



The
University
Of
Sheffield.

**Uncertainty Propagation in
Nonlinear Systems**

by

Daley Chetwynd

A thesis submitted for the degree of Doctor of
Philosophy in the Faculty of Engineering

Department of Mechanical Engineering

December 2005

Acknowledgements

Thanks to Graeme and Keith for all their help. Thanks to Davide Vige and David Storer at Centro Ricerche Fiat.

Abstract

This thesis examines the effects of uncertainty on a variety of different engineering systems. Uncertainty can be best described as a lack of knowledge for a particular system, and can come from a variety of different sources. Within this thesis the possibilistic branch of uncertainty quantification is used.

A combination of simulated and real-life engineering systems are studied, covering some of the most popular types of computational models. An outline of various background topics is presented first, as these topics are all subsequently used within the thesis. The most important of these is the transformation method, a possibilistic uncertainty approach derived from fuzzy arithmetic.

Most of the work here examines uncertain systems by implementing Ben-Haim's information gap theory. Uncertainty is deliberately introduced into the parameters of the various computational models to use the concept of "opportunity". The basic rationale is that if some degree of tolerance can be accepted on a model prediction of a system, it is possible to obtain a lower value of prediction error than with a standard crisp-valued model.

For the use of interval-valued computational models there is generally a trade-off to be made between minimising the prediction error of the model and minimising the range of predicted outputs, to reduce the tolerance on the solution. The studied models all use a "degree of uncertainty" parameter that allows any user to select the suitable trade-off level for their particular application.

The thesis then concludes with a real-life engineering study, undertaken as a nine month placement on a European Union project entitled MADUSE. The work was done at Centro Ricerche Fiat, and examined the dynamic effects of uncertainties related to automotive spot welds. This study used both finite element modelling and experimental modal testing of manufactured specimens.

Contents

1. Introduction and Background.....	1
1.1 Overview.....	1
1.2 Background.....	5
1.3 Information-gap Decision Theory.....	5
1.4 Computational Information.....	7
2. Interval Techniques.....	9
2.1 Interval Arithmetic.....	9
2.2 Affine Arithmetic.....	12
2.3 Quadratic Arithmetic.....	16
3. Possibilistic Sampling Methods.....	17
3.1 Vertex Two-Level Full Factorial Method.....	17
3.2 Vertex Three-Level Full Factorial Method.....	19
3.3 Fuzzy Arithmetic.....	20
3.4 Reduced Transformation Method.....	25
3.5 General Transformation Method.....	26
3.6 Modified Reduced Transformation Method.....	28
4. Evolutionary Computation.....	31
4.1 Genetic Algorithms.....	31
4.2 Differential Evolution.....	34
5. Autoregressive Models.....	39
5.1 ARMAX and ARX Models.....	39
5.2 One-step-ahead and Long Range Prediction.....	40
5.3 NARX Models.....	41
5.4 The Pseudoinverse Approach.....	43
5.5 Training an Autoregressive Model.....	44

6. Artificial Neural Networks.....	45
6.1 Introduction.....	45
6.2 Multilayer Perceptron Networks.....	47
6.3 Radial Basis Function Networks.....	49
7. Uncertain ARX Modelling of a Linear System.....	51
7.1 Initial Crisp ARX Model.....	51
7.2 Proportional Radii Expansion.....	55
7.3 Differential Evolution Radii Training.....	56
7.4 Results and Discussion.....	59
8. On Replacing a Nonlinear System with an Uncertain Linear Model.....	66
8.1 Introduction.....	66
8.2 Crisp Models.....	67
8.3 Crisp ARX Input Scaling.....	68
8.4 Uncertain ARX Modelling.....	68
8.4.1 Variable Input RMS.....	69
8.4.2 Variable Output Noise.....	70
8.5 Results and Discussion.....	70
9. An Application of Interval-Valued Neural Networks to a Regression Problem.....	75
9.1 Background Work.....	75
9.2 Introduction.....	77
9.3 Interval Weight Radius Training.....	78
9.4 Results and Discussion.....	80
10. Uncertainty Propagation through Radial Basis Function Networks.....	86
10.1 Introduction.....	86
10.2 Interval Network Training.....	87
10.3 Results and Discussion.....	89

11. Interval-Valued Regression Models of Bifurcating Nonlinear Systems.....	95
11.1 Introduction.....	95
11.2 The Crisp System.....	95
11.3 An Interval-Valued Model.....	99
11.4 Results and Discussion.....	103
12. The Dynamic Effects of Uncertainties within Automotive Spot Welds.....	111
12.1 Introduction.....	111
12.2 Resistance Spot Welding.....	112
12.3 The Spot-Welded Component.....	114
12.4 Spot Weld Uncertainties.....	117
12.5 The Finite Element Model.....	121
12.6 Modal Analysis of the Spot-Welded Component.....	123
12.6.1 The Modal Assurance Criterion.....	123
12.6.2 The Frequency Domain Assurance Criterion.....	126
12.6.3 Modal Testing.....	127
12.7 Results.....	131
12.7.1 All Welds Present.....	135
12.7.1.1 Finite Element Modelling.....	137
12.7.1.2 Experimental Testing of Manufactured Specimens.....	139
12.7.2 Any Single Weld Missing.....	142
12.7.3 Welds 2 and 8 Missing.....	147
12.7.3.1 Finite Element Modelling.....	151
12.7.3.2 Experimental Testing of Manufactured Specimens.....	153
12.7.4 Any Two Welds Missing.....	156
12.8 Discussion and Conclusions.....	160
13. Conclusions and Future Work.....	170
13.1 Conclusions.....	170
13.2 Future Work.....	185
References.....	188

Appendix A.....193
Appendix B.....194

1. Introduction

1.1 Overview

This thesis looks at a variety of different ways of modelling uncertainty within engineering systems. A number of different systems are examined, some of them simulated by computational routines, and others given by real-life industrial processes. The characteristics of the different systems also vary here; between linear and nonlinear; also between the time and frequency domains.

The first question that arises is why is uncertainty important to consider? How does it justify the study that has been undertaken here?

Uncertainty was best defined by Oberkampf (2001). It can be separated into two distinct types – aleatory uncertainty (often referred to as “variability”) and epistemic uncertainty. Aleatory uncertainty is described as the inherent variation of a particular model property, such as the differences in a specific measured quantity within a number of nominally identical specimens. Epistemic uncertainty originates from incorrect or incomplete measured data, overly simplistic system models or lack of knowledge about a system. It can be reduced by obtaining additional knowledge of the system, such as finding additional data, or using a more suitable model to represent the system.

Uncertainty is very important to consider, not just in engineering systems, but in all systems. In a given situation, there can be a large number of uncertainties present, and of differing types. These uncertainties can come from a variety of different sources and their effects on the system may differ in magnitude. It is unwise, however, to simply ignore these effects. Examples of uncertainties on a typical engineering system are uncertainties in the parameters of the system, uncertainties in the loading/operating conditions, uncertainties in the material or geometric properties, measurement uncertainties and uncertainties relating to lack of knowledge about the system.

In the past, the approach has often been to use a nominal model, where all uncertainties within a system are viewed as being fixed at their mean values. This has the danger that the model is then not designed to deal with all operating regimes that the system is likely to encounter in its real-life usage. Many designers simply use a factor of safety for the system, where the system is shown not to fail at some operating regime far from

its original design regime. The disadvantage of this approach is that it can be excessively conservative. The system may never be likely to approach the point where the calculations have been made, meaning that the system has been over-designed. For an engineering component, this would generally mean that it is much heavier/stronger than it needs to be, leading to a corresponding increase in costs.

A more efficient method would be to consider the uncertainties associated with the system in its operation, and calculate the effect that these would have on the output. A range prediction could be made for the system, showing the possible variation of the output in relation to the uncertainties that are present. This prediction would be less conservative than that arising from the use of a safety factor, with a subsequent reduction in costs for the system.

Uncertainty within engineering systems is a reasonably active field of research, and therefore a large number of approaches have been proposed in the past. These generally separate into two groups - possibilistic methods and probabilistic methods. Probabilistic methods use statistics regarding the distributions of the uncertain quantities, so frequently involve the use of probability density functions (PDFs). They therefore usually depend on having some prior knowledge of the system, which can be estimated. The predictions obtained from this method also take the form of distributions, so it can be seen where the "most likely" response is to fall. Probabilistic methods are often quite computationally expensive, and examples are Bayesian statistics and the Monte Carlo method.

Of these two different uncertainty sub-groupings, the majority of past research in the literature has been done on the probabilistic approach. This thesis concentrates on using possibilistic methods, where there is more opportunity for fresh work to be attempted.

Following this introductory chapter, chapters 2-6 of this thesis can be classed as theory sections. Together these outline the computational and mathematical techniques that will be used within the remainder of the thesis. Chapter 2 describes interval techniques, and in particular, interval arithmetic. Interval arithmetic is one of the most well known possibilistic modelling methods, and its basic operation is explained in this chapter. The main purpose of the chapter is to illustrate the deficiencies that interval arithmetic and the other included interval techniques have, and why they are unsuitable for the applications undertaken here.

Chapter 3 then details a number of methods that are all based on the design of experiments methodology, grouped together under the loose heading of “possibilistic sampling methods”. The first of these methods is the vertex method. This then leads on to the transformation method, a possibilistic approach that is used heavily within this thesis. The transformation method was proposed as a means of conducting fuzzy arithmetic, so this branch of mathematics is briefly outlined beforehand.

Chapter 4 examines evolutionary computation, an umbrella term for a number of different algorithm types that mimic the process of natural selection. The well known genetic algorithm type is introduced first, and its characteristic features are explained. This chapter serves to introduce the differential evolution algorithm, an evolutionary strategy that is used several times within this thesis.

Chapter 5 introduces autoregressive models for representing discrete time systems. Many of the subsequent thesis chapters examine problems within the time domain, and this type of computational model forms a significant part of the modelling. The pseudoinverse approach is also introduced within this chapter; a mathematical technique for minimising a time series regression error with respect to certain parameters.

Chapter 6 gives an outline of artificial neural networks, with a brief description of their origin and methodology. It then goes on to outline the two most popular types of neural network – the multilayer perceptron and the radial basis function network, as well as information on how they can be trained using a particular set of data. Both of these two neural network types will be used in subsequent thesis chapters.

Chapters 7-11 then apply the theory introduced up to that point. In each of these chapters, a different uncertain computational model is used to represent a particular simulated system. Between them, a number of the most popular models for time series regression are used. The models used in these chapters successively increase in complexity and detail.

Chapter 7 uses one of the simplest kinds of time series regression model – the ARX model. This model is used with uncertain parameters to represent a simulated SISO linear system with added output noise. Chapter 8 then introduces the NARX model, the successor to the conventional ARX model. In this chapter, a nonlinear model (given by

an NARX model) is replaced by an uncertain linear ARX model. The system being represented is a SISO simulation of a Duffing oscillator.

Chapters 9 and 10 both involve the use of artificial neural networks, and between them they feature the two most common types of neural network – the multilayer perceptron and the radial basis function network. In both chapters an inverse representation of a SISO pre-sliding friction system is modelled. Due to the hysteresis exhibited by this data set, ARX inputs are used for these two models. In chapter 9, interval numbers are used for the adaptive weight parameters of the multilayer perceptron network. In chapter 10, the output layer weights of the radial basis function network are set as interval numbers. These two chapters also enable a comparison to be made between the multilayer perceptron and radial basis function networks, for the case of the studied data set.

Chapter 11 uses an interval-valued NARX model to represent a nonlinear system susceptible to bifurcation. The system in question was a harmonically-driven discrete-time Duffing oscillator where an additional steady-state limit cycle was seen to appear within a particular range of forcing frequency.

Chapter 12 then presents a real-life example of uncertainty modelling, with an application to automotive spot welds. This work was done at Centro Ricerche Fiat in Turin, as part of a European Union research and training network entitled MADUSE (Modelling Product Variability and Data Uncertainty in Structural Dynamics Engineering). A spot-weld component based upon an actual automotive component was designed and a number of specimens were manufactured. These specimens include differences in certain parameter values, to represent uncertainties. An uncertain dynamic response of the component is calculated using finite element modelling, and the results from this computational model are compared to experimental modal testing results.

Chapter 13 then groups together the conclusions from chapters 7-12, followed by a discussion of these results and the work undertaken here as a whole. This discussion includes an assessment of the work done here, with suggestions for improvements and future work. Much of the work in this thesis has been presented at academic conferences, the details of which are given in appendix A.

1.2 Background

Possibilistic uncertainty modelling approaches can be traced back to the 19th century, with the first use of design of experiments by Sir John B. Lawes. That study used the design of experiments methodology to examine the effects of different factors on growing crops at the well-known Rothamsted agricultural research institution. Sir Ronald A Fisher then released a book entitled "The Design of Experiments" (1935), considered to be very important from a retrospective point of view.

The interval technique of Interval Arithmetic (1966) was first introduced by Ramon Moore as his PhD thesis, and is well used within the field of uncertainty analysis. The third main area of possibilistic methods, fuzzy set theory, was introduced by Zadeh (1965). This is also a widely used technique, and tends to use subjective measures for uncertain quantities, as opposed to the objective approach used in probabilistic methods.

For previous related work that the author is aware of, interval arithmetic has been used with neural networks by Drago and Ridella (1998). That application was for a classification problem though, unlike the regression applications used for neural networks here. The differential evolution algorithm outlined in Chapter 4 has been used to train neural networks by several people, including Abbas (2002), Masters and Land (1997), and Plagianakos et al (2001), although not in the case of interval-valued networks.

1.3 Information-gap Decision Theory

Information-gap theory was proposed by Ben-Haim (2001), and is an alternative uncertainty modelling approach. It is specifically designed for decisions made under severe uncertainty, i.e. where there is a large information gap between the information that is required to undertake a given task, and the information that is actually known. It is therefore not intended as a replacement for the conventional probabilistic uncertainty modelling approaches, but as an alternative approach for problems where a probabilistic approach would be unsuitable.

Probabilistic and possibilistic uncertainty modelling approaches can be described as "distribution-based", since they usually rely on the generation of probability density functions or fuzzy membership functions for the uncertainty quantities present.

Information-gap theory is generally used for problems where there is insufficient information available for these distributions to be estimated.

An information-gap model can be defined as $U(\alpha, \tilde{u})$, where α is called the “uncertainty parameter” and \tilde{u} is the nominal model for the uncertain system. The above model can be thought of as the set of all possible models u that differ from the nominal model \tilde{u} by less than α . The info-gap model $U(\alpha, \tilde{u})$ is therefore described as being “centred around \tilde{u} ”. There are many different forms of info-gap model available, depending on the type of system that is being modelled.

The performance of the uncertain system can be obtained by specifying some “reward” function $R(q, u)$, where q is a decision vector representing the choices made by the user, and u is the system info-gap model $U(\alpha, \tilde{u})$. If the system is assessed by several different performance measures, each of these measures will have its own individual reward function $R_i(q, u)$.

The two principal tools of information-gap theory are the two “immunity” functions of robustness and opportunity. It is recognised that uncertainty can be either favourable or adverse, i.e. it can provide behaviour that is either worse or better than that anticipated by the user. The robustness and opportunity functions take these two contrasting factors into account, and are often the measures by which an uncertain system is assessed.

Robustness is the immunity of the uncertain system to failure, where failure is the inability of the system to satisfy some minimum level of performance. This minimum performance level is specified by the user as some critical value r_c of the reward function $R(q, u)$. The robustness value $\hat{\alpha}$ for a given system is:

$$\hat{\alpha}(q, r_c) = \max \left\{ \alpha : \min_{u \in U(\alpha, \tilde{u})} R(q, u) \geq r_c \right\} \quad (1.1)$$

In other words, robustness is the maximum level of ambient uncertainty that the system can tolerate so that the specified failure criteria r_c is never breached. A high level of robustness is preferable.

Opportunity is the immunity of the uncertain system to a “windfall success”, i.e. some user-defined performance level that is significantly better than the level that the user

would be prepared to accept. This windfall success is specified by the user as some value r_w of the reward function $R(q,u)$. By definition, $r_w > r_c$. The opportunity value $\hat{\beta}$ for a given uncertain system is:

$$\hat{\beta}(q, r_w) = \min \left\{ \alpha : \max_{u \in U(\alpha, \bar{u})} R(q, u) \geq r_w \right\} \quad (1.2)$$

In other words, opportunity is the minimum level of ambient uncertainty that is needed for the windfall success r_w to be attainable. Opportunity therefore assesses the favourable aspects of the system ambient uncertainty, and a low opportunity level is preferable.

An indication of the behaviour of an uncertain system can therefore be given by its robustness and opportunity attributes. A system with a high value of robustness and a low value of opportunity can tolerate a large amount of uncertainty without the output becoming extremely bad, and has the possibility of producing an exceptionally good result when only a small amount of uncertainty is present.

