

Durham E-Theses

Bayesian inspection planning for large industrial systems

Gavin Hardman

How to cite:

Hardman, Gavin (2007) Bayesian inspection planning for large industrial systems. Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/1998/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Bayesian inspection planning for large industrial systems

Gavin Hardman

Submitted for the degree of Doctor of Philosophy

May 2007

Abstract

The implementation of consistent and repeatable methods for inspection planning is a problem faced by a wide range of industries. The theory of Bayesian design problems provides a well established method for the treatment of inspection planning problems, but is often difficult to implement for large systems due to its associated computational burden. We develop a tractable Bayesian method for inspection planning. The use of Bayes linear methods in the place of traditional Bayesian techniques allows us to assess properties of proposed inspection designs with greater computational efficiency. This improvement in efficiency allows a greater range of designs to be assessed and the design space to be searched more effectively. We propose a utility based criterion for the identification of designs that offer improved prediction for future system properties. Designs with good typical performance are identified through utility maximisation. The viability of the method is demonstrated by application to an example based on data from a real industrial system.

Declaration

The work in this thesis is based on research carried out at the Statistics and Probability Group, Department of Mathematical Sciences, Durham University, England. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text.

Copyright © 2007 by Gavin Hardman.

The copyright of this thesis rests with the author. No quotation from it should be published in any format, including electronic and the Internet, without the author's prior written consent. All information derived from this thesis must be acknowledged appropriately.

Acknowledgements

This work has made possible by the support of a range of individuals and organisations. I'd like to thank the Department of Mathematical Sciences at the University of Durham for giving me the opportunity to do a PhD, and for providing a working environment which allowed for progress to be made (reasonably) swiftly. I'd also like to thank the Management Science department at the University of Strathclyde, for tolerating my somewhat distracted contribution to working there whilst writing up. The financial support provided by the EPSRC, and the CASE partners Shell is also gratefully acknowledged. I'd like to thank my supervisor, Michael Goldstein, for enduring my frequent bouts of incompetence and tardiness, and for providing an invaluable level of assistance and advice, without which the thesis could not have been completed. I'd also like to thank the Shell Statistical Consulting team, for supplying their time and expertise to help with the thesis, and in particular Phil Jonathan, whose positivity about the work was always refreshing. I'm obliged to acknowledge the support of my parents throughout the last three years, whose insistence on asking questions like, "Haven't you finished yet?", and, "When are you going to get a proper job?", has been strangely encouraging. Finally, I'd like to thank my friends from Durham, both from my undergraduate and postgraduate days, who have had the pleasure of experiencing my sullen disposition, and have all, in their way, contributed equally to keeping me well balanced. Although some have contributed more equally than others.

Contents

Abstract	i
Declaration	ii
Acknowledgements	iii
1 Inspection planning problems	1
1.1 The role of inspection and maintenance within industry	2
1.2 Current practice	3
1.2.1 Invasive inspection	4
1.2.2 Non-invasive inspection	4
1.2.3 Pigs [1]	6
1.2.4 Risk-based inspection	6
1.3 Problem structure	7
1.4 Literature review	9
1.4.1 Corrosion modelling literature	9
1.4.2 Bayesian design literature	11
1.5 Thesis outline	14
2 Modelling large industrial systems	16
2.1 Impact of corrosion damage	17
2.1.1 How corrosion occurs	18
2.1.2 Types of corrosion	19
2.1.3 Modelling corrosion behaviour	20
2.2 Modelling corrosion through wall thickness	21

2.2.1	Modelling wall thickness in a single component	22
2.2.2	Modelling wall thickness over many components	24
2.3	The observation process	26
2.3.1	Inspection designs	28
2.4	System model: Linear growth DLM	28
2.4.1	Linear growth DLM	29
2.4.2	Local variation	32
2.4.3	System Model	33
2.4.4	Graphical representation of corrosion surfaces	34
2.5	Observation model	36
2.6	Building covariance structures	39
2.6.1	Representing uncertainty	40
2.6.2	Defining weights $\lambda_{v_{h_k}}$	42
2.7	Model specification	43
2.7.1	Specifying covariances	43
2.7.2	Specifying initial conditions	44
2.7.3	Observation model specification	47
2.8	Simulation approach	47
2.8.1	A simulation realisation	48
3	Modelling large systems in practice: Site A data	53
3.1	Site A data	54
3.2	Factors relevant to corrosion	57
3.2.1	Substance	57
3.2.2	'Nearby' corrosion	58
3.2.3	Age	58
3.2.4	Flow	59
3.2.5	Environmental conditions	59
3.2.6	Material	60
3.3	Observation procedures	60
3.3.1	Ultrasonic testing	61
3.3.2	Radiographic testing	62

3.3.3	Modelling implications	63
3.4	Data analysis for the Site A data	64
3.4.1	Primary Data - Circuit A	65
3.4.2	Auxiliary data	69
3.4.3	Regression Analysis	75
3.5	Illustrative model, based on the Site A data	83
3.5.1	Parameter assessment	83
3.5.2	Local Variation	88
3.5.3	Spatial effects	93
4	Bayes linear updating for the corrosion model	98
4.1	Updating the Model	98
4.2	The Bayes linear approach	101
4.3	Adjusted assessments	104
4.3.1	Data for Circuit A	106
4.3.2	Estimating the necessary means and variance for updating	108
4.3.3	Updating for Circuit A	111
5	Evaluating the worth of an inspection design	115
5.1	Inspection design	116
5.2	Problem structure	118
5.2.1	Sources of loss	120
5.3	Design utility	121
5.3.1	The Bayes decision function	122
5.3.2	Calculating $E[L(d)]$	123
5.4	Approximating posterior failure probabilities	125
5.4.1	Assessment of posterior moments	126
5.4.2	Evaluating I_1	129
5.4.3	Evaluating I_2	131
5.5	Evaluating the expected loss criterion	133
5.5.1	Simulating future system behaviour	134
5.5.2	Specifying loss values	136

5.6	Designs	138
5.6.1	Lower baseline: the no inspection case	139
5.6.2	Upper baseline: the total inspection case	146
5.6.3	Partial inspection	150
6	Selecting inspection designs	158
6.1	Identifying good designs	159
6.1.1	Inspection blocks	159
6.2	Search routines	160
6.2.1	Stepwise Add	160
6.2.2	Stepwise Delete	162
6.2.3	Combining stepwise add and stepwise delete methods	164
6.2.4	Identifying d_0	165
6.3	Design selection for Site A example	167
6.3.1	Selection using the stepwise add algorithm	168
6.3.2	Selection using the stepwise delete algorithm	178
6.3.3	Selecting initial designs for Circuit A	182
6.4	Analysis of design performance	183
6.5	Choice of loss function	192
7	Discussion	198
7.1	Modelling	200
7.1.1	System Modelling	200
7.1.2	Observation Modelling	201
7.1.3	Modelling for the example	202
7.2	Updating	203
7.3	Design Selection	206
	Bibliography	208
A	Tables	216
B	Plots	223

Contents

ix

C Code

228

D Data

232

List of Figures

1.1	An articulated pig for use in pipe cleaning	6
2.1	Drawing of a cross section through an uncorroded and corroded pipe	22
2.2	Mapping a regular lattice to a pipe	34
2.3	Visual representation of system degradation	36
3.1	Representation of Ultrasonic Testing Procedure [9]	61
3.2	Representation of radiographic testing procedure [9]	62
3.3	Plan of Circuit A	66
3.4	Boxplots of Nominal Wall Thickness and Pipe Diameter variables for Circuit A	68
3.5	Distance (adjacency) matrix for modelled Circuit A components	68
3.6	Members of zero/nonzero subsets displayed by component type	72
3.7	Members of zero/nonzero subsets displayed by PMC level	73
3.8	Members of zero/nonzero subsets displayed by Material type	73
3.9	Members of zero/nonzero subsets displayed by Piping Schedule	74
3.10	χ^2 test values	74
3.11	Distribution of residuals	79
3.12	Residual plotted by Component type	80
3.13	Residuals plotted by PMC level	81
3.14	Residuals plotted against Nominal Wall Thickness	82
3.15	Distribution of WLR values of corroding components at $t = 0$	85
3.16	Mean wall loss rates for different variable levels, ± 2 standard deviations	87
3.17	Time series of observations for component 27	88
3.18	Distribution of residuals	91

3.19 Identifying appropriate standard deviation values for the local variation term	92
3.20 Plan of Circuit C	95
3.21 Adjacency matrices for Circuit C	96
3.22 Correlation values in Circuits B, C and E against distance between components	96
4.1 Observational data for $t = 5$, plus simulated expected range for observations	108
4.2 Observational data for $t = 9$, plus simulated expected range for observations	109
4.3 Observational data for $t = 10$, plus simulated expected range for observations	109
4.4 Observational data for $t = 11$, plus simulated expected range for observations	110
4.5 Observational data for $t = 12$, plus simulated expected range for observations	110
4.6 Updating beliefs for x_{12} given y_{d5}	111
4.7 Updating beliefs for x_{12} given y_{d9}	112
4.8 Updating beliefs for x_{12} given y_{d10}	112
4.9 Updating beliefs for x_{11} given y_{d11}	113
4.10 Updating beliefs for x_{12} given y_{d12}	113
4.11 Updating beliefs for α_t given y_h	114
5.1 Decision tree representation of component replacement problem . . .	119
5.2 Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 8 from Circuit A	126
5.3 Histograms of adjusted expectation values for component 8 from Circuit A based on simulated realisations of observations at 4 different time points	128
6.1 Correlations for the Circuit A system	169
6.2 Expected utility values for single half component inspection	170

6.3	Expected loss values for two half components inspection	174
6.4	Expected utility values for best designs using stepwise add	174
6.5	Detail from Figure 6.4 showing turning point at step 13	174
6.6	Proposed design using stepwise add routine	175
6.7	Variance Reduction achieved by proposed inspection plan	178
6.8	Expected utilities for search steps	180
6.9	Detail of Figure 6.8 showing increase in expected utility value	181
6.10	Expected loss values for each component (1)	181
6.11	Omega values for all components	182
6.12	Plot showing the behaviour the expected utility with respect to com- ponent failure loss	193
6.13	Plot showing the behaviour the expected utility of full inspection and best partial inspection with respect to component failure loss	194
6.14	Number of components involved in a design chosen via stepwise search for different failure losses	195
6.15	Selected components in designs for different failure losses	195
B.1	Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 16 from Circuit A	224
B.2	Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 24 from Circuit A	224
B.3	Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 32 from Circuit A	225
B.4	Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 40 from Circuit A	225
B.5	Histograms of adjusted expectation values for component 16 from Circuit A based on simulated realisations of observations at 4 different time points	226
B.6	Histograms of adjusted expectation values for component 24 from Circuit A based on simulated realisations of observations at 4 different time points	226

-
- B.7 Histograms of adjusted expectation values for component 32 from
Circuit A based on simulated realisations of observations at 4 different
time points 227
- B.8 Histograms of adjusted expectation values for component 40 from
Circuit A based on simulated realisations of observations at 4 different
time points 227

List of Tables

3.1	Component type summary for Circuit A	67
3.2	Categorical variable summary for Circuit A	69
3.3	Component type summary for Auxiliary Data	70
3.4	Categorical variable level summary for Auxiliary Data	71
3.5	Adjusted R^2 values for regression model fit for use in variable selection	76
3.6	Variables and their associated coefficients for second regression model	78
3.7	Adjusted R^2 values for second regression model fit	79
3.8	Mean Percentage wall loss at $t = 0$	84
3.9	Key to variable levels in Figure 3.16	86
3.10	Variable information for Circuit A	89
3.11	Variable information for Circuit A	90
3.12	Order of magnitude estimates for σ_{ζ}^2	93
3.13	WLR variance values by descriptive variable level	97
4.1	Model Time Steps	107
5.1	Loss Function	118
5.2	Model Time Steps	134
5.3	Losses incurred through component replacement	137
5.4	Losses associated with severity of component failure	138
5.5	Losses associated with component failure in Circuit A	138
5.6	Changes in variance values for different inspection designs	141
5.7	Changes in variance values for different inspection designs	142

5.8	Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information	144
5.9	Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information	145
5.10	Necessary elements of the expected loss calculation for each component in the full inspection case, excluding inspection cost information	151
5.11	Necessary elements of the expected loss calculation for each component in the full inspection case, excluding inspection cost information	152
5.12	Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information	156
5.13	Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information	157
6.1	Expectations and critical wall thickness values for Circuit A system based on data to time $t = 12$	171
6.2	Losses having inspecting half of component 3 for selected components	172
6.3	Losses having inspecting half of component 3 for selected components	173
6.4	Expected decision losses by component using the proposed design . .	176
6.5	Expected decision losses by component using the proposed design . .	177
6.6	Number of component failures in simulation experiment	185
6.7	Number of component failures in simulation experiment	186
6.8	Performance of the no inspection design	187
6.9	Performance of the recommended partial inspection design	188
6.10	Performance of the saturated inspection design	188
6.11	Effect of inspection design on decision making ability	189
6.12	Effect of inspection design on decision making ability	191
6.13	Test values for varying failure cost	193

A.1	Losses for the $c1$ update	217
A.2	Losses for the $c1$ update	218
A.3	Losses for the $c3$ update	219
A.4	Losses for the $c3$ update	220
A.5	Values for calculating the Ω criterion	221
A.6	Values for calculating the Ω criterion	222
D.1	Data for the circuit A example	233
D.2	Data for the circuit A example	234
D.3	Data for the circuit A example	235
D.4	Data for the circuit A example	236

Chapter 1

Inspection planning problems

The objective of this thesis is to develop a rational method for tackling problems encountered when planning inspections for systems consisting of many different, but related, components, performing a variety of functions. Systems of this kind occur in a range of areas. Examples of complex systems consisting of a number of subsystems carrying out different tasks can be found in fields such as computing, manufacturing, transport planning and the environmental sciences. We consider developing methods for applications in industry which provide a structured and defensible strategy for the planning of an inspection.

Inspection planning problems can be described as the problem of determining which parts of a system we should inspect to learn most about the overall system condition. The exact interpretation of 'system condition' will depend on the system under consideration, but will in general be some measurable property of the system - or the components within a system - that provides a quantification of the system's current level of performance. Identifying sets of points that allow us to learn more about the system improves the quality of the inspection; our aim is to produce a method which allows us to consistently select 'improved' design sets.

In this chapter we introduce the problems we address in the rest of the thesis. We describe the uses of inspection in industry in section 1.1. In section 1.2 we consider current inspection practice and inspection planning methods. The particular type of inspection problems we will be considering will be explained in section 1.3. Existing



inspection planning methods within the Bayesian literature are reviewed in section 1.4, and finally we outline the structure of the rest of the thesis in section 1.5.

1.1 The role of inspection and maintenance within industry

Inspection and maintenance plays an important role in most industries. Inspection is the means by which industries learn about the performance/condition of the system. Maintenance decisions are then made on the basis of the inspection results.

Maintenance is defined as actions taken to ensure the good performance/condition of the system [38]. Maintenance actions include repairing or replacing of some or all of the system, or less invasive treatments such as cleaning or lubricating components. The aim of ensuring systems are kept in a good condition is to maximise system efficiency and guard against potential safety hazards. Many industrial systems are expected to degrade over time. System degradation can be caused by mechanisms such as corrosion of the materials from which the system is constructed, build up of dirt or sediment within the system or weakening of system components through use. These factors, and others, can reduce the ability of the system to function properly. Systems that run at a reduced level of functionality are at best an inconvenience. More often, underperforming systems result in a serious loss of income, or can present a safety risk. The goal of maintenance is to ensure good system performance, and therefore maximise the system's productivity, or minimise the risk to the system's environment (including operators) introduced by system degradation.

Industrial inspection and maintenance is motivated by a range of considerations. In addition to the financial benefits of operating a system in good condition there are also safety factors. System owners have a legal responsibility to minimise both the individual and societal risks of operating a system. Individual risks are defined to be the chances of harm to the individual of a certain activity, and societal risk is treated as the risk of incidents which, for example, cause "widespread or large scale detriment or the occurrence of multiple fatalities in a single event"

[26]. The Health and Safety Executive (HSE) is responsible for the enforcement of Health and Safety legislation in the U.K., and for reviewing the risk assessments conducted by system owners. Inspection is not explicitly requested by the Government's Health and Safety regulations, but operators are required to produce a "major accident prevention policy document" in which they must specify methods for performance monitoring and identifying major hazards [25], a regulation that, in practice, makes inspection unavoidable. The motivation for complying with health and safety regulations is frequently the 'loss in public and employee confidence' that follows major system failures. This can greatly affect an organisation's profitability, and consequently provides a strong reason for avoiding system failure through effective inspection and maintenance.

Good quality inspection information accurately describes current system performance and allows the system owner to make informed decisions about necessary maintenance work. Effective inspection and maintenance allows the system to operate in a good condition and avoid system failure. However, the down side is that inspections are not necessarily cheap to implement or easy to plan. Poor quality inspection information can lead to bad maintenance decisions, resulting in expensive and ineffectual maintenance. Consequently, there is a need to plan inspections so that the information obtained allows us to make better maintenance decisions. We also have to be willing to trade off the quality of information we receive against the cost of the inspection. If inspection is cheap, we can inspect more thoroughly, however if inspecting the system is expensive, relative to the cost of repairing the system, thorough inspection will not be a sensible option.

1.2 Current practice

Many different types of inspection are available to industry. The particular choice of an inspection method is determined by a combination of the suitability of the method for the system in question and the cost of implementing inspection using this technique. Inspection methods can be split into two major groups, invasive

inspection techniques and non-invasive techniques. Invasive inspection involves the system being shut down and inspected internally by the inspectors. Non-invasive inspection (NII) or non-destructive testing (NDT) methods do not necessarily require system shutdown, and inspection is conducted from outside the system.

1.2.1 Invasive inspection

Invasive inspection techniques require internal access to the system, which necessitates system shutdown. Invasive inspection techniques are, in general, thorough and engender high customer confidence. The techniques themselves are frequently quite simple, with visual inspection of the system (or enhanced images of the system), being primarily favoured [40].

The perception of invasive techniques is that they offer a more ‘complete’ picture of current system behaviour, as shutdown enables almost total inspection and all ‘significant’ problems will be detected. The disadvantages of invasive techniques are held to be the length of time required to conduct the inspection (and the associated loss of income through reduced production) and the potential safety risks to inspectors caused by the disassembly process and the requirement to work in confined areas.

However, as invasive inspection is frequently dependent on the ability of the inspectors to identify system defects, it may not be as reliable as it is perceived to be. Information gathered through invasive inspection techniques is frequently qualitative rather than quantitative, which may not be sufficient to satisfy the legal system monitoring requirements placed on the system owner. Consequently, alternative techniques, which offer cheaper, and possibly more defensible, solutions, have been sought.

1.2.2 Non-invasive inspection

Non-invasive inspections are not, in general, as extensive as invasive inspections.

NII techniques involve sampling sections of the system, as total inspection using most NII techniques is infeasible due to either the time or cost involved in inspection. This clearly creates the additional problems of determining where to sample, and how the sampled locations relate to the unsampled locations.

There are many different NII techniques, although the majority have the same format. [67] identifies the following five stages in any NII technique [72]:

1. A suitable form and distribution of energy is supplied to the test object from an external source.
2. The supplied energy distribution is modified within the test object as a result of discontinuities, flaws and other variations in material properties.
3. The modifications to the energy distribution are detected by a sensitive detector.
4. The energy measurement from the detector are recorded in a form useful for interpretation.
5. The recorded values are interpreted and used to assess the current system state.

The objective of such procedures is to learn about structural differences in the system. Different inspection techniques are better at detecting different types of defects or damage. We discuss specific inspection techniques relevant to our examples in section 3.3. Commonly used techniques include [9]:

- Ultrasonic techniques - use of sound waves as the energy source.
- Radiographic techniques - use of X-rays as the energy source.
- Eddy current techniques - use of electromagnetic fields to detect defects.
- Penetrant techniques - use of dye (or similar) to seep into the surface, which is then absorbed by defects.
- Magnetic particle techniques - use of magnetic fields to highlight defects.



Figure 1.1: An articulated pig for use in pipe cleaning

1.2.3 Pigs [1]

Pigs (or pipeline inspection gauges - a retrospectively defined acronym), are tools used by industries that rely on the use of pipelines, such as the petrochemical and natural gas industries. Pigs have a variety of uses, including pipeline cleaning and fluid separation functions. An example of a typical pig is shown in figure 1.1.

Pigs are designed to fit inside a pipeline and almost fill the pipe. They are moved along the pipeline by the pressure built up by the pipe contents. [31] discusses the use of 'intelligent' or 'smart' pigs in inspection. Pigs can be fitted with sensors and other forms of inspection tools to detect internal pipe damage and can record the condition of the pipeline as they are moved along. This is an example of a compromise between invasive and non-invasive inspection. The pig is monitoring the entire pipe surface, but there is no need to shutdown the pipeline. However, this method is expensive to implement and is therefore not widely considered.

1.2.4 Risk-based inspection

The choice of technique is dependent on the type of defect about which the system owner is most worried and, more often, the cost of implementing the technique. Risk-based inspection (RBI) methods are generally favoured for the determination of inspection locations. There is an extensive literature surrounding techniques in

risk-based inspection, which develops RBI methods for a variety of applications. The RBI approach is summarised by [23], and some recommendations for best practice in RBI methods are given by [71].

RBI has traditionally treated inspection planning as a univariate problem. Under this approach the risks associated with component failure are calculated and the components that pose the greatest ‘threat’ are inspected. Risk is usually defined to be the probability of component failure multiplied by the cost of the consequences of component failure. This quantity is calculated for each component within a system, allowing a ‘risk ranking’ of components to be produced. Components are then selected for inspection based on their risk ranking. Baker and Descamps [4] propose considering not just the risk, but also the probability of component failure and the cost of component failure to create a combined ranking that assigns greater importance to components that are likely to fail or have greater financial consequences if they fail.

Recent work in RBI studies ([56]) has treated the calculation of risks as a joint problem, taking into account the influence that components have on each other when evaluating the failure probabilities. However, we do not know how widely this approach is currently being implemented.

1.3 Problem structure

Inspection planning problems occur across a wide range of industries. For the purposes of this thesis we will restrict to a specific subset of problems, that of monitoring corrosion damage within the petrochemical industry. In considering this class of problems we illustrate a large number of issues that are of relevance to a wider class of inspection planning problems. However, the methods developed within this thesis are constructed specifically with corrosion applications in mind.

Throughout this thesis we will use the following terminology to describe our problems:

Component - A region (of variable size), that is consistent with respect to its

physical properties and corrosion behaviour.

System - An organised set of components that acts to perform a function.

Our definition of a component is intended to be sufficiently flexible to allow for different parts of a system with the same structure and corrosion behaviour to be considered as a single component even if they are in different physical locations. The intention is to allow components to be defined by their corrosion properties. In practice, we treat components as contiguous regions defined by their physical properties due to the difficulties in distinguishing between corrosion processes in operating systems. In addition to systems and components, we may also consider *subsystems*. Subsystems are defined as smaller subsets of system components (containing at least 2 components) that satisfy our definition of a system. By requiring that our systems are both organised and functional we ensure that not every set of components will form either a system or a subsystem.

The petrochemical industry offers a particularly challenging set of inspection problems. The systems used within the petrochemical industry are often large and complex, with many subsystems performing a variety of different functions. There are also a range of different installation types used by the petrochemical industry, including offshore platforms, refineries and other processing plants. All of these carry out complex operations involving large systems, the condition of which needs to be monitored. There is a clear need for good inspection practice within the petrochemical industry motivated by the consequences of system failure. Petrochemical system failure can have extremely severe consequences, therefore the careful monitoring of these systems to ensure successful detection of potential problems is an important consideration for companies and regulators within this industry.

We will focus on the problems of inspecting pipework systems for corrosion damage. Pipework systems are an important part of the petrochemical industry, providing the means of transporting hydrocarbons through the extraction and processing procedures. The pipework used is generally metallic and is therefore subject to corrosion, which can weaken the pipe. Corrosion is a time dependent process, in which the damage progressively increases over time, and can therefore be addressed if detected sufficiently early. Corrosion can spread over an entire surface, or can

remain localised in specific points. Therefore we need to consider the impact of both general and localised corrosion behaviour.

Our objective is to develop a procedure for identifying locations within a corroding system that will tell us most about the corrosion behaviour throughout the entire system, and thereby identify components in need of repair or other appropriate maintenance. To achieve this we will need to model system-wide corrosion behaviour in such a way that the relationship between corrosion levels at different locations is captured, and subsequently establish a consistent and repeatable method for identifying which points of a system are informative about corrosion behaviour.

1.4 Literature review

In this section we review the existing literature within areas related to inspection. We consider these to be the fields of corrosion modelling and Bayesian design, and we review each of these in turn.

1.4.1 Corrosion modelling literature

One key difficulty of inspection planning problems is that they consist of two substantial problems, that of modelling the degradation of the system, and then using the degradation model as a basis for design selection. [20] consider this to be the fundamental structure of maintenance problems, and review existing ‘deterioration models’. They classify existing models into failure rate, reliability index, Markov and renewal models.

The available engineering literature details some of the physical models used to describe corrosion behaviour. [49] review models for water pipelines damage. Water pipelines are frequently buried and therefore subject to slightly different degradation mechanisms than the exposed systems we will consider, but the review offers an insight into typical water pipeline failure behaviour. A number of papers address specific types of corrosion behaviour: [50] offer a detailed description of corrosion

fatigue modelling in unfavourable conditions; [43] and [29] address CO₂ driven corrosion; [70] tackle the problem of pitting corrosion; [65] propose a model for crevice corrosion and [41] deals with the problems of modelling marine corrosion damage. This illustrates the potential range of corrosion mechanisms that could affect a system.

A relatively large amount of the existing literature for corrosion modelling considers reducing the problem to a univariate situation, by treating components as independent. For large systems, such as those found in the petrochemical industry, this is a pragmatic assumption that allows the modeller to handle potentially intractable modelling problems. However, this is not usually a realistic assumption. [37], [36] consider the use of a spatio-temporal dynamic linear model to model the corrosion behaviour over an entire surface of a system (furnace). [57] and [21] consider partial inspection problems using models that relate components to each through a covariance structure.

There is a range of different system properties that have been modelled. One of the most commonly modelled quantities is the time to component failure. [61] propose using a conjugate gamma process approach to model time to failure, which can be easily updated given inspection data. [21] treat the system problem as one of modelling a set of correlated binary variables, which are either operating correctly or have failed, and conditions can be set on the number of components required to be operating for the system to be operating successfully. Both models require the specification of failure rates as initial parameters. [6], [62] and [33] choose to model 'deterioration' - assessed as the difference between the stresses (or loads) applied to the system and the remaining 'strength' of the system based on its material properties.

Alternatively, related quantities for the corroding surface can be modelled directly. The spatial approach of [37] models wall thickness behaviour, and [48] model the changing mass of the corroding components. The conjugate gamma process approach is also adapted to measure defect size in [63]. In all of these cases it is possible to relate the modelled quantity to component failure by a deterministic rule concerning minimum acceptable standards. Arguably, these approaches have more

relevance for inspection models than the time to failure methods.

A common approach in the corrosion modelling literature is to treat a system as a set of unrelated components, in which corrosion properties remain approximately constant. [48] advocate using a dynamic model that reflects changes in the rate of corrosion which are consistent with the properties of the type of corrosion occurring at a particular site. The use of a dynamic linear model in [37] also allows for control of variation in the corrosion rate. [53] proposes the use of a fuzzy logic based method to account for the uncertainty in the choice of corrosion model and [73] suggest a method for modelling multiple site crack corrosion using different competing models.

The literature contains a variety of proposed corrosion modelling methodologies. Many of these do not consider the situation of correlated components within systems, but this is usually as a consequence of difficulties this creates in solving the subsequent design problem. The geostatistics literature contains examples of spatial modelling with a view to design. [14] and [5] both present methods for selecting designs for learning about the behaviour of random spatial fields. However both authors comment on the computational intensity of their suggested approaches.

1.4.2 Bayesian design literature

There is also a well developed literature for the Bayesian approach to maintenance optimisation. A number of articles focus on determining optimal time intervals for inspection, such as [16], [55] and [28], in which the aim is to identify the optimal point to inspect the system so that the useful life of all components is maximised.

The risk-based inspection (RBI) techniques - favoured in industry for inspection planning - lend themselves naturally to use with Bayesian deterioration models. In RBI procedures, the aim is to identify the components with the greatest risk of failure (overviews of the RBI methodology are provided by [23], [44]), and use the risk ranking as a guide to which components are in need of inspection. All RBI procedures require a model for the probability of failure, which is re-stated after the observation of the system. Bayesian methods provide a coherent framework for the updating of failure probabilities given new inspection data, as illustrated by [28],

[61] and [51].

In a number of industries, component failure probabilities are difficult to estimate due to a lack of available historical data. This can be because the consequences of failure are deemed to be so severe that extremely conservative maintenance policies are adopted and failure data is therefore scarce. Alternatively, this can be a result of the use of a new design of component, for which there is no historical data. In such cases it is necessary to employ expert judgement based methods to establish beliefs about failure probabilities and component failure. Such methods are also readily implemented in Bayesian models. [53] and [47] contain examples about the incorporation of Bayesian expert judgement into design selection problems.

The theoretical basis for Bayesian design problems is well established. An extensive review of Bayesian experimental design literature is presented by [10], who comment on the lack of applications to support the theory. The benefits of a Bayesian approach to design are discussed in [46], which concludes that Bayesian methods allow for model uncertainties to be accounted for more straightforwardly. The principles of Bayesian design are summarised in [10], in which the expected utility of a design is maximised over all values of the data that could potentially be observed. The optimal design is the design that returns the greatest expected utility value. This approach is based on the work of Raiffa and Schlaifer (1961) and Lindley (1972) [34], who suggested the use of Bayesian decision theory as a means of optimisation.

Bayesian decision theory is widely used as a means of solving Bayesian design problems. The basis for the decision is the utility function, which must be carefully constructed to reflect our design aims. The importance of carefully constructed utility functions is emphasised by [10], [60] and [47]. [14] illustrate how the utility criterion changes as the inspection goals change from prediction to estimation.

The decision theoretic approach has been applied to a number of different examples recently. [11] demonstrate the use of Bayesian decision theory in selecting inspection plans for the reliability of aircraft wings, [66] applies the method to the problem of building metamodels of complex computer simulation models (i.e. identifying and selecting influential components). Applications in clinical trials, [32], and the placement of security sensors, [59], have also been considered. [28] consid-

ers applying the method to corrosion within pipework, but restricts to the problem of recommending component replacement times, rather than inspection locations.

A frequently occurring problem in Bayesian inspection/experimental design is that of computational burden. Evaluating the expected utility of a design over all potential observation sets is often a problem that has to be solved using numerical or simulation methods. [10], [14], and [39] all comment on the computational intensity of Bayesian design methods. [15] develop a Bayesian design method for regression variable selection, but comment on its intractability for non-linear problems. Improvements in available computational power over recent years have reduced this problem, but for large complex systems of the type in which we are interested, the computational burden of traditional approaches to Bayesian design is still a very large problem.

The computational difficulties of Bayesian decision problems are also discussed in the RBI literature. [55] develops a tractable method for assessing reasonable inspection intervals by restricting to simple inspection policies. The assumptions made in adopting the approach are tested in [16]. [58] proposes an interpolation approach to reduce the computational burden of a full Bayesian design procedure. It is suggested that the ‘value’ of a number of generic designs can be evaluated prior to the design process, and that the value of specific designs can be established by interpolating between the values of the generic designs. [52] suggests the use of maximum entropy sampling as a means of identifying sensible designs.

Maximum entropy sampling is appropriate whenever the entropy of the observations does not functionally depend on the design. This makes it useful for linear and variable selection problems, but impractical for spatial problems. [3] describes how entropy methods can be used to identify optimal designs for spatio-temporal problems, when we are working with a multivariate Gaussian field; however, these methods are still reasonably computationally intense.

[17] consider the optimal repair problem from an adequate safety/optimal cost perspective. A Bayesian decision theoretic approach is developed in which inspection is treated as beyond the modeller’s control. [30] and [7] illustrate the use of cost-benefit analysis to assess inspection worth in the software industry, but the issues

addressed are slightly different to those found in the petrochemical industry.

[37], [36] and [35] develop a method for computationally efficient inspection planning using Bayes linear methods for applications in industry. Bayes linear methods have also been applied, in theory, to other design problems in [13], [19] and [54], where methodology offers a tractable approach to complex design problems.

[56] deals with the problem of planning inspection assuming correlated components. The generic RBI method of [58] is adopted to reduce computational load and correlations are included between identified corrosion ‘hotspots’ via a covariance matrix which accounts for the ‘similarity’ of the hotspots. Methods for building covariance structures for large, complicated systems are discussed in [18].

1.5 Thesis outline

In chapter 2 we outline our proposed method for modelling corrosion in large systems, using a dynamic linear model to take account of global trends and a separate sub-model to describe local behaviour. In chapter 3 we illustrate how we would apply the model to a real system with an example taken from an industrial site. The model is developed for illustrative purposes, so certain characteristics of the data have been emphasised to provide a more effective illustration. This example will be used throughout the thesis to illustrate both the modelling and inspection planning procedures in practice, and the initial modelling, which is discussed in chapter 3, forms the basis for the simulations used to assess design performance in later chapters. Updating is considered in chapter 4. With a view to proposing a tractable design selection procedure we make use of the Bayes linear approach to updating; the motivation behind this choice and the details of the Bayes linear method are discussed here. The use of Bayes linear methods allows us to evaluate properties of the design more rapidly than a fully Bayesian update would permit, leading to a tractable design methodology, even for large systems. In chapters 5 and 6 we deal with the design selection process. Chapter 5 considers the problems of identifying a suitable criterion for choosing between designs and details how we can

calculate a utility based criterion for design ‘worth’. Chapter 6 discusses how we can then apply this criterion to the problem of selecting sensible inspection designs. The criterion balances inspection cost against the benefit of the inspection in terms of information gained. Therefore a design that performs well according to our criterion will improve our understanding of the system, and our decision making ability, so that the savings made from making better decisions outstrip the cost of performing the inspection. This process is then illustrated using the example from chapter 3. We conclude the thesis with a review of the completed work and consider the scope for further research in chapter 7.

Chapter 2

Modelling large industrial systems

In this section we describe the modelling approach that will be used throughout. As stated in chapter 1 we will be concentrating on the problem of corrosion damage to large pipework systems. We develop the model with these systems and corrosion in mind, but it is possible to generalise the method to handle many situations in which we wish to model a multiple component system evolving in time.

Ultimately, we will be interested in planning inspections for large industrial systems. However, informed inspection planning requires an understanding of the system we are planning to inspect and its current operating state. Modelling the system gives us a way of representing what we believe to be happening within the system. We can then use data from previous inspections to update these beliefs - in line with standard Bayesian practice.

The model provides the mechanism for us to interpret our observations in terms of system behaviour. It will also be used to describe our beliefs about future system behaviour, which is a crucial part of inspection planning. Therefore the model should accurately characterise the corrosion process in which we are interested. However, we are concentrating on planning inspections for systems with many different components, each with distinct properties, so any model we consider will have to retain sufficient flexibility to account for these different operating environments. We must also consider the scale of the systems we wish to model and remember that, in addition to building a model which adequately characterises component corrosion

behaviour, the calculations required for inspection planning should remain computationally tractable for large numbers of components. The data associated with such systems can often consist of complicated functions of the corrosion state, themselves requiring additional modelling which adds to the computational load of the inspection planning problems.

This chapter describes how we construct such a model. Section 2.1 details the necessary properties of any corrosion damage model by giving an outline description of the chemical process of corrosion and expert judgement on typical corrosion behaviour. Section 2.2 gives a broad description of the model structure and sets up our notation. Section 2.3 presents the need for a separate measurement model, which is described in detail in section 2.5. Section 2.4 explains in detail how each model term behaves with respect to evolution in time and interaction between components. Defining the model covariance structures is described in section 2.6 and all other aspects of model specification are discussed in section 2.7. Section 2.8 explains how we can use a simulation approach to obtain full covariance structures between any model term and our observations for any time point of interest.

2.1 Impact of corrosion damage

Corrosion is a process mostly affecting metallic materials in which the material reacts with its environment to produce a new material. In general, this new material is less effective for its intended function than the original uncorroded material. For the situation of pipeline corrosion, we will consider that each of our components is initially constructed from a known material which is subject to corrosion. A corroded component will usually not function as well as an uncorroded component, and in extreme cases corrosion can even be a cause of component failure. Component failure is an expensive consequence of corrosion, resulting in unwanted expenditure on emergency repair work and loss of income due to related parts of the system not functioning. More commonly, corrosion damage is repaired before failure occurs, but a loss of income can still be incurred due to the inefficient operation of corroded

components. The loss of income, and potential threat posed by the release of chemicals from damaged components provides the motivation for tackling the corrosion problem. The scale of the problem is highlighted by [2], which estimated the annual cost of corrosion effects in the U.S.A. in 1998 to be \$ 276 billion. Furthermore, [45] claim that 13% of all hydrocarbon releases are due to corrosion/erosion.

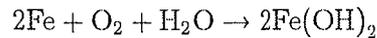
2.1.1 How corrosion occurs

Corrosion is a chemical process driven by environmental factors causing a (generally unwanted) change in the nature of the material being acted on. [42] states that 5 steps are involved in the corrosion of metallic materials:

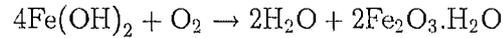
1. The metal has to give up electrons (and thus create positively charged ions) in order to begin the process.
2. Ions require a medium in which to move (usually water).
3. There must be a supply of oxygen present.
4. The reaction requires a driving force.
5. A new material is formed. This may react further with the environment.

In most situations the reaction driving force will be something very simple, such as the presence of oxygen ions in water. To illustrate these 5 steps we will consider the example of iron rusting. This is a very familiar example of a corrosive process in which iron reacts with water and oxygen in the environment to form (hydrated) iron oxide. The process begins with the iron (Fe) atoms giving up electrons to the surrounding water and therefore becoming positively charged ions ($\text{Fe}^{2/3+}$). Rusting takes place in wet conditions, so the transfer medium for the ions (step 2) is water. Water contains a lot of dissolved oxygen, so step 3 is also satisfied. In this reaction, the driving force is provided by the oxygen atoms (O_2) (in the water) taking on the released electrons to form negatively charged oxygen ions. The positively charged iron ions then react with the negatively charged oxygen ions and the water (step 5)

as follows:



forming iron hydroxide ($\text{Fe}(\text{OH})_2$). In this case the iron hydroxide then reacts again with the oxygen:



to give water and hydrated iron oxide. The hydrated iron oxide is the flaky brown substance seen in a typical rusting process.

Rusting is typical of many corrosion processes in that it follows the five defined steps. This gives us an idea of the type of reaction which is typically taking place. We now move on to considering more typical forms of corrosion behaviour.

2.1.2 Types of corrosion

There are many different types of corrosion. Two important classes are uniform and localised corrosion. The main difference between these types of corrosion are the extent to which they affect the surface upon which they are acting. Uniform corrosion occurs over the majority of the surface and consequently is usually easy to detect. The rate of uniform corrosion is usually stable and in isolation it is usually easy to handle by including sufficient redundancy into the system to allow for typical uniform corrosion. However, when occurring in conjunction with forms of localised corrosion, the impact of uniform corrosion can be much more damaging, providing an already weakened surface for the localised corrosion to act upon, and also camouflaging the effects of localised corrosion during inspection.

Localised corrosion only occurs at specific points on the component surface and could have one of several different causes. Localised corrosion behaviour such as pitting, crevice or intergranular corrosion is caused by corrosion attacking weak points of the surface, such as those where protective coatings have been worn away, points at which components are joined or simply defects in the component surface. Problems such as galvanic corrosion and selective attack occur as a result of using alloys in which one constituent part is more susceptible to corrosion than others,

resulting in the metal being structurally weakened. The function of a component is also important in determining localised corrosion behaviour. Fatigue, stress corrosion cracking and fretting corrosion are all brought about by strains repeatedly put on the material of the component, either by repetitive motion or by static tensile stress. These strains can cause weak points in the surface of the material, resulting in increased susceptibility to corrosion. Localised corrosion is much harder to detect than uniform corrosion and can occur at much higher rates. It therefore poses a much greater threat to components and accounts for the majority of component failures.

Both uniform and localised corrosion will be affected by environmental factors such as location, temperature and pressure. Also, we are dealing with systems in which there is a directed flow. This can create further corrosion damage by the abrasive action of substances flowing through the system and also the further effect of corroded materials being swept along by the flow.

2.1.3 Modelling corrosion behaviour

We know corrosion is a process that causes component degradation. We also know that these changes due to corrosion are irreversible, i.e. once part of the component surface has reacted it does not become part of the component again at some future time point. So any model we choose must be strictly monotonic.

The perception amongst corrosion experts is that corrosion rates are approximately constant once corrosion has initiated (see [42], [40], [64]). This applies to both uniform and localised corrosion although the uncertainty associated with this constant rate is much higher for localised corrosion cases. Due to the different natures of uniform and localised corrosion it is reasonable to propose a model structure that has separate terms for each type of corrosion. The consensus of expert opinion in favour of an approximately constant corrosion rate also suggests our model should incorporate this feature.

Factors such as material, function and environmental conditions can influence corrosion behaviour (particularly localised corrosion behaviour). Our model should

reflect this knowledge and allow us to build in any auxiliary information at our disposal. Finally, we want to retain a sufficiently simple structure so that computational tractability is ensured for systems with large numbers of components. To summarise, we require a model that:

1. Allows us to use auxiliary information to distinguish between components.
2. Allows us to separate global and local corrosion effects.
3. Gives an approximately constant corrosion rate.
4. Allows us to correlate components with similar locations/characteristics.
5. Remains computationally tractable for high dimensions.

2.2 Modelling corrosion through wall thickness

Modelling the underlying process, in this case corrosion, is an important stage in solving the inspection planning problem. The model will be used as the basis for our predictions for future system behaviour and should therefore be carefully considered. In this section we will outline our modelling strategy. Further details are given in sections 2.3 - 2.5. Following the definitions of system and component given in section 1.3 we will start by considering how to model corrosion within a single component and then move on to the many component (system) problem.

We must establish which quantity we are modelling. We are interested in component (and consequently system) integrity with respect to corrosion. Corrosion is a difficult quantity to measure directly, so we will have to use a related quantity which tells something about the effect of any corrosion processes taking place within the system. One such quantity is *wall thickness*, which is a commonly measured quantity in the inspections performed by our industrial collaborators.

The interpretation of *wall thickness* for a pipe is illustrated by Figure 2.1. We can relate wall thickness to corrosion by attributing decreases in wall thickness (wall loss) to degradation caused by corrosion.

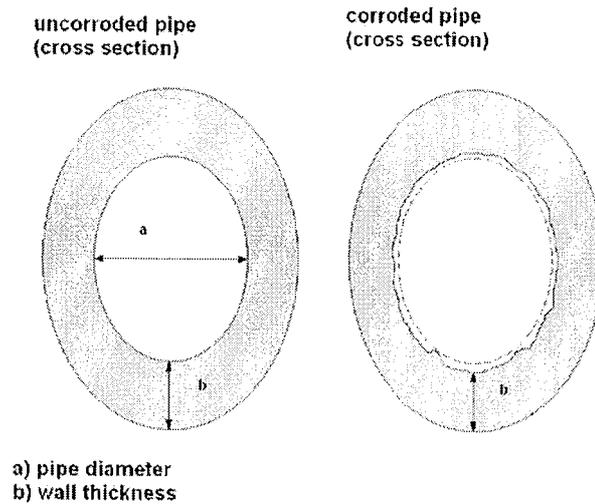


Figure 2.1: Drawing of a cross section through an uncorroded and corroded pipe

In section 2.1 we stated that corrosion is an irreversible process. In terms of wall thickness this means that only wall loss can occur, i.e. our model should not allow wall thickness to increase. Given that expert judgement informs us corrosion is approximately constant at typical inspection frequency, we will therefore choose to model wall thickness as having an approximately linear decreasing trend.

2.2.1 Modelling wall thickness in a single component

Each component is characterised by beliefs about the properties of the corrosion process occurring at sites within the component and auxiliary information about the component. The exact nature of the auxiliary information will depend on the specific data set, but we can expect to have some knowledge of the design and function of each modelled component. This information can be used to shape our beliefs about the type and rate of the corrosion process(es) within the component.

We will model an individual component as a grid of s_c locations. The size of s_c will be determined by the characteristics of the component. The size, design and function of a component could all play a role in influencing the extent of local

variation seen in a component - and therefore the size of the grid necessary to model this variability. We refer to points on the grid as 'locations'. For each location, l , at time t , we assume there is an underlying true wall thickness value, denoted u_{lt} . We choose to model the u_{lt} values using two terms, the global term, x_t and the local term, r_{lt} . That is, $u_{lt} = x_t + r_{lt}$. Note that the global term (x_t) does not depend on location. A corrosion based interpretation would be to view x_t as modelling global corrosion and r_{lt} as modelling localised corrosion.

Separating the model into global and local terms also allows us to distinguish (where possible) between the different corrosion types and it also provides the flexibility to treat different types of auxiliary information differently. Characteristics such as component function, material or design could reasonably be expected to affect all areas of the component in a similar fashion and should therefore be modelled by the global term, x_t . More dynamic characteristics such as temperature, weathering or system flow effects could be expected to have a less consistent influence on the wall loss and be more susceptible to local fluctuations, thereby making them more suitable to model using the local term r_{lt} . However, such ideas would have to be confirmed either by expert judgement or data before being implemented.

In principle, the x_t and r_{lt} terms can take any form. Indeed, as much or as little structure as is felt to be necessary or appropriate can be built into each term. We have chosen to concentrate most of the structure into the global term, x_t , which we will be modelling with a linear trend dynamic linear model (DLM). Full details of the construction and evolution of this model are given in section 2.4, but the key properties of the model are described by the coupled equations:

$$x_t = x_{t-1} + \alpha_t + \varepsilon_{xt} \quad (2.1)$$

$$\alpha_t = \alpha_{t-1} + \varepsilon_{\alpha t} \quad (2.2)$$

α_t is the trend term and $\varepsilon_{xt}, \varepsilon_{\alpha t}$ are independent deviations from the global term value. The deviations associated with the x_t (ε_{xt}) and α_t ($\varepsilon_{\alpha t}$) terms are independent of each other and are drawn from separate populations with Normal distributions:

$$\varepsilon_{xt} \sim N(0, \sigma_{\varepsilon_{xt}}^2)$$

$$\varepsilon_{\alpha t} \sim N(0, \sigma_{\varepsilon_{\alpha t}}^2)$$

We will use the notation μ_{v_b} to denote the expectation of a random variable v_b throughout, where the variable will be written including any relevant subscripts, b . Similarly, all variances will be labelled as $\sigma_{v_b}^2$ and all covariance matrices as Σ_{v_b} .

The r_{lt} values are chosen to have a simple form and are modelled as:

$$r_{lt} = r_{l,t-1} + \zeta_{lt} \quad (2.3)$$

where,

$$\zeta_{lt} \sim N(0, \sigma_{\zeta_{lt}}^2), \quad (2.4)$$

so in principle each location could have its own specific distribution, although in practice this may be difficult to specify. In effect this means the r_{lt} values are simply the sum of t draws from Normal populations with mean 0 and variance $\sigma_{\zeta_{lt}}^2$. By choosing to build a trend only into the global term we are emphasising the impact of global factor over local ones. More emphasis could be placed on localised corrosion by building more structure into the local term, possibly to include its own decreasing trend and also an initiation factor. Parameters for the local trend term should be chosen to ensure there is no possibility of increasing wall thickness.

2.2.2 Modelling wall thickness over many components

We want to model systems consisting of large numbers of components. In this section we describe how we move from the single component model to the many component model. The system is to be modelled as n components. We use the index c to distinguish between these. For each component we have a grid of s_c locations. The underlying true wall thickness values are now given by:

$$u_{lct} = x_{ct} + r_{lct}. \quad (2.5)$$

The model structure corresponds directly to the single component case. The global term (x_{ct}) depends only on component and time and not on location, meaning we have a separate trend for each component. Treating the the global terms as n distinct linear trend DLMS would, in general, allow us to simplify the inspection design problem. However, this would not allow us to model interactions between

components. We instead consider each global term (x_{ct} for $c = 1, \dots, n$) as an element of a multivariate DLM [24]. The multivariate DLM provides a natural framework for modelling correlated components as they evolve in time. The only difference this makes to modelling the global term is that the coupled evolution equations (2.1) and (2.2) are now written in terms of n -vectors in which the c th element of each vector corresponds to the global term for the c th component:

$$\underline{x}_t = \underline{x}_{t-1} + \underline{\alpha}_t + \underline{\varepsilon}_{xt} \quad (2.6)$$

$$\underline{\alpha}_t = \underline{\alpha}_{t-1} + \underline{\varepsilon}_{\alpha t} \quad (2.7)$$

$\underline{\varepsilon}_{xt}, \underline{\varepsilon}_{\alpha t}$ are now draws from a multivariate Normal population with distribution:

$$\underline{\varepsilon}_{xt} \sim N(\underline{0}, \Sigma_{\varepsilon_{xt}})$$

$$\underline{\varepsilon}_{\alpha t} \sim N(\underline{0}, \Sigma_{\varepsilon_{\alpha t}})$$

Correlation is introduced via the deviations; these will be correlated across components at each time point. Therefore the covariance matrices $\Sigma_{\varepsilon_{\alpha t}}$ are the means by which our beliefs about the system correlation structure are entered into the model.

We continue to use minimal structure on the local term r_{lct} . Following the approach adopted for a single component, we have the value of r_{lct} given by equations (2.3) and (2.4). We are therefore assuming the extent of local variation is specific to the component and location being modelled. We also assume there is no correlation between the local variation in different components, i.e. that any between component correlation is accounted for by the global term. We can construct $\Sigma_{\zeta_{ct}}$ as the $s_c \times s_c$ matrix $\Sigma_{\zeta_{ct}} = (\sigma_{\zeta_{l'l'ct}})_{l=1}^{s_c}{}_{l'=1}^{s_c}$, where $\sigma_{\zeta_{l'l'ct}}$ is the covariance between locations l and l' . We would require an $s_c \times s_c$ matrix for each component c .

Our overall aim is to produce an efficient inspection design methodology which will remain tractable for very large systems. The model sketched here offers a structure that is both sufficiently flexible to characterise different aspects of the problem in question, yet sufficiently simple to scale up to many components. The model is not proposed as a definitive corrosion model, but as a modelling framework for tackling the large system problem with a view to inspection planning.

2.3 The observation process

In addition to building a model which describes the nature of the corrosion process taking place within a system we also need to build models to describe the observation process.

In any physical process there is a gap between what can be observed and the actual system state. In general, measurements will be made with error. This introduces a level of uncertainty into the inspection process that has to be modelled. This uncertainty stems from not knowing exactly what our observations are telling us about the true system behaviour.

We treat the observation process as a distinct part of the model. We do this because the observation process is different to the corrosion process. The source of the uncertainty being modelled is associated with the measurement methods used. Although the choice of measurement method may be influenced by the characteristics of the system, it is not determined by the extent of corrosion damage within the system. This approach allows us to model many different observation procedures without alteration to the system model. Consequently we will have two models, the system model and the observation model. The system model tracks the evolution of corrosion damage to the system in time, as outlined in section 2.2. The observation model provides our method for interpreting observations and how they relate to system behaviour.

Observations could be any computable function of the system. As a general structure we will use the form $\underline{y}_{ct} = f(\underline{u}_{ct} + \underline{\xi}_{ct})$, in which f is the observation function and $\underline{\xi}_{ct}$ is a location specific measurement error vector. By adopting this structure we can account for a range of different observation processes. The error term is included inside the function to represent our belief that any measurement device will take measurements with some degree of inaccuracy, so the observation will not simply be a transformation of the underlying true wall thickness value but in fact a transformation of a value which is in some sense 'near' to the underlying true wall thickness (where 'near' is controlled by $\underline{\xi}_{ct}$). The observation function then operates on this inaccurate value.

The form of the observations \underline{y}_{ct} will be determined by the observation function, f - the choice of which will be determined by the inspection aims. For example if the goal of the inspection is to learn about corrosion behaviour over the whole surface then f should return a value for each observed location, giving \underline{y}_{ct} as a vector. However, if we are interested in identifying components that are experiencing greater levels of corrosion then a summary statistic such as the mean or minimum wall thickness value may be more appropriate, thus \underline{y}_{ct} would take a single value for each component. If the inspection was only interested in overall system performance with respect to corrosion then it might be suitable to use a function, f , which only returns a single y_t value for the entire system.

To illustrate the observation model we consider how it could be used to generate two plausible observation processes:

1. Total observation with Normal measurement error.
2. Minimum component wall thickness with under-estimated error.

The first process requires a value to be returned for each observed location, it represents simple observation with error. For this case we would model ξ_{lct} as being Normally distributed with mean 0 and variance $\sigma_{\xi_{lct}}^2$, i.e. $\xi_{lct} \sim N(0, \sigma_{\xi_{lct}}^2)$. This is then sufficient to model the observation process we want, so for this example the observation function f would simply be the identity, therefore return a vector of length s_c for each component c as our observations, \underline{y}_{ct} . We can control the size of the measurement error through $\sigma_{\xi_{lct}}^2$. For example, perfect inspection would correspond to $\sigma_{\xi_{lct}}^2 = 0$.

