection techniques

¹the costs in this table are fictive; all other data is actual plant data.

during the last decades. Usually, inspection results which were obtained more than 12 years ago are disregarded as being too old and unreliable. Also, the nominal material thickness supplied with the product by the constructor is not the exact thickness of the material. The actual thickness at service start can be up to 12% larger than specified. DNV currently advises all its customer to perform a null-measurement at service start to determine how much thicker the material is than the specification indicates. This helps avoid confusion about the corrosion rate later on in the life of the component. With all this in mind, we will disregard the measurement done in 1982 and assume that the initial wall thickness was $15 \times 1.12 = 16.8\text{mm}$. This is of course a very subjective choice, but it is also very conservative, which is key in this business. It might not seem to be a conservative choice when you assume a much higher material thickness, but it directly results in a very high corrosion rate. Since we take the last measurement as the actual wall thickness (including a small error), the initial thickness is irrelevant.

4.5.1 Without uncertainty

Before the final result is presented, we first take a look at the case where there is no uncertainty. This means that we assume that the pressure, material strength and the corrosion rate are exactly as given in table 4.1. The only variability is that of the gamma process itself, which is determined by the coefficient of variation ν . Table 4.2 summarizes some of the key values for this problem. The individual elements of the result are depicted in figure

time unit:	1	yr
coefficient of variation ν :	0.25	-
safety margin m :	11.268	mm
replacement percentage ρ :	0.3994	%
expected time to replacement:	45.5	yr
expected time to failure:	113.2	yr
optimal inspection interval Δk :	104	yr
cost per year for $\Delta k = 104$:	581	\$

Table 4.2: Results for the expected average costs per year in the case of no uncertainty.

4.5. The expected average costs per year are given by figure 4.5c and this is the ratio of the expected costs per cycle (figure 4.5b) over the expected cycle length (figure 4.5a). From the figure in (c) we obtain the optimal inspection interval $\Delta k = 104$ years, which results in a cost of \$581 per year. This represents a total cost of $104 \times \$581 = \$60,424$, which can be split up in one inspection of \$10,000 and one preventive replacement of \$50,000. Since the drum is currently 25 years old, we would have to plan the next inspection

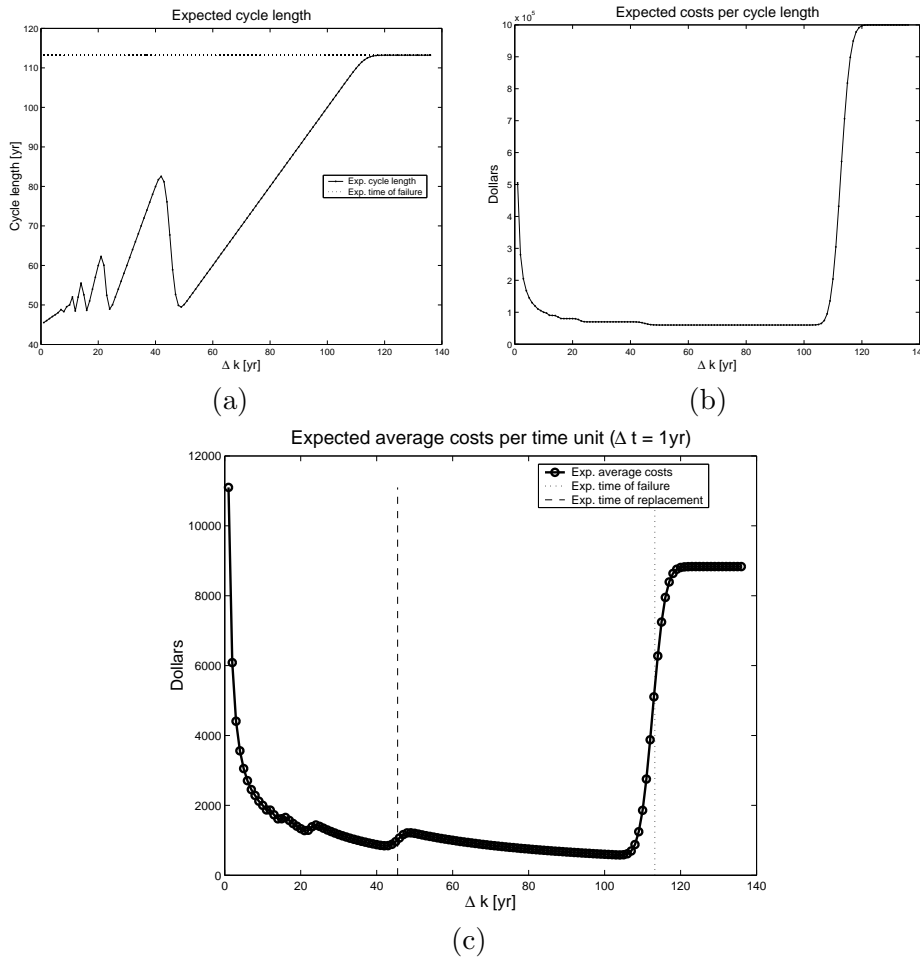


Figure 4.5: The expected cycle length (a), the expected costs per cycle (b) and the expected average costs per year (c) for the case with no uncertainty.

in the year 2081. Nothing is more unrealistic than this, but this is of course due to the fact that we have assumed all variables to be known. This enables us to plan the inspection right before the expected time of failure, when we know that we'll have to replace it. The expected replacement time and the expected time of failure are represented by vertical lines in figure 4.5c.

The interpretation of the expected costs per cycle in figure 4.5b is very simple: inspecting the component too often results in a very high cost, which is also the case when we plan the inspection too late. The first figure 4.5a requires a little more explanation. It is obvious that the expected cycle length should in the end converge to the expected time of failure for the component. The oscillating behavior of the graph is due to the fact that certain inspection intervals will 'land' right after the replacement time (low

dips) and some will end up closer to the failure time (high peaks). For example, if we choose $\Delta k = 40$ years, then we will replace the component at the second inspection, therefore the cycle length will be 80 years. On the other hand, if we choose $\Delta k = 50$ years, then we will replace the component at the first inspection, which makes the cycle length 50 years, because the material has passed the replacement level at 45.5 years. The last remark concerning the expected cycle length is that when we choose $\Delta k = 1$ year, then the expected cycle length will be (about) the same as the expected time of preventive replacement. The shorter we choose the inspection interval, the higher the chance that we will replace the component exactly at the moment when the material has degraded up to the corrosion allowance.

4.5.2 With uncertainty

Now we consider the variables p and s to be uncertain. To include these together with the posterior of μ determined by the inspection results, we use Monte Carlo simulation. This means that we sample each variable N times. For the pressure p and material strength s we use a normal distribution with the means given by the input in table 4.1 and standard deviations determined with the coefficients of variation, which were introduced in section 2.1.1.

In contrast to when we have data obtained using perfect inspections, we do not have a continuous posterior density for the corrosion rate μ . Section A.4 includes a short explanation on how we can sample from a discrete distribution. With the inspection data from table 4.1, we calculate the

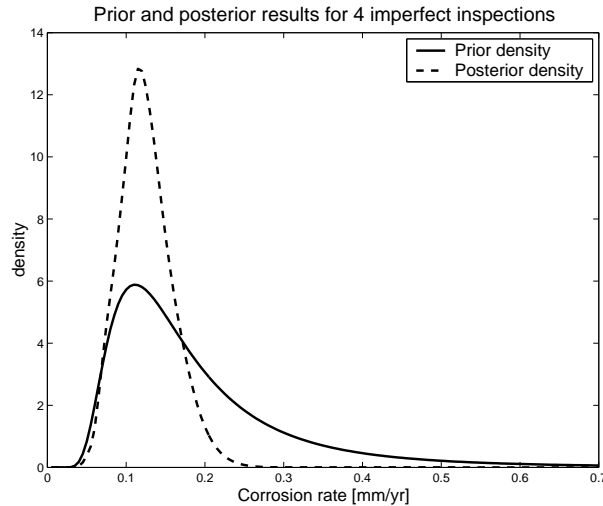


Figure 4.6: The prior and posterior density for μ given the inspection data of the hydrogen dryer.

posterior density for μ with equation (A.17). The result is given in figure 4.6. There are now two ways of calculating the expected average costs per year given by equations (A.24) and (A.25) in section A.4 of the appendix. The difference between the two is that for the first one we calculate the average for the expected costs per cycle and the average for the expected cycle length, after which we take the ratio equivalent to (4.7). These are shown in figures 4.7a and 4.7b respectively. These results were obtained using $N = 200$,

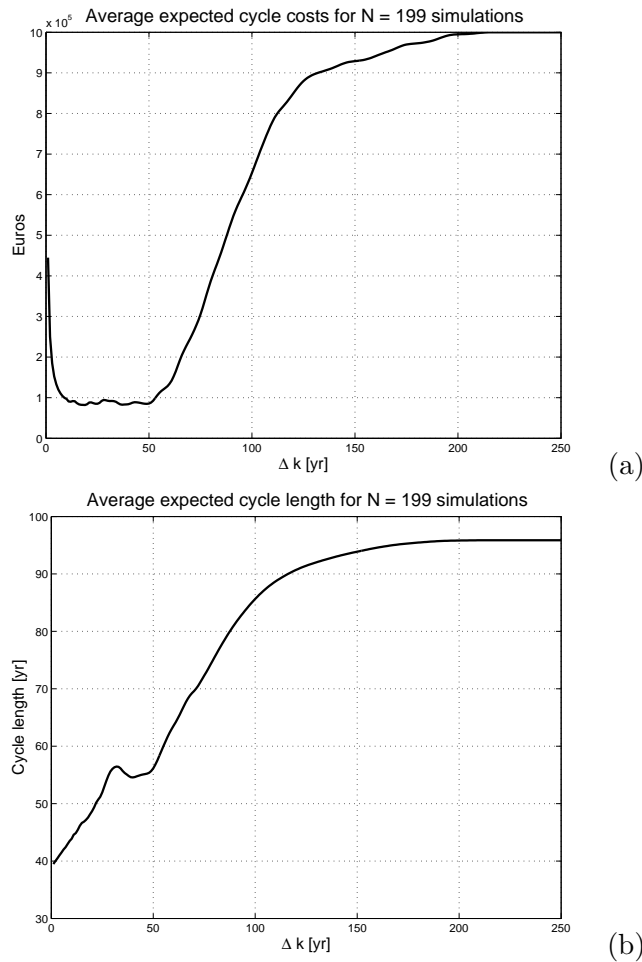


Figure 4.7: Expected costs per cycle (a) and the expected cycle length (b) calculated using 199 simulations.

but one calculation turned out to be incorrect, therefore we are left with $N = 199$ samples for each variable. This is usually not an adequate number of simulations and this is confirmed by the small ‘vibrations’ in the plots. Unfortunately these simulations take a very long time to perform, therefore these results should be considered as purely demonstrational. The final

decision will not change considerably, but the result will become smoother.

The second way of calculating the expected average costs per year is to calculate the ratio (4.7) N times and then take the average of these results. Mathematically, these two approaches can be interpreted as: 1) the ratio of the expectation of the expected cycle costs over the expectation of the expected cycle length, and 2) the expectation of the expected average costs per year respectively. Due to the fact that the data of a component is not used for a new component which took its place, we should use the first approach (ratio of averages) and not the second (average of ratios). Both results for the expected average costs per year are shown in figure 4.8 and we can conclude that they do not differ very much.

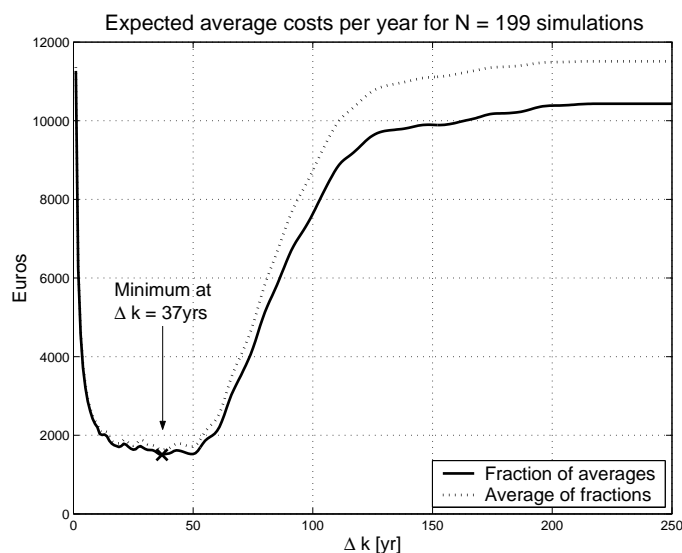


Figure 4.8: The expected average costs for the hydrogen dryer using two approaches with $N = 199$ simulations.