Chapters 7-11 of this thesis use an approach similar to that of the opportunity concept within information-gap theory. In these chapters, uncertainty is deliberately introduced into the parameters of an initially deterministic computational model. These uncertainties mean that a whole range of results are now possible from the models, so that the model outputs are now themselves uncertain. If the user of the model is prepared to accept some level of tolerance on the prediction obtained, then the prediction error of the model in relation to the true target outputs can be lower. In almost all real-life engineering applications there is some degree of output tolerance that can be accepted, so it makes sense to use this tolerance to best advantage.

1.4 Computational Information

A number of different computer workstations were used for the work here. These included:

- A 1.6GHz Dell 4300 Windows XP PC with 768MB RAM
- A 3.2GHz Packard Bell Windows XP laptop with 512MB RAM

- A 2.8GHz Windows 2000 desktop PC with 1024MB RAM, later converted to a Fedora core 3 Linux workstation
- An IRIX64 Unix workstation
- A HP-Unix workstation

The finite element models at Centro Ricerche Fiat were solved with dedicated IBM servers, each containing four processors.

The majority of the computational work in this thesis was done by self-written Java routines, using Java SDK 1.4.2 and the Eclipse 3.0 IDE. Source code is available upon request. Matlab 6.1 was also used on occasions, particularly for the calculation of pseudoinverse solutions, and for plotting of results. Chapter 10 used the NETLAB neural networks toolbox for MATLAB, a free source code package available online.

Chapter 12 used Hypermesh 6.0 as a finite element pre-processor and post-processor, and Nastran 2004 as the finite element solver. Specialist Centro Ricerche Fiat software packages were also used for the calculation of certain modal analysis results.

2. Interval Techniques

Interval techniques are used for handling uncertain quantities, where only the possible range of each quantity is known. The most basic interval technique is interval arithmetic, proposed by Ramon Moore (1966).

2.1 Interval Arithmetic

Interval arithmetic stores each uncertain quantity that it uses in the form $[low, high]$. It considers the extreme values of the uncertain quantities, and is an inclusive approach, so it states the entire range of values that the variable can take.

Unlike a PDF, there is no information about the probability or likelihood of the variable taking a particular value within the interval range. The interval range cannot be viewed as being a uniform distribution of that variable, since there is no information available with which to make that statement. It is simply a case that nothing is known about the variable, except the extreme values that it can take.

Within the original interval arithmetic definition by Ramon Moore, a set of rules were also defined for all the standard mathematical operations on interval numbers with interval arithmetic, and some of these are given below:

$$[a, b] + [c, d] = [(a + c), (b + d)] \quad (2.1)$$

$$[a, b] - [c, d] = [(a - d), (b - c)] \quad (2.2)$$

$$[a, b] * [c, d] = [\min(ac, ad, bc, bd), \max(ac, ad, bc, bd)] \quad (2.3)$$

$$C * [a, b] = [Ca, Cb] \forall C > 0 \quad (2.4)$$

$$F([a, b]) = [F(a), F(b)] , F \text{ is non-decreasing within } [a, b] \quad (2.5)$$

$$F([a,b]) = [F(b), F(a)] , F \text{ is non-increasing within } [a,b] \quad (2.6)$$

Interval numbers are both associative and commutative, i.e. for the interval numbers X, Y and Z

$$X + (Y + Z) = (X + Y) + Z \quad (2.7)$$

$$X + Y = Y + X \quad (2.8)$$

$$X * Y = Y * X \quad (2.9)$$

Interval arithmetic can be used to give an interval result for a series of calculations involving interval numbers. In the case of a mathematical model with interval-valued parameters or uncertain inputs, this cumulative calculation ability would initially appear to be beneficial, as it could potentially be used to estimate the range of model outputs subject to these uncertainties.

Interval numbers are not distributive, however. For example:

$$I(J + K) \neq IJ + IK \quad (2.10)$$

They do exhibit sub-distributivity though:

$$I(J + K) \subseteq IJ + IK \quad (2.11)$$

The lack of distributivity leads to an inherent conservatism. It is the well-known and serious deficiency of interval arithmetic, and is generally known as "dependency" or "the bound explosion effect". Consider, for example, the interval number

$$Z = [-1,1]$$

The simple arithmetical expression $Z - Z$ is evaluated by interval arithmetic as:

$$Z - Z = [-1,1] - [-1,1] = [(-1-1), (1+1)] = [-2,2]$$

However, basic mathematical knowledge states that any number subtracted from itself equals zero, so the range of the expression has been significantly overestimated. The problem is that interval arithmetic considers all quantities that it uses to be completely independent of each other, which is untrue in situations such as the above arithmetic expression. In the above expression, interval arithmetic does not recognise that the two operands of the expression are the same number, but views them as being separate numbers which just happen to take the same interval range. This bound explosion effect becomes an issue when dealing with recursive models, i.e. any model in which outputs from the model are in some way "fed back" into the model as inputs. The consequence of this recursion is a continual, progressive widening of the interval result, to the point where the calculated range becomes useless for a real-life application. An example of this behaviour is shown in Figure 2.1.

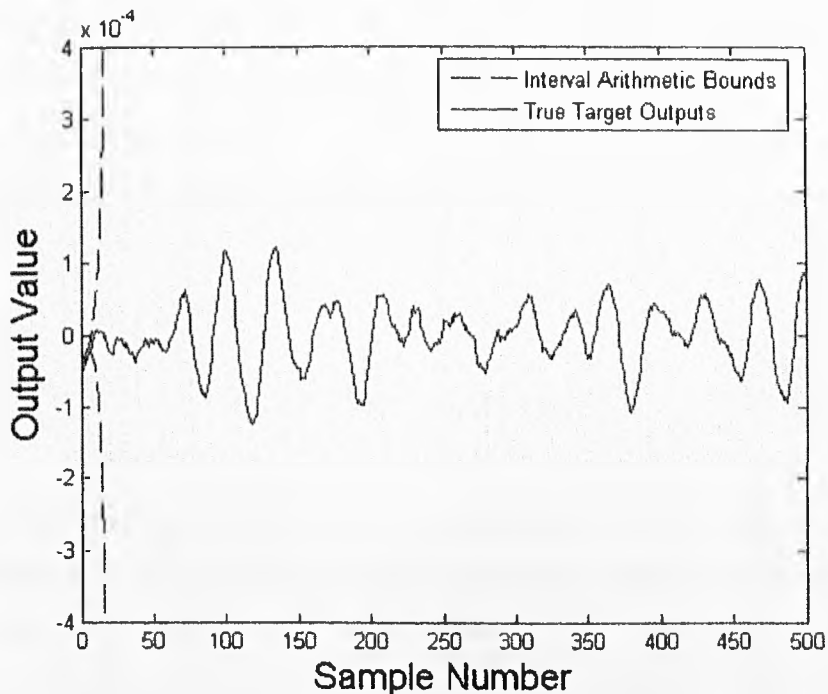


Figure 2.1: Interval arithmetic long-range ARX prediction bounds

Figure 2.1 shows a time series prediction from an interval-valued ARX model, the operation of which is defined later in Chapter 5. The target output data is also displayed for comparison. It can be seen that for each discrete time step, the predicted output range by

interval arithmetic becomes successively wider, with the overall effect of exponentially diverging output bounds. This divergence will continue until the bounds reach the limits of machine precision. Even if the initial ranges of the uncertain quantities are kept very narrow, this bound explosion is still shown to occur for a recursive model. The logical conclusion is, therefore, that interval arithmetic simply isn't suitable for modeling any kind of recursive model, so an alternative approach is needed.

2.2 Affine Arithmetic

Affine arithmetic was proposed by Comba and Stolfi (1993) with an application to computer graphics. The technique is also designed for uncertain quantities and, like interval arithmetic, can give the total uncertainty after a series of calculation steps. Affine arithmetic uses a different calculation approach though, to try and overcome the dependency issues that plague interval arithmetic.

Each source of uncertainty within the system is considered separately and assigned its own index. Any uncertain quantity A is denoted as a linear sum of its mean or "central" value a_0 and all of the uncertainty sources that the quantity is subject to, weighted by their respective magnitudes a_i . Each uncertainty source is denoted by ε_i .

For example,

$$A = 2.3 + 0.4\varepsilon_1 - 0.6\varepsilon_2 + 0.5\varepsilon_3$$

indicates that the uncertain quantity A has a central value of 2.3, and is subject to the uncertain sources 1, 2 and 3, with respective magnitudes of 0.4, (-0.6) and 0.5. Each ε_i coefficient can take a value between -1 and 1. Hence,

$$A_{\max} = 2.3 + 0.4 + 0.6 + 0.5 = 3.8$$

$$A_{\min} = 2.3 - 0.4 - 0.6 - 0.5 = 0.8$$

This affine form gives much more information about A than the corresponding interval

arithmetic form of

$$A = [0.8, 3.8]$$

As with interval arithmetic, the original affine arithmetic definition by Comba and Stolfi also provided rules for arithmetic operations with affine numbers:

$$A + B = (a_0 + b_0) + (a_1 + b_1)\epsilon_1 + (a_2 + b_2)\epsilon_2 + \dots + (a_n + b_n)\epsilon_n \quad (2.12)$$

$$A - B = (a_0 - b_0) + (a_1 - b_1)\epsilon_1 + (a_2 - b_2)\epsilon_2 + \dots + (a_n - b_n)\epsilon_n \quad (2.13)$$

Multiplication of two affine numbers is more complex. Affine arithmetic only considers linear uncertainty terms, whereas any multiplication has the ability to generate second order uncertainty terms, i.e. $\epsilon_i \epsilon_j$. This is circumvented at the end of each affine multiplication by summing the absolute values of any second order terms that are present, and adding a new uncertainty source with that magnitude to the result.

$$A * B = a_0 b_0 + (a_0 b_1 + a_1 b_0)\epsilon_1 + (a_0 b_2 + a_2 b_0)\epsilon_2 + \dots \quad (2.14)$$

$$+ (a_0 b_n + a_n b_0)\epsilon_n + \left(\sum_{i=1}^n \sum_{j=1}^n |a_i b_j| \right) \epsilon_{n+1}$$

There are in fact two modifications that can be added to this original affine arithmetic multiplication definition, in order to reduce conservatism on the result. The first of these modifications relates to the appearance of squared uncertainty terms in a calculated quantity. In equation (2.14), $a_i \epsilon_i^2$ is considered to take the range $[-a_i, a_i]$. Although ϵ_i symbols within affine arithmetic can vary between -1 and 1, ϵ_i^2 can only vary between 0 and 1, so $a_i \epsilon_i^2$ has the true range of $[0, a_i]$. This can be overcome by writing a squared uncertainty term as

$$a_i \epsilon_i^2 = \frac{a_i}{2} + \frac{a_i}{2} \epsilon_i^2 \quad (2.15)$$

prior to summing the moduli of the second order terms, so that the true range of the term is maintained.

The second modification that can be made regards the cancellation of second order uncertainty terms arising from a product of two different linear terms, i.e. $a\varepsilon_1\varepsilon_2$. In equation (2.14) the absolute values are taken and the summation is made without any collection of similar terms, so $b\varepsilon_i\varepsilon_j - c\varepsilon_i\varepsilon_j$ becomes $(b+c)\varepsilon_i\varepsilon_j$. In this second modification, all similar second order uncertainty terms are collected together and simplified before being summed. Both of these two modifications to the original affine arithmetic multiplication definition can reduce conservatism of the result obtained.

Affine arithmetic is useful for considering any interactions between two uncertain terms. If two affine numbers share one or more uncertainty sources then there is a correlation between the two, and the technique takes this correlation into account. If the quantity Z in section 2.1 is rewritten as

$$Z = 0 + \varepsilon_1$$

Then,

$$Z - Z = 0 - 0 + (1 - 1)\varepsilon_1 = 0$$

which is the true result.

The additional information stored by affine arithmetic leads to it being significantly more computationally expensive than interval arithmetic. Any multiplication of two affine quantities will generate a new uncertainty term, so a large calculation can end up with thousands of these terms. Trigonometric functions (including hyperbolic functions) can be represented within affine arithmetic by using the Chebyshev approximation, but each of these operations also generates a new uncertainty term for the approximation error.

Although affine arithmetic attempts to overcome the dependency problems present in interval arithmetic, for a recursive model the same “bound explosion” behaviour can be

seen to occur. This divergence of the prediction bounds is dependent on the ranges of the uncertain parameters however, as for small uncertainties the bounds can remain stable.

Figure 2.2 shows an affine arithmetic prediction for the same interval ARX time series model as used in Figure 2.1. The interval arithmetic bounds are superimposed on the same axis for comparison purposes. It can be seen that although the affine arithmetic prediction bounds are not expanding at the same rate as the interval arithmetic bounds, they are still expanding.

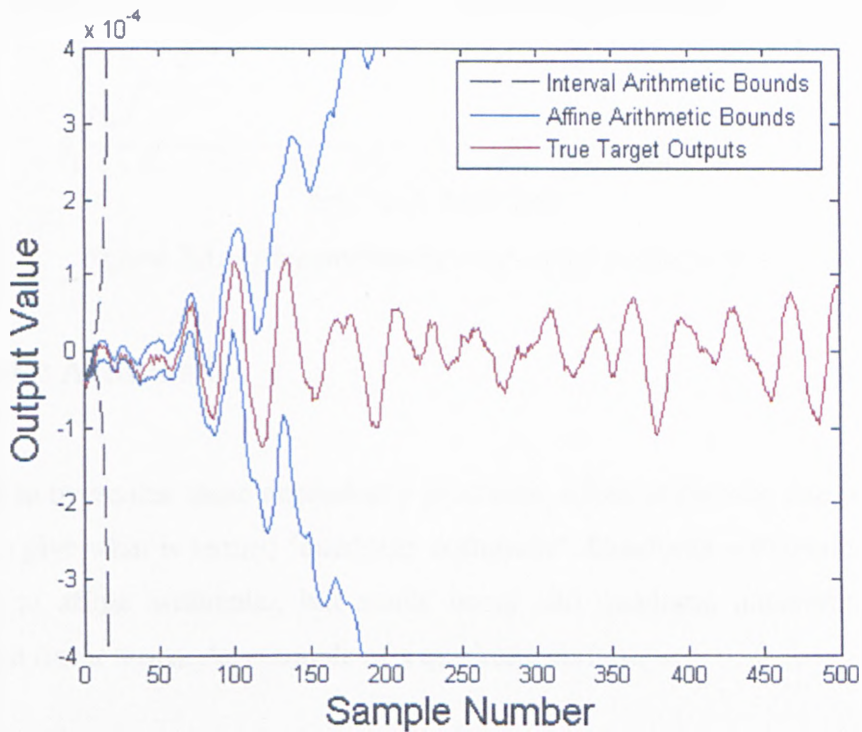


Figure 2.2: Interval and affine arithmetic long-range ARX prediction bounds

Figure 2.3 shows the width of each predicted sample range for successive discrete time instants. The increase in bound width is not monotonic, but there is an undeniable widening effect with time. The unpredictability of knowing when the predicted output range of the model will become unstable, coupled with the computational expense of the method means that, as with interval arithmetic, affine arithmetic is also unsuitable for predicting an uncertain recursive model.




Figure 2.3: Affine arithmetic long-range output widths

2.3 Quadratic Arithmetic

In an attempt to overcome these dependency problems, affine arithmetic can be taken one step further to give what is termed "quadratic arithmetic". Quadratic arithmetic works on a similar basis to affine arithmetic, but stores linear and quadratic uncertainty terms as opposed to just linear terms. An example of a quadratic variable is

$$C = c_0 + c_1 \varepsilon_1 + c_2 \varepsilon_2 + c_{1,2} \varepsilon_1 \varepsilon_2 + c_{1,1} \varepsilon_1^2 + c_{2,2} \varepsilon_2^2 \quad (2.16)$$

where C has a central value of c_0 , and is susceptible to uncertainty sources 1 and 2. Just as moving from interval to affine arithmetic gives a large increase in computational effort, going from affine to quadratic arithmetic means an even larger increase, often to beyond the realms of practicality.

Not only is quadratic arithmetic hugely computationally expensive, but it is also still susceptible to the unstable "bound explosion" behaviour exhibited by interval and affine arithmetic, albeit at an even slower rate of divergence.

3. Possibilistic Sampling Methods

This chapter outlines the theory behind several different possibilistic sampling methods for modelling uncertainty. These methods are all used to calculate a system output value when it is subject to one or more uncertain parameters, and are based on the design of experiments methodology. In each case, the uncertain parameters of the system form a multi-dimensional hypercube within the parameter space. A number of sample points are deterministically selected from within this hypercube, each one representing a different crisp set of parameter values. From the set of points that are calculated, the maximum and minimum output values are returned as an estimate for the system output range.

The sample points can be denoted as

$$\{c_i\} , 1 \leq i \leq N \quad (3.1)$$

where N is the total number of points used. If the function or system output value is denoted by f , the set of sampled values is given by:

$$\{f(c_i)\} , 1 \leq i \leq N \quad (3.2)$$

An estimate F for the possible output range of the function or system is then given as:

$$F = [\min\{f(c_i)\}, \max\{f(c_i)\}] , 1 \leq i \leq N \quad (3.3)$$

Equations (3.1), (3.2) and (3.3) are used in all of the possibilistic sampling methods that will be described here. These methods differ only in the number of sampling points used, and how they are distributed within the uncertain parameter hypercube.