The second observation process is more complicated. Firstly, we require only one value to be returned for each component (the minimum wall thickness), so our observations y_{ct} will be scalar for this observation process. Secondly, we are told that this value is always underestimated. To model either under- or over-estimation we choose ξ_{lct} to have an appropriate one-sided distribution (e.g. $\tilde{\Gamma}$, Beta); the degree of over- or under-estimation can be controlled through the parameters of this distribution. In this case the observation function f is minimisation with respect to

location, l , of the wall thickness values observed with error. That is:

$$y_{ct} = \min_l \{u_{lct} + \xi_{lct}\}$$

for each component, c .

2.3.1 Inspection designs

The observation model will also depend on the inspection plan, d . d tells us how to inspect the system by specifying which locations and which components to measure and also which measurement technique to use. The model structure gives us a representation of the system at each time point through the underlying true wall thickness values (u_{lct}). Our observations will take the form of a transformation of some or all of these u_{lct} values which have been measured with error. We use the subscript d to identify which u_{lct} are to be included in an inspection. This takes the form of a set of components which are to be inspected (denoted C_d) and set of locations to be inspected (L_{c_d} for all $c_d \in C_d$), individual locations are denoted l_d . So the wall thickness value for a particular location included in the design would be denoted $u_{l_d c_d t}$; for notational simplicity we will refer to the inspection sites as those included in the vector \underline{u}_{dt} . Similarly, our observations will be denoted by \underline{y}_{dt} (where d is as described and t is time), and are defined to be:

$$\underline{y}_{dt} = f(\underline{u}_{dt} + \underline{\xi}_{dt}). \quad (2.8)$$

2.4 System model: Linear growth DLM

A brief description of the model was given in section 2.2. We now expand on that description to provide more detail on the evolution of the model in time.

The role of the system model is to produce representations of the behaviour of the surface of each component as we progress in time. As has been described, we expect corrosion to act on components at a roughly constant rate and therefore components to become more corroded as we move further forward in time. We model wall thickness (an observable quantity related to corrosion damage) at multiple locations

within a component as the sum of a global trend term (x_{ct}) and an independent local variation term (r_{lct}).

Although this structure allows us a great deal of flexibility and offers the potential to construct very detailed models of the component surface, we have chosen to work with a relatively simple model. This is motivated by our ultimate goal of producing an efficient means of designing an inspection plan for large systems. To be able to evaluate the worth of a single designed inspection, d , we will have to update our model for each possible set of inspection data \underline{y}_{dt} . For large systems performing this becomes a computationally intensive calculation, particularly if the measurements \underline{y}_{dt} have a complicated form. On top of this we will want to assess many designed inspections, so we will have to carry out multiple updates to account for the range of potential observation for each designed inspection, d . Consequently, the amount of the computation involved in comparing inspection plans can grow very rapidly. We have made the following simplifying assumptions in our model:

1. We are modelling a system in which corrosion has initiated, so the issue of corrosion initiation becomes a retrospective one. It still has to be modelled, but can now be handled as part of defining initial conditions, and not directly as part of the model.
2. Localised variation is not given a downward trend. We assume any trends can be modelled by the global terms.

A discussion of how these assumptions could be relaxed is included in chapter 7.

2.4.1 Linear growth DLM

As stated in Section 2.2 we will be modelling the global \underline{x}_t term using a linear trend dynamic linear model (DLM), as given by equations (2.6) and (2.7):

$$\underline{x}_t = \underline{x}_{t-1} + \underline{\alpha}_t + \underline{\varepsilon}_{xt}$$

$$\underline{\alpha}_t = \underline{\alpha}_{t-1} + \underline{\varepsilon}_{\alpha t}$$

The assumption of approximate linearity is modelled by the trend term $\underline{\alpha}_t$. $\underline{\alpha}_t$ controls the rate of change in \underline{x}_t , and remains roughly constant, subject to the random deviation controlled by $\underline{\varepsilon}_{\alpha t}$.

A Gaussian DLM is, in general, defined by the equations [24]:

$$\underline{y}_t = \mathbf{F}_t \underline{\theta}_t + \underline{\xi}_t, \quad \underline{v}_t \sim N(\underline{0}, \Sigma_\nu) \quad (2.9)$$

$$\underline{\theta}_t = \mathbf{G}_t \underline{\theta}_{t-1} + \underline{\varepsilon}_t, \quad \underline{\varepsilon}_t \sim N(\underline{0}, \Sigma_\varepsilon) \quad (2.10)$$

where (2.9) is the observation equation and (2.10) is the system equation. The DLM is defined by the quadruple $\{\mathbf{F}, \mathbf{G}, \Sigma_\nu, \Sigma_\varepsilon\}_t$ and its initial conditions, which will be specified as a mean $\underline{\mu}_{\theta_0}$ and variance Σ_{θ_0} of the initial state vector $\underline{\theta}_0$, such that $(\underline{\theta}_0 | D_0) \sim N(\underline{\mu}_{\theta_0}, \Sigma_{\theta_0})$ where D_0 is the information available at time $t = 0$. More generally D_t is the information available at time t . This general DLM structure can be related to the linear growth DLM we will be using as follows. We can define the state vector as $\underline{\theta}_t = [\underline{x}_t^T, \underline{\alpha}_t^T]^T$, the $2n$ -vector constructed from the n -vectors for system level and system slope. The error vector is given by $\underline{\varepsilon}_t = [\underline{\varepsilon}_{xt}^T, \underline{\varepsilon}_{\alpha t}^T]^T$. We can recover equations (2.6) and (2.7) through the choice of \mathbf{G} . The choice of \mathbf{G} which gives us a linear trend model is:

$$\mathbf{G} = \begin{bmatrix} \mathbf{I}_n & \mathbf{I}_n \\ \mathbf{0}_n & \mathbf{I}_n \end{bmatrix}$$

for all t . Σ_ε can be constructed using the covariance matrices for $\underline{\varepsilon}_{xt}$ and $\underline{\varepsilon}_{\alpha t}$:

$$\Sigma_\varepsilon = \begin{bmatrix} \Sigma_x + \Sigma_\alpha & \Sigma_\alpha \\ \Sigma_\alpha & \Sigma_\alpha \end{bmatrix}.$$

Our observation equation will typically be more complicated than the standard DLM observation equation given in (2.9), and for this we will use a separate observation model, as described in section 2.5. However, if we were interested in observing the exact system level values we could do this by choosing:

$$\mathbf{F} = \begin{bmatrix} \mathbf{I}_n \\ \mathbf{0}_n \end{bmatrix}, \quad \Sigma_\nu = \mathbf{0} \quad \forall t.$$

This ensures that our DLM ‘observation equation’ simply returns our system level

values, \underline{x}_t . In summary, the DLM we will be using has the form:

$$\underline{\gamma}_t = \mathbf{F}\underline{\theta}_t = \underline{x}_t \quad (2.11)$$

$$\underline{\theta}_t = \mathbf{G}\underline{\theta}_{t-1} + \underline{\varepsilon}_t, \quad \underline{\varepsilon}_t \sim N(\underline{0}, \Sigma_\varepsilon) \quad (2.12)$$

and is defined by the quadruple $\{\mathbf{F}, \mathbf{G}, \mathbf{0}, \Sigma_\varepsilon\}$ for all values of t . This stage identifies the values of \underline{x}_t from the system vector $\underline{\theta}_t$ for use as the global term; we then have to bring in the local variation term r_{lct} before considering observations. However, we can use this DLM framework to learn about the distribution of \underline{x}_t .

For any DLM, the k time step ahead forecast distributions are well defined. These are given by [24] to be:

$$(\underline{\theta}_{t+k}|D_t) \sim N[\underline{a}_t(k), \mathbf{V}_t(k)] \quad (2.13)$$

for the system state distribution. The forecast distribution for the system “observations” is given as:

$$(\underline{\gamma}_{t+k}|D_t) \sim N[\underline{f}_t(k), \mathbf{Q}_t(k)]. \quad (2.14)$$

The quantities $\underline{a}_t(k)$, $\mathbf{V}_t(k)$, $\underline{f}_t(k)$, $\mathbf{Q}_t(k)$ are defined recursively using the defining quadruple of the DLM. For a general DLM, these recursion relations are [24]:

$$\begin{aligned} \underline{a}_t(k) &= \mathbf{G}_{t+k}\underline{a}_t(k-1) \\ \mathbf{V}_t(k) &= \mathbf{G}_{t+k}\mathbf{V}_t(k-1)\mathbf{G}'_{t+k} + \Sigma_{\varepsilon_{t+k}} \\ \underline{f}_t(k) &= \mathbf{F}'_{t+k}\underline{a}_t(k) \\ \mathbf{Q}_t(k) &= \mathbf{F}'_{t+k}\mathbf{V}_t(k)\mathbf{F}_{t+k} + \Sigma_{\nu_{t+k}} \end{aligned}$$

where $\underline{a}_t(0) = \underline{\mu}_{\theta_t} = \begin{pmatrix} \underline{\mu}_{x_t} \\ \underline{\mu}_{\alpha_t} \end{pmatrix}$, the system level and system slope expectations at time t , and $\mathbf{V}_t(0) = \Sigma_{\theta_t} = \begin{bmatrix} \Sigma_{x_t} & \Sigma_{x_t, \alpha_t} \\ \Sigma_{\alpha_t, x_t} & \Sigma_{\alpha_t} \end{bmatrix}$, the system level (Σ_{x_t}) and system slope (Σ_{α_t}) variance matrices, plus the system level and slope covariance matrix (Σ_{α_t, x_t}). In the case where the defining quadruple is constant in time (as applies to

our model), these can be simplified to:

$$\underline{a}_t(k) = \mathbf{G}^k \underline{\mu}_{\theta_t} \quad (2.15)$$

$$\mathbf{V}_t(k) = \mathbf{G}^{k \Sigma_{\theta_t}} \mathbf{G}'^k + \sum_{j=0}^{k-1} \mathbf{G}^j \Sigma_{\epsilon} \mathbf{G}'^j \quad (2.16)$$

$$\underline{f}_t(k) = \mathbf{F}' \mathbf{G}^k \underline{\mu}_{\theta_t} \quad (2.17)$$

$$\mathbf{Q}_t(k) = \mathbf{F}' \mathbf{V}_t(k) \mathbf{F} + \Sigma_{\nu}$$

and for our model, in which $\{\mathbf{F}, \mathbf{G}, \Sigma_{\nu}, \Sigma_{\epsilon}\}_t$ are given by $\{\mathbf{F}, \mathbf{G}, 0, \Sigma_{\epsilon}\}$ for all t , the definition of $\mathbf{Q}_t(k)$ further simplifies to:

$$\mathbf{Q}_t(k) = \mathbf{F}' \mathbf{V}_t(k) \mathbf{F}. \quad (2.18)$$

So we can describe our beliefs about future system level behaviour in terms of:

$$\begin{aligned} E(\underline{x}_t) &= \underline{f}_0(t) \\ \text{var}(\underline{x}_t) &= \mathbf{Q}_0(t) \end{aligned}$$

We can see from equations (2.16) and (2.18) that Σ_{ϵ} is important in determining the forecast system variance. Therefore the constituent elements of Σ_{ϵ} , Σ_{ϵ_x} and $\Sigma_{\epsilon_{\alpha}}$ are influential in establishing covariance between components. (2.16) shows how this covariance structure is introduced over time. Our assumption is that we can use these matrices to capture the full ‘‘between component’’ covariance structure, and therefore not have to build a trend or any correlations into the local variation term.

2.4.2 Local variation

For each time point, t , the local variation terms r_{lct} form an $s \times n$ matrix

$$\mathbf{R}_t = (r_{lct})_{l=1}^s \quad c=1}^n$$

where

$$s = \max_c \{s_c\}.$$

Each column represents the values at the s_c locations in component c . We model $r_{lct} - r_{lc(t-1)}$, the difference in local variation between successive times, as $\zeta_{lct} \sim N(0, \sigma_{\zeta_c}^2)$, an independent random draw from a normal population. For any t , these differences form an $s \times n$ matrix, $\mathbf{Z}_t = (\zeta_{lct})_{l=1}^s_{c=1}^n$. Each column of \mathbf{Z}_t is a draw from the multivariate normal distribution $N(\mathbf{0}, \sigma_{\zeta_c}^2 \mathbf{I}_s)$. With ζ_{lct} defined as:

$$\zeta_{lct} = r_{lct} - r_{lc(t-1)}$$

we can write the evolution of the matrix \mathbf{R}_t in time to be:

$$\mathbf{R}_t = \mathbf{R}_{t-1} + \mathbf{Z}_t.$$

The initial conditions \mathbf{R}_0 will be determined by beliefs about the system state at $t = 0$. For details see section 2.7

We have adopted the simplest local variation structure, that of uncorrelated random deviation from the mean. It would be possible to allow for greater between component correlation by allowing the off-diagonal elements of Σ_{ζ} to take non-zero values. Each row of \mathbf{Z}_t would then be a draw from multivariate Normal with covariance structure given by Σ_{ζ} . If we had reason to believe certain regions of a component were behaving differently to others we could model this by allowing those regions to take different variance values. In effect this would mean giving each component a full $s_c \times s_c$ covariance matrix, Σ_{ζ_c} , so each column of \mathbf{Z}_t would be drawn from $N(\mathbf{0}, \Sigma_{\zeta_c})$.

2.4.3 System Model

This gives us a method for modelling the underlying true wall thickness value u_{lct} as $u_{lct} = x_{ct} + r_{lct}$. In matrix form this would be written as:

$$\mathbf{U}_t = \mathbf{1} \cdot \mathbf{x}_t^T + \mathbf{R}_t \quad (2.19)$$

where $\mathbf{1}$ is an s -vector of ones.

At each time point we represent the underlying true wall thickness values as an $s \times n$ matrix \mathbf{U}_t in which each column corresponds to the values taken at each of the

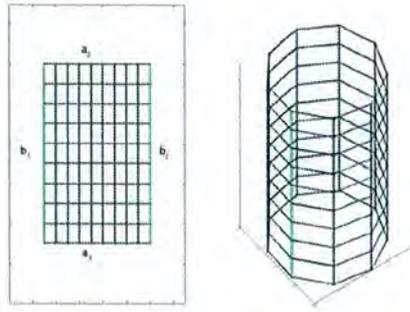


Figure 2.2: Mapping a regular lattice to a pipe

locations within a component c . The expected value and variance at each of these locations can be evaluated as the sum of the expectation and variance of the relevant global and local terms. From the DLM theory [24], we know \underline{x}_t is distributed as $N[\underline{f}_0(t), \mathbf{Q}_0(t)]$, where $\underline{f}_0(t)$ and $\mathbf{Q}_0(t)$ are defined as in Section 2.4.1. The local variation term is distributed as $N(\underline{0}, t\Sigma_C)$. So for any location l in a component c , the expected value of u_{lct} is $\underline{f}_0(t) + \underline{0} = \underline{f}_0(t)$. Similarly, the variance of u_{lct} is given by $\mathbf{Q}_0(t) + t\Sigma_C$.

2.4.4 Graphical representation of corrosion surfaces

The matrix \mathbf{U}_t can be easily converted in to a visual representation of the system in terms of the underlying true wall thickness values. For each component we would be able to produce a plot of the wall thickness behaviour. This is done on a componentwise basis simply for ease of interpretation. We model the surface of a component as a regular 2-dimensional $l_a \times l_b$ lattice, where $l_a \times l_b = s$ (see Figure 2.2). In most cases the surface we model will not be rectangular, so we will have to define a mapping between the lattice and the real surface shape. In the example of pipework (with which we will mostly be dealing) we can consider the rectangular lattice as corresponding to the cylindrical pipe by identifying the edges b_1 and b_2 of the lattice with each other, so that these correspond to adjacent locations on the cylinder. Essentially, we ‘fold the surface around’ to form a cylinder.

Having related the lattice to the component we are modelling, we can then build the visual representation. We identify each of the s_c locations for which we have a value with a point on the component surface. So for each location we will have co-ordinate pairs (a, b) , which tell us where the location l lies on the lattice. For each l , we then have the u_{lct} which tells us the modelled wall thickness value at that point. So at a given time t for location l in component c we can write down a vector $(a, b, u_{lct}) = \underline{V}_{lct}$ that allows us to identify where on the component's surface a location is and what the wall thickness value is at that location, l . By repeating this for all locations, l , within component c at time t , we obtain a matrix \mathbf{V}_{ct} , in which each row can be seen as a set of co-ordinates to be plotted in 3-dimensional space. By plotting each row of \mathbf{V}_{ct} we can obtain a 3-dimensional representation of the surface of the component.

As we always map our component surface on to a regular lattice, we can create a 2-dimensional representation of the component surface, by colouring each location, l , according to its u_{lct} value. 2-dimensional representations can be easier to interpret than 3-dimensional ones, which provides the motivation for this extra step. We define a fixed colour scale for the wall thickness, so that any value of u_{lct} has a colour associated with it. Our new representation uses vectors $(a, b, COL(u_{lct}))$ - for each location l in component c at time t - to generate the plot. a and b define the co-ordinates of the lattice associated with location l and $COL(u_{lct})$ tells us how we colour the location. Having coloured each location, we colour the rectangles between locations by interpolation based on the colours at the 4 locations corresponding to the 4 vertices of the rectangle.

For example, if we consider a new component for which the surface can be represented by 30 locations arranged as a 6×5 grid, the evolution of the surface of this component over the first 8 time steps can be represented as shown in Figure 2.3. The colour scheme chosen here attaches red colours to low wall thickness values and dark blues to high wall thickness values. Problem areas are therefore coloured in red, and in general we can track the evolution of the wall thickness values of the new component from all being equal (as they must be in a new component) at time $t = 0$, through each time step until $t = 8$, at which point we can see a few reddish

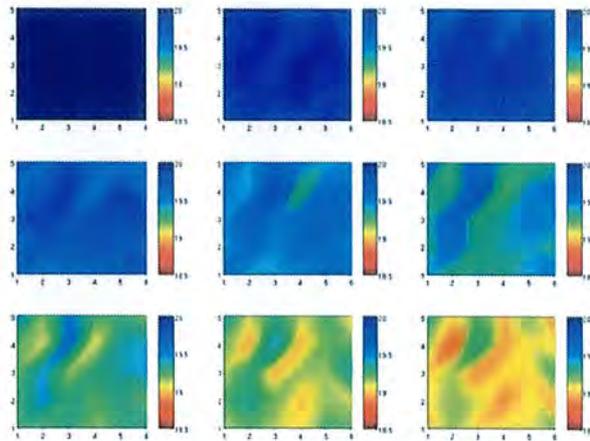


Figure 2.3: Visual representation of system degradation

areas starting to develop.

Visualisation is important for circumstances in which modelling the whole component surface is the aim. It offers a quick reference for viewing the model output which may be useful for both model validation and identifying regions of the system which may be in need of more careful analysis.

2.5 Observation model

We follow the modelling approach discussed in section 2.3. That is, for a designed inspection, d , our observations $\underline{y}_{dt} = f(\underline{u}_{dt} + \underline{\xi}_{dt})$. As stated in 2.3, the choice of observation function is determined both by our inspection aims and the available inspection tools, so to understand our observation function we must refer back to our problem structure (section 1.3).

The aim of inspection for the systems we will be considering is to identify those components at risk of imminent breakdown and to help inform decisions about whether or not the remaining components will be able to function safely up to the next inspection. Breakdown is defined as the component's wall thickness having fallen below some known acceptable minimum value, henceforth referred to as the

critical wall thickness W_C (for component c). Systems are assumed to operate as efficiently as their weakest component and in turn components are deemed to be as strong as their weakest point. Therefore a system ceases to operate effectively if a single component, c is operating inefficiently. Reduced wall thickness clearly increases the likelihood of a component failure, but it can also act to reduce component efficiency. We denote by W_C the critical wall thickness value - below which a component's ability to function efficiently cannot be guaranteed and the probability of component failure is greatly increased. If any point on its surface (u_{lct}) falls below W_C , we shall treat the component as though it has failed. This simplified view of component and system failure is in line with typical inspection and maintenance practice, where the consequences of actual component failure are seen as so severe that these 'breakdowns' are used as a trigger for maintenance. Component breakdown is not the same as component failure - in which the component completely ceases to function - but should be regarded as the step that immediately precedes component failure. We can thus formulate a definition of component breakdown in terms of the minimum wall thickness value; a component, c , is said to have broken down if and only if its minimum wall thickness value is less than W_C .

The observation function could, in general, be any function of the underlying wall thickness values. For this account we restrict discussion to the case of minimisation over a component, or more precisely over the locations within a component. We model the 'true' component minimum as:

$$m_{ct} = \min_l \{u_{lct}\} = \min_l \{x_{ct} + r_{lct}\} = x_{ct} + \min_l \{r_{lct}\} \quad (2.20)$$

and denote by \underline{m}_t the vector of component minima for the full system. Note that the minimisation is carried out over the locations, l , so the global trend term, which does not depend on l , can be taken out of the minimisation.

The distribution of these component minima is not easy to write down analytically. We know the minimisation does not act on the global term, \underline{x}_t , which has distribution $N[\underline{f}_0(t), \mathbf{Q}_0(t)]$. However, the distribution function of the minimum of independent draws from a Normal population is given by the distribution of the population to the power of the number of locations over which we are minimising. In the limiting case, where s_c is sufficiently large, this will be a Gumbel distribution.

However in most situations we consider, s_c will not be sufficiently large. This is a consequence of our location model. We model the component surface using a grid in which locations are sufficiently spaced so as to be independent of each other, so the value of s_c is relatively small. Increasing the number of locations will violate the assumption of spatial independence, also meaning we cannot use the limiting results. The exact distributional form will lie somewhere between a Normal and a Gumbel distribution. As the exact distribution depends on the value of s , we cannot write down a concise expression for the distribution, but, using the independence of the global and local terms, the expected value and variance can be written as:

$$\begin{aligned} E(\underline{m}_t) &= \underline{f}_0(t) + E\left(\min_l \{r_{lct}\}\right) \\ \text{var}(\underline{m}_t) &= \mathbf{Q}_0(t) + \text{var}\left(\min_l \{r_{lct}\}\right). \end{aligned}$$

We assume we are using an imperfect measurement device and will therefore need to account for measurement error. Measurement error could take any form, but we use Gaussian errors. Consequently, our ξ_{lct} term is distributed as $\xi_{lct} \sim N(0, \sigma_{\xi_c}^2)$, where $\sigma_{\xi_c}^2$ is the measurement error variance for component c . We allow this variance to depend on c to reflect the belief that some components, due to operational factors, environmental effects or physical inaccessibility may be more difficult to observe accurately than others. Our observation equation therefore takes the form:

$$\begin{aligned} y_{c_{dt}} &= f\left(\underline{u}_{c_{dt}} + \underline{\xi}_{c_{dt}}\right) \\ &= \min_{l_d} \{x_{c_{dt}} + r_{l_d c_{dt}} + \xi_{l_d c_{dt}}\} & \xi_{l_d c_{dt}} &\sim N(0, \sigma_{\xi_c}^2) \\ &= x_{c_{dt}} + \min_{l_d} \{r_{l_d c_{dt}} + \xi_{l_d c_{dt}}\} & \xi_{l_d c_{dt}} &\sim N(0, \sigma_{\xi_c}^2) \\ &= x_{c_{dt}} + \tilde{\omega}_{c_{dt}} \end{aligned} \tag{2.21}$$

for each $c_d \in C_d$ and where $\tilde{\omega}_{c_{dt}}$ is the minimum of the $r_{l_d c_{dt}} + \xi_{l_d c_{dt}}$ terms.

This distribution is as complicated as that of the true component minima. We have the same problem as before of the precise form depending on the number of locations over which we have minimised. For the case of historical data, the problem is compounded by uncertainty over the number of inspected locations. We have no way of checking how many locations were inspected to find the reported minimum, and therefore no way of writing down the full distribution for the observation model.

For simulated observations, the problems faced are the same as those for the ‘true’ minimum case.

These distributional uncertainties lead us into using a simulation based approach to assess the relationship between observations and system behaviour, which is detailed in Section 2.8.

2.6 Building covariance structures

We will need to define covariance matrices for the evolution errors of the DLM (Σ_ϵ), the extent of local variation between time steps (Σ_ζ) and measurement error (Σ_ξ). Recalling the structure of Σ_ϵ : $\Sigma_\epsilon = \begin{bmatrix} \Sigma_{\epsilon_x} & 0_n \\ 0_n & \Sigma_{\epsilon_\alpha} \end{bmatrix}$, we can see that we must specify Σ_{ϵ_x} and Σ_{ϵ_α} . As Σ_{ϵ_x} , Σ_{ϵ_α} , Σ_ζ and Σ_ξ are all constructed similarly, each being defined by multiple inputs, we propose a general framework for defining covariance matrices.

For large inspection planning problems we will usually have auxiliary information which is relevant to the construction of the correlation structure. Characteristics such as component design, usage and location can be expected to have an influence on corrosion behaviour and should therefore be used for guidance in quantifying Σ_{ϵ_x} , Σ_{ϵ_α} and Σ_ζ . Σ_ξ will have a different set of relevant characteristics, but information concerning the inspection procedure is also likely to be available. Precisely what these characteristics are will be determined by the specific problem and data set being analysed, but for any application it is reasonable to assume the existence of some useful auxiliary information.

We want a method for defining Σ_{ϵ_x} , Σ_{ϵ_α} and Σ_ζ in terms of auxiliary information relevant to system behaviour and Σ_ξ in terms of information relating to the inspection procedure.

2.6.1 Representing uncertainty

We define $H = \{h_k\}_{k=1}^N$ to be the set of characteristics described by the auxiliary information. These can be used to define the behaviour of a system with respect to each of ε_x , ε_α and ζ and therefore help in quantifying the respective covariance matrices. We assume that the deviations ε_x , ε_α and ζ can be modelled as a linear combination of independent terms - one term for each characteristic h_k , each with a known weighting $\lambda_{\varepsilon_x h_k}$ (respectively $\lambda_{\varepsilon_\alpha h_k}$, $\lambda_{\zeta h_k}$ or $\lambda_{\zeta h_k}$ depending on which deviation is being calculated) - and an independent Gaussian noise term $\underline{\psi}_x$, with variance Σ_{ψ_x} (the specific variation). We can then write ε_x as:

$$\varepsilon_x = \sum_{k=1}^N \lambda_{\varepsilon_x h_k} h_k + \psi_{\varepsilon_x}. \quad (2.22)$$

Assuming each of these characteristics to be Normally distributed with mean $\underline{0}$ and variance Σ_{h_k} , we can then write Σ_{ε_x} as:

$$\Sigma_{\varepsilon_x} = \sum_{k=1}^N \lambda_{\varepsilon_x h_k}^2 \Sigma_{h_k} + \Sigma_{\psi_{\varepsilon_x}} \quad (2.23)$$

and similarly $\Sigma_{\varepsilon_\alpha} = \sum_{k=1}^N \lambda_{\varepsilon_\alpha h_k}^2 \Sigma_{h_k} + \Sigma_{\psi_\alpha}$ and $\Sigma_\zeta = \sum_{k=1}^N \lambda_{\zeta h_k}^2 \Sigma_{h_k} + \Sigma_{\psi_\zeta}$.

We will be interested in the relationship between an auxiliary variable and the component corrosion rate. Considering rate allows us to handle the observational data on a standard scale. Our covariances will be built up from the corrosion rate behaviour associated with each characteristic. To illustrate, we consider a simplified version of the characteristics component type (CT) and pipe diameter (PD). Component type is a categorical variable with three levels: Straight (CT1), Bend (CT2) and Tee (CT3). These describe the shape of pipework in the section of system we are modelling. We have data from similar systems that allows us to estimate the standard deviation of the corrosion rate for each level of the characteristic, giving us quantities: σ_{CT1} , σ_{CT2} , σ_{CT3} . Pipe Diameter is a continuous variable and the standard deviation, σ_{PD} , for a component is taken to be a function of its pipe diameter. For this example, the ε_x term for an individual component with type bend and Pipe Diameter 3mm would have covariance given by:

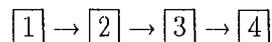
$$\sigma_{\varepsilon_x} = \lambda_{\varepsilon_x CT}^2 \sigma_{CT2} + \lambda_{\varepsilon_x PD}^2 \sigma_{PD}(3) + \sigma_{\psi_{\varepsilon_x}}$$

Each of the Σ_{h_k} terms represents our uncertainty over system corrosion behaviour with respect to characteristic h_k . These terms are common uncertainties, characteristics which affect every component in some way. We use the covariance matrix to describe the extent of the relationship between components. For an n component system, these matrices will be of size $n \times n$, where the (c, c') th entry is the covariance between component c and component c' . We build up the individual characteristic covariance matrices, Σ_{h_k} , elementwise, using the function:

$$\sigma_{h_k c c'} = \sigma_{h_k c} \sigma_{h_k c'} \exp\{-\tau_{h_k} \|c - c'\|^2\} \quad (2.24)$$

where the quantities $\sigma_{h_k c}$ represent the standard deviation of characteristic h_k with respect to the value of h_k taken in component c . The $\|c - c'\|^2$ term is a measure of distance between components c and c' where ‘distance’ may depend on factors such as adjacency rather than physical distance, and τ_{h_k} (≥ 0) is an importance weighting associated with distance. We can control the impact of distance through τ_{h_k} . If we believe a characteristic to be unaffected by distance (i.e. association is based entirely on the value of h_k taken within components c and c') we can choose τ_{h_k} close to 0 to reflect this view. Similarly if we believe distance to be important in determining the extent of a characteristic’s influence on the covariance structure we can choose τ_{h_k} to be large. In general, the τ_{h_k} values will have to be elicited and then validated with data from related systems. The inclusion of a distance measure into the evaluation of the covariance matrices allows us to introduce spatial correlation into the model.

So for the component type example, the full covariance would be constructed elementwise using (2.24). Suppose we have a four component system, with structure:



where $c = 1$ is a tee joint $c = 2$ and $c = 3$ are straights and $c = 4$ is a bend. If we let our distance measure be defined as the difference between c , the resulting covariance

matrix would have the form:

$$\begin{pmatrix} \sigma_{CT3}^2 & \sigma_{CT3}\sigma_{CT1}e^{-\tau_{CT}} & \sigma_{CT3}\sigma_{CT1}e^{-4\tau_{CT}} & \sigma_{CT3}\sigma_{CT2}e^{-9\tau_{CT}} \\ \sigma_{CT1}\sigma_{CT3}e^{-\tau_{CT}} & \sigma_{CT1}^2 & \sigma_{CT1}^2e^{-\tau_{CT}} & \sigma_{CT1}\sigma_{CT2}e^{-4\tau_{CT}} \\ \sigma_{CT1}\sigma_{CT3}e^{-4\tau_{CT}} & \sigma_{CT1}^2e^{-\tau_{CT}} & \sigma_{CT1}^2 & \sigma_{CT1}\sigma_{CT2}e^{-\tau_{CT}} \\ \sigma_{CT2}\sigma_{CT3}e^{-9\tau_{CT}} & \sigma_{CT2}\sigma_{CT1}e^{-4\tau_{CT}} & \sigma_{CT2}\sigma_{CT1}e^{-\tau_{CT}} & \sigma_{CT2}^2 \end{pmatrix}$$

The variance values $\sigma_{h_k c}^2$ depend on the value of characteristic h_k taken by component c . This structure is a natural one for categorical and ordinal variables, in which groups already exist within the data. This interpretation is less readily applied for continuous variables, where it will be necessary either to find a suitable way to discretise the continuous variable and apply the same treatment, or simply use a continuous covariance measure, similar to that used for distance.

We would have to estimate matrices for each characteristic $h_k \in H$. Having evaluated all of these covariance matrices, we must also consider the specific variation term $\underline{\psi}_{\varepsilon_x} \sim N(\underline{0}, \Sigma_{\psi_{\varepsilon_x}})$. This term accounts for the variation we can expect to see in an individual component distinct from the ‘explained’ variation from the common causes of uncertainty.

2.6.2 Defining weights $\lambda_{\nu_{h_k}}$

In general only a subset of characteristics will contribute to each covariance matrix. We can control the influence of characteristics h_k through their weightings, $\lambda_{\varepsilon_x h_k}$ (respectively $\lambda_{\varepsilon_\alpha h_k}$ and $\lambda_{\zeta h_k}$). For each covariance matrix we wish to define there exists a set of weights Λ_x , Λ_α and Λ_ζ , in which the individual elements λ_{h_k} are the weighting associated with characteristic h_k . These values will either be estimated from data, or more likely elicited from experts. Weights will be assigned non-negative values with a weight of 0 being equivalent to the belief that the characteristic does not contribute to the final covariance structure. Higher weights correspond to characteristics that are considered to be more influential.

Overall, this provides us with a structured way of incorporating auxiliary information and using it to build justifiable covariance structures. It provides a

structured method for taking variables we know are related to corrosion and/or measurement properties and using the information they provide to inform our beliefs about system covariances. Making decisions as to which variables contribute to which covariance matrices is a difficult question with no definitive answer. Where it is possible to identify a variable that clearly forms part of one of the four matrices Σ_{ε_x} , $\Sigma_{\varepsilon_\alpha}$, Σ_ζ or Σ_ξ then that variable should be used to form part of that matrix. Usually we will have to rely on expert judgement.

2.7 Model specification

For any particular system we wish to model, we will need to know how to move from the general model, as described so far in this chapter, to the specific model for that system. The system model is completely specified by the evolution covariance matrices for the global term, Σ_{ε_x} and $\Sigma_{\varepsilon_\alpha}$, the local deviation covariance matrices Σ_{ζ_c} and the initial conditions for system level, \underline{x}_0 , system slope, $\underline{\alpha}_0$, and local deviation, \mathbf{R}_0 . Having specified each of these, we will then be able to evolve the model in time to obtain forecasts of future system behaviour. We will also need to specify parameters for the observation model. The two elements we need to determine the observation model will be the observation function f and the observational error variance Σ_ξ .

This gives us 8 quantities we are required to specify: 3 for system model covariances, 3 sets of initial conditions and 2 for observation model parameters. Each of these could have multiple elements, so we describe here exactly what specifying these quantities entails. In chapter 3 we carry out such a specification for an example based on data from an industrial installation provided by our collaborators.

2.7.1 Specifying covariances

We use the factor model approach described in section 2.6. This provides us with each of the covariances we need for evolving the model in time. We have to rely on

auxiliary data to establish a set of characteristics that we believe influence corrosion behaviour. These beliefs will be based on a combination of trends observed in the data and prior judgements of experts validated through data analysis. As stated previously, determining where the influence of an individual characteristic $h_k \in H$ should be built into the model may be difficult. This will always be a subjective process dependent on the modeller and his/her interpretation of the system.

In general, particular characteristics will often be useful in defining more than one of the covariance matrices, Σ_{ϵ_x} , Σ_{ϵ_α} , Σ_{ζ_c} . This is because certain characteristics, such as location or component function, are important in learning about more than one type of corrosion behaviour. How these matrices are determined and exactly how the factor modelled is constructed will depend on the system being modelled, but we can make some broad statements about Σ_{ϵ_x} , Σ_{ϵ_α} , Σ_{ζ_c} .

Σ_{ϵ_x} is the covariance matrix for the evolution deviation in system level at each time step. Therefore characteristics that affect the system directly, and so make the impact of corrosion more unpredictable, should be given more weight in factor model for Σ_{ϵ_x} . Σ_{ϵ_α} is the covariance matrix for the evolution deviation in system slope at each time step. This corresponds to the change in the wall loss rate between successive time points, so characteristics that affect the rate of change in system level rather than system level should be incorporated here. Σ_{ζ_c} represents the corrosion behaviour at individual locations within a component, and can be viewed as a measure of how predictable the localised corrosion behaviour is within a component. Components with high local variability can be interpreted as being those more susceptible to localised corrosion, so characteristics important in explaining localised corrosion behaviour should be given higher weighting in the factor model for these matrices.

2.7.2 Specifying initial conditions

The nature of the initial conditions again depends on the system being modelled, and the data available. We have assumed we will be modelling older systems in which corrosion has initiated. In general, components do not begin corroding instantly,

but there is instead a period between a component's installation T and the point at which it begins to corrode. We use the notation t_{ci} to indicate the point at which corrosion initiated in component c (corrosion is unlikely to initiate in all components simultaneously), and $t = 0$ as the first point in our model. Therefore both T and t_{ci} will, in general, be negative. We restrict ourselves to considering systems in which corrosion has initiated, and therefore the model structure as described can be applied without modification.

If we assume we know when the system started corroding (and at what rate), then we can use our model to obtain estimates for the mean and variance of \underline{x}_0 and $\underline{\alpha}_0$ by using the DLM forecast equations for the t_i time steps for which each component has been corroding. Similarly, we can estimate the mean and variance of \mathbf{R}_0 as the sum of t_{ci} Normal $N(\underline{0}, \Sigma_{c_c})$ distributions, for each component, c .

However, we will not usually know at which time point $t_i \in [T, 0)$ the system started corroding, or the initial corrosion rate. Therefore the wall thickness values at $t = 0$, and the associated rates $\underline{\alpha}_0$, will be random. We can estimate wall loss rates from data for corroding systems, and by using our assumption that wall loss due to corrosion occurs at an approximately constant rate, these estimates will be directly comparable to the initial wall loss rates. More precisely, we use data to obtain estimates for the mean and variance of the wall loss rate, where the wall loss rate in component c at time t is modelled as:

$$\alpha_{ct} = \alpha_{c,t_{ci}} + \sum_{j=t_{ci}}^t \varepsilon_{\alpha_j} \quad \varepsilon_{\alpha_j} \sim N(0, \sigma_{\varepsilon_{\alpha_c}}^2) \quad (2.25)$$

and $\alpha_{c,t_{ci}}$ is constructed as the weighted sum of characteristic wall loss rates (i.e. our factor model for wall loss rate):

$$\alpha_{c,t_{ci}} = \sum_{k=1}^N \lambda_{\alpha_{h_k}} h_k(c) + \psi_c \quad \psi_c \sim N(0, \sigma_{\psi_c}^2) \quad (2.26)$$

where $h_k(c)$ is the contribution of characteristic h_k to the wall loss rate for the value of characteristic h_k taken by component c .

We denote the vector of initial wall loss rates by $\underline{\alpha}_{t_i}$ where we acknowledge that the initiation time could be different for each component, although this is not explicit in the notation. Our expectation for $\underline{\alpha}_{t_i}$, denoted $\mu_{\underline{\alpha}_{t_i}}$, is given by the expectation of

our factor model (2.26), the sum of expected contributions from each characteristic. Similarly, the variance is given by the sum of the variances for the factor model for wall loss rate, which is denoted Σ_α .

It is likely that we will have some historical data from the system we are modelling, in which case we may be able to identify the initiation point, or at least narrow the range in which it lies. We may also be able to use the historical data to provide improved estimates for our initial wall thickness values, removing the need to estimate these. In situations where we have good historical data, then it is preferable to use these values as our initial conditions. However, for situations in which we have no data from the system we are modelling, methods for estimating initial conditions using auxiliary data from other systems will have to be considered.

We estimate initiation time by considering historical data. In situations where the data set has a mixture of corroding and non-corroding components, we can use the age of the system and the number of new initiations within the time frame covered by the data to provide a plausible estimate of the initiation rate, β . As we are only considering initiation as part of the initial conditions, and not as part of the full model, we make the simplifying assumption that the time of corrosion initiation is exponentially distributed and independent for each component, with parameter, β . That is:

$$p(t_{ci} = t) = \beta \exp\{-\beta t\} \quad \forall c.$$

However, it would be possible to use a more involved model without compromising tractability, whilst the initiation modelling is treated as part of the initial conditions. We can then randomly generate the vector of initiation times \underline{t}_i as a draw from this exponential distribution. The combination of the initiation time vector, \underline{t}_i , and initial wall loss rate distributed as $N(\mu_{\underline{\alpha}_{t_i}}, \Sigma_{\epsilon_\alpha})$ will allow us to estimate a mean and variance for \underline{x}_0 and $\underline{\alpha}_0$ using the DLM forecast equations, as we now have an estimated initial rate and time. \mathbf{R}_0 can be estimated from $N(0, t_i \Sigma_\zeta)$, once the initiation time is known.

2.7.3 Observation model specification

The observation model is the most straightforward aspect of model specification. The observation function is determined by the inspection aims and the available inspection equipment, and should therefore follow directly from an understanding of these two things. Although the observation function f may itself be complicated, specifying it should be simple, as the relationship between the observation and the true system values should, in principle, be known. The observational error variance could be constructed using the factor model as described in Section 2.6, however it is more likely that there will be detailed prior information about the typical accuracy of inspection tools used. This tolerance information can be used to produce a final estimate for the observation error variance matrix.

2.8 Simulation approach

Due to the issues presented by the approximation of component surfaces introduced by the location model and the difficulties caused by complicated observation functions, we have chosen not to rely on an analytic approach, but instead use simulation based methods for the model inference and updating. The goal of the simulation approach is to provide an understanding of the relationship between the elements of the system model and the elements of the observation model, thereby allowing us to use observational data to update the system model in an informed way, and also to make predictions about future system behaviour based on typical model behaviour as understood through the simulations. The introduction of a simulation based approach is motivated by the goal of producing a method for comparing different inspection designs. Comparing designs is a computationally intensive process, and using a direct simulation approach allows us to avoid evaluating computationally demanding integrals and therefore makes the process computationally more tractable.

2.8.1 A simulation realisation

A simulation realisation is one run of the simulation algorithm. We now describe how to generate one such realisation. Each realisation of our simulation will produce one simulated system evolved over a specified number of time steps, T . So, we simulate a wall thickness value for every modelled location, l , in all components, c , for all times $t \in [1, \dots, T]$. In addition to this, we can then simulate observations of this simulated system for any designed inspection, d , thus allowing us to build up information about the relationship between system evolution and observations.

The first stage of the simulation process is setting up the simulation algorithm. The key steps of ours are listed below:

1. Defining simulation inputs (i.e. initial conditions)
 - 2a. Generate $\underline{\theta}_t$ for $t \in [1, \dots, T]$ using DLM.
 - 2b. Generate \mathbf{R}_t for $t \in [1, \dots, T]$.
3. Generate u_{lct} values using $u_{lct} = x_{ct} + r_{lct}$.
- 4a. Generate observations of the system for designed inspection d , using $\underline{y}_{dt} = f(\underline{u}_{dt} + \underline{\xi}_{dt})$.
- 4b. Find ‘true’ system values for the observation function f , i.e. $\underline{y}_t = f(u_{lct})$.
5. Store required output for $t \in [1, \dots, T]$.

This provides us with a set of instructions for what needs to be defined to set up the simulation. We work through these steps in order by way of explaining the simulation procedure.

1. Defining simulation inputs

We will be using our model structure for the simulation procedure. Therefore our inputs will be the same as the initial conditions for the model for the system of interest. That is, we will need to specify distributions for \underline{x}_0 , $\underline{\alpha}_0$, \mathbf{R}_0 and covariance matrices Σ_{ϵ_x} , Σ_{ϵ_α} , Σ_ζ for the system simulation, and give an observation function,

f , and an observational error covariance matrix Σ_ξ for the simulated observation process. Specification of these parameters is discussed in section 2.7. For those parameters for which an initial distribution is given, the value used for a particular realisation of the simulation will be drawn randomly from the specified distribution.

2. Generating the system

Using the model structure discussed in section 2.4, we generate the $\underline{\theta}_t$ and \mathbf{R}_t values independently. The s_c -vectors \underline{r}_{ct} are generated as the sum of independent draws, $\underline{\zeta}_{ct}$, from the Normal population $N(\underline{0}, \Sigma_{\zeta_c})$, using the relationship $\underline{r}_{ct} = \underline{r}_{ct-1} + \underline{\zeta}_{ct}$. This gives us an \underline{r}_{ct} vector for each component, and n in total for each time point, t , each of potentially different length. We store these as an $n \times s$ matrix, \mathbf{R}_t , where $s = \max_c \{s_c\}$, the greatest number of locations in a single component. The c th column of \mathbf{R}_t contains the s_c location values for component c . For those components in which $s_c < s$, the location values are entered as elements $1, \dots, s_c$ of column c and the remaining elements of \mathbf{R}_t are assigned a non-numeric identifier. This identifies these elements of the matrix as not corresponding to modelled location on the system and excludes these elements of \mathbf{R}_t from any further analysis.

For each time step, t , $\underline{\theta}_t$ is generated as:

$$\underline{\theta}_t = \mathbf{G}\underline{\theta}_{t-1} + \underline{\varepsilon}_t \quad \text{where} \quad \mathbf{G} = \begin{bmatrix} \mathbf{I}_n & \mathbf{I}_n \\ \mathbf{0}_n & \mathbf{I}_n \end{bmatrix}$$

and $\underline{\varepsilon}_t$ is an independent random draw from $N(\underline{0}, \Sigma_\xi)$. We store the \underline{x}_t values as the $n \times s$ matrix \mathbf{X}_t , which is defined as $\mathbf{X}_t = \underline{\mathbf{1}}_s \underline{x}_t^T$, where $\underline{\mathbf{1}}_s$ is a vector of ones of length s .

This gives us \mathbf{X}_t and \mathbf{R}_t as matrices of equivalent size. Each element of \mathbf{R}_t which takes a numeric value corresponds to a modelled location on the system. We will have both an \mathbf{X}_t and an \mathbf{R}_t matrix for every $t \in [1, \dots, T]$.

3. Generating \mathbf{U}_t

Our system realisations are determined by the \mathbf{X}_t and \mathbf{R}_t values, in line with the model. We obtain \mathbf{U}_t as an $n \times s$ matrix by taking $\mathbf{U}_t = \mathbf{X}_t + \mathbf{R}_t$. The non-numeric identifiers will appear in the \mathbf{U}_t matrix in the same places as they appeared

in \mathbf{R}_t , and will retain the same interpretation. So we obtain a matrix \mathbf{U}_t for each $t \in [1, \dots, T]$. Overall, \mathbf{U} , \mathbf{X} and \mathbf{R} will be stored as 3 dimensional arrays in which the 3rd dimension is indexed by time.

4. Generating observations

Simulating the observation process allows us to learn about how our observation model relates to our system model. Rather than build the observation process fully in to the simulation, which would require us to run a full simulation for each different designed inspection we wish to consider, a more practical approach would be to partially incorporate the inspection process into the simulation. Whilst generating the underlying system values \mathbf{U}_t , we can also generate a representation of the surface with observation error, by adding the error matrix Ξ_t (size $n \times s$), generated as s draws from the n -dimensional distribution $N(\mathbf{0}, \Sigma_\xi)$. We will call the resulting matrix $\mathbf{W}_t = \mathbf{U}_t + \Xi_t$, our simulation of every location in the system observed with error. Once we have the \mathbf{W}_t matrix of locations observed with error, we can apply any observation function, f , to any number of elements of \mathbf{W}_t , we choose. This gives us the option of comparing many different observation processes without needing to simulate many times.

It is reasonable to assume we will be interested in the output of the observation function for both the system observed with error, and the system observed without error (i.e. the ‘true’ value of the observation function for our system). For any real system, this is unobservable, but within the simulation framework, we can easily find both the observed value \underline{y}_{dt} and the true value \underline{m}_{ct} . For example, if the observation function is component minima then we can find the true component minimum for our simulated system by minimising \mathbf{U}_t with respect to the location index (in this case, by taking the column minima), but our observations \underline{y}_{dt} will be given by columnwise minimisation over only the locations and components of \mathbf{W}_t specified by d . Recording both of these quantities allows us to learn about the relationship between our observation model and system model. For our simulation out put, the vector of true surface minima at time t by:

$$\underline{m}_t = \min\{\mathbf{U}_t\}$$

and the vector of observed surface minima at time t for all c is defined to be:

$$\underline{y}_t = \min\{\mathbf{W}_t\}$$

5. Storing output

The simulation output will be determined by exactly what we are trying to learn about the system and its observations. However, it will not be more than:

- The system vector $\underline{\theta}_t$.
- The local variation matrix \mathbf{R}_t .
- The system plus error matrix \mathbf{W}_t .

each for all $t \in [1, \dots, T]$. All other variables discussed can be recovered as functions of these three quantities.

In practice, we may only be interested in a subset of these values. The complete simulation output above allows us to describe full system behaviour for a particular system as defined by our initial conditions. We are not especially interested in the behaviour of a single realisation of the system, but more so in typical system and model behaviour. In order to develop an understanding of model behaviour in general we have to consider many realisations of the system and the observations. By repeating steps 1-5 many times we will be able to build up a large number of output sets and by examining the typical properties of these sets learn about the behaviour of all aspects of the system.

In particular, we may wish to learn about the distributional behaviour of the simulation output. Empirical distributions can be created directly from the simulation output for each of \underline{x}_t , $\underline{\alpha}_t$, \underline{m}_t , \underline{y}_t . We could either use these immediately or fit a standard distributional form to the empirical distribution to allow us to make probability statements about system behaviour. We will also be interested in the covariance structure between pairs of variables, and we can also use the simulation output to learn about this.

The information we will be interested in will be the variances of the trend, true minimum and observed minimum within components at each time point, the means

of these quantities and the covariances between them. So for each time point we would have the following output:

- Means:
 - $E(\underline{\theta}_t)$ - system ($2n \times 1$);
 - $E(\underline{m}_t)$ - true minimum ($n \times 1$);
 - $E(\underline{y}_t)$ - observed minimum ($n \times 1$);
- Variances:
 - $\text{var}(\underline{\theta}_t)$ ($2n \times 2n$);
 - $\text{var}(\underline{m}_t)$ ($n \times n$);
 - $\text{var}(\underline{y}_t)$ ($n \times n$);
- Covariances:
 - $\text{cov}(\underline{y}_t, \underline{\alpha}_t)$, $\text{cov}(\underline{y}_t, \underline{x}_t)$, $\text{cov}(\underline{y}_t, \underline{M}_t)$;
 - $\text{cov}(\underline{m}_t, \underline{\alpha}_t)$, $\text{cov}(\underline{m}_t, \underline{x}_t)$;
 - $\text{cov}(\underline{x}_t, \underline{\alpha}_t)$ all $n \times n$

Chapter 3

Modelling large systems in practice: Site A data

In this chapter we focus on applying the ideas of chapter 2 to a real world example. Our objective is to use the Site A data set as an illustration of the modelling process described in chapter 2. We go through each aspect of model building for the particular system we are modelling both to indicate how to implement the model structure discussed in chapter 2 and also to highlight potential difficulties encountered when modelling real world systems. Where difficulties are encountered, we demonstrate by example how these can be negotiated.

The development of the model for this data set is important for illustrating the methods of chapter 2, but it will also be used as our standard example throughout the thesis. The model developed for the Site A data set will be the model used to illustrate our updating approach and the inspection planning procedure proposed in chapters 4, 5 and 6.

We begin by discussing the Site A data set (section 3.1) in terms of its structure and contents. Section 3.2 outlines what behaviour expert judgement would lead us to expect. Section 3.3 discusses the observation techniques used and how this should influence our modelling of the observation process. In section 3.4 we carry out an analysis of the Site A data set with the aim of identifying which elements of the data set are to be used in our model and verifying the expert assessments of 3.2.

How we can use the framework set up in chapter 2 to construct the model and which quantities we will use to specify our model parameters is also discussed in section 3.4. Section 3.5 then deals with the practicalities of specifying these parameters.

3.1 Site A data

The Site A data set consists of wall thickness measurements taken from the pipework elements of a large industrial system belonging to Shell. The data set is made up of observations made on 4910 different components. These measurements take the form of an irregular time series of between 1 and 9 observations for each component. Only components that have been observed are included in the data set. The 4910 components belong to a single very large system (Site A) that can be split into smaller units called ‘Corrosion Circuits’. The very large systems are built up from individual components by firstly constructing smaller sections by welding groups of components together to form a single unit of pipework consisting of joined components. These sections are then bolted together to form the very large system. A corrosion circuit corresponds to one of these sections of welded components, and these are the smallest units for which replacement is economically feasible. Individual corrosion circuits vary greatly in size, but can still contain many components. There are 63 corrosion circuits in the Site A data set, containing between 1 and 908 observed components.

The data set contains 6203 observations made over the 4910 distinct components during the period spanning 16/04/1998 to 29/09/2004. Multiple observations in time have been made at 828 of the components. We have auxiliary information about each component in the form of 113 descriptive variables. For this data set, no component takes values for all 113 variables. The descriptive variables can be categorised in one of 3 ways:

1. Variables which describe component properties.
2. Variables which describe the observation process.

3. Variables which describe expected corrosion behaviour within a component.

Variables of type 3 are in general expert subjective judgements.

Type 1 Variables: component properties

Variables of this type account for 26 of the 113 descriptive variables, and describe properties of the component that are at least semi-permanent, i.e. those which either cannot be changed or require changes to the whole system to change. These include variables related to the design specifications of each component, such as *Component Type*, *Internal/External Pipe Diameter* and *Material Type*. The age of an individual component can be measured from the *Commissioning Date* variable. Component function is summarised by the *Piping Schedule* and *Product* variables. These tell us what substance is travelling through the pipes and the regularity of component usage. Function is a good example of a semi-permanent component property; it is something which could be changed, unlike component design. However, there are no recorded instances of a component changing function throughout the span of the data.

Type 2 Variables: observation process

45 of the 113 variables are of this type. These variables contain information relating to both the inspection output and also describing how the observations were made. There are 3 clear subsets of type 2 variables, each one corresponding to a different inspection technique. The 3 observation techniques are given as Ultrasonic (*UT*), Radiographic (*GAMMA*) and Time of Flight (*TOFD*). For the Site A data set, only the ultrasonic and radiographic techniques were used to make observations, so we can immediately discount the Time of Flight subset. Only one technique is used for any single observation. Detailed descriptions of the observation techniques that were used to take the measurements for the Site A data set will be given in section 3.3.