The optimal inspection interval is $\Delta k = 37$ years, which is associated with an expected average cost of about \$1500 per year when using (A.24). This amount is obtained when we divide the expected cycle costs of \$82,781 by the expected cycle length of 55.078 years. Since the hydrogen dryer is already 25 years old, the next inspection should be performed 12 years from now, i.e. in the year 2014. 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. On the other hand, if the inspection results in a contradicting measurement, then there is a big chance that the result

of this model will give an optimal interval which ended before the current date. When this happens, the item should be inspected as soon as possible or immediately replaced.

4.5.3 Influence of the replacement level

As a last remark in this chapter, I would like to explain the influence of the replacement level on the overall result. Since we assumed ρ to be fixed by the ratio of the corrosion allowance over the safety margin, this remark is not too important to make. However, if we would also be interested in optimizing the replacement level, then the influence of this parameter is very important to understand.

Most of the individual results for the expected cycle costs will look like the plot in figure 4.5b. Table 4.2 shows that $\rho \approx 40\%$. This means that the amount of corrosion, which can take place before the item is ready to be replaced, is less than half the amount of corrosion needed to result in failure of the component. Some of the results in the simulation have a $\rho > 50\%$, as is the case with simulation 36 for which the expected cycle costs are shown in figure 4.9. In this figure $\rho \approx 60\%$ and we see a large

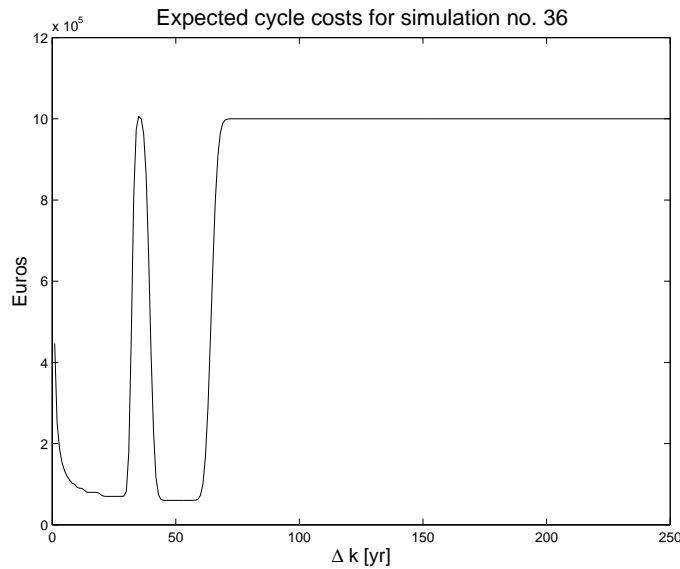


Figure 4.9: Expected cycle costs for simulation 36, where $\rho = 0.6078$.

peak at $\Delta k = 35$ years. The reason why this peak is there, can be found by looking at the expected time of replacement and the expected time of failure. These are approximately at 40 and 65 years respectively. If we would use $\Delta k = 35$ years, then the first inspection would be performed right before the corrosion allowance has been reached and by the time the

second inspection is performed at 70 years, the component has failed already. With a $\rho > 50\%$ there is a probability that the inspection interval is chosen such that the whole interval between the replacement level and the failure level is skipped. The component will in this case fail before we had the chance to inspect it a second time.

4.6 Model discussion

Overall, we can say that the gamma deterioration process and the expected average costs per unit time cost-criterion work very well for the problem at hand. The optimal inspection interval given by the model is not too short, because this would be too expensive, nor is it too long as it avoids the risk of failure becoming too large. A confirming inspection will enable us to inspect the component again in the future, whereas contradicting measurements could result in the necessity to replace the component immediately. The model even enables an expansion to include the possibility of modelling non-linear degradation. This means that we could model a corrosion rate which is not constant, but variable in time. Cases where this might arise are components, which have a coating applied or carry insulation to protect the surrounding equipment and personnel from the heat. These extra layers will corrode slower (coating) or faster (insulation) than the construction material itself. Therefore, a component with faulty coating will start corroding faster than before. The only downside of the model is that it is computationally very expensive. The simulation with 200 samples took about 3 hours to calculate on a Microsoft Windows XP computer with an AMD Athlon 1.7+ processor (clocks at 1.47GHz) and 256MB internal memory. This is therefore the only aspect in which the model fails to live up to its specifications which we defined at the beginning of this chapter. However, if one decides to actually implement the model into a software package in order to use in practice, then the following considerations will considerably increase the efficiency of the model:

- my implementation is in Matlab, which is a higher order programming language specifically geared towards mathematical calculations. It is convenient for this purpose due to the many available statistical and probabilistic functions. However, this program is notoriously slow with loops, which means that an implementation using a lower level programming language like Visual Basic (or C/C++, Pascal, Java etc.) will perform much faster.
- consider only constant corrosion and not non-linear degradation. This will enable us to go over to a so-called l_1 -isotropic grid by choosing $\Delta t = 1/a = \nu^2$. This will reduce the gamma density to an exponential density, which in turn results in a Poisson distribution for the proba-

bility of failure. These characteristics significantly simplify the calculations. More information on this topic can be found in van Noortwijk [15].

The model will however never come close to the speed with which ORBIT™ can perform calculations on hundreds of components. A small anecdote illustrates that ORBIT™ has also been in this position: the large project involving around 4000 components² that DNV Rotterdam did in 1998 for the Terneuzen plant owned by Dow Chemical, took a week to do all the calculations on what was their best computer at the time! As computers become more powerful, the gamma deterioration model will become more and more appealing.

²not all components were considered in the detailed calculations; only those that were considered to be high-risk items during an initial screening.

Final conclusions

There are roughly two parts to this thesis: one concerning the ORBIT™ inspection planning tool and a second part in which we consider a totally different approach to the inspection and maintenance planning problem. The ORBIT™ package has proven to be a very useful tool to get an indication of when certain components should be inspected such that the risk of a failure and its consequences is controlled. A sensitivity and uncertainty analysis using the new UNICORN software package has illustrated some of the dependencies in the corrosion and cracking models. The most interesting of these is that the pressure does not have any influence in the corrosion model, unless the component has a very large diameter. A change in operating pressure in pipelines will therefore not influence the result noticeably.

One of the most intensely investigated aspects of the methodology is how the probability of failure is approximated. In ORBIT™, this is done using the well known MVFOSM technique and works very well for the corrosion model which is described by a state function. The technique slightly underestimates the probability of failure in case of the cracking model, but not so much that it can be considered incorrect. A number of alternative techniques have been discussed, namely: FORM and directional sampling. FORM is a common reliability index method, whereas directional sampling belongs to the family of simulation methods. I have implemented the FORM technique to compare its performance to a crude Monte Carlo simulation method. This implementation has been done using the most common iteration scheme known as the Rackwitz–Fiessler algorithm. Surprisingly, this algorithm performed exceptionally bad for the corrosion model. This is mostly likely due to the almost linearity of the failure plane at the design point, which makes the algorithm ‘jump’ around while moving slowly towards this critical point. As an alternative, the problem is considered as a standard non-linear equality-constrained optimization problem. The theory of linear programming prescribes the use of Lagrange multipliers to find the optimal solution to this problem. This approach reduces the problem to finding a solution of a system of non-linear equations. One of the fastest and easiest schemes to solve this problem is the Newton–Raphson scheme. The number of iterations which are required in this case to achieve convergence, is astonishingly small compared to the Rackwitz–Fiessler algorithm. This

is especially true when the objective function of the optimization problem is slightly rewritten, which makes the resulting system of equations almost fully linear. This approach is expected to always outperform the Rackwitz–Fiessler algorithm and if not, then it will certainly be equally fast. Considering this, it is advised to always use the (simplified) Lagrange multiplier approach.

The implementation of the directional sampling technique to approximate the probability of failure has resulted in the conclusion that it is not an attractive alternative to the FORM technique for the corrosion model. However, more clarity has been brought to the much underestimated aspect of choosing the right radius to use in the equations for the failure probability. This only holds for the model with only normally distributed variables. The fact that it was necessary to create a decision tree for this purpose, shows that the technique has not yet grown past its initial phase. Better performance is expected from the model with exponential or Rayleigh densities, but the use of these densities would make the results incomparable to the FORM results.

Last but not least, a new approach (for the process industry) to the optimal inspection decision problem has been considered. Inspired by the successful application of the gamma stochastic deterioration process to the maintenance problem of hydraulic structures in the Netherlands, the very same model has been applied to model the deterioration of pressurized components in the process and refining industry. The applicability of the gamma process is remarkable, certainly considering the fact that the model has been expanded to also incorporate measurement data obtained using imperfect inspections. Since any kind of inspection will never be able to exactly measure the wall thickness of a component, it was critical to build this uncertainty into the Bayesian updating model. Finally the cost-based criterium of expected average costs per unit time has been used to find an optimal inspection interval. This criterium ensures that we maximize the safety while simultaneously minimizing the cost of inspecting the component. Altogether this approach fits extremely well with the problem at hand due to the mathematically elegant model for the degradation, but also due to the fact that the model does not require any more (complicated) input from the user. Besides the cost of a preventive replacement, the input from ORBIT™ only has to be supplemented by the cost of an inspection and the standard deviation of its measurement. The latter is known by the inspection engineer or can be found in numerous sources on non-destructive testing (NDT) techniques.

Currently, the time required for a decent amount of simulations is too long to cover hundreds of components. However, its efficiency can be greatly improved and its application does not have to be limited to the process industry. There are many more situations where uncertain inspections are present, which promises a bright future for this model or any derived models.

Appendix A

Mathematical details

A.1 FORM

The following calculations are the implementation of three approaches to FORM for the thinning model. The calculations for the stress corrosion cracking model are essentially the same, but are a little more elaborate due to the fact that there are two more variables in this model compared to the thinning model.

For all three methods we need to write the state function (2.1) in the standard normal space as given by the definition (3.5):

$$g(\sigma \mathbf{y} + \mu) = (\sigma_S Y_1 + \mu_S) \left[1 - (\sigma_C Y_2 + \mu_C) \frac{t}{th} \right] - (\sigma_P Y_3 + \mu_P) \frac{d}{2th}. \quad (\text{A.1})$$

Rackwitz–Fiessler

First we need to calculate the partial derivatives of the transformed state function (A.1):

$$\begin{aligned} \frac{\partial g}{\partial Y_1} &= \sigma_S \left[1 - (\sigma_C Y_2 + \mu_C) \frac{t}{th} \right] \\ \frac{\partial g}{\partial Y_2} &= -(\sigma_S Y_1 + \mu_S) \sigma_C \frac{t}{th} \\ \frac{\partial g}{\partial Y_3} &= -\sigma_P \frac{d}{2th} \end{aligned} \quad (\text{A.2})$$

Using these derivatives, it's easy to implement the unit vector $\boldsymbol{\alpha}$ as in (3.7) and to determine $\mathbf{y}^{(i+1)}$ using equation (3.6).