3.1 Vertex Two-Level Full Factorial Method

This is the simplest of the possibilistic methods used here, and was originally proposed by Dong and Shah (1987). It is analogous to a design of experiments test with full factorial

level 2. The method states that if the function/system output in question is monotonic over the parameter hypercube being used, the output range can be calculated by considering only the extreme points of the hypercube, i.e. the vertices. These vertex points are illustrated in red on Figure 3.1 for the case of three uncertain parameters. For a system with n uncertain parameters, the number of sample points used in this method will be 2^n .




Figure 3.1: Vertex two-level full factorial hypercube sampling points

If the function being modelled is not monotonic over the parameter range, the two-level vertex method can incorrectly estimate the range of the result. For example, consider the function:

$$f(x) = 2x^2 - x \quad , \quad 0 \leq x \leq 1$$

The function is plotted in Figure 3.2, along with the sampling points of the method. Basic calculus can be used to ascertain the minimum point at $x = 0.25$. However, the two-level vertex method only considers the points $x = 0$ and $x = 1$:

$$f(0) = 0$$

$$f(1) = 1$$




Figure 3.2: Vertex two-level sampling points for $f(x) = 2x^2 - x$

The range of the function $f(x)$ over the interval $[0, 1]$ is thus estimated as $[0, 1]$. At the minimum point, $f(0.25) = -0.125$, so the true range of $f(x)$ is $[-0.125, 1]$. The two-level vertex method has overlooked the minimum point at $x = 0.25$, and therefore incorrectly estimated the range of the function.

The original paper suggested placing extra sample points at the location of local maxima and minima, in order to correct this behaviour. This can only be applied if the function is known, however. For a problem that estimates the outputs of a particular system subject to uncertain parameters, the output function of the system will not be known, and hence, neither will the locations of any extremum points in the parameter space.

3.2 Vertex Three-Level Full Factorial Method

This method is analogous to a design of experiments approach with full factorial level 3. Figure 3.3 shows the sample points in red for a problem with three uncertain parameters.

For a problem with n uncertain quantities, the number of sample points used for the three-level vertex method will be 3^n . For an equal number of uncertain parameters, this three-level vertex method increases the number of sample points when compared to the two-level

method by a factor of 1.5^n . If only a few uncertain parameters are present, this difference is not particularly significant, but for a large number of these parameters it can mean a huge increase in computational expense.




Figure 3.3: Vertex three-level full factorial hypercube sampling points

3.3 Fuzzy Arithmetic

Fuzzy arithmetic and fuzzy sets are not directly used within this thesis. The transformation method technique however (which will be outlined in sections 3.4 - 3.6), was developed as a practical means of conducting fuzzy arithmetic, so some explanation of fuzzy set theory is needed here.

Fuzzy set theory was originally proposed by Zadeh (1965). Fuzzy arithmetic subsequently arrived, and forms one of the main areas within the possibilistic branch of uncertainty quantification. The original mathematical concept of set theory defines a set A as:

$$A = \{x|A(x)\} \quad , \quad x \in X \quad (3.4)$$

$A(x)$ can be thought of as a truth operator, where a true or false value (1 or 0) is returned with no ambiguity. Each x value is assessed for belonging to set A by the characteristic function $\mu_A(x)$. This is defined as:

$$\mu_A : \{x \in X\} \rightarrow \{0,1\} \quad (3.5)$$

Although classical set theory is a well established branch of mathematics, there are certain applications where it can be ineffective. For example, consider the linguistic variable “x is around 5”. This is clearly a subjective quantity, containing a large amount of ambiguity. The definition of “around 5” is likely to vary between different people, so the characteristic function $\mu_A(x)$ is no longer suitable.

Fuzzy set theory was proposed to overcome these shortcomings. It extends the concept of classical set theory to allow a "subjective belief" for a given value x fulfilling some condition/criterion. Gradations of membership are used, to allow for different beliefs that the set criterion has been satisfied. The characteristic function becomes a membership function, so for a fuzzy set B :

$$\mu_B : \{x \in X\} \rightarrow [0,1] \quad (3.6)$$

Figure 3.4 shows one particular fuzzy set representation $B(x)$ of the statement “x is around 5”.




Figure 3.4: Fuzzy membership function for “x is around 5”

An important concept in fuzzy arithmetic is that of an α -cut. An α -cut is the interval range

of values for which the membership function is greater than the specified value α . A_α denotes an α -cut for fuzzy set A , and is given by:

$$A_\alpha = \{x \in X \mid \mu_A(x) \geq \alpha\} \quad (3.7)$$




Figure 3.5: α -cut of 0.5 for fuzzy set $A(x)$

Figure 3.5 shows an α -cut of 0.5 for the fuzzy set $A(x)$. A convex fuzzy set is one for which any level of α -cut will always give a continuous interval number. Figure 3.6 is an example of a non-convex fuzzy set $A(x)$, as illustrated by taking an α -cut of 0.5.




Figure 3.6: A non-convex fuzzy set

Another important concept in fuzzy arithmetic is the extension principle, also introduced by Zadeh. The extension principle is designed for calculating the output of a function f , where the arguments of the function are fuzzy numbers, i.e.

$$C(y) = f(A_1(x_1), A_2(x_2), A_3(x_3), \dots, A_n(x_n)) \quad (3.8)$$

The extension principle defines the membership function of the resulting fuzzy set $C(y)$ as:

$$\mu_C(y) = \begin{cases} \sup_{y=f(x_1, x_2, \dots, x_n)} \min\{\mu_{A_1}(x_1), \mu_{A_2}(x_2), \dots, \mu_{A_n}(x_n)\} & , \quad \text{if } \exists y = f(x_1, x_2, \dots, x_n) \\ 0 & , \quad \text{otherwise} \end{cases} \quad (3.9)$$

In short, the extension principle constructs an output fuzzy set $C(y)$ for the given input fuzzy numbers and function f . Each output value y is assessed for its degree of belonging to the set C and assigned a related membership value. For a given value of y , the entire n -dimensional input domain is sought for any input vectors (x_1, x_2, \dots, x_n) that produce y as the response from function f . Where several of these input vectors exist, a single vector is deterministically selected by considering all choices within the sup-min assessment of equation (3.9). Where no such suitable input vectors exist, a value of 0 is returned for $\mu_C(y)$.

The main drawback of the extension principle is that for the elementary mathematical operations, continuous input fuzzy sets give an infinite number of possible combinations to obtain a given output value. The extension principle can therefore only be realistically implemented for the case of discrete fuzzy sets containing a finite number of values.

One way of implementing the extension principle is to discretise the input domain, as shown in Figure 3.7. The alternative approach is to discretise the membership value range and take a selection of α -cuts at these discrete membership levels. This membership discretisation is shown in Figure 3.8, and the fuzzy number $A(x)$ can then be given by a set of interval numbers representing the ranges of the different α -cuts.




Figure 3.7: Discretising a fuzzy number by its input domain




Figure 3.8: Discretising a fuzzy number by its membership function

For the case of systems with uncertain parameters, a fuzzy output can be given for a number of fuzzy inputs that are present in the system. This use of fuzzy arithmetic has the same drawbacks as interval arithmetic, in that all uncertain quantities present are considered to be completely independent of each other. As with interval arithmetic, this can lead to significant widening of the predicted output bounds.

The transformation method (Hanss, 2002) works by combining the principles of the membership discretisation approach with the theory behind the vertex method. Whereas the vertex method considers only the hypercube at the limits of the multi-dimensional uncertain variable space, the transformation method also uses a number of hypercubes contained

within it. In relation to fuzzy numbers, the centre of an uncertain parameter hypercube corresponds to the point within each input dimension where the membership value equals unity (called the “core” in fuzzy arithmetic terminology).

Using the transformation method, each fuzzy input variable is discretised into m α -cuts. In the parameter space this gives m hypercubes of increasing size, including the outermost limits of the hypercube. This is illustrated in Figure 3.9 for three uncertain variables.




Figure 3.9: Discretising a hypercube into inner hypercubes by α -cuts

In the original paper Hanss proposed two variants of the transformation method, termed the reduced and general transformation methods. In this thesis a third variant is also used, referred to here as the modified reduced transformation method.

3.4 Reduced Transformation Method

For this transformation method variant, the vertices of each of the m hypercubes are used as calculation points. For the α -cut level of 1, the corresponding hypercube is actually just the central point of the parameter space, leaving another $(m - 1)$ hypercubes to consider. For a problem with three uncertain variables, Figure 3.10 shows the distribution of sample points in the reduced transformation method. The various α -cut levels are in different colours, with the sample points shown in yellow. For m α -cut levels and n uncertain variables, the number of points k to calculate is given by:

$$k = 1 + (m - 1)2^n \quad (3.10)$$




Figure 3.10: Reduced transformation method sampling points in three dimensions

If the function is non-monotonic over any of the n input variable ranges, it cannot be guaranteed that the true range of the function output will be returned. Some extrema points of the function may fall between the hypercube sample points, and therefore be overlooked. The effect of this oversight will be much less severe than with the use of the two-level full factorial method, since many more sample points have been used within the uncertain variable hypercube and not just on the outermost hyperplanes.

The effect of missing extrema of the function is minimised by using a greater number of α -cut levels m , since this will mean a greater density of sample points within the hypercube. Increasing the value of m has an almost linear effect on the number of calculations that need to be made, and thus the overall computational time.

3.5 General Transformation Method

The general version of the transformation method is significantly more computationally expensive than the reduced version described previously. The overall uncertain variable hypercube is still divided into m embedded hypercubes, defined by incremental α -cut levels of the n fuzzy number inputs. The distribution of sample points over successive inner hypercubes is no longer constant though.

Starting from the α -cut level of 1, this corresponds to a single crisp value within the overall

uncertain hypercube. For each incremental decrease in the α -cut level, the number of sample points in each hypercube dimension is increased by one, shown in Figure 3.11. This distribution of sample points is shown for the three-dimensional case in Figure 3.12, with some of the points on the near face omitted for clarity.




Figure 3.11: Distribution of general transformation method sampling points by fuzzy α -cuts




Figure 3.12: General transformation method sampling points in three dimensions

The general transformation method differs from the reduced version in that as well as using the vertices of the various embedded n -dimensional hypercubes, points are also taken on their $(n-1)$ -dimensional hypersurfaces. For a triangular fuzzy number, this has the effect

that on any α -cut level within the overall region, the sample points in a given dimension are always separated by the same distance. This equidistant effect is not true for a non-triangular fuzzy number (i.e. one with arbitrary L-R shape (Dubois and Prade, 1978, 1979)), but the general transformation method still gives a much more even distribution of sample points than the reduced version.

The number of sample points k for the general transformation method is given by

$$k = \sum_{p=1}^m p^n \quad (3.11)$$

For a large number of uncertain variables, the number of calculations can be several orders of magnitude greater than with the reduced transformation method.

3.6 Modified Reduced Transformation Method




Figure 3.13: Modified reduced transformation method sampling points in three dimensions

The approach outlined here, termed as the modified reduced transformation method, is a compromise between the reduced and general versions, with respect to computational expense. As with these two original versions, the membership function range of each fuzzy input variable is divided by m α -cuts, giving m embedded hypercubes within the uncertain region. Whereas the reduced transformation method considers the end points of each

hypercube dimension for a given α -cut, this modified version considers the centre and end points. This distribution of modified reduced transformation method points is shown for the three-dimensional case in Figure 3.13.

The number of sample points k for this modified version of the reduced transformation method is given by:

$$k = 1 + (m - 1)3^n \quad (3.13)$$

As shown in chapter 2, the interval methods of interval arithmetic, affine arithmetic and quadratic arithmetic all have dependency problems. In contrast, the possibilistic sampling methods in this chapter do not suffer from these bound explosion effects. The uncertain output range they predict, however, is not guaranteed to be the true range with respect to the uncertain variables used. If the output quantity is monotonic in all directions over the multi-dimensional variable hypercube, the predicted output range will be the true range, but in many cases this monotonicity will not exist.

Without this monotonicity, it cannot be known for sure if a particular output surface extremum has been missed by the hypercube sample points, one which would give an output value outside the predicted bounds. The general, reduced, and modified reduced versions of the transformation method make it less likely that one of these features has been missed, but it still cannot be known with any certainty. The true range could only be stated for definite if the entire output surface topology was known in advance. This however, would completely eliminate the need to use these sampling methods in the first place.

Figure 3.14 shows a prediction from the interval-valued ARX model of Figures 2.1 and 2.2, with the interval bounds calculated by the three-level vertex method. It can be seen that the output bounds do not explode as with the interval techniques in Chapter 2, but remain at sensible and practical widths.

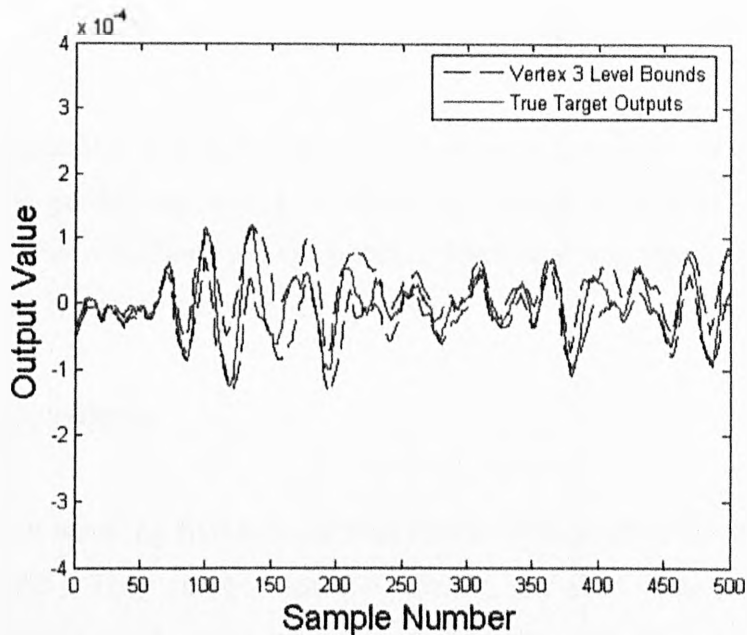


Figure 3.14: ARX long-range output bounds by the vertex three-level method

The transformation method is thus used here as a way of circumventing the dependency problems of other interval techniques, for the case of recursive models. The general transformation method uses a higher frequency of sample points at lower levels of fuzzy number α -cuts, i.e. more sample points towards the exterior faces of the uncertain variable hypercube. In practice, it is often too computationally expensive to use in many applications.

The reduced and modified reduced versions of the transformation method both use embedded inner hypercubes within the original uncertain variable hypercube, in order to set the positions of the deterministic sample points. The sizes of these respective hypercubes were set here by discretising the input variable space.

The choice of which vertex method or transformation method variant to use for a problem is principally dependent on the complexity of the problem, and is usually constrained by the available computational time and power. For a system with a large number n of uncertain parameters, changing from two-level to three-level in the vertex method can have a huge increase in the computational effort, as it is the difference between 2^n and 3^n sample points. Likewise, an increase in the number of α -cut levels for the reduced or modified reduced transformation methods will also increase the required computational time, although this will be by a linear factor.

4. Evolutionary Computation

Evolutionary computation is a subset of artificial intelligence, and is generally used as an umbrella term for genetic algorithms, evolutionary strategies, genetic programming and similar methods. These methods are all global optimisation approaches that are based on the "survival of the fittest" principle within nature.

4.1 Genetic algorithms

Genetic algorithms were the first type of evolutionary computation method, proposed by John Holland (1975). They can be used for virtually any kind of application where the parameters of a system need optimising to obtain the best result. These parameters need to be suitably represented within the algorithm, and an objective function is required to appropriately quantify the problem at hand. If these two challenges are met correctly, genetic algorithms are capable of giving effective solutions to a problem.

Genetic algorithms were inspired by the principle of natural selection, and are stochastic algorithms. The standard gradient descent optimisation method (or one of its numerous variants) considers only a single solution point within the entire parameter space, and is highly susceptible to local extrema. For a problem where the output surface is high-dimensional and multi-modal (i.e. a nonlinear problem), this can lead to difficulties in finding the global optimum. Simulated annealing (Metropolis et al, 1953) improves upon the gradient descent technique by its ability to avoid local minima, and is therefore a global optimization approach, but still only considers one solution point at a time.

In contrast, genetic algorithms simultaneously consider a whole group of solution points, known as a "population". In keeping with the genetic allegory, this population can be viewed as a "gene pool", and each solution is a "chromosome", or "genome". By scattering these solutions through the parameter space, a range of local regions can be explored simultaneously. The rationale behind genetic algorithms states that although the global optimum solution may not exist in the population at a given time, its constituent parts may be scattered amongst the genetic material of the solutions that are present. The task is therefore to find the global optimum by rearranging this genetic material, which is done by

the genetic operators of selection, mutation and recombination. The population solutions are encoded within the algorithm as binary strings.