Type 3 Variables: corrosion process

There are 39 variables that could be viewed as describing the corrosion properties of a component. Expected corrosion rates (not based on observed values) are given

for components and corrosion circuits. These have been specified using a risk based method, which will be covered in section 3.2. There are also 16 likelihood ratings for different types of corrosion behaviour. The relevance of these quantities to actual corrosion behaviour will have to be verified by data analysis.

Type 3 variables include a set of consequence ratings describing the consequences of component failure. These will not be of direct use when building the system model, but they could be used during the inspection planning process. For inspection planning we will require a way of assigning a cost to component degradation and these consequence ratings can help us to do this.

In addition to the Site A data set, we also have access to corrosion rate information from two further large systems. However we only have time series information for Site A. We can also gain spatial information for system layout in the form of schematic diagrams of the system. The data set is a potentially rich source of information, but we have some initial problems to work around. The scale of the system is a potential problem; the intricacies involved in modelling the entire Site A system are such that modelling the whole system would require a far more detailed treatment than it would be feasible to present here as an example. Another issue is our confidence in the auxiliary information.

We intend to resolve the issue of scale by modelling only a subset of the data - we will identify a modelling set, which will consist of the components we intend to model - and a training set, which we will use to learn about the corrosion behaviour in Site A. The model is developed for illustrative purposes, to show how the method could be applied, and to develop a non-trivial model on which to base our design procedure. With this in mind we will choose a modelling set that has sufficient complexity to illustrate the different aspects of the modelling process, but which remains interpretable.

Consultation with experts in the field of industrial inspection and maintenance will provide us with greater insight into how we should interpret the Site A data set. An illustrative analysis of the data set will allow us to detect any trends or patterns within the data. This analysis (detailed in section 3.4) will help us identify a conceptually more manageable subset of descriptive variables on which to base

our modelling. It will also allow us to assess the validity of the subjective expert judgements we have (either directly or from type 3 variables) and learn how we should interpret these for modelling purposes.

3.2 Factors relevant to corrosion

In this section we discuss corrosion behaviour we would expect to observe in the Site A system. We then offer an assessment of these judgements based on data analysis in section 3.4. Experts believe the following factors to be important in determining a component's corrosion behaviour:

- substance within the component;
- presence of corrosion in adjacent components;
- age of a component;
- flow of substances through components;
- temperature and pressure within the component;
- material from which a component is made.

We consider each of these in turn, providing an explanation of how we should interpret these factors.

3.2.1 Substance

It seems reasonable to assume that the substance a component is carrying will affect its corrosion behaviour. We would expect a component carrying a more reactive substance to be at more risk of severe corrosion than a component carrying a relatively inert substance. There is a wide range of substances carried by the components of the Site A installation, including both very inert substances such as water and air and very reactive substances like nitric acid and hydrogen sulphide.

Information on these factors can be described by the product variable, which gives details on how the component is being used and which substances are being moved through the component. *Product* is a potentially dynamic variable, as the use of a component could change with time, however no such changes are observed during the time frame covered by the data.

3.2.2 'Nearby' corrosion

Presence of other corroding components nearby being deemed as an important factor indicates a belief in some degree of spatial association involved in corrosion. We can check this by informally using the schematic representations of the systems to assess if there is historical evidence to support the notion that components experiencing high levels of corrosion occur near to each other. It may be that the effect we are expected to observe within the components is that of adjacent components being in similar physical locations and consequently being exposed to similar environmental conditions.

3.2.3 Age

The age of a component can make it more susceptible to certain types of corrosion. Therefore we would expect older components to be more corroded than newer ones. In particular, we would expect the minimum wall thickness to be lower for older components as more types of corrosion begin to take effect. Our beliefs concerning the initiation phase of the corrosion process imply that new components will not, in general, be subject to high levels of corrosion. The age of components can be measured using the *Commissioning Date* variable. This gives the date at which the component was first put into operation, so the component's age at each inspection point can be found relative to this. There is no record in the data of any component replacement, so we assume the *Commissioning Date* gives an authentic representation of the component age.

3.2.4 Flow

Flow of product through a component, and in particular disruption to flow, is considered to be an influential factor in determining the extent of corrosion. However, flow is a difficult and/or an expensive quantity to measure directly, so we have to make use of what we believe to be associated variables to learn about the effects of flow. Variables providing information on the effects of flow are expected to be the *Component Type*, *Product* and *Piping Schedule* variables. *Component Type* tells us about the design of a component, which can inform our views on how much flow disruption we can expect. For example, a straight section of pipe will probably be subject to fewer disruptions in flow than a T-shaped component. The *Product* variable tells not only what is flowing through the component, but also gives details on the nature of the process taking place within the component. Components involved in processes such as water injection and drainage are likely to suffer greater flow disruption than those involved in simple conveying processes. *Piping Schedule* details how substances are pumped through the system. This is a categorical variable in which the different piping schedules can tell us how the components are used, in particular whether the usage is continual or more intermittent. This information can again be used to learn about the flow within a component.

3.2.5 Environmental conditions

All components are exposed to internal and external environmental conditions. Temperature and pressure are two such examples of these conditions. We can learn about the internal pressure from the *Piping Schedule* variable, but we have little information on temperature. This is typical of most environmental factors; although they are believed to exert an influence, we have no variables that provide good information about them. This lack of information, and resulting inability to account for the effects of environmental factors in an informed way means we exclude them from the analysis. If reliable data on these factors were available, we could model

them explicitly as a term in the factor model. The factors are implicitly modelled by the spatial measures used in the model. By including information on the spatial relationships between components we can learn about which components are likely to be experiencing similar environmental conditions as a result of being in similar locations.

3.2.6 Material

Different materials respond to attack by corrosion in different ways, therefore we would expect material to play an important role in determining corrosion behaviour, and to see differences in corrosion behaviour between components made from different materials. However, the *Material* variable shows us that all parts of the Site A system included in the data set are made from carbon steel. Different types of carbon steel are used throughout the system, with different treatments and linings applied to them, but the underlying material used for each component remains essentially the same. The *PMC* (pipe material code) variable provides information on any additional treatments and/or linings that have been used on a component, and as such we could expect to see differences between the corrosion behaviour for the different levels of this variable.

3.3 Observation procedures

The observation method used can have a significant influence on the observation values. For the Site A data set two different measurement techniques are employed; these are ultrasonic and radiographic (or Gamma) testing. Both methods are part of a larger group of inspection techniques known as ‘non-destructive testing’ (NDT) or ‘non-intrusive inspection’ (NII) methods. This is because either can be carried out without the need for disassembly of the system, and therefore cause less disruption to the operation of the system during inspection. We provide a description of each and discuss the modelling implications of using either technique.

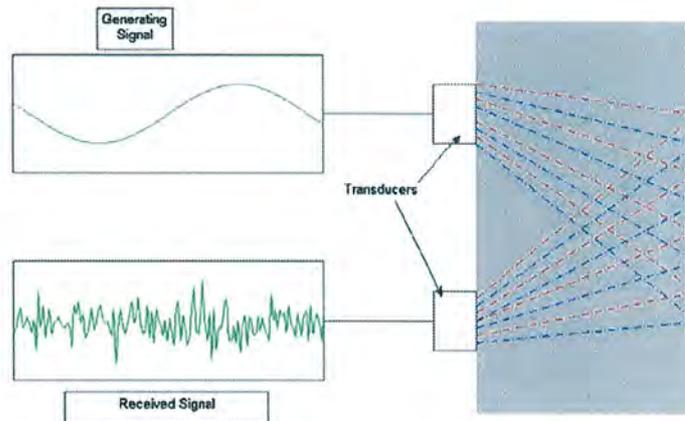


Figure 3.1: Representation of Ultrasonic Testing Procedure [9]

3.3.1 Ultrasonic testing

Ultrasonic testing methods are based on the use of sound waves to measure wall thicknesses. Sound waves are forced through the object to be inspected via a transducer. The reflected waves are then received by another transducer that displays the results. The first transducer converts a generating signal of known frequency into the corresponding sound wave, which is then introduced into the component at a known point. The sound wave travels through the component and is received by the second transducer that converts the information from received sound energy and the time it is received into an output signal. The output signal can be compared to the known input signal and other reference readings to assess wall thickness.

Ultrasonic testing can be used to provide wall thickness readings based on the sound velocity and attenuation measurements taken by the receiving transducer, and can be used to detect defects within materials and changes to material properties. It also requires minimal access to components to test them, needing only points at which the transducers can be attached, which is advantageous when inspecting complex systems. Such flexibility means ultrasonic testing is widely used within industrial inspection. However it is not without disadvantages. Interpreting results can be difficult and requires extensive training. Even for highly skilled inspectors the

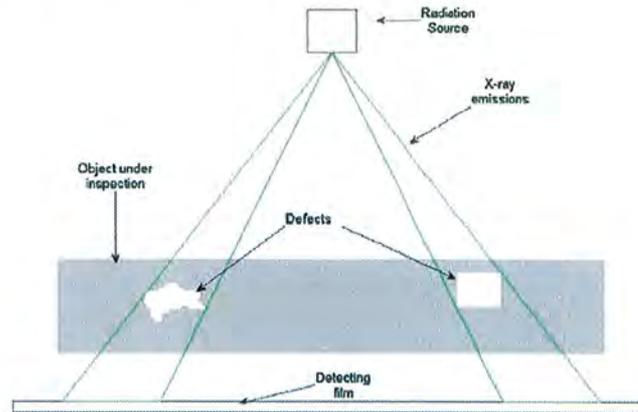


Figure 3.2: Representation of radiographic testing procedure [9]

use of reference readings are necessary to distinguish between normal and abnormal behaviour. The technique is less efficient if the wall thickness is low or if the defects run parallel to the line of inspection (although this is a problem faced by the majority of inspection techniques).

3.3.2 Radiographic testing

Radiographic (or Gamma) testing can be regarded as taking an X-ray of parts of the system. As with medical X-rays, the part we wish to observe is placed between a radiation source and a film or plate that is sensitive to radiation. The amount of radiation reaching the film creates an image, with the darker areas corresponding to those areas in which more radiation has penetrated through to the film. The extent of radiation penetration is controlled by the thickness and density of the material being observed. Therefore areas experiencing degradation and wall loss will show up as darker patches on the resulting output.

Radiographic testing can be used to detect defects within a component or to inspect components that are hidden inside larger components. As an NII technique it requires little disruption to system operation to perform inspection and is suitable for use on most materials. The results are more easily interpretable than those for ultrasonic testing. However, the method requires access to both sides of a compo-

ment, which may make it impractical for some systems. Operation of the equipment requires extensive training as defect detection can depend on the orientation of the emissions beam. Particularly thick components require longer exposure times for image generation, making the process more costly and increasing radiation exposure for the inspectors. However, the method works well for components with thin walls and offers the potential to inspect multilayered systems without direct contact with all components.

3.3.3 Modelling implications

An understanding of the way the observation process works is a vital part of modelling the observation error. The descriptions of each technique tell us that both methods have potential advantages and disadvantages, and both are subject to different sources of error, such as mismeasurement by the device itself, and human error due to the difficulty of interpreting the results. The inspection techniques are both quite different and should therefore be modelled separately. We choose to use a Gaussian observation error term. We assume the distribution of the measurement error can be written as: $N(0, \Sigma_\xi)$, where $\Sigma_\xi = \{\sigma_{\xi_c}^2\}_{c=1}^n$ is an $n \times n$ diagonal matrix in which n is the number of components and the values of $\sigma_{\xi_c}^2$ are determined by the inspection technique used in component c . This simple structure allows us to illustrate the principle of modelling measurement error in a rational way, but retain computational tractability for high dimensions.

For our example we will base our estimates for $\sigma_{\xi_c}^2$ on elicited values for typical measurement error. Expert judgement informs us that ultrasonic testing is accurate to within .5mm and radiographic testing is accurate to within 1mm when used to predict wall thicknesses. Validating these values from the available data would be a difficult process, requiring different data from that which is available to estimate the extent of measurement error, and in particular to determine different sources of error. From the available minimum wall thickness data it will be difficult to dissociate variation due to genuine wall loss and the variation due to measurement error. Due to these difficulties, and in the interest of obtaining a plausible solution to the

inspection planning problem, we have decided to trust these values. We will interpret these values as being equivalent to 3 standard deviations of the Normal distribution representing measurement error. Consequently, we set $\sigma_{\xi_c}^2 = (.5/3)^2$ if component c was inspected using the ultrasonic method, so that, under the assumption of Normality, 99% of observation errors fall within 0.5mm of the true underlying values. Similarly, we set $\sigma_{\xi_c}^2 = (1/3)^2$ if radiography was used.

It would be possible to implement a more careful model for the measurement error that takes into account what we know about the inspection techniques and the associated sources of error. A factor model, based on characteristics $h_{obs_k} \in H_{obs}$, similar to that used previously, could be constructed specifically for the measurement error:

$$\xi_{lct} = \sum_{k=1}^{N_{ins}} \lambda_{obs_k} h_{obs_k} + \psi_{obs_{lct}}.$$

Under this model the measurement error would be constructed as the sum of independent terms describing different characteristics of component inspection. These could include equipment inaccuracy and human error and can be used to allow the measurement to depend on component characteristics. We know that ultrasonic inspection is less accurate for lower wall thicknesses; given more precise information we could use the factor model to reflect this understanding, and similarly represent the reduced effectiveness of radiographic testing for large wall thicknesses. However, despite the advantages of a more careful measurement model, the difficulty of estimating the influence of each characteristic (via $\lambda_{obs_k}, h_{obs_k}$) leads us to the more pragmatic choice of the simple model described earlier.

3.4 Data analysis for the Site A data

We identify a subset of the Site A data set which we wish to model. A subset of auxiliary data is also identified for use in model verification. Firstly, we give an overview of the subsets identified for modelling and model assessment. We then illustrate how we could fit the two-term model of chapter 2 to our chosen subset. Our intention is to demonstrate the feasibility of such a model for this type of data.

The model introduced here will be used throughout later sections as the basis for our inspection planning examples. For this account we will not attempt to plan an inspection for the entire Site A system, but instead focus on a conceptually more manageable subset. Each corrosion circuit can be treated as a contiguous, complete unit and can therefore be considered separately. We have selected Circuit A to be the primary data as the wall loss behaviour within the circuit can be seen to be typical of wall loss behaviour throughout the system, but it is also a well observed circuit featuring several components with multiple observations in time, and few of the potentially anomalous readings seen elsewhere in the data. A further 12 circuits (accounting for approximately 40% of the total data) have been selected for use as the auxiliary data. These circuits - B, C, D, E, F, G, H, I, J, K, L and M - have been selected to provide a data set containing components with similar characteristics to the components of Circuit A.

The aim of this section is to construct an example based on the properties of a real world system for which we can then propose an inspection plan designed using a rational method. As such, the primary objective of our data analysis will be to ensure that the initial values we will specify for the model are of a similar order of magnitude as is seen for those quantities in the data set, rather than to provide exact parameter estimates.

3.4.1 Primary Data - Circuit A

Circuit A will be treated as our system. Circuit A consists of 46 distinct features and 73 welds, as depicted in Figure 3.3. We will restrict our model to the observed components. In general, components do not have to have been observed to be modelled, but in this situation we have no available data for the unobserved components, hence the need for this restriction. Figure 3.3 also provides us with the available spatial information for Circuit A, which tells us about the connectivity of components, but not the physical scale or even relative proportions of components.

The components of Circuit A belong to 7 different component types, summarised in Table 3.1. The division of the components into different material types, pipe

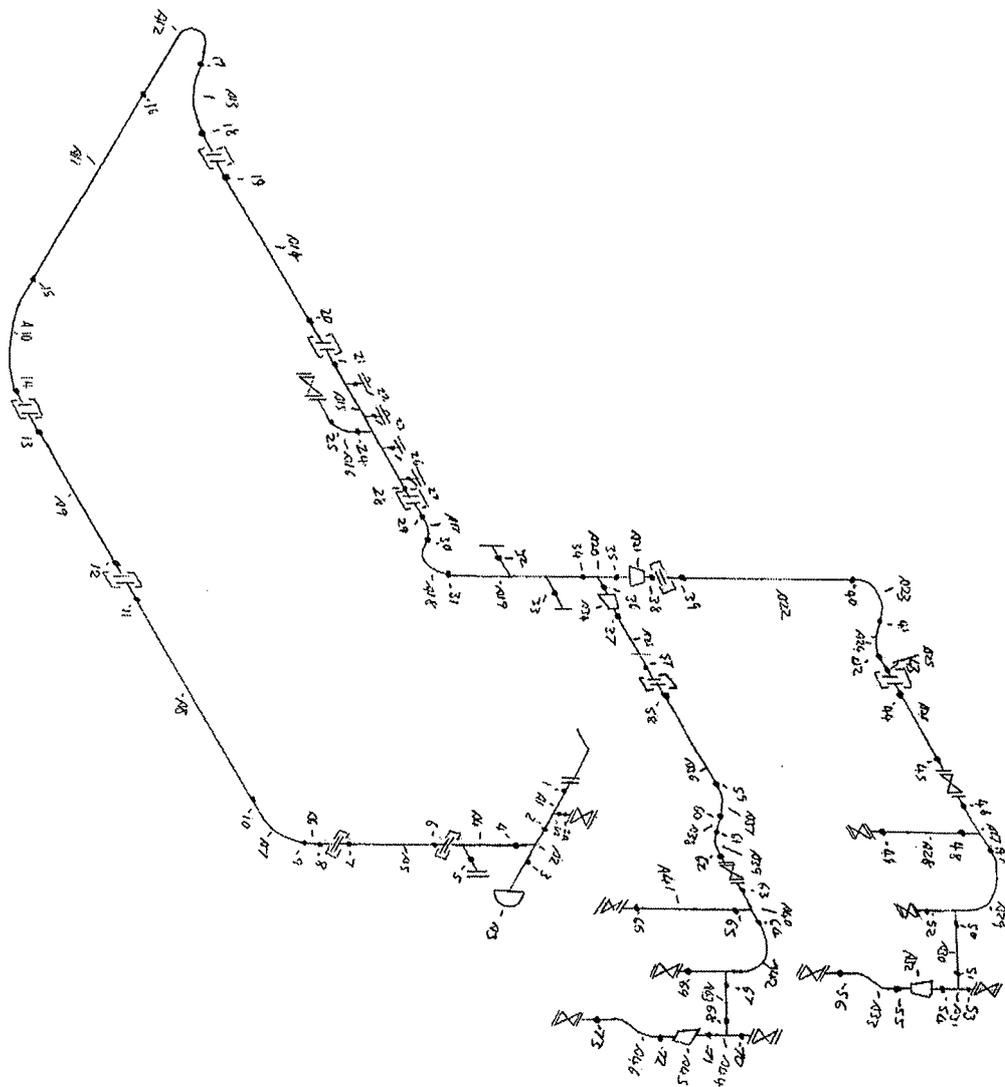


Figure 3.3: Plan of Circuit A

Component Type	Component Type Code	Frequency
Wellhead	KeyPoint(2)	11
Bend (small radius)	Bes	8
End Cap	Cap	1
Straight	Str	6
Weld	Weld	13
Tee, equal sized	TEQ	1
Branch connections	KeyPoint(4)	6

Table 3.1: Component type summary for Circuit A

material codes (*PMC*) and piping schedules is shown in Tables 3.2. The continuous variable information is summarised in Figure 3.4 in which the lower plot shows the spread of the initial wall thickness values and the upper plot shows the spread of the pipe diameter variable. The model components are numbered in the order in which they appear in Circuit A. The *Product*, *Confidence Rating* and *Commissioning Date* variables take the same value for all components within Circuit A, these are the level P, Medium and 18/11/1978 respectively. These are all common levels of the *Product*, *Confidence Rating* and *Commissioning Date* variables seen frequently in the data set as a whole. The adjacency matrix Adj , is shown in Figure 3.5, where nearer components are denoted by blues and red colours represent large distances between components. $a_{cc'}$, the adjacency value between components c and c' is defined to be the number of component boundaries we cross in moving between component c and c' . Under this system if $c = c'$, $a_{cc'} = 0$, and if c and c' are immediately next to each other $a_{cc'} = 1$, and so on.

The number of components experiencing wall loss within Circuit A is 40, therefore the value of n in our system model is 40. So to set-up a model for Circuit A we will need to define and verify 40-vectors \underline{x}_0 , $\underline{\alpha}_0$ for the initial wall thickness values and wall loss rate (WLR) values. We will also need to define and validate the 40×40 variance/covariance matrices Σ_{e_x} , Σ_{e_α} for the correlated DLM evolution deviations and Σ_ζ for the local variation terms. For this process we will use the auxiliary data set.

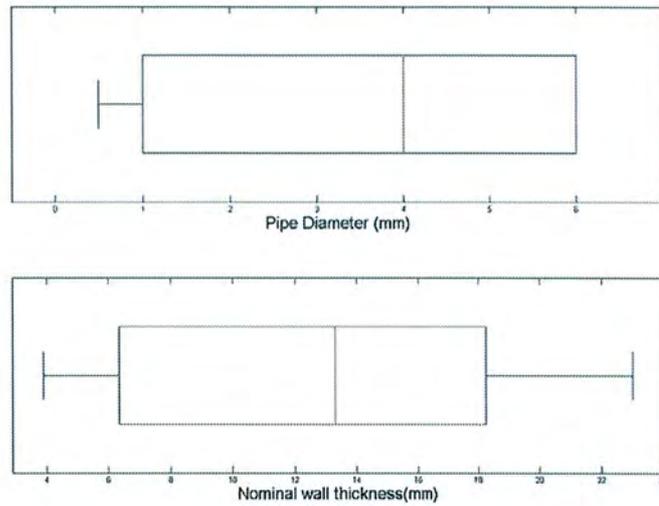


Figure 3.4: Boxplots of Nominal Wall Thickness and Pipe Diameter variables for Circuit A

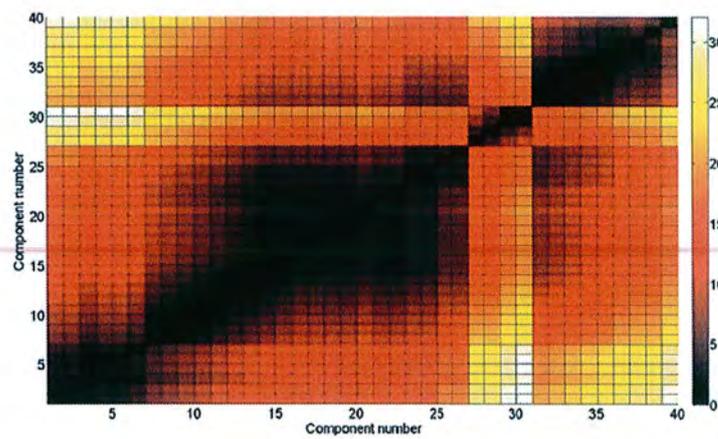


Figure 3.5: Distance (adjacency) matrix for modelled Circuit A components

Material	Frequency
1	17
2	29
Pipe Maintenance Code	Frequency
1	17
2	29
Piping Schedule	Frequency
PS2	9
PS3	29
PS1	8

Table 3.2: Categorical variable summary for Circuit A

3.4.2 Auxiliary data

The auxiliary (or training) data set consists of the information relating to circuits B, C, D, E, F, G, H, I, J, K, L and M. In total, this gives us information for 2009 different components, observed 2728 times during the time frame covered by the data. However, analysis of the data reveals that a key feature of the data set is the large number of components in which no wall loss has been recorded. Of the 2009 observed components, there are 1166 components in which the observed wall thickness value never falls below its initial value, and 843 in which some wall loss is observed. This is a pattern reflected over the entire Site A data set, in which approximately 64% of components have no recorded loss in wall thickness. As we are mostly interested in modelling the behaviour of components experiencing wall loss, the number of components showing no wall loss in Circuit A is considerably lower than this typical value (6 of 46). However, consistency with this value was harder to achieve with the much larger auxiliary data set. Due to the large number of components with a wall loss rate (WLR) of 0 we choose to treat the data set as 2 groups: the subset of components in which $WLR = 0$, and the subset in which $WLR > 0$.

The levels of the categorical variables represented in the $WLR > 0$ subset of the

Component Type	Component Type Code	Frequency
Wellhead	KeyPoint(2)	47
Bend (small radius)	Bes	40
End Cap	Cap	17
Straight	Str	48
Weld	Weld	106
Tee, equal sized	TEQ	8
Branch connections	KeyPoint(4)	8
Others		569

Table 3.3: Component type summary for Auxiliary Data

auxiliary data are given in Tables 3.3 and 3.4. As can be seen here all the values taken by these variables in Circuit A are well represented in the auxiliary data set, with the exceptions of the T-joint and End Cap, which are not well represented in the data. The Branch Connections variable is also under represented in the $WLR > 0$ subset of the auxiliary data. This is because a large number of these components have no recorded wall loss values. The large number of 'other' *PMC* levels seen is due to the large number of *PMC* levels (45) found in the Site A data set.

To help establish which variables should be treated as characteristics involved in determining wall loss behaviour we examine whether or not there are significant differences between the zero and nonzero subsets of the auxiliary data. Figures 3.6 - 3.9 show the number of components belonging to each subset for different levels of the categorical variables *Component Type*, *PMC*, *Material* and *Piping Schedule*. The number of components we would 'expect' to see in each of these groups is also displayed. We make the assumption that if there is no difference between the levels of the categorical variables the ratio of components belonging to the zero and nonzero *WLR* subsets will be the same as that seen in the whole data set. Approximately 64% of components in the Site A data set belong to the subset experiencing no wall loss (the $WLR = 0$ set), so we would expect to see approximately 64% of the components for each categorical variable belonging to the $WLR = 0$ subset. Under this hypothesis any significant deviations from this ratio tell us that a variable level

Material	Frequency
1	547
2	107
Others	186
Pipe Maintenance Code	Frequency
1	43
2	49
Others	751
Piping Schedule	Frequency
PS2	26
PS3	520
PS1	41
Others	256
Product	Frequency
P	192
Others	651

Table 3.4: Categorical variable level summary for Auxiliary Data

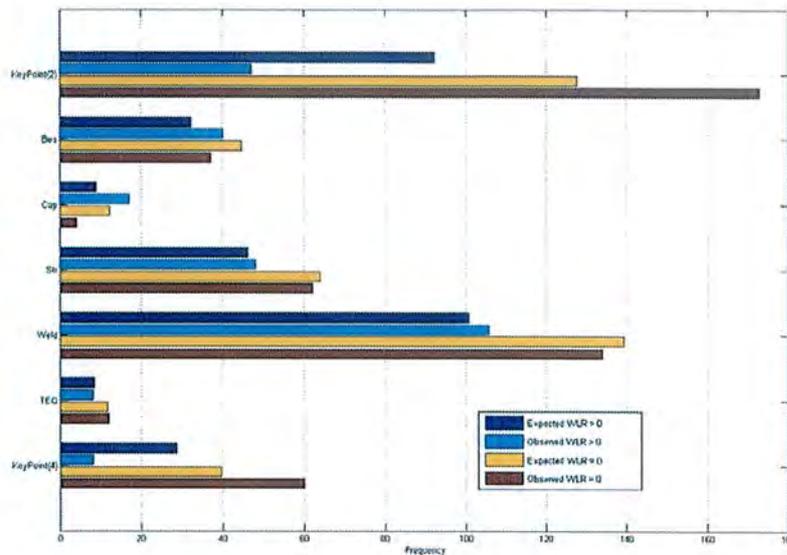


Figure 3.6: Members of zero/nonzero subsets displayed by component type

may be associated with a components wall loss behaviour.

There are considerable differences between the observed and the expected values for some of the levels of each of these categorical variables. The lack of consistency displayed across the different levels, with some following the expected pattern very closely and others deviating greatly, indicates that it is reasonable to assume that the illustrated variables do play some role in determining a components wall loss behaviour. That is, the subset to which a component belongs is not independent of the values it takes for the variables *Component Type*, *PMC*, *Material* and *Piping Schedule*.

The significance of these deviations can be tested using a χ^2 test, the results of which are displayed in Figure 3.10. Figure 3.10 plots the χ^2 values for each level of the variables in question. The three black lines indicate .95, .99 and .999 significance levels for a χ^2 test on one degree of freedom. Many of the differences are not significant. However, there is at least one very large value associated with each variable, suggesting that the variables *Component Type*, *PMC*, *Material*, *Piping Schedule* and *Product* at least have a role in determining whether or not wall loss initiates within a component, if not in influencing the rate of wall loss.

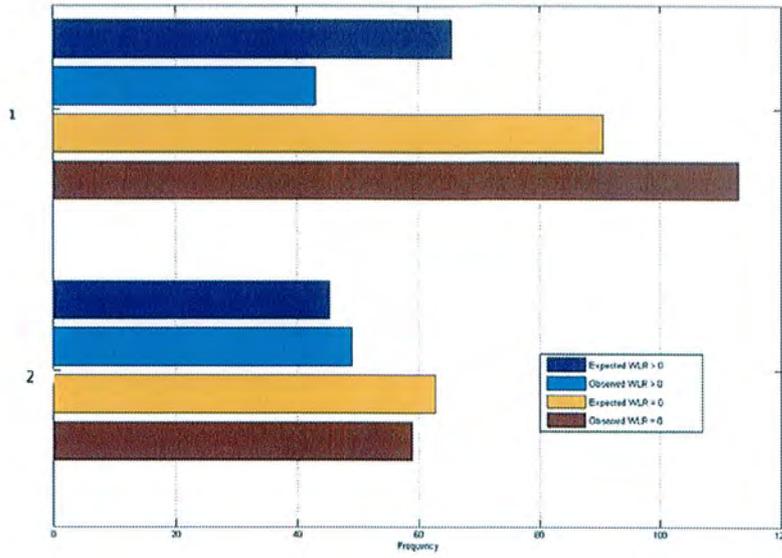


Figure 3.7: Members of zero/nonzero subsets displayed by PMC level

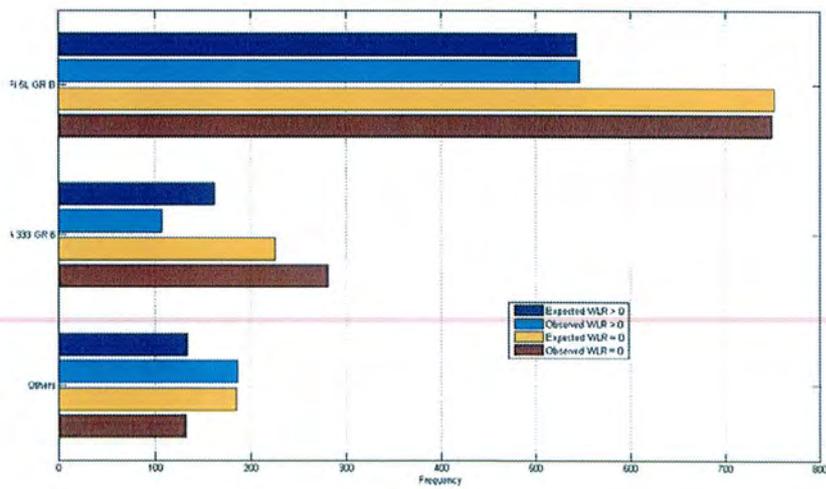


Figure 3.8: Members of zero/nonzero subsets displayed by Material type

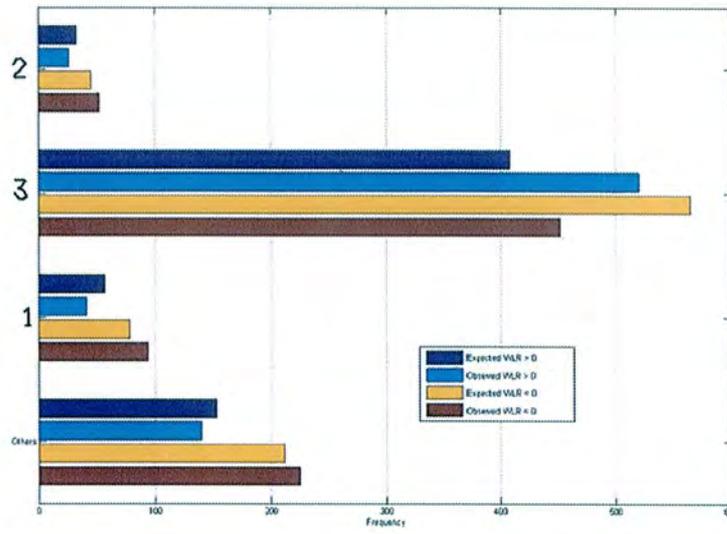


Figure 3.9: Members of zero/nonzero subsets displayed by Piping Schedule

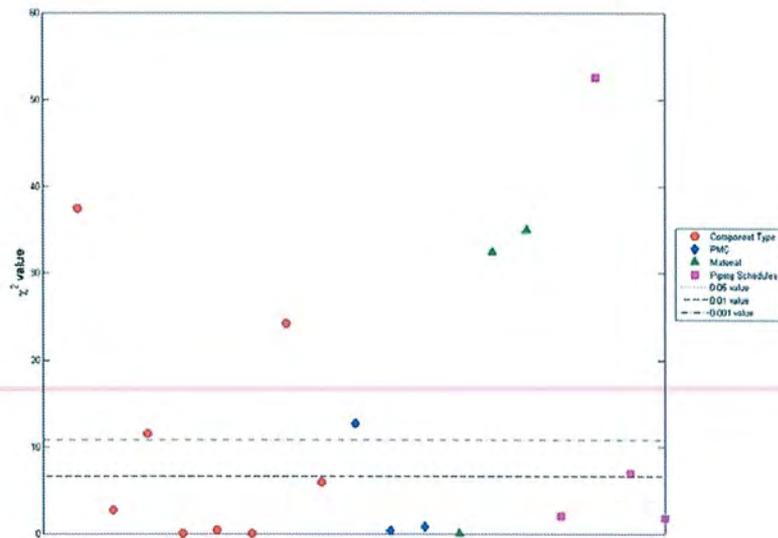


Figure 3.10: χ^2 test values

3.4.3 Regression Analysis

To further test which variables we should regard as important characteristics in our model we conduct a regression analysis on the $WLR > 0$ subset of the auxiliary data. We build a model using the following variables:

1. Component Type
2. PMC
3. Material
4. Piping Schedule
5. Pipe Diameter
6. Nominal Wall Thickness
7. Commissioning Date/Inspection Date (i.e. Age)
8. Confidence Rating
9. Criticality

These are the variables from the Site A data set we consider to be relevant to the wall loss process, rather than the observation process, which will be of limited importance in determining wall loss behaviour. The difficulties posed by this set of variables for regression stem from the different types of variables we have. Variables 5 - 7 are continuous and can therefore be used without alteration. Variables 1 - 4 are categorical and will therefore have to be coded into the regression model as dummy variables, i.e. a categorical variable ν with N_ν levels will first have to be converted into $N_\nu - 1$ 0-1 valued variables. We also have the ordinal variables, *Criticality* and *Confidence Rating*. We choose to model these as categorical variables. Ultimately, this gives us a set of predictor variables X_{reg} consisting of 46 dummy variables and 3 continuous variables.

The response variable y_{reg} is defined to be the percentage wall loss at time of inspection. We choose this quantity as it allows us to put our response variable

Model fitted	Adjusted R^2 value
All variables	0.4072
All except Component Type	0.3461
All except PMC	0.3788
All except Material	0.4082
All except Piping Schedule	0.4040
All except Pipe Diameter	0.4071
All except Nominal Wall Thickness	0.4077
All except Age	0.3959
Component Type, PMC and Age	0.4005

Table 3.5: Adjusted R^2 values for regression model fit for use in variable selection

on the same scale for all components. Components have a range of different wall thicknesses, so suffering a loss of 1mm may be severe wall loss for some components, but relatively minor damage for another component. Using percentage wall loss we can compare directly across components. \underline{y}_{reg} is defined as:

$$\underline{y}_{reg} = 1 - \frac{WT_{ins}}{WT_{nom}}$$

where WT_{ins} is the vector of wall thicknesses at inspection and WT_{nom} is the vector of nominal (or initial) wall thickness values. We then perform a linear least squares fit of the predictor variables X_{reg} on to \underline{y}_{reg} . To assess the importance of each variable we then refit the model omitting each variable in turn. For the categorical variable we omit all dummy variables associated with the original categorical variable. We compare the model fit using the R^2 statistic representation of residual error. The results are shown in Table 3.5.

The reduction in the R^2 statistic when the variables *Component Type*, *Age* and *PMC* are taken out of the model indicates that these are relatively important variables in determining wall loss behaviour. Indeed, if we fit a model based solely on these terms we obtain an R^2 value of 0.4005 - almost the same as that of the full model. However, it is also clear that there is considerable residual error in the model. The lack of fit can be partially explained by the fact that our set of \underline{y}_{reg} val-

ues are based on measurements taken with error. Therefore in addition to the error caused by lack of fit of the model, the R^2 values are also affected by the influence of measurement error, so the true error for the regression model in predicting actual percentage wall loss values is lower than that which is observed.

A further potential problem with the preliminary model is that it overlooks the initiation phase. The modelled values are based on percentage wall loss values at time t in components which we know experience wall loss at some stage. By using *Age* as our time predictor, we fail to take into account at which point a component begins to experience wall loss. To address this we fit a second regression model based on the *Component Type* and pipe material code (*PMC*) variables, in which the response variable, \underline{y}_{reg2} , values are defined to be the change in wall thickness between two points on a component which has already shown evidence of wall loss. That is, the difference between wall thickness (*WT*) readings at $t+k$ and t , given that the wall thickness reading at t is less than the *Nominal Wall Thickness* value for that component.

$$\underline{y}_{reg2} = \underline{WT}_{t+k} - \underline{WT}_t \quad \forall k \geq 1$$

This definition removes the dependence of our data \underline{y}_{reg2} on the *Nominal Wall Thickness* variable, present in the first regression, so *Nominal Wall Thickness* should be brought back into the predictor set. As it is a continuous variable, this is a simple process. The time predictor *Age* is redefined to be equivalent to the value of k , the number of time steps between observations. We aim to model change in wall loss between observations as a function of *Component Type*, *PMC*, *Nominal Wall Thickness* and time between inspections. The predictor set can be written as a matrix, \mathbf{X}_{reg2} in which each column corresponds to a variable shown in Table 3.6. The model takes the form:

$$\underline{y}_{reg2} = \underline{b}^T \mathbf{X}_{reg2} + \underline{e} \quad (3.1)$$

where \underline{b} is the vector of coefficients (Table 3.6) obtained via linear least squares fit and \underline{e} is the residual (error) vector. The levels of the *Component Type* and *PMC* variables which are poorly represented (fewer than 10 observations) in the predictor set are modelled as part of the intercept. The value of the adjusted R^2 statistic for this model, and the models resulting from leaving out individual variables are shown

Column in \mathbf{X}_{reg2}	Variable	Regression Coefficient (b_i)
1	Intercept	1.2445
2	Bend (small)	0.3129
3	Cap	1.2645
4	Straight	0.2003
5	Weld	-0.5229
6	Reducer	-0.8968
7	Key Point (9)	0.3482
8	PMC2	-0.5037
9	PMC3	0.0585
10	PMC4	-0.1938
11	PMC5	-0.2929
12	PMC6	0.1146
13	PMC7	1.0571
14	Age	-0.1834
15	Nominal Wall Thickness	-0.1620

Table 3.6: Variables and their associated coefficients for second regression model

in Table 3.7. These indicate an improvement in model fit.

The model structure of (3.1) mirrors the model for the minimum underlying wall thickness value, \underline{m}_t , given by (2.20) in chapter 2:

$$\underline{m}_t = \underline{x}_t + \tilde{r}_t \quad (3.2)$$

where \underline{x}_t is a trend term and \tilde{r}_t is a vector of deviations from the trend. Similarly in (3.1) we have:

$$\underline{y}_{reg2} = \underline{b}^T \mathbf{X}_{reg2} + \underline{e}$$

in which $\underline{b}^T \mathbf{X}_{reg2}$ models the expected change in wall thickness between observations and \underline{e} measures the observed deviation from this trend. As such, we will use the fitted regression model to inform our initial beliefs about wall loss rate, and the

Model fitted	Adjusted R^2 value
Initial regression model	0.4005
All variables (Second regression)	0.4632
All except Component Type	0.4329
All except Pipe Maintenance Code	0.4432
All except Nominal Wall Thickness	0.4150
All except Age	0.4237

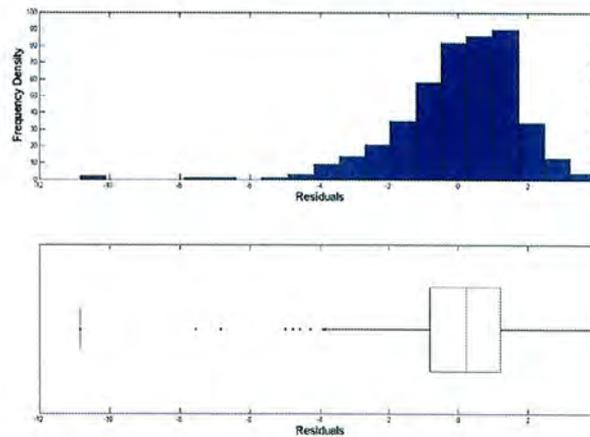
Table 3.7: Adjusted R^2 values for second regression model fit

Figure 3.11: Distribution of residuals

residuals \hat{e} to learn about the order of magnitude of the local variation term. We will describe fitting the model in detail in section 3.5.

If we consider the distribution of the residuals (see Figure 3.11), we notice there is a considerable positive skewness. This indicates that the regression model will in general predict greater wall loss than is observed. This skewness in the data can be expected as a result of the nature of our response variable. We are considering minimum wall thickness readings, and fitting a linear regression model, consequently, skewness should be expected in the distribution of the residuals. However, the lower tail of the distribution is very long, and these values represent where the model has underestimated the extent of wall loss. As can be seen from Figure 3.11, when the

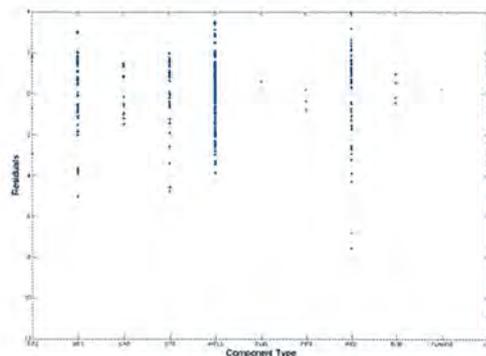


Figure 3.12: Residual plotted by Component type

wall loss is particularly extreme, the residuals become relatively very large meaning the model is likely to be very inaccurate.

The regression analysis highlights the importance of the *Component Type*, *PMC*, *Nominal Wall Thickness* and *Age* variables, which we now discuss in greater detail.

Component Type The *Component Type* variable contains information which tells us about the structural design of a component. This is a fixed variable which will not change throughout the life of the component. The component type cannot be changed without changing the component. *Component Type* is a categorical variable with 22 levels. 10 of these levels belong to the set of Key Points. Unlike other levels of the component type variable, Key Points are not defined by their physical characteristics, but are instead registered as points which it is important to inspect. We can recover the meaning of the *Component Type* variable by grouping together Key Points which share the same design to form new levels of the variable. These new levels are labelled Key Point (i) for $i = 0, \dots, 9$ and are treated in the same way as the original levels.

Despite the poor fit of the initial regression model, the *Component Type* variable could be seen to be influential, and the difference plots (Fig. 3.6) highlighted there were particularly unusual patterns for the levels of the variables which occur in the modelling Circuit Circuit A. Therefore *Component Type* will be used in constructing our model inputs.

When we consider the plot of residuals of the regression model fitted against

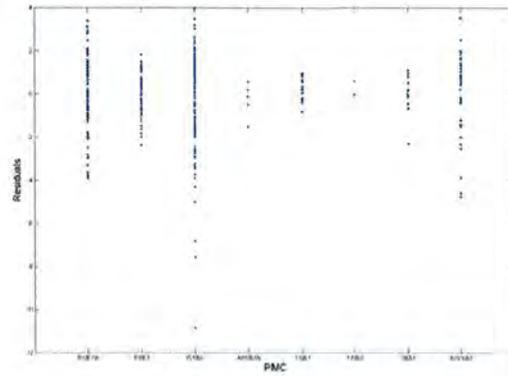


Figure 3.13: Residuals plotted by PMC level

their component type values (Figure 3.12), we can see that in general the spread of the distribution remains around the same order of magnitude for those variable which have been well observed.

PMC The *PMC* (pipe material code) variable describes from what the component is built. We know that all components are effectively a form of carbon steel, but this variable allows us to distinguish between different types of treatments and linings which have been applied to the steel. This is also a fixed variable which will not change its value unless the component is changed. There are 45 different *PMC* levels taken by components within the Site A data set, each well represented. Consequently, by the time we have restricted ourselves to considering corroding components which have been observed at least twice - a necessary condition for being able to estimate wall loss rate using our model - the amount of available data is not that large (see Table 3.13).

However, we can see from the regression modelling that *PMC* is an influential factor in determining wall loss behaviour, and so we will be using it to help determine our initial corrosion rate values. The residual plot for this variable indicates there are noticeable changes in variation between the well observed variable levels, with certain codes behaving in a much more predictable fashion than others.

Nominal Wall Thickness *Nominal Wall Thickness* is a continuous variable which tells us the initial wall thickness value for a new component. The *Nomi-*

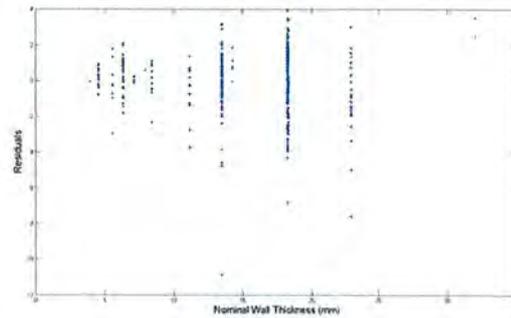


Figure 3.14: Residuals plotted against Nominal Wall Thickness

nal Wall Thickness values range from 3.91mm to 35mm and are required in the model to describe the scale information. The residual plot (Figure 3.14) indicates that the wall loss behaviour becomes less predictable as Nominal Wall Thickness increases.

Product It is reasonable to assume that the *Product* variable may be an influential variable. However, because all components in Circuit A are involved in the same production process, the effect of this variable will not be built into the model.

In a more detailed model, it would be reasonable to consider the effects of the interactions between the descriptive variables. In our case, we are restricting ourselves to producing a model which is of the correct order of magnitude, and based on clear trends in the auxiliary data, as our main objective is to provide a procedure for solving the inspection planning problem. For any detailed application, in which one was intending to apply the methods in practice, a more involved modelling strategy should be used which would take into account the full range of subtleties introduced by the complexity of the system.

We have seen that factors relating to a components design (*Component Type*) and construction (*PMC*) show evidence of influencing wall loss behaviour. Component age at inspection was also highlighted as a relevant factor by the regression. These were factors identified as potentially influential by the expert judgement. A more detailed analysis would allow us to determine exactly how influential each of these factors would be in fixing the wall loss rate for a particular component. Each of

these will be built into our model for the wall loss. The key aspect of estimating our model parameters will be determining the rate of wall loss, and the relationship between the wall loss rates of different components.

3.5 Illustrative model, based on the Site A data

The model will initially run over the time frame for which we have observations, which for Circuit A means 12/6/1998 to 08/04/2004. Each time step will be of length 6 months and no distinction is made as to which point within a time step observations are made. We will start our model with the half year beginning 01/07/1998, giving us a total of 12 time steps until the end of the data period in 2004. The dates corresponding to each time step are shown in Table 4.1. Time step $t = 0$ represents our initial belief state. We use any observations from this period to quantify the initial wall thickness values. All time steps referred to when discussing this model can be assumed to be of length 6 months. Similarly, all wall loss rates are given as their 6 monthly rate, i.e. the amount of wall loss we expect to occur in a single model time step.

3.5.1 Parameter assessment

Our objective is to fit an illustrative model for the purposes of demonstrating the inspection planning procedure. However, we still wish to fit a model which possesses similar characteristics to the real system, and behaves in a superficially similar fashion. We therefore use the auxiliary data to provide us with a guide to the scale of the quantities we want to model, rather than to estimate quantities directly. Clearly if such a process was to be implemented in practice, the model fitting stage would require a great deal more attention, but, for the purposes of this account, a model which has similar order of magnitude behaviour will be sufficient.

We first obtain values for our initial wall thickness values \underline{x}_0 . We have readings for 15 of the 40 corroding components at $t = 0$ and for these components we simply use

Nominal Wall Thickness	Estimated Mean Percentage Wall Loss	Standard Deviation of estimate	No. of Observations
$NWT \leq 7.2$	0.0771	0.0228	23
$7.2 < NWT \leq 12$	0.1492	0.0283	12
$12 < NWT \leq 18$	0.0495	0.0156	11
$18 < NWT$	0.0364	0.0290	10

Table 3.8: Mean Percentage wall loss at $t = 0$

these readings as our initial values. However, the remaining 25 values are calculated based on the typical wall loss behaviour seen at this time point in the auxiliary data set. The auxiliary data set contains 169 measurements taken during the $t = 0$ time step, of which 58 belong to components which are either experiencing wall loss at $t = 0$ or will go start suffering wall loss before $t = 12$. The percentage wall loss seen in the corroding components at time 0 has mean 0.0799 and variance 0.0102. The small number of observations means that considering any further breakdown of this subset into component types or materials offers little meaningful information. However we can split the data into groups by *Nominal Wall Thickness*, the results are shown in Table 3.8. The standard deviations of these estimates are generally large, relative to the estimated means, indicating that relying on the mean estimates may not be reasonable (Figure 3.16). If we consider the distribution of the percentage wall loss observed in components at $t = 0$ (Figure 3.15), we can see there is a heavy skew to the right, with a large number of components experiencing less than 5% wall loss. Given this information, it would be unsuitable to use the mean estimate (of 7.99%) to predict how much wall loss we can expect to have taken place at $t = 0$. Instead, we group the wall loss at $t = 0$ into three categories: no wall loss as yet, moderate wall loss and severe wall loss. We estimate the unknown \underline{x}_0 values as:

$$x_c = (1 - \pi_c)WT_{nom}$$

where

$$\pi_c = \left\{ \begin{array}{l} 0 \text{ with probability } 0.45 \\ 0.08 \text{ with probability } 0.33 \\ 0.27 \text{ with probability } 0.22 \end{array} \right\}$$

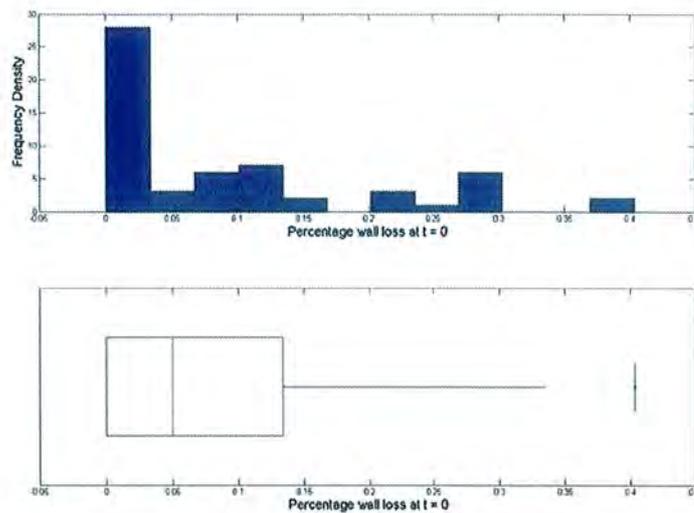


Figure 3.15: Distribution of WLR values of corroding components at $t = 0$

and the three values correspond to each of the wall loss states. These values are estimated from the auxiliary data, as are the probabilities of belonging to each group. Which group a component belongs to will be decided at random. [Expectation and variance of this constructed distribution are .0858 (against .0799 from data) and .0108 (against .0102).]

More data is available for the assessment of the wall loss rate. Our assumption that once wall loss has initiated, it continues at an approximately constant rate allows us to use data between any pair of observations for the same component. The rate of wall loss typically seen over the levels of the variables found in Circuit A is shown in Figure 3.16. These WLR values are calculated from the data as:

$$WLR_i = \frac{WT_{t+k} - WT_t}{k}$$

for all components in which the wall thickness (WT) is measured at least twice.

This is a similar quantity to that which we were using as the response variable in our regression analysis. Therefore it seems reasonable to make use of the regression model here. We use the model described earlier to predict our initial trend values based on the *Component Type*, pipe material code *PMC* and *Nominal Wall Thickness* variables. The *Age* will always be 1, as we are interested in predicting the half year rate for a component. For example, if we consider component 27 in Circuit C.