Lagrange multipliers

Regular objective function The Lagrangean function is given by

$$L(\mathbf{y}, \lambda) = \|\mathbf{y}\| + \lambda g(\mathbf{y})$$

and the partial derivatives of this function are:

$$F_i = \frac{\partial L}{\partial y_i} = \frac{y_i}{\|\mathbf{y}\|} + \lambda \frac{\partial g}{\partial y_i}, \text{ for } i = 1, 2, 3$$

$$F_4 = \frac{\partial L}{\partial \lambda} = g(\mathbf{y})$$

where the partial derivatives of g are given by (A.2). We have to set these derivatives equal to zero and solve for (\mathbf{y}, λ) . Due to the norm and the nonlinearity of g , this is a system of nonlinear equations. We solve this problem using the iterative *Newton-Raphson scheme*:

$$\mathbf{y}^{(i+1)} = \mathbf{y}^{(i)} - J^{-1}(\mathbf{y}^{(i)}) \cdot \mathbf{F}(\mathbf{y}^{(i)}),$$

where $\mathbf{F}(\mathbf{y}^{(i)})$ and $J^{-1}(\mathbf{y}^{(i)})$ are the system of nonlinear equations and the inverse Jacobian matrix respectively evaluated at $\mathbf{y}^{(i)}$. The i -th row of the *Jacobian matrix* is defined as

$$J(\mathbf{x})_i = \left(\frac{\partial F_1}{\partial y_i}, \dots, \frac{\partial F_n}{\partial y_i} \right).$$

For the thinning model the Jacobian is as follows:

$$J = \begin{pmatrix} \frac{y_2^2 + y_3^2}{\|\mathbf{y}\|^3} & -\frac{y_1 y_2}{\|\mathbf{y}\|^3} - \lambda \sigma_S \sigma_C \frac{t}{th} & -\frac{y_1 y_3}{\|\mathbf{y}\|^3} & \sigma_S \left[1 - \frac{t}{th} (\mu_C + \sigma_C y_2) \right] \\ -\frac{y_1 y_2}{\|\mathbf{y}\|^3} - \lambda \sigma_S \sigma_C \frac{t}{th} & \frac{y_1^2 + y_3^2}{\|\mathbf{y}\|^3} & -\frac{y_2 y_3}{\|\mathbf{y}\|^3} & -\sigma_C \frac{t}{th} [\sigma_S y_1 + \mu_S] \\ -\frac{y_1 y_3}{\|\mathbf{y}\|^3} & -\frac{y_2 y_3}{\|\mathbf{y}\|^3} & \frac{y_1^2 + y_2^2}{\|\mathbf{y}\|^3} & -\frac{\sigma_P d}{2th} \\ \sigma_S \left[1 - \frac{t}{th} (\mu_C + \sigma_C y_2) \right] & -\sigma_C \frac{t}{th} [\sigma_S y_1 + \mu_S] & -\frac{\sigma_P d}{2th} & 0 \end{pmatrix}$$

We see that in this case the starting point \mathbf{y}^0 should not be the zero vector, because this will result in a division by zero.

Simplified objective function Now the Lagrangean function is given by

$$L(\mathbf{y}, \lambda) = \mathbf{y}^T \mathbf{y} + \lambda g(\mathbf{y}),$$

which results in the following partial derivatives:

$$\frac{\partial L}{\partial y_i} = 2y_i + \lambda \frac{\partial g}{\partial y_i}$$

$$\frac{\partial L}{\partial \lambda} = g(\mathbf{y})$$

The Jacobian matrix is now

$$J = \begin{pmatrix} 2 & -\lambda \sigma_S \sigma_C \frac{t}{th} & 0 & \sigma_S \left[1 - \frac{t}{th} (\mu_C + \sigma_C y_2) \right] \\ -\lambda \sigma_S \sigma_C \frac{t}{th} & 2 & 0 & -\sigma_C \frac{t}{th} (\sigma_S y_1 + \mu_S) \\ 0 & 0 & 2 & -\frac{\sigma_P d}{2th} \\ \sigma_S \left[1 - \frac{t}{th} (\mu_C + \sigma_C y_2) \right] & -\sigma_C \frac{t}{th} (\sigma_S y_1 + \mu_S) & -\frac{\sigma_P d}{2th} & 0 \end{pmatrix}$$

All the elements in this symmetric matrix are linear and overall the matrix is much more simple compared to the previous Jacobian matrix for the regular objective function. Also, we can now choose any initial vector as a starting point. It therefore really pays off to use the squared norm for the objective function.

A.2 Directional sampling

Thinning model. Finding the solution r^* such that $g(r^*, \mathbf{a}) = 0$ is not too difficult for this model. Transforming the original variables to polar coordinates is done as follows:

$$\begin{aligned} Y_1 &= r \cos(\theta) \cos(\phi) \\ Y_2 &= r \cos(\theta) \sin(\phi) \\ Y_3 &= r \sin(\theta) \end{aligned}$$

It is not possible to simply sample θ and ϕ uniformly on $(0, 2\pi)$, because this will not result in a uniform distribution on the 3D-sphere. Sampling these angles such that the directions are uniformly distributed on this sphere is not straightforward (see chapter 7 in [9] for details) and therefore this approach is rarely used. A more convenient simulation method is to sample the independent standard normal variables Z_1, Z_2, Z_3 and to use the unit directional vector \mathbf{A} defined in (3.12). Then the variables can be rewritten as:

$$\begin{aligned} Y_1 &= rA_1 \\ Y_2 &= rA_2 \\ Y_3 &= rA_3 \end{aligned} \tag{A.3}$$

Substituting these into the state function $g(\mathbf{Y})$, where Y_i is the transformation given by (3.2), gives

$$\begin{aligned} g(r\mathbf{A}) &= (\sigma_1 r A_1 + \mu_1) \left[1 - (\sigma_2 r A_2 + \mu_2) \frac{t}{th} \right] - (\sigma_3 r A_3 + \mu_3) \frac{d}{2th} \\ &= - \left(\sigma_1 \sigma_2 A_1 A_2 \frac{t}{th} \right) r^2 + \dots \\ &\quad + \left[\sigma_1 A_1 \left(1 - \mu_2 \frac{t}{th} \right) - \sigma_2 \mu_1 A_2 \frac{t}{th} - \sigma_3 A_3 \frac{d}{2th} \right] r + \dots \\ &\quad + \mu_1 \left(1 - \mu_2 \frac{t}{th} \right) - \mu_3 \frac{d}{2th} \\ &\equiv ar^2 + br + c. \end{aligned}$$

This is a quadratic equation with unknown radius r , which we can solve by first calculating the discriminant $d = b^2 - 4ac$. If $d > 0$, then there

are two solutions for the radius: $r = (b^2 \pm \sqrt{d})/2a$. If $d = 0$, then there is one radius: $r = b^2/2a$ and if $d < 0$ there are two complex solutions for the radius, which we immediately disregard and assume there is no feasible radius for which $g = 0$. The next step is to choose the correct radius and calculate the corresponding failure probability according to the decision tree in figure 3.8.

For the thinning example, the situation where $d = 0$ never occurs and $d < 0$ only occurs when $g(\mathbf{0}) < 0$. When the center of the multidimensional standard normal distribution is located in the failure region, some sampled directions will pass exactly through the narrow opening between the two L-shaped surfaces which can clearly be observed in figure 3.7. This is the only time when there will be no radius for which $g = 0$, because there is no intersection with the failure surface. The complex radii should therefore be rejected and P_j set to 1.

SCC model. A more difficult problem is given by the stress corrosion cracking model, because now we are dealing with a state function for which we will not be able to determine the zero analytically. This is interesting to look at since most applications of directional sampling will involve models which require a numerical solution.

Equivalent to (A.3), we transform the variables to the polar coordinate space: $(Y_1, \dots, Y_5) = (rA_1, \dots, rA_5)$. Then we substitute these again in $g(\mathbf{Y})$:

$$g(r\mathbf{A}) = (\sigma_1 r A_1 + \mu_1) - \dots \\ + (\sigma_2 r A_2 + \mu_2) \left[\frac{(\sigma_3 r A_3 + \mu_3) d}{2th} + (\sigma_4 r A_4 + \mu_4) \right] \sqrt{\pi (\sigma_5 r A_5 + \mu_5)}.$$

This is a nonlinear equation, which we need solve using a numerical scheme. The fastest iterative scheme is the Newton–Raphson scheme which, for a single equation, is defined as

$$r^{(k+1)} = r^{(k)} - \frac{g(r^{(k)})}{g'(r^{(k)})},$$

where $k = 0, 1, 2, \dots$ is the iteration step and $g'(r^{(k)})$ is the partial derivative of $g(r\mathbf{A})$ to the radius r . The iteration is repeated until $|r^{(k+1)} - r^{(k)}| < \epsilon$, where ϵ is small (e.g. I've used $\epsilon = 0.01$) or until a maximum number of iterations has been reached (e.g. 300 iterations).

There are more than one solution for r such that $g(r, \mathbf{A}) = 0$, but the iteration scheme will only find one of them. Which one it finds depends on the starting value $r^{(0)}$. Since we want to find the closest (feasible) radius, it makes sense to take $r^0 = 0$ or, if this not possible (e.g. due to division by zero or a negative value under the square root), $r^0 \ll 1$.

This particular model requires the evaluation of $\mathbf{Y}^{(k)} = r^{(k)} \mathbf{A}$ at each step. As soon as one of the $Y_i^{(k)}$'s become smaller than $-\mu_i/\sigma_i$ the iteration

should be stopped. The radius in this case is considered to be infeasible. The condition can be relaxed to only the crack depth parameter ($i = 5$): if $\sigma_5 r^k A_5 + \mu_5 < 0$, then the square root will result in a complex radius. More general, if the model under consideration does not have this problem, then the radius can be rejected after the scheme has converged.

A smarter implementation of this scheme would include a number of different startvalues for the radius, which is followed by choosing the smallest feasible radius out of all the results. This will substantially increase the calculation effort, but it will result in a more dependable solution. While ensuring high accuracy, smart programming will enable the practitioner to avoid the directional sampling technique becoming less efficient than the regular sampling technique.

A.3 Gamma process inspection model

A.3.1 Gamma process

In this text I will use the following definition for the *gamma density* with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$:

$$\text{Ga}(x|\alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp\{-\beta x\} \text{ for } x \geq 0 \quad (\text{A.4})$$

A gamma process with stationary increments is defined as follows. 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 all $\tau > t \geq 0$,
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

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

The following is a proof for equations (4.1) and (4.2). In order to find the first (expectation) and second (variance) moments of $X(t)$, we use the *moment generating function*:

$$\begin{aligned} M_{X(t)}(\tau) &= Ee^{\tau X(t)} \\ &= \int_0^\infty e^{\tau x} f_{X(t)}(x) dx \\ &= \int_0^\infty \frac{b^{at}}{\Gamma(at)} x^{at-1} \exp\{(\tau - b)x\} dx \end{aligned}$$

Now we can use the substitution $-y = (\tau - b)x \Rightarrow dx = \frac{dy}{b-\tau}$ to solve the above integral:

$$\begin{aligned} M_{X(t)}(\tau) &= \int_0^\infty \frac{b^{at}}{\Gamma(at)} \left(\frac{y}{b-\tau} \right)^{at-1} \exp\{-y\} \frac{dy}{b-\tau} \\ &= \frac{b^{at}}{(b-\tau)^{at}\Gamma(at)} \underbrace{\int_0^\infty y^{at-1} \exp\{-y\} dy}_{=\Gamma(at)} \\ &= \frac{1}{(1-\tau/b)^{at}} \end{aligned}$$

We can use the first and second order derivatives of the so-called *cumulant generating function*

$$R_{X(t)}(\tau) = \ln(M_{X(t)}(\tau)) = -at \ln(1 - \tau/b)$$

with

$$\frac{\partial R}{\partial \tau} = \frac{at/b}{1-\tau/b} \quad \text{and} \quad \frac{\partial^2 R}{\partial \tau^2} = \frac{at/b^2}{(1-\tau/b)^2}$$

to determine the mean and variance of $X(t)$. These are now given by

$$\mathbb{E}(X(t)) = R'(0) = \frac{a}{b}t \quad \text{and} \quad \text{Var}(X(t)) = R''(0) = \frac{a}{b^2}t \quad (\text{A.5})$$

Let $X(t)$ be the cumulative deterioration at time t with $\mathbb{E}(X(t)) = \mu t$ and $\text{Var}(X(t)) = \sigma^2 t$, then

$$\frac{a}{b} = \mu \quad \text{and} \quad \frac{a}{b^2} = \sigma^2,$$

or

$$a = \frac{\mu^2}{\sigma^2} \quad \text{and} \quad b = \frac{\mu}{\sigma^2}, \quad (\text{A.6})$$

therefore $X(t) \sim \text{Ga}\left(x \mid \frac{\mu^2}{\sigma^2}t, \frac{\mu}{\sigma^2}\right)$.