Genetic algorithms use an iterative methodology, where each iteration is called a "generation", and the first phase of each iteration is the calculation of fitness values. The fitness of a solution is a measure of its quality (in relation to the problem at hand), and is the score obtained from the objective function of the problem. Viewing the fitness values of an entire population gives an instant guide as to which are the superior problem solutions, and which ones are inferior.

The first of the three genetic operators to be used in each iteration is the selection operator. This determines which current solutions are to be used for the subsequent mutation and recombination operations, analogous to choosing which animals are to breed in a natural environment. The solutions with higher fitness values are those which contain the better genetic material, and it is therefore more beneficial for these to be selected. If only the best solutions are selected at each iteration, however, the population as a whole will quickly converge to their values, losing the population diversity that is so essential for the algorithm to have. For this reason it is important that the selection operator retains a stochastic aspect. Of the numerous selection schemes that are available, the one that best achieves this is "roulette wheel selection", where the likelihood of a solution being selected is proportional to its fitness value.




Figure 4.1: A bit-flipping mutation scheme

Once a number of population solutions have been selected, the mutation operator follows. As with selection, there are a variety of different mutation schemes available for use. Some

of these adjust the length of the solution vectors (destructive and generative schemes), whereas others modify elements of the vectors whilst keeping the length unchanged ("bit-flipping" schemes). An example of a bit-flipping mutation scheme is shown in Figure 4.1. The purpose of the mutation operator is to introduce randomness into the search process, which the subsequent recombination operator is not able to give.

Recombination (also known as crossover) combines the genetic material of two "parent" solution vectors to produce "child" vectors. Crossover schemes generally produce two child vectors from two parent vectors, although sometimes the operation may be used to produce only a single child vector. The most popular crossover schemes are one point crossover, two point crossover and uniform crossover.

One point crossover involves randomly selecting a position index for the two selected parent vectors. All elements beyond this point are then "switched", to produce two new child vectors. This operation is illustrated graphically in Figure 4.2. Two point crossover selects two random position indices for the parent vectors, then switches all elements between these two position limits. Uniform crossover exchanges individual parent vector elements with a given probability, as opposed to the segment exchange that takes place with one point and two point crossover. All three of these crossover schemes maintain the length of the population vectors, although there exist other crossover schemes that do not.




Figure 4.2: One point crossover

After the recombination phase is completed, the final step of the iteration is the replacement phase. In this part, the newly generated child vectors are inserted into the population at the expense of certain existing solutions, as the population size always remains constant. A

number of schemes exist for specifying how this replacement operation is to be done. The two most common of these schemes are to replace the current solutions with the lowest fitness values, or to replace those that are most similar to the child vectors (generally their parent vectors). The former scheme helps to improve the overall fitness of the population, whereas the latter helps to maintain population diversity. A genetic algorithm where the number of generated child vectors equals the population size is called a "simple state" implementation, or if less children are generated than the population size, is called a "steady state" or "overlapping populations" implementation. The latter is the more commonly used of the two types.

It needs to be noted that in a genetic algorithm there is no rejection of child vectors, they are always inserted into the population at each step. It may be that a child vector actually gives a worse fitness value than its parents, especially if the algorithm is nearing convergence. Genetic algorithms are therefore capable of "uphill moves", i.e. the overall population fitness may actually decrease between two successive iterations (known as generations). A convergence criteria needs to be devised for the algorithm, and this is usually either a maximum number of iterations being exceeded, or the growth rate over a number of iterations (eg ten) falling to below an arbitrarily specified level.

Not only are genetic algorithms able to avoid local sub-optimal extrema, but they can also be used where the output surface contains discontinuities, whereas a gradient descent-based algorithm could not be. Although genetic algorithms are global optimisation algorithms, their stochastic methodology means that they are not guaranteed to find the global optimum. For this reason it is recommended to run a genetic algorithm application several times, to increase the chance of the global optimum being found.

More general information on genetic algorithms can be found in the book by Goldberg (1989), considered a seminal publication in the evolutionary computation field.

4.2 Differential Evolution

In this thesis the global optimisation algorithm of differential evolution was used; an example of an evolutionary strategy. The technique was proposed by Storn and Price (1996). Evolutionary strategies share a common origin with genetic algorithms and operate

on the same natural selection principle, but there are differences in methodology between the two.

The first of these differences relates to how parameter values are stored. Within genetic algorithms the parameters are usually encoded binary strings, whereas in differential evolution real number values are used throughout. This avoids decoding the binary representations to calculate the objective function value, and then re-encoding them afterwards. It also gives the parameters to a much higher precision. Although differential evolution still uses the three genetic operators of selection, mutation and recombination, they are applied differently than in standard genetic algorithms, and in a different order.

The three parameters that define the operation of the differential evolution algorithm are the population size NP, the weighting factor F, and the crossover constant CR. A general heuristic is that NP should be 5-10 times greater than the number of parameters being trained by the algorithm. A value of between 0.4 and 1.2 is recommended for F, and CR takes a value between 0 and 1 inclusive.

Within differential evolution the existing population solutions are called target vectors, and the child vectors are termed trial vectors. The number of child vectors generated at each iteration is equal to the population size, analogous to a simple state implementation of a genetic algorithm.

The algorithm is initialised by randomly generating a population of target vector solutions. As with genetic algorithms, each iteration starts by calculating the fitness values of all the population solutions. The mutation operator is then the first genetic operator to be used. For each current target vector in turn, three other solution vectors are randomly selected from the target population. The vector difference between the first two of these vectors is calculated, and then multiplied by the weighting factor F. The third random vector is then added to this weighted difference vector to give a "noisy vector". One noisy vector is therefore generated for each solution in the target population. The algorithm name originates from this use of vector differentials during the mutation process.

The use of vector differentials for mutation has benefits for the search process. The first of these benefits is that it allows differing ranges for different parameters, i.e. different

dimension lengths in the multi-dimensional parameter space. The second benefit is that the range of the mutation operation will decrease as the solution approaches convergence. At the start when the solutions are further apart, the differences between them will be greater, so the mutation will cause a greater "jump" in the solution space. Towards the end of the algorithm, the solutions should be gathered together in smaller local regions, so mutations are likely to stay close to these regions.

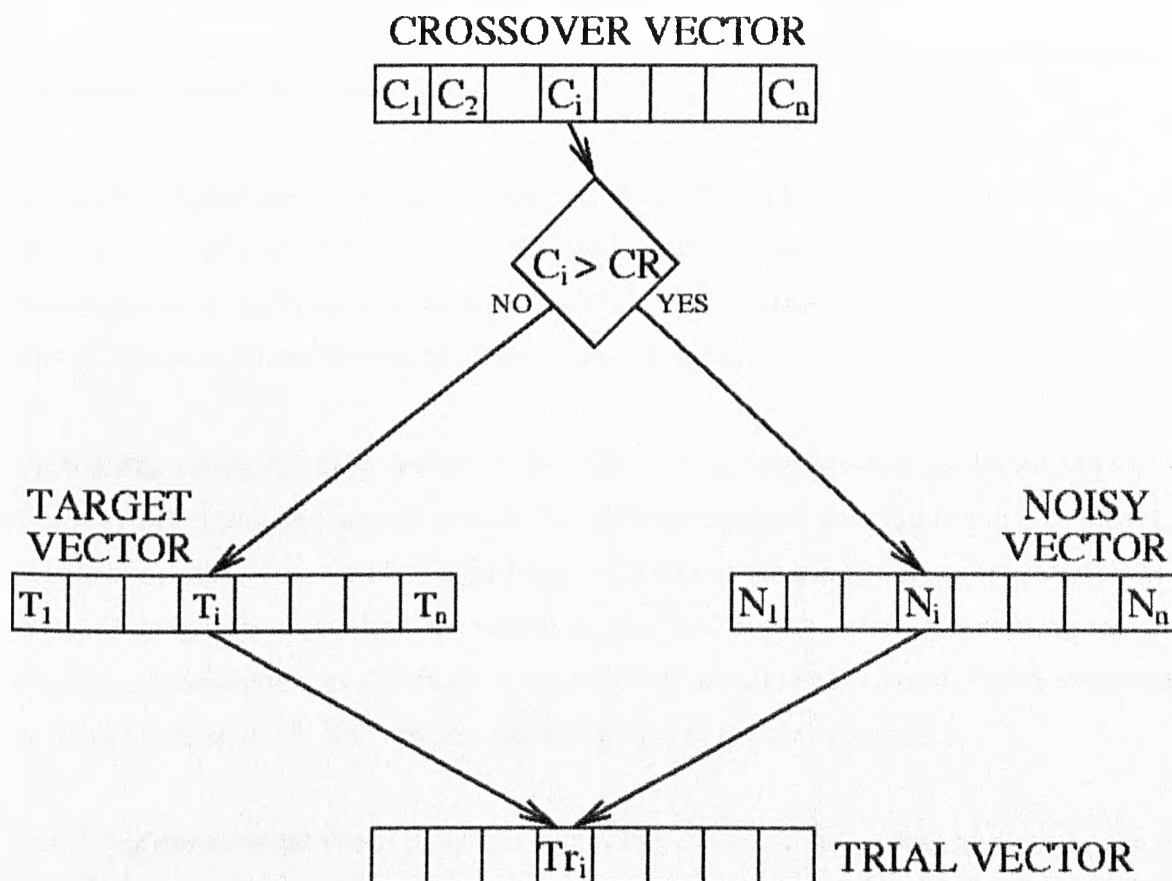


Figure 4.3: Crossover (recombination) within differential evolution

The next genetic operator to be used in the iterative process is recombination. Each target vector and its respective noisy vector are combined to make a single child (trial) vector, where the crossover process is governed by the value of the crossover constant CR . One randomly selected element of the trial vector is automatically taken from the noisy vector, to ensure that each trial vector is in some way different from its parent target vector. The remaining elements of the trial vector are then filled by uniform crossover operations between the corresponding elements of the target and noisy vectors. For each of these element positions, a random number is uniformly generated between 0 and 1. If this number

is greater than CR, the trial vector value is taken from the noisy vector, otherwise it is taken from the target vector. This recombination operation is illustrated graphically in Figure 4.3.

The CR parameter therefore has the duty of specifying the variation between target and trial vectors. For a low CR value, the trial vector elements are more likely to come from the noisy vector, meaning that each trial vector will tend to be significantly different from its parent target vector. Conversely, a high CR value means that each trial vector will tend to be very similar to its target vector. CR can hence be thought of as determining the exploration rate of the solution space.

In genetic algorithms it is often stated that recombination is a much more important operator than mutation, in relation to finding optimum solutions. In differential evolution this statement is less true; it is the combination of the mutation and crossover parameters F and CR that determines the effectiveness of the algorithm.

Once a trial vector has been generated for each existing target vector, the fitness values of this new trial population are calculated. The selection operator then follows to conclude the iteration, but this is used with a significant difference to its genetic algorithm counterpart. Whereas in genetic algorithms the selection operator chooses which solutions to use for breeding, in differential evolution all of the solutions are allowed to breed, and the selection operator chooses which child vectors should be kept in the next generation.

Unlike the replacement phase in genetic algorithms, where child vectors always survive to the next generation, trial vectors in differential evolution only advance if they have a higher fitness than their parent target vector. Hence differential evolution does not feature "uphill moves", as the populations can only be improved over successive generations, or at worst, remain the same. This selection operation is illustrated graphically in Figure 4.4.




Figure 4.4: The selection operation within differential evolution

As with genetic algorithms, differential evolution needs some form of convergence criterion to terminate the algorithm. This is generally either an arbitrary maximum iteration limit being exceeded, or the growth in overall population fitness over a number of iterations falling to below some specified value.

5. Autoregressive Models

Figure 5.1 shows a general representation of a discrete time system. This system can be physical, such as a car engine, or completely theoretical, such as a predictive model for simulating trade dynamics of the stock exchange. $u(t_i)$ is the system input signal, $y(t_i)$ is the measured output signal, and $e(t_i)$ is the disturbance signal, often assumed to be a zero-mean white noise process.




Figure 5.1: A general representation of a discrete time system

5.1 ARMAX and ARX Models

Autoregressive models are often used in system identification, with the most general form being the ARMAX model (Auto Regressive Moving Average with eXogenous Inputs). The ARMAX model estimates the value of the system output at a given time instant by a linear combination of previous system values. These values are referred to as "lags", and an arbitrary number of them can be used within a model. The form of an ARMAX model is

$$y(t) = a_1 y(t-1) + a_2 y(t-2) + \dots + a_{na} y(t-na) + b_0 u(t-nk) + b_1 u(t-nk-1) \quad (5.1) \\ + \dots + b_{nb} u(t-nk-nb) + e(t) + c_1 e(t-1) + c_2 e(t-2) + c_{nd} e(t-nd)$$

where na is the number of output lags used, nk is the system delay between input and output, nb is the number of input lags, and nd is the number of disturbance lags. The coefficient sets $\{a_i\}$, $\{b_i\}$ and $\{c_i\}$ are real number sets.

The output lag terms are referred to as autoregressive terms and the disturbance lags as moving average terms, hence the origin of the ARMAX name, with “exogenous inputs” referring to the system input lags. A model with $nb = 0$ (no input lag terms) is called an ARMA model, and with $nd = 0$, is called an ARX model. With $nb = nd = 0$, the model is simply an AR model. The individual RHS terms in (5.1) are also often called "regressor terms"

An ARX model can be written as

$$\hat{y}(t) = a_1 y(t-1) + a_2 y(t-2) + \dots + a_{na} y(t-na) + b_0 u(t-nk) + b_1 u(t-nk-1) + \dots + b_{nb} u(t-nk-nb) \quad (5.2)$$

where $\hat{y}(t)$ is the estimated output from the model. As $e(t)$ is zero-mean, the relationship

$$y(t) = \hat{y}(t) + e(t) \quad (5.3)$$

shows that the model is equally likely to underestimate or overestimate the true measured output at a particular time instant.

In relation to the standard transfer function form, na gives the number of system poles and nb is the number of system zeros. The number of output lags na is also called the order of the model. The previous “lagged” values used by autoregressive models make them particularly suitable for systems that exhibit hysteresis.

5.2 One-step-ahead and Long Range Prediction

For an autoregressive model, the output lag terms can either be taken directly from the data set of measured system values or from previously estimated outputs by the model itself. The former is called the one-step-ahead prediction, and the latter is known as the long range prediction, or the model predicted output. In terms of minimising $e(t)$, the one-step-ahead prediction is much the easier of the two to use, since any incorrectly estimated system outputs by the model will not affect the subsequent time intervals. If the model is to be used in a real-life application, however, it needs to be shown that it can accurately predict the

given system without the need for a reference output, so the long range prediction error is the one on which the quality of the model should be assessed.

When used in the one-step-ahead prediction mode, an autoregressive model is linear-in-the-parameters with respect to the regressor coefficients, but in the long range prediction mode it is not. For the linear-in-the-parameters case, the optimum regressor coefficients can be found very easily by using the pseudoinverse approach

5.3 NARX Models

The NARX model was originally proposed by Leontaritis and Billings (1985), and is a type of autoregressive model specifically for nonlinear systems. NARX stands for Nonlinear AutoRegressive with eXogenous inputs, and is a subset of the more general NARMAX model. Its general form is given by:

$$\hat{y}(t) = \sum_{i=1}^R a_i T_i \quad (5.4)$$

UNIVERSITY
OF SHEFFIELD
LIBRARY

where R is the number of terms used in the model. Whereas the regressor terms within an ARX model are linear, the $\{T_i\}$ terms in (5.4) are nonlinear functions of previous time input and output lags. The standard approach is to use polynomial functions up to a particular order, eg. $y(t-1)$, $u(t)u(t-1)$, $u(t)^2 y(t-1)$ or $u(t)^3$, although any nonlinear function can be selected. As with the ARX model, the set of regressor coefficients $\{a_i\}$ are real number values. For the one-step-ahead prediction case, the NARX model is therefore linear-in-the-parameters with respect to the regressor coefficients.

As the NARX model uses nonlinear regressor terms, it is therefore much more effective than the ARX model at representing discrete time systems that have nonlinear components. It is consequently more difficult to train, however. Along with the NARX model, Korenberg and Billings proposed an algorithm for training such a model, called the forward regression orthogonal estimator (FROE) algorithm (Korenberg et al 1988, Billings et al, 1989). For a maximum specified number of input and output lags and an upper limit on the complexity of the regressor terms (usually a maximum polynomial order), the FROE can

return the optimum set of regressor terms and coefficients for a given data set.