	Variable	Level
1	Component Type	Key Point(2)
2	Component Type	Bends (small)
3	Component Type	Cap
4	Component Type	Straight
5	Component Type	Weld
6	Component Type	Key Point (4)
7	PMC	1
8	PMC	2
9	Piping Schedule	PS2
10	Piping Schedule	PS3
11	Piping Schedule	PS1
12	Nominal Wall Thickness	$NWT \leq 4.6$
13	Nominal Wall Thickness	$4.6 < NWT \leq 6.4$
14	Nominal Wall Thickness	$6.4 < NWT \leq 7.2$
15	Nominal Wall Thickness	$7.2 < NWT \leq 12$
16	Nominal Wall Thickness	$12 < NWT \leq 18.2$
17	Nominal Wall Thickness	$18.2 < NWT \leq 18.3$
18	Nominal Wall Thickness	$18.3 < NWT$
19	ALL	ALL

Table 3.9: Key to variable levels in Figure 3.16

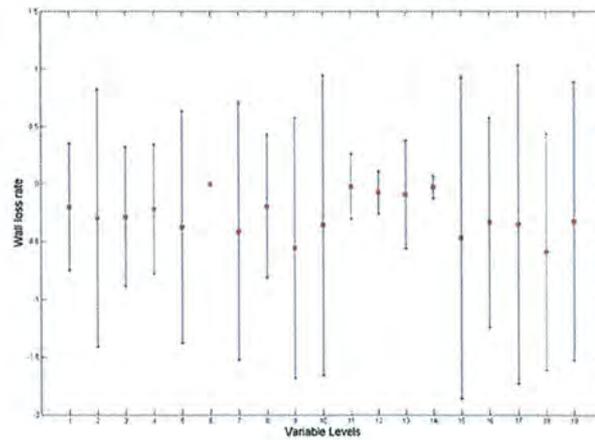


Figure 3.16: Mean wall loss rates for different variable levels, ± 2 standard deviations

This is a weld, with pipe maintenance code (*PMC*) 2 and nominal wall thickness 18.3mm. For this component, our predictor vector would be:

$$\underline{X}_{27} = [1 \quad \overbrace{000100}^{\text{ComponentType}} \quad \underbrace{100000}_{\text{PMC}} \quad \underbrace{1}_{\text{Age}} \quad \overbrace{18.3}^{\text{NominalWallThickness}} \quad]'$$

and we would obtain our prediction for α_{c0} for $c = 27$ from:

$$\alpha_{c0} = \underline{b}'\underline{X}_{27} = -2.9298 \quad (3.3)$$

where \underline{b} is the vector of regression coefficients given in Table 3.6. A rate of -2.9298mm per half year would be assigned to component 27 of Circuit C. This is a particularly high rate of wall loss, but if we consider the time series of observations made on 27 (shown in Figure 3.17), we see that this component is prone to large drops in wall thickness over short time periods. The anomalous upward jumps in wall thickness could be a consequence of changes in measurement policy between observations, or component replacement, or measurement error. However, this is difficult to confirm.

The initial wall loss rate values for Circuit A, and the variable information on which the calculations are based are shown in Tables 3.10 and 3.11.

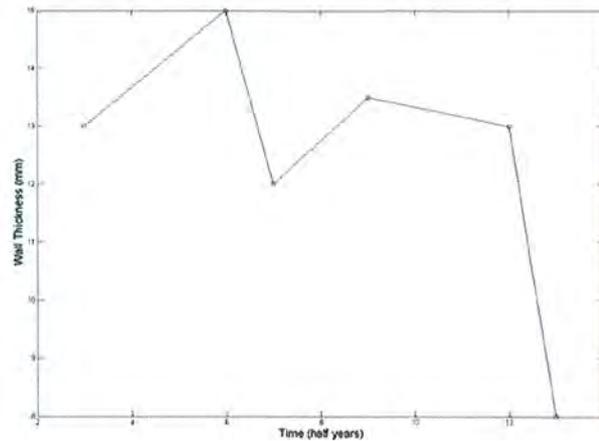


Figure 3.17: Time series of observations for component 27

3.5.2 Local Variation

The local variation term $\tilde{\epsilon}_t$ of our model is controlled through the variance matrix Σ_{ζ} . The value of $\tilde{\epsilon}_t$ for a given component c is defined to be the minimum of s_c independent sums of t independent draws from a Normal population $N(0, \sigma_{\zeta c}^2)$. For an n component system the matrix Σ_{ζ} is constructed as the diagonal $n \times n$ matrix in which the (c, c) -th element is $\sigma_{\zeta c}^2$. For simplicity of specification, we will not evaluate a value for all n components, but rather use a single estimate for the influence of local variation based on the regression residual values.

To assess the order of magnitude of the local variation matrix Σ_{ζ} , we have to consider the differences between what we would expect to happen if there is no local variation and what is observed in the data. The regression modelling provides us with a means of considering a related quantity in the form of the residuals. To relate the residual error to the local variation term we use the following procedure:

1. Identify a set of components (in the auxiliary data) which are observed at t and again at $t + k$. These components should be sufficiently far apart to minimise spatial interaction effects.
2. Extract the residuals $(\underline{e}_{t,t+k})$ for this subset from the full residual vector (\underline{e}) .

Component ID	Component Type	PMC level	Nominal Wall Thickness	α_{e0}
1	Bes	2	23	-2.4188
2	KP2	2	18.24	-1.6323
3	Bes	2	23	-2.1488
4	Bes	2	23	-2.1488
5	KP2	2	18.24	-1.6323
6	Bes	2	6.35	-0.2605
7	Bes	2	18.24	-1.6528
8	TEQ	2	23	-2.4875
9	Cap	2	18.24	-0.6816
10	Str	2	13.49	-0.9962
11	Str	1	13.11	-0.7874
12	Bes	1	13.11	-0.6796
13	Str	1	13.11	-0.7874
14	Bes	1	18.24	-1.5051
15	Str	1	23	-2.3788
16	Bes	2	23	-2.1488
17	KP2	2	18.24	-1.6323
18	KP2	2	18.24	-1.6323
19	KP4	2	4.78	-0.4444
20	KP2	2	18.24	-1.6323

Table 3.10: Variable information for Circuit A

Component ID	Component Type	PMC level	Nominal Wall Thickness	α_{e0}
21	KP4	2	8.74	-0.1928
22	KP4	2	8.74	-0.1928
23	Weld	2	6.35	-0.5863
24	Weld	2	6.35	-0.5863
25	KP4	2	7.10	-0.0711
26	KP4	2	7.10	-0.0711
27	KP4	2	4.78	-0.4444
28	KP2	2	18.24	-1.6323
29	KP2	2	18.24	-1.6323
30	Weld	2	4.78	-0.3337
31	Weld	2	8.74	-0.9709
32	Weld	1	13.50	-1.5892
33	Weld	2	7.14	-.07135
34	Weld	1	13.50	-1.5892
35	Weld	1	5.56	-0.3115
36	Weld	1	5.56	-0.3115
37	Weld	1	13.50	-1.5893
38	Weld	1	13.50	-1.5892
39	Weld	1	5.56	-0.3115
40	Weld	1	5.56	-0.3115

Table 3.11: Variable information for Circuit A

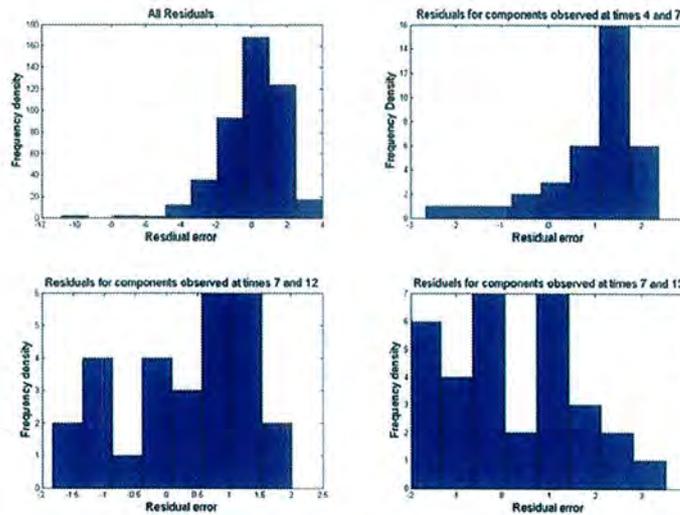


Figure 3.18: Distribution of residuals

3. Calculate the value for model local variation ($\hat{\sigma}_\zeta^2$) which would have produced the observed residual variation.
4. Check scale of $\hat{\sigma}_\zeta^2$ against estimates from other sets.

Three suitable subsets of observations were identified. These are the subsets of components observed at times 7 and 12, 7 and 13 and 4 and 7. A large number of components were observed at time step 7, of which some were observed again at time point 12 and some at time point 13. The subsets have been constructed so that components observed at time step 7 and then again at time points 12 and 13 only appear in one of the 7 - 12 or 7 - 13 subsets. The distribution of residuals in these subsets is shown in Figure 3.18.

The variance we are interested in learning about is not immediately related to the variance of the residuals. The residual variance is telling us about the variability we expect to see in the difference between observed wall thickness values at known time points. In our model this would be written (for a single component) as:

$$\begin{aligned} \tilde{\omega}_{t+k} - \tilde{\omega}_t &= \min_l \{r_{lt+k} + \xi_{lt+k}\} - \min_l \{r_{lt} + \xi_{lt}\} \\ &= \min_l \left\{ \sum_{i=1}^{t+k} \zeta_{li} + \xi_{lt+k} \right\} - \min_l \left\{ \sum_{i=1}^t \zeta_{li} + \xi_{lt} \right\}. \end{aligned}$$

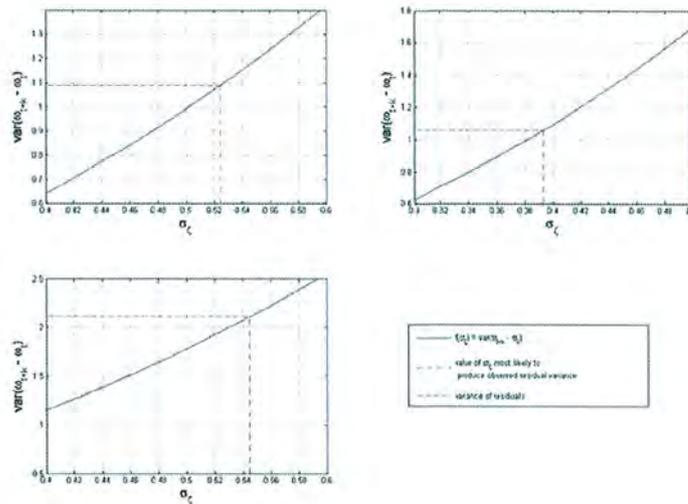


Figure 3.19: Identifying appropriate standard deviation values for the local variation term

We are interested in the behaviour of ζ_{li} . We have information for the approximate scale of the measurement error (ξ_t), so we need to know what value of σ_ζ^2 would produce a similar level of variation as that observed in the residuals. A simulation study allows us to obtain estimates for the level of variability in the quantity $\tilde{w}_{t+k} - \tilde{w}_t$ expected for different values of σ_ζ^2 , and hence obtain a most likely value for σ_ζ^2 . It is necessary to repeat this simulation for each different pair of time steps considered.

We generate realisations of the value of the quantity $\tilde{w}_{t+k} - \tilde{w}_t$ for a given σ_ζ^2 by simulating draws s from the Normal populations $N(0, t\sigma_\zeta^2 + \sigma_\xi^2)$ and $N(0, (t+k)\sigma_\zeta^2 + \sigma_\xi^2)$. We then take the difference between the minima of these two samples as a realisation of $\tilde{w}_{t+k} - \tilde{w}_t$. We repeat this process a large number of times and take the variance of our sample as an estimate for the variance we would expect to see for our given value of σ_ζ^2 . We then repeat this for many different values of σ_ζ^2 . From this simulated data we can identify the value of σ_ζ^2 that lies closest to our observed variance.

Using this procedure we obtain estimates for the value of σ_ζ^2 for each of the sets to be as shown in Table 3.12 (see also Figure 3.19).

There is a similarity between the values taken for the 4 - 7 and 7 - 13 subsets.

Subset	Residual variance	$\hat{\sigma}_\zeta^2$
4 - 7	1.0891	.2750
7 - 12	1.0588	.1547
7 - 13	2.1099	.2968

Table 3.12: Order of magnitude estimates for σ_ζ^2

If further data were available, it would be possible to assess whether there is any reason for this similarity, or if it is simply a chance occurrence. The observed residual variance value for the 7 - 12 subset seems small with respect to the other two values, but given the lack of available information (sample sizes are between 20 and 34), we would expect variation. We would expect the variance of the residual to increase the further apart measurements are taken in time, so the fact that the variance of the residuals in the 7 - 12 subset is slightly smaller than that of the 4 - 7 subset makes us suspicious of that value. Given our concerns about the validity of the variance for the 7 - 12 subset, the similarity of the values for the two sets, and the absence of further data to allow greater clarification, we choose our value by drawing it from the Normal population $N(.285, .0375^2)$. This is in line with the data, and is set to be smaller than the typical variance of the global term, so that the correlated global term dominates, which should produce more interesting decision problems.

3.5.3 Spatial effects

The spatial information available to us for Circuit A, is also available for the circuits which make up the auxiliary data. Testing for spatial association over the entire data set would be impractical, so we instead focus on the spatial association between components of some of the smaller corrosion circuits in the auxiliary data - B, C and E. We have suitable observational data (i.e. time series of at least two points with no increases in wall thickness) for 16 points for B, 23 points for C and 10 points for E. As stated previously, the schematic diagrams for the circuits are topologically accurate, but offer no information on scale or actual location of an individual circuit. We have no physical measure of distance, but we can ascertain

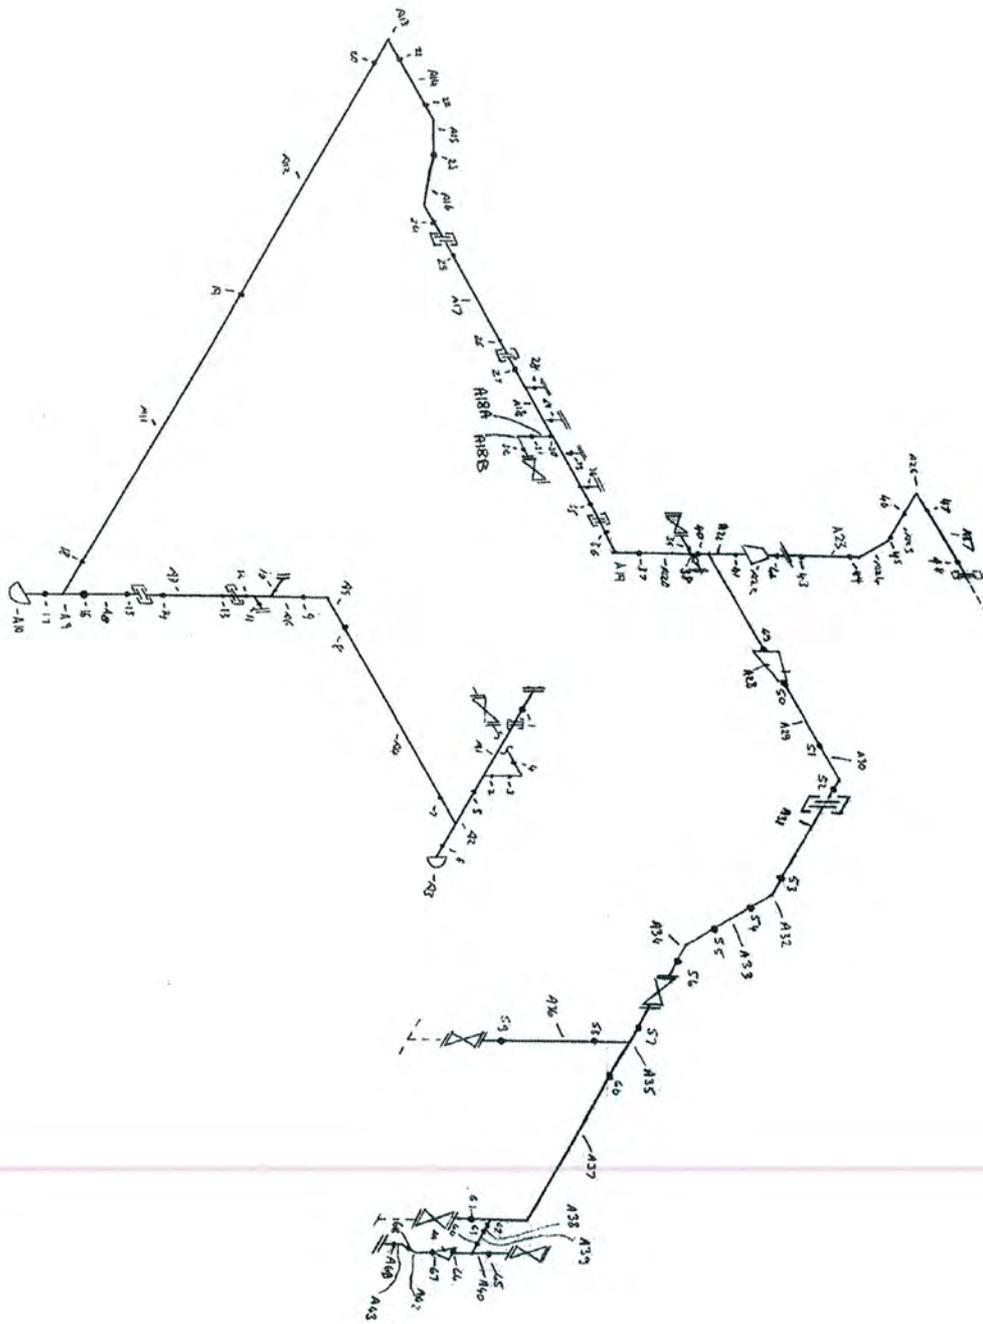
where components are placed in relation to each other, so we use adjacency as our measure of space. The adjacency matrix, A_{adj_C} (for Circuit C) is defined as previously and illustrated in Figure 3.21.

We test for spatial association by calculating the correlation between all components a known distance from each other. We begin by evaluating the correlation between all immediately adjacent components, the move to components separated from each other by an intermediate component, and progressively increase the distance. As we move further away, observations are pooled due to lack of data, so our final groups are assessed as, for example, the correlation between all components between 15 and 20 components removed from each other rather than as all components separated by 20 components. The effect of distance on correlation for Circuits B, C and E is illustrated in Figure 3.22

These results imply that as we increase the distance between components the correlation between their wall thickness readings decreases, eventually descending to approximately 0. The relationship between correlation and distance implies there may be a spatial pattern to wall loss, and as such it is not unreasonable to base our correlation/covariance structures on distance, as described in section 2. We proposed a model for the correlated evolution deviation matrix of the form (2.24):

$$\sigma_{h_k cc'} = \sigma_{h_k c} \sigma_{h_k c'} \exp\{-\tau_{h_k} \|c - c'\|^2\}$$

where h_k was a characteristic determining wall loss behaviour. In our model these will be the *Component Type* and *Pipe Material Code* variables. The values of $\sigma_{h_k c}$ are equivalent to the variance for a particular level of the variable. These are given in Table 3.13 and can be estimated directly from the data. The order of magnitude of the covariance between variables separated by equivalent distances has been calculated as part of finding the correlation between the same sets. This provides us with a guide to scale we should be restricting to when defining $\sigma_{h_k cc'}$. The values of τ_{h_k} can be used to achieve an appropriate rescaling.



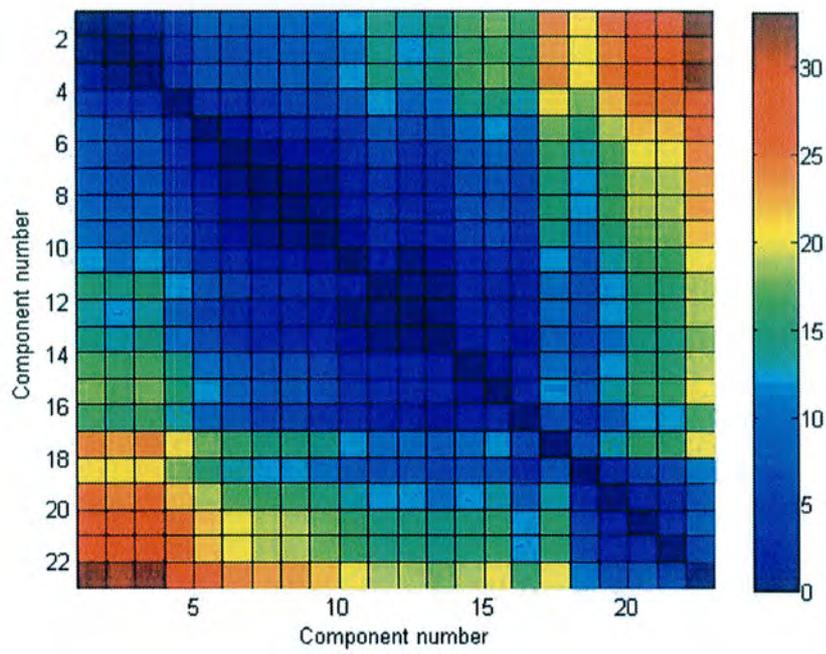


Figure 3.21: Adjacency matrices for Circuit C

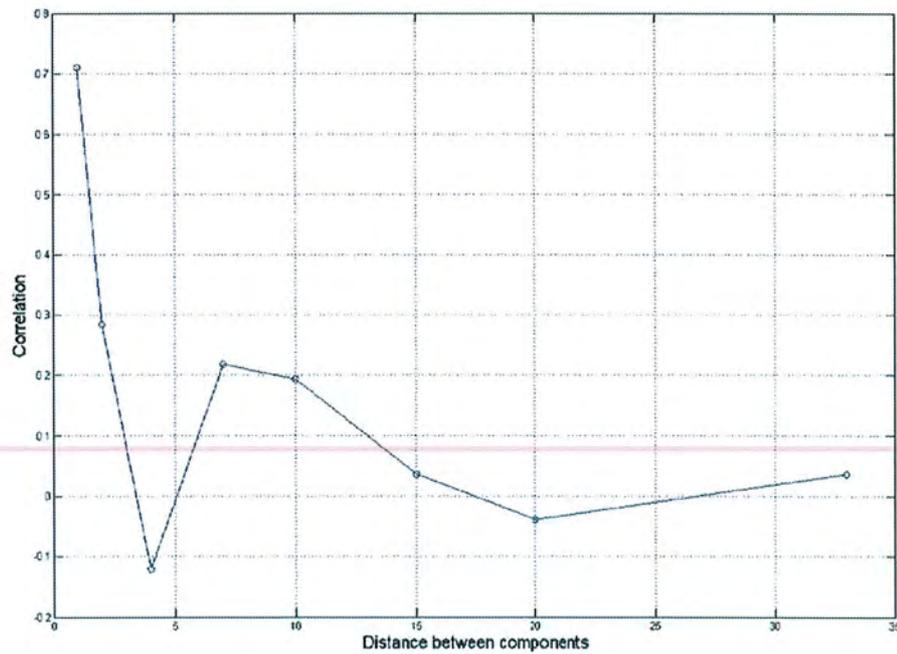


Figure 3.22: Correlation values in Circuits B, C and E against distance between components

Component Type	Variance of <i>WLR</i>	No. of Observations
KeyPoint(2)	0.0750	5
Bends (small)	0.3116	23
Cap	0.0908	7
Straight	0.0781	20
Weld	0.2506	51
TEQ	NaN	0
Branch	0	2
PMC Level	Variance of <i>WLR</i>	No. of Observations
1	0.3105	20
2	0.0954	8
Nominal Wall Thickness	Variance of <i>WLR</i>	No. of Observations
$NWT \leq 4.6$	0.0081	9
$4.6 < NWT \leq 6.4$	0.0544	11
$6.4 < NWT \leq 7.2$	0.0022	11
$7.2 < NWT \leq 12$	0.4838	17
$12 < NWT \leq 18.2$	0.2044	27
$18.2 < NWT \leq 18.3$	0.4742	105
$18.3 < NWT$	0.2624	10

Table 3.13: *WLR* variance values by descriptive variable level

Chapter 4

Bayes linear updating for the corrosion model

We are adopting a Bayesian approach to modelling and inspection planning as it provides a natural framework for the combination of informed expert judgement and observed data, both of which we believe to offer valuable information for modelling wall loss behaviour. In this chapter we discuss how we update our beliefs as we receive new data. Section 4.1 considers the wall loss problem and the necessary updates. In section 4.2 we offer an alternative solution to a full Bayesian update - the Bayes linear update - and consider its relevance to this problem. Section 4.2 discusses the different ways in which the Bayes linear approach could be interpreted for the wall loss problem and the advantages of using Bayes linear methods for inspection planning problems. We detail how we will apply the updating strategy to our wall loss example (Circuit A), and the implications for our simulations in section 4.3.

4.1 Updating the Model

The model takes observation values, \underline{y}_{dt} , defined as in section 2.3.1, which are functions of a subset of the underlying true wall thickness values, \underline{u}_{dt} , and an ob-

ervation error $\underline{\xi}_{dt}$, where each term is a vector over components and the locations and components to be observed are specified by design d . Recalling the model form from section 2.3.1, an observation for a particular component, c , is defined as:

$$\begin{aligned} y_{ct} &= f(u_{l_dct} + \xi_{l_dct}) \\ u_{l_dct} &= x_{ct} + r_{l_dct} \end{aligned}$$

We want to use the observable quantity y_{ct} to learn about the behaviour of u_{lct} . Our model for u_{lct} consists of the global trend term x_{ct} and the local deviation term r_{lct} , about which we have made certain distributional assumptions (described fully in chapter 2). These assumptions provide us with beliefs about the behaviour of the underlying true wall thickness u_{lct} at location l in component c at time t , and it is our beliefs about this quantity that we wish to update.

We will usually observe more than one component. The set of observations made at a particular inspection are denoted \underline{y}_{dt} , where \underline{y}_{dt} is either a vector, containing an element for each $c_d \in C_d$ (the set of components included in design d), or a matrix, in which each row describes the observations made for component c_d , depending on the choice of observation function. We use the subscript d as a shorthand for l_dc_d , which would provide a complete, but more cumbersome notation. The quantity \underline{y}_{dt} is therefore defined as a vector over c_d :

$$\underline{y}_{dt} = f(\underline{u}_{l_d t} + \underline{\xi}_{l_d t}) \quad (4.1)$$

where $\underline{u}_{l_d t}$ and $\underline{\xi}_{l_d t}$ are vectors over c_d

In general, we may want to learn about the relationship between our observations and any other function of the underlying wall thickness values. Commonly, this function will be the same as the observation function, but without measurement error. In our case f is componentwise minimisation, and we will be interested in the relationship between our observations, \underline{y}_{dt} , and the true minimum wall thickness value for each component, \underline{m}_t , defined as:

$$\underline{m}_t = f(\underline{u}_t) \quad (4.2)$$

where \underline{u}_t is a vector over all c .

This relationship can be described by the (posterior) distribution $p(\underline{m}_t|\underline{y}_{dt})$, evaluated as:

$$p(\underline{m}_t|\underline{y}_{dt}) = \frac{p(\underline{y}_{dt}|\underline{m}_t)p(\underline{m}_t)}{p(\underline{y}_{dt})} \quad (4.3)$$

The distributions of $p(\underline{y}_{dt}|\underline{m}_t)$, $p(\underline{m}_t)$ and $p(\underline{y}_{dt})$ will depend on the observation function f and the measurement error term ξ_{dt} , in addition to the distributions of the model terms, $p(\underline{x}_t)$, and $p(\underline{r}_t)$. The model structure gives us the distributions for $p(\underline{x}_t)$, and $p(\underline{r}_t)$. However, this will not always be straightforward to relate to $p(\underline{y}_{dt}|\underline{m}_t)$, $p(\underline{m}_t)$ and $p(\underline{y}_{dt})$. As discussed for the minimisation case in chapter 2, the distribution of our observations \underline{y}_{dt} and component minimum \underline{m}_t are difficult to write down analytically. Similarly, the likelihood function $p(\underline{y}_{dt}|\underline{m}_t)$ for minimisation is not easy to write down. Consequently, for the case of choosing component wise minimisation as the observation function, f , we are forced into using a simulation approach to evaluate the posterior distribution, as we cannot derive any of the components of equation (4.3) analytically.

We may also be interested in updating our beliefs about individual model terms, such as the trend term, \underline{x}_t or the minimum of the local deviation terms \tilde{r}_t . This allows us to use new data to learn about model behaviour rather than a particular function of the model. This will be useful if we are interested in producing an improved model for forecasting future wall thickness behaviour rather than a particular function of wall thickness at a given time. The equations for calculation of the relevant posterior distributions for these updates are:

$$p(\underline{x}_t|\underline{y}_{dt}) = \frac{p(\underline{y}_{dt}|\underline{x}_t)p(\underline{x}_t)}{p(\underline{y}_{dt})} \quad (4.4)$$

$$p(\tilde{r}_t|\underline{y}_{dt}) = \frac{p(\underline{y}_{dt}|\tilde{r}_t)p(\tilde{r}_t)}{p(\underline{y}_{dt})}. \quad (4.5)$$

For the case of the componentwise minimisation we encounter the same problems in evaluating the posterior distributions as we had for (4.3). Performing these updates also requires numerical/simulation methods.

This is a consequence of the choice of observation function. If the observation function is sufficiently simple, then it may be possible to perform full Bayesian updates in closed analytic form. However, for many choices of observation function,

f , we will not be able to carry out our updates analytically, and instead have to rely on numerical/simulation methods to evaluate the posterior distributions.

This reliance on simulation methods for updating distributions will affect the tractability of the inspection planning process. Given that we want to consider models for complicated systems with large numbers of components, evaluating potential benefits for a large number of different inspection designs would become very time consuming, as, for each design, we would need to generate many samples and for each such sample we must evaluate the posterior distribution by simulation. In section 4.2 we discuss an alternative Bayesian updating approach and review its merits for inspection design applications

4.2 The Bayes linear approach

The Bayes linear approach to updating is a suitable way of handling problems in which beliefs have only been partially specified. These partial belief specifications arise because it is either unnecessary or inappropriate to specify a complete prior probability density. This could be because our partial belief specifications are sufficient to characterise the behaviour of the quantity of interest, or because it is felt that assigning a full density to our random quantity is making a more definite statement about our beliefs than the available information allows. For an overview of the motivations and principles of the Bayes linear approach see [22].

Bayes linear methods use expectation as a basis for making quantitative statements about random quantities instead of probability. A Bayes linear prior specification consists of the first and second order moments for the vector of interest, B , namely, an expectation, $E(B)$ and variance $\text{var}(B)$ for the random quantity. We will also have to be able to specify values for the covariance structure between our random quantity, B , and the observable vector, D . The general form of the Bayes linear updating equations, for a specific observation, d , when $\text{var}(D)$ is positive definite, is:

$$E_D(B) = E(B) + \text{cov}(B, D) \text{var}^{-1}(D) [d - E(D)] \quad (4.6)$$



$$\text{var}_D(B) = \text{var}(B) - \text{cov}(B, D) \text{var}^{-1}(D) \text{cov}(D, B). \quad (4.7)$$

In general, the Bayes linear equations adjust our beliefs about a random quantity, B , linearly, by data, D . B and D can be scalar or vector quantities. To illustrate this, we consider a small example. Suppose we have a collection of 3 related quantities \underline{B} , with current expectation and variance:

$$E(\underline{B}) = \begin{bmatrix} 7 \\ 4 \\ 2 \end{bmatrix} \quad \text{var}(\underline{B}) = \begin{bmatrix} 4 & 0.8 & 0.05 \\ 0.8 & 1 & 0.2 \\ 0.05 & 0.2 & 0.25 \end{bmatrix}$$

Suppose further that there exists a related observable quantity, \underline{D} , with expectation $E(\underline{D})$ and variance $\text{var}(\underline{D})$:

$$E(\underline{D}) = \begin{bmatrix} 7 \\ 5 \\ 1 \end{bmatrix} \quad \text{var}(\underline{D}) = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 1.25 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

whose relationship to B can be characterised via the covariance matrix,

$$\text{cov}(\underline{B}, \underline{D}) = \begin{bmatrix} 1.5 & 0.4 & -0.3 \\ 0.4 & 0.6 & 0.1 \\ -0.3 & 0.1 & 0.2 \end{bmatrix}$$

\underline{D} can be considered as an observation of \underline{B} made using an inaccurate measuring tool, or as an observation of a different, but dependent quantity (for example, an individual's height can be used as a predictor of his weight). The covariance matrix $\text{cov}(\underline{B}, \underline{D})$ is required to quantify the relationship between the two quantities.

How would observing a particular realisation, \underline{d} , of the related observable quantity, \underline{D} , change our beliefs about \underline{B} ? Assume we observe,

$$\underline{d} = \begin{bmatrix} 8 \\ 6 \\ 0 \end{bmatrix}$$

we can evaluate the adjusted expectation (using (4.6)) as:

$$E_D(\underline{B}) = E(\underline{B}) + \text{cov}(\underline{B}, \underline{D}) \text{var}^{-1}(\underline{D}) [\underline{d} - E(\underline{D})]$$

$$\begin{aligned}
 E_{\underline{D}}(\underline{B}) &= \begin{bmatrix} 7 \\ 4 \\ 2 \end{bmatrix} + \begin{bmatrix} 1.5 & 0.4 & -0.3 \\ 0.4 & 0.6 & 0.1 \\ -0.3 & 0.1 & 0.2 \end{bmatrix} \begin{bmatrix} 0.5 & 0 & 0 \\ 0 & 0.8 & 0 \\ 0 & 0 & 1 \end{bmatrix} \left[\begin{pmatrix} 8 \\ 6 \\ 0 \end{pmatrix} - \begin{pmatrix} 7 \\ 5 \\ 1 \end{pmatrix} \right] \\
 &= \begin{bmatrix} 8.37 \\ 4.58 \\ 1.73 \end{bmatrix}.
 \end{aligned}$$

We can see that the Bayes linear update has acted here to increase the expectation for those elements for which a larger than expected value was observed, and reduce the expectation for the element in which the observed value was lower than expected. Using (4.7) we find the adjusted variance to be:

$$\text{var}_{\underline{D}}(\underline{B}) = \begin{bmatrix} 2.657 & 0.338 & 0.303 \\ 0.338 & 0.622 & 0.192 \\ 0.303 & 0.192 & 0.157 \end{bmatrix}.$$

Note that the calculation of the adjusted variance does not depend on the particular values of \underline{d} observed. So given any set of observations in which we see a value for each element of our collection, \underline{D} , the value of the adjusted variance, $\text{var}_{\underline{D}}(\underline{B})$ will be the same.

Suppose we only see the observations for the first and third elements of our collection. To update our beliefs, we identify the parts of $\text{var}(\underline{D})$ and $\text{cov}(\underline{B}, \underline{D})$ which relate to the observed elements. So for a different set of observations, \underline{d} , in which the first and third elements are observed, the expectation $E(\underline{D})$, variance $\text{var}(\underline{D})$ and covariance $\text{cov}(\underline{B}, \underline{D})$ are taken to be:

$$\begin{aligned}
 E(\underline{D}) &= \begin{bmatrix} 7 \\ 1 \end{bmatrix} \\
 \text{var}(\underline{D}) &= \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix} \\
 \text{cov}(\underline{B}, \underline{D}) &= \begin{bmatrix} 1.5 & -0.3 \\ 0.4 & 0.1 \\ -0.3 & 0.2 \end{bmatrix}.
 \end{aligned}$$

These values are defined by identifying and extracting the first and third elements of the full mean vector, the first and third columns of the full covariance matrix and the values of the full variance matrix relating to the first and third elements.

Updating for this data, we find:

$$\begin{aligned} E_{\underline{D}}(\underline{B}) &= \begin{bmatrix} 7 \\ 4 \\ 2 \end{bmatrix} + \begin{bmatrix} 1.5 & -0.3 \\ 0.4 & 0.1 \\ -0.3 & 0.2 \end{bmatrix} \begin{bmatrix} 0.5 & 0 \\ 0 & 1 \end{bmatrix} \left[\begin{pmatrix} 8 \\ 0 \end{pmatrix} - \begin{pmatrix} 7 \\ 1 \end{pmatrix} \right] \\ &= \begin{bmatrix} 8.05 \\ 4.10 \\ 1.65 \end{bmatrix} \end{aligned}$$

and

$$\text{var}_{\underline{D}}(\underline{B}) = \begin{bmatrix} 2.785 & 0.530 & 0.335 \\ 0.530 & 0.910 & 0.240 \\ 0.335 & 0.240 & 0.165 \end{bmatrix}.$$

We can see that the updated values for this data show similar behaviour to those of the full update. Our beliefs about the second element of \underline{B} have been updated, despite us having not observed it. In comparison to the results achieved from the update based on the full set of observations, we have not been able to reduce the variance as much, which is consistent with us not having seen as much.

4.3 Adjusted assessments

In this section we illustrate the updating calculations we advocate using the Circuit A example. We update the model for all available data by performing a sequence of updates based upon the order in which we receive data. For a general data set, observed over the time period $[0, t]$, we receive data at times $\{t_1, \dots, t_\omega\}$. These data will, in general, be a partial observation of the system at time t_i , denoted \underline{y}_{dt_i} . Depending on the goals of the update we will either want to learn about the underlying model behaviour (through updating our beliefs about $\underline{\alpha}_{t_i}$ and \underline{x}_{t_i}), or learn about the system minima (by updating for \underline{m}_{t_i}). For the initial process of

model improvement, we will be interested in updating $\underline{\alpha}_{t_i}$ and \underline{x}_{t_i} , using the Bayes linear updating equations:

$$E_{\underline{y}_{dt_i}}(\underline{x}_{t_i}) = E(\underline{x}_{t_i}) + \text{cov}(\underline{x}_{t_i}, \underline{y}_{dt_i}) \text{var}^{-1}(\underline{y}_{dt_i}) [\underline{y}_{dt_i} - E(\underline{y}_{dt_i})] \quad (4.8)$$

$$E_{\underline{y}_{dt_i}}(\underline{\alpha}_{t_i}) = E(\underline{\alpha}_{t_i}) + \text{cov}(\underline{\alpha}_{t_i}, \underline{y}_{dt_i}) \text{var}^{-1}(\underline{y}_{dt_i}) [\underline{y}_{dt_i} - E(\underline{y}_{dt_i})] \quad (4.9)$$

$$\text{var}_{\underline{y}_{dt_i}}(\underline{x}_{t_i}) = \text{var}(\underline{x}_{t_i}) - \text{cov}(\underline{x}_{t_i}, \underline{y}_{dt_i}) \text{var}^{-1}(\underline{y}_{dt_i}) \text{cov}(\underline{y}_{dt_i}, \underline{x}_{t_i}) \quad (4.10)$$

$$\text{var}_{\underline{y}_{dt_i}}(\underline{\alpha}_{t_i}) = \text{var}(\underline{\alpha}_{t_i}) - \text{cov}(\underline{\alpha}_{t_i}, \underline{y}_{dt_i}) \text{var}^{-1}(\underline{y}_{dt_i}) \text{cov}(\underline{y}_{dt_i}, \underline{\alpha}_{t_i}). \quad (4.11)$$

This requires us to have values for $\text{var}(\underline{y}_{dt_i})$, $\text{cov}(\underline{x}_{dt_i}, \underline{y}_{dt_i})$, $\text{cov}(\underline{y}_{dt_i}, \underline{\alpha}_{dt_i})$, plus the initial expectations and variances for \underline{x}_{dt_i} ($E(\underline{x}_{dt_i})$ and $\text{var}(\underline{x}_{dt_i})$) and $\underline{\alpha}_{dt_i}$ ($E(\underline{\alpha}_{dt_i})$ and $\text{var}(\underline{\alpha}_{dt_i})$). We obtain assessments of these quantities by simulation. Using the process described in section 2.8, and the initial conditions obtained in section 3.5, we simulate S realisations of the system for the time interval $[0, t+k]$, where $[0, t]$ covers the period for which data are available and $(t, t+k]$ is the period into the future we wish to consider. We can then use the simulation output to provide assessments of the required quantities, and thus carry out our updates.

The simulations provide S realisations based on our initial beliefs for all inspection times t_i . That is, S values of each model term we want to update, \underline{x}_t and $\underline{\alpha}_t$, and S realisations of our system observations, \underline{y}_{-t_i} . This simulation output can be used to provide assessments of the covariances between the observations and the underlying model terms, as described in section 2.8.

We are interested in updating our beliefs about the underlying model at time t , the final point for which we have data. Therefore we must assess the values of the covariances between the observations in the components for which we have observations and the system level at time t , $\text{cov}(\underline{x}_t, \underline{y}_{dt_i})$. Similarly, we must assess the covariance between our observations and the system slope, at time t , $\text{cov}(\underline{\alpha}_t, \underline{y}_{dt_i})$. It is also necessary to provide an assessment of the variance matrix of the observations, $\text{var}(\underline{y}_{dt_i})$, which describes the extent of the variation for a particular observation and also the relationships between the different observations. The process can be summarised as follows:

1. Simulate S realisations of the system from time 0 to time $t+k$, using initial simulation conditions as described in section 2.8.

2. Obtain assessments of the necessary expectations, $E(\underline{\alpha}_t)$, $E(\underline{x}_t)$ and $E(\underline{y}_{dt_i})$, as the vector with elements defined as:

$$E(\alpha_{ct}) = \frac{\sum_{k=1}^S \alpha_{ctk}}{S}, \quad E(x_{ct}) = \frac{\sum_{k=1}^S x_{ctk}}{S}, \quad \forall c$$

$$E(y_{c_d t_i}) = \frac{\sum_{k=1}^S y_{c_d t_i k}}{S} \quad \forall c_d \in C_d$$

3. Assess values of the required variance/covariance matrices $\text{var}(\underline{\alpha}_t)$, $\text{var}(\underline{x}_t)$, $\text{var}(\underline{y}_{dt_i})$, $\text{cov}(\underline{y}_{dt_i}, \underline{\alpha}_t)$ and $\text{cov}(\underline{y}_{dt_i}, \underline{x}_t)$ as the matrices with (j, j') -th element defined to be:

$$\sigma_{\nu_j \nu_{j'}} = \frac{\sum_{k=1}^S \nu_{jk} \nu_{j'k}}{S} - \left(\frac{\sum_{k=1}^S \nu_{jk}}{S} \cdot \frac{\sum_{k=1}^S \nu_{j'k}}{S} \right) \quad (4.12)$$

where $\nu_j, \nu_{j'}$ represent the pair of variables under assessment.

4. Calculate $E_{\underline{y}_{dt_i}}(\underline{x}_t)$ and $E_{\underline{y}_{dt_i}}(\underline{\alpha}_t)$ using (4.8) and (4.9).
 5. Calculate $\text{var}_{\underline{y}_{dt_i}}(\underline{x}_t)$ and $\text{var}_{\underline{y}_{dt_i}}(\underline{\alpha}_t)$ using (4.10) and (4.11).

4.3.1 Data for Circuit A

We wish to update Circuit A for all available data using the procedure outlined above. Circuit A is a corrosion circuit consisting of 40 corroding components, 6 non-corroding components and 73 components for which there is no data. We are modelling wall loss behaviour in the 40 corroding components using a dynamic linear model with time steps of length 6 months. Our model starts running at 01/01/98 at steps in half yearly intervals to 01/01/04 (see Table 4.1). We have observational data for Circuit A taken during time steps $\{5, 9, 10, 11, 12\}$. In total there are 46 observations and every corroding component is observed at least once. Full details of Circuit A and the wall loss model for this example are given in section 3.4 and section 3.5.

We are interested in updating our beliefs about the expectations and variances of \underline{x}_t and $\underline{\alpha}_t$. t is taken to be equivalent to the final point for which real inspection data are available, which for this example is $t = 12$. These adjusted belief specifications

Time Step	Start Date	End Date
0	01/01/1998	30/06/1998
1	01/07/1998	31/12/1998
2	01/01/1999	30/06/1999
3	01/07/1999	31/12/1999
4	01/01/2000	30/06/2000
5	01/07/2000	31/12/2000
6	01/01/2001	30/06/2001
7	01/07/2001	31/12/2001
8	01/01/2002	30/06/2002
9	01/07/2002	31/12/2002
10	01/01/2003	30/06/2003
11	01/07/2003	31/12/2003
12	01/01/2004	30/07/2004

Table 4.1: Model Time Steps

will then be used in subsequent simulations as our best estimates for the true values of \underline{x}_{12} and $\underline{\alpha}_{12}$. Ideally, we would also like to improve our estimates for the covariance structures Σ_x and Σ_α through updating. However, this is a more involved process than updating for expectations, requiring us to construct our simulation differently, and store different information in the simulation output. [69], [68] provide a method for updating covariance structures using Bayes linear methods, but for this thesis we limit ourselves to updating expectations, as this will form the basis for a tractable inspection methodology for large systems.

The observational data \underline{y}_{dt_i} for Circuit A are shown in Figures 4.1 - 4.5. Each plot shows the value of the observed data (and which components were observed) at a different observation time. The observed data (\underline{y}_{dt_i}) are marked as red circles. The black line shows the initial wall thickness values for each component and the green line shows the expected value of \underline{y}_{dt_i} according to the model. These points are marked with a grey area indicating a distance of 3 standard deviations either side of the expectation. These indicate the plausibility of the observations given the

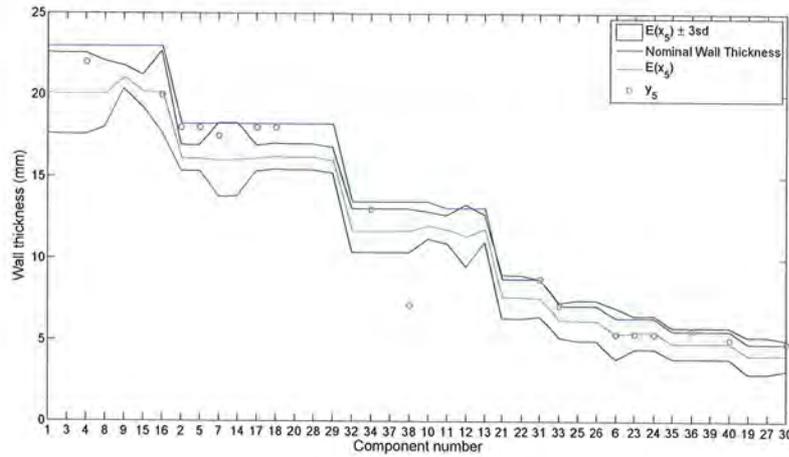


Figure 4.1: Observational data for $t = 5$, plus simulated expected range for observations

model. We assume components to have been 100% observed and the minimum wall thickness value recorded. This assumption cannot be verified from the available data, but is assumed to be a plausible assumption about the nature of the observation procedure.

4.3.2 Estimating the necessary means and variance for updating

At the first observation time, $t = 5$, 17 observations are taken at different components as indicated in Figure 4.1. We call this set C_d , in line with previous notation, and denote by c_{d_j} the j -th member of C_d . Our simulation output takes the form of $n \times S$ matrices \mathbf{Y}_t , \mathbf{M}_t , \mathbf{A}_t and \mathbf{X}_t for each time point. In this example, the number of components, $n = 40$, and the number of simulation realisations, $S = 100000$. We construct the 40×17 matrix $\text{cov}(\underline{\alpha}_{12}, \underline{y}_{d5})$ elementwise, where the (j, j') -th element of $\text{cov}(\underline{\alpha}_{12}, \underline{y}_{d5})$ is defined as in equation (4.12), with $\nu_j = \alpha_{e12}$ and $\nu'_{j'} = y_{c_{d5}}$. The full 40×46 covariance structure $\text{cov}(\underline{\alpha}_{12}, \underline{y}_{dt_i})$ is constructed similarly, using the information from all other time points.

The construction of $\text{cov}(\underline{x}_{12}, \underline{y}_{d5})$ follows the same procedure and the variance ma-

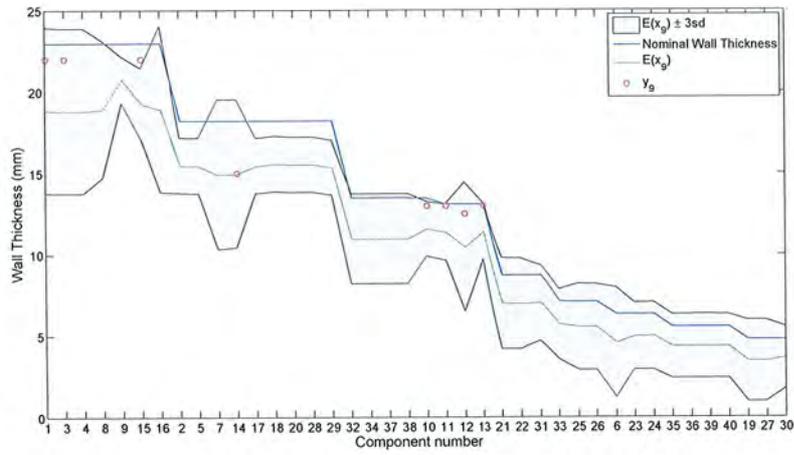


Figure 4.2: Observational data for $t = 9$, plus simulated expected range for observations

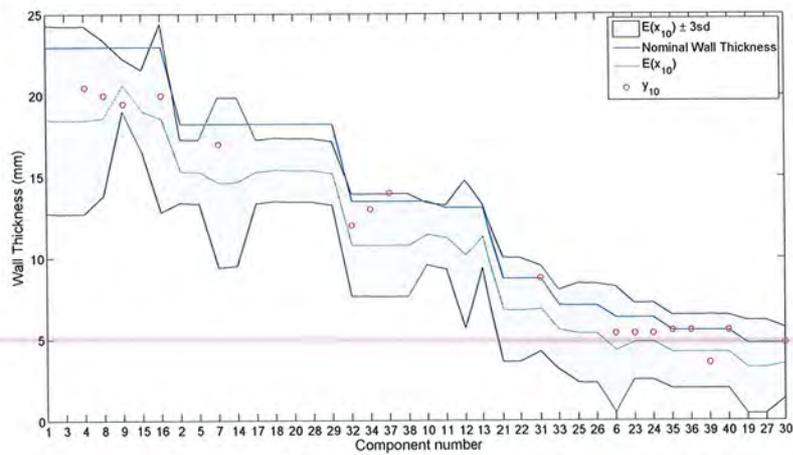


Figure 4.3: Observational data for $t = 10$, plus simulated expected range for observations

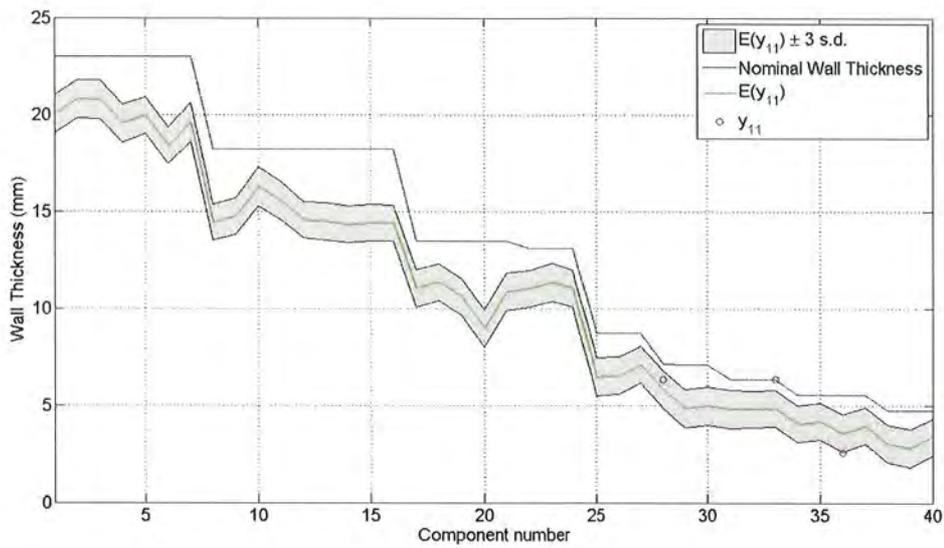


Figure 4.4: Observational data for $t = 11$, plus simulated expected range for observations

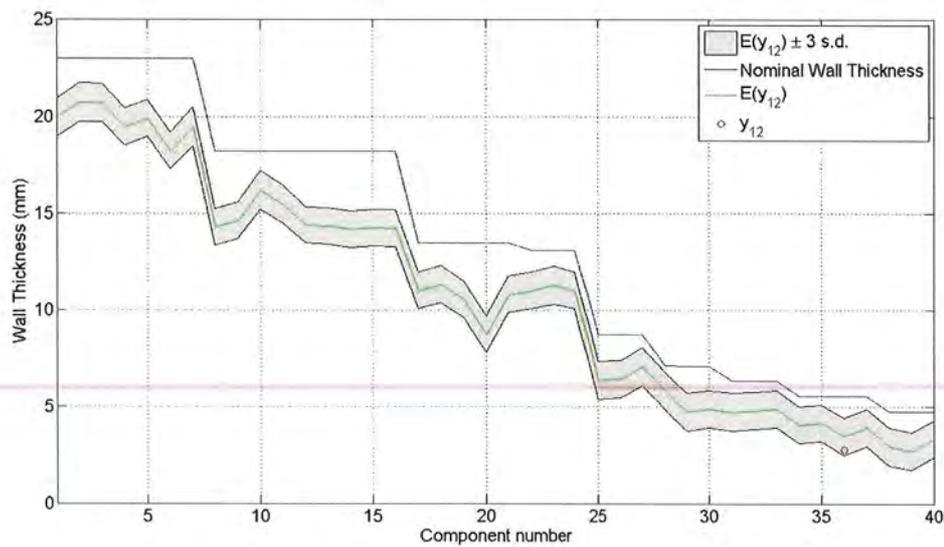


Figure 4.5: Observational data for $t = 12$, plus simulated expected range for observations

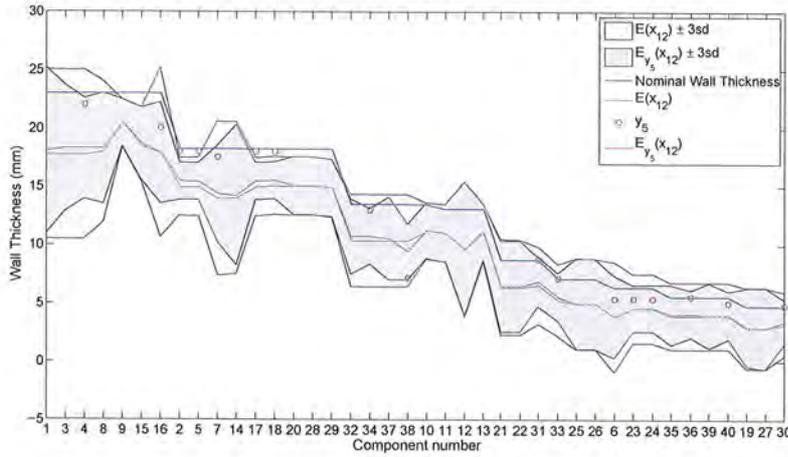


Figure 4.6: Updating beliefs for x_{12} given y_{d5}

trix $\text{var}(\underline{y}_{dt_i})$ is also evaluated elementwise using (4.12). It is necessary to construct a full 46×46 matrix for $\text{var}(\underline{y}_{dt_i})$ so that the relationships between observations made at different times are taken into account. This is particularly relevant when a component has been inspected twice. $\text{var}(\underline{\alpha}_{12})$ and $\text{var}(\underline{x}_{12})$ are also estimated using the same method.