A.3.2 Perfect inspection(s)

In sections 2.1.1 and 2.1.2 the coefficients of variation were introduced for the variables of the thinning and SCC models respectively. For example, in the thinning model, ORBIT assumes $\sigma_C = 0.1\mu_C$. The ratio σ/μ is therefore fixed for all model variables. Let $\nu^2 = (\sigma/\mu)^2$, then

$$a = \frac{1}{\nu^2} \quad \text{and} \quad b = \frac{1}{\mu\nu^2},$$

where ν is the coefficient of variation. Since this coefficient is fixed, the only uncertain variable left in this model is the mean deterioration μ , therefore the likelihood of a measurement x given μ is

$$l(x|\mu) = \text{Ga} \left(x \mid \frac{t}{\nu^2}, \frac{1}{\mu\nu^2} \right).$$

Although the use of this coefficient can be disputed, fixing the standard deviation of a variable relative to the user input for the mean is common practice in industrial applications. For us it is very handy, because we are only left with the uncertainty over the degradation rate μ , which is exactly the variable of interest. The user or the software program can express the uncertainty over μ by using a representative prior which is denoted with $\pi(\mu)$. Based on the discussion in chapter 4 we use the inverted gamma density (4.5) for the prior: $\pi(\mu) = \text{Ig}(\mu|\alpha, \beta)$.

If we have one perfect inspection which supplies us with a measurement x , then

$$\begin{aligned} l(x|\mu)\pi(\mu) &= \frac{\left[\frac{1}{\mu\nu^2}\right]^{t/\nu^2}}{\Gamma\left(\frac{t}{\nu^2}\right)} x^{\left(\frac{t}{\nu^2}-1\right)} \exp\left\{\frac{-x}{\mu\nu^2}\right\} \times \frac{\beta^\alpha}{\Gamma(\alpha)} \left(\frac{1}{\mu}\right)^{1+\alpha} \exp\left\{\frac{-\beta}{\mu}\right\} \\ &= \frac{\beta^\alpha \left[\frac{1}{\nu^2}\right]^{t/\nu^2}}{\Gamma\left(\frac{t}{\nu^2}\right)\Gamma(\alpha)} x^{\left(\frac{t}{\nu^2}-1\right)} \mu^{-\left(\frac{t}{\nu^2}+\alpha\right)-1} \exp\left\{-\frac{1}{\mu}\left(\frac{x}{\nu^2} + \beta\right)\right\} \\ &\equiv C \times \mu^{-A-1} \times \exp\left\{-\frac{1}{\mu}B\right\} \end{aligned} \quad (\text{A.7})$$

Now we take the integral over this result:

$$\int_0^\infty l(x|\mu)\pi(\mu)d\mu = C \int_0^\infty \left(\frac{1}{\mu}\right)^{A+1} \exp\left\{-\frac{1}{\mu}B\right\} d\mu \quad (\text{A.8})$$

With the following substitution:

$$y = B/\mu \Rightarrow d\mu = -y^{-2}Bdy,$$

(A.8) transforms to

$$\begin{aligned} \int_0^\infty l(x|\mu)\pi(\mu)d\mu &= C \int_\infty^0 \left(\frac{y}{B}\right)^{A+1} \exp\{-y\} (-y^{-2}) Bdy \\ &= \frac{C}{B^A} \int_0^\infty y^{A-1} \exp\{-y\} dy \\ &= \frac{C}{B^A} \Gamma(A). \end{aligned}$$

Using these two results, we can calculate the posterior density:

$$\begin{aligned}
\pi(\mu|x) &= \frac{l(x|\mu)\pi(\mu)}{\int_0^\infty l(x|\mu)\pi(\mu)d\mu} \\
&= \frac{B^A}{\Gamma(A)} \left(\frac{1}{\mu}\right)^{A+1} \exp\left\{-\frac{1}{\mu}B\right\} \\
&= \text{Ig}(\mu|A, B) \\
&= \text{Ig}\left(\mu \left| \frac{t}{\nu^2} + \alpha, \frac{x}{\nu^2} + \beta\right.\right). \tag{A.9}
\end{aligned}$$

To extend the model to allow for multiple inspection measurements x_i ($i = 1, \dots, n$), we need to multiply the likelihoods for the individual independent increments:

$$l(x_1, \dots, x_n|\mu) = \prod_{i=1}^n l_{X(t_i)-X(t_{i-1})}(x_i - x_{i-1}|\mu).$$

If we have only one measurement available, then this likelihood becomes

$$l(x|\mu) = l_{X(t)-X(0)}(x - 0|\mu) = l_{X(t)}(x|\mu),$$

where we used the property $X(0) = 0$. This is of course the same as the likelihood we used previously to calculate the posterior with only one inspection measurement. Identical to how we determined (A.7), we can perform the same calculations for multiple inspections:

$$\begin{aligned}
&l(x_1, \dots, x_n|\mu)\pi(\mu) \\
&= \prod_{i=1}^n l_{X(t_i)-X(t_{i-1})}(x_i - x_{i-1}|\mu) \times \pi(\mu) \\
&= \frac{\beta^\alpha}{\Gamma(\alpha)} \left[\frac{1}{\nu^2}\right]^{\sum_{i=1}^n \frac{t_i - t_{i-1}}{\nu^2}} \prod_{i=1}^n \left\{ \frac{(x_i - x_{i-1})^{\left(\frac{t_i - t_{i-1}}{\nu^2} - 1\right)}}{\Gamma\left(\frac{t_i - t_{i-1}}{\nu^2}\right)} \right\} \times \dots \\
&\times \left[\frac{1}{\mu}\right]^{\left(\sum_{i=1}^n \frac{t_i - t_{i-1}}{\nu^2} + \alpha\right) + 1} \exp\left\{-\frac{1}{\mu} \left(\frac{\sum_{i=1}^n x_i - x_{i-1}}{\nu^2} + \beta\right)\right\} \\
&\equiv C \times \mu^{-A-1} \times \exp\left\{-\frac{1}{\mu}B\right\}
\end{aligned}$$

We now have exactly the same result as (A.7), only now C is a different factor and

$$A = \frac{\sum_{i=1}^n t_i - t_{i-1}}{\nu^2} + \alpha \quad \text{and} \quad B = \frac{\sum_{i=1}^n x_i - x_{i-1}}{\nu^2} + \beta.$$

The posterior density is now given by:

$$\pi(\mu|x_1, \dots, x_n) = \text{Ig}\left(\mu \left| \frac{\sum_{i=1}^n t_i - t_{i-1}}{\nu^2} + \alpha, \frac{\sum_{i=1}^n x_i - x_{i-1}}{\nu^2} + \beta\right.\right). \tag{A.10}$$

The rest of the proof is identical to the proof for one inspection starting at (A.8).

A.3.3 Imperfect inspection(s)

Assuming that $Y(t) = X(t) + \epsilon$, where $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon)$, the likelihood of the measurement y given the degradation rate μ is given by the convolution:

$$l(y|\mu) = f_{Y(t)}(y) = f_{X(t)+\epsilon}(y) = \int_{-\infty}^{\infty} f_{X(t)}(y-\epsilon) f_\epsilon(\epsilon) d\epsilon, \quad (\text{A.11})$$

Here $f_{X(t)}(y-\epsilon) = \text{Ga}(y-\epsilon|at, b)$ is the likelihood of the gamma increment $X(t)$ with the parameters a and b as given by (A.6). Remember that the parameter b depends on the uncertain degradation rate μ . Unfortunately, it is not possible to solve the above integral analytically. We can discretize this integral to obtain an easier to evaluate sum of densities:

$$f_{Y(t)}(y) \approx \sum_{j=1}^m \text{Ga}(y-\epsilon_j|at, b) p(\epsilon_j), \quad (\text{A.12})$$

which is the discrete equivalent of $\mathbb{E}[h(\epsilon)]$, where $h(\epsilon) = \text{Ga}(y-\epsilon|a, b)$ and $\sum_{j=1}^m p(\epsilon_j) = 1$. Discretization means that we approximate the Riemann integral by a sum of ‘discrete’ parts:

$$\int_x f(x) dx \approx \sum_{j=1}^m f(x_j) \Delta x = \sum_{j=1}^m p(x_j),$$

which is demonstrated in figure A.1. Again, we want to determine the posterior density for the degradation factor μ . In this case it is given by

$$\pi(\mu|y) = \frac{l(y|\mu)\pi(\mu)}{\int_{\mu=0}^{\infty} l(y|\mu)\pi(\mu)d\mu}. \quad (\text{A.13})$$

If we now substitute the parameters of the gamma density such that the dependence on μ becomes clear, then the denominator is given by

$$\int_{\mu=0}^{\infty} l(y|\mu)\pi(\mu)d\mu = \int_{\mu=0}^{\infty} \left\{ \int_{\epsilon=-\infty}^{\infty} \text{Ga}\left(y-\epsilon \left| \frac{1}{\nu^2}t, \frac{1}{\mu\nu^2}\right.\right) f_\epsilon(\epsilon) d\epsilon \right\} \pi(\mu) d\mu.$$

These integrals can also be approximated by discretization:

$$\int_{\mu=0}^{\infty} l(y|\mu)\pi(\mu)d\mu \approx \sum_{i=1}^n \left\{ \sum_{j=1}^m \text{Ga}\left(y-\epsilon_j \left| \frac{1}{\nu^2}t, \frac{1}{\mu_i\nu^2}\right.\right) p(\epsilon_j) \right\} p(\mu_i),$$

where $p(\mu_i) = \Pr\{\mu = \mu_i\}$ is the discrete density of μ . Using Fubini’s theorem we can interchange the integrals in the above equations and integrate

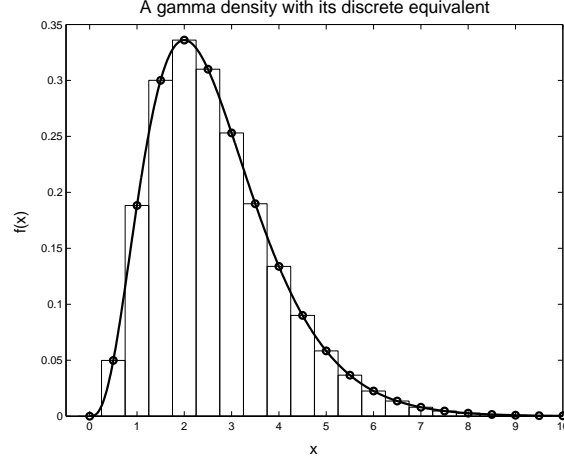


Figure A.1: A comparison between a continuous density function and its discrete equivalent for the purpose of numerical integration.

out the μ . This is only possible when we consider only one inspection. With these results, the posterior (A.13) is approximated by

$$p(\mu_i|y) \approx \frac{p(\mu_i) \left\{ \sum_{j=1}^m \text{Ga}(y - \epsilon_j | t/\nu^2, 1/[\mu_i\nu^2]) p(\epsilon_j) \right\}}{\sum_{i=1}^n p(\mu_i) \left\{ \sum_{j=1}^m \text{Ga}(y - \epsilon_j | t/\nu^2, 1/[\mu_i\nu^2]) p(\epsilon_j) \right\}}, \quad (\text{A.14})$$

for $i = 1, \dots, n$. Special care should be taken to ensure that the posterior is indeed a density, i.e. $\sum_{i=1}^n p(\mu_i|y) = 1$. Depending on how these results are implemented, it might be necessary to normalize the approximated density given by (A.12) such that the total mass equals one.

To extend the model to multiple imperfect inspections, we again need the product of individual likelihoods for the measurement increments:

$$l(y_1, \dots, y_k | \mu) = \prod_k l_{Y(t_k) - Y(t_{k-1})}(y_k - y_{k-1} | \mu).$$

Using the integral convolution like in (A.11), the above likelihood can be rewritten as

$$\begin{aligned} l(y_1, \dots, y_k | \mu) &= \\ &= \int_{-\infty}^{\infty} \dots \int_{-\infty}^{\infty} \prod_k f_{X(t_k) - X(t_{k-1})}(y_k - y_{k-1} - \delta_k) f(\delta_1, \dots, \delta_k) d\delta_1 \dots d\delta_k, \end{aligned}$$

where $\delta_k = \epsilon_k - \epsilon_{k-1}$ ¹. The δ 's are not independent, therefore we are left

¹Note the change of notation here: ϵ_k is the error term associated with the k -th inspection, whereas $p(\epsilon_j)$ in equation (A.14) refers to the j -th element of the discrete probability density function for ϵ .