The regressor terms in an NARX model are not generally orthogonal, since these terms all use the same data values and several terms can share one or more input or output lags. From the regressor terms available for selection within the model, the FROE transforms these into an equivalent set of orthogonal regressor terms, and finds the optimum set of coefficients for these terms. The algorithm uses an iterative procedure to successively add new regressor terms into the model, using the one-step-ahead prediction error as an objective function to be minimised. At each iteration, the algorithm examines the remaining unused regressor terms, to see which one best minimises the one-step-ahead error when added to the model. If none of these terms can improve the model, the algorithm is seen to be converged. The resulting orthogonal solution is then transformed into its equivalent set of the original non-orthogonal regressor terms and their associated coefficients.

As the FROE algorithm only considers one-step-ahead prediction errors, the results obtained can be sub-optimal. As the long range prediction error of the model is not being considered, the algorithm is in danger of overfitting the data, since a good one-step-ahead error does not always correspond to a good long range prediction error. In order to overcome this, the simulation error minimisation with pruning (SEMP) algorithm was proposed by Piroddi and Spinelli (2003).

The SEMP algorithm also uses an equivalent orthogonal representation of the NARX model, but minimises the long range prediction error. It uses the same iterative methodology of the FROE algorithm, but also includes a "backward step" at each iteration. Therefore as well as seeking to add a new regressor term to the model at each step, it is also examined if the removal of one of the existing regressor terms can improve the long range prediction error.

Both the FROE and SEMP algorithms exponentially increase in computational effort if the maximum number of input and output lags or the order of regressor term complexity are increased, since at each iteration there will be many more possible regressor terms to consider. However, if the regressor terms within an NARX model are fixed at a particular set, the optimum coefficients can be found for these regressors (in respect to the one-step-ahead outputs) by the pseudoinverse approach.

5.4 The Pseudoinverse Approach

The model error for a given time instant is given by

$$e(t) = y(t) - \hat{y}(t) \quad (5.5)$$

If all of the model inputs at each time instant are arranged into a matrix X , called the system matrix, and the vector of regressor coefficients is given by β , the least squares error is given by:

$$E = \frac{1}{N} (Y - X\beta)^2 \quad (5.6)$$

where Y is the vector of measured system outputs and N is the number of data samples

Differentiating (5.6) with respect to the regressor coefficients β :

$$\frac{\partial E}{\partial \beta} = -\frac{2X}{N} (Y - X\beta) \quad (5.7)$$

Setting $\frac{\partial E}{\partial \beta} = 0$ gives the optimum regressor coefficients β^* for minimising the least squares error:

$$-\frac{2X}{N} (Y - X\beta^*) = 0 \quad (5.8)$$

$$X^T Y = X^T X \beta^* \quad (5.9)$$

$$\beta^* = (X^T X)^{-1} X^T Y \quad (5.10)$$

(5.10) can be calculated quickly and easily in any maths software package, such as MATLAB

5.5 Training an Autoregressive Model

For the general case of training an autoregressive model to a particular set of data, the convention is to use three separate data sets, labelled as the training, validation and test sets. This in line with the recommended best practice by Tarrashenko (1998). Initially the model hyperparameters need to be found. In the case of an ARX model, the hyperparameters only consist of the number of input and output lags for the model, although other types of computational model have additional hyperparameters.

Once the optimum hyperparameters have been selected, the optimum parameters of the model are then found. In the case of an ARX or NARX model, these parameters are the coefficients of the individual regressor terms. The solution model from this phase is then tested on the previously unseen test set. The use of a computational model on unseen data from the same process is to test for generalisation, an important requirement of any time series simulation model.

A frequently used measure for assessing the quality of a time series prediction is the mean squared error (MSE):

$$MSE = 100 * \frac{\sum_{i=1}^N (y_i - \hat{y}_i)^2}{Nv} \quad (5.11)$$

where \hat{y}_i and y_i are the respective predicted and true output values for a given data sample, N is the total number of samples, and v is the variance of the target output data. The mean squared error is simply the conventional least squares error measure normalised by the target data variance, so using that mean as a constant predicted output value would give a MSE value of 100%. A general heuristic is that 5% MSE is a good model prediction, and 1% MSE is a very good prediction.

6. Artificial Neural Networks

6.1 Introduction

Artificial neural networks are a collection of computational models that were inspired by the operation of the human brain, and in particular, by its ability to learn adaptively. The applications of neural networks can principally be grouped into two different areas – classification and regression. Classification is the more popular of the two, and covers areas such as pattern recognition, medical diagnosis and damage detection. In this thesis only regression applications are encountered however, so neural network classification usage will not be covered here.

Artificial neural networks can be used when it is necessary to deduce a relationship between a set of system inputs and outputs, for the case where the system is nonlinear and an analytical solution is not possible or practical. Neural networks are often described as “black box” models, due to them predicting a set of system outputs from a given set of inputs without explicitly giving the relationship between the two.

Neural networks consist of a selection of processing nodes arranged together in a particular structure. The original concept of a neural processing node was inspired by the operation of an individual neuron within the human brain (McCulloch and Pitts, 1943). Each neuron is connected to a number of adjacent neurons by synaptic connections of varying strengths. If the summed electrical signals into the neuron exceed a particular threshold value, the neuron “fires” and generates a new electrical impulse of its own. The operation of the brain at any one time is defined by the electrical activity present and the strengths of all the synaptic connections.

The human brain is able to adapt to new experiences and, through practice, improve a person's proficiency at performing the given task. This is done by forming new synaptic connections and adjusting the strengths of existing connections. Neural networks can be “trained” to a particular system by a similar process. A data set is required of measured system values, but this data has to meet certain constraints, otherwise the resulting neural network model will be poor. The data must be representative of the system being modelled,

there must be enough data available, and it must be of a suitable quality, i.e. without excessive measurement noise, incorrectly recorded values and omissions.

By using the measured data set, the parameters of the neural network model are continually adjusted until the network model is suitably able to predict the system outputs in question. This type of training procedure, where the true system outputs are available as a reference, is called supervised learning. It is important, however, that the network model is able to represent the general system in question, and not just the measured data. The latter case is known as “overfitting”.

Based on this, neural network applications always involve separating the measured data into at least two sets; sometimes three. The parameters of the neural network are trained iteratively, for which there are a number of different algorithms available, and by using only the first of the data sets. The resulting solution network model is then used on the second “unseen” data set, in a process known as generalisation. If the network is able to effectively model this unseen data set then it has mapped the underlying relationship of the system, and not just fitted to the initial training set.




Figure 6.1: A single-layer perceptron network

The earliest type of neural network was called a perceptron (Rosenblatt, 1962), and is shown in Figure 6.1. The functions ϕ_i are processing elements that transform raw input data into a more usable form, and the parameters $\{w_i\}$ are adjustable weights. The function

g is called the activation function, and is analogous to the assessment that determines whether a brain neuron should fire. The activation functions used in these early neural networks were simply anti-symmetric step functions. The predictive capabilities of the perceptron were poor however, and so improved neural network models were necessary.

By using additional layers of processing nodes the multilayer perceptron network was created; one of the two main types of neural network in use today. The second type is the radial basis function network, which differs in operation to the multilayer perceptron and is the lesser used of the two types.

6.2 Multilayer Perceptron Networks

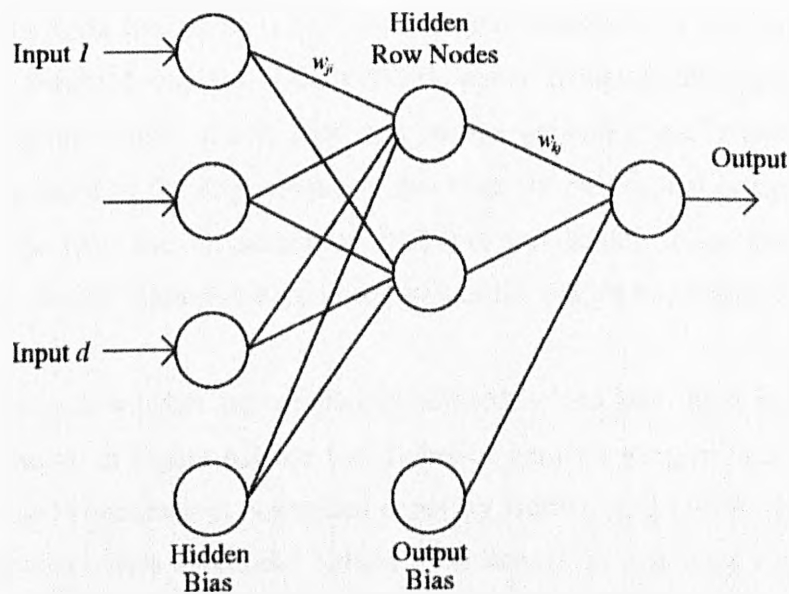


Figure 6.2: A multilayer perceptron neural network

The multilayer perceptron (MLP) network arose because it was observed that single layer networks were limited in the range of functions that they could represent. The general form of an MLP network is shown in Figure 6.2. The number of network inputs and outputs is determined by the problem being modelled, and information passes through the network in a forward direction only. An arbitrary number of hidden row nodes can be selected for use, although excessive amounts are likely to be limited by computational expense. An MLP network is often referred to by the number of adjustable weight layers that it uses; for example, the MLP network in Figure 6.2 is a two-layer MLP.

The operation of a two-layer MLP network is given by:

$$y_k = \sum_{j=0}^M w_{kj}^{(2)} g \left(\sum_{i=0}^d w_{ji}^{(1)} x_i \right) \quad (6.1)$$

where M is the number of hidden row nodes, d is the number of network inputs, $w_{kj}^{(2)}$ is a second layer weight and $w_{ji}^{(1)}$ is a first layer weight. Equation (6.1) assumes that the network does not use nonlinear output layer activation functions. Each of the hidden row and output nodes are subject to additional bias parameters, given by $w_{j0}^{(1)}$ and $w_{k0}^{(2)}$ respectively, where the inputs to these biases are permanently set to 1.

Each hidden row node uses an activation function g to transform its summed inputs, which are themselves weighted versions of the network inputs. Although this type of network still uses the perceptron name, the hidden row nodes generally use sigmoidal activation functions, as opposed to the step activation functions of the original perceptron networks. The output node may use an activation function, but usually does not for regression applications, so simply calculates a weighted sum of the hidden node outputs and bias node.

All layers of network weights are completely adjustable, and take the form of real number values. The network in Figure 6.2 has two layers of network weights, although additional layers can be used if necessary. A seminal paper by Hornik et al (1989) showed that two-layer MLP networks with sigmoidal hidden node activation functions can represent any continuous function to an arbitrary degree of accuracy if a sufficient number of hidden row nodes are used. Based on this, two layer multilayer perceptron networks are frequently used for a wide range of modelling problems.

The most common way of training MLP network weights by the error back-propagation algorithm, proposed by Rumelhart et al (1986). This can be described as the standard gradient-descent optimisation technique implemented for neural networks. Error back-propagation operates by differentiating backwards through the network, so cannot be used if the hidden node activations are step functions. Other training methods can be used to find the optimum weights set, however, including the use of genetic algorithms. The first use of genetic algorithms to train neural networks was by Montana and Davis (1989).

6.3 Radial Basis Function Networks

The radial basis function neural network was originally proposed by Broomhead and Lowe (1988). It is based on the concept of a generalised linear discriminant, where the input variables are transformed by nonlinear functions, then the model output is given by a linear sum of these transformed inputs. An example of a radial basis function network is shown in Figure 6.3




Figure 6.3: A radial basis function neural network

The network takes its name from the basis function nodes used in the hidden row. A number of different functions can be used by these nodes, but the Gaussian function is usually the most popular choice. Whereas the hidden node activations in a multilayer perceptron operate on a weighted sum of the network inputs, the basis function activations here use the distance between the network input vector x and some specified “centre” vector μ . Each basis function node output is given by:

$$\phi_j = \exp\left(\frac{-\|x - \mu_j\|^2}{2\sigma_j^2}\right) \tag{6.2}$$

where σ_j is a width parameter.

Unlike the multilayer perceptron, no weights are used prior to the hidden basis function nodes. The network outputs are therefore given by:

$$y_k = \sum_{j=1}^M \phi_j w_j + w_0 \quad (6.3)$$

where M is the number of basis function nodes. As with the multilayer perceptron, each output node uses a bias w_0 , where the input to this weight is permanently set to 1.

The main advantage of a radial basis function over a multilayer perceptron is that it can be trained much faster. The training phase generally takes place in two steps, the first of which finds the optimum basis node parameters by generating a mixture model for the input data. This is an example of unsupervised network learning, and one of the most popular choices for this first training phase is to use the Expectation Maximisation algorithm (Dempster et al, 1977).

The Expectation Maximisation (EM) algorithm was proposed as a solution for incomplete data problems. Approximating the probability density function of a given data set by a mixture model is an example of an incomplete data problem, since for each data sample it is not known which component density function the sample can be attributed to. The EM algorithm is an iterative algorithm where each iteration consists of two separate steps – the expectation step and the maximisation step. It enables the optimum parameters to be found for the Gaussian mixture model, in relation to minimising the error function of the model.

The second training stage for a radial basis function network then trains the output layer weights $\{w_j\}$. By using the basis function parameters already found with the EM algorithm, the optimum output layer weights can simply be found by the pseudoinverse approach for the training data set.

7. Uncertain ARX Modelling of a Linear System

7.1 Initial Crisp ARX Model

In this first piece of work, an interval-valued ARX model was used to represent a linear system. The system in question was a simple SISO time series representing a SDOF vibrating structure, and was defined by:

$$m\ddot{x} + c\dot{x} + kx = F(t) \quad , m = 1\text{kg}, c = 20\frac{Ns}{m}, k = 10000\frac{N}{m} \quad (7.1)$$

where $F(t)$ was a broadband random excitation. 24000 measured samples were recorded from the system at a sampling frequency of 500Hz. The resonance of the system was at a frequency of 16Hz. In line with section 5.5, the data was divided equally into training, validation and test data sets, in preparation for fitting a discrete-time model to the input-output data.

As the ARX model is linear-in-the-parameters (for the one-step-ahead prediction case), the optimum coefficients for a given number of input and output lags can be found by using the pseudoinverse approach. For each combination of input and output lags up to four of each, the optimum regressor coefficients were found with respect to the training data set. Each of these models was then used in a long range prediction on the previously unseen validation set, the MSE results of which are shown in Table 7.1 .

		Output Lags			
		1	2	3	4
Input Lags	1	92.0916	0.9583	0.0490	0.0083
	2	96.4519	0.0100	0.0056	0.0057
	3	99.6925	0.0062	0.0052	NaN
	4	103.6112	0.0058	Infinity	NaN

Table 7.1: Long-range ARX MSE values for original crisp data set

As can be seen from Table 7.1, for the use of at least two output lag terms it was relatively easy to obtain an extremely good prediction for the data. A small number of these models gave unstable long range predictions however, as can be seen by the NaN and infinity values returned. NaN is a quantity used in Matlab and Java to denote “not a number”, and in this case indicates that the solution has become unstable and thus

unable to return a real number for the model error value. In order to make the system more of a challenge, noise was added to the measured outputs in different amounts. This noise was Gaussian distributed with zero mean, and so a number of new data sets were generated for the system. $n\%$ output noise signifies that the noise added to the measured outputs had a standard deviation that was $n\%$ of the original data standard deviation.

Table 7.2 shows the long range validation MSE results for 6% added output noise. As with the original “clean” data sets, the regressor coefficients of the models were found by using the pseudoinverse approach for the training set one-step-ahead outputs.

		Output Lags			
		1	2	3	4
Input Lags	1	90.9277	63.7377	36.3432	29.8930
	2	92.8368	58.0626	21.2457	6.5752
	3	95.3397	57.0087	21.4767	5.8002
	4	98.6355	56.1415	21.4902	6.0170

Table 7.2: Long-range ARX MSE values for 6% added output noise

For the interval modelling done here, a crisp ARX model was initially required for the data. After consulting Table 7.2, the model with two input lags and four output lags was selected. At 6.6%, the long range validation set MSE of this model was reasonably close to the threshold that could be considered as “good”, so the system was a suitably challenging one to use. The exact form of this model was:

$$\hat{y}_i = 9.83 \times 10^{-7} u_i + 2.77 \times 10^{-6} u_{i-1} - 4.14 \times 10^{-6} u_{i-2} + 9.79 \times 10^{-1} y_{i-1} + 4.15 \times 10^{-1} y_{i-2} - 8.87 \times 10^{-2} y_{i-3} - 3.97 \times 10^{-1} y_{i-4} \quad (7.2)$$

where u refers to a system input term, y refers to a system output term, and each subscript regressor index relates to a particular time instant.

A time series plot of the validation set long range prediction from this optimum crisp ARX model is shown in Figure 7.1, alongside the target data for comparison. For ease of viewing purposes, only the first 500 of the 8000 data samples are shown.