Estimation of the necessary expectations, $E(\underline{\alpha}_{12})$, $E(\underline{x}_{12})$ and $E(\underline{y}_{dt_i})$ is more straightforward. We estimate $E(\underline{\alpha}_{12})$, $E(\underline{x}_{12})$ as the sum of the columns of \mathbf{A}_{12} and \mathbf{X}_{12} divided by $S = 100000$. $E(\underline{y}_{dt_i})$ is estimated as the sum of the relevant columns of \mathbf{Y}_{t_i} divided by $S = 100000$. For this example, $E(\underline{y}_{dt_i})$ will be a vector of length 46.

4.3.3 Updating for Circuit A

We illustrate the effects of updating by the available historical data, y_h , with Figures 4.6 - 4.10. Each plot has the same format, the blue line represents the initial wall thickness, the green line the expected value of the \underline{x}_t at $t = 12$, the magenta line shows the effect of adjusting for all the available data up to the time stated in the figure. The darker grey region indicates the value of $E_{y_5}(\underline{x}_{12}) \pm 3\sqrt{\text{var}_{y_5}(\underline{x}_{12})}$ and the lighter grey region shows the values covered by $E(\underline{x}_{12}) \pm 3\sqrt{\text{var}(\underline{x}_{12})}$. The plots

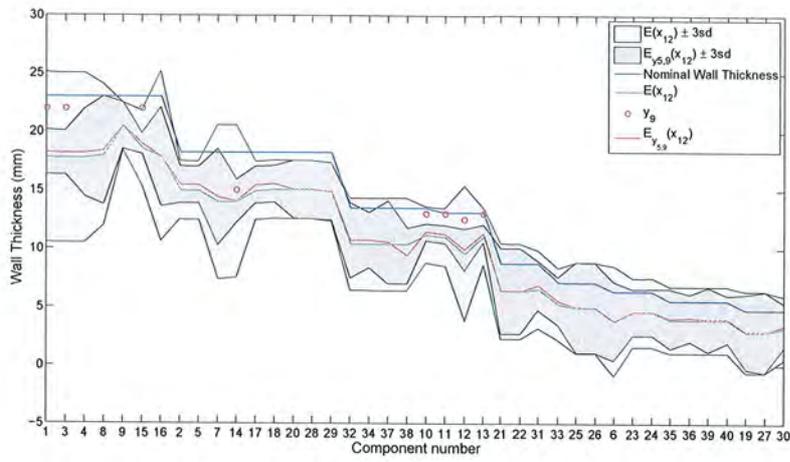


Figure 4.7: Updating beliefs for x_{12} given y_{d9}

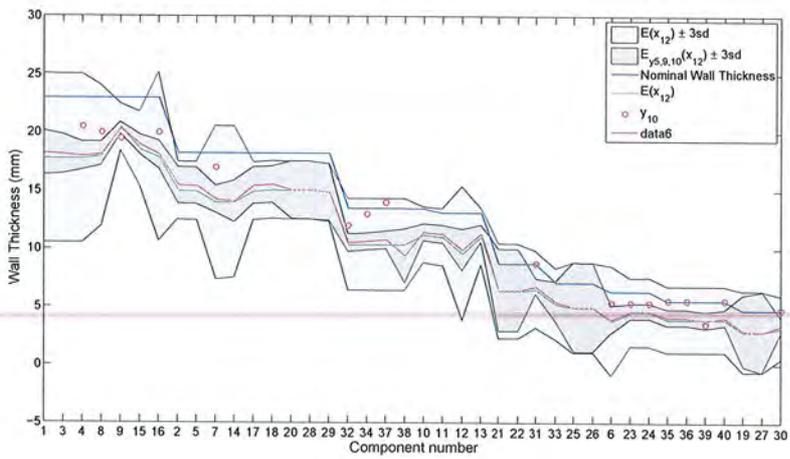
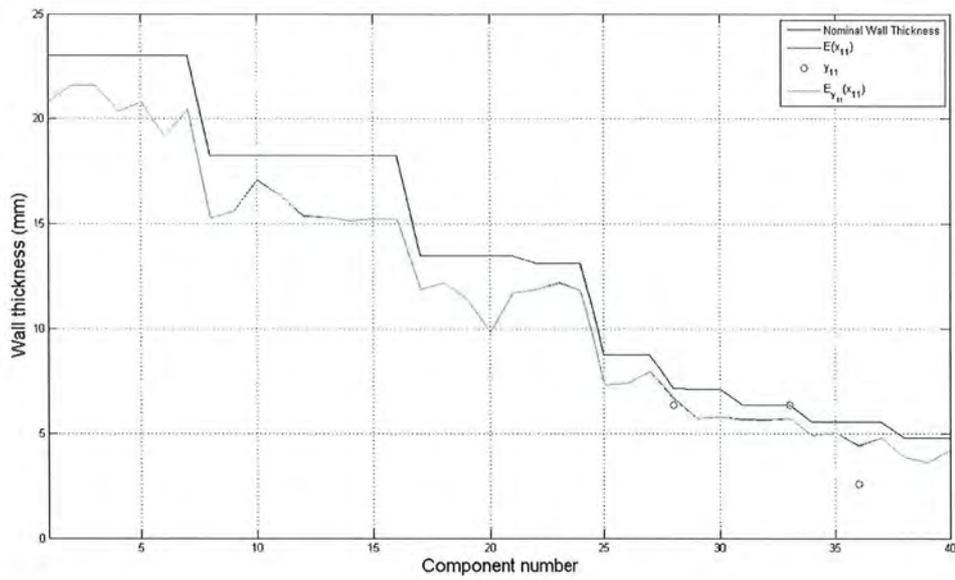
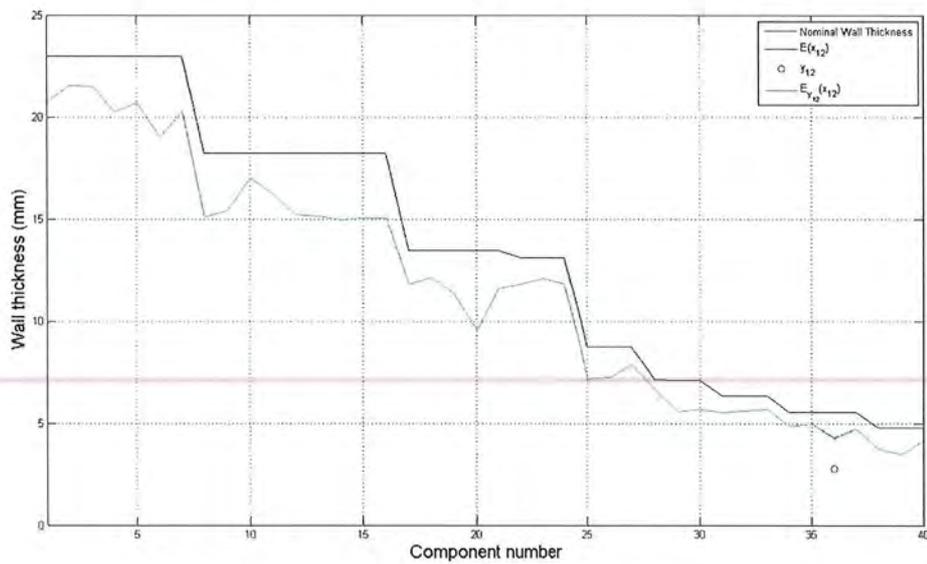


Figure 4.8: Updating beliefs for x_{12} given y_{d10}

Figure 4.9: Updating beliefs for x_{11} given y_{d11} Figure 4.10: Updating beliefs for x_{12} given y_{d12}

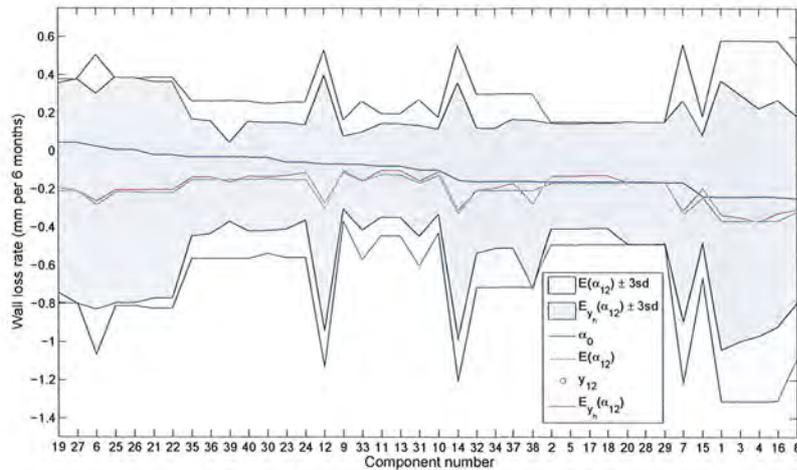


Figure 4.11: Updating beliefs for α_t given y_h

should be considered sequentially as a way of demonstrating the effect of including progressively more observation points in the set y_h .

We can see from the sequence of plots that the adjusted standard deviation band narrows at each inspection point. This is a necessary consequence of Bayes linear updating. As the system is observed our variance must be reduced. We also note that the general trend in the observations is for them to be greater than the expected value as predicted by our model. That is, the system is not degrading as rapidly as we would have expected. This can be seen as the magenta line showing the adjusted expectation lies, in general, above the green line showing our unadjusted beliefs.

It is noticeable that the majority of observations lie above the expected global value, indicating some inaccuracy in the model. The model may be over-estimating the severity of the global corrosion rate as a consequence of emphasising its role as driving wall loss, but basing parameter estimation on minima data. This discrepancy will have to be addressed before implementing the method in practice.

Figure 4.11 shows the effect of updating our beliefs about α_{12} by the available data y_h . We can see that, in line with the patterns in the historical data, the wall loss rate is generally closer to 0 than is predicted by the model. However, we observe there is a lot of variability in the wall loss rate, so our ability to predict future behaviour with confidence is not good.

Chapter 5

Evaluating the worth of an inspection design

So far we have developed methods for modelling large systems of components and updating these models for observational data. In this chapter we propose a rational quantitative criterion for assessing how ‘good’ a particular inspection design is. The problems of selecting inspection designs are tackled in chapter 6.

Our objective is to provide a sensible method for planning future inspections. Consequently, we need to develop design criteria that will allow us to discriminate between superior and inferior inspection designs. Such criteria should be defined in relation to our inspection aims. It is necessary to establish what we hope to achieve by inspecting the system in order to determine whether or not a particular inspection design has performed well in helping us to achieve the inspection goals.

We introduce the concepts of inspection design and design criteria in section 5.1, discuss the problem structure we will be tackling and propose a method for evaluating the worth of an inspection design in section 5.2. In section 5.3 we describe how we assign utility values to designs, and we illustrate how we can explicitly evaluate this criterion under certain distributional assumptions in section 5.4. Finally, we give examples of calculating the criterion for a specific inspection design based on our real world system in section 5.5

5.1 Inspection design

We consider an inspection design to be a set of instructions detailing how to perform a particular inspection. We denote by D the set of all possible inspection designs and an individual inspection design is labelled d . For this account, an individual inspection design, d , will consist of instructions for each component indicating if that component is to be inspected, which locations within the component are to be inspected, and which inspection technique is to be used. This is consistent with the definition given in section 2.3.1.

Our aim is to provide a method for comparing inspection designs to identify those which are, in some sense, 'better' than others. To achieve this we require a rational way of comparing different inspection designs. The choice of design criteria should depend on the aims of the inspection procedure. The ultimate goal of inspection is to use the observational data to ensure that the system is well maintained, and that any components approaching failure are identified. A 'good' inspection design is one that helps the inspectors to achieve these goals. For the case of corrosion, during any inspection we will want to learn about the current wall thickness (i.e. the current system level) and also update our beliefs about the expected wall thickness at future time points (future system level). Our objectives are to use the updated beliefs about the state of the system to determine what maintenance work is required.

Frequently, design criteria are based on either the variance of a related quantity or maximising the 'information' gained [12], [8]. Variance minimisation criteria seek to achieve the maximum reduction in the variance (and therefore decreased uncertainty) of the objective function over all locations/components based on the observation of a subset of locations. The subset that yields the smallest overall variance is then the optimal design set. Maximal entropy (ENT) [52] and the Shannon Information Index (SII) [3] criteria are examples of information maximisation criteria. These aim to find the set of sites that yield the most information about the uninspected sites for a stated observation function. Both criteria are maximised by the design, d , which gives the greatest overall information for all sites, conditional on the observation sites.

We would prefer to use a criterion on a more readily interpretable scale. By considering the outcome of the inspection as a decision problem, in which we must choose what maintenance to perform on the basis of our inspection results, we can measure design performance in terms of utility. We can assign utilities $U(d)$ based on the consequences of our decisions, δ , which are based on the observations \underline{y}_{dt} . In this framework, 'good' designs are those which allow us to make the 'right' decisions most often.

We have to compare expected performance. Our criterion will assess expected inspection plan performance over all possible values of the observations, \underline{y}_{dt} , which will in general be written as y . Our focus is comparing typical design performance over a large design space, D , so we would like to ensure that any design criterion can be easily calculated. If possible, we would like to avoid conducting computationally intensive simulation experiments for each choice of design, d , simply to evaluate the design criterion as this will restrict how many $d \in D$ we will be able to compare. Making our criterion straightforward to calculate will allow us to search D more thoroughly for inspection plans d with 'good' typical performance. By taking the utility of the expected outcome over all values of \underline{y}_{dt} we will obtain a measure of typical inspection performance.

We need to develop a criterion with the following properties:

1. Prefers designs that reduce 'uncertainty' (in some measure).
2. Prefers designs that minimise loss (in some measure).
3. Easily computable to allow more designs to be compared.
4. Easily interpreted output (preferably utility).

In general, designs that are cheap to implement will not be good at reducing uncertainty, and designs that are good at reducing uncertainty will not be cheap to implement, so the criterion will have to balance these conflicting requirements.

We describe how to define one such suitable criterion, the expected loss of a design where loss is taken to be equivalent to negative utility. Adopting this approach requires us to structure our inspection problem more carefully.

5.2 Problem structure

We can formulate inspection choice as a Bayesian decision problem. For each component, we can choose to perform the appropriate level of maintenance work based on our updated beliefs about the current and future state of the system. All maintenance work will have an associated loss, as will component failure. These losses will form the loss function for our action space, defined as the set of possible decisions, Δ , we can make with respect to a component, c .

We consider the situation of planning maintenance for the finite time interval $[t, t+k]$, where t is the current time and $t+k$ is the time of the next major inspection. The decision problem is treated as the simple case in which there are only two maintenance policies - either replacing component c (action a) or doing nothing (action \bar{a}) and only two potential outcomes - component failure, F , before $t+k$, and component survival to $t+k$, \bar{F} . This gives us four outcome/decision pairs. The losses associated with each pair are shown in Table 5.1. We have made the further assumption that a replaced component will not fail during the time interval under consideration, that is before $t+k$.

	F	\bar{F}
a	$L_R(c)$	$L_R(c)$
\bar{a}	$L_F(c)$	0

Table 5.1: Loss Function

For each component, we can represent this problem as a simple decision tree as shown in Figure 5.1. The chance node associated with 'replacement' has 2 branches. These represent possible behaviour of the unreplaced component. This gives us a way of taking into account the effect of unnecessary replacement. We treat a replacement action as a single, unrepeated cost. Our formulation considers that a replaced component will not fail before $t+k$. This is a simplification introduced to make the decision problem more tractable, which we acknowledge may not be realistic. Working from top to bottom, the four leaves of this tree represent:

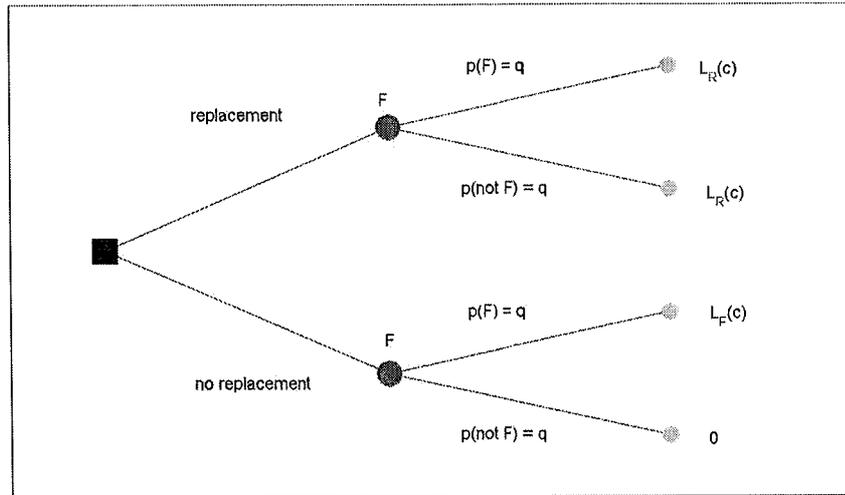


Figure 5.1: Decision tree representation of component replacement problem

1. Component replaced, if unreplaced would have failed. Loss $L_R(c)$.
2. Component replaced, if unreplaced would not have failed. Loss $L_R(c)$.
3. Component not replaced, and fails. Loss $L_F(c)$.
4. Component not replaced, and does not fail. Loss 0.

Decisions 1 and 4 can be said to represent 'correct' decisions that minimise maintenance expenditure, and decisions 2 and 3 illustrate 'wrong' decisions in which unnecessary component replacement was performed and failing components were not replaced. In the decision problem framework, inspection plans with 'good' typical performance are those which allow us to make the 'right' decision more often. We aim to produce criteria that allow us to select inspection designs which will, in general, help inspectors make more correct decisions than incorrect decisions, by giving high probability of observing inspection data which leads to improved decision making for many components as a result of increased confidence over the probability of component failure (q) for many components.

5.2.1 Sources of loss

The loss associated with any given inspection design is an accumulation of the costs incurred by component replacement, component failure and the cost of making the inspection. For the component replacement problem, replacing a component that would not have failed or not replacing a component that then fails are perceived to be 'wrong' choices of action. The costs associated with these 'wrong' actions are greater than those associated with the complementary action for a particular outcome. The cost of each action-outcome pair (Table 5.1), is henceforth written as L_R, L_F and q in place of $L_R(c), L_F(c)$ and $q(c)$ to represent the replacement cost, failure cost and failure probability of a generic component c . We also assume that replacement is a necessary consequence of failure, so failure incurs a loss of at least L_R (i.e. $L_F \geq L_R$). Each c has its own replacement cost, failure cost and failure probability and actions are chosen componentwise.

Describing inspection losses (i.e. the costs of making the inspection) is open to many different models. We choose to model L_I (the inspection losses) as an initial cost and an incremental cost.

$$L_I = L_{SU} + n_d L_{IC}. \quad (5.1)$$

The initial cost models the price of 'setting up' the inspection. The losses incurred by transportation of the inspectors, preparing the site for inspection, and any other activities caused by the decision to inspect, would be included in this part of the model. The incremental cost measures the extra loss incurred as the scale of the inspection increases, and could, for example, be used to account for factors such as the cost of materials used during inspection or losses due to reduced operating capacity. We base the incremental cost on the number of components inspected as part of design d (n_d). This provides a measure of time and effort expended on making the inspection.

The loss associated with making the inspection is independent of the other losses. This allows us to incorporate it into the expected loss criterion simply by adding it to the other losses. Accounting for inspection cost in the criterion means it is possible

to compare designs that cost different amounts to implement. For example, consider a design which allows us to predict failure perfectly, but is extremely expensive to carry out. Although we would prefer to inspect perfectly, the extra cost of doing so may force us to choose a less effective design. By adding the cost of making the inspection into the criterion we have a method of assessing whether or not it is worth paying the extra for the gain in predictive power a more expensive design offers.

5.3 Design utility

The strength of our preferences for particular inspection designs is measured by the utility of the design, $U(d)$. We choose to work with loss that is defined as negative utility because this offers a more natural interpretation for assigning. Consequently, we will seek to minimise loss, which will in turn maximise utility. The value of L is a function of the outcome, O , and our decision, $\delta(y)$. Our decision $\delta(y)$, is determined by the values of y we observe, and, therefore, by the posterior probability of failure, $P(F|y)$. For an inspection planning problem, the values of y are random, and therefore our decision is random. Our aim is to choose the decision procedure $\delta(y)$, from the set of all possible decision procedures, Δ , such that the loss associated with our decision procedure is minimised. So, the utility of a design, d , is given by the loss associated with our decisions plus loss incurred by making the inspection:

$$\begin{aligned} U(d) &= -\{E_{\min\{\delta \in \Delta\}} [L(O, \delta(y))] + L_I\} \\ &= -\{E[L(O, \delta^*(y))] + L_I\} \end{aligned} \quad (5.2)$$

where δ^* is the Bayes decision function, the decision function that minimises (5.2).

For any decision $\delta(y)$, the expected loss associated with a particular outcome, $o \in O$, is given by the product of the value of the loss function associated with decision $\delta(y)$ and outcome o , $(L[o, \delta(y)])$, and the probability of outcome o , given observations, y . The expected loss, conditional on y , for a design d is the sum of the outcome expected losses for all possible outcomes. For the simple case under

consideration, there are only two outcomes and the expected loss for a design can be calculated as:

$$\begin{aligned} E[L(O, \delta(y))] &= E\{E[L(O, \delta(y))|y]\} \\ &= E\{L(F, \delta(y))P(F|y) + L(\bar{F}, \delta(y))P(\bar{F}|y)\} \end{aligned} \quad (5.3)$$

We may use (5.3) to identify the Bayes decision function.

5.3.1 The Bayes decision function

The decision to replace or not to replace is taken componentwise on the basis of our observations y and our posterior probability of failure, updated for y . Losses associated with our decisions depend on component failure and are therefore random. We will always choose the action that minimises the expected loss. Having observed y , we choose a (to replace) or \bar{a} (not to replace) a component depending on $\min\{E[L(O, a)|y], E[L(O, \bar{a})|y]\}$. These can be calculated using Table 5.1 as:

$$\begin{aligned} E[L(O, a)|y] &= L(F, a)P(F|y) + L(\bar{F}, a)P(\bar{F}|y) \\ &= L_R P(F|y) + L_R(1 - P(F|y)) \\ &= L_R \end{aligned} \quad (5.4)$$

$$\begin{aligned} E[L(O, \bar{a})|y] &= L(F, \bar{a})P(F|y) + L(\bar{F}, \bar{a})P(\bar{F}|y) \\ &= L_F P(F|y) + 0 \cdot P(\bar{F}|y) \\ &= L_F P(F|y) \end{aligned} \quad (5.5)$$

which tells us to choose a (i.e. to replace) if the posterior probability of failure is greater than the proportion of the failure cost contributed by the replacement cost.

That is, to replace if:

$$P(F|y) > \frac{L_R}{L_F} = \rho$$

and not to replace otherwise. Therefore, for any realisation of the observations, the Bayes decision function is:

$$\delta^*(y) = a \text{ if } p(F|y) > \rho \quad (5.6)$$

$$\delta^*(y) = \bar{a} \text{ if } p(F|y) \leq \rho \quad (5.7)$$

5.3.2 Calculating $E[L(d)]$

The utility of an inspecting a component can be calculated as the sum of the expected value of the loss function associated with our Bayes decision function, $\delta^*(y)$, and the inspection cost, L_I . The expected loss of our decisions for a design, d , is denoted $E[L(d)]$ and is evaluated, using (5.3), as:

$$\begin{aligned}
 E[L(d)] &= E[L(F, \delta^*(y))P(F|y) + L(\bar{F}, \delta^*(y))P(\bar{F}|y)] \\
 &= E[L(F, a)P(F|y) + L(\bar{F}, a)P(\bar{F}|y)|\delta^*(y) = a]P(\delta^*(y) = a) \\
 &\quad + E[L(F, \bar{a})P(F|y) + L(\bar{F}, \bar{a})P(\bar{F}|y)|\delta^*(y) = \bar{a}]P(\delta^*(y) = \bar{a}) \\
 &= E[L(F, a)P(F|y) + L(\bar{F}, a)P(\bar{F}|y)|P(F|y) > \rho]P(P(F|y) > \rho) \\
 &\quad + E[L(F, \bar{a})P(F|y) + L(\bar{F}, \bar{a})P(\bar{F}|y)|P(F|y) \leq \rho]P(P(F|y) \leq \rho)
 \end{aligned}$$

where each term corresponds to the different decisions. The first term describes the expected loss due to replacement, and the second the expected loss due to inaction. Recalling the expected values of the loss function for each decision (5.4), (5.5), this can be written as:

$$E[L(d)] = L_R P(P(F|y) > \rho) + L_F E[P(F|y)|P(F|y) \leq \rho] P(P(F|y) \leq \rho) \quad (5.8)$$

We see, from (5.8), that the expected loss of our decisions for a design, d , is determined by the probability that, using design d , we will observe data y for which $P(F|y) > \rho$ and by the expected value of $P(F|y)$ over all y for which $P(F|y) \leq \rho$. For example, consider the best conceivable design. This would be the design in which our updated beliefs allow for perfect prediction, i.e. $P(F|y) = 0$ or 1 , and so we could identify failing components with certainty. In this case, $E[L(d)]$ would be given by $n_R L_R$, where n_R is the number of components in need of replacement. The expected value $E[P(F|y)|P(F|y) \leq \rho] = 0$ under perfect prediction, as we will never incur a failure loss. Components will be replaced (with loss L_R) if $P(F|y) = 1$, and we therefore know they will fail before $t+k$, or they will not be replaced if $P(F|y) = 0$ and we know, with certainty, they will not fail before $t+k$, and therefore cause no loss in utility. Such a design minimises unnecessary expenditure.

In practice, even using complete inspection will not provide us with perfect information, so our $E[L(d)]$ values will instead include some unavoidable random loss. These are losses introduced by replacing components that would not have failed or by failing to replace components that will fail.

Calculating $E[L(d)]$ requires us to be able to evaluate the posterior probability of component failure, which for notational simplicity we set to be $P(F|y) = q(y)$. The posterior probability of component failure $q(y)$ is also a random quantity, as it depends on the unknown quantity y , and has its own probability distribution, which we will also need to describe. The evaluation of $E[L(d)]$ requires the evaluation of:

1. $P(P(F|y) > \rho) = p(q(y) > \rho)$, which is equivalent to the integral:

$$p(q(y) > \rho) = \int_{\rho}^1 p(q(y))dq(y) = I_1. \quad (5.9)$$

2. $E[P(F|y)|P(F|y) \leq \rho]P(P(F|y) \leq \rho) = E[q(y)|q(y) \leq \rho]p(q(y) \leq \rho)$, which can be expressed as:

$$\begin{aligned} E[q(y)|q(y) \leq \rho]p(q(y) \leq \rho) &= \int_0^{\rho} q(y) \frac{p(q(y))}{p(q(y) \leq \rho)} dq(y) \cdot p(q(y) \leq \rho) \\ &= \int_0^{\rho} q(y)p(q(y))dq(y) = I_2. \end{aligned} \quad (5.10)$$

3. The expected loss of our decisions for a design can thus be written as:

$$E[L(d)] = L_R I_1 + L_F I_2 \quad (5.11)$$

and the utility of a design is:

$$U(d) = -\{L_R I_1 + L_F I_2 + L_I\}$$

To evaluate these integrals we require an expression for $q(y) = p(F|y)$.

5.4 Approximating posterior failure probabilities

From (5.8), assessment of $E[L(d)]$ for an inspection design, d , depends on assessment of the posterior probability of component failure, conditioned on our inspection observations, y . This is given by:

$$P(F|y) = q(y) = P(m_{t+k} < W_C|y) \quad (5.12)$$

Expressing this probability explicitly requires assumptions to be made about the distributions of the random quantity m_{t+k} , which is the minimum wall thickness. The method could equally be applied to any well defined observation function; for this model, the observation function is the component minimum wall thickness. We choose to approximate the behaviour of m_{t+k} given y with a Normal distribution:

$$m_{t+k} \sim N[\mu_{t+k}(y), \sigma_{t+k}^2(y)] \quad (5.13)$$

In general this may not be an appropriate assumption, and the validity of any distributional assumptions should be checked. In our situation we are modelling minimum wall thickness as $m_{t+k} = x_{t+k} + \tilde{r}_{t+k}$, where from data analysis, expert elicitation and system simulation, the global trend term, x_{t+k} , is Normally distributed and \tilde{r}_{t+k} is expected to be distributed with, approximately, an extreme value distribution. The exact form will depend on the choices made with respect to the number of locations within a component, but the expected distribution will lie somewhere between a Normal and a Gumbel distribution. Figure 5.2 illustrates the distributions we can expect to be associated with the minimum wall thickness m_t for a single component of the Circuit A system for simulations of length $t = \{1, 5, 10, 15\}$ steps. Plots for other components in Circuit A follow a similar pattern (see Appendix). These indicate that the assumption of Normality is not grossly unreasonable for these quantities. The determination of appropriate values for $\mu_{t+k}(y)$ and $\sigma_{t+k}^2(y)$ will be described in section 5.4.1.

The assumption of Normality allows us to evaluate $q(y)$ from the standard Nor-

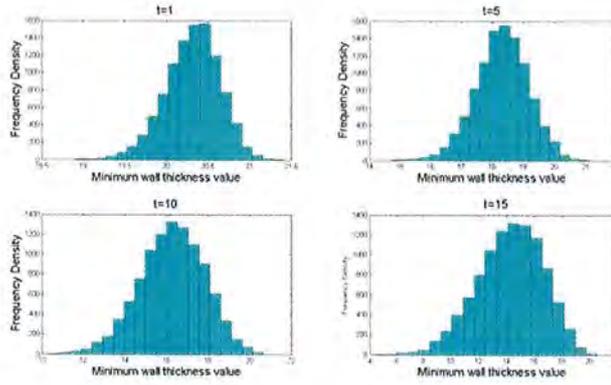


Figure 5.2: Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 8 from Circuit A

mal cdf. Standardising by $\mu_{t+k}(y)$ and $\sigma_{t+k}^2(y)$, we have:

$$\begin{aligned}
 q(y) = P(m_{t+k} < W_C | y) &= P\left(\frac{m_{t+k} - \mu_{t+k}(y)}{\sigma_{t+k}(y)} < \frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}(y)} \middle| y\right) \\
 \text{so that } q(y) &= \Phi\left(\frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}(y)}\right) \quad (5.14)
 \end{aligned}$$

where both $\mu_{t+k}(y)$ and $\sigma_{t+k}(y)$ are random quantities depending on the inspection design. While (5.14) will often give a reasonable order of magnitude assessment for the posterior failure probability in practice, observe that we could replace the assumption of posterior Normality by any alternative pdf which was fully characterised by the posterior mean and variance.

5.4.1 Assessment of posterior moments

We described in chapter 4 the Bayes linear method for updating beliefs. We again choose to employ the Bayes linear approach to update our beliefs about the minimum wall thickness, m_{t+k} , given the observational data, y . Characterising the behaviour of both m_{t+k} and y with a specified distribution is not straightforward, and also not necessarily appropriate. Simulation allows us to learn about the expectation, variance and covariance structures of these quantities directly from the underlying model. To update beliefs about m_{t+k} given a particular realisation of the data, y^* ,

we will use the following equations:

$$\mu_{t+k}(y^*) = E_y(m_{t+k}) = E(m_{t+k}) + \text{cov}(m_{t+k}, y) \text{var}^{-1}(y) [y^* - E(y)] \quad (5.15)$$

$$\sigma_{t+k}^2(y^*) = \text{var}_y(m_{t+k}) = \text{var}(m_{t+k}) - \text{cov}(m_{t+k}, y) \text{var}^{-1}(y) \text{cov}(y, m_{t+k}) \quad (5.16)$$

Assessing the value of $\sigma_{t+k}^2(y)$ for a particular inspection design using Bayes linear adjustment allows us to remove the dependence of our estimate on the particular values of the observations, y . This can be seen from (5.16), in which the specific values of the observations, denoted y^* , are not involved in the calculation, but expectation and variance properties of the general quantity, y , are used. Therefore, under this approximation, the value of $\sigma_{t+k}^2(y)$ is fixed for a given design, and we have $\sigma_{t+k}^2(y) = \sigma_{t+k}^2 \forall y$.

The Bayes linear update of $\mu_{t+k}(y)$ does depend on the particular values of y . Therefore we have to regard $\mu_{t+k}(y)$ as a random quantity. $\mu_{t+k}(y)$ has an unknown distributional form (about which we can learn through simulation). Recalling $\mu_{t+k}(y) = E_y(m_{t+k})$, assessments of the expectation of $\mu_{t+k}(y)$ can be found from:

$$\begin{aligned} E[\mu_{t+k}(y)] &= E[E_y(m_{t+k})] \\ &= E(E(m_{t+k} + \text{cov}(m_{t+k}, y) \text{var}^{-1}(y) [y - E(y)])) \\ &= E(m_{t+k}) + \text{cov}(m_{t+k}, y) \text{var}^{-1}(y) [E(y) - E(y)] \\ &= E(m_{t+k}) \end{aligned} \quad (5.17)$$

Similarly, an assessment of the variance of $\mu_{t+k}(y)$ is given by:

$$\text{var}(\mu_{t+k}(y)) = \text{var}(E_y(m_{t+k}))$$

We then use the relation:

$$m_{t+k} = E_y(m_{t+k}) + [m_{t+k} - E_y(m_{t+k})]$$

to rewrite m_{t+k} as the sum of two uncorrelated terms. Taking the variance of this expression gives:

$$\begin{aligned} \text{var}(m_{t+k}) &= \text{var}(E_y[m_{t+k}]) + \text{var}(m_{t+k} - E_y[m_{t+k}]) \\ &= \text{var}(\mu_{t+k}(y)) + \text{var}_y(m_{t+k}) \end{aligned}$$

yielding $\text{var}(E_y[m_{t+k}]) = \text{var}(m_{t+k}) - \text{var}_y(m_{t+k}) \quad (5.18)$

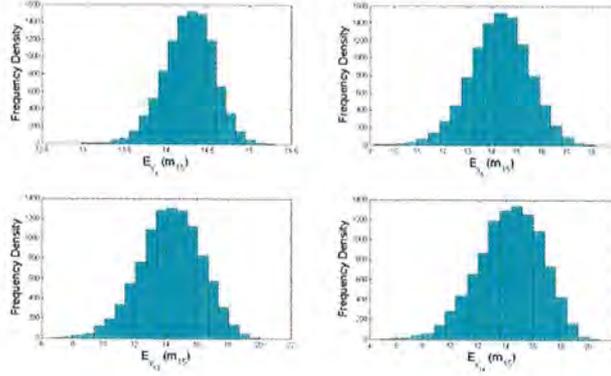


Figure 5.3: Histograms of adjusted expectation values for component 8 from Circuit A based on simulated realisations of observations at 4 different time points

where $\text{var}_y(m_{t+k}) = \sigma_{t+k}^2$, as given by (5.16).

To be able to perform calculations using this quantity, we have to make further assumptions about the distribution of $\mu_{t+k}(y)$. For this account, we approximate $\mu_{t+k}(y)$ by a Normal distribution, with moments as given by (5.17, 5.18). Figure 5.3 shows the behaviour of the adjusted expectation for the minimum at time 15, $E_{y_t}(m_{15})$, in a component of Circuit A based on total observation of the system at times $t = \{1, 5, 10, 14\}$. The plots are typical of those produced for this quantity in most components, and indicate that the choice of a Normal approximation is again not grossly unreasonable. It would be possible to adopt any approximating distribution for this quantity, which could be summarised by its expectation and variance, but this account only considers Normal approximations. This choice provides us with sufficient structure for evaluating I_1 and I_2 .

In summary, we assume that m_{t+k} is Normally distributed with mean $\mu_{t+k}(y)$ (itself a random quantity, dependent on y) and variance σ_{t+k}^2 . We also assume that our beliefs about the random quantity $\mu_{t+k}(y)$ follow a Normal distribution with mean $E(m_{t+k})$ and variance $\text{var}(m_{t+k}) - \text{var}_y(m_{t+k})$, i.e.

$$m_{t+k} \sim N[\mu_{t+k}(y), \sigma_{t+k}^2] \quad (5.19)$$

$$\mu_{t+k}(y) = E_{y_t}(m_{t+k}) \sim N[E(m_{t+k}), \text{var}(m_{t+k}) - \text{var}_y(m_{t+k})] \quad (5.20)$$

from equations (5.17, 5.18).

The Bayes linear updating equations allow us to resolve our problem to one of uncertainty over the behaviour of m_{t+k} , which is controlled by the random quantity $\mu_{t+k}(y)$. We have obtained expressions for the expectation and variance of both quantities that depend on the updated expectation and variance values. It is still necessary to assign appropriate distributions to these moments to allow the determination of the values of $q(y)$ and $p(q(y))$, and even given such assumptions the integrals I_1 and I_2 may not be easy to evaluate. However, this provides a clear method for evaluating a basic property of an inspection design that incorporates the updated variance and presents the results on an interpretable scale - the expected loss per component of the decisions made as result of using design d .

To convert the individual component expected loss values into a design expected loss, we sum the expected loss over all components. This simplification follows as we are making decisions about component replacement independently of the other components. It would be possible to construct a more realistic model of the decision making process, which allowed for joint decision making and the adjusting of utilities as the scale of necessary repair work increases (thereby allowing us to account for potential non-linearity of utilities), but the proposed method offers a plausible order of magnitude assessment of design quality.

5.4.2 Evaluating I_1

Following from the assumptions of Normality made previously, the standardised quantity,

$$z(y) = \frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}}, \quad (5.21)$$

on which we base our probability of failure estimate,

$$q(y) = \Phi(z) \quad (5.22)$$

can be assumed to follow a Normal distribution with expected value:

$$\begin{aligned}
 E(z) &= E\left(\frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}}\right) \\
 &= \frac{W_C - E[\mu_{t+k}(y)]}{\sigma_{t+k}} \\
 &= \frac{W_C - E(m_{t+k})}{\sigma_{t+k}} = \mu_z
 \end{aligned} \tag{5.23}$$

which we denote as μ_z , and variance given by:

$$\begin{aligned}
 \text{var}(z) &= \text{var}\left(\frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}}\right) \\
 &= \frac{\text{var}(\mu_{t+k}(y))}{\sigma_{t+k}^2} \\
 &= \frac{\text{var}(m_{t+k}) - \text{var}_y(m_{t+k})}{\sigma_{t+k}^2} = \sigma_z^2
 \end{aligned}$$

labelled σ_z^2 . We also define:

$$\sigma_\mu^2 = \text{var}(m_{t+k}) - \text{var}_y(m_{t+k}) \tag{5.24}$$

to represent the variance associated with $\mu_{t+k}(y)$. Consequently,

$$\sigma_z^2 = \left(\frac{\sigma_\mu}{\sigma_{t+k}}\right)^2 \tag{5.25}$$

We will use this notation to condense our expressions for I_1 and I_2 .

Recall from (5.11) that $E[L(d)] = L_F(\rho I_1 + I_2)$ where:

$$I_1 = \int_\rho^1 p(q(y))dq(y)$$

This integral is equivalent to calculating the probability that the posterior failure probability will be greater than ρ , the ratio of replacement and failure costs. We can write I_1 as:

$$I_1 = \int_\rho^1 p(q(y))dq(y) = p(q(y) > \rho) = p[p(m_{t+k} \leq W_C|y) > \rho]$$

Using our assumptions of Normality for m_{t+k} (5.19) and $\mu_{t+k}(y)$ (5.20), we can then

say:

$$\begin{aligned}
 p [p(m_{t+k} \leq W_C | y) > \rho] &= p \left[\Phi \left(\frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}} \right) > \rho \right] \\
 &= p [\mu_{t+k}(y) < W_C - \Phi^{-1}(\rho) \sigma_{t+k}] \\
 &= \Phi \left(\frac{W_C - \Phi^{-1}(\rho) \sigma_{t+k} - E[\mu_{t+k}(y)]}{\sqrt{\text{var}(m_{t+k}) - \text{var}_y(m_{t+k})}} \right) \\
 &= \Phi \left(\frac{\sigma_{t+k} \left\{ \frac{W_C - E(m_{t+k})}{\sigma_{t+k}} - \Phi^{-1}(\rho) \right\}}{\sigma_\mu} \right)
 \end{aligned}$$

which can be written more succinctly in terms of μ_z (5.23) and σ_z (5.25) as:

$$p(q(y) > \rho) = \Phi \left[- \left(\frac{\Phi^{-1}(\rho) - \mu_z}{\sigma_z} \right) \right]. \quad (5.26)$$

We can then evaluate I_1 from a standard (cumulative) Normal distribution function.

5.4.3 Evaluating I_2

This is a more involved calculation. The required integral is given by:

$$I_2 = \int_0^\rho q(y) p(q(y)) dq(y)$$

We have assumed that $q(y)$ can be written as $q(y) = \Phi \left(\frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}} \right) = \Phi(z)$ (see section 5.4 for details). Changing variables from $q(y)$ to z is a change of variable $q(y) \mapsto f_1(\mu_{t+k}) \mapsto f_2(z)$, but we will treat it as setting $q(y) = \Phi(z) \Rightarrow dq = \phi(z) dz$, where $\phi(z)$ is used to denote the derivative of $\Phi(z)$ (in this case $\phi(z)$ would be a Normal pdf). Denoting the pdf of the posterior probability of failure, $p(q)$ by f_q , we can express I_2 as an integral in z :

$$I_2 = \int_{-\infty}^{\Phi^{-1}(\rho)} \Phi(z) f_q(\Phi(z)) \phi(z) dz$$

$f_q(\Phi(z))$ can be found as the derivative of F_q . $F_q = p(q(y) < x)$, arguing as for I_1 (section 5.4.2) gives:

$$\begin{aligned}
 p [p(m_{t+k} \leq W_C | y) \leq x] &= p \left[\Phi \left(\frac{W_C - \mu_{t+k}(y)}{\sigma_{t+k}} \right) \leq x \right] \\
 &= p [\mu_{t+k}(y) > W_C - \Phi^{-1}(x) \sigma_{t+k}] \\
 &= 1 - \Phi \left(\frac{W_C - \Phi^{-1}(x) \sigma_{t+k} - E(m_{t+k})}{\sigma_\mu} \right)
 \end{aligned}$$

Setting:

$$f_q = \frac{dF_q}{dx} = \frac{d}{dx} \left[1 - \Phi \left(\frac{W_C - E(m_{t+k}) - \sigma_{t+k} \Phi^{-1}(x)}{\sigma_\mu} \right) \right]$$

and using the chain rule with $J = \Phi^{-1}(x)$:

$$\begin{aligned} \frac{dF_q}{dx} &= \frac{dF_q}{dJ} \cdot \frac{dJ}{dx} \\ \frac{dF_q}{dJ} &= \left(\frac{-\sigma_{t+k}}{\sigma_\mu} \right) \left(-\phi \left[\frac{W_C - E(m_{t+k}) - \sigma_{t+k} J}{\sigma_\mu} \right] \right) \\ J &= \Phi^{-1}(x) \Leftrightarrow x = \Phi(J) \\ \frac{dx}{dJ} &= \phi(J) \end{aligned}$$

we find:

$$f_q = \left(\frac{-\sigma_{t+k}}{\sigma_\mu} \right) \left(-\phi \left[\frac{W_C - E(m_{t+k}) - \sigma_{t+k} J}{\sigma_\mu} \right] \right) (\phi(J))^{-1}.$$

Substituting for J , this can be written as:

$$f_q = \phi \left(\frac{W_C - E(m_{t+k}) - \sigma_{t+k} \Phi^{-1}(x)}{\sigma_\mu} \right) \frac{1}{\phi(\Phi^{-1}(x))} \cdot \frac{\sigma_{t+k}}{\sigma_\mu}.$$

Which, when x is replaced by the variable of interest $\Phi(z)$, gives the expression required to evaluate I_2 .

$$\begin{aligned} I_2 &= \int_{-\infty}^{\Phi^{-1}(\rho)} \Phi(z) \phi \left(\frac{W_C - E(m_{t+k}) - \sigma_{t+k} \Phi^{-1}(\Phi(z))}{\sigma_\mu} \right) \frac{\sigma_{t+k} \phi(z) dz}{\phi(\Phi^{-1}(\Phi(z))) \sigma_\mu} \\ &= \int_{-\infty}^{\Phi^{-1}(\rho)} \Phi(z) \phi \left(\frac{W_C - E(m_{t+k}) - \sigma_{t+k} z}{\sigma_\mu} \right) \frac{\phi(z) \sigma_{t+k} dz}{\phi(z) \sigma_\mu} \\ &= \int_{-\infty}^{\Phi^{-1}(\rho)} \Phi(z) \phi \left(\frac{\sigma_{t+k} \left(\frac{W_C - E(m_{t+k})}{\sigma_{t+k}} - z \right)}{\sigma_\mu} \right) \frac{\sigma_{t+k}}{\sigma_\mu} dz \end{aligned}$$

Using (5.23, 5.25) this can be rewritten as:

$$I_2 = \int_{-\infty}^{\Phi^{-1}(\rho)} \Phi(z) \phi \left(\frac{\mu_z - z}{\sigma_z} \right) \frac{dz}{\sigma_z} \quad (5.27)$$

which will be evaluated numerically.

For I_2 , we are left with an integral that depends on the three quantities μ_z , σ_z and ρ , and must be evaluated using numerical methods.

5.5 Evaluating the expected loss criterion

In this section we illustrate how to apply the expected loss criterion to a real system. The system will again be the Circuit A corrosion Circuit from the Site A platform. Our intention is to consider the inspection planning problem for this system, which we demonstrate in chapter 6. Here we show how to calculate the expected loss for specific inspection designs.

The expected loss criterion can be used as a way of planning the next inspection, based upon the need to run the system for a given time frame, so we are required to specify which time frame we will be considering. The UK Health and Safety Executive Regulations require that a *safety case* for the running of industrial systems is prepared at least every 5 years [27]. The safety case provides details on the safety implications of running the system, and provision made for handling failures. Therefore planning an inspection within around five years of the final update should be considered suitable for the Circuit A problem. The last available real data point is at $t = 12$ in our initial model, which covers the period 01/01/1998 - 30/06/2004 (see Table 4.1). We will be interested in planning an inspection before 31/12/2009, which translates to 11 further times steps in our model - specified in Table 5.2. Therefore the value of k , the number of steps into the future which we will consider will be 11 and the final time point for inspection planning procedure is $t + k = 12 + 11 = 23$, equivalent to 31/12/2009.

To carry out an expected loss calculation it is also necessary to specify other key aspects of the problem:

1. Initial beliefs - our beliefs about the future state of the system will be based on the current system state. Therefore it will be necessary to specify our current beliefs.
2. Losses - to formulate the decision problem it is necessary to specify values for the losses associated with component failure, component replacement and performing the inspection.

Time Step	Start Date	End Date
13	01/07/2004	31/12/2004
14	01/01/2005	30/06/2005
15	01/07/2005	31/12/2005
16	01/01/2006	30/06/2006
17	01/07/2006	31/12/2006
18	01/01/2007	30/06/2007
19	01/07/2007	31/12/2007
20	01/01/2008	30/06/2008
21	01/07/2008	31/12/2008
22	01/01/2009	30/06/2009
23	01/07/2009	31/12/2009

Table 5.2: Model Time Steps

3. Design - we must stipulate the design for which we want to evaluate the expected loss value.

5.5.1 Simulating future system behaviour

When conducting the initial simulation we should consider how far into the future we will want to plan, and simulate realisations up to this time. Adopting this approach means no further simulation will be required. The output from the initial simulation will be sufficient to estimate the variances and covariances required.

The information required to perform the updates consists of the values of the system property of interest - which to us will be the minimum wall thickness - from the final time point, m_{t+k} , and the observation function values from every time point, $y_{t+j} \forall j \in \{1, \dots, k\}$. It is preferable to generate realisations for all observation functions that may be used as part of the initial simulation, as this avoids the need for further simulation if the observation function is changed. We will consider the options for our observation functions to be minimisation over different amounts of the component surface. It is therefore necessary to establish a plausible set of

observation functions before simulating. In general these could be any function of the underlying true wall thickness values (u_{lct}) for which it is possible to obtain a value via some observation technique currently in use.

In the case of Circuit A, we are concentrating on minimum observed values, but would like to consider the possibility of partial component observation. Partial component observation offers the possibility of ‘missing’ the true minimum during inspection, and is also a slightly more realistic representation of the inspection process, where inaccessibility and other practical constraints are likely to apply to parts of components as well as whole components. For this illustration we consider the two observation functions of complete component inspection, in which we take the minimum over every location observed with error, and half component inspection, where the minimum is taken over half the locations observed with error. Due to the lack of correlation between locations within components in our simulation, it is acceptable to select those locations that form the ‘observed half’ of the component at random.

Using a simulation approach, we must ensure that our simulation output includes all information that will be required to calculate the values of the observation functions we wish to consider, to allow computation to proceed.

To forecast behaviour at the time point of interest, we simulate S system realisations at time $t + k$, using the simulation approach described in section 2.8. For Circuit A, our updated beliefs about the global term level \underline{x}_{12} and slope $\underline{\alpha}_{12}$ values are used as the basis for our simulation. We then generate S realisations of the system to time 23 (31/12/2009), and observations of the system for each time step, for both observation functions. To account for all design possibilities, it will be necessary to store the S \underline{m}_{t+k} realisations, and the S realisations of the two observation functions for each time point. We denote by $\underline{y}_{t+j}^{(1)}$ the complete component observations and by $\underline{y}_{t+j}^{(0.5)}$ the half component observations. Simulation output is stored in the $S \times 40$ array \mathbf{M}_{t+k} , in which the c -th column contains S simulated values for the minimum of component c at time $t + k$, and the $S \times 40 \times 11$ arrays $\mathbf{Y}^{(1)}$ and $\mathbf{Y}^{(0.5)}$, which contains and $S \times 40$ array for each time step.

To evaluate the expected utility criterion we need to estimate the variance of

the minimum wall thickness at $t + k$, adjusted for the planned observations. We will concentrate on planning an inspection at the next available time step, so the adjustment we need to consider is, in general:

$$\text{var}_{\underline{y}_{dt+1}}(\underline{m}_{t+k}) = \text{var}(\underline{m}_{t+k}) - \text{cov}(\underline{m}_{t+k}, \underline{y}_{dt+1}) \text{var}^{-1}(\underline{y}_{dt+1}) \text{cov}(\underline{y}_{dt+1}, \underline{m}_{t+k})$$

and for the Circuit A case it is:

$$\text{var}_{\underline{y}_{d12+1}}(\underline{m}_{23}) = \text{var}(\underline{m}_{23}) - \text{cov}(\underline{m}_{23}, \underline{y}_{d12+1}) \text{var}^{-1}(\underline{y}_{d12+1}) \text{cov}(\underline{y}_{d12+1}, \underline{m}_{23}).$$

All of these quantities, and the unadjusted expectation $E(\underline{m}_{t+k})$ - which will also be required, can be estimated from the simulation output as described in section 4.3.2.

For a large system, with many simulation realisations, it would be impractical to store all the simulation output. However, as our interest is in the mean and variance structure of the variables, we can recover this information by recording the cumulative means and the sums of squares and cross products for each of the simulation outputs.

Using the simulation output we will be able to evaluate these quantities for any inspection design. This, in combination with the loss information, will allow us to evaluate the expected loss criterion.

5.5.2 Specifying loss values

Specifying the losses associated with each outcome is a fundamental part of solving the decision problem. Table 5.3 shows the values assigned to component replacement, which are based around the cost of replacing a *Straight* component:

The losses reflect the cost of a new component and the effort required to install the new component. Higher costs reflect those components that are either more expensive to construct, difficult to replace or both. For example, a Cap would be a cheap component that is easy to replace, hence its low value. Obtaining estimates for replacement losses will, in general, be reasonably straightforward. The cost of a new component, and the labour required to install such a component, will usually be approximately known for most practical problems.

Component Type	Replacement Cost
Bend (short radius)	1.5L
Branch	2L
Cap	0.1L
Straight	L
Weld	0.6L
Wellhead	2L

Table 5.3: Losses incurred through component replacement

Failure costs are constructed using the consequence data from the Site A data set. The consequence data takes the form of subjective assessments of failure consequences for different groups. Each component has the following ratings:

1. *Location* (L_{F1}) - where on the platform the system is located, and how easily failure could be contained.
2. *Fluid Characteristics* (L_{F2}) - indicator of threat posed by system contents.
3. *Fluid Pressure* (L_{F3}) - highly pressurised fluids are potentially more dangerous.
4. *Criticality* (L_{F4}) - effect of component failure on larger system.
5. *Environmental* (L_{F5}) - damage caused to surroundings by failure.
6. *Commercial* (L_{F6}) - lost revenue caused by failure.