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 above likelihood can be written as

$$\begin{aligned} l(y_1, \dots, y_k | \mu) &= \mathbb{E} \left[\prod_k f_{D_k}(d_k - \delta_k) \right] \\ &\approx \frac{1}{N} \sum_{j=1}^N \prod_k f_{D_k}(d_k - \delta_k^j) \text{ as } N \rightarrow \infty \end{aligned} \quad (\text{A.15})$$

where $D_k = X(t_k) - X(t_{k-1})$ and $d_k = y_k - y_{k-1}$. In the last step we used the strong law of large numbers (see for example Ross [20] or any other textbook on probability theory) to approximate the expectation of a function by use of simulation.

For each k we sample ϵ_k^j ($j = 1, \dots, N$) and calculate $\delta_k^j = \epsilon_k^j - \epsilon_{k-1}^j$. Since the gamma distributed $f_{D_k}(x) = \text{Ga}(x | a[t_k - t_{k-1}], b)$ is not defined for $x < 0$, we need to make sure that $d_k - \delta_k^j \geq 0$. In my implementation I have solved this by using

$$l(y_1, \dots, y_k | \mu) \approx \frac{1}{N} \sum_{j=1}^N \prod_k \text{Ga} \left(d_k - \min \{ \delta_k^j, d_k \} \mid \frac{t_k - t_{k-1}}{\nu^2}, \frac{1}{\mu\nu^2} \right) \quad (\text{A.16})$$

The equivalent of the posterior in (A.14) for multiple imperfect inspections is now given by

$$p(\mu_i | \mathbf{y}) = \frac{p(\mu_i) \frac{1}{N} \sum_{j=1}^N \prod_k \text{Ga} \left(d_k - \min \{ \delta_k^j, d_k \} \mid \frac{\Delta t_k}{\nu^2}, \frac{1}{\mu\nu^2} \right)}{\sum_{i=1}^n p(\mu_i) \frac{1}{N} \sum_{j=1}^N \prod_k \text{Ga} \left(d_k - \min \{ \delta_k^j, d_k \} \mid \frac{\Delta t_k}{\nu^2}, \frac{1}{\mu\nu^2} \right)} \quad (\text{A.17})$$

for $i = 1, \dots, n$ and $\Delta t_k = t_k - t_{k-1}$.

We now have a Bayesian updating model which accepts the results of multiple imperfect inspections. These imperfect inspections have a measurement error which is given by a standard deviation σ_ϵ . The advantage of using the simulation technique to calculate the posterior density over the degradation rate is that we can use different errors for each inspection. This allows us to insert the data from inspections which are different in precision and accuracy. During its operational life, a component might have been inspected using X-ray techniques and have undergone ultrasonic or magnetic flux leakage (MFL) testing as well. The results from all these measurements can be used and the (in)accuracy of each individual method can be accounted for.

A.4 Expected inspection and maintenance costs

Identical to (4.3), we define the safety margin m as

$$m \equiv m(p, s) = th - \frac{p \times d}{2s}$$

This safety margin is derived from the state function for corrosion (2.1). It's easiest to assume that the pressure p and material strength s are given. If we want to incorporate uncertainty over these variables, then this can be done by use of simulation. Given the cost of an inspection c_I , the cost of a preventive replacement c_P and the cost of failure c_F , we can write the expected cycle costs as (see chapter 7 in Noortwijk [15]):

$$\begin{aligned} & \sum_{i=1}^{\infty} c_i(\rho, \Delta k | p, s) p_i(\rho, \Delta k, \mu | p, s) = \\ & = \sum_{j=1}^{\infty} [j c_I + c_P] \Pr \{ X((j-1)\Delta k) \leq \rho m, \rho m < X(j\Delta k) \leq m \} + \\ & \quad \sum_{j=1}^{\infty} [(j-1)c_I + c_F] \Pr \{ X((j-1)\Delta k) \leq \rho m, X(j\Delta k) > m \} \end{aligned}$$

The first part of this equation can be interpreted as the cost incurred at time $j\Delta k$ where an inspection was followed by a replacement of the component. The second part corresponds to the costs due to inspections up until time $(j-1)\Delta k$ and a failure during the time period $((j-1)\Delta k, j\Delta k]$. This means that no inspection is required to determine if the item has failed, which is the case in the process industry. The replacement level is given by a percentage ρ ($0 < \rho < 1$) of the actual safety margin. For example, we can choose to replace the component when the wall loss has reached 70% of its critical level m . If desired, the model can be solved for multiple replacement levels such that the practitioner can also decide on the optimal replacement level.

In the process industry, each component has a maximal amount of corrosion assigned to it, up to which it is assumed to operate safely. This amount is called the *corrosion allowance* and it should always be less than the safety margin m . If the corrosion allowance c_{max} is known, then the safety level ρ is fixed by the ratio $\rho = c_{max}/m$.

To be able to implement the expected cycle costs, we need to expand

the previous equation as follows:

$$\sum_{i=1}^{\infty} c_i(\rho, \Delta k | p, s) p_i(\rho, \Delta k, \mu | p, s) = [c_I + c_P] \Pr \{ \rho m < X(\Delta k) \leq m \} + c_F \Pr \{ X(\Delta k) > m \} \quad (\text{A.18})$$

$$+ \sum_{j=2}^{\infty} [j c_I + c_P] \Pr \{ X((j-1)\Delta k) \leq \rho m, \rho m < X(j\Delta k) \leq m \} \quad (\text{A.19})$$

$$+ \sum_{j=2}^{\infty} [(j-1)c_I + c_F] \Pr \{ X((j-1)\Delta k) \leq \rho m, X(j\Delta k) > m \} \quad (\text{A.20})$$

Here we've taken the term with $j = 1$ outside the summation, which results in the equation (A.18) on the first line. This line can be further simplified as

$$\begin{aligned} [c_I + c_P] \Pr \{ \rho m < X(\Delta k) \leq m \} + c_F \Pr \{ X(\Delta k) > m \} &= \\ &= [c_I + c_P] (\Pr \{ X(\Delta k) \leq m \} - \Pr \{ X(\Delta k) \leq \rho m \}) \\ &+ c_F (1 - \Pr \{ X(\Delta k) \leq m \}) = \\ &= [c_I + c_P] (F_{X(\Delta k)}(m) - F_{X(\Delta k)}(\rho m)) \\ &+ c_F (1 - F_{X(\Delta k)}(m)), \end{aligned}$$

where $F_{X(\Delta k)}(x) = \Pr \{ X(\Delta k) \leq x \}$ is the cumulative distribution function of the increment $X(\Delta k) \sim \text{Ga}(x | a\Delta k, b)$ with parameters a and b defined by (A.3.2). The last line given by (A.20) in the equation for the expected cost is calculated as

$$\begin{aligned} \Pr \{ X((j-1)\Delta k) \leq \rho m, X(j\Delta k) > m \} &= \\ &= \Pr \{ X((j-1)\Delta k) \leq \rho m, [X(j\Delta k) - X((j-1)\Delta k)] > m - X((j-1)\Delta k) \} = \\ &= \int_{\delta=0}^{\rho m} \int_{\theta=m-\delta}^{\infty} f_{X(j\Delta k) - X((j-1)\Delta k)}(\theta) f_{X((j-1)\Delta k)}(\delta) d\theta d\delta = \\ &= \int_{\delta=0}^{\rho m} \int_{\theta=m-\delta}^{\infty} \text{Ga}(\theta | a\Delta k, b) \text{Ga}(\delta | a(j-1)\Delta k, b) d\theta d\delta = \\ &= \int_{\delta=0}^{\rho m} \frac{\Gamma(a\Delta k, b[m-\delta])}{\Gamma(a\Delta k)} \text{Ga}(\delta | a(j-1)\Delta k, b) d\delta \quad (\text{A.21}) \end{aligned}$$

The gamma function $\Gamma(a)$ and the incomplete gamma function $\Gamma(a, x)$ were introduced in equations (3.14) and (3.15) respectively. The second part of the expected costs, given by (A.19), can be written as a subtraction of two probabilities of which the second half can be solved in the same fashion as (A.21):

$$\begin{aligned} \Pr \{ X((j-1)\Delta k) \leq \rho m, \rho m < X(j\Delta k) \leq m \} &= \\ &= \Pr \{ X((j-1)\Delta k) \leq \rho m, X(j\Delta k) > \rho m \} \\ &- \Pr \{ X((j-1)\Delta k) \leq \rho m, X(j\Delta k) > m \}. \end{aligned}$$

The first half of this result can be considered as the probability of the gamma process surpassing the corrosion allowance ρm between the times $(j-1)\Delta k$ and $j\Delta k$:

$$\Pr \{X((j-1)\Delta k) \leq \rho m, X(j\Delta k) > \rho m\} = F_{X((j-1)\Delta k)}(\rho m) - F_{X(j\Delta k)}(\rho m)$$

It is therefore better and easier to solve any probabilities of this form as above and not like (A.21). Next we look at the expected cycle length:

$$\begin{aligned} \sum_{i=1}^{\infty} ip_i(\rho, \Delta k, \mu|p, s) &= \\ &= \sum_{j=1}^{\infty} j\Delta k \Pr \{X((j-1)\Delta k) \leq \rho m, \rho m < X(j\Delta k) \leq m\} \\ + \sum_{j=1}^{\infty} \sum_{n=(j-1)\Delta k+1}^{j\Delta k} n \Pr \{X((j-1)\Delta k) \leq \rho m, X(n-1) \leq m, X(n) > m\}, \end{aligned}$$

from which we can again take the part for $j=1$ out of the summation:

$$\begin{aligned} \sum_{i=1}^{\infty} ip_i(\rho, \Delta k, \mu|p, s) &= \\ &= \Delta k [\Pr \{X(\Delta k) \leq m\} - \Pr \{X(\Delta k) \leq \rho m\}] + \\ &+ \Pr \{X(1) > m\} + \sum_{n=2}^{\Delta k} n \Pr \{X(n-1) \leq m, X(n) > m\} + \\ &+ \sum_{j=2}^{\infty} [j\Delta k \Pr \{X((j-1)\Delta k) \leq \rho m, \rho m < X(j\Delta k) \leq m\} + \\ &+ [(j-1)\Delta k + 1] \Pr \{X((j-1)\Delta k) \leq \rho m, X((j-1)\Delta k + 1) > m\} + \\ &+ \sum_{n=(j-1)\Delta k+2}^{j\Delta k} n \Pr \{X((j-1)\Delta k) \leq \rho m, X(n-1) \leq m, X(n) > m\}] \end{aligned} \tag{A.22}$$

All the probabilities, except for the last one, in the above equation can be determined as with the expected cycle costs. The last probability can be simplified substantially by splitting it into parts. The interpretation of this probability is: the probability that the corrosion has not surpassed the corrosion allowance ρm at time $(j-1)\Delta k$ and failed during the interval $(n-1, n)$. Therefore the probability is split up as

$$\begin{aligned} \Pr \{X((j-1)\Delta k) \leq \rho m, X(n-1) \leq m, X(n) > m\} &= \\ &= \Pr \{X((j-1)\Delta k) \leq \rho m, X(n) > m\} + \\ &- \Pr \{X((j-1)\Delta k) \leq \rho m, X(n-1) > m\}. \end{aligned}$$

Both parts can be calculated using (A.21). The average cost per year is now found by calculating

$$C(\rho, \Delta k|p, s) = \frac{\int_{\mu=0}^{\infty} \sum_{i=1}^{\infty} c_i(\rho, \Delta k) p_i(\rho, \Delta k, \mu|p, s) \pi(\mu|\mathbf{y}) d\mu}{\int_{\mu=0}^{\infty} \sum_{i=1}^{\infty} i p_i(\rho, \Delta k, \mu|p, s) \pi(\mu|\mathbf{y}) d\mu} \approx \frac{\sum_{l=1}^n \{ \sum_{i=1}^{\infty} c_i(\rho, \Delta k) p_i(\rho, \Delta k, \mu|p, s) \} p(\mu_l|\mathbf{y})}{\sum_{l=1}^n \{ \sum_{i=1}^{\infty} i p_i(\rho, \Delta k, \mu|p, s) \} p(\mu_l|\mathbf{y})}, \quad (\text{A.23})$$

where $\pi(\mu|\mathbf{y})$ is the posterior density ($p(\mu_l|\mathbf{y})$ is the discrete version of the posterior density) over μ given the inspection measurements $\mathbf{y} = y_1, \dots, y_k$, which is given by equation (A.17). The choice of integrating out the degradation rate μ individually over the numerator and denominator as opposed to integration over the full fraction, is supported by the fact that we are only interested in one single cycle. After a component has been replaced it has to be inspected at least twice at intervals which are prescribed by the Dutch Rules for Pressure Vessels [21]. Also, the component will very often have undergone a wall thickness measurement at service start, therefore there will be around three sets of measurement data available. These results will be used to calculate the optimal time to the next inspection. The fundamental issue is that we do not use the data from the previous cycle/component to update the information of the current component.