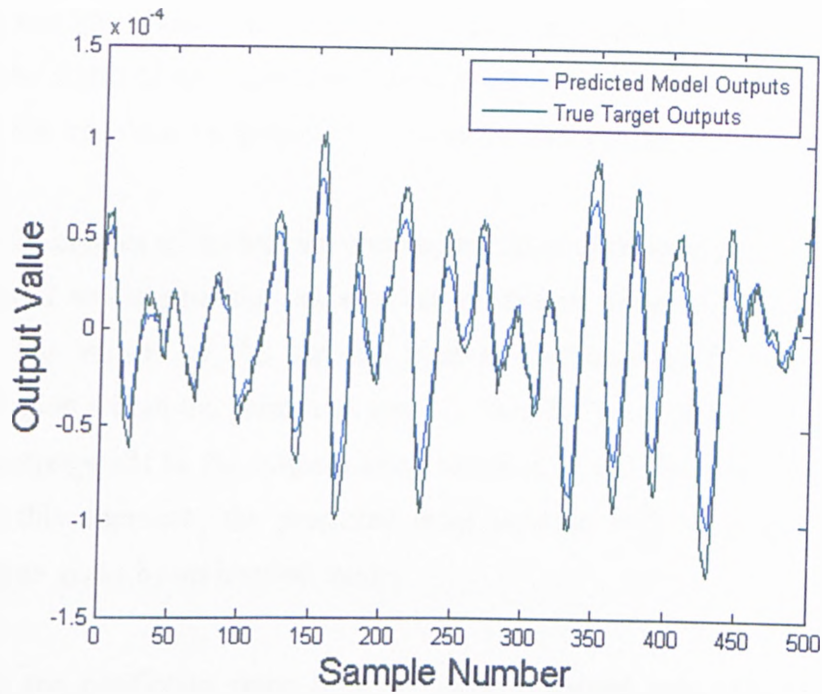


Figure 7.1: Long-range crisp ARX model prediction and target data

An uncertain ARX model was then created by using interval numbers for the regressor coefficients, with the same regressor terms as in equation (7.2). Whereas the conventional interval form is:

$$X = [a, b] \quad , \quad a \leq b \quad (7.3)$$

here each interval number was stored in an alternative “centre and radius” form. This can be illustrated by:

$$[a, b] \rightarrow \frac{(a+b)}{2} \pm \frac{(a-b)}{2} \quad , \quad a \leq b \quad (7.4)$$

The centre of each interval regressor coefficient was held fixed at its original crisp value in equation (7.2), so that only the interval radii values were being trained. As the widths of the ARX regressor coefficients become wider, this means that more uncertainty is being introduced into the model. Intuitive reasoning states that this will lead to a wider range of possible outputs from the model, so the predicted output bounds will also become wider.

The ideal interval model solution is one for which all of the target data is contained within the predicted model bounds, and the range between these bounds is of the

minimum possible width. For Figure 7.1, this corresponds in graphical terms to increasing the width of the output prediction at all the points where the predicted curve strays from the true data, i.e. generally at the target data peaks.

By holding the centres of the interval coefficients fixed at their original crisp values, this can be viewed as “expanding” an interval prediction about the initial crisp model prediction. The ranges of the various interval regressor coefficients will form a hypercube region within the parameter space, where the point at the very centre of this hypercube corresponds to the original crisp solution. It can therefore be expected that when using this approach, the predicted crisp solution will lie at the centre of the bounded region given by an interval model.

To calculate the prediction error from an interval-valued autoregressive model, an alternative approach is needed to the conventional MSE measure in equation (5.11). The error per sample for an interval model can be given by:

$$E_i = \begin{cases} 0 & , \hat{y}_{i,low} \leq y_{it} \leq \hat{y}_{i,high} \\ (\hat{y}_{i,low} - y_{it})^2 & , y_{it} < \hat{y}_{i,low} \\ (y_{it} - \hat{y}_{i,high})^2 & , y_{it} > \hat{y}_{i,high} \end{cases} \quad (7.5)$$

where $\hat{y}_{i,low}$ and $\hat{y}_{i,high}$ are the respective predicted lower and upper bounds for that sample, and y_{it} is the target sample output. This interval error function is analogous to the error function used in support vector machines (Cristianini & Shawe-Taylor, 2000).

Using equation (7.5), the MSE for an interval model can then be given by:

$$MSE = 100 * \frac{\sum_{i=1}^N E_i}{Nv} \quad (7.6)$$

where N is the number of samples and v is the variance of the target output data.

In order to obtain the best set of regressor coefficient interval radii there are two different approaches that can be used – the “proportional radii expansion” approach or a differential evolution training approach.

7.2 Proportional Radii Expansion

Each interval regressor coefficient within the model can be written in the form:

$$w_j = w_{jc} \pm w_{jr} \quad (7.7)$$

where w_{jc} and w_{jr} are the respective centre and radius values of the interval.

In the proportional radii expansion approach, every interval coefficient radius in the model is set as the same relative proportion ρ of its respective central value. i.e:

$$w_j = w_{jc} \pm \rho w_{jc} \quad (7.8)$$

Using this proportional radius form maintains the relative magnitudes of the coefficients with respect to each other. From viewing equation (7.2) it can be seen that the regressor coefficients of the initial crisp ARX model differ by several orders of magnitude. If the same absolute quantity was used for the radius of each interval regressor coefficient, this radius would mean a far greater relative width for some interval coefficients than for others.

After a suitable initial crisp model has been found for the data, this proportional radii expansion approach increases the value of the radius proportion ρ by successive regular increments. All of the interval regressor coefficients in the model are therefore being widened (or “expanded”) together, and by the same relative proportion. For each different ρ increment the bounds of the interval ARX model can be calculated, along with the long range prediction MSE on the relevant data set. As described in chapters 2 and 3, the long range bounds for an uncertain autoregressive model can only be calculated using one of the described possibilistic sampling methods.

When used on the test data set, the original crisp ARX solution gave a long range MSE value of 7.0%. Four different methods were used to calculate the output bounds for each

proportional expansion increment of the radii values. These were the two-level full factorial vertex method, the three-level full factorial vertex method, the reduced transformation method and the modified reduced transformation method, all described in chapter 3. For the two transformation methods, an arbitrary ten α -cut levels were used to discretise the parameter space, excluding the original central crisp solution point. For each of these four different possibilistic sampling methods, the same set of proportional radius increments was used.

7.3 Differential Evolution Radii Training

The proportional radii expansion approach of section 7.2 is relatively quick to implement, although this is dependent on the number of uncertain parameters in the model and the type of possibilistic sampling method used to calculate the interval model bounds.

The disadvantage that the approach has, however, is that it does not consider the relative importance of each individual regressor coefficient. Although a crisp ARX model may have a low prediction error, the majority of this error may be due to one particular incorrectly valued regressor coefficient. In order to reduce the error to a suitable value, the other regressor coefficients would have to be widened enough to overcome the contribution of this “bad” coefficient. They may therefore be widened more than is necessary, leading to wider model output bounds than could be obtained otherwise.

A more beneficial approach would be to train the interval regressor coefficient radii independently of each other, through the use of an optimisation algorithm. The differential evolution algorithm of section 4.2 would be ideal for this. It is designed for real number values, which the regressor coefficient radii are, and is a global optimisation approach, so is able to avoid local minima in the parameter space. As the goal here is to obtain an interval-valued ARX model with a low error, this is a minimisation problem and not a maximisation one. The aim is therefore to minimise the cost of the solution population, not to increase its fitness.

As with all optimisation algorithms, an objective function is needed to assess the quality of each solution point. A good solution for an interval-valued ARX model is one which suitably satisfies the trade-off between minimising the prediction error of the model and minimising the widths of the output bounds, to reduce the tolerance on the solution. The

differential evolution objective function used here needs to take this trade-off into account.

If this objective function was solely to consider the long range MSE of the interval network, the model bounds would deliberately be set very wide so that all the target data was contained within, and a zero error value would be returned. Once a particular coefficient radii solution was found with a zero MSE value, it could never be improved within the population, and would therefore remain there forever. It is likely that this first instance of a zero solution for each population vector would not be a very tight fit to the target data. The end result would be a population of radii solutions with wide bounds, and thus very vague predictions for the data.

The addition of a width penalty term to the objective function solves this problem. As the width of the output bounds is directly linked to the widths of the interval regressor coefficients, an additional term can be used to penalise these coefficient widths. As with the proportional radii expansion approach, it is the relative proportion of the interval coefficient radii that is important, and not their absolute values. Incorporating the width penalty term, the cost of a particular population solution k for a model with M regressor terms can be given as:

$$Cost_k = MSE_{INTERVAL,k} + \beta \sum_{j=1}^M \left| \frac{w_{jr}}{w_{jc}} \right| \quad (7.9)$$

The width penalty term therefore uses the sum of the relative coefficient radii proportions. β is a positive real number value that alters the magnitude of the width penalty term. It allows the user of the model to specify the emphasis placed on minimising the widths of the output bounds.

As interval-valued ARX models such as these could conceivably be used for any SISO time series application, this variable width penalty is important. Certain applications of an uncertain ARX model may require a high solution precision, so a high β value would be recommended. For other applications a much higher tolerance on the solution may be sufficient, so a lower value of β can be used.

For differential evolution training of the regressor coefficient radii, the algorithm parameters $F = 0.5$, $CR = 0.5$ and $NP = 35$ were used. This corresponded to using a

population size equal to five times the number of parameters to be trained; a general rule of thumb for differential evolution.

The width penalty coefficient β was varied in regular increments. Due to the stochastic nature of the differential evolution algorithm, ten tests were ran for each β value. Two different criteria were used to terminate the algorithm in each test, one of which was a maximum iteration number of 1000. The second termination criterion involved the decay of the total population cost. As the algorithm uses floating point numbers, the total population cost would take an exceptionally long time to decay to constant for all decimal places. To all intents and purposes it may be considered as converged well before this. The second criterion therefore used the relative decay rate of the total population cost over ten successive iterations. If this rate fell to below 0.0001, the algorithm was terminated and the best solution selected.

The problem with differential evolution interval radii training is the large computational expense that it requires. It would be much more beneficial to train the interval-valued ARX model based on its long range predictions, but due to computational expense this is simply not possible. As outlined in chapters 2 and 3, a possibilistic sampling method has to be used to calculate long range interval bounds. For this particular small ARX model with 7 uncertain parameters, the two level vertex method (the simplest of these methods) uses 128 hypercube sample points. This gives 128 predictions of an 8000 sample time series for each population target vector, 35 of these target vectors per iteration, an additional 35 similar trial vector predictions per iteration, and hundreds of iterations per test. There are then 10 tests for each β increment, and a suitably large number of these increments are needed to generate the results. In regard to the required computational time, this is well beyond practical limits. This excessive computational time therefore limits the differential evolution algorithm to the optimisation of interval one-step-ahead outputs only, with the use of interval arithmetic.

For each different incremental value of β , the ten tests were used to optimise the one-step-ahead interval outputs in conjunction with the width penalty term. Each converged coefficient radii solution was then used for a long range prediction of the validation data set, returning the interval MSE only. This error value was calculated separately using the same four possibilistic sampling methods used previously in the proportional radii expansion approach. For each different possibilistic sampling method, the solution with

the lowest value of validation set interval MSE was selected as the best solution, and then used for a long range model prediction of the test data set.

7.4 Results and Discussion

For the proportional radii expansion approach, Figure 7.2 shows the relationship between the proportional expansion coefficient ρ and the long range test set interval MSE, for each of the four different possibilistic sampling methods used here. The curves for the three-level vertex method and the modified reduced transformation method are extremely similar, to the point where the former is masked by the latter.

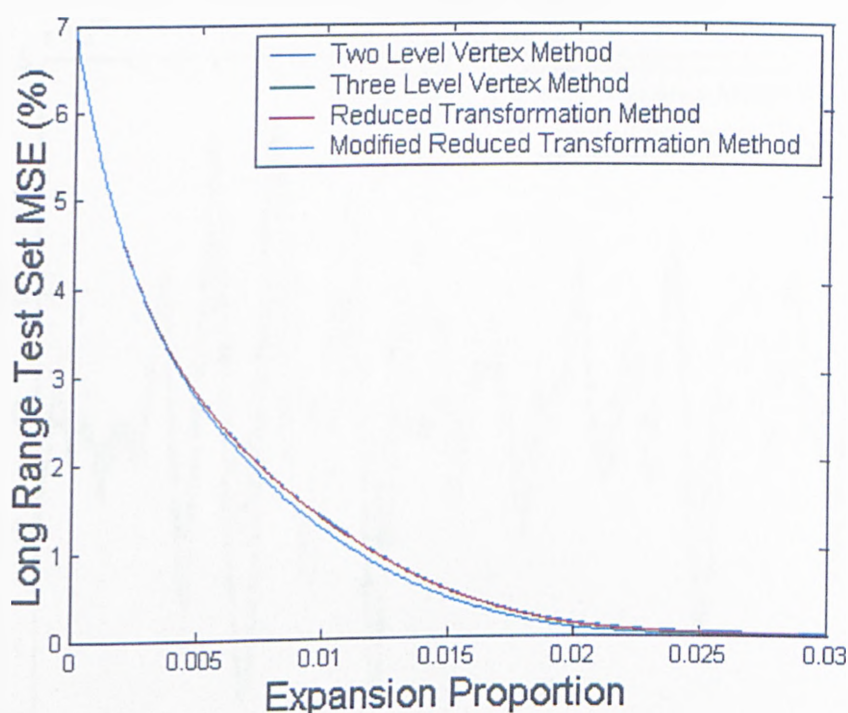


Figure 7.2: Proportional expansion long-range MSE values

All four curves show that as the interval ARX regressor coefficients are expanded about their crisp values, the long-range prediction error of the model falls monotonically. Increasing the width of the regressor coefficients introduces more uncertainty into the model, and creates a corresponding increase in the widths of the model output bounds. These bounds then contain more of the target output data, and so the prediction error of the model falls.

If the proportion ρ of regressor coefficient expansion is kept small, the resulting parameter hypercube can be viewed as a local region of the parameter space around the original crisp ARX coefficient solution. The four different possibilistic sampling

methods used have different capabilities in detecting non-monotonicity of the output surface. As the four curves of Figure 7.2 are relatively similar, this suggests that at each different coefficient expansion increment, there is very little output surface non-monotonicity over that specific parameter hypercube. Within the parameter hypercubes generated here at each step, the output surface is almost monotonic.

Figure 7.3 shows an interval model prediction obtained from the proportional expansion approach. The interval bounds were estimated by the reduced transformation method with 11 α -cut levels, and this prediction has a long range interval MSE value of 1% on the test set. For clarity of viewing, only the first 500 samples of the full 8000 are shown.

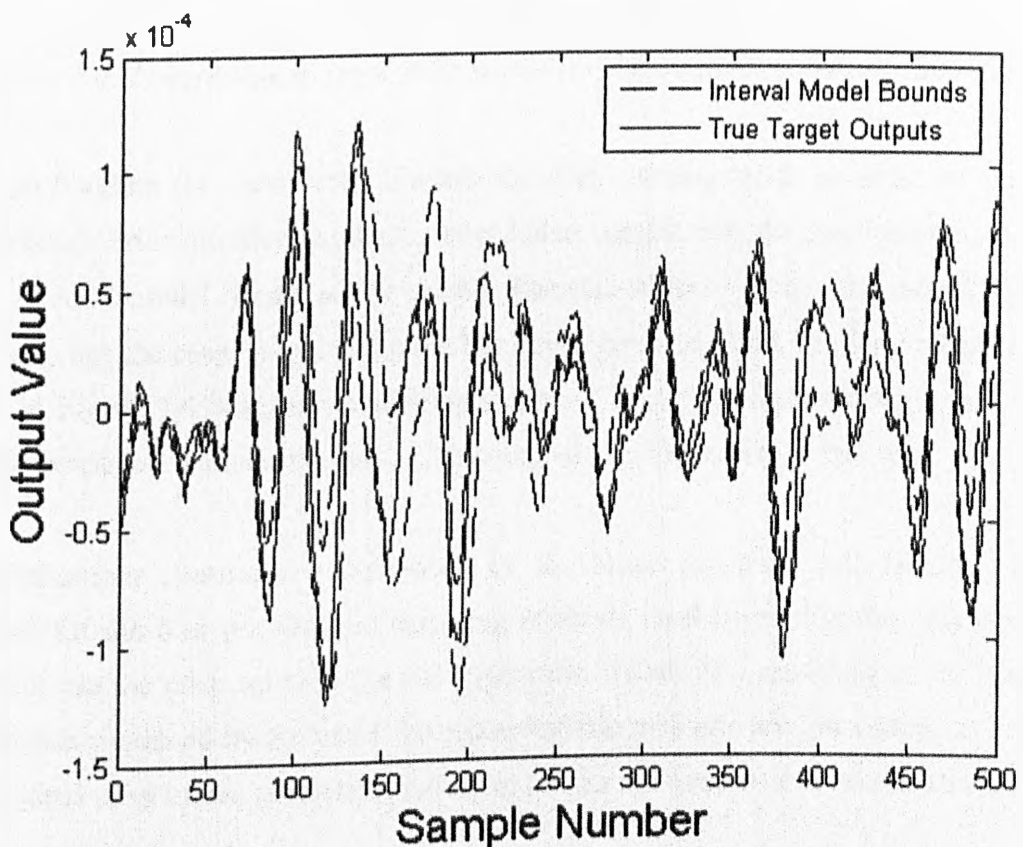


Figure 7.3: 1% long-range MSE solution obtained from proportional expansion

For the differential evolution coefficient radii training, Figure 7.4 shows the relationship between width penalty coefficient β and long range test set interval MSE, for the two-level vertex method. It can be seen that this plotted curve is very erratic. At low β values the prediction bounds become unstable and the model returns an error value of NaN, as can be seen by the absence of the plotted curve in this region.