Qualitative ratings are assigned to each consequence ranging from 'Low' to 'Extreme'. These are assigned values as shown in Table 5.4. The 'Extreme' values have been adjusted to allow for a more interesting design example. In practice it is reasonable to assume that the 'Extreme' consequence cost of failure would be much more than 8 times larger than the 'Low' consequence cost. For this example, we adopt a simple linear model for combining the ratings to obtain a single value for component failure cost:

$$L_F = \sum_{i=1}^6 l_i L_{Fi} \quad (5.28)$$

where the l_i are weights allowing us to rescale factors according to importance.

Consequence	Failure Cost
Low	L
Medium	2L
High	4L
Extreme	8L

Table 5.4: Losses associated with severity of component failure

Consequence	Rating	Loss	Weight (l_i)
Location	High	4L	l_1
Fluid Characteristics	Extreme	8L	l_2
Fluid Pressure	Medium	2L	l_3
Criticality	Extreme	8L	l_4
Environmental	Extreme	8L	l_5
Commercial	Extreme	8L	l_6

Table 5.5: Losses associated with component failure in Circuit A

Within Circuit A all components are processing the same substances and performing very similar functions. As a result, all consequence ratings are determined to be equal over the system, meaning we can calculate a single value for the loss due to failure for all of Circuit A. Using Table 5.5 and (5.28), the loss incurred by component failure will be given by:

$$L_F = 2L(l_3 + 2(l_1 + 2(l_2 + l_4 + l_5 + l_6))) \quad (5.29)$$

Inspection losses are modelled as in section 5.2.1, with values $L_{SU} = .5$ and $L_{IC} = .025$. These values have been set to be comparatively smaller than the L_R and L_F values so that the 'value' of reducing uncertainty is higher.

5.6 Designs

As explained in section 2.3.1, an inspection design will consist of instructions stating which components are to be inspected, what percentage is to be inspected

(now restricted to either 50% or 100%), and when they are to be inspected. For our discussion we are restricting to designs in which the smallest unit is half a component. That is a component can be half inspected, completely inspected or not inspected. Even under this restriction we are still left with 2^{80} potential designs, as we have 80 half components that are either inspected or not inspected.

We will show how to calculate the utility of the following designs. Firstly, the case in which there is no further inspection. In this situation we base our replacement decisions on our beliefs about the system at time $t = 12$. Secondly, we will consider the design in which we inspect all components. This will give us an idea of the most it is possible to learn about the system at time $t = 13$, but we must also consider the expense of inspecting. To obtain a scale on which to compare expected loss values it is advisable to first consider the scores obtained for the two extreme cases - those of no inspection and total inspection. Assessing the performance of these baseline cases is important for making expected loss values more interpretable. We work through the expected loss calculations for these designs and then show how to evaluate the expected loss of a partial inspection design. In chapter 6 we will consider how to identify potentially 'good' inspection designs, but here we illustrate the calculations involved in evaluating design utility.

We must specify a time point for the inspection. To learn most about the system at $t + k = 23$ we would like to inspect close to $t + k$, but delaying too long before inspecting may result in failures prior to inspection and losses being incurred. For our illustration we will consider a total inspection at $t = 13$, near the start of the $[t, t + k] = [12, 23]$. In general, the problem of when to inspect should be treated as carefully as that of where to inspect, as fixing the inspection point constrains us to being able to find the best inspection plan for this time step, which is not necessarily the same as the best inspection plan, but for the purposes of our illustration we omit the further considerations required for delaying inspection.

5.6.1 Lower baseline: the no inspection case

The value of the no inspection case allows us to establish the minimum level

of performance we require from a design to make it worth implementing. The no inspection case will always have the lowest inspection loss ($L_I = 0$) and the least influence on our predictive ability (no improvement). In certain situations, it may still prove to be the most effective design. For example, when the cost of inspection is high but the benefits minimal or when the cost of replacement is approximately equal to the cost of failure ($L_F \approx L_R$), then choosing not to inspect and replace components as they fail will be the best option. However, in many situations $L_F \gg L_R$, and so no inspection becomes a less attractive option.

To evaluate the expected loss criterion for the no inspection case we need to calculate, from (5.8):

$$\begin{aligned} E[L(d)] &= L_R P(P(F|y) > \rho) + L_F E[P(F|y) | P(F|y) \leq L] P(P(F|y) \leq \rho) \\ L_I &= L_{SU} + n_d L_{IC}. \end{aligned} \quad (5.30)$$

As we will not be inspecting, there is no uncertainty over the probability of choosing action a or \bar{a} . Our beliefs will not be updated, because we receive no new data, so we can make all replacement decisions based on our initial beliefs (for Circuit A, beliefs at time $t = 12$), so for the no inspection case, we do not have to deal with the more complicated integration that usually forms part of the expected loss calculation. Instead it is sufficient to consider:

$$z = \frac{W_C - E_{y_h}(m_{23})}{\sqrt{\text{var}_{y_h}(m_{23})}}$$

and $\Phi(z)$ where y_h is the available historical data. The critical wall thickness values are set as being 40% of the nominal wall thickness or 3mm, whichever is larger. The values used in calculating z for this example are given in Tables 5.6 and 5.7. Note that a positive value for $W_C - E_{y_h}(m_{23})$ indicates that we are expecting the component to have failed by $t = 23$.

$\Phi(z)$ gives us a vector of failure probabilities, in which each element corresponds to the failure probability of a particular component. We then construct the vector $p(\delta^*(y) = a)$ (the vector of probabilities of choosing to replace a component) elementwise as:

$$p(\delta_c^*) = \begin{cases} 1 & \text{if } \Phi(z_c) > \frac{L_R(c)}{L_F(c)} \\ 0 & \text{otherwise} \end{cases}$$

Component Number	Difference between critical and expected minimum wall thickness ($W_C - E_{y_h}(m_{23})$)	Variance of min wall thickness No Inspection $\text{var}_{y_h}(m_{23})$	Variance of min wall thickness Full Inspection $\text{var}_{y_h,13}(m_{23})$	Variance of min wall thickness Partial Inspection $\text{var}_{y_h,d13}(m_{23})$
1	-4.6776	11.8499	5.1172	11.6023
2	-6.6283	2.4654	0.6394	0.6411
3	-4.4668	10.5392	4.9645	7.5551
4	-4.1153	9.2645	4.9271	4.9517
5	-6.5509	2.4212	0.6356	2.4010
6	1.2695	2.0453	1.5085	1.5097
7	-2.9031	8.0491	4.6223	6.6256
8	-5.1638	6.4560	3.1338	3.1825
9	-9.7274	0.9586	0.5337	0.9207
10	-4.6723	1.3541	0.6210	0.6223
11	-4.6538	1.5631	0.7220	1.4237
12	-1.1529	7.4594	4.1057	4.1103
13	-4.6758	1.5771	0.7246	1.4069
14	-2.7634	9.8367	4.7685	4.7740
15	-7.4621	2.3040	0.9443	2.1068
16	-4.6273	9.2711	4.9363	4.9375
17	-6.6482	2.4530	0.6357	2.4311
18	-6.7314	2.4146	0.6333	0.6337
19	1.6827	1.5804	0.9213	1.4575
20	-5.8960	3.8255	0.6874	0.6888

Table 5.6: Changes in variance values for different inspection designs

Component Number	Difference between critical and expected minimum wall thickness $(W_C - E_{y_h}(m_{23}))$	Variance of min wall thickness No Inspection $\text{var}_{y_h}(m_{23})$	Variance of min wall thickness Full Inspection $\text{var}_{y_h,13}(m_{23})$	Variance of min wall thickness Partial Inspection $\text{var}_{y_h,d13}(m_{23})$
21	-0.4959	5.3345	2.3208	4.8294
22	-0.4673	5.2698	2.3153	2.3190
23	-0.0354	1.4637	0.9775	1.3604
24	-0.2010	1.3393	0.9841	0.9923
25	0.2271	4.1741	1.9117	4.1053
26	0.2689	4.0869	1.8845	1.8875
27	1.7683	1.7323	0.9442	1.7320
28	-5.8235	3.8319	0.6971	0.6972
29	-5.6854	3.7369	0.6843	3.7183
30	1.0838	1.0707	0.7507	0.7510
31	-1.3110	1.9394	1.1618	1.7639
32	-2.6150	2.8071	1.4482	1.4736
33	-0.5081	1.4885	1.0579	1.3349
34	-2.8080	2.7079	1.4478	1.4761
35	0.5704	1.3857	0.9267	1.2391
36	0.5509	1.3635	0.9240	0.9343
37	-3.1682	2.8960	1.4655	2.5595
38	-0.7959	5.4764	1.6344	1.6370
39	0.8535	1.0882	0.9321	1.0774
40	0.5318	1.3425	0.9354	0.9384

Table 5.7: Changes in variance values for different inspection designs

So $p(\delta_c^*)$ takes value 1 if the component is to be replaced and 0 if no action is deemed necessary. For the no further inspection case, (5.8) can be rewritten in terms of $p(\underline{\delta}^*(y) = a)$, L_R , L_F and $\Phi(\underline{z})$ as:

$$E[L(d)] = \underline{L}_R' \cdot p(\underline{\delta}^*(y)) + \underline{L}_F' \cdot [(1 - p(\underline{\delta}^*(y))) \cdot \Phi(\underline{z})] \quad (5.31)$$

The inspection loss, L_I , for the no inspection case is always 0.

The necessary numbers for the explicit expected loss calculation are given in Tables 5.8 and 5.9. Recalling the definition of the expected utility:

$$U(d) = -(L_R P(P(F|y) > \rho) + L_F E[P(F|y)|P(F|y) \leq \rho] P(P(F|y) \leq \rho) + L_I)$$

we can see we are left with two terms to calculate ($L_I = 0$ in this case). The 2nd and 3rd columns of Tables 5.8 and 5.9 correspond to the values of $P(P(F|y) > \rho)$ and $E[P(F|y)|P(F|y) \leq \rho] P(P(F|y) \leq \rho)$ respectively. We then multiply the first of these quantities by the component replacement loss and the second by the failure loss to obtain an overall loss per component. These losses are then summed to give the total decision based loss for the design. The cost of implementing the design ($L_I = 0$) is then added to this quantity to give the full loss value for the design.

The presence of a 1 in column 2 of Tables 5.8 and 5.9 indicates we would choose to replace a component. In the no further inspection case we can say with certainty which components we would replace. For this example we would recommend replacing 26 of the 40 components, at a loss of 30.3. The 14 unreplaced components generate an expected loss of 1.8003, most of which is attributable to component 16, in which the probability of failure falls just below the accepted standard. This expected loss is obtained as the probability of component failure multiplied by the cost of failure. For the no further inspection case components can only have either a replacement cost or a failure cost, as there is no uncertainty over our decisions.

Summing the values in Tables 5.8 and 5.9 we find the utility value for the no inspection case to be

$$\begin{aligned} U(\emptyset) &= -\{E[L(\emptyset)] + L_I\} = -(L_R I_1 + L_F I_2 + L_I) \\ &= -(30.3 + 1.8003 + 0) = -32.1003 \end{aligned}$$

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) P(F y) \leq \rho] \times P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1	1	0	1.5	0	1.5
2	0	0.0000	0	0.0002	0.0002
3	1	0	1.5	0	1.5
4	1	0	1.5	0	1.5
5	0	0.0000	0	0.0003	0.0003
6	1	0	1.5	0	1.5
7	1	0	1.5	0	1.5
8	0	0.0211	0	0.4213	0.4213
9	0	0.0000	0	0.0000	0.0000
10	0	0.0000	0	0.0006	0.0006
11	0	0.0001	0	0.0020	0.0020
12	1	0	1.5	0	1.5
13	0	0.0001	0	0.0020	0.0020
14	1	0	1.5	0	1.5
15	0	0.0000	0	0.0000	0.0000
16	0	0.0643	0	1.2858	1.2858
17	0	0.0000	0	0.0002	0.0002
18	0	0.0000	0	0.0001	0.0001
19	1	0	2	0	2
20	0	0.0013	0	0.0257	0.0257

Table 5.8: Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) $ $P(F y) \leq \rho \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
21	1	0	2	0	2
22	1	0	2	0	2
23	1	0	0.6	0	0.6
24	1	0	0.6	0	0.6
25	1	0	2	0	2
26	1	0	2	0	2
27	1	0	2	0	2
28	0	0.0015	0	0.0293	0.0293
29	0	0.0016	0	0.0327	0.0327
30	1	0	0.6	0	0.6
31	1	0	0.6	0	0.6
32	1	0	0.6	0	0.6
33	1	0	0.6	0	0.6
34	1	0	0.6	0	0.6
35	1	0	0.6	0	0.6
36	1	0	0.6	0	0.6
37	1	0	0.6	0	0.6
38	1	0	0.6	0	0.6
39	1	0	0.6	0	0.6
40	1	0	0.6	0	0.6

Table 5.9: Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information

5.6.2 Upper baseline: the total inspection case

Calculating the expected loss value for the total inspection case provides us with a way of establishing the maximum we should pay for an inspection. It is impossible to learn more about the system state by using another design, but total inspection will always be the most expensive inspection (under our choice of cost function). Total inspection is unlikely to be a practical choice of inspection design in practice, but by evaluating the worth of a complete inspection we can ascertain how successful other potential designs are in improving our predictive ability relative to the best performing design in terms of information gain.

To calculate the expected loss for total inspection at $t = 13$ we need to evaluate (5.8) and (5.30). In this case, unlike the no further inspection case, we are uncertain about our choice of action, as we have not yet seen the inspection data. Therefore we will need to use the method of section 5.4 to calculate $E[L(1, \dots, 40)]$. To do this we need to update our beliefs about the variance of \underline{m}_{23} given a total inspection at $t = 13$.

For the total design we are considering full inspection of all components, so our design set $d = \{1, 2, \dots, 39, 40\}$. To find the adjusted variance for this design we need to evaluate:

$$\text{var}_{\underline{y}_{h,13}}(\underline{m}_{23}) = \text{var}(\underline{m}_{23}) - \text{cov}(\underline{m}_{23}, \underline{y}_{h,13}) \text{var}^{-1}(\underline{y}_{h,13}) \text{cov}(\underline{y}_{h,13}, \underline{m}_{23}) \quad (5.32)$$

The chosen d has 40 elements, so the variance matrix of our design observations, $\text{var}(\underline{y}_{13})$ is of size 40×40 . The covariances, $\text{cov}(\underline{m}_{23}, \underline{y}_{h,13})$ and $\text{cov}(\underline{y}_{h,13}, \underline{m}_{23})$, are both of size 40×40 , as are the $\text{var}_{\underline{y}_{h,d13}}(\underline{m}_{23})$ and $\text{var}_{\underline{y}_h}(\underline{m}_{23})$ matrices. We can estimate the required quantities from the simulation output. To estimate the covariance structure, we require the output relating to the whole component observations at time $t = 13$ for every component. This is given by the columns of $\mathbf{Y}_{13}^{(1)}$ corresponding to the design elements, in this case, all columns. We will also need the simulation output for the component minima at time $t = 23$ for all components. This is given by \mathbf{M}_{23} . The covariance matrix is then defined to be the matrix in which the (i, j) -th

element is equal to:

$$\sigma_{m_i y_j} = \frac{\sum_{k=1}^S m_{ik} y^{(1)}_{jk}}{S} - \frac{\sum_{k=1}^S m_{ik}}{S} \cdot \frac{\sum_{k=1}^S y^{(1)}_{jk}}{S} \quad (5.33)$$

where i ranges from $1, \dots, 40$ and j ranges over the elements of d .

Using the Bayes linear formula for adjusted variance and the simulation output for observations at $t = 13$ and system minima at $t = 23$, we can estimate $\text{var}_{y_{13}}(m_{23})$ and thus evaluate $\underline{\mu}_z$ and $\underline{\sigma}_z^2$ as:

$$\underline{\mu}_z = \frac{W_C - E(m_{23})}{\sqrt{\text{var}_{y_{h,13}}(m_{23})}} \quad (5.34)$$

$$\underline{\sigma}_z^2 = \frac{\text{var}(m_{23})}{\text{var}_{y_{h,13}}(m_{23})} - 1 \quad (5.35)$$

which is sufficient to allow the evaluation of the integrals I_1 and I_2 .

$$I_1 = \Phi \left\{ - \left(\frac{\Phi^{-1}(\rho) - \mu_z}{\sigma_z} \right) \right\} \quad \text{and} \quad I_2 = \int_{-\infty}^{\Phi^{-1}(\rho)} \Phi(z) \phi \left(\frac{\mu_z - z}{\sigma_z} \right) \frac{dz}{\sigma_z}$$

thereby calculating the expected loss value.

We can use a standard Normal cdf to estimate the value of I_1 for all c . I_2 can be estimated, either by using software packages designed for integral evaluation, or by using simulation techniques. We adopted a simulation approach, based on a change of variable in integral I_2 from z to g , where:

$$g = \frac{z - \mu_z}{\sigma_z}$$

allows us to write I_2 as:

$$I_2 = \int_{-\infty}^{\rho_g} \Phi(\sigma_z g + \mu_z) \phi(g) dg \quad (5.36)$$

where,

$$\rho_g = \frac{\Phi^{-1}(\rho) - \mu_z}{\sigma_z}.$$

We can treat (5.36) as evaluating the conditional expectation of a function of g , thus:

$$I_2 = E[\Phi(\sigma_z g + \mu_z) | g < \rho_g] p(g < \rho_g). \quad (5.37)$$

$g \sim N(0, 1)$, by construction, so we can estimate the value of I_2 using the following simulation routine:

1. Generate a sample G of size $|G|$ independent draws from a standard Normal distribution with mean 0 and variance 1.
2. Find all $g \in G$ such that $g < \rho_g$, call this set Γ .
3. Let $|\Gamma|$ be the size of set Γ . The expectation we want is estimated by:

$$E[\Phi(\sigma_z g + \mu_z) | g < \rho_g] = \frac{\sum_{g \in \Gamma} g}{|G|}$$

4. Estimate the probability of $g < \rho_g$ by:

$$p(g < \rho_g) = \frac{|\Gamma|}{|G|}$$

the proportion of $g \in G$ less than ρ_g .

5. Estimate I_2 by

$$I_2 = \frac{\sum_{g \in \Gamma} g}{|\Gamma|} \cdot \frac{|\Gamma|}{|G|} = \frac{\sum_{g \in \Gamma} g}{|G|}.$$

It may be necessary to make the size of $|G|$ very large in order to ensure convergence, particularly when $\frac{L_R}{L_F} = \rho$ is small.

Tables 5.10 and 5.11 contains all the necessary information for calculating the expected loss value for the total inspection case. The table has the same layout as that of Tables 5.8 and 5.9. In this case, we are not certain about our decision, so every component takes a value for both the I_1 and I_2 terms, even if this value is negligible in some cases. Using the values in Tables 5.10 and 5.11 we can show the expected loss from decisions $E[L(1, \dots, 40)] = 26.1073$, which can be split into an expected loss of 23.4375 from replacement and 2.6698 from failure. As a consequence of the uncertainty introduced into our decision making through inspecting, the loss due to replacement can no longer be definitely stated, but is now replaced by an expected loss evaluated as the sum of the expected loss due to replacement for each component. This is calculated as the product of the probability of choosing to replace a component multiplied by the cost of replacing that component. The expected loss due to failure is calculated as previously. In this case we also need to specify the inspection cost. Using the specified values for $L_{SU} = .5$ and $L_{IC} = .0125$, the total inspection loss is:

$$L_I = .5 + 80 * .0125 = 1.5$$

and the overall expected loss for this design is therefore:

$$\begin{aligned} U(1, \dots, 40) &= -\{E[L(1, \dots, 40)] + L_I\} = -(L_R I_1 + L_F I_2 + L_I) \\ &= -(23.4375 + 2.6698 + 1.5) = -27.6073 \end{aligned}$$

The main reason for the reduction in expected loss we see here is the decrease in the number of components we believe it is necessary to replace. A definite replacement policy will depend on the particular data observed. However, the probabilities of our deciding to replace each component are shown in column 2 of Tables 5.10 and 5.11. The reduction in variance achieved by fully inspecting at $t = 13$ can be seen by comparing the 3rd and 4th columns of Tables 5.6 and 5.7.

The values of Table 5.10 and 5.11 are presented without adjustment for inspection cost. We see that the extra certainty we have gained in our predictive ability, because of the reduction in variance brought about by updating, ensures that the loss in utility is at most as large as when not inspecting, and in general reduced. In particular, if we consider component 16, which contributed most of the failure cost in the no further inspection case, due to it having an expected probability of failure close to the value of ρ , we can see that its loss value has been halved. This occurs as a consequence of introducing the possibility of replacing component 16, depending on what we observe during our inspection. The updated distribution of the probability of failure for component 16 is such that we would expect to replace it 24.63% of the time. Allowing the possibility of replacement - dependent on the actual observations - has for this component reduced the expected loss. For components such as 1, 3 and 4, allowing the possibility of not replacing has caused the reduction in expected loss. These components also have an expected probability of failure close to the value ρ (as for component 16), but slightly greater meaning without further inspection these components would always be replaced. Further inspection allows us to confirm situations in which this is the most appropriate action.

The benefits of inspection can be seen here to be that of allowing decisions to be delayed until more relevant information is available. For those components with an expected probability of failure close to the value of ρ , further inspection is particularly valuable. Inspection does not necessarily change the number of components

that must be replaced, but it allows a more appropriate decision to be made on the basis of stronger evidence.

From evaluating $E[L(d)]$ when d is equivalent to total inspection we know that greatest possible gain in information we can achieve from any design is that which allows us to get close to -26.1073. Partial designs will not allow us to gain a greater value of information than $32.1003 - 26.1073 = 5.9030$, but they may allow us to learn nearly as much for lower L_I .

5.6.3 Partial inspection

The procedure for evaluating the expected loss criterion for a partial inspection design is similar to that of evaluating the criterion for a full inspection design. The same calculations are performed using the same methods outlined in section 5.6.2. We consider inspecting all the even number components fully at time 13. This means we take the minimum over every location on the surface observed with error for 20 components within Circuit A. At this stage we make no attempt to select these 20 components in an informed way. The purpose at this stage is to illustrate how to use the criterion for a given partial design, not to select an optimal design. The set of components to be inspected is $C_d = \{2, 4, \dots, 38, 40\}$, components 1, 3, 5, \dots , 37, 39 will not be inspected.

Evaluating expected loss for any design relies on calculating the adjusted variance. For our partial design we are considering full inspection of the even numbered components, so our design set $d = \{2, 4, 6, \dots, 38, 40\}$. To find the adjusted variance for this design we need to evaluate:

$$\text{var}_{\underline{y}_{h,d13}}(\underline{m}_{23}) = \text{var}_{\underline{y}_h}(\underline{m}_{23}) + \text{cov}(\underline{m}_{23}, \underline{y}_{h,d13}) \text{var}^{-1}(\underline{y}_{d13}) \text{cov}(\underline{y}_{h,d13}, \underline{m}_{23}) \quad (5.38)$$

$$\text{var}_{\underline{y}_{h,d13}}(\underline{m}_{23}) = \text{var}(\underline{m}_{23}) - \text{cov}(\underline{m}_{23}, \underline{y}_{h,d13}) \text{var}^{-1}(\underline{y}_{h,d13}) \text{cov}(\underline{y}_{h,d13}, \underline{m}_{23}) \quad (5.39)$$

The chosen d has 20 elements, so the variance matrix of our design observations, $\text{var}(\underline{y}_{d13})$ is of size 20×20 . The covariances, $\text{cov}(\underline{m}_{23}, \underline{y}_{h,d13})$ and $\text{cov}(\underline{y}_{h,d13}, \underline{m}_{23})$, are of size 40×20 and 20×40 respectively. Both the $\text{var}_{\underline{y}_{h,d13}}(\underline{m}_{23})$ and $\text{var}_{\underline{y}_h}(\underline{m}_{23})$

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) P(F y) \leq \rho] \times P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1	0.2919	0.0112	0.4379	0.2244	0.6623
2	0.0000	0.0000	0.0000	0.0003	0.0004
3	0.2969	0.0121	0.4453	0.2427	0.6880
4	0.3293	0.0134	0.4940	0.2686	0.7626
5	0.0000	0.0000	0.0000	0.0003	0.0004
6	1.0000	0.0000	1.5000	0.0001	1.5000
7	0.5413	0.0132	0.8119	0.2632	1.0751
8	0.0561	0.0091	0.1122	0.1829	0.2951
9	0.0000	0.0000	0.0000	0.0000	0.0000
10	0.0000	0.0000	0.0000	0.0006	0.0006
11	0.0002	0.0001	0.0002	0.0018	0.0020
12	0.8323	0.0059	1.2484	0.1171	1.3655
13	0.0002	0.0001	0.0002	0.0018	0.0020
14	0.5670	0.0111	0.8506	0.2213	1.0719
15	0.0000	0.0000	0.0000	0.0000	0.0000
16	0.2463	0.0129	0.3694	0.2587	0.6281
17	0.0000	0.0000	0.0000	0.0003	0.0003
18	0.0000	0.0000	0.0000	0.0002	0.0002
19	0.9998	0.0000	1.9997	0.0002	1.9999
20	0.0032	0.0004	0.0064	0.0073	0.0137

Table 5.10: Necessary elements of the expected loss calculation for each component in the full inspection case, excluding inspection cost information

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) $ $P(F y) \leq \rho] \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
21	0.7993	0.0083	1.5985	0.1655	1.7640
22	0.8058	0.0082	1.6117	0.1634	1.7751
23	0.9956	0.0001	0.5973	0.0017	0.5990
24	0.9974	0.0001	0.5984	0.0011	0.5995
25	0.9081	0.0042	1.8162	0.0848	1.9009
26	0.9141	0.0040	1.8283	0.0809	1.9092
27	0.9997	0.0000	1.9993	0.0003	1.9996
28	0.0036	0.0004	0.0073	0.0079	0.0152
29	0.0041	0.0005	0.0081	0.0092	0.0173
30	1.0000	0.0000	0.6000	0.0000	0.6000
31	0.7917	0.0028	0.4750	0.0557	0.5307
32	0.3815	0.0045	0.2289	0.0903	0.3192
33	0.9851	0.0003	0.5911	0.0053	0.5964
34	0.3137	0.0047	0.1882	0.0946	0.2828
35	0.9998	0.0000	0.5999	0.0001	0.5999
36	0.9998	0.0000	0.5999	0.0001	0.5999
37	0.2281	0.0040	0.1368	0.0809	0.2178
38	0.7941	0.0018	0.4765	0.0360	0.5125
39	1.0000	0.0000	0.6000	0.0000	0.6000
40	0.9999	0.0000	0.5999	0.0001	0.6000

Table 5.11: Necessary elements of the expected loss calculation for each component in the full inspection case, excluding inspection cost information

matrices are 40×40 . We can estimate the required quantities from the simulation output. To estimate the covariance structure, we require the output relating to the whole component observations at time $t = 13$ for the components in the design. This is given by the columns of $Y_{13}^{(1)}$ corresponding to the design elements, i.e. the 2nd column of $Y_{13}^{(1)}$ contains the output for component 2, the 4th the output for component 4 etc. We will also need the simulation output for the component minima at time $t = 23$ for all components. This is given by M_{23} . The covariance matrix is then defined to be the matrix in which the (i, j) -th element is defined using (5.33) where i ranges from $1, \dots, 40$ and j ranges over the elements of d .

Having estimated the relevant variance and covariance structures, it is possible to calculate the adjusted variance matrix for design d . This is sufficient to allow the computation of $\underline{\mu}_z$ and $\underline{\sigma}_z$, as in (5.34) and (5.35), which allows us to evaluate the expected loss criterion for this design. Comparing the adjusted variances for the partial design to those from the full and no inspection designs (Tables 5.6 and 5.7), we can see that the partial design performs nearly as well as full inspection in the components we have observed, indicating that the majority of explainable variation in most components is learned about through direct observation of the component. However, the design also achieves variance reduction in the components we do not inspect, most notably in components 3 and 7. Learning about uninspected components through the system covariances is an important aspect of inspection planning. Partial inspection designs are cheaper to implement than total inspections. Therefore if we can successfully target our inspections to the components that are most informative about not just themselves, but also the system as a whole, we will be able to develop more cost effective inspection plans.

The necessary information for calculating the expected loss values for this partial inspection case is shown in Tables 5.12 and 5.13. As for full inspection, every component must take a value for both I_1 and I_2 , as partial inspections act to reduce uncertainty over the whole system, not just in the inspected components. Assessing the loss values componentwise we can see that the partial design compares favourably with the full design in the inspected components, achieving similar reductions in loss, but less well in the uninspected components. However, it still

outperforms not inspecting, achieving a loss less than or equal to that seen in the no further inspection case (as must happen). For example, if we again consider component 16, the partial design achieves an almost identical reduction utility to the full design, implying that information from other components is of less importance when learning about component 16. If we assess the performance of components 1 and 3 - neither inspected under the partial design, we can see a noticeable difference in degree of reduction achieved. Component 1 has an associated loss of 1.5 given no further inspection, but this drops to 0.6623 under full inspection. Our partial design manages to reduce the loss to just 1.4198, which is not much of an improvement, indicating that we have to consider components outside of the current partial design to learn effectively about component 1. Component 3 also has an associated loss of 1.5 given no further inspection, and achieves a similar reduction to 0.6880 under full inspection. However, for this component, the partial design allows us to reduce loss to 0.9433, a considerable improvement on not inspecting, but not as much as is offered by full inspection. This tells us component 3 is a component we can learn a non-trivial amount about through inspecting other components.

The partial design does not offer as much improvement in our decision making ability as the full design, and consequently both the losses due to replacement (25.2752) and due to failure (2.7750) are higher than those for the full inspection case, however, the design is cheaper to implement, as we inspect half as many units (40 half components as opposed to 80). The losses due to inspection are:

$$L_I = .5 + 40 * .0125 = 1$$

and the overall expected loss for this design is therefore:

$$\begin{aligned} U(d) &= -\{E[L(d)] + L_I\} = L_R I_1 + L_F I_2 + L_I \\ &= -(25.2752 + 2.7750 + 1) = -29.0502 \end{aligned}$$

So of these three designs, the full inspection is preferable. This will not always be the case. The arbitrary fashion in which this design was selected meant the quality of its performance could not be predicted. We will demonstrate in the next chapter that the use of sensible design selection procedures will allow the inspection planner

to identify partial inspection designs which, due to their reduced cost, offer better value information than total inspection.

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) $ $P(F y) \leq \rho] \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1	0.6750	0.0204	1.0125	0.4074	1.4198
2	0.0000	0.0000	0.0000	0.0003	0.0004
3	0.3839	0.0184	0.5759	0.3674	0.9433
4	0.3303	0.0135	0.4954	0.2695	0.7649
5	0.0000	0.0000	0.0000	0.0003	0.0003
6	1.0000	0.0000	1.5000	0.0001	1.5000
7	0.7493	0.0116	1.1240	0.2327	1.3567
8	0.0559	0.0093	0.1117	0.1865	0.2983
9	0.0000	0.0000	0.0000	0.0000	0.0000
10	0.0000	0.0000	0.0000	0.0006	0.0006
11	0.0000	0.0001	0.0000	0.0020	0.0020
12	0.8327	0.0058	1.2490	0.1168	1.3658
13	0.0000	0.0001	0.0000	0.0020	0.0020
14	0.5674	0.0111	0.8511	0.2210	1.0721
15	0.0000	0.0000	0.0000	0.0000	0.0000
16	0.2463	0.0129	0.3694	0.2588	0.6283
17	0.0000	0.0000	0.0000	0.0002	0.0002
18	0.0000	0.0000	0.0000	0.0002	0.0002
19	1.0000	0.0000	2.0000	0.0000	2.0000
20	0.0032	0.0004	0.0064	0.0074	0.0137

Table 5.12: Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) $ $P(F y) \leq \rho] \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
21	0.9995	0.0000	1.9989	0.0006	1.9995
22	0.8062	0.0082	1.6125	0.1635	1.7759
23	1.0000	0.0000	0.6000	0.0000	0.6000
24	0.9977	0.0001	0.5986	0.0010	0.5997
25	1.0000	0.0000	2.0000	0.0000	2.0000
26	0.9144	0.0041	1.8289	0.0811	1.9100
27	1.0000	0.0000	2.0000	0.0000	2.0000
28	0.0036	0.0004	0.0073	0.0079	0.0152
29	0.0000	0.0016	0.0000	0.0328	0.0328
30	1.0000	0.0000	0.6000	0.0000	0.6000
31	0.9977	0.0001	0.5986	0.0011	0.5998
32	0.3869	0.0046	0.2322	0.0923	0.3244
33	1.0000	0.0000	0.6000	0.0000	0.6000
34	0.3187	0.0048	0.1912	0.0959	0.2872
35	1.0000	0.0000	0.6000	0.0000	0.6000
36	0.9999	0.0000	0.5999	0.0001	0.6000
37	0.3919	0.0095	0.2351	0.1891	0.4243
38	0.7944	0.0018	0.4767	0.0361	0.5127
39	1.0000	0.0000	0.6000	0.0000	0.6000
40	0.9999	0.0000	0.5999	0.0001	0.6000

Table 5.13: Necessary elements of the expected loss calculation for each component in the no further inspection case, excluding inspection cost information

Chapter 6

Selecting inspection designs

For the inspection of any system, there will be a choice of inspection design that is optimal under our modelling assumptions. In the decision problem framework, this is the design which allows us to make correct decisions most often, and, in terms of our formulation of the problem, it is the design which yields the greatest improvement in decision making ability offset against the cost of implementing the design. We have discussed previously some of the computational issues involved in planning inspection for large systems, and these are again relevant when considering how best to select an inspection design. The number of possible inspection designs for any real system is likely to be too large for it to be feasible to analyse all designs, unless stringent constraints are placed on how inspection is carried out. For example, even if we restrict to the simplest case of recommending whole components within an n component system are either inspected or not inspected, the number of potential designs will be 2^n , so, even for this simple procedure, we would have to abandon exhaustive search routines for comparatively small n . Given that inspection designs are, in general, more complex, it becomes clear that attempting to search the whole design space is not a sensible policy.

Exhaustive analysis of the design space is guaranteed to find the optimal design. However, it will in general be too computationally intensive to be practical. Consequently, we need to consider rational methods for selecting designs, which reduce the number of designs it is necessary to compare whilst identifying designs that perform

well in terms of improving our decision making ability.

In this chapter we discuss choosing between inspection designs (section 6.1). In section 6.2 we outline the different search strategies we will use and how we select initial designs from which to start searching. In section 6.3 we illustrate the use of the different routines to plan an inspection for the Site A subsystem Circuit A, and we analyse the performance of the recommended designs in section 6.4. In section 6.5 we assess the impact of varying the failure costs.

6.1 Identifying good designs

The design criterion described in chapter 5 is used to choose between designs. The expected utility criterion has been defined in such a way that both the losses incurred as a result of the system behaviour and those involved in making the inspection are taken into account within the criterion. This means we can compare different designs (made at the same time point) which cost different amounts to implement directly through the expected utility criterion, without the need for further calculation.

We always prefer designs which have a higher expected utility value. Therefore the optimal design is the single design $d^* \in D$ which has the minimum associated expected loss in utility (or simply expected loss), and ‘good’ designs are those which take expected loss values close to this minimum. The cost of making the inspection plays an important role here, and the optimal design is the one which balances the reduction in loss due to increased confidence in our decision making against the price paid to achieve that gain in confidence.

6.1.1 Inspection blocks

Throughout the rest of this chapter we will refer to ‘inspection blocks’. An *inspection block*, is defined to be smallest unit of a system which can be inspected. Blocks should be regarded as a partitioning of the system into regions for inspection

purposes. It should be assumed that there is no overlap between blocks, and that every point on the surface of the system can be identified with one (and only one) inspection block. Individual blocks are denoted, b_i , and the set of all inspection blocks contained in the system is B .

6.2 Search routines

We will use a stepwise search method to identify designs. Stepwise search is a simple method which adds/removes blocks from a design one at a time until there is no further benefit from adding or deleting more blocks, or until the maximum inspection budget is reached, whichever occurs first. 'Benefit' in our search algorithm will be determined by the design utility values.

We consider 3 stepwise approaches; stepwise addition of inspection blocks, stepwise deletion of inspection blocks and finally a combination of the two. We outline the algorithms used in sections 6.2.1 - 6.2.3. These searches follow the methods used in regression analysis for variable selection.

6.2.1 Stepwise Add

In implementing the stepwise procedures we will use the notation d_j to indicate the design at the j -th step of the process. Furthermore, we denote the initial design to be d_0 and the final recommended design as d .

To illustrate the process consider a system with n inspection blocks. $d_0 = \emptyset$, i.e. there are initially no inspection blocks included in the design. To find d_1 we must identify which block, $b_i^* \in B$, we should inspect if we can only inspect one. To do this we evaluate the expected loss from our decisions, $E[L(\{b_i\})]$, for the design inspecting only b_i , for all $b_i \in B$. We then let $L(b_i)$ be the loss associated with inspecting b_i and let b_1^* be the choice of block b_i which minimises $E[L(b_i)]$ over all choices $b_i \in B$. We then set:

$$d_1 = \{b_1^*\}$$

The recommended design at the next step consists of the previous recommended design, plus the inspection block b_i that yields the minimum expected loss (i.e., the maximum effect on all the remaining uninspected sites). Having identified the next proposed point we check:

$$\begin{aligned} U(d_1) &> U(d_0) \\ \Leftrightarrow -\{E[L(d_1)] + L_I(d_1)\} &> -\{E[L(d_0)] + L_I(d_0)\} \end{aligned} \quad (6.1)$$

$$\Leftrightarrow E[L(d_1)] + L_I(d_1) < E[L(d_0)] \quad (6.2)$$

i.e. that the total loss of inspecting the new set is less than that incurred from inspecting the old set. If (6.1) is true, we recommend using d_1 and proceed to add another block. However, if (6.1) does not hold there is no benefit from adding an inspection block to the design, so we stop the stepwise process and recommend using d_0 as our inspection design. Using this method, our design set at step j will consist of the set of points $\{b_1^*, b_2^*, \dots, b_j^*\}$, where $b_j^* = b_i$ such that:

$$\max_{b_i \in B} \{U(b_1^*, b_2^*, \dots, b_{j-1}^*, b_i)\} = U(b_1^*, b_2^*, \dots, b_{j-1}^*, b_j^*) \quad (6.3)$$

The stepwise add algorithm (at the j -th step) is summarised below:

Stepwise Add

1. Using $d_{j-1} = \{b_1^*, b_2^*, \dots, b_{j-1}^*\}$, identify the set of uninspected blocks, $B - \{b_1^*, b_2^*, \dots, b_{j-1}^*\}$.
2. Evaluate $U(b_1^*, b_2^*, \dots, b_{j-1}^*, b_i) \forall b_i \text{ in } B - \{b_1^*, b_2^*, \dots, b_{j-1}^*\}$ using the expected loss criterion.
3. Find the block, b_j^* so that (6.3) is satisfied, and set:

$$d_j = \{d_{j-1}, b_j^*\} = \{b_1^*, b_2^*, \dots, b_{j-1}^*, b_j^*\}$$

4. If $E[L(d_j)] + L_I(d_j) > E[L(d_{j-1})] + L_I(d_{j-1})$ repeat from 1, else stop.

This process should be repeated until either the stopping point is reached, i.e. when $U(d_j) < U(d_{j-1})$, or the maximum inspection budget is reached.

Using a stepwise approach has the advantage of reducing the number of designs for which the utility must be evaluated, relative to exhaustive search. For any one-way stepwise search in which components can only be inspected once, the number of designs for consideration decreases by one at each step, as once an inspection block has been added to the design, it remains part of the design. Therefore, the number of designs considered in finding the preferred design for system with n inspection blocks using a one-way stepwise search is at most:

$$n + (n - 1) + (n - 2) + \dots + 2 = \sum_2^n i = \frac{n(n+1)}{2} - 1 = \frac{1}{2}(n+2)(n-1)$$

and

$$\frac{1}{2}(n+2)(n-1) \ll 2^n \text{ as } n \text{ increases}$$

However, adding blocks to a design sequentially does not take account of relationships between components particularly well. A search strategy which helps to address this problem is stepwise delete

6.2.2 Stepwise Delete

The stepwise delete procedure is essentially the same as stepwise add, but instead of adding a block to the design, at each step one is removed. The computational load is the same as that of stepwise add but the delete approach allows us to take into account relationships between components differently. For example, consider the situation in which we have two inspection blocks which are both, individually, relatively uninformative, but in combination tell us a great deal about system behaviour. Under stepwise add these blocks would both be late additions to our design and their joint predictive power overlooked. However, using stepwise delete the strong joint effect would mean both were retained as part of the design for longer, strengthening the designs influence on our decision making ability.

We consider applying a stepwise delete procedure to a system which has been partitioned into n inspection blocks. For this process we will take the initial design,

d_0 , to be the saturated design, in which every block is inspected (i.e. $d_0 = B$). We write B as the set of selected blocks, b_j^* , $B = \{b_1^*, \dots, b_j^*, \dots, b_n^*\}$. At the first stage we identify b_{-j}^* , as the block, $b_i \in B$, we can remove with the least impact on the expected loss. That is, the block b_i satisfying:

$$\max_{b_i \in B} \{U(\{b_1^*, b_2^*, \dots, b_n^*\} - b_i)\} \quad (6.4)$$

The block, b_{-j}^* , for which the maximum utility is attained is then removed from the design d_0 to give d_1 , so we have:

$$d_1 = \{\{b_1^*, b_2^*, \dots, b_n^*\} - b_{-j}^*\}.$$

We then repeat this process until we reach a suitable stopping point.

As for stepwise add, the stopping point (i.e. the point at which we derive no benefit from continuing to delete points) occurs when:

$$U(d_j) < U(d_{j-1})$$

However, if we are working to an inspection budget, we may reach this point whilst still in excess of our budget. In which case it will be necessary to continue deleting points until we have a design which satisfies the budgetary constraints.

The stepwise delete algorithm, for the j -th step, is:

Stepwise Delete

1. Identify the inspected blocks, $d_{j-1} = \{b_1^*, b_2^*, \dots, b_{j-1}^*\}$.
2. Evaluate $U(\{b_1^*, b_2^*, \dots, b_{j-1}^*\} - b_i) \forall b_i \in d_{j-1}$ using the expected loss criterion.
3. Identify b_{-j}^* to be the point satisfying 6.4, and define the new design set, d_j to be:

$$d_j = \{d_{j-1} - b_{-j}^*\} = \{\{b_1^*, b_2^*, \dots, b_{j-1}^*\} - b_{-j}^*\}$$

4. If $U(d_j) > U(d_{j-1})$ repeat from 1, else stop.

Both stepwise add and stepwise delete provide a structured way of searching the design space. However, they do not necessarily find the optimal design. The designs found using stepwise searches should perform well with respect to our criterion, and the use of a structured search procedure greatly reduces the computational burden, which compensates for the loss of optimality in our design.

We have stated that a stepwise delete search is more likely to identify jointly informative subsets of inspection blocks, which makes it preferable to a stepwise add approach. However, for very large systems it will often be impractical to use a stepwise delete procedure from a saturated design, particularly if the scale of the design we wish to (or can afford to) implement is very much smaller than the system.

6.2.3 Combining stepwise add and stepwise delete methods

If we are planning inspections for a very large system, adopting a stepwise delete approach is unlikely to be sensible. A less time consuming method would be to use a combination of the stepwise add and delete techniques. In general, we will be inspecting a proportion of the system which is relatively small. Therefore it makes more sense to start from nothing and sequentially add blocks to the design, as we will reach the most cost-effective design more efficiently. However, we may not take account of any important joint relationships, so to counteract this we suggest using stepwise add followed by stepwise delete, as below:

Stepwise Add and Delete

1. Identify an initial design, d_0
2. Add blocks sequentially to the design using stepwise add algorithm, up to stopping point (either authentic or budget) at step j^* .
3. Add a further j^* blocks to the design using stepwise add.
4. Remove blocks sequentially from design d_{2j^*} using stepwise delete algorithm, until a stopping point is reached.

Using a combination method should improve on simply using the stepwise add procedure, but allows for a design search to be started from a more appropriate initial design - or at least one which is nearer in size to the design which we can afford to implement. However, in some very large systems it may not be appropriate to start a stepwise search from either an empty or a saturated design, so we will need a method for identifying reasonable starting points for our searches.

6.2.4 Identifying d_0

For very large systems stepwise searches from either empty or saturated designs are likely to remain time consuming. Therefore we propose the following method to identify plausible starting designs, d_0 , which can then be improved using the stepwise search procedures.

To identify a reasonable starting point we must identify blocks which are worth inspecting, namely those blocks for which we are most unsure of the outcome, and those which tell us most about system behaviour overall. Identifying blocks that allow us to learn about overall system behaviour is not straightforward, but we can identify blocks where we can achieve the greatest improvement in understanding individual component behaviour by considering the one dimensional updates for each block. That is, we only take into account what data from an inspection block tells us about the component of which it forms a part, and not what it tells us about the whole system.

For each component we have beliefs about the expectation and variance of the wall thickness values, which we use to describe the probability of component failure through the assumption of Normality. To update beliefs about a single component, c , we use the equations:

$$E_{y_{ct}}(m_{c(t+k)}) = E(m_{c(t+k)}) + \text{cov}(m_{c(t+k)}, y_{ct}) \text{var}^{-1}(y_{ct}) [y_{ct} - E(y_{ct})]$$

$$\text{var}_{y_{ct}}(m_{c(t+k)}) = \text{var}(m_{c(t+k)}) - \text{cov}(m_{c(t+k)}, y_{ct}) \text{var}^{-1}(y_{ct}) \text{cov}(y_{ct}, m_{c(t+k)})$$

in which all quantities are scalars. In the inspection design situation, where we deal with potential data, rather than observed data, $E[E_{y_{ct}}(m_{c(t+k)})] = E(m_{c(t+k)})$, (see

section 5) so the only changes to our beliefs come from updating $\text{var}(m_{c(t+k)})$. From the updated variances we can obtain a measure of the importance of inspecting the component.

Consider the variance of $m_{c(t+k)}$. This can be written as:

$$\text{var}(m_{c(t+k)}) = E[m_{c(t+k)} - E(m_{c(t+k)})]^2$$

which is equivalent to:

$$E[m_{c(t+k)} - E(m_{c(t+k)})]^2 = E[m_{c(t+k)} - E_{y_{ct}}(m_{c(t+k)}) + E_{y_{ct}}(m_{c(t+k)}) - E(m_{c(t+k)})]^2.$$

Using independence, we can write:

$$\begin{aligned} E[m_{c(t+k)} - E(m_{c(t+k)})]^2 &= E[m_{c(t+k)} - E_{y_{ct}}(m_{c(t+k)})]^2 + E[E_{y_{ct}}(m_{c(t+k)}) - E(m_{c(t+k)})]^2 \\ &= \text{var}_{y_{ct}}(m_{c(t+k)}) + \text{var}(E_{y_{ct}}(m_{c(t+k)})) \end{aligned}$$

This allows us to write $\text{var}(m_{c(t+k)})$ as the sum of the adjusted variance of $m_{c(t+k)}$ and the variance of the adjusted expectation of $m_{c(t+k)}$, and thereby obtain an expression for the adjusted expectation of $m_{c(t+k)}$

$$\begin{aligned} \text{var}(m_{c(t+k)}) &= \text{var}_{y_{ct}}(m_{c(t+k)}) + \text{var}(E_{y_{ct}}(m_{c(t+k)})) \\ \text{var}(m_{c(t+k)}) - \text{var}_{y_{ct}}(m_{c(t+k)}) &= \text{var}(E_{y_{ct}}(m_{c(t+k)})) \\ \sqrt{\text{var}(m_{c(t+k)}) - \text{var}_{y_{ct}}(m_{c(t+k)})} &= \text{SD}(E_{y_{ct}}(m_{c(t+k)})) \end{aligned}$$

Considering the standard deviation of the adjusted variance allows us to assess the effect of the one dimensional update. In addition, we weight this by the difference between the expected value of $m_{c(t+k)}$ and the critical wall thickness, W_C :

$$\Omega = \left| \frac{\sqrt{\text{var}(m_{c(t+k)}) - \text{var}_{y_{ct}}(m_{c(t+k)})}}{E(m_{c(t+k)}) - W_C} \right| = \left| \frac{\sigma_\mu}{\mu_{t+k} - W_C} \right|. \quad (6.5)$$

Ω provides a measure of the impact observing y_{ct} has on our beliefs about c .

As a consequence of using the Bayes linear updating equations $\text{var}_{y_{ct}}(m_{c(t+k)}) < \text{var}(m_{c(t+k)})$, so the numerator of Ω is always positive. For a component expected to be in acceptable condition, $W_C < \mu_{t+k}$, so the denominator will, in general, be

positive. However Ω can also take a negative value if $W_C > \mu_{t+k}$, i.e. if the component wall thickness is expected to have fallen below the acceptable level. Our interest is in identifying those components expected to be close to failure, and looking at those Ω with a large absolute values allows us to do this.

Ω is a criterion for measuring which components we learn most about through individual component inspection. The difference between $\text{var}_{y_{ct}}(m_{ct+k})$ and $\text{var}(m_{ct+k})$ quantifies the increased confidence in our probability of failure value, and weighting this by $\mu_{t+k} - W_C$ allows us to measure the relevance of improving beliefs about this component. Reducing uncertainty in components which are either certain to fail or certain to survive is not particularly important.

Evaluating Ω for each component allows us to identify the set of components which tell us most about themselves when inspected. We select the number required to form a sensible initial design d_0 by selecting those with the highest Ω values. This criterion makes no attempt to measure what inspecting component c tells us about the uninspected components and we rely on the assumption that the components which are most informative about themselves will also be informative about other components. Therefore we do not propose using Ω as a way of identifying final recommended designs, but instead suggest it as a sensible method for obtaining a starting point from which the stepwise selection algorithms can then be used to improve the design.

6.3 Design selection for Site A example

We now apply the methods outlined in section 6.2 to the Site A subsystem Circuit A. Circuit A is sufficiently small that we can legitimately perform a stepwise add routine taking $d_0 = \emptyset$ and similarly a stepwise delete starting from a saturated design. We illustrate both procedures here (stepwise add in section 6.3.1, stepwise delete in section 6.3.2), and also demonstrate how to use the Ω criterion to suggest a rational search starting point in section 6.3.3. We assume that our designs are not subject budgetary constraints. Our inspection blocks will correspond to half a

component, and we will therefore have 80 blocks for this system. The blocks will be labelled as b_{c1} or b_{c2} , where c is the component number and 1,2 distinguishes between different halves of the component.

We will restrict ourselves to identifying the best possible inspection design for time $t = 13$, which we consider as the 'now'. In practice, we would need to consider also recommending an inspection time, but for this illustration we deal with the constrained problem of identifying the best possible design for immediate implementation.

6.3.1 Selection using the stepwise add algorithm

We apply the stepwise add algorithm with an initial design of no inspection $d_0 = \emptyset$. We simulate 100,000 realisations of the Circuit A system for the 23 time steps which cover the time period 01/01/1998 - 31/12/2009, and 100,000 observed minimum values (i.e. the simulated component minimum plus a simulated observation error) based on simulated errors for observing half the locations within each component and simulated errors for observing the whole component. These simulations will provide us with the necessary information for evaluating the expectations, variances and covariances required to update our beliefs about components, and therefore to calculate the utility values for any given inspection design, and correspond to the results generated in section 2.

At the first stage of our additive stepwise search we have to evaluate the utility of all designs which only consider a single inspection block. Let us start by considering how inspecting half of component 1 alters our beliefs about the probability of component failure. We know from section 5.6.1 that the design with no inspection has an expected utility of 32.1003. The expected minimum wall thickness, adjusted by the historical data, $\underline{y}_h (E_{\underline{y}_h}(n_{23}))$ and the critical wall thickness values (\underline{W}_C) will remain constant over all designs, these values are given in Table 6.1.

Table 6.2 shows the components of the expected loss term for components 1 to 10. This table should be compared to Tables 5.8 and 5.9 to see how the half component observation has influenced our beliefs about the probability of failure

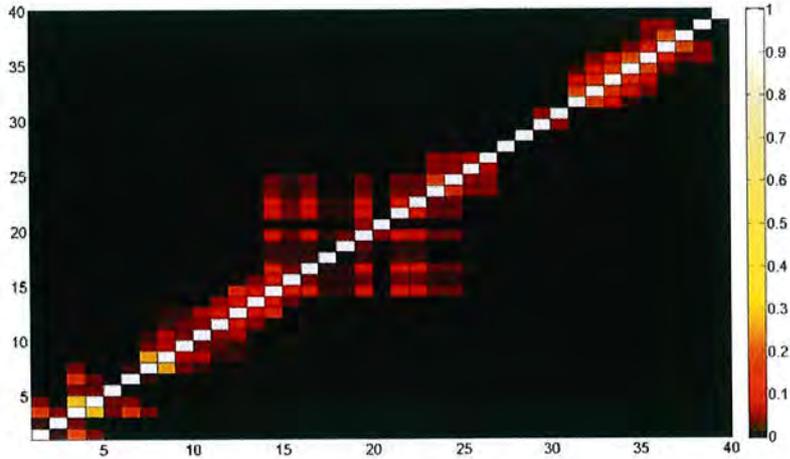


Figure 6.1: Correlations for the Circuit A system

and consequently our decision about component replacement. As we might expect from such a small inspection the overall effect is negligible. The greatest effect is on component 1, and there are also noticeable effects on components 3 and 4, but the effects on other components diminish as we move further from component 1 and a large number of components retain the same variance loss values and the same fixed replacement decisions. Recalling the correlation structure (Figure 6.1) for deviations at a single time step from chapter 3, we can see that component 1 is most strongly associated with components 3 and 4 and has effectively no association with components further away than component 6, so the effects of this inspection plan can be seen to be consistent with the correlation structure known to be present.