If we also want to include the uncertainty over the pressure p and the material strength s , then we can sample these variables from their respective probability distribution functions f_P and f_S . The best approach is to also include the uncertainty over the degradation rate in the simulation, instead of using the discrete solution (A.23). In the case of perfect inspections, we can do this very simply by sampling from the inverted gamma distribution which we obtained in (A.10). Unfortunately we only have the posterior (A.17) in discrete form, when we consider multiple imperfect inspections. Therefore we need to know how to sample from this discrete (cumulative) distribution F . We start by sampling a uniformly distributed number u between 0 and 1 after which we can find the sample as follows:

$$F^{-1}(u) = \begin{cases} x_1 & \text{if } u \leq p_1 \\ x_2 & \text{if } p_1 < u \leq p_1 + p_2 \\ \vdots & \\ x_{n-1} & \text{if } \sum_{i=1}^{n-2} p_i < u \leq \sum_{i=1}^{n-1} p_i \\ x_n & \text{if } \sum_{i=1}^{n-1} p_i < u \leq \sum_{i=1}^n p_i \end{cases}, \quad u \in (0, 1)$$

The advantage of using the simulation technique is that the degradation rate μ can be sampled before we go into the loops to calculate the expected average costs per year. Also, there will be no extra loop needed if we assumed that the pressure p and material strength s are also uncertain. The solution

will in this case be given by

$$\mathbb{E}C(\rho, \Delta k) \approx \frac{\frac{1}{N} \sum_{j=1}^N \left\{ \sum_{i=1}^{\infty} c_i(\rho, \Delta k) p_i(\rho, \Delta k | \mu^{(j)}, p^{(j)}, s^{(j)}) \right\}}{\frac{1}{N} \sum_{j=1}^N \left\{ \sum_{i=1}^{\infty} i p_i(\rho, \Delta k | \mu^{(j)}, p^{(j)}, s^{(j)}) \right\}}, \quad (\text{A.24})$$

or

$$\begin{aligned} \mathbb{E}C(\rho, \Delta k) &\approx \frac{1}{N} \sum_{j=1}^{\infty} \frac{\sum_{i=1}^{\infty} c_i(\rho, \Delta k) p_i(\rho, \Delta k | \mu^{(j)}, p^{(j)}, s^{(j)})}{\sum_{i=1}^{\infty} i p_i(\rho, \Delta k | \mu^{(j)}, p^{(j)}, s^{(j)})} \quad (\text{A.25}) \\ &= \frac{1}{N} \sum_{j=1}^n C(\rho, \Delta k | \mu^{(j)}, p^{(j)}, s^{(j)}), \end{aligned}$$

as $N \rightarrow \infty$. The first result (A.24) corresponds to (A.23), where we assume that the data is used for only 1 cycle. If we would use the data from a replaced component to estimate the corrosion rate of the new component, then we would use (A.25), where the expectation is calculated over the full fraction. In the first case we first average the expected cycle costs and the expected cycle length independently before we calculate the ratio, whereas in the second case we calculate the average of N ratios representing the expected average costs per unit time. As we discussed on the previous page, we will not use (A.25) due to the fact that old data can not be used for a new component. Nonetheless, the comparison between the two results is interesting and can be found in figure 4.8.

Appendix B

Gamma process model implementation

This appendix contains all the Matlab code, which forms the implementation of the gamma process inspection decision model. Comments are included in the code itself, as well as throughout this appendix in order to make the procedures clear.

Matlab is a command line driven program, where a series of commands can be put into separate files, which can in turn be run from the command line. There is also the possibility of having calculations performed by separate functions, which return the results to the main programs. Comment lines are preceded by a % and ... splits long lines onto the next line.

B.1 modelinput: input data

The following code prepares all the input data for the model. The data which you can see here is for the case study of the hydrogen dryer.

```
1  dm = 1180;                % component diameter [mm]
   th = 15*1.12;            % thickness at service start [mm]
   TS = 413.69;            % tensile strength [MPa]
   YS = 206.84;            % yield strength [MPa]
5
   mu_sf = min(1.1*(TS+YS)/2,TS); % mean strength of material [MPa]
   mu_cr = 0.1;            % estimated corrosion rate [mm/yr]
   mu_pr = 3.2;            % mean operating pressure [MPa]
   pressure_cov = 0.05;    % coefficient of variation = 5%
10  flowstress_cov = 0.20;  % coefficient of variation = 20%
   si_sf = flowstress_cov*mu_sf; % standard deviation strength
   si_pr = pressure_cov*mu_pr;  % standard deviation pressure
   cov = 0.25;            % cov for corrosion rate = 25%
   CorrAllowance = 4.5;   % replacement level [mm]
15
   c_fail = 1000000;      % cost of failure + corrective repl.
```

```

c_insp = 10000;           % cost of a single inspection
c_prev = 50000;          % cost of preventive replacement

20  t = [86-77 90-86 94-90 98-94]; % times between past inspections
    D = [(th-15.6) (15.6-14.6) ... % amount of corrosion during these times
          (14.6-14.2) (14.2-13.8)];
    K = length(t);        % total number of inspections

25  InspCOV = [0.5; 0.5; 0.5; 0.5]; % cov for each inspection
    SigmaEpsilon = InspCOV.*D';    % standard deviation for each inspection

    % define the grid for the normal density of the measurement error:
    NormalLimit = round(max(SigmaEpsilon)*1000*4)/1000;
30  eps = [-NormalLimit:0.001:NormalLimit];

```

B.2 corrgamma: main program

The main program first runs the `modelinput` program to define the input data. It then calls the function `invgamfit` to find the shape parameter α and the scale parameter β for the prior. It does this by sampling the ORBIT™ discrete prior $n = 10,000$ times and fitting the inverted gamma density to this data. See the code listing of this function later on in this appendix for more details. The next step is to create a grid x at which points the prior and posterior are calculated.

```

1  clear           % clear the memory
   modelinput;    % load the input data

   % fit a continuous inverted gamma density to the discrete ORBIT prior:
5  n = 10000;
   par = invgamfit(mu_cr);
   a = par(1);
   b = 1/par(2);

10 % define the grid over which the densities are calculated:
    GridDist = mu_cr/20;
    x = GridDist:GridDist:7*mu_cr;
    N = length(x);

15 % the inverted gamma distributed prior is given by:
    Prior = exp(a*log(b)-gammaln(a)+(-a-1)*log(x)-b./x);

    % inverted gamma posterior for 1 perfect inspection:
    A = a + t(K)/cov^2;
20  B = b + D(K)/cov^2;
    PostPerfInsp = exp(A*log(B)-gammaln(A)+(-A-1)*log(x)-B./x);

```

At this stage, the prior and the posterior for K perfect inspections have been defined using the definition of the inverted gamma density and equation (A.10). A trick has been used here to avoid numbers which are too large

for Matlab to handle. Especially when the mass of the densities is very concentrated, there is a probability that the program will result in ‘NaN’, which means Not-A-Number. We therefore first take the natural logarithm of the densities, followed by taking the exponential of this: $x = \exp(\ln(x))$. The Matlab function `gammaIn(a)` is a special implementation of $\ln(\Gamma(a))$. For the inverted gamma density this results in

$$\frac{b^a}{\Gamma(a)} \left(\frac{1}{x}\right)^{(a+1)} \exp\left\{-\frac{b}{x}\right\} = \exp\left\{a\ln(b) - \ln(\Gamma(a)) - (a+1)\ln(x) - \frac{b}{x}\right\}$$

for $x > 0$. Notice that the grid `x` does not include $x = 0$, because this would result in division by zero. The commands `./`, `.^` and `.*` are element-wise versions of division, raising to the power and multiplication respectively. These are used with vectors when the operations have to be performed for each element of these vectors or matrices independently. In the next stage of the main program, the posterior for K imperfect inspections is calculated. This is done using simulation in which we use $n = 2000$ samples for normally distributed measurement error:

```

% posterior is found by simulation over the measurement error:
n = 2000;

25 E = zeros(n,K);      % matrix with sampled errors
   d = zeros(n,K);      % d_k = e_k - e_(k-1) with e_0 = 0
   for k=1:K
       % each column holds the samples of an inspection:
       E(:,k) = normrnd(0,SigmaEpsilon(k),n,1);
30   if k==1
       d(:,k) = E(:,k);
       else
       d(:,k) = E(:,k) - E(:,k-1);
       end
35   end

   likelihood = zeros(N,K);
   LikeliProd = zeros(N,1);

40   for j=1:N          % loop over grid for degradation rate
       for k=1:K      % loop over inspection number
           likelihood(j,k) = (1/n)*sum( ...
               exp(-(t(k)/cov^2)*log(x(j)*cov^2) - ...
               gammaIn(t(k)/cov^2) + (t(k)/cov^2 - 1)*log(D(k) - ...
45               min(D(k)-0.001,d(:,k))) - ...
               (D(k) - min(D(k),d(:,k)))/(x(j)*cov^2));
           end
       end
       % take the product of the individual likelihoods:
50   LikeliProd = prod(likelihood,2);
   PostImpInsp = Prior'.*LikeliProd/(Prior*LikeliProd*GridDist);
   % make the cumulative distribution for the sampling procedure:
   PostImpInspCDF = cumsum(PostImpInsp)*GridDist;

```

```

55 % open a new figure and plot the prior + posteriors:
    figure
    plot(x,Prior,'g-',x,PostPerfInsp,'m-',x,PostImpInsp,'r-', 'LineWidth',2);
    grid
    legend('prior density','posterior 1 perf. insp.',...
60     ['posterior ',num2str(K),' imp. insp.'],0);
    title(['Prior and posterior results (\sigma_\epsilon = ',...
        num2str(mean(SigmaEpsilon)),'; n = ',num2str(n),'')]);
    xlabel('Corrosion rate [mm/yr]');
    ylabel('density');
65     test = input('Ready to start simulation. Press ENTER to continue.');
```

The last part of this stage has opened a figure window and plotted the prior, the posterior for K perfect inspections and the posterior for K imperfect inspections. The likelihood inside the FOR-loops is the likelihood given by equation (A.16).