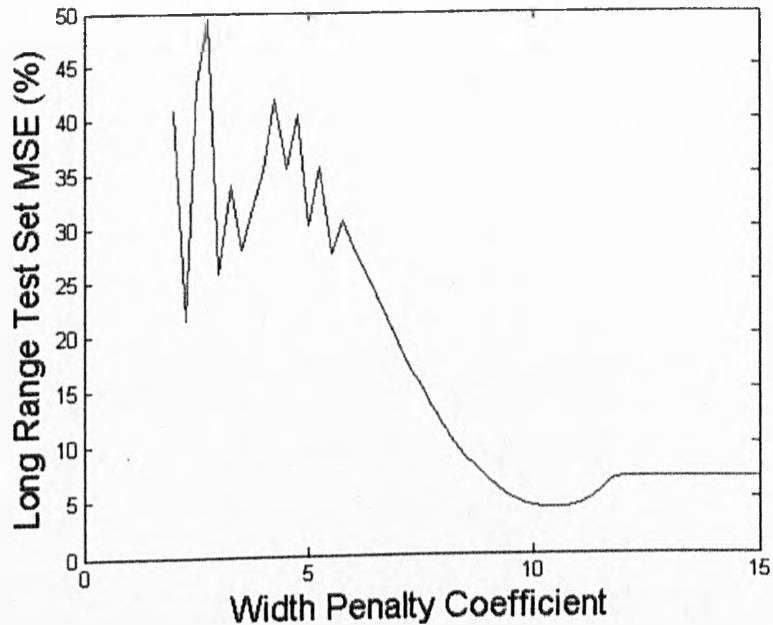


Figure 7.4: Long-range interval MSE values for differential-evolution-trained models

At high β values the curve tends towards the crisp solution MSE, as could be expected beforehand. Prior intuitive reasoning would also suggest that the prediction error of an interval ARX model should not be greater than that of its original crisp model, but this is clearly not the case. When using the two-level vertex method, the long range test set MSE in Figure 7.4 becomes significantly greater than the original crisp model error, which completely violates the reason for using an interval model in this way.

This behaviour illustrates a deficiency of the basic two-level full factorial vertex method. Of the four possibilistic sampling methods used here, it is the only one that does not use the crisp solution (i.e the hypercube centre) as a sampling point. Whereas the bounds produced by the other three possibilistic methods are guaranteed to contain the original crisp model prediction, this is untrue for the two-level vertex method.

A low width penalty coefficient β means that a low penalty is placed on the widths of the interval regressor coefficients. The coefficient radii values returned by the differential evolution algorithm will thus be relatively large here, and hence so will be the resulting uncertain parameter hypercubes. The outer vertices of these hypercubes will be far from the crisp solution point, so the corresponding model prediction curves will likely be significantly different from it. Thus for a low β value with the two-level vertex method, not only could the model prediction bounds not contain the crisp solution prediction, but they could be also located far from this crisp prediction. A particular interval solution exhibiting this behaviour is shown in Figure 7.5.




Figure 7.5: Vertex two-level ARX long-range bounds for a differential-evolution-trained model with low β value

The relationship between the differential evolution width penalty coefficient and the interval long range MSE of the model on the test set is shown for the remaining three possibilistic sampling methods in Figure 7.6. For the reduced and modified reduced transformation methods, an arbitrary 11 α -cut levels were used for each, including the crisp central point. The predicted curves for these two transformation methods are very similar, so they cannot be distinguished from one another on Figure 7.6.

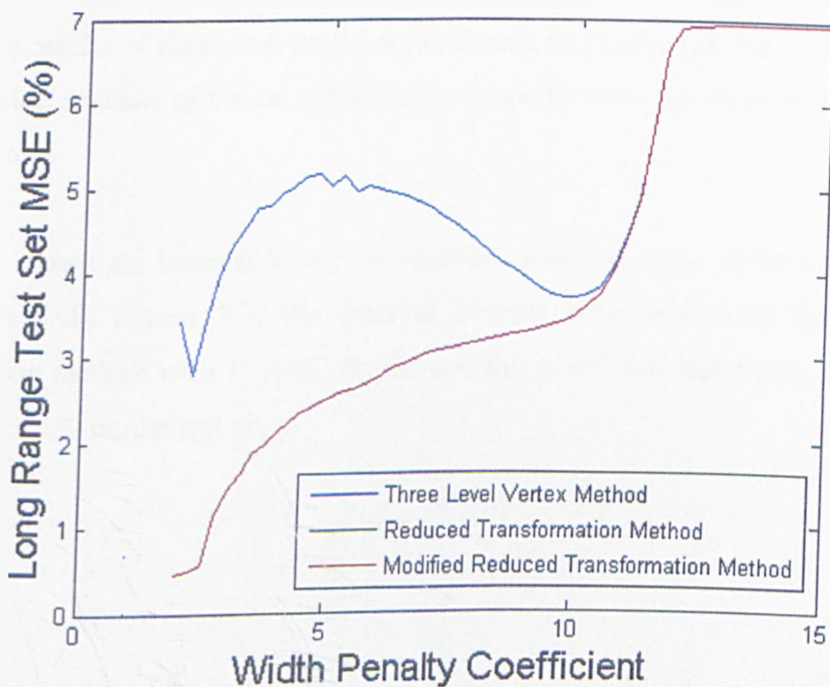


Figure 7.6: Differential-evolution-trained long-range MSE values

As with the two-level vertex method in Figure 7.4, these three possibilistic methods are also all unstable at low β levels. Unlike the two-level vertex method however, the prediction error of the interval ARX model does remain below the crisp solution value for all three of these methods. For the three-level vertex method, the relationship between β and interval model error is not monotonic however, which would be the ideal situation.

As already mentioned in section 7.3, an important feature of interval-valued autoregressive models is that they allow a trade-off between model prediction error and the prediction width. A monotonically increasing relationship between β and model error leads to a unique error value for each selected β value, enabling any user of the model to easily select a suitable trade-off level for their application. For this reason, the reduced or modified reduced transformation method seems appropriate for assessing the differential evolution solutions obtained here.

The differences between the curves in Figures 7.4 and 7.6 suggest that, unlike the proportional expansion approach, the output surfaces here are non-monotonic with respect to the uncertain parameter ranges returned. The three level vertex method is more effective at detecting surface non-monotonicity than the two level method, but still only considers the hypercube central point and a number of points on the outer surfaces. The reduced and modified reduced transformation methods consider a large number of sample points within the interior of the hypercube, which affects the bounds that are predicted, and hence the interval error of the prediction. Due to the extreme similarities between the results of these two methods, as shown in Figure 7.6, the modified reduced transformation method offers no significant advantage over the reduced method in this particular case.

Figure 7.7 shows an interval model prediction obtained from differential evolution training. As with Figure 7.3, the interval bounds were estimated by the reduced transformation method with 11 α -cut levels, and this prediction has a long range interval MSE value of 1% on the test set.




Figure 7.7: 1% long-range MSE solution obtained from differential evolution training

Figures 7.3 and 7.7 allow a direct comparison to be made between the proportional expansion approach and differential evolution training for this particular problem. It can be seen from these two figures that there is a staggering difference between the widths of the predicted model bounds. The bounds in Figure 7.3 are reasonably tight in relation to the target data, but this certainly does not apply to Figure 7.7. The differential evolution bounds in Figure 7.7 are very wide and do not fit the target data with any reasonable degree of precision.

Figure 7.7 is therefore the first indication that the differential evolution approach has not been effective for this studied system. Figure 7.6 shows that there is a monotonic relationship between the width penalty coefficient and the interval prediction error when the transformation method is used, but gives no indication of the prediction bound widths in each case.

As outlined in section 7.3, it is not possible to conduct differential evolution training of the coefficient radii based on the model long range outputs because of the computational expense involved. The assumption was therefore initially made here that optimising the one-step-ahead predictions using interval arithmetic would also optimise the long range

predictions using a possibilistic sampling approach. Figure 7.7 shows that this assumption does not hold for this uncertain ARX system, as the differential evolution training has failed to provide a tight set of long range prediction bounds in relation to the target data.

Both the proportional radii expansion approach and the differential evolution training approach used here can be viewed as implementations of the “opportunity” concept in Ben-Haim's information gap theory. In both cases, the windfall reward is the desired long range interval MSE of the model prediction. The opportunity level is the minimum uncertainty that needs to be introduced into the model (in the form of interval regressor coefficients) for this windfall reward to be obtained.

Figure 7.2 and the two transformation method variants in Figure 7.6 all show that there is a continuous monotonic trade-off available between the long range prediction error of the model, and the widths of the model output bounds. By selecting a suitable value of either the expansion proportion ρ or the width penalty coefficient β , the user of these techniques can select the correct level of this trade-off for their particular application of an uncertain ARX model. As shown in Figures 7.3 and 7.7, however, the proportional expansion approach is the more effective of the two in this case, despite being much less computationally expensive than the differential evolution training method.

8. On Replacing a Nonlinear System with an Uncertain Linear

Model

8.1 Introduction

Within control engineering, equivalent linearisation of a nonlinear system is a well-known and long-established technique. Real-life systems are generally nonlinear, but many controllers require a linear model in order to be used. Linearisation provides this model by giving a linear approximation to the system in the local region of a particular operating point. The disadvantage of the approach is that the model is only effective close to this operating point. If the system operating conditions stray significantly from this regime, the model predictive capabilities will progressively worsen.

In this chapter, the concept of equivalent linearisation is advanced by using an uncertain linear model. A nonlinear model is able to give a good prediction for the system over a wide range of operating points, so is not subject to the local constraints of a linear model. The aim here is to obtain a model that has the same potential range of operation as a nonlinear model, but with the modelling simplicity associated with a linear model. Using an uncertain linear model enables this to be achieved.

The system being modelled here is a simulation of a symmetric Duffing oscillator, given by:

$$m\ddot{y} + c\dot{y} + ky + k_3y^3 = F(t) \quad (8.1)$$

with $m = 1\text{kg}$, $c = 20\frac{Ns}{m}$, $k = 10^4\frac{N}{m}$ and $k_3 = 5 \times 10^9\frac{N}{m^3}$. This type of system is also often referred to as a cubic stiffness system, and its operation is well known. It has been chosen here because it contains a nonlinear component - the cubic stiffness term, but does not have too high a degree of complexity. If this nonlinear term is removed, the result is the standard mass-spring-damper system.

The data was simulated by a Fortran routine that used a fourth-order Runge-Kutta integration scheme and a sampling frequency of 250Hz. This routine enabled the input RMS value of the data to be specified. A Gaussian white noise excitation was used, but

then low-pass filtered to ensure that the integration scheme remained stable. Several different Duffing data sets with varying levels of input RMS were generated, where each data set contained 24000 sample points. These sets were each equally divided into a training, validation and test set, in line with the standard practice for training time series models.

8.2 Crisp Models

The work here examined the effects of varying the operating point for time series models, where this operating point was given in terms of the system input RMS value. A "base" operating point was needed for use as a reference point. By using the Fortran routine, a base data set was generated with an input RMS value of 2.37, after filtering of the data had taken place.

An initial nonlinear model was given for the system in the form of an NARX model, and this was then subsequently replaced by a linear ARX model. For an NARX model the base data set was not a particularly challenging system to model, so it was not necessary to use the FROE or SEMP algorithms. A number of arbitrary regressor combinations were selected, and their optimum coefficient values were found for the base training set by using the pseudoinverse approach. Each of these models was then tested on the previously unseen base validation set. Several very good models were found, and the one selected for the work here was:

$$\hat{y}_i = 1.97 \times 10^{-6} u_i + 1.17 \times 10^{-5} u_{i-1} + 1.84 \times 10^{-6} u_{i-2} + 1.77 y_{i-1} - 0.92 y_{i-2} - 70619.0 y_{i-1}^3 - 2669.0 y_{i-2}^3 \quad (8.2)$$

When used on the base validation set, this NARX model gave a long range validation set MSE of 0.04%, an extremely good prediction value. This, of course, raises the question as to why an uncertain linear model is necessary here, when an NARX model can give such an excellent prediction. The purpose of this chapter is to illustrate the use of an uncertain linear model for a simple example system. The technique can then be applied to other nonlinear systems, where the presence of much greater nonlinear components means that the initial nonlinear model doesn't give such a low prediction error.

The approximate linearisation of the system here was achieved by removing the final two cubic terms from equation (8.2). The regressor coefficients were retrained on the training set using the pseudoinverse approach, to give the crisp ARX model:

$$\hat{y}_i = 1.97 \times 10^{-6} u_i + 1.17 \times 10^{-5} u_{i-1} + 1.82 \times 10^{-6} u_{i-2} + 1.76 y_{i-1} - 0.92 y_{i-2} \quad (8.3)$$

As could be expected, this linear model was less effective at representing the system than the original nonlinear NARX model. When used on the base validation set, the model gave a long range prediction error of 3.1%. The work then splits into two distinct sections, labelled as "crisp ARX input scaling" and "uncertain ARX modelling".

8.3 Crisp ARX Input Scaling

In this section, the crisp ARX model in equation (8.3) was used to model the system. A large number of data sets were generated by the Fortran routine, covering a large range of input RMS values. Using a scaling technique for the system inputs, the effects of changing the system input RMS value were examined.

For each different data set A used, the inputs of that SISO set were scaled by the factor μ , where

$$\mu = \frac{RMS_{base}}{RMS_{data}} \quad (8.4)$$

The resulting data set then had an input RMS value equal to that of the original base data set. Once the entire set of long range model outputs had been calculated, these outputs were scaled by the reciprocal of μ , and the MSE value was calculated. Over an incremental range of input RMS values, the long range prediction errors of the model were calculated when this input scaling approach was used.

8.4 Uncertain ARX Modelling

The crisp ARX model in equation (8.3) was then transformed into an interval-valued model, by using interval numbers for some or all of the regressor coefficients. The regressor terms remained the same. As in the previous chapter, each of these interval coefficients was written in the centre and radius form given in equation (7.4).

Two different capabilities of interval-valued ARX models were then assessed here. The first of these was the robustness of the model to changes in input RMS value, and the second was robustness to system output noise.

8.4.1 Variable Input RMS

Changing the system input RMS value is analogous to changing the operating point of the linear model. It was wished to examine the robustness of an interval-valued linear model to changes in system operating point, as opposed to the conventional crisp linear model.

The input RMS value of the data was changed by adjusting the input parameters of the Fortran routine, so this also affected the measured outputs of each data set. As changes in input RMS value therefore affected every one of the five regressor terms in the crisp ARX model of equation (8.3), it was decided to introduce uncertainty into all of the regressor coefficients for these terms.

The centre of each interval regressor coefficient was held fixed at its original crisp value in equation (8.3), and the proportional radius form of equation (7.8) was used. The radius proportion ρ was increased in regular successive increments, which had the effect of widening the interval regressor coefficients of the model. The long range prediction bounds of the model were calculated using the modified reduced transformation method with 11 α -cut levels, since the model only contained five input variables. The MSE of this long range prediction was given by the interval error function in equations (7.5) and (7.6).

For each ρ expansion increment, a number of different Duffing data sets were used by the model, each with a differing value of input RMS. It was calculated the range of input RMS that could be tolerated by each proportional coefficient expansion ρ , in order to give an interval long range prediction error of up to 1% on the respective test data set. As already stated in section 5.5, 1% MSE can be judged as a very good model prediction.

8.4.2 Variable Output Noise

As stated in section 8.2, the “base” input RMS level here was considered to be 2.37. When examining the effects of varying output noise, zero-mean Gaussian noise was added to the outputs of this base data set, so that the inputs remained unchanged. This therefore only affected the final two regressor terms in the crisp ARX model of equation (8.3). $n\%$ output noise refers to the added Gaussian output noise having a standard deviation that is $n\%$ of the base set output standard deviation.

Due to only two uncertain regressor coefficients being used here, the computational expense was much lower. This enabled the general transformation method to be used; the operation of which was detailed in section 3.5. 11 α -cut levels were used for each of the two uncertain regressor coefficients, including the crisp central point corresponding to $\alpha = 1$.

As with section 8.4.1, the centres of the two uncertain regressor coefficients were held fixed at their crisp values in equation (8.3), and the proportional radius form of equation (7.8) was used. The two coefficient radii were expanded in successive increments using the expansion proportion ρ . For each different ρ increment, different percentage levels of Gaussian noise were added to the outputs of the base test data set. The percentage level of added noise was calculated that could be tolerated by each ρ increment in order to still give an interval MSE of up to 1%. The interval MSE function of equations (7.5) and (7.6) was again used here.