The expected loss in utility for this design is:

$$\begin{aligned} U(b_{c1}) &= -\{(L_R I_1 + L_F I_2) + L_{SU} + n_d L_{IC}\} \\ &= -(30.9219 + 0.5 + 1 \times 0.0125) = -31.4344 \end{aligned}$$

larger than the utility for not inspecting. This decrease in loss is caused by two factors; firstly, as we learn about the system the decisions we make about whether or not to repair may change. For example, in this case we are no longer choosing to replace component 1 every time, but instead around 30% of the time, depending

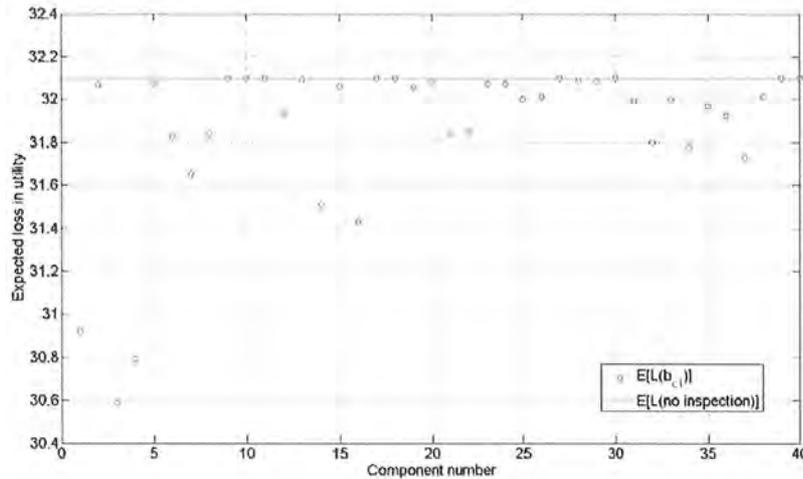


Figure 6.2: Expected utility values for single half component inspection

on the value of y_{c1} . Secondly, our beliefs about the probability of component failure change, so when we choose not to replace, we do so with greater confidence in the eventual outcome. This is reflected by the conditional expectation of failure, having chosen not to replace ($E[P(F|y)|P(F|y) \leq \rho].P(P(F|y) \leq \rho)$), being closer to 0 in components 1, 3 and 4. This is itself a consequence of updating our beliefs about the system variance, which has the effect of shortening the tails of the posterior probability of failure distribution, given our assumptions of Normality.

Consequently, the inspection is expected to be less costly than doing nothing, even accounting for the set up cost L_{SU} , and the incremental cost L_{IC} , which tell us how expensive it is to perform the inspection. For this component, we are paying the inspection cost for appreciable information gain, so the utility increases, telling us this is potentially a sensible design.

We then have to evaluate $U(b_{c1})$ for all other components. The results are shown in Figure 6.2. The figure clearly shows the expected utility value to be lowest for half inspection of component 3. Table 6.3 shows the elements of the expected loss from decisions calculations, which should again be compared to Tables 5.8 and 5.9 to establish where improvements have been made. The full version of the table is in the appendix, components for which the expected losses are noticeably affected are given in Table 6.3. The critical wall thickness (W_C) and expectations are the

Component Number	Critical Wall Thickness (W_C)	Expected minimum wall thickness ($E_{y_h}(m_{L23})$)	Component Number	Critical Wall Thickness (W_C)	Expected minimum wall thickness ($E_{y_h}(m_{L23})$)
1	9.2000	13.8776	21	3.4960	3.9919
2	7.2960	13.9243	22	3.4960	3.9633
3	9.2000	13.6668	23	3.0000	3.0354
4	9.2000	13.3153	24	3.0000	3.2010
5	7.2960	13.8469	25	3.0000	2.7729
6	3.0000	1.7305	26	3.0000	2.7311
7	7.2960	10.1991	27	3.0000	1.2317
8	9.2000	14.3638	28	7.2960	13.1195
9	9.2000	18.9274	29	7.2960	12.9814
10	5.3960	10.0683	30	3.0000	1.9162
11	5.2440	9.8978	31	3.4960	4.8070
12	5.2440	6.3969	32	5.4000	8.0150
13	5.2440	9.9198	33	3.0000	3.5081
14	7.2960	10.0594	34	5.4000	8.2080
15	9.2000	16.6621	35	3.0000	2.4296
16	9.2000	13.8273	36	3.0000	2.4491
17	7.2960	13.9442	37	5.4000	8.5682
18	7.2960	14.0274	38	5.4000	6.1959
19	3.0000	1.3173	39	3.0000	2.1465
20	7.2960	13.1920	40	3.0000	2.4682

Table 6.1: Expectations and critical wall thickness values for Circuit A system based on data to time $t = 12$

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) P(F y) \leq \rho] \times P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1	0.2928	0.0113	0.4392	0.2253	0.6644
2	0.0000	0.0000	0.0000	0.0002	0.0002
3	0.4855	0.0240	0.7282	0.4801	1.2083
4	0.7482	0.0163	1.1222	0.3264	1.4486
5	0.0000	0.0000	0.0000	0.0003	0.0003
6	1.0000	0.0000	1.5000	0.0000	1.5000
7	1.0000	0.0000	1.5000	0.0000	1.5000
8	0.0000	0.0211	0.0000	0.4213	0.4213
9	0.0000	0.0000	0.0000	0.0000	0.0000
10	0.0000	0.0000	0.0000	0.0006	0.0006

Table 6.2: Losses having inspecting half of component 3 for selected components

same as those given in Table 6.1. We again learn most about the component we have inspected. However, in this case we also learn about components 1 and 4. By inspecting these components we improve our ability to decide how to treat them, and therefore reduce our expected losses as a result of our decisions noticeably. This results in overall reduction in the expected loss in utility of 1.5132 for this component. The design with the highest utility value at this first step is the design which recommends inspecting half of component 3, so our new preferred design is:

$$\begin{aligned} d_1 &= \{d_0, b_1^*\} \\ &= \{\emptyset, b_{3,1}\} = \{b_{3,1}\}. \end{aligned} \tag{6.6}$$

We now carry out the same process to identify the second inspection block to be included in the design. We use the same procedure of evaluating the expected utility for every design including one additional block to d_1 , this time we identify inspecting half of component 16 (see Figure 6.3) as the inspection block which allows

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) P(F y) \leq \rho] \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1	0.4979	0.0231	0.7468	0.4625	1.2093
2	0.0000	0.0000	0.0000	0.0002	0.0002
3	0.3014	0.0124	0.4521	0.2487	0.7008
4	0.4429	0.0206	0.6643	0.4119	1.0762
5	0.0000	0.0000	0.0000	0.0003	0.0003
6	1.0000	0.0000	1.5000	0.0000	1.5000
7	1.0000	0.0000	1.5000	0.0000	1.5000
8	0.0000	0.0211	0.0000	0.4213	0.4213
9	0.0000	0.0000	0.0000	0.0000	0.0000
10	0.0000	0.0000	0.0000	0.0006	0.0006

Table 6.3: Losses having inspecting half of component 3 for selected components

us to achieve the greatest reduction in expected utility, so we set:

$$d_2 = \{b_1^*, b_2^*\} = \{b_{3,1}, b_{16,1}\}.$$

This process is then repeated until we reach the point at which adding another point causes an increase in expected loss in utility. Figure 6.4 shows the general trend of the expected loss value as the number of steps increases. The solid blue line illustrates the behaviour of the expected loss and the broken red line shows how the criterion would behave if we did not adjust for inspection cost. We can see both drop initially and then begin to level off, the red will continue to decrease as we include more components. However, Figure 6.5 (a close up of a section of Figure 6.4) illustrates that the expected loss has already starts to increase as of step 19. The slight increase we see here indicates that the reduction in loss we can achieve by learning more about the system is now less than the loss incurred through inspecting the extra block.

The components included in the proposed design $d = d_{19}$ are shown in Figure

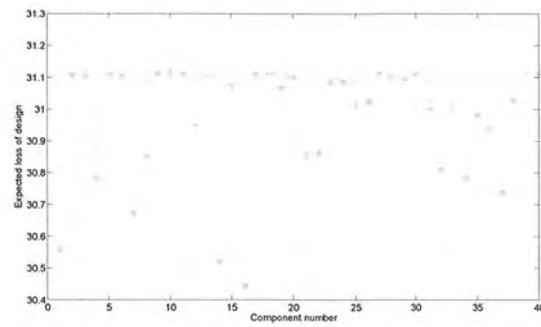


Figure 6.3: Expected loss values for two half components inspection

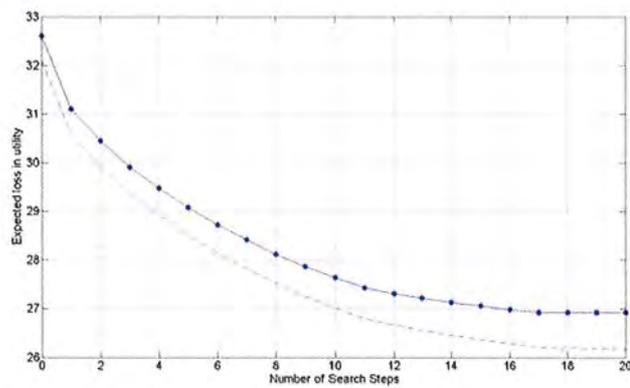


Figure 6.4: Expected utility values for best designs using stepwise add

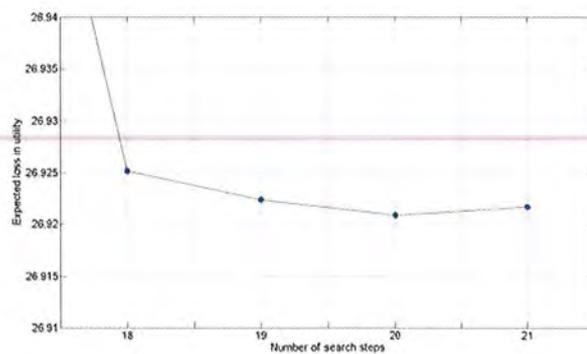


Figure 6.5: Detail from Figure 6.4 showing turning point at step 13

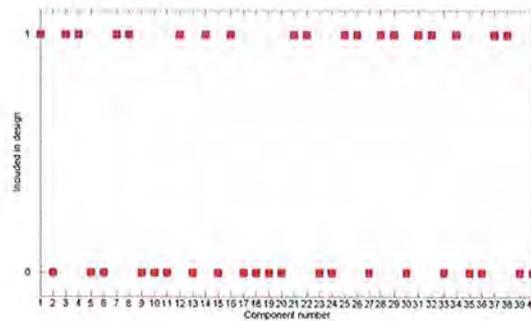


Figure 6.6: Proposed design using stepwise add routine

6.6, components included in the design take value one and those points not included in the design take value 0. Tables 6.4 and 6.5 show the comparison between the variances and failure probabilities of the no inspection case and the proposed design. The layout is the same as Table 6.3, with the index associated with the component number taking values 0 if a component is not inspected, 1 if it is half inspected and 2 if it is fully inspected. We can see that no components are fully inspected, indicating that, for this problem, under the current parameterisation, there is no benefit from full inspection over partial inspection. That is, the further reduction in uncertainty offered by full inspection is not sufficiently large to overcome the cost of performing the inspection.

The components included in the proposed design for stepwise add are generally those in which our initial uncertainty - measured as the variance for that component - is large. This can be seen in Figure 6.7. The Figure plots the initial variance for each component, adjusted for all data to $t = 12$ (the red line), the reduction which can be achieved through total inspection (green line), and the reduction achieved by using the design recommended by the stepwise search procedure (the blue line). The components have been reordered to show the variances in decreasing order of size, this new order is shown on the x -axis. The design selection in this case has targeted the components about which we are most uncertain, which is a reasonable policy. In the cases where the selection has preferred to inspect components with smaller variances, this is because the maintenance decision in the overlooked component cannot be influenced by inspection (because the component is either certain to fail

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) P(F y) \leq \rho] \times P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1 (1)	0.2925	0.0113	0.4388	0.2252	0.6640
2 (0)	0.0000	0.0000	0.0000	0.0002	0.0002
3 (1)	0.2985	0.0123	0.4478	0.2452	0.6930
4 (1)	0.3306	0.0135	0.4959	0.2697	0.7655
5 (0)	0.0000	0.0000	0.0000	0.0003	0.0003
6 (0)	1.0000	0.0000	1.5000	0.0000	1.5000
7 (1)	0.5432	0.0132	0.8148	0.2633	1.0782
8 (1)	0.0560	0.0092	0.1119	0.1850	0.2969
9 (0)	0.0000	0.0000	0.0000	0.0000	0.0000
10 (0)	0.0000	0.0000	0.0000	0.0006	0.0006
11 (0)	0.0000	0.0001	0.0000	0.0020	0.0020
12 (1)	0.8339	0.0058	1.2508	0.1161	1.3669
13 (0)	0.0000	0.0001	0.0000	0.0020	0.0020
14 (1)	0.5691	0.0111	0.8536	0.2217	1.0753
15 (0)	0.0000	0.0000	0.0000	0.0000	0.0000
16 (1)	0.2472	0.0130	0.3708	0.2607	0.6315
17 (0)	0.0000	0.0000	0.0000	0.0002	0.0002
18 (0)	0.0000	0.0000	0.0000	0.0001	0.0001
19 (0)	1.0000	0.0000	2.0000	0.0000	2.0000
20 (0)	0.0000	0.0013	0.0000	0.0258	0.0258

Table 6.4: Expected decision losses by component using the proposed design

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) P(F y) \leq \rho] \times P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
21 (1)	0.8004	0.0083	1.6008	0.1650	1.7658
22 (1)	0.8070	0.0082	1.6139	0.1633	1.7772
23 (0)	1.0000	0.0000	0.6000	0.0000	0.6000
24 (0)	1.0000	0.0000	0.6000	0.0000	0.6000
25 (1)	0.9092	0.0042	1.8183	0.0836	1.9020
26 (1)	0.9153	0.0040	1.8306	0.0805	1.9112
27 (0)	1.0000	0.0000	2.0000	0.0000	2.0000
28 (1)	0.0036	0.0004	0.0073	0.0081	0.0153
29 (1)	0.0041	0.0005	0.0081	0.0094	0.0175
30 (0)	1.0000	0.0000	0.6000	0.0000	0.6000
31 (1)	0.8002	0.0027	0.4801	0.0536	0.5338
32 (1)	0.3893	0.0047	0.2336	0.0930	0.3266
33 (0)	1.0000	0.0000	0.6000	0.0000	0.6000
34 (1)	0.3242	0.0049	0.1945	0.0975	0.2921
35 (0)	1.0000	0.0000	0.6000	0.0000	0.6000
36 (0)	1.0000	0.0000	0.6000	0.0000	0.6000
37 (1)	0.2339	0.0043	0.1403	0.0851	0.2254
38 (1)	0.7965	0.0018	0.4779	0.0361	0.5140
39 (0)	1.0000	0.0000	0.6000	0.0000	0.6000
40 (0)	1.0000	0.0000	0.6000	0.0000	0.6000

Table 6.5: Expected decision losses by component using the proposed design

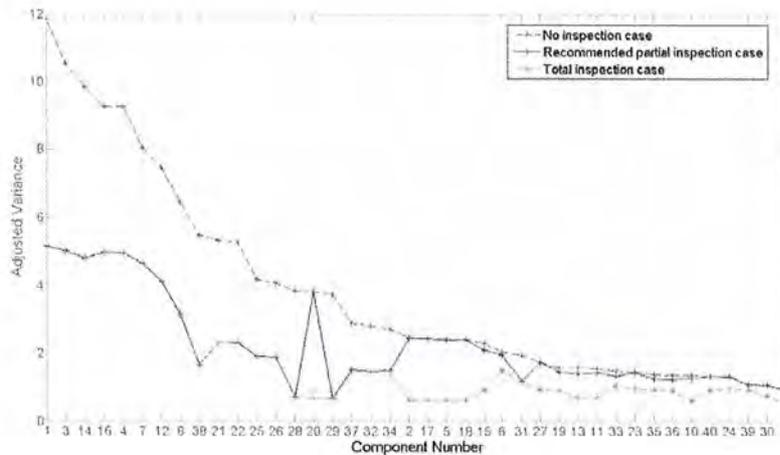


Figure 6.7: Variance Reduction achieved by proposed inspection plan

or certain not to fail). Inspecting according to this design gives an expected loss from decisions (not including inspection costs) of $E(L(d)) = 26.9209$. The value of this design is lower than the value associated with not inspecting (32.1003), and, as we showed in section 5.6.2, the best we can achieve in terms of improved decision making is inspecting fully, which gives us an expected loss from our decision of 26.1073 at a cost of 1.5 to make the inspection. Using d from the stepwise add procedure we can reduce losses from our decision to 26.1834, at a cost of just 0.7375.

6.3.2 Selection using the stepwise delete algorithm

To apply the stepwise delete algorithm we follow essentially the same process we used when applying the stepwise add algorithm. We use the simulation data to provide estimates of the expectations, variances and covariances we require and we use repeated evaluation of the expected loss in utility criterion to establish the best design at every stage. The starting point for our stepwise delete process will be the saturated design, in which every component of Circuit A is inspected fully. We again restrict to the case of half or full component inspection as our only methods of inspecting.

The aim of stepwise delete is to see which inspection block we can remove from

our design with least effect on our ability to predict. We know from section 5.6.2 that saturated design (i.e. $d = B$) has a utility of:

$$\begin{aligned} U(B) &= -(L_R I_1 + L_F I_2 + L_{SU} + n_d L_{IC}) \\ &= -(26.1073 + .5 + 80 \times .0125) = -27.6073. \end{aligned}$$

We want to identify inspection blocks such that the increase in loss due to decision making (separate from that of inspecting) that must occur as we reduce the scope of the design is less than the cost of inspecting that block. The utilities for removing one block from each component individually are negligible at the first step, so, in this case, any action we choose will be an improvement. The range of values covered by the expected losses is very small, but component 3 is deemed to be the least influential (as the design without half of component 3 has the highest utility), so we remove inspection of one half of this component from the design set. The new design, d_1 , instructs us to inspect everywhere fully except component 3 (where we will only inspect half the component).

We now repeat the process using d_1 as the starting design. Again we evaluate the utility for every design which can be found by removing a single inspection block from d_1 . The results from this step have a very similar pattern to the first step, with every value below the best value from step 1 and all utilities close to each other. In this case the design with the highest utility is the one which omits half of component 5, so we remove this block from the design and set d_2 to be the design which inspects everything but half of components 3 and 5.

This process is then repeated until the utility value stops increasing. For this example, 62 steps have to be taken before we see an increase in utility. Figure 6.8 shows the behaviour of the minimum expected loss (i.e. the loss which corresponds to the selected design) at each of the 62 steps. The broken red line shows the behaviour of the loss from our decisions. This quantity can be thought of as describing our performance with respect to information gained or lost, without taking into account the cost of inspection. The amount of information it is possible to learn remains approximately constant through the first 50 steps, causing the expected loss to decrease smoothly at a gradient approximately equivalent to the cost of inspecting

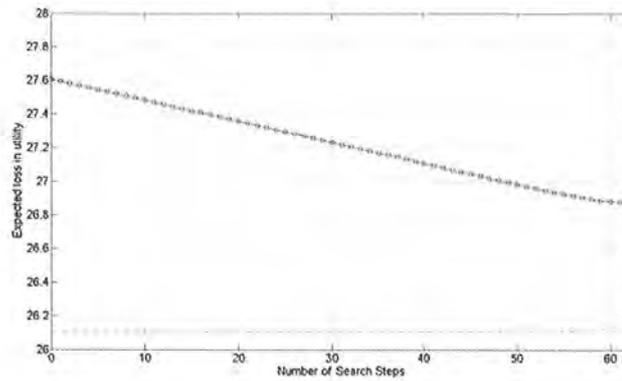


Figure 6.8: Expected utilities for search steps

a single block. When the amount of information lost through not inspecting starts to increase, the expected loss curve begins to level off, and eventually increase, as highlighted in Figure 6.9.

For this example, the recommended design using the stepwise delete algorithm coincides with the design recommended by the stepwise add algorithm. In general this is unlikely to be the case, due to the reasons outlined in sections 6.2.1 and 6.2.2. The reduction in expected loss which can be achieved for this system via inspection at $t = 13$ is based on learning about the distribution of the posterior failure probabilities of the components. When the expected component failure probability is large, reducing uncertainty will push the expected component failure probability closer to 1, but small expected component failure probabilities will move closer to 0 as uncertainty is reduced. This has the effect of increasing our confidence in our decision making, causing us to make the ‘correct’ decision (i.e. replacing failing components and not replacing safe components) more often. Choosing the correct maintenance policy more frequently, based on the observation data, allows us to reduce the expected loss for a component. Figure 6.10 shows the reduction in loss our design achieves for each component, relative to non inspection.

The large number of steps (relative to the maximum number which could have been taken) required to reach the preferred design using the stepwise delete procedure is an indication of the importance of choosing sensible starting points for searches. Poor selection of initial designs can result in large amounts of unneces-

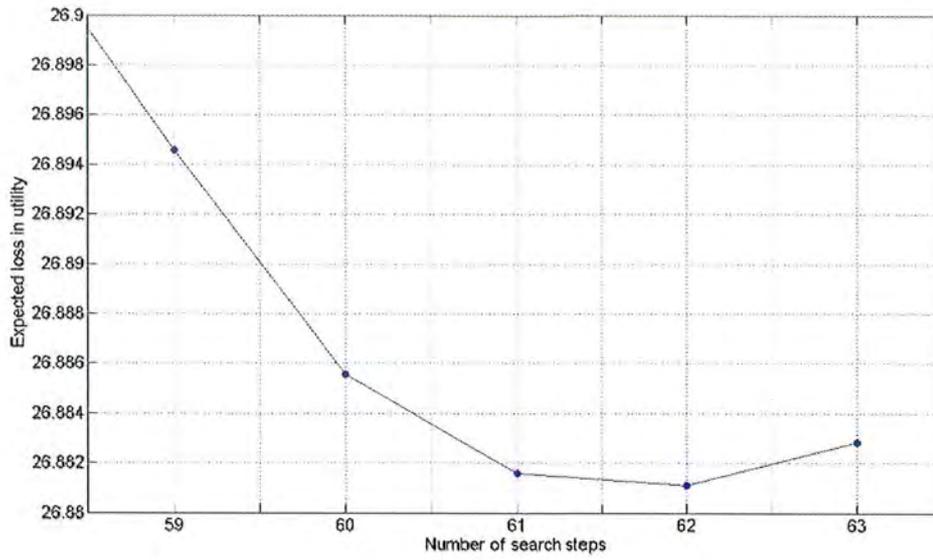


Figure 6.9: Detail of Figure 6.8 showing increase in expected utility value

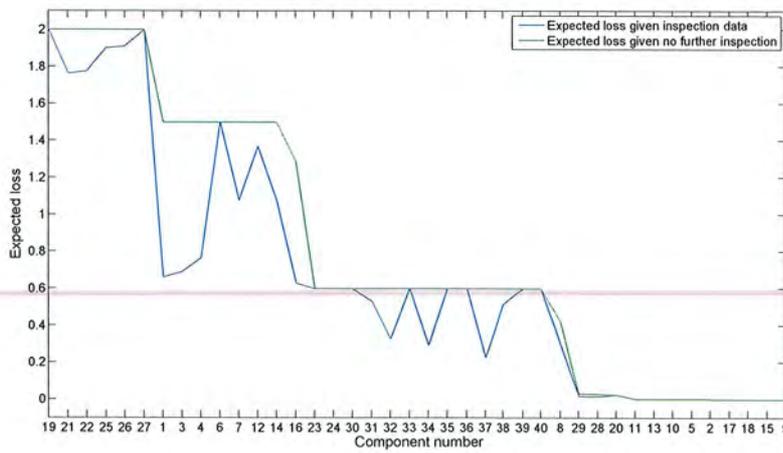


Figure 6.10: Expected loss values for each component (1)

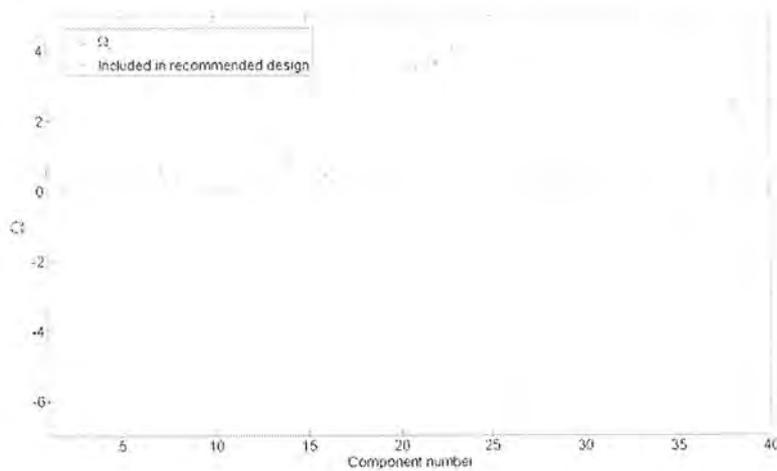


Figure 6.11: Omega values for all components

sary computation being required, or inferior design selections, and consequently bad decisions, being made.

6.3.3 Selecting initial designs for Circuit A

We now illustrate how to apply the Ω criterion outlined in section 6.2.4 to the Circuit A system, and use the resulting design as a starting point for stepwise search routines.

To calculate the values of Ω for our system, we need to evaluate the 1-dimensional variance updates for each component in our system. We use the simulated observations based on whole component observation, but the choice here does not matter, as long as the same form of observations is used for all updates. The values used to calculate Ω are given in the Appendix, where we used the values from Table 6.1 as part of the calculation. Figure 6.11 shows the resulting Ω scores for each component.

Our recommended design from both the stepwise add and stepwise delete procedures consisted of half inspecting 19 components, in no situation was the benefit from fully inspecting a component deemed worthwhile. We can check the Ω score of these components on Figure 6.11 (they are the components marked with an ad-

ditional dot), we see that many of the components are amongst the highest scoring components. Note that components 25 and 26 have the lowest Ω score of all components, but are included in the recommended design by both the stepwise add and stepwise delete method. This exemplifies the need to consider the absolute value of this quantity. The Ω criterion emphasises the relevance of looking at a component by dividing the potential reduction in variance by the difference between the a component's critical wall thickness and its expectation. Therefore those components which are expected to be close to failure at $t = 23$, but not to have failed return large positive scores. However, if a component has an expected minimum at $t = 23$ which is slightly less than the critical wall thickness it receives a large negative score. Consequently, it is worth considering the two extremes of the Ω scale, and this can be done through looking at the largest absolute values.

6.4 Analysis of design performance

To assess how well our design performs it is necessary to establish how often it allows us to make the correct maintenance decision. That is, on the basis of the data available to us, how often do we replace components which go on to fail and how many times are we right not to replace a component. We use a simulation study, based on our previous modelling assumptions, to estimate our typical prediction accuracy using the design proposed by both the stepwise add and stepwise delete.

In section 6.3 we showed that for the Circuit A system the design preferred by both the search methods coincided. The recommended design was to inspect half of each of the following components, 1, 3, 4, 7, 8, 12, 14, 16, 21, 22, 25, 26, 28, 29, 31, 32, 34, 37 and 38 . This design had an expected loss of 26.9209 (to 4 dp), of which 0.7375 can be attributed to inspection costs. We will compare the performance of this design to that of the two baseline cases - no inspection and total inspection. When comparing performance we must consider which decisions a design instructs us to make. The particular choice of inspection design and consequently maintenance decisions depends on the particular data set observed. The expected

loss provides a guide to designs which typically perform well, but does not guarantee good performance for all data sets.

Tables 6.6 and 6.7 show the probabilities of choosing to replace each component under the designs of no inspection, total inspection, and the recommended partial design.

This probability of choosing to replace changes as we inspect, this is a result of increased certainty in component behaviour. The fact that this quantity is random, except in the no inspection case, means we can only establish the typical performance properties of a design. To further assess the quality of the design we have identified, we carry out a simulation study.

We generate 10000 realisations of our system, Circuit A, from time $t = 12$ to $t = 23$. The simulation method used is the same as that employed previously. Our initial conditions for the system slope, $\underline{\alpha}_{12}$ and system level \underline{x}_{12} are taken to be draws from MVN populations with distributions:

$$\underline{\alpha}_{12} \sim N[E_y(\underline{\alpha}_{12}), \text{var}_y(\underline{\alpha}_{12})] \quad (6.7)$$

$$\underline{x}_{12} \sim N[E_y(\underline{x}_{12}), \text{var}_y(\underline{x}_{12})] \quad (6.8)$$

where y represents all the observed data up to $t = 12$. For every realisation, we use the simulated data from time 13 to update our beliefs about the wall loss behaviour by inspecting according to each of our three considered designs. On the basis of these updates, we then make our decisions about the need for component replacement. We then compare how often our designs allowed us to make the correct decision for every component within the system, by recording how many times each component has failed by $t = 23$, and we use this to assess whether or not our decisions were good. Our test procedure can be summarised as:

1. Generate a new realisation of the system for times $t = 12$ to $t = 23$.
2. Sample simulated data from time $t = 13$ according to the recommended partial design, d_p and the saturated design, d_s .
3. Update beliefs using simulated sample data.

Component Number	Prob. of replacement No inspection	Prob. of replacement Partial inspection	Prob. of replacement Total Inspection
1	1	0.2925	0.2919
2	0	0.0000	0.0000
3	1	0.2985	0.2969
4	1	0.3306	0.3293
5	0	0.0000	0.0000
6	1	1.0000	1.0000
7	1	0.5432	0.5413
8	0	0.0560	0.0561
9	0	0.0000	0.0000
10	0	0.0000	0.0000
11	0	0.0000	0.0002
12	1	0.8339	0.8323
13	0	0.0000	0.0002
14	1	0.5691	0.5670
15	0	0.0000	0.0000
16	0	0.2472	0.2463
17	0	0.0000	0.0000
18	0	0.0000	0.0000
19	1	1.0000	0.9998
20	0	0.0000	0.0032

Table 6.6: Number of component failures in simulation experiment

Component Number	Prob. of replacement No inspection	Prob. of replacement Partial inspection	Prob. of replacement Total Inspection
21	1	0.8004	0.7993
22	1	0.8070	0.8058
23	1	1.0000	0.9956
24	1	1.0000	0.9974
25	1	0.9092	0.9081
26	1	0.9153	0.9141
27	1	1.0000	0.9997
28	0	0.0036	0.0036
29	0	0.0041	0.0041
30	1	1.0000	1.0000
31	1	0.8002	0.7917
32	1	0.3893	0.3815
33	1	1.0000	0.9851
34	1	0.3242	0.3137
35	1	1.0000	0.9998
36	1	1.0000	0.9998
37	1	0.2339	0.2281
38	1	0.7965	0.7941
39	1	1.0000	1.0000
40	1	1.0000	0.9999

Table 6.7: Number of component failures in simulation experiment

	Component would have failed	Component would not have failed	Total
Replacement recommended	101908 (25.48%)	158092 (39.52%)	260000 (65%)
Replacement not recommended	523 (0.13%)	139477 (34.87%)	140000 (35%)
Total	102431 (25.61%)	297569 (74.39%)	400000 (100%)

Table 6.8: Performance of the no inspection design

4. Record maintenance decisions (replace/do nothing) for the system components.
5. Record simulated system condition at $t = 23$.
6. Compare maintenance decisions and outcomes.

The output we will be interested in is which decisions we make, and how often those are the correct decisions.

The simulation results are summarised in Tables 6.8 - 6.10. We summarise the results by recording how many times we recommend replacement for components which would have failed, how many replacements are recommended for components which would not have failed, how many components we do not replace which would then fail and how many unreplaced components would not have failed.

We can see that the results for the full design and the recommended partial design give similar results. There is a noticeable decrease in the number of unnecessary replacements carried out under the designs which involve further inspection. This is offset by an increase in the number of component failures not prevented by replacement. There are still a large number of unnecessary replacements ($\approx 1/4$ of all components), but this is due to the uncertainty we cannot resolve, even under full inspection. The results indicate that scheduling a second inspection, or considering inspecting closer to $t = 23$ may be beneficial, as this would allow us to learn about the system. However, the proposed method does not allow quantitative compari-

	Component would have failed	Component would not have failed	Total
Replacement recommended	99517 (24.88%)	99997 (25.00%)	199514 (49.88%)
Replacement not recommended	2914 (0.73%)	197572 (49.39%)	200486 (50.12%)
Total	102431 (25.61%)	297569 (74.39%)	400000 (100%)

Table 6.9: Performance of the recommended partial inspection design

	Component would have failed	Component would not have failed	Total
Replacement recommended	99511 (24.88%)	99626 (24.91%)	199137 (49.79%)
Replacement not recommended	2920 (0.73%)	197943 (49.49%)	200863 (50.21%)
Total	102431 (25.61%)	297569 (74.39%)	400000 (100%)

Table 6.10: Performance of the saturated inspection design

	No inspection	Partial inspection	Full inspection
Correct decisions	241385 (60.35%)	297089 (74.27%)	297454 (74.36%)
Wrong decisions	158615 (29.65%)	102911 (25.73%)	102546 (25.64%)

Table 6.11: Effect of inspection design on decision making ability

son between designs at different time points, so a meaningful comparison cannot be made without further work.

Tables 6.8 - 6.10 provide us with all the necessary information for calculating how often our designs allow us to make the correct decision. Our simulation was conducted for 10000 realisations of a 40 component system. This gives us a total of 400000 decisions to compare. The decisions to replace a failing component and not to replace a component which would not have failed are deemed to be 'correct'. Using the data from Tables 6.8 - 6.10, we find the number of 'correct' decisions to be as shown in Table 6.11

Table 6.11 shows a considerable empirical improvement in our decision making as a result of further inspection. We are getting 3 in 4 decisions correct as oppose to 3 in 5 under no inspection. However, we are also observing more failures, so how do the inspection policies compare financially? We consider the average cost of 'wrong' decisions for a realisation of the decision making process.

For the no inspection cost we have 158092 unnecessary replacements at a cumulative cost of 175929, and 523 unexpected failures at a cumulative cost of 10460. The average cost of wrong decisions for a system realisation based on not inspecting (based on our simulation) is:

$$\frac{175929 + 10460}{10000} = 18.6389.$$

The total inspection design has just 99626 unnecessary replacements at a much lower cumulative cost of 114571, but many more unexpected failures, contributing a cost of 58400 - much more than seen for the no inspection case. This is a consequence of allowing more components to not be replaced depending on the data observed.

The average cost of wrong decisions per system realisation given full inspection is:

$$\frac{114571 + 58400}{10000} = 17.2971.$$

The typical cost is lower, showing our improved decision making ability offers a financial benefit. However, when we consider the cost of total inspection, we see that the combined cost of wrong decisions and making an inspection is $17.2971 + 1.5 = 18.7971$, which is greater than the cost of the wrong decisions made when not inspecting, so for this situation we would prefer not to inspect rather than inspect fully.

Our recommended partial inspection design compares favourably with the full inspection design in terms of information gain. We have a similar number of unnecessary replacements (99997), with a cumulative cost of 115128, and a similar number of unexpected failures (2914) with a cost of 58280. The average cost of wrong decisions per system realisation given partial inspection is:

$$\frac{115128 + 58280}{10000} = 17.3408.$$

This is slightly higher than the cost associated with full inspection. However, the reduced cost of implementing a partial inspection scheme means that the overall average cost of wrong decisions and inspection for this design is: $17.3408 + 0.7375 = 18.0783$, lower than the cost of total inspection and of no inspection. This shows that the improved decision making ability that comes from reducing uncertainty, when obtained at a more reasonable cost via implementing a partial inspection design, offers financial benefits over not inspecting.

If we allow the failure cost to vary we see that the performance of the total inspection design relative to the no inspection design and recommended partial designs improves towards the extremes of the failure costs. Table 6.12 tells us that when L_F is small ($L_F = 2$) we see that there is a convergence in empirical design performance, with all three designs producing similar levels of correct decisions. For medium values, such as those explored initially, we saw a great improvement for both total inspection and partial inspection in terms of correct decision making over not inspecting. When we set L_F to be much larger than our replacement cost, we observe complete agreement between the partial and the total design. This is

$L_F = 20$	No inspection	Partial inspection	Full inspection
Correct decisions	241385 (60.35%)	297089 (74.27%)	297454 (74.36%)
Wrong decisions	158615 (39.65%)	102911 (25.73%)	102546 (25.64%)
$L_F = 2$	No inspection	Partial inspection	Full inspection
Correct decisions	308839 (77.21%)	315856 (78.96%)	316407 (79.10%)
Wrong decisions	91161 (22.79%)	84144 (21.04%)	83593 (20.90%)
$L_F = 2 \times 10^6$	No inspection	Partial inspection	Full inspection
Correct decisions	122256 (30.56%)	222250 (55.56%)	222250 (55.56%)
Wrong decisions	277744 (69.44%)	177750 (44.44%)	177750 (55.56%)

Table 6.12: Effect of inspection design on decision making ability

probably because for such a large failure cost, most replacement decisions are unlikely to be influenced by inspection. The partial design will inspect all of the locations where it can have some influence. So the agreement of these numbers is evidence of the redundancy in the total inspection, in this case.

Clearly, all of these outcomes are sensitive to the replacement, failure and inspection costs (as well as to the underlying model). We have attempted to use values which are of a similar order of magnitude to those found in real systems, but different systems with different inputs may not show the same improved decision making ability. However, the method established here will allow inspection planners to ascertain whether or not inspection is beneficial in a structured, rational manner.

6.5 Choice of loss function

Varying the loss function will cause the design selection process to produce different results. In this section we consider how design selection and expected loss change if we allow the loss associated with component failure L_F to vary. L_F is likely to be the most unpredictable element of the loss calculations. It is reasonable to assume that good information about the losses involved in replacing components and making an inspection will be available. These are planned events which occur at regular intervals, with clearly defined primary (e.g. materials) and secondary (e.g. loss of production) losses. Losses for component failure are more difficult to specify, which can often follow from the consequences of failure being difficult to specify. Within large, complex systems, component failure will often have implications for the performance of other components, and possibly also have some external effects (such as damage to the environment, or personnel). As such, combining the different sources of loss to give a reliable estimate of the loss due to failure is not a straightforward task and such estimates can be inaccurate.

We allowed L_F to take 15 values in the range [2, 2,000,000] (these are shown in Table 6.13), all other aspects of the example system remain unchanged. Our aim was to establish if there were any obvious changes in the behaviour of either the

Test	L_F value	Test	L_F value	Test	L_F value
1	2	6	20	11	200
2	4	7	25	12	2000
3	8	8	50	13	20000
4	10	9	75	14	200000
5	16	10	100	15	2000000

Table 6.13: Test values for varying failure cost

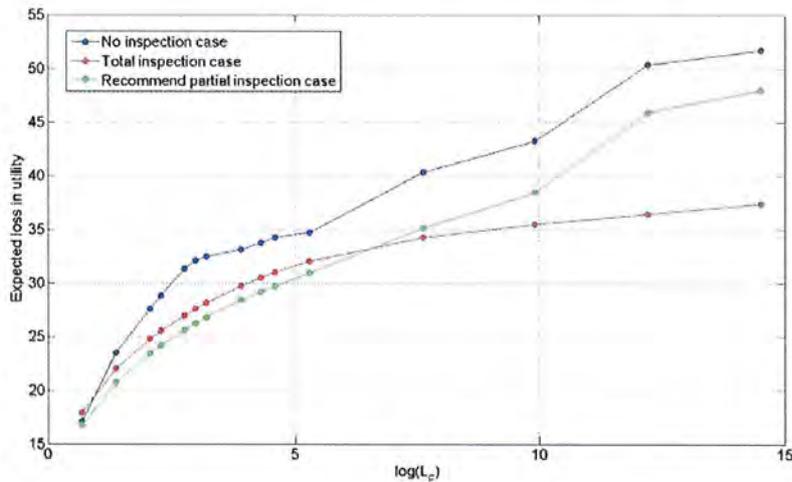


Figure 6.12: Plot showing the behaviour the expected utility with respect to component failure loss

design selection process or the expected utility criterion. One expected change is shown in Figure 6.12, there is an increase in the expected loss for all designs as the loss associated with component failure, L_F , increases.

This can be explained by both the increased loss of component failure, and also the increased number of replacements it becomes necessary to carry out as L_F grows. Our decision rule is determined by the ratio of replacement and failure cost. As failure cost increases, the acceptable probability of failure decreases and so, typically, more components will be replaced.

Figure 6.12 shows the behaviour of the expected loss criterion for the two baseline designs of no inspection and complete inspection, and also the performance of our

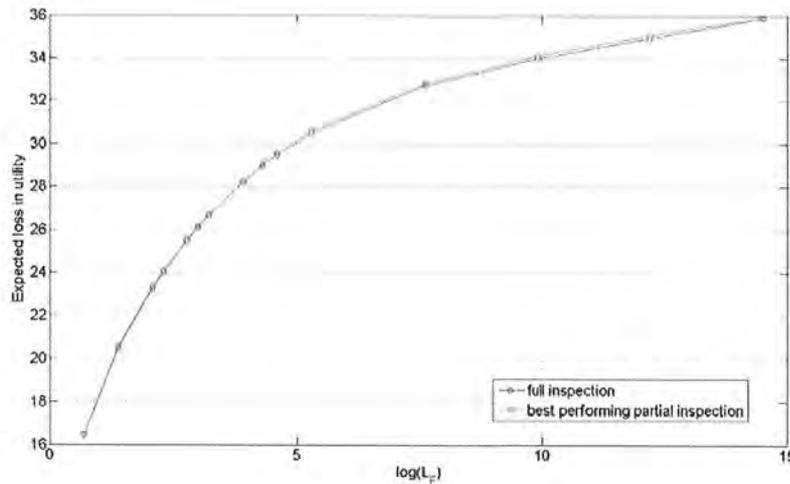


Figure 6.13: Plot showing the behaviour the expected utility of full inspection and best partial inspection with respect to component failure loss

recommended inspection design as L_F is varied. The recommended design, was identified for $L_F = 20$, and so should perform well around this value. The plotted values are adjusted for inspection cost, so they reflect the ‘value’ offered by each inspection design. Figure 6.12 shows us that the benefits of inspection become more pronounced as the consequences of failure become more severe. This can be seen from the red line showing the expected loss associated with total inspection increasing less rapidly than the expected loss incurred by not inspecting (the blue line). We can also see that our recommended inspection design performs well for smaller values of L_F , but is progressively less effective as L_F grows, indicating that the design is no longer relevant to the areas of interest, due to the changing costs.

To assess design behaviour, we consider the performance of the recommended design (selected using the stepwise add algorithm with each of the 15 failure values in turn). Figure 6.13 shows the performance of the best design in comparison to complete inspection. In this case inspection costs have not been taken into account. It can be seen that the green line (the recommended partial inspections) compares favourably with the red line (the full inspection). Indeed, when inspection costs are taken into account, the partial designs will almost always be the preferred choice.

There is no clear pattern as to the effect varying inspection cost has on the size of

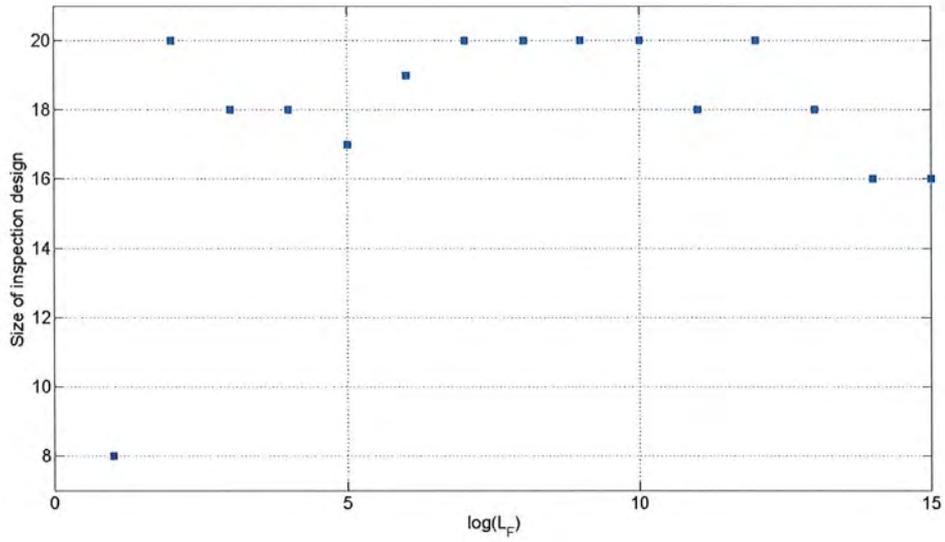


Figure 6.14: Number of components involved in a design chosen via stepwise search for different failure losses

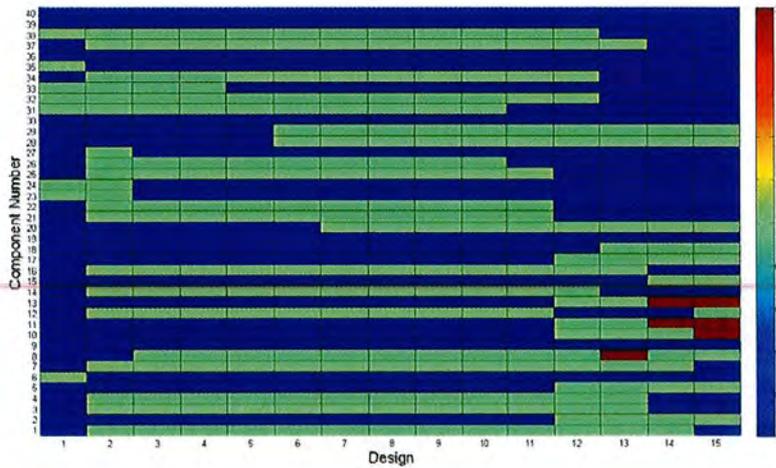


Figure 6.15: Selected components in designs for different failure losses

the recommended design (see Figure 6.14). When $L_F = 2$, the recommended design is considerably smaller. This is probably because at $L_F = 2$, the failure costs and replacement cost will be equal (or nearly equal) for many of components, therefore we are indifferent between replacing now or replacing at failure, and inspection will not change this, so the number of blocks available for inclusion in the design is reduced. For all other costs, the size of the design is within the range 16-20 blocks, indicating that for this system we can expect an optimal design to involve looking at between 20% - 25% of the available inspection blocks. The size of the designs for $L_F = 200000$ and $L_F = 2000000$ is lower than that seen for smaller values. It would make sense that as L_F increases to being very much larger than the replacement cost L_R , the amount of inspection necessary decreases. This is a consequence of our decision rule, in which any component with a probability of failure greater than $\frac{L_R}{L_F}$ is replaced. As $L_F \rightarrow \infty$, $\frac{L_R}{L_F} \rightarrow 0$, so a large number of components could be counted as 'certain' failures without needing to inspect them.

The recommended partial inspection design changes as L_F is increased. The 15 recommended designs are shown in Figure 6.15. Each square of the displayed grid gives inspection instructions for a component given a failure cost. For example, the square in row 12 and column 4, tells us what we should do with component 12 if $L_F = 10$, the 4th failure loss under consideration. In this case, the square is green, which means we should inspect half of the component. Maroon squares mean full inspection and navy means no inspection for that component.

The focus of inspection shifts as the value of L_F increases. Partial inspection designs observe components about which our decisions are most uncertain, with the aim of improving our decision making ability. For example, if we consider two components, A and B . A has a probability of failure $p_A = 0.12$, and B has probability of failure $p_B = 3 \times 10^{-6}$, based on our beliefs updated with available data up to time $t = 12$. When $\frac{L_R}{L_F} \approx 0.1$, B will be deemed to be not worth inspecting, as we are confident that $p_{B|y_d}$ will be less than 0.1. However, p_A is close to the value at which we decide to replace, so we would be interested in improving our knowledge about p_A . Consequently a design for $\frac{L_R}{L_F} \approx 0.1$, would favour inspecting components which help us to learn about p_A over those which tell us about p_B . Furthermore, when

$\frac{L_R}{L_F} \approx 10^{-6}$, component A would almost always be replaced, as we would be confident that $P_{A|y_d}$ will exceed 10^{-6} , but learning about p_B may be useful. Therefore, a design for $\frac{L_R}{L_F} \approx 10^{-6}$, would favour inspecting components which help us to learn about p_B over those which tell us about p_A .

An example of the latter can be seen in the design for $L_F = 2 \times 10^6$ (column 15 in Figure 6.15). This design concentrates almost exclusively on components with small failure probabilities (see Tables 6.4 and 6.5), but also includes half of component 12 - a component which is expected to fail almost 1 in 3 times. Component 12 is included with a view to learning about components 10,11 and 13 through the covariance structure. Similarly, if we look at the components for the design when $L_F = 4$, we see that components 32 and 34 are included in the design, despite having expected failure probabilities of 0.0593 and 0.0440, both much less than the value of $\frac{L_R}{L_F} = 0.15$ for these components. However, they are included in the design to help learn, through the covariance structure, about components 31 and 33, which have initial failure probabilities of 0.1733 and 0.3385 - both closer to the cut-off point of 0.15.

We can conclude that the focus of the recommended inspection design shifts as the value of L_F increases. The expected utility criterion concentrates on identifying inspection blocks that allow us to learn about components that have a probability of failure close to the critical value of $\frac{L_R}{L_F}$ (the value which determines whether a component is replaced or not), both through direct observation of the component and also by proposing the observation of related components.

Chapter 7

Discussion

Within this thesis we have developed an approach for planning inspections for large industrial systems with correlated components. The aim of the work has been to demonstrate that it is possible to establish a tractable Bayesian method for inspection planning for large systems, given a non-trivial system model. We have proposed a structure for planning inspections for large industrial systems which allows for system components to be jointly modelled and for the quantitatively defensible selection of inspection designs given a suitable system model. This improves on most current industrial practice, by explicitly representing dependencies between components, and also offers a practical solution to the computational problems of using Bayesian design ideas to address real world examples.

The standard Bayesian approach to inspection planning problems is to formulate the inspection planning problem as a Bayesian decision problem. In the decision theory context our inspections yield observations, y , which allow us to identify the Bayes decision, $\delta^*(y)$. The Bayes decision is defined to be the best available decision from a set of potential decisions, Δ , having observed y . All potential decisions have an associated reward, or utility, that depends on the state of the system. The best available decision is the decision which maximises this utility. We must take expectations over all data values, y , that we could potentially observe, and therefore determine the expected utility of inspecting. Evaluating this expectation is frequently a computationally intensive calculation, and therefore assessing many

different inspection designs using fully Bayesian techniques is impractical. Bayes linear variance adjustment depends on the location of the inspected components, and not on the specific values of y observed. This reduces the computation required to evaluate the expected utility for a single design, thereby making it feasible to compare a much larger number of designs.

Adopting a Bayes linear approach allows us to specify beliefs about quantities of interest in terms of expectations and variance/covariance structures, without the need to specify a full probability distribution. Specifying meaningful coherent joint distributions for systems with large numbers of correlated components performing a variety of different tasks is very difficult. Specifying covariance matrices to describe associations between components without the need to further constrain joint behaviour is advantageous, particularly when the joint distribution is uncertain. Modelling correlations is an essential part of describing system behaviour, and the use of Bayes linear methods to update variance structures retains tractability whilst allowing us to learn about components we do not inspect, as well as those we do inspect. Evaluating expected utility for models with correlated components using traditional Bayesian methods is extremely computationally intensive, so the implementation of a Bayes linear approach, allowing a more computationally efficient method of inspection design selection for systems represented by non-trivial models, is very useful.

The proposed method offers a structured approach to tackling inspection design problems. The building blocks of the method are:

1. Modelling - developing a model which accurately characterises the system properties of interest.
2. Developing suitable criteria for our inspection goals - it is important to use criteria which measure inspection performance in terms of what we are trying to achieve through inspection. Criteria must be chosen so that they identify the designs which allow us to learn most about the quantity of interest. Therefore successfully identifying the inspection goals is also important.
3. Producing tractable methods for evaluating design criteria over all potential

observations, and many possible designs.

The use of Bayes linear methods allows us to tackle the computational problems usually associated with Bayesian design problems. The application of Bayes linear methods to design problems, with a view to establishing a tractable design methodology is the most important aspect of the thesis. The choice of modelling technique and inspection goals will depend on the particular system being modelled.

We now discuss ways in which the methodology could be extended.

7.1 Modelling

7.1.1 System Modelling

We have considered the use of a simple, but flexible structure for the modelling of corrosion behaviour. We have chosen to model corrosion using a semi-spatial model for component ‘wall thickness’. The model is split into a ‘global term’, which describes the average wall thickness level throughout the component, and a ‘local term’ which describes the deviations from the global term at specific locations within the component.

We have chosen to use a linear trend dynamic linear model (DLM) to describe the behaviour of the global term. A linear trend was adopted following consultation with the experts at Shell. The linear trend DLM allows for the trend to vary in time via the introduction of evolution errors to the trend at each time step. The DLM has a term for each component being modelled. The use of a correlated evolution error structure allows us to correlate component behaviour at the global term level.

The local terms were modelled as independent Gaussian noise. The choice to enforce independence on these terms, both within components and between components, is why the model is referred to as ‘semi-spatial’. There is an acknowledgement of different corrosion behaviour at different locations within a component, but there is an absence of a defined relationship between locations.