Now that we have the discrete posterior for K imperfect inspections, we can start calculating the expected average costs per time unit. As explained in section A.4, we can best do this by sampling the pressure p , the material strength s and the posterior for the corrosion rate c . The posterior is sampled using the technique described in the section A.4.

```

    % sample N times the variables p, s and c + perform simulations:
    N = input('Please input the number of simulations desired: ');
    p = normrnd(mu_pr,si_pr,N,1); % normal distributed samples for pressure
70    s = normrnd(mu_sf,si_sf,N,1); % normal distributed samples for strength
    m = th - p*dm./(2*s); % vector of (2*s) safety margins
    c = zeros(N,1); % initialize the samples vector for c
    for i=1:N
        u = unifrnd(0,1); % sample a uniformly distributed in (0,1)
75        c(i) = x(min(find(PostImpInspCDF > u)));
    end
    MaxTime = max(m./c);
    % the replacement percentage rho is given by the following ratio:
    Rho = CorrAllowance./m;
80
    TimeUnit = input('Please input the desired time unit: ');
    PrepData = zeros(N,3);
    for i=1:N
        PrepData(i,:) = simulprep(cov,m(i),Rho(i),c(i),TimeUnit,MaxTime);
85    end
    % determine maximum TimeHorizon:
    TimeHorizon = max(PrepData(:,3))
    ExpAveCost = zeros(length(TimeUnit:TimeUnit:TimeHorizon),2,N);
90
    LoopTimer = zeros(N,1);
    StartTime = clock; % save current date and time
    % display the start time and date to the user:
    disp(['Calculations started at ',num2str(round(StartTime(4))),'h', ...
        num2str(round(StartTime(5))),'m on ', ...
```



```

95         num2str(round(StartTime(3))), '/' , ...
           num2str(round(StartTime(2))), '/' , ...
           num2str(round(StartTime(1))), '.' ]);

for i=1:N                                     % loop over number of sample sets
100     disp(['Currently calculating exp. average cost for simulation ', ...
            num2str(i), ' out of ', num2str(N), ' ...']);
    % plot individual results for the last simulation:
    if i == N
        tic                                   % start timer
105         ExpAveCost(:, :, i) = eacptu(cov, m(i), Rho(i), c(i), ...
            TimeUnit, TimeHorizon, c_insp, c_prev, c_fail, 1);
    else
        tic                                   % start timer
110         ExpAveCost(:, :, i) = eacptu(cov, m(i), Rho(i), c(i), ...
            TimeUnit, TimeHorizon, c_insp, c_prev, c_fail, 0);
    end
    LoopTimer(i) = round(toc);                % stop timer
    % average time per loop is:
    CalcDuration = N*sum(LoopTimer)/i;
115     % determine when the calculations will be done + output to the user:
    FinishTime = calcfinish(N, StartTime, CalcDuration);
    disp(['Estimated end time: around ', num2str(FinishTime(2)), 'h', ...
          num2str(FinishTime(3)), 'm, ', num2str(FinishTime(1)), ...
          ' days later.']);
120 end

```

This part of the program uses a number of functions: `simulprep`, `calcfinish` and `eacptu`. The first one calculates the expected time of replacement, the expected failure time and the so-called time horizon for each sample set. The time horizon is the expected failure time plus 20% and is used as the limit up to which the inspection intervals are summed up. This horizon replaces the infinity in the sum of equation (4.7) and is chosen such that the summation is performed far enough in order to ensure that all failure probabilities are fully accounted for. This is done for each sample set and then the largest is used for the calculation of the expected average costs per time unit. The most important function is `eacptu`, which calculates the expected cycle costs and the expected cycle length for each sample set. The third function `calcfinish` is a simple function which estimates the time and day at which the calculations will be done.

The result of this program is a matrix containing the vectors with the expected cycle costs and the expected cycle length for each sample set. The calculations for the hydrogen dryer case study contained one faulty result, therefore the data has to be checked before it is further used to find the expected average costs per time unit.

B.3 invgamfit: fitted inverted gamma prior

This function is used by the main program to estimate the parameters of the inverted gamma density when fitted to the ORBIT™ prior. For the case study of the hydrogen dryer $\alpha = 3.58$ and $\beta = 0.51$. The parameters are estimated using the Matlab function `gamfit(x)`, which estimates the gamma density parameters¹ given the data \mathbf{x} . We use the fact that $1/\mathbf{x} \sim \text{Ga}(\alpha, \beta)$. This is not the best way of fitting the inverted gamma density to the ORBIT™ prior, but since this prior is highly subjective and has the sole purpose of adding conservatism, extreme precision is not necessary.

```

1  function y = invgamfit(c)
    % INVGAMFIT parameter estimation for the Inverted Gamma density
    % using 50% for 1*c, 30% for 2*c and 20% for 4*c. This is the
    % prior (discrete) density which is used in DNV's ORBIT package.
5  %
    % [a b] = invgamfit(c), where c is the degradation rate (e.g. mm/yr)
    % a and b are the parameter used by MATLAB for the Gamma distribution
    % NOTE: check the help file for the Gamma density to see how these
    % parameters are defined!
10
    n = 100;
    U = unifrnd(0,1,n,1);
    G = zeros(n,1);

15  for i=1:n
        if U(i)<=0.5
            G(i)=1*c;
        elseif U(i)>0.5 & U(i)<=0.8
            G(i)=2*c;
20        else
            G(i)=4*c;
        end
    end

25  GInv = 1./G;
    y = gamfit(GInv);

```

Once the parameters have been found, they can be hard coded into the model, because only the scale parameter β depends on the estimated corrosion rate. For example, $\beta \approx 0.5$ if $\mu_C = 0.1\text{mm/yr}$ and $\beta \approx 5$ when $\mu_C = 1\text{mm/yr}$.

B.4 simulprep: simulation data preparation

This function determines the expected time of replacement, the expected time of failure and the necessary time horizon over which the inspection

¹Note that the gamma density is defined differently in Matlab, than we have in (A.4). The scale parameter is the inverse of our β .

intervals have to be summed. For the replacement time and the failure time it uses the probability p_i of replacement/failure in unit time i :

$$p_i = F_{X(i-1)}(\rho m) - F_{X(i)}(\rho m) \quad (\text{replacement})$$

$$p_i = F_{X(i-1)}(m) - F_{X(i)}(m) \quad (\text{failure})$$

The expectations are then calculated using $\sum_{i=1}^m ip_i$:

```

1  function y = simulprep(cov,m,rho,mu,dT,MaxT)
   % y = simulprep(cov,m,rho,mu,dT,MaxT)
   %
   % This function calculates the expected time of preventive
5  % replacement, the expected time of failure and the time
   % horizon over which the calculations are performed
   % given the coefficient of variation, the margin m,
   % the preventive replacement percentage rho, the sampled
   % average deterioration rate mu and the time unit dT.
10
   a = 1/cov^2;                % shape parameter
   b = 1/(mu*cov^2);          % scale parameter

   ExpFailTime = 1 - gammainc(b*m,a*dT);
15  ExpReplTime = 1 - gammainc(b*rho*m,a*dT);
   for i=2*dT:dT:round(1.5*MaxT)
       ExpFailTime = ExpFailTime + ...
           i*(gammainc(b*m,a*(i-1*dT)) - gammainc(b*m,a*i));
       ExpReplTime = ExpReplTime + ...
20      i*(gammainc(b*rho*m,a*(i-1*dT)) - gammainc(b*rho*m,a*i));
   end
   % horizon over which calculations are done:
   TimeHorizon = round(1.2*ExpFailTime);

25  y = [ExpReplTime ExpFailTime TimeHorizon];

```

B.5 eacptu: exp. ave. costs per time unit

The key function in this program is the following, which calculates the expected cycle costs and the expected cycle length for each sample set. The last part of this function is only active when the last parameter *Fig* is set to one. This tells the function to plot the expected cycle costs, the expected cycle length and their ratio: the expected average costs per time unit. The calculations for the expected cycle costs and length are split up in two parts: $j = 1$ and $j > 1$. This is equivalent to equations (A.18)-(A.20) and (A.22).

```

1  function y = eacptu(cov,m,rho,mu,TimeUnit,TimeHorizon,Ci,Cp,Cf,Fig)
   % y = eacptu(cov,m,rho,TimeUnit,TimeHorizon,Ci,Cp,Cf,Fig)
   %
   % Given the corrosion safety margin m and the replacement
5  % percentage rho + a sample of the margin m, this function
   % determines the expected average costs per time unit.

```

```

% The cost of an inspection, a preventive replacement and
% of failure are given by Ci, Cp and Cf respectively.
% If separate figures are desired, set Fig = 1, otherwise
10 % set Fig = 0.

a = 1/cov^2;           % shape parameter
b = 1/(mu*cov^2);     % scale parameter

15 ExpCost = zeros(TimeHorizon/TimeUnit,1);
ExpCycle = zeros(TimeHorizon/TimeUnit,1);

D = 50;               % number of discretization points
M = m*[0:D-1]'/(D-1); % range (0,m) split in D-1 parts
20 dM = m/(D-1);      % distance between grid points
Pdf = zeros(D,1);    % initialize array for discrete density

index = 0;
for dk=TimeUnit:TimeUnit:TimeHorizon
25   index = index + 1;
   % expected cost for j=1: -----
   ExpCost(index) = (Ci + Cp)*(gammainc(m*b,a*dk) - ...
       gammainc(rho*m*b,a*dk)) + Cf*(1-gammainc(m*b,a*dk));
   % expected cycle length for j=1: -----
30   ExpCycle(index) = dk*(gammainc(m*b,a*dk) - ...
       gammainc(rho*m*b,a*dk)) + 1 - gammainc(m*b,a*TimeUnit);
   n = [2*TimeUnit:TimeUnit:dk]';
   ExpCycle(index) = ExpCycle(index) + ...
       n*(gammainc(m*b,a*(n-1*TimeUnit)) - gammainc(m*b,a*n));
35   for j=2:round(TimeHorizon/dk)
       % expected cost for j>1: -----
       Pdf(2:D) = rho*dM*exp(a*(j-1)*dk*log(b) - ...
           gammaln(a*(j-1)*dk) + (a*(j-1)*dk-1)*log(rho*M(2:D)) - ...
40           b*rho*M(2:D));
       ExpCost(index) = ExpCost(index) + ...
           (j*Ci + Cp)*(gammainc(rho*m*b,a*(j-1)*dk) - ...
               gammainc(rho*m*b,a*j*dk) - sum((1 - ...
                   gammainc(b*(m-rho*M),a*dk)).*Pdf)) + ...
45           ((j-1)*Ci + Cf)*sum((1-gammainc(b*(m-rho*M),a*dk)).*Pdf);
       % expected cycle length for j>1: -----
       ExpCycle(index) = ExpCycle(index) + ...
           j*dk*(gammainc(rho*m*b,a*(j-1)*dk) - ...
               gammainc(rho*m*b,a*j*dk) - sum((1 - ...
                   gammainc(b*(m-rho*M),a*dk)).*Pdf)) + ...
50           ((j-1)*dk+1*TimeUnit)*sum((1 - ...
               gammainc(b*(m-rho*M),a*TimeUnit)).*Pdf);
       n = [(j-1)*dk+2*TimeUnit:TimeUnit:j*dk]';
       ExpCycle(index) = ExpCycle(index) + ...
55       n*((1-gammainc(b*(m-rho*repmat(M',[length(n) 1])), ...
           a*(repmat(n,[1 length(M)])-(j-1)*dk)))*Pdf - ...
           (1-gammainc(b*(m-rho*repmat(M',[length(n) 1])), ...
           a*(repmat(n,[1 length(M)]-1*TimeUnit-(j-1)*dk)))*Pdf);
   end
60 end

```

```

if Fig == 1
    % first calculate expected time to replacement and failure:
    ExpFailTime = 1 - gammainc(b*m,a*TimeUnit);
65    ExpReplTime = 1 - gammainc(b*rho*m,a*TimeUnit);
    for i=2*TimeUnit:TimeUnit:200
        ExpFailTime = ExpFailTime + ...
            i*(gammainc(b*m,a*(i-1*TimeUnit)) - gammainc(b*m,a*i));
        ExpReplTime = ExpReplTime + ...
70        i*(gammainc(b*rho*m,a*(i-1*TimeUnit)) - ...
            gammainc(b*rho*m,a*i));
    end

    figure
75    x = TimeUnit:TimeUnit:TimeHorizon;
    plot(x,ExpCycle,'b.-',x,ExpFailTime*ones(length(x)),'r:');
    legend('Exp. cycle length','Exp. time of failure',0);
    title('Expected cycle length');
    xlabel('\Delta k [yr]');
80    ylabel('Cycle length [yr]');

    figure
    plot(x,ExpCost,'r.-');
    title('Expected costs per cycle length');
85    xlabel('\Delta k [yr]');
    ylabel('Euros');

    figure
    plot(x,ExpCost./ExpCycle,'go-','LineWidth',2);
90    hold
    plot([ExpFailTime ExpFailTime],[0 max(ExpCost./ExpCycle)],'r:',...
        [ExpReplTime ExpReplTime],[0 max(ExpCost./ExpCycle)],'b--');
    legend('Exp. average costs','Exp. time of failure', ...
        'Exp. time of replacement',0);
95    title(['Expected average costs per time unit (\Delta t = ', ...
        num2str(TimeUnit),'yr)']);
    xlabel('\Delta k [yr]');
    end
100 y = [ExpCost ExpCycle];