The chosen model characterises the experts’ views about corrosion to a sufficient

level of detail to suggest efficient overall inspection schemes for the system. However, the adopted model should not be seen as a definitive corrosion model. There is provision within the model structure to allow for trends, and correlations, within the local term, if these are deemed necessary. The global term is not restricted to a linear trend. Dynamic linear models allow for a large number of different behaviour types to be modelled. The flexibility of dynamic linear models and the range of potential applications are covered by [24]. If it was felt that a particular component needed further structure, it would be possible to introduce an intermediate model term to describe behaviour only occurring in certain regions of the component. In our model the balance between global and local terms is in favour of the global term. That is, the component wall thickness behaviour is controlled by the linear trend DLM, and the local terms are regarded as noise. However, this balance can be adjusted as is felt to be appropriate through the structuring and parameterisation of the model, while still allowing for a tractable design methodology.

7.1.2 Observation Modelling

It is necessary to model observations as a distinct process from system modelling. Firstly, this is a more realistic representation of the observation process, especially in the context of detecting corrosion damage. Secondly, observation procedures are likely to change more frequently than the system evolution will change. Having a separate observation model means that in the event of a different observation technique being used, there is minimal change to the modelling process.

We have considered two different observation techniques, but have modelled them in the same way. Our inspection aim was to detect failing components. A component fails when the wall thickness becomes too thin to remain structurally sound. Consequently a component's minimum wall thickness is a quantity of interest to inspectors. The available data consisted of minima observations, and our inspection aims directed us to consider minima. Therefore our observation model identified the minimum wall thickness values within a component.

A simple model was adopted, in which a proportion of the locations within

a component were observed with error (usually half or all of the locations), and the minimum of these values was returned. That is, the minimum observed wall thickness, y_c for a component c at time t is given by:

$$y_{ct} = \min_l \{u_{lct} + \xi_{lct}\} \quad (7.1)$$

where u_{lct} is the ‘true’ value (from the model), and ξ_{lct} is the observation error.

The ultrasonic and radiographic measurement techniques were used to collect the available data. We adopted the minima observation model for both techniques, using Gaussian errors. The variance of the errors differed between the two models, but otherwise they were the same. A more detailed analysis may highlight systematic differences between the two measurement techniques that necessitates more careful modelling of the two practices, but the adopted model provided a pragmatic approach, which was suitable both for the level of available information about inspection procedures, and for the purposes of establishing a non-trivial model for analysis.

7.1.3 Modelling for the example

The illustrative model provided for the example was intended to provide a correct order of magnitude assessment of the relevant parameters. For any serious implementation of the method, a more careful analysis of the available data would be required. The model as implemented demonstrates the plausibility of the design methodology for a model of this type, but does not fully address the issues of specifying model parameters accurately.

In particular, the assessment of the covariance structures could be improved. We have adopted a standard model for our covariance structures which ensures they will be positive definite. To use the Bayes linear approach, it is useful to ensure our variance/covariance matrices are at least positive semi-definite (see [22]). By using covariance structures where the (i, j) -th elements are of the form:

$$\sigma_{ij} = \sigma_i \sigma_j \exp\{-\tau \|i - j\|^2\} \quad (7.2)$$

where $\|i - j\|$ is a measure of distance between i and j , we can guarantee positive definiteness, but this does not necessarily provide the most realistic representation of the actual covariances between components.

The available data merits more careful consideration, and a more detailed investigation of the data will allow the successful identification and verification of trends in the data which were not detected by the analysis in this thesis. A detailed approach to model parameterisation will provide a better representation of the real system, and therefore inspection plans which are more directly applicable, but it should not affect the efficiency of the design selection procedure.

The illustrative model provides an example on which we can base our design selection procedure. By concentrating on achieving a parameterisation that is of a similar order of magnitude to the effects which we can observe from an analysis of the data, we can demonstrate that the procedure is tractable for models of this type. Changing the model will change the design outputs, but the design procedure will operate consistently given a coherently specified model.

7.2 Updating

Bayesian methods are the natural approach for handling inspection planning problems. In most inspection problems we are seeking to acquire further knowledge about a system for which we already have some understanding, either from expert judgement or historical data. The Bayesian framework provides a clear method of combining of historical data with new information through the updating process.

The traditional Bayesian method requires the specification of complete probability distributions for our priors and likelihood functions. By following the Bayes linear approach we can reduce the extent of specification required, and also remove the need to specify a particular distribution. Updating via Bayes linear methods is conducted using the Bayes linear updating equations (chapter 4). The use of the Bayes linear update allows us to avoid dealing with the potentially intractable integrals sometimes found in full Bayesian analyses, and the difficulties that can arise

from being unable to obtain closed form expressions for our posterior distributions.

However, it is necessary to be able to express beliefs about the mean, variance and covariance structures between all quantities of interest. We learn about these quantities through simulation. To achieve an accurate representation of these beliefs, a large initial simulation covering the entire time frame in question has to be run. This initial simulation could potentially be time consuming. However, we only need to conduct this simulation once, to establish our belief structure. To implement a full Bayes strategy we would have to run a similarly large simulation as part of the evaluation of every design we wish to consider. So although there is a significant simulation task as part of our proposed procedure, this is still a much lower level of computation than would be necessary under a full Bayes approach.

We have adopted a full simulation approach to updating, in which we perform a single large simulation to obtain estimates for all the covariance structures we may need to consider over the entire time interval we are interested in. However, for very large systems it may be beneficial to consider a sequential simulation approach. Under this approach, our single initial simulation would be split into a series of smaller simulations which run between the times of historical inspections up to the present, and a final step which runs from the present to a designated point in the future. Using this procedure we would adjust our simulation parameters after each historical observation and run the simulation with the new parameters to the time of the next historical observation, at which stage we would then update again.

In summary, the sequential procedure, with real historical observations made at times t_i , $i = \{1, \dots, \omega\}$ would be:

1. Simulate S realisations of the system from time $t = 0$ to time t_1 , using initial simulation conditions as described in section 2.8.
2. Estimate $\text{var}(\underline{y}_{dt_i})$, $\text{cov}(\underline{x}_{dt_i}, \underline{y}_{dt_i})$, $\text{cov}(\underline{y}_{dt_i}, \underline{\alpha}_{dt_i})$, $E(\underline{x}_{dt_i})$, $\text{var}(\underline{x}_{dt_i})$, $E(\underline{\alpha}_{dt_i})$ and $\text{var}(\underline{\alpha}_{dt_i})$ from simulation output, for $i = 1$.
3. Calculate $E_{\underline{y}_{dt_i}}(\underline{x}_{t_i})$ and $E_{\underline{y}_{dt_1}}(\underline{\alpha}_{t_i})$, using the Bayes linear updating equations, for $i = 1$.

4. Simulate S realisations of the system from time t_i to time t_{i+1} , using the adjusted expectations as part of the initial simulation conditions, for $i = 1$.
5. Repeat steps 2 - 4 for $i = 2, \dots, \omega$.

Sequential updating may provide updated values which are of greater relevance to the particular system being modelled, by adjusting to simulate around the observed values at more regular intervals. Sequential updating was considered as it is computationally more appealing. For a 40 component system (as modelled in this thesis), the difference in computational load between the two approaches is negligible. However, for very large systems, the sequential method may be a more efficient simulation which concentrates computational resources on simulation realisations that lie closer to the observed data. The sequential method is an approximation to the full simulation, so the effects of implementing a sequential updating technique on the simulation accuracy need to be investigated further before definite conclusions can be drawn.

Our current updating strategy allows us to learn about the means and variances of the model outputs, such as the minimum component wall thickness, or the corrosion rate. However, we would like to be able to update the underlying model parameters, such as the DLM evolution deviation covariance matrix. This quantity drives the correlation structure in the model, and we would like to be able to learn about it from the historical data. Currently, our model retains the same covariance structure throughout its life, which is not ideal. To be able to update the covariance structure we need to record information about the 4th order moments of the system. Recording higher order moments would also allow us to consider skewness in the data. This could be handled within the Bayes linear context by fitting our adjusted moments to appropriate non-central distributions. Bayes linear updating of covariance matrices is discussed in [69], and including the ability to learn about the structure of the underlying model in this way is an important requirement of the modelling strategy. As with the necessity for an improved data analysis before considering a serious implementation of the design selection procedure, this is an additional modelling consideration, which should not affect the way in which inspection designs are selected.

7.3 Design Selection

We have shown by illustration that our chosen inspection design selection method leads to the identification of inspection designs that perform well in terms of improving our decision making ability. There are several potential developments which would be beneficial.

We would like to be able to consider a wider class of decision problems. We have currently considered the simplest case of using the inspection data to determine whether or not a component should be replaced. For most real inspection planning applications, there will more than two potential maintenance decisions. Expanding to consider a large number of maintenance decisions should be reasonably straightforward in theory, as this resolves to specifying utilities for each maintenance decision, and constructing decision rules such that the Bayes decision can be chosen. However, such considerations will further complicate the calculation of the design criterion, so it may be necessary to restrict to a small set of potential decisions to maintain efficient computation.

The evaluation of the expected loss criterion depends on the assumption of Normality at two key points. Firstly, we assume the probability of component failure (defined as a component's minimum wall thickness being less than a specified critical value) can be described by a Normal distribution with expectation and variance equal to the expectation and variance of the component's minimum wall thickness at the final time of the interval, adjusted for the unseen inspection data. Secondly, we further assume that the adjusted expectation - which is itself a random quantity in inspection problems, as it depends on the unseen data - has a Normal distribution. It would be useful to investigate the appropriateness of these assumptions, and to what extent they could be relaxed. Both assumptions could be replaced by any form of distribution solely parameterised by location and scale parameters, with similar analysis to that which we have described.

Calculating the expected loss criterion requires the evaluation of two integrals. The first of these simply corresponds to a cumulative Normal distribution function. However, the second is less tractable and requires evaluation by numerical methods.

This difficulty can be resolved by the use of specialist software packages for evaluating integrals, or by the use of simulation routines. Given different distributional assumptions for the probability of component failure and the adjusted expectation, the resulting integrals may be more difficult to evaluate. Consequently, the need to evaluate this integral numerically for each inspection design could hamper the computational tractability of the method. In the case of the integrals being intractable, it would be possible to tabulate results for generic values of the quantities on which the integrals depend. Calculating these tables would be a time consuming operation, but again it would only be necessary to perform these calculations once.

The current method does not explicitly allow for the comparison of inspection designs made at different times. To be able to compare inspections made at different times, we must factor in the costs incurred by delaying the inspection. Under the current decision model, the only costs of delaying inspection would be the losses due to component failures prior to inspection. By incorporating this additional loss into the expected loss criterion it would be possible to compare all potential inspection designs within the interval of interest, and thus determine an improved inspection policy which maximises utility taking all potential time points into consideration.

Implementing a selection procedure across multiple time points is not straightforward. It would probably be necessary to identify the best available design at each time point individually, and then carefully compare the performance of these best available designs in terms of improved decision making ability and expected cost of delaying inspection.

The quality of search routines used in the design selection process could also come in for scrutiny. We have used stepwise search routines, which are efficient for the scale of example considered. However, it may be more sensible for larger systems to investigate alternative methods. In particular, consideration could be given to the intelligent choice of starting point for design searches.

Many of the issues outlined in the discussion can be tackled within the framework laid out in this thesis. There are still many interesting related research problems, and potential extensions to the method, but we have shown, in principle, the viability of our approach to inspection planning for large systems.

Bibliography

- [1] Pipeline inspection gauges. <http://en.wikipedia.org/>, January 2007.
- [2] Federal Highways Administration. Corrosion costs and preventive strategies in the united states. Report Number: FHWA-RD-01-156, 2002.
- [3] J. M. Angulo and M. C. Bueso. Random perturbation methods applied to multivariate spatial sampling design. *Environmetrics*, 12:631–646, 2001.
- [4] M. J. Baker and B. Descamps. Reliability-based methods in the inspection planning of fixed offshore steel structures. *Journal of Constructional Steel Research*, 52:117–131, 1999.
- [5] Milena Banjevic. *Optimal Network Designs in Spatial Statistics*. PhD thesis, Stanford University, 2004.
- [6] B. Bhattacharya and B. Ellingwood. A new CDM-based approach to structural deterioration. *International Journal of Solids and Structures*, 36:1757–1779, 1999.
- [7] S. Biffi and M. Halling. Investigating the defect detection effectiveness and cost benefit of nominal inspection teams. *IEEE Transactions on Software Engineering*, 29:385–397, 2003.
- [8] G. E. P. Box. *Statistics for experimenters: an introduction to design, data analysis, and model building*. Wiley, Chichester, 1978.
- [9] NDT Resource Center. NDT method summary. <http://www.ndt-ed.org/>.

- [10] K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10:273–304, 1995.
- [11] T. J. Chen and A. C. Thornton. Quantitative selection of inspection plans. In *Proceedings of the 1999 ASME Design Engineering Technical Conference*, pages DEC/DTM–8759, 1999.
- [12] D. R. Cox and N. Reid. *The theory of the design of experiments*. Chapman & Hall, Boca Raton, 2000.
- [13] P. S. Craig, M. Goldstein, J. C. Rougier, and A. H. Seheult. Bayesian forecasting for complex systems using computer simulators. *Journal of the American Statistical Association*, 96:717–729, 2001.
- [14] P. J. Diggle and S. Lophaven. Bayesian geostatistical design. Working paper 42, Johns Hoskins University, Dept of Biostatistics, 2004.
- [15] M. L. Eaton, A. Giovagnoli, and P. Sebastiani. A predictive approach to the bayesian design problem with application to normal regression models. *Biometrika*, 83:111–125, 1996.
- [16] S. Englund, J. D. Sorensen, M. H. Faber, and A. Bloch. Approximations in inspection planning. In *8th ASCE Speciality Conference on Probabilistic Mechanics and Structural Reliability*, pages PMC200–198, 2000.
- [17] A. C. Estes and D. M. Frangopol. Minimum expected cost-oriented optimal maintenance planning for deteriorating structures: application to concrete bridge decks. *Reliability Engineering and System Safety*, 73:281–289, 2001.
-
- [18] M. Farrow. Practical building of subjective covariance structures for large complicated systems. *Statistician*, 52:553–573, 2003.
- [19] M. Farrow and M. Goldstein. Bayes linear methods for grouped multivariate repeated measurement studies with application to crossover trials. *Biometrika*, 80:39–59, 1993.

- [20] D. M. Frangopol, M.J. Kallen, and J. M. van Noortwijk. Probabilistic models for life-cycle performance of deteriorating structures: review and future directions. *Prog. Struct. Engng Mater.*, 6:197–212, 2004.
- [21] J. Gasemyr and B. Natvig. Bayesian inference based on partial monitoring of components with applications to preventive system maintenance. *Naval Research Logistics*, 48:551–577, 2001.
- [22] M. Goldstein. Bayes linear analysis. In S. Kotz et al, editor, *Encyclopedia of Statistical Sciences*, volume 3. Wiley, 1999.
- [23] J. Goyet. Integrated approach for RBI offshore installations. Technical briefing, Bureau Veritas, 2001.
- [24] J. Harrison and M. West. *Bayesian Forecasting and Dynamic Models*. Springer-Verlag, New York, 1989.
- [25] Health and Safety Executive. *Control of Major Accident Hazards Regulations 1999*. Queen’s Printer of Acts of Parliament, 1999.
- [26] Health and Safety Executive. *Reducing risk, protecting people*. Her Majesty’s Stationery Office, 2001.
- [27] UK Health and Safety Executive. *The Offshore Installations (Safety Case) Regulations 2005*. Queen’s Printer of Acts of Parliament, 2005.
- [28] S. G. Hellevik, I. Langen, and J. D. Sorensen. Cost optimal reliability based inspection and replacement planning of piping subjected to CO₂ corrosion. *International Journal of Pressure Vessels and Piping*, 76:527–538, 1999.
- [29] S. Hernandez and S. Nestic. Use of artificial neural network for predicting crude oil effect on CO₂ corrosion of carbon steels. Technical report, Ohio University, 2003.
- [30] P. M. Johnson. Reengineering inspection. In *Communications of the ACM*, volume 41, pages 49–52, 1998.

- [31] Daniel Lewandowski. Gas pipelines corrosion data analysis and related topics. Master's thesis, Delft University of Technology, 2002.
- [32] R. J. Lewis and D. A. Berry. Group sequential clinical trials: A classical evaluation of bayesian decision-theoretic designs. *Journal of the American Statistical Association*, 89:1528–1534, 1994.
- [33] C. Q. Li. Computation of the failure probability of deteriorating structural systems. *Computers and Structures*, 56:1073–1076, 1995.
- [34] D. V. Lindley. *Bayesian Statistics - A review*. SIAM, Philadelphia, 1972.
- [35] J. Little, M. Goldstein, and P. Jonathan. Efficient bayesian sampling inspection for industrial processes based on transformed spatio-temporal data. *Statistical Modelling*, 2004.
- [36] J. Little, M. Goldstein, P. Jonathan, and K. den Heijer. Spatio-temporal modelling of corrosion in an industrial furnace. *Applied Stochastic Modelling in Business and Industry*, 2004.
- [37] John Little. *Bayesian spatio-temporal modelling for inspection and prediction of complex problems in the petrochemical industry*. PhD thesis, University of Durham, 2003.
- [38] H. Lofsten. Management of industrial maintenance - economic evaluation of maintenance policies. *International Journal of Operations & Production Management*, 19(7):716–737, 1999.
- [39] T. Loredó. Bayesian adaptive exploration. In press: *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, 2003.
- [40] Lisa McConnell. Alignment and calibration of s-nii methodology with s-rbi methodology and safety cases. Master's thesis, University of Aberdeen, 2003.
- [41] R. E. Melchers. Corrosion uncertainty modelling for steel structures. *Journal of Constructional Steel Research*, 52:3–19, 1999.

- [42] B. Nimmo and G. Hinds. The beginners guide to corrosion. National Physical Library, 2003.
- [43] M. Nordsveen, S Nestic, R. Nyborg, and A. Stangeland. A mechanistic model for carbon dioxide corrosion of mild steel in the presence of protective iron carbonate films - part 1: Theory and verification. *Corrosion*, 59:443–465, 2003.
- [44] T. Onoufriou and D. M. Frangopol. Reliability-based inspection optimization of complex structures: a brief retrospective. *Computers and Structures*, 80:1133–1144, 2002.
- [45] R. Patel and J. Rudlin. Analysis of corrosion/erosion incidents in offshore process plants and implications for non-destructive testing. *Insight*, 42, 1997.
- [46] D. F. Percy. Bayesian enhanced strategic decision making for reliability. *European Journal of Operations Research*, 139:133–145, 2002.
- [47] H. Procaccia, R. Cordier, and S. Muller. Application of bayesian statistical decision theory for a maintenance optimization problem. *Reliability Engineering and System Safety*, 55:142–149, 1997.
- [48] S. Qin and W. Cui. Effect of corrosion models on the time-dependent reliability of steel plated elements. *Marine Structures*, 16:15–34, 2003.
- [49] B. Rajani and Y. Kleiner. Comprehensive review of structural deterioration of water mains: physically based models. *Urban Water*, 3:151–164, 2001.
- [50] D. V. Ramsamooj and T. A. Shugar. Modelling of corrosion fatigue in metals in an aggressive environment. *International Journal of Fatigue*, 23:S301–S309, 2001.
- [51] D. T. Rusk, K. Y. Lin, D. D. Swartz, and G. K. Ridgeway. Bayesian updating of damage size probabilities for aircraft structural life-cycle management. Technical briefing, Naval Air System Command, 2001.
- [52] P. Sebastiani and H. P. Wynn. Maximum entropy sampling and optimal bayesian experimental design. *J. Royal Statistical Society B*, 62:145–157, 2000.

- [53] A. Sergaki and K. Kalaitzakis. A fuzzy knowledge based method for maintenance planning in a power system. *Reliability Engineering and System Safety*, 77:19–30, 2002.
- [54] S. C. Shaw and M. Goldstein. Simplifying complex designs: Bayes linear experimental design for grouped multivariate exchangeable systems. In A. P. Dawid J. M. Bernardo, J. O. Berger and A. F. M. Smith, editors, *Bayesian Statistics 6*. Oxford University Press, 1998.
- [55] J. D. Sorensen, M. H. Faber, and A. Bloch. Simplified approach to inspection planning. In *8th ASCE Speciality Conference on Probabilistic Mechanics and Structural Reliability*, pages PMC200–188, 2000.
- [56] D. Straub. *Generic Approaches to Risk Based Inspection Planning for Steel Structures*. PhD thesis, Swiss Federal Institute of Technology, Zurich, 2004.
- [57] D. Straub and M. H. Faber. On the relation between inspection quantity and quality. In *3rd European-American Workshop on NDE Reliability*, page <http://www.ndt.net/article/v09n07/straub/straub.htm>, 2002.
- [58] D. Straub and M. H. Faber. Computational aspects of risk-based inspection planning. *Computer-Aided Civil and Infrastructure Engineering*, 21:179–192, 2006.
- [59] K. Tanaka and G. J. Klir. A design condition for incorporating human judgment into monitoring systems. *Reliability Engineering and System Safety*, 65:251–258, 1999.
- [60] K. Tang, R. Plante, and H. Moskowitz. Multiattribute bayesian acceptance sampling plans under nondestructive inspection. *Management Science*, 32:739–750, 1986.
- [61] Kallen/Delft University. Inspection and maintenance decisions based on imperfect inspections.

- [62] J. M. van Noortwijk, R. M. Cooke, and M. Kok. A bayesian failure model based on isotropic deterioration. *European Journal Operational Research*, 82:270–282, 1995.
- [63] J. M. van Noortwijk and H. E. Klatter. Optimal inspection designs for the block mats of the eastern-scheldt barrier. *Reliability Engineering and System Safety*, 65:203–211, 1999.
- [64] A. Volker, Frits Dijkstra, Sieger Terpstra, Jan Heerings, and Max Lont. Modelling of NDE reliability. Technical report, TNO, Rontgen Technische Dienst, Shell Global Solutions International, 2004.
- [65] J. C. Walton, G. Cragolino, and S. K. Kalandros. A numerical model of crevice corrosion for passive and active metals. *Corrosion Science*, 38:1–18, 1996.
- [66] G. G. Wang and S. Shan. Review of metamodelling techniques in support of engineering design optimization. *ASME Transactions, Journal of Mechanical Design*, page To appear, 2006.
- [67] S. A. Wenk and R. C. McMaster. *Choosing NDT: applications, costs and benefits of non-destructive testing in your quality assurance program*, page 96. Columbus, 1987.
- [68] D. J. Wilkinson. Bayes linear variance adjustment for locally linear DLMs. *Journal of Forecasting*, 16:329–342, 1997.
- [69] D. J. Wilkinson and M. Goldstein. Bayes linear adjustment for variance matrices. In A.F.M. Smith J. O. Berger, A.P. Dawid, editor, *Bayesian Statistics 5*, pages 791–799. OUP, 1996.
- [70] D. E. Williams, C. Westcott, and M. Fleischmann. Stochastic model of pitting corrosion of stainless steels. *Journal of the Electrochemical Society*, 132:1796–1804, 1985.
- [71] J. B. Wintle, B. W. Kenzie, G.J. Amphlett, and S. Smalley. *Best practice for risk based inspection as a part of plant integrity management*. Health and Safety Executive. Her Majesty's Stationery Office, 2001.

-
- [72] S. Yella, M. S. Dougherty, and N. K. Gupta. Artificial intelligence techniques for the automatic interpretation of data from non-destructive testing. *Insight*, 48:10–20, 2006.
- [73] R. Zhang and S. Mahadevan. Model uncertainty and bayesian updating in reliability-based inspection. *Structural Safety*, 22:145–160, 2000.

Appendix A

Tables

Component Number	$P(P(F y) > \rho)$ (I_1)	$E\{P(F y) $ $P(F y) \leq \rho\} \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1.0000	0.2928	0.0113	0.4392	0.2253	0.6644
2.0000	0	0.0000	0	0.0002	0.0002
3.0000	0.4855	0.0240	0.7282	0.4801	1.2083
4.0000	0.7482	0.0163	1.1222	0.3264	1.4486
5.0000	0	0.0000	0	0.0003	0.0003
6.0000	1.0000	0	1.5000	0	1.5000
7.0000	1.0000	0	1.5000	0	1.5000
8.0000	0	0.0211	0	0.4213	0.4213
9.0000	0	0	0	0	0
10.0000	0	0.0000	0	0.0006	0.0006
11.0000	0	0.0001	0	0.0020	0.0020
12.0000	1.0000	0	1.5000	0	1.5000
13.0000	0	0.0001	0	0.0020	0.0020
14.0000	1.0000	0	1.5000	0	1.5000
15.0000	0	0.0000	0	0.0000	0.0000
16.0000	0	0.0643	0	1.2860	1.2860
17.0000	0	0.0000	0	0.0002	0.0002
18.0000	0	0.0000	0	0.0001	0.0001
19.0000	1.0000	0	2.0000	0	2.0000
20.0000	0	0.0013	0	0.0257	0.0257

Table A.1: Losses for the c1 update

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) $ $P(F y) \leq \rho] \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
21.0000	1.0000	0	2.0000	0	2.0000
22.0000	1.0000	0	2.0000	0	2.0000
23.0000	1.0000	0	0.6000	0	0.6000
24.0000	1.0000	0	0.6000	0	0.6000
25.0000	1.0000	0	2.0000	0	2.0000
26.0000	1.0000	0	2.0000	0	2.0000
27.0000	1.0000	0	2.0000	0	2.0000
28.0000	0	0.0015	0	0.0293	0.0293
29.0000	0	0.0016	0	0.0327	0.0327
30.0000	1.0000	0	0.6000	0	0.6000
31.0000	1.0000	0	0.6000	0	0.6000
32.0000	1.0000	0	0.6000	0	0.6000
33.0000	1.0000	0	0.6000	0	0.6000
34.0000	1.0000	0	0.6000	0	0.6000
35.0000	1.0000	0	0.6000	0	0.6000
36.0000	1.0000	0	0.6000	0	0.6000
37.0000	1.0000	0	0.6000	0	0.6000
38.0000	1.0000	0	0.6000	0	0.6000
39.0000	1.0000	0	0.6000	0	0.6000
40.0000	1.0000	0	0.6000	0	0.6000

Table A.2: Losses for the c1 update

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) $ $P(F y) \leq \rho] \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
1.0000	0.4979	0.0231	0.7468	0.4625	1.2093
2.0000	0.0000	0.0000	0.0000	0.0002	0.0002
3.0000	0.3014	0.0124	0.4521	0.2487	0.7008
4.0000	0.4429	0.0206	0.6643	0.4119	1.0762
5.0000	0.0000	0.0000	0.0000	0.0003	0.0003
6.0000	1.0000	0.0000	1.5000	0	1.5000
7.0000	1.0000	0.0000	1.5000	0	1.5000
8.0000	0.0000	0.0211	0.0000	0.4213	0.4213
9.0000	0.0000	0.0000	0.0000	0	0
10.0000	0.0000	0.0000	0.0000	0.0006	0.0006
11.0000	0.0000	0.0001	0.0000	0.0020	0.0020
12.0000	1.0000	0.0000	1.5000	0	1.5000
13.0000	0.0000	0.0001	0.0000	0.0020	0.0020
14.0000	1.0000	0.0000	1.5000	0	1.5000
15.0000	0.0000	0.0000	0.0000	0.0000	0.0000
16.0000	0.0000	0.0643	0.0000	1.2861	1.2861
17.0000	0.0000	0.0000	0.0000	0.0002	0.0002
18.0000	0.0000	0.0000	0.0000	0.0001	0.0001
19.0000	1.0000	0.0000	2.0000	0	2.0000
20.0000	0.0000	0.0013	0.0000	0.0257	0.0257

Table A.3: Losses for the c3 update

Component Number	$P(P(F y) > \rho)$ (I_1)	$E[P(F y) $ $P(F y) \leq \rho] \times$ $P(P(F y) \leq \rho)$ (I_2)	Replacement Costs $L_R \cdot I_1$	Failure Costs $L_F \cdot I_2$	Component Utility $L_R I_1 + L_F I_2$
21.0000	1.0000	0.0000	2.0000	0	2.0000
22.0000	1.0000	0.0000	2.0000	0	2.0000
23.0000	1.0000	0.0000	0.6000	0	0.6000
24.0000	1.0000	0.0000	0.6000	0	0.6000
25.0000	1.0000	0.0000	2.0000	0	2.0000
26.0000	1.0000	0.0000	2.0000	0	2.0000
27.0000	1.0000	0.0000	2.0000	0	2.0000
28.0000	0.0000	0.0015	0.0000	0.0293	0.0293
29.0000	0.0000	0.0016	0.0000	0.0327	0.0327
30.0000	1.0000	0.0000	0.6000	0	0.6000
31.0000	1.0000	0.0000	0.6000	0	0.6000
32.0000	1.0000	0.0000	0.6000	0	0.6000
33.0000	1.0000	0.0000	0.6000	0	0.6000
34.0000	1.0000	0.0000	0.6000	0	0.6000
35.0000	1.0000	0.0000	0.6000	0	0.6000
36.0000	1.0000	0.0000	0.6000	0	0.6000
37.0000	1.0000	0.0000	0.6000	0	0.6000
38.0000	1.0000	0.0000	0.6000	0	0.6000
39.0000	1.0000	0.0000	0.6000	0	0.6000
40.0000	1.0000	0.0000	0.6000	0	0.6000

Table A.4: Losses for the c3 update

Component Number	$\text{var}_{y_h}(m_{c23})$	$\text{var}_{y_h, c_{13}}(m_{c23})$	$E_{y_h}(m_{c23}) -$ W_C	Ω
1	11.8499	5.1172	4.6776	0.5539
2	2.4654	0.6394	6.6283	0.0238
3	10.5392	4.9645	4.4668	0.5261
4	9.2645	4.9271	4.1153	0.5042
5	2.4212	0.6356	6.5509	0.0247
6	2.0453	1.5085	-1.2695	-0.2121
7	8.0491	4.6223	2.9031	0.6353
8	6.4560	3.1338	5.1638	0.3514
9	0.9586	0.5337	9.7274	0.0189
10	1.3541	0.6210	4.6723	0.0591
11	1.5631	0.7220	4.6538	0.0759
12	7.4594	4.1057	1.1529	1.5840
13	1.5771	0.7246	4.6758	0.0882
14	9.8367	4.7685	2.7634	0.8121
15	2.3040	0.9443	7.4621	0.0590
16	9.2711	4.9363	4.6273	0.4479
17	2.4530	0.6357	6.6482	0.0223
18	2.4146	0.6333	6.7314	0.0199
19	1.5804	0.9213	-1.6827	-0.2040
20	3.8255	0.6874	5.8960	0.0130

Table A.5: Values for calculating the Ω criterion

Component Number	$\text{var}_{y_h}(m_{c23})$	$\text{var}_{y_h, c_{13}}(m_{c23})$	$E_{y_h}(m_{c23}) -$ W_C	Ω
21	5.3345	2.3208	0.4959	3.4948
22	5.2698	2.3153	0.4673	3.6721
23	1.4637	0.9775	0.0354	4.0250
24	1.3393	0.9841	0.2010	0.6164
25	4.1741	1.9117	-0.2271	-6.6076
26	4.0869	1.8845	-0.2689	-5.5030
27	1.7323	0.9442	-1.7683	-0.0111
28	3.8319	0.6971	5.8235	0.3035
29	3.7369	0.6843	5.6854	0.3065
30	1.0707	0.7507	-1.0838	-0.0293
31	1.9394	1.1618	1.3110	0.6640
32	2.8071	1.4482	2.6150	0.4397
33	1.4885	1.0579	0.5081	0.7817
34	2.7079	1.4478	2.8080	0.3904
35	1.3857	0.9267	-0.5704	-0.6098
36	1.3635	0.9240	-0.5509	-0.6243
37	2.8960	1.4655	3.1682	0.3710
38	5.4764	1.6344	0.7959	2.4572
39	1.0882	0.9321	-0.8535	-0.0122
40	1.3425	0.9354	-0.5318	-0.0272

Table A.6: Values for calculating the Ω criterion

Appendix B

Plots

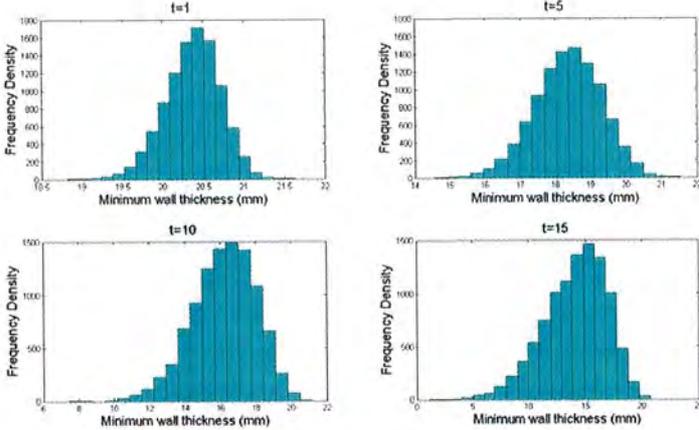


Figure B.1: Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 16 from Circuit A

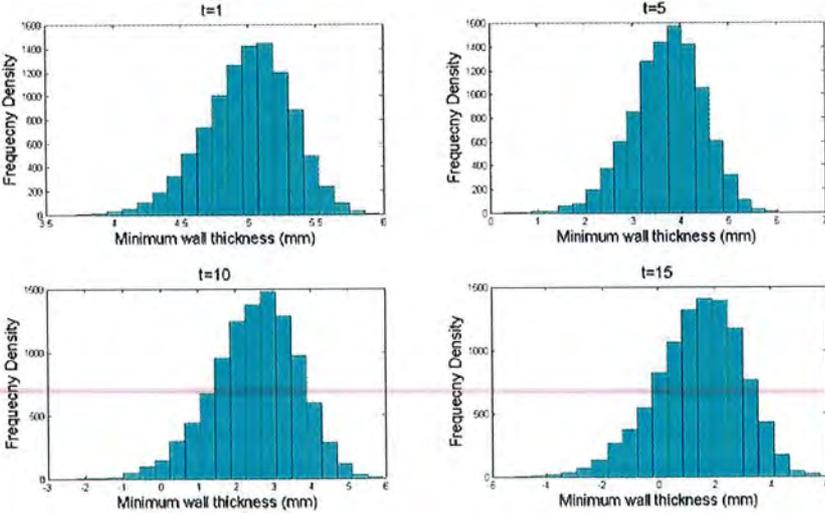


Figure B.2: Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 24 from Circuit A

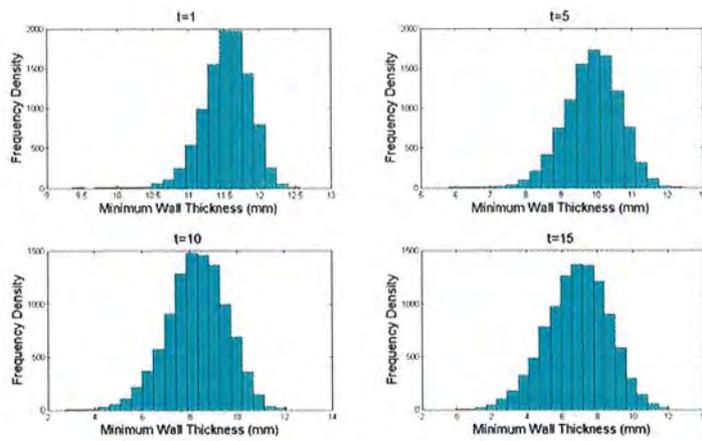


Figure B.3: Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 32 from Circuit A

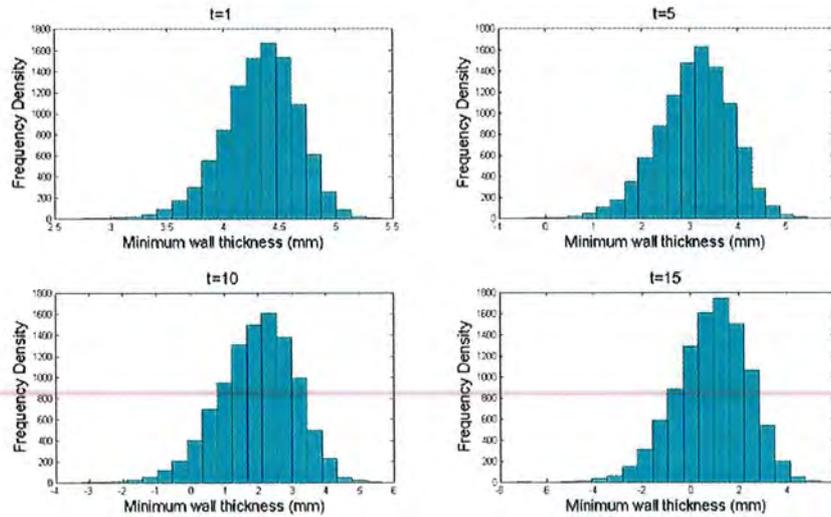


Figure B.4: Histogram of simulated minimum wall thickness values for 4 different time points in the life of component 40 from Circuit A

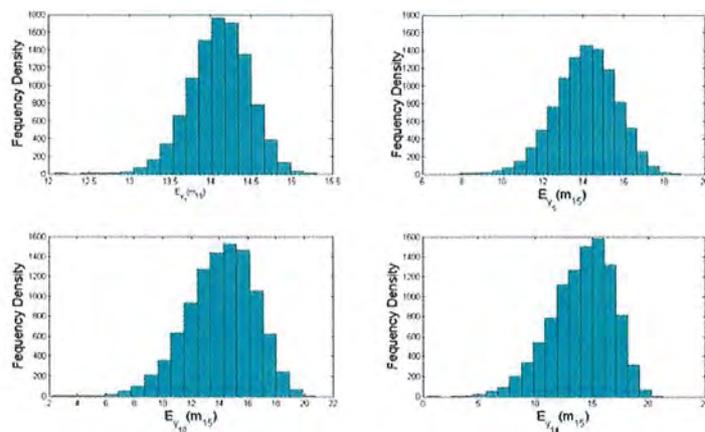


Figure B.5: Histograms of adjusted expectation values for component 16 from Circuit A based on simulated realisations of observations at 4 different time points

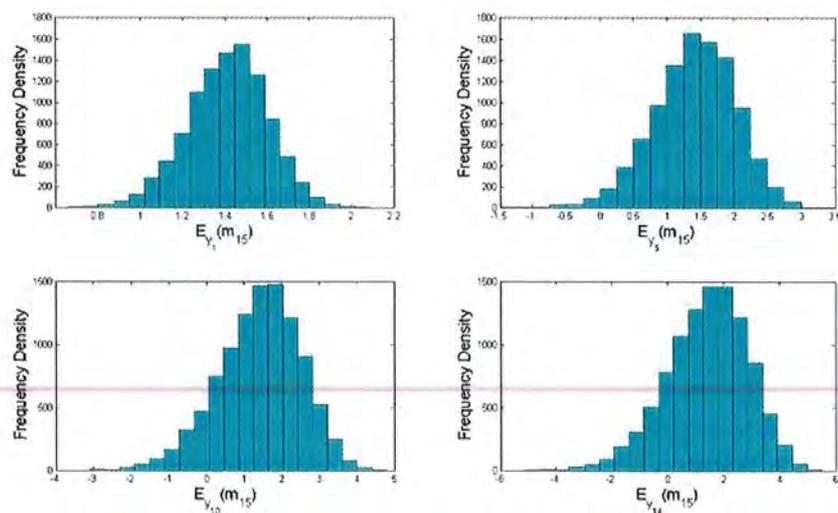


Figure B.6: Histograms of adjusted expectation values for component 24 from Circuit A based on simulated realisations of observations at 4 different time points

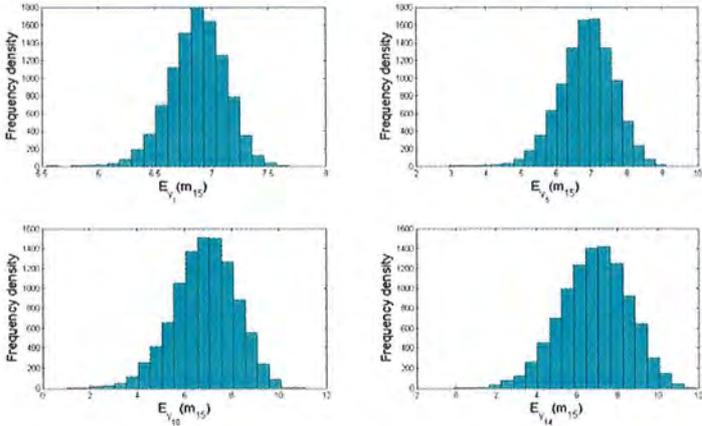


Figure B.7: Histograms of adjusted expectation values for component 32 from Circuit A based on simulated realisations of observations at 4 different time points

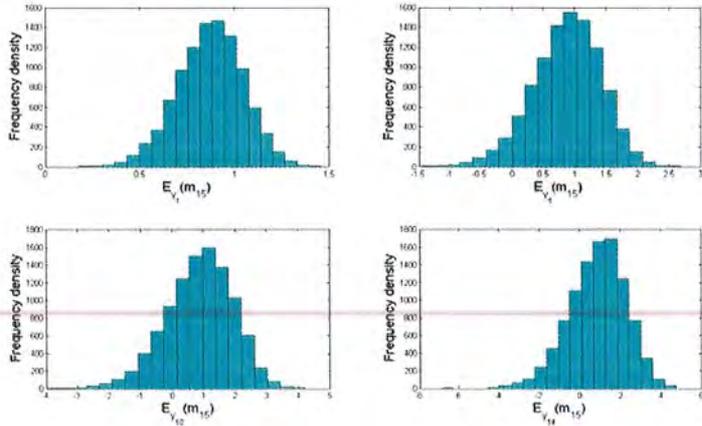


Figure B.8: Histograms of adjusted expectation values for component 40 from Circuit A based on simulated realisations of observations at 4 different time points

Appendix C

Code

```
function [A,X,M,Y1,Y2] = SimJoint(x0,a0,nTim,nRls,nLoc,Sx,Sa,Sr,Sobs)
%simulates the joint distributions between:
% - X (dlm level);
% - Y (observed minima);
% - A (dlm slope);
% - M (true minima).
%by running the forward model

%x0 - nStr vector defining initial wall thickness
%a0 - nStr vector defining initial wall loss rate
%nTim - number of time points to run model (scalar)
%nRls - number of simulation realisations (scalar)
%nLoc - no. of locations within a component (scalar)
%Sx - nStr x nStr correlated standard deviation matrix wall thickness values
%Sa - nStr x nStr correlated standard deviation matrix wall loss rates
%Sr - nStr x nStr diagonal matrix of local variation standard deviations
%Sobs - nStr x nStr observation error standard deviation matrix (diagonal)

nStr = length(x0); %define value of nStr for future calculations
Xo = x0*ones(1,nRls);%Wall thickness values at t=0
Ao = a0*ones(1,nRls);%Wall loss rate at t=0
```

```
%create arrays for storing output
X = NaN*ones(nRls, nStr, nTim+1);
A = NaN*ones(nRls, nStr, nTim+1);
R = zeros(nRls, nStr, nLoc);
M = NaN*ones(nRls, nStr, nTim+1);
Y1 = NaN*ones(nRls, nStr, nTim+1);
Y2 = NaN*ones(nRls, nStr, nTim+1);

%define initial conditions
X(:,:,1) = Xo';
A(:,:,1) = Ao';

%step model forward in time
for iTim = 2:nTim+1
    A(:,:,iTim) = A(:,:,iTim-1) + randn(nRls,nStr)*Sa; %evolve trend term
    A(find(A>0))=0; %stop trend term being positive
    X(:,:,iTim) = X(:,:,iTim-1) + A(:,:,iTim) + randn(nRls,nStr)*Sx; %evolve level
    X(find(X<0))=0; %stop level term being negative

    Rt = zeros(nRls, nStr, nLoc); %
    E = zeros(nRls, nStr, nLoc);

    for iLoc = 1:nLoc
        Rt1(:,:,iLoc) = randn(nRls,nStr)*Sr; %Generates min local deviation term
        E(:,:,iLoc) = randn(nRls,nStr)*Sobs; %Generates observation error term
    end

    R = Rt + R; %cumulative local deviation effect
    Er = R+E; %cumulative local deviation + observation error
```

```

M(:, :, iTim) = X(:, :, iTim) + min(R, [], 3); %true surface minima
Y1(:, :, iTim) = X(:, :, iTim) + min(Er, [], 3); %observed surface minima
Y2(:, :, iTim) = X(:, :, iTim) + min(Er(:, :, 1:5), [], 3); %observe half surface min
end

function [EL, I1, I2] = ExpLoss(Lr, Lf, Wc, mu0, sigD, sig0, RndDat)
%evaluates the expected loss for a design with known adjusted variance

% Lr - vector of replacement costs
% Lf - vector of failure costs
% Wc - vector of critical wall thicknesses
% mu0 - expected wall loss
% sigD - updated standard deviations
% sig0 - prior standard deviations
% RndDat - random sample from N(0,1) - ensures same sampling errors for each run

LrOvrLf = Lr./Lf; %define ratio of costs

Mz = (Wc - mu0)./sigD; %estimate expectation of z

Sz = sqrt((sig0./sigD).^2 - 1); %estimate standard deviation of z

Lmt = norminv(LrOvrLf); %define integration limits

for iC = 1:length(Sz); %loop over components
    if Sz(iC) > 0; %check for consistency of variances
        I1(iC) = normcdf((Mz(iC)-Lmt(iC))./Sz(iC)); %evaluate prob of replacing c
    elseif Mz(iC) > Lmt(iC); %component has failed
        I1(iC) = 1;
    else;
        I1(iC) = 0; %component
    end;
end;

```

```
I2(iC) = EstI2(Mz(iC),Sz(iC),Lmt(iC),RndDat);
%evaluate prob of failure if not replacing c
end;

EL = sum(Lr.*I1' + Lf.*I2'); %return overall expected loss from decisions

function I2=EstI2(Mz,Sz,Lmt,RndDat);
%approximates the I2 integral as part of the expected loss calculation

%Mz, Sz, Lmt - parameter of the integral
%RndDat - N(0,1) sample, read in to ensure consistency of estimation error

X=RndDat;
nX=size(RndDat,1);

if Sz>0;
tLmt=(Lmt-Mz)/Sz;
tInt=normcdf(X.*Sz+Mz);
Kep=X<=tLmt;
I2=sum(tInt(Kep))/nX;
else;
if Mz>Lmt;
I2=0;
else;
I2=normcdf(Mz);
end;
end;
```

Appendix D

Data

Corrosion Circuit Code	Comp. No.	Comp. Type	PMC	Product	Material Type	Pipe Diam.	Piping Schedule	Nominal Wall Thickness	Inspection Date	Inspection Method	Inspection Reading
A	1	Bes	2	P	2	6	PS3	23	10/11/2002	UT	22
A	2	Wellhead	2	P	2	6	PS3	18.24	12/12/2000	UT	18
A	3	Bes	2	P	2	6	PS3	23	10/11/2002	UT	22
A	4	Bes	2	P	2	6	PS3	23	17/05/2003	UT	20.5
A	4	Bes	2	P	2	6	PS3	23	12/12/2000	UT	22
A	5	Wellhead	2	P	2	6	PS3	18.24	12/12/2000	UT	18
A	6	Bes	2	P	2	1	PS3	6.35	12/12/2000	UT	5.35
A	6	Bes	2	P	2	1	PS3	6.35	18/05/2003	UT	5.35
A	7	Bes	2	P	2	6	PS3	18.24	17/05/2003	UT	17
A	7	Bes	2	P	2	6	PS3	18.24	12/12/2000	UT	17.5
A	8	TEQ	2	P	2	6	PS3	18.24	17/05/2003	UT	20
A	9	Cap	2	P	2	6	PS3	18.24	12/06/1998	UT	18.24
A	9	Cap	2	P	2	6	PS3	23	17/05/2003	UT	19.5
A	10	Str	1	P	1	0.75	PS1	3.91	18/05/2003	UT	3.91
A	11	Str	2	P	2	4	PS3	13.49	10/11/2002	UT	13
A	12	Str	1	P	1	4	PS2	11.13	10/11/2002	UT	13
A	13	Bes	1	P	1	4	PS2	11.13	10/11/2002	UT	12.5

Table D.1: Data for the circuit A example

Corrosion Circuit Code	Comp. No.	Comp. Type	PMC	Product	Material Type	Pipe Diam.	Piping Schedule	Nominal Wall Thickness	Inspection Date	Inspection Method	Inspection Reading
A	14	Str	1	P	1	4	PS2	11.13	10/11/2002	UT	13
A	15	Bes	1	P	1	4	PS2	11.13	10/11/2002	UT	15
A	16	Str	1	P	1	4	PS2	11.13	10/11/2002	UT	22
A	17	Str	1	P	1	0.75	PS1	3.91	18/05/2003	UT	3.91
A	18	Bes	2	P	2	6	PS3	18.24	17/05/2003	UT	20
A	18	Bes	2	P	2	6	PS3	18.24	12/12/2000	UT	20
A	19	Wellhead	2	P	2	6	PS3	18.24	12/12/2000	UT	18
A	20	Wellhead	2	P	2	6	PS3	18.24	12/12/2000	UT	18
A	21	Wellhead	2	P	2	6	PS3	18.24	12/12/2000	UT	18.24
A	22	Wellhead	2	P	2	6	PS3	18.24	12/12/2000	UT	18.24
A	23	Branch	2	P	2	0.5	PS3	4.78	12/06/1998	UT	4.78
A	24	Wellhead	2	P	2	6	PS3	18.24	12/06/1998	UT	18.24
A	25	Branch	2	P	2	2	PS3	8.74	12/06/1998	UT	8.74
A	26	Branch	2	P	2	2	PS3	8.74	12/06/1998	UT	8.74
A	27	Weld	2	P	2	1	PS3	6.35	18/05/2003	UT	5.35
A	27	Weld	2	P	2	1	PS3	6.35	12/12/2000	UT	5.35
A	27	Weld	2	P	2	1	PS3	6.35	12/06/1998	UT	6.35

Table D.2: Data for the circuit A example

Corrosion Circuit Code	Comp. No.	Comp. Type	PMC	Product	Material Type	Pipe Diam.	Piping Schedule	Nominal Wall Thickness	Inspection Date	Inspection Method	Inspection Reading
A	28	Weld	2	P	2	1	PS3	6.35	18/05/2003	UT	5.35
A	28	Weld	2	P	2	1	PS3	6.35	12/12/2000	UT	5.35
A	28	Weld	2	P	2	1	PS3	6.35	12/06/1998	UT	6.35
A	28	Weld	2	P	2	1	PS3	6.35	31/08/2003	UT	6.35
A	29	Branch	2	P	2	1.5	PS3	7.1	12/06/1998	UT	7.1
A	30	Branch	2	P	2	1.5	PS3	7.1	12/06/1998	UT	7.1
A	31	Branch	2	P	2	0.5	PS3	4.78	12/06/1998	UT	4.78
A	32	Wellhead	2	P	2	6	PS3	18.24	12/06/1998	UT	18.24
A	33	Wellhead	2	P	2	6	PS3	18.24	12/06/1998	UT	18.24
A	34	Weld	2	P	2	0.5	PS3	4.78	12/06/1998	UT	4.78
A	35	Weld	2	P	2	0.5	PS3	4.78	12/12/2000	UT	4.78
A	35	Weld	2	P	2	0.5	PS3	4.78	18/05/2003	UT	4.78
A	36	Weld	2	P	2	2	PS3	8.74	12/06/1998	UT	8.74
A	36	Weld	2	P	2	2	PS3	8.74	18/05/2003	UT	8.74
A	36	Weld	2	P	2	2	PS3	8.74	12/12/2000	UT	8.74
A	37	Weld	1	P	1	4	PS2	13.5	17/05/2003	UT	12
A	38	Weld	2	P	2	1	PS3	6.35	31/08/2003	UT	6.35

Table D.3: Data for the circuit A example

Corrosion Circuit Code	Comp. No.	Comp. Type	PMC	Product	Material Type	Pipe Diam.	Piping Schedule	Nominal Wall Thickness	Inspection Date	Inspection Method	Inspection Reading
A	38	Weld	2	P	2	1	PS3	6.35	12/06/1998	UT	7.1
A	38	Weld	2	P	2	1	PS3	6.35	12/12/2000	UT	7.1
A	39	Wellhead	1	P	1	1.5	PS1	7.14	12/12/2000	UT	7.14
A	40	Weld	1	P	1	4	PS2	13.5	12/12/2000	UT	13
A	40	Weld	1	P	1	4	PS2	13.5	17/05/2003	UT	13
A	41	Weld	1	P	1	0.75	PS1	5.56	18/05/2003	UT	5.56
A	42	Weld	1	P	1	0.75	PS1	5.56	18/05/2003	UT	5.56
A	42	Weld	1	P	1	0.75	PS1	5.56	12/12/2000	UT	5.56
A	43	Weld	1	P	1	4	PS2	13.5	17/05/2003	UT	14
A	44	Wellhead	1	P	1	1.5	PS1	7.14	12/12/2000	UT	7.14
A	45	Weld	1	P	1	0.75	PS1	5.56	31/08/2003	UT	2.6
A	45	Weld	1	P	1	0.75	PS1	5.56	08/04/2004	UT	2.8
A	45	Weld	1	P	1	0.75	PS1	5.56	18/05/2003	UT	3.56
A	46	Weld	1	P	1	0.75	PS1	5.56	12/12/2000	UT	5
A	46	Weld	1	P	1	0.75	PS1	5.56	18/05/2003	UT	5.56

Table D.4: Data for the circuit A example