```

B.6 calcfinish: simulation timer

The last function simply adds the expected total simulation duration to the current date and outputs the expected time when the simulation is finished. In case the calculations are spread over more than one day, it also outputs the amount of days are required.

```

1 function y = calcfinish(N,StartTime,CalcDuration)
    % y = calcfinish(N,StartTime,CalcDuration)
    %
    % this function calculates the time at which the simulation

```

```
5  % is expected to be done. It also indicates how many days
   % ahead from the current date this is in case the calculations
   % are spread over multiple days.

   CalcDays = floor(CalcDuration/86400);
10  leftover = mod(CalcDuration,86400);
   CalcHours = floor(leftover/3600);
   leftover = mod(leftover,3600);
   CalcMinutes = floor(leftover/60);

15  FinishMin = mod((StartTime(5) + CalcMinutes),60);
   FinishHour = mod((StartTime(4) + CalcHours + ...
   floor((StartTime(5)+CalcMinutes)/60)),24);
   CalcDays = CalcDays + floor((StartTime(4) + CalcHours + ...
   floor((StartTime(5)+CalcMinutes)/60))/24);
20  y = [CalcDays FinishHour FinishMin];
```

References

- [1] API581 (a.k.a. Base Resource Document), 2nd edition, American Petroleum Institute, October 2000
- [2] ASME Boiler and Pressure Vessel Code, American Society of Mechanical Engineers, New York, 2001.
- [3] Cornell, C.A.: A probability-based structural code. *Journal of the American Concrete Institute*. 66(12):p974–985, 1969.
- [4] Ditlevsen, O. and Madsen, H.O.: *Structural Reliability Methods*. Chichester, John Wiley & Sons, 1996.
- [5] Ellingwood, B.R. & Zheng, R.: Role of Non-destructive evaluation in Time-dependent Reliability Analysis. In *Structural Safety*, 20:p.325–339, Elsevier Science, 1998.
- [6] European Union directive 97/23/EC, Official Journal L181, 09/07/1997:p0001–0055, 1997.
- [7] Haldar, A. and Mahadevan, S.: First-Order and Second-Order Reliability Methods. In *Probabilistic Structural Mechanics Handbook: theory and industrial applications*, Sundararajan, C. (ed.), New York, Chapman and Hall, p27–52, 1995.
- [8] Hasofer, A.M. and Lind, N.C.: Exact and invariant second moment code format. *Journal of the Engineering Mechanics Division of American Society of Civil Engineers*. 100(part 1):p111–121, 1974.
- [9] Johnson, M.E.: *Multivariate Statistical Simulation*. John Wiley & Sons, New York, 1987.
- [10] Madsen, H.O., Krenk, S. and Lind, N.C.: *Methods of Structural Safety*. Englewood Cliffs NJ, Prentice-Hall, 1986.
- [11] Mahadevan, S. & Zhang, R.: Model Uncertainty and Bayesian Updating in Reliability-Based Inspection. In *Structural Safety*, 22:p.145–160, Elsevier Science, 2000.

-
- [12] Melchers, R.E.: *Structural Reliability Analysis and Prediction*. 2nd edition, Wiley, Chichester, 1999.
- [13] Newby, M. and Dagg, R.: Optimal Inspection Policies in the Presence of Covariates. *Proceedings of ESREL 2002*, Lyon (France), 2002.
- [14] Noortwijk, J.M., Cooke, R.M., Kok, M.: A Bayesian Failure Model based on Isotropic Deterioration. *European Journal of Operational Research*, 82(2):p270–282, 1995.
- [15] Noortwijk, J.M.: *Optimal Maintenance Decisions for Hydraulic Structures under Isotropic Deterioration*. Ph.D. thesis, Delft University of Technology, 1996.
- [16] ORBIT Technical Manual, Det Norske Veritas Inc., Houston U.S.A., May 2000.
- [17] OREDA-97: *Offshore Reliability Data*. 3rd edition, Det Norske Veritas, Høvik Norway, 1997.
- [18] Pandey, M.D.: Probabilistic Models for Condition Assessment of Oil and Gas Pipelines. In *NDT&E International*, Vol.31 no.5 p.349–358, Elsevier Science, 1998.
- [19] Rackwitz, R. and Fiessler, B.: Structural reliability under combined random load sequences. *Computers and structures*. 9:p489–494, 1978.
- [20] Ross, S.M.: *Introduction to Probability Models*, 7th edition, Harcourt/Academic Press, 2000.
- [21] Rules for Pressure Vessels, Vol.1–3, Sdu Publishers, The Hague, 1997.
- [22] Robinson, D.G.: *A survey of probabilistic methods used in reliability, risk and uncertainty analysis*. Sandia National Laboratories report SAND98–1189, Albuquerque New Mexico, June 1998.
- [23] Thoft–Cristensen, P. and Baker, M.J.: *Structural Reliability Theory and its Applications*. Springer–Verlag, 1982.
- [24] Vo, T.V. & Balkey, K.R.: Risk–Based Inspection and Maintenance. In *Probabilistic Structural Mechanics Handbook: theory and industrial applications*, Sundararajan, C. (ed.), New York, Chapman and Hall, p388–415, 1995.
- [25] Wagner, H.M.: *Principles of Operations Research*. Second edition, Prentice–Hall, 1975
- [26] Zhao, Z., Haldar, A., Breen, F.L.: Fatigue–Reliability Updating through Inspections of Steel Bridges. In *Journal of Structural Engineering*, Vol.120 no.5 p.1624–1643, 1994.

Symbols, notations and definitions

α, β	shape and scale parameter resp. of the gamma and inverted gamma densities.
$\epsilon, \sigma_\epsilon, \epsilon_k$	ϵ is a normally distributed variable with mean 0 and standard deviation σ_ϵ representing the error in a thickness measurement. ϵ_k is the error term associated with the k -th inspection.
δ_k	difference between errors of two successive inspections: ϵ_k and ϵ_{k-1} .
$\lambda, \boldsymbol{\lambda}$	Lagrange multiplier and a vector of Lagrange multipliers respectively, see section 3.2.2 for details.
μ, μ_X	mean and mean of the random variable X respectively. In the chapter on the gamma deterioration process, μ is the uncertain mean of the gamma density.
σ, σ_X	standard deviation and standard deviation of the random variable X respectively. The variance of X is the square of the standard deviation: $\text{Var}(X) = \sigma_X^2$.
ν	coefficient of variation which fixes the standard deviation of the gamma process to the mean of this stochastic process.
ρ	replacement percentage representing the percentage of the safety margin m at which the component should be replaced. ρ is usually fixed by the corrosion allowance: $c_{max} = \rho m$.
$\Gamma(a)$	gamma function, see definition (3.14).
$\Gamma(a, x)$	incomplete gamma function, see definition (3.15).
$\Phi(x)$	standard normal distribution: $X \sim \mathcal{N}(0, 1)$, which means the distribution is symmetrically distributed around mean 0 with a standard deviation of 1.
\mathbf{x}	vector, where x_i represents the i -th element in the vector of length n : $\mathbf{x} = (x_1, x_2, \dots, x_n)$.

$\mathbf{u} \cdot \mathbf{v}$ *inner product* of vectors \mathbf{u} and \mathbf{v} ; this notation is referred to the vector notation and is equivalent to the matrix notation $\mathbf{u}^T \mathbf{v}$:

$$\mathbf{u}^T \mathbf{v} = [u_1 \ u_2 \ \cdots \ u_n] \begin{bmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{bmatrix} = u_1 v_1 + \cdots + u_n v_n,$$

where \mathbf{u}^T is the transpose of \mathbf{u} .
 $\|v\|$ *norm* of vector \mathbf{v} , where

$$\|v\| = \sqrt{\mathbf{v} \cdot \mathbf{v}} = \sqrt{v_1^2 + v_2^2 + \cdots + v_n^2}.$$

c_{max} *corrosion allowance*; this value is the maximum cumulative amount of thinning which may have occur for the component to operate safely. The actual failure level m is much higher than the corrosion allowance, therefore this value is used for the replacement level.

$c_i(\rho, \Delta k)$ costs incurred during unit time $(i-1, i]$ as a function of the percentage ρ and the inspection interval length Δk .

$C(\rho, \Delta k)$ expected average costs per time unit as a function of the replacement percentage ρ and the inspection interval length Δk . See definition (4.7).

d_k, D_k d_k is the measured increment between inspection k and $k-1$: $d_k = y_k - y_{k-1}$; D_k is the increment of the stochastic process $X(t)$ between times t_k and t_{k-1} : $D_k = X(t_k) - X(t_{k-1})$.

$L(\mathbf{x}, \boldsymbol{\lambda})$ Lagrangean function, see definition (3.8).

$f_X(x), F_X(x)$ probability density function (PDF) and cumulative distribution function (CDF) resp. of the random variable X .

$k, K, \Delta k$ $k = 1, 2, \dots, K$ is the inspection number and Δk is the length of the time interval between two inspections.

m safety margin determined by the limit state of the state functions (2.1) and (2.2). This is amount of degradation (corrosion or cracking) at which the component is assumed to fail.

p operating pressure.

s material strength.

th material thickness.

$t_k, \Delta t_k$ t_k is the time since the service start of the component at which the k -th inspection has been performed; Δt_k is the interval between the k -th inspection and the previous inspection: $\Delta t_k = t_k - t_{k-1}$.

$\mathbb{E}(X)$ expectation of the variable X , given by

$$\mathbb{E}(X) = \int_{x=-\infty}^{\infty} x f_X(x) dx.$$

$\Pr\{X \leq x\}$ probability of the variable X being less than or equal to the value x . It holds that $\Pr\{X \leq x\} = F_X(x)$.

$f'(x) = \frac{\partial f}{\partial x}$ the (partial) derivative of the function f to x . The second derivative is $f''(x)$.

$\nabla_{\mathbf{x}} f(\mathbf{x})$ gradient of function f to the vector \mathbf{x} :

$$\nabla_{\mathbf{x}} f(\mathbf{x}) = \left(\frac{\partial f}{\partial x_1}, \dots, \frac{\partial f}{\partial x_n} \right).$$

$p(x_i)$ the discrete probability of X : $p(x_i) = \Pr\{X = x_i\}$.

p_i probability of failure in unit time i : $p_i = F(t_i) - F(t_{i-1})$, where $F(t_i)$ is the probability of failure up to time t_i .

$l(\mu|x)$ likelihood of the variable μ given the measurement x .

$\pi(\mu)$ prior density for the variable μ .

$\pi(\mu|x)$ posterior density for the variable μ given x .

$X, X(t)$ random variable X has a distribution assigned to it (e.g. by a density $f_X(x)$). $X(t)$, $t \geq 0$ is the notation for the stochastic process.

$Y(t)$ stochastic process consisting of an underlying stochastic process $X(t)$ and an error ϵ : $Y(t) = X(t) + \epsilon$.

$\text{Ga}(x|\alpha, \beta)$ gamma density of X as a function of x with the shape parameter α and the scale parameter β . See definition (A.4).

$\text{Ig}(x|\alpha, \beta)$ same as above only for the inverted gamma density. See definition (4.5).

Index

- cobwebs, 29
- corrective replacement, 59
- correlation ratio, 28
- corrosion allowance, 64, 88, 106
- cumulant generating function, 82
- cycle, 65

- degradation states, 24
- directional simulation, 46

- FORM sensitivity measure, 28
- FOSM, 42

- gamma density, 81
- gamma function, 47

- imperfect inspection, 57
- incomplete gamma function, 47
- inner product, 38, 106

- Jacobian matrix, 78

- Lagrange multipliers, 39
- Lagrangean function, 39
- likelihood, 23
- limit state, 20

- moment generating function, 81
- Monte Carlo simulation, 46
- MVFOSM, 42

- Newton-Raphson scheme, 78
- norm, 38, 106

- PHAST, 19
- preventive replacement, 59
- prior, 23
- product moment correlation, 27

- PWHT, 21

- reliability index, 36
- RIMAP project, 13
- risk, 17

- safety index, 43
- simplified Lagrange multipliers, 40
- state function, 20
- stochastic process, 58
- Stress corrosion cracking, 20

- Taylor series, 37
- technical modules, 18
- Thinning, 19
- TMSF, 18
- transition temperature, 23

- unit normal vector, 38