8.5 Results and Discussion

Figure 8.1 shows the relationship between the input RMS value of the data, and the long range test set MSE of the model prediction when using with the crisp ARX input scaling technique. When used on the base test set, the crisp ARX model gave a long range MSE value of 1.9%. As already stated, a standard crisp linear model becomes less effective as its operating point moves further from its original linearisation point. This agrees with Figure 8.1, where the minimum long range prediction error of the model lies close to the base input RMS value of 2.37. One interesting feature of the curve is that its minimum is not exactly at this base RMS value however, but at an input RMS value of 2.12. Therefore, by using the crisp ARX model at an input RMS slightly lower than its

linearisation RMS value, the prediction error of the model can actually be improved slightly.




Figure 8.1: Input RMS vs long-range MSE for "crisp ARX input scaling" technique




Figure 8.2: Crisp ARX prediction from "crisp ARX input scaling" technique

Figure 8.2 shows a model prediction for an input RMS value of 1.34, where a long range test set MSE of 4.0% was obtained. It shows that by using this crisp ARX input scaling method, a good model prediction can still be obtained if the input RMS is within a reasonable distance of the linearisation RMS value. Here the input RMS of the system differs from the operating point by approximately a unit RMS value, but a long range prediction error of 4% MSE is still able to be obtained by the model.

Figure 8.3 shows the long range test set MSE values obtained from the NARX model in equation (8.2), for the same range of system input RMS values as in Figure 8.1. For this NARX model, no scaling of the model inputs or subsequent unscaling of the model outputs was used. As stated previously, a nonlinear model is able to give a good prediction over a range of input RMS values, so does not have the same local requirements as a linear model. Even at the upper limit of the operating range in Figure 8.3, the model is still able to give an exceptionally good long range prediction error of approximately 0.3% MSE.




Figure 8.3: Input RMS vs long-range MSE for crisp NARX model with no input scaling

For the uncertain ARX model of section 8.4.1, the relationship between coefficient expansion proportion ρ and tolerable system input RMS range is shown in Figure 8.4. This tolerable RMS range is the system input RMS for which the interval ARX model is able to give a long range MSE of up to 1% on the test data set.




Figure 8.4: Interval ARX expansion proportion vs tolerable input RMS range for 1% MSE

The original crisp ARX model gave a test set long range MSE value of 1.9%, so some expansion of the regressor coefficients was necessary to obtain a prediction error of 1% MSE. Figure 8.4 shows that if the coefficient expansion proportion ρ is less than 8×10^{-4} , it is not possible to obtain a 1% MSE model prediction, regardless of what the input RMS value is. The remainder of the figure then shows that as the expansion of the regressor coefficients increases, a greater range of system input RMS can be tolerated by the interval-valued ARX model.




Figure 8.5: Interval ARX expansion proportion vs tolerable output noise for 1% MSE

Figure 8.5 shows the relationship between ρ and maximum tolerable output noise from section 8.4.2. The figure shows the level of added Gaussian output noise that can be tolerated by the interval model in order to still obtain an interval MSE of up to 1%. As with Figure 8.4, some initial level of regressor coefficient expansion is necessary for 1% MSE to be obtained, as the long range MSE of the initial crisp model is 1.9% on the test data set. It can be seen from Figure 8.5 that as the coefficients of the two output lag regressor terms are expanded further, a greater range of added output noise can be tolerated.

Figures 8.4 and 8.5 are both examples of the “opportunity” concept in Ben-Haim’s information-gap theory. The degree of uncertainty parameter here is the regressor coefficient expansion proportion ρ , and the windfall reward is a long range interval prediction MSE of 1%. Figure 8.5 illustrates that it is not obligatory for all of the ARX regressor coefficients to be uncertain; some of them may be kept as fixed crisp values.

As in chapter 7, increasing the widths of the ARX regressor coefficients introduces more uncertainty into the model, which means a greater range of possible model outputs, and thus wider output bounds. These output bounds will contain more of the target output data, and thus the prediction error of the interval model will fall. Although Figures 8.4 and 8.5 show that introducing more uncertainty into the interval model allows a greater robustness against both input RMS variation and output noise variation, this comes at the cost of wider prediction bounds. The trade-off between minimising the prediction error of the interval model and minimising the prediction widths has to be considered at all times. The user of this approach will need to assess the suitable level for this trade-off in relation to their particular application.

This chapter has therefore shown that it is possible to satisfactorily replace a nonlinear time series model with an uncertain linear model. As with a nonlinear model, the uncertain linear model is able to be used over a much wider range of system operating points than a crisp linear model. This operational range can be increased by introducing more uncertainty into the model coefficients, although this will cause a corresponding increase in the widths of the model prediction bounds.

9. An Application of Interval-Valued Neural Networks to a Regression Problem

9.1 Background work

In this chapter an interval-valued neural network was used to obtain a regression model for an engineering system. The work here uses results from a background paper by Parlitz et al (2004), in which a number of different modelling approaches were used for a pre-sliding friction process. A specifically designed experimental setup was used to represent the friction process, and this setup was used to generate experimental data. Some of the modelling approaches were physics-based and others used a “black box” approach, including the use of NARX models and NARX multilayer perceptron (MLP) neural networks. It is the results from the NARX MLP network that are used here.

Pre-sliding friction refers to the contact regime where an applied tangential force causes relative motion between two surfaces, but where some of the surface asperity contacts remain intact. It therefore occurs prior to gross sliding, and within the pre-sliding regime a fixed force will provide a fixed displacement, which is untrue for gross sliding. When the applied force is removed, the relative displacement between the two surfaces will not return to its original zero value, signifying a hysteretic loss. Whilst in the pre-sliding regime, the friction force between the two surfaces is a function of the displacement, and the relationship between the two is highly hysteretic.

Figure 9.1 shows a schematic of the experimental setup used in the background paper to obtain the pre-sliding friction data. The moving and fixed mirrors acted as a displacement sensor, and were located in line with the contact line. 24000 measured data samples were taken from this setup at a sampling frequency of 250Hz. This data took the form of a simple SISO time history of force vs displacement data pairs. Due to the hysteretic displacement-dependent force however, the aim was to obtain an inverse representation of the system, so that the force that created a given displacement value could be calculated.




Figure 9.1: Schematic of experimental setup used in the background paper

Figure 9.2 shows the input-output relationship of this inverse system, where the hysteretic relationship between displacement and friction force can clearly be seen. As mentioned in section 5.1, autoregressive models can be effective for systems that exhibit hysteresis, and therefore the selected MLP neural network model used autoregressive input nodes.




Figure 9.2: Hysteretic input-output relationship of the inverse system

In line with the recommended best practice for neural network training, the measured data in the background paper had been equally divided into three sets of 8000 samples,

respectively called the training, validation and test sets. After a suitable training procedure, the optimum architecture for this multilayer perceptron was found to be a two layer network with two input lags, three output lags and two hidden row nodes using hyperbolic tangent activation functions. This procedure considered only the training and validation sets, and a subsequent long range prediction of the network model on the previously unseen test set gave a MSE value of 6.5%.

One of the principal aims of this chapter is to see if the NARX MLP model from the background paper can be improved by the use of an interval-valued neural network.

9.2 Introduction

The training procedure for a multilayer perceptron network generally involves finding the optimum network architecture as well as the optimum adjustable weight parameters. In this work, however, the same network architecture was used as for the crisp NARX MLP network in the background paper. As with the background paper, the measured experimental data was equally divided into three sets here, labelled as the training, validation and test sets. Both the input and output values of the data were normalised, in order to improve the performance of the neural network model.

Although the architecture of the crisp NARX MLP network was unchanged from the background work, the weight parameters of that network were not automatically assumed to be optimal. Using the differential evolution algorithm, tests were undertaken to see if an improved weights solution could be found. These parameters were trained on the one-step-ahead outputs from the training data set, with algorithm parameters of population size $NP=85$, weighting factor $F=0.5$ and crossover constant $CR=0.1$. After several hundred training runs, an optimum weights solution was found that gave a long-range prediction on the validation set of 4.1% MSE, an improvement on the crisp NARX MLP network with 6.5% MSE from the background work.

As stated in section 5.5, 4.1% MSE is below the threshold that can be considered as a good model prediction. The rationale behind the work in this chapter was to see if the use of an interval-valued network could reduce this model error even further, possibly beyond the 1% threshold for a very good prediction.

Each network weight in the model was hence transformed into an interval number, where the centre and radius form of equation (7.4) was used. The centre of each weight was held fixed at its value within the optimum 4.1% MSE crisp solution, so that only the weight radii values were being trained here. By using interval numbers for the MLP network weights, this allowed uncertainty to propagate through the model, and an interval range was thus predicted as a model output. This used the same reasoning as chapters 7 and 8; that deliberately introducing uncertainty into a model can allow an improved model prediction to be obtained.

As with chapters 7 and 8, although the uncertain model parameters here took the form of interval numbers, interval arithmetic was not used to calculate the long range prediction bounds of the interval model. This is due to the bound explosion behaviour of interval arithmetic when used for any kind of uncertain autoregressive model, as detailed in section 2.1. Therefore a possibilistic sampling method was needed to estimate the output bounds from the interval network.

Chapters 7 and 8 both used versions of the transformation method to calculate the output bounds from uncertain time series models, but the models in those two chapters were both much smaller than the NARX MLP network used here. This network model contains 17 weights which, for the case of the interval training here, requires 17 interval weight radii to be trained. In relation to available computational time, this ruled out all possibilistic sampling methods except for the two-level full factorial vertex method. As detailed in section 3.1, this possibilistic method uses only the 2^n vertices of an n-dimensional uncertain parameter hypercube.

9.3 Interval Weight Radius Training

The most common way of training a multilayer perceptron network to a particular set of data is by the error back-propagation algorithm. For the case of an interval MLP network however, the network prediction error for a given sample cannot be differentiated backwards due to the interval weights, so the error back-propagation algorithm cannot be used here.

The prediction error of the uncertain MLP model was given by the same interval MSE

function used in chapters 7 and 8, where the error for a single sample is:

$$E_i = \begin{cases} 0 & , \hat{y}_{i,low} \leq y_{it} \leq \hat{y}_{i,high} \\ (\hat{y}_{i,low} - y_{it})^2 & , y_{it} < \hat{y}_{i,low} \\ (y_{it} - \hat{y}_{i,high})^2 & , y_{it} > \hat{y}_{i,high} \end{cases} \quad (9.1)$$

The total prediction error is then given by:

$$MSE = 100 * \frac{\sum_{i=1}^N E_i}{N \nu} \quad (9.2)$$

As with the uncertain ARX regressor coefficients in chapter 7, two different methods were used to train the interval network weight radii values here. The first of these was the proportional radii expansion approach, where each weight radius is given as the same proportion ρ of its respective central value, and this proportion is increased in successive incremental steps. This proportional expansion approach equates to expanding an uncertain hypercube around the point in the network weight space corresponding to the initial crisp MLP solution.

The proportional expansion approach is not computationally expensive since no actual optimisation of the interval weight radii is taking place. Its disadvantage, however, is that it does not consider the relative effects of each individual network weight. It may be that certain network weights are much more important than others in obtaining a good network solution. An optimisation approach may thus be advantageous in finding the optimum weight radii values for the interval network.

The second interval network training approach therefore used the differential evolution algorithm to optimise the weight radii values of the MLP network. As in chapter 7, the objective function of this algorithm needed an additional width penalty term for the interval weights, in order to minimise the width of the network output bounds. The same differential evolution objective function was used as in chapter 7, where the cost of a particular interval

weights solution is given by:

$$Cost_k = MSE_{INTERVAL,k} + \beta \sum_{j=1}^M \left| \frac{w_{jr}}{w_{jc}} \right| \quad (9.3)$$

where M is the number of network weights. β is referred to as the width penalty coefficient, and allows different emphasis to be placed on minimising the widths of the MLP network output bounds.

For the proportional expansion approach, the interval network weights were expanded about the original crisp MLP solution at regular increasing ρ increments. For each increment, the long range prediction error of the interval network was calculated on the test data set.

For the differential evolution training, a number of different width penalty coefficient values were used. Ten different runs of the algorithm were undertaken for each of these β values on the training data set. As with chapter 7, the computational expense of this approach meant that the network could only be trained on its one-step-ahead outputs by using interval arithmetic.

Each of the interval weight radii solutions for each β value were then used for a long range prediction on the previously unseen validation set. Each of these predictions was assessed only on the interval MSE contribution, so the width penalty term of equation (9.3) was not used in the second phase. The weight radii solution with the lowest long range validation set MSE was selected as the best solution, and then used in a long range prediction on the test data set.

9.4 Results and Discussion

For the proportional radii expansion approach, Figure 9.3 shows the relationship between the weight expansion proportion ρ and the test set long range interval MSE of the MLP network. As with the same method in chapter 7, it can be seen that expanding the interval network weights about the initial crisp solution causes the prediction error of the model to

fall monotonically.




Figure 9.3: Weight expansion proportion vs long-range MSE

Increasing the widths of the network weights introduces more uncertainty into the MLP network. This uncertainty propagates through the network, with the effect of increasing the output bound widths. As these output bound widths become wider they therefore contain more of the target data, and the prediction error of the model falls.

The two-level vertex method uses to estimate the network output bounds here is not guaranteed to return the true bounds, since the true relationship between the network weights and predicted outputs is not known in its entirety. By definition, however, the true network output bounds can only be wider than those estimated by the two-level vertex method, since all possibilistic sampling points are taken from within the uncertain weights hypercube. Wider network output bounds will give a lower prediction error, so the interval network errors obtained here can be viewed as upper bounds on the true error values.

Figure 9.4 shows an interval network prediction obtained by the proportional expansion approach, with a long range interval MSE of 1% on the test data set. Figure 9.5 shows an alternative 1% test set network prediction also obtained by the proportional expansion approach, but starting from a different crisp MLP solution with 9.4% validation set MSE.

By comparing these two figures it can be seen that if an inferior initial crisp weights solution is used for the interval network, the proportional expansion bounds to obtain a given interval MSE will need to be wider.




Figure 9.4: 1% long-range MSE proportional expansion bounds from 4.1% MSE crisp solution




Figure 9.5: 1% long-range MSE proportional expansion bounds from 9.4% MSE crisp solution

For the differential evolution training approach, Figure 9.6 shows the relationship between the width penalty coefficient β and the long range test set MSE of the interval network. It can be seen that as the width penalty coefficient increases, the prediction error of the interval MLP network increases monotonically. Increasing β causes a larger penalty to be placed on the interval weight widths during training, so narrower bounds will be obtained from the differential evolution algorithm. These bounds will thus contain less of the target data, and so the network prediction error increases. As β tends to high values, the prediction error of the interval network tends to that of the original 4.1% MSE crisp solution.




Figure 9.6: Differential-evolution-trained long-range MSE values

Figure 9.7 shows an interval network prediction obtained from differential evolution training, where the long range prediction MSE is 1% on the test data set. By comparing Figure 9.7 with Figure 9.4 it can be seen that for the same value of long range prediction error, the differential evolution training approach is able to give output bounds marginally narrower than those from the proportional expansion approach.




Figure 9.7: 1% long-range MSE differential-evolution-trained interval network

This is in contrast to the results from chapter 7, where the studied uncertain model in that chapter was also trained by both the proportional expansion approach and the differential evolution algorithm. In that chapter the prediction bounds from differential evolution were much wider than from proportional expansion, even though proportional expansion is the simpler of the two approaches. In both chapter 7 and here, the initial assumption was made that optimising the uncertain parameter radii by one-step-ahead interval arithmetic would also optimise the long range bounds obtained from a possibilistic sampling approach. In chapter 7 that assumption appeared not to hold, but in this chapter it appears to be true, as the long range prediction bounds in Figure 9.7 are tight in relation to the target data. It can be argued, however, that as the differential evolution algorithm has provided only a marginal improvement over the proportional expansion approach in return for a significant increase in computational expense, its use is not totally justifiable here.

Both of the two interval network training methods used here are examples of the “opportunity” concept in Ben-Haim's information gap theory being implemented. The “degree of uncertainty” parameter is given by the expansion proportion ρ and the width penalty coefficient β respectively. The windfall reward in this case is the acquisition of an interval network with a long range prediction error of 1% MSE, a value that can be

considered to be a very good prediction. Figures 9.3 and 9.6 show the required level of network uncertainty to obtain this error, or indeed any other arbitrarily selected MSE value.

As already discussed in chapter 7, interval-valued models such as the NARX MLP network used here allow a trade-off to be made between minimising the prediction error of the model and minimising the widths of the predicted output bounds. The level for this trade-off is application dependent, since some regression network applications may require a much higher solution precision than others.

Based on this, the ρ and β parameters allow a suitable level to be selected for this trade-off, in relation to the particular application of the interval MLP network. The monotonicity of Figures 9.3 and 9.6 means that a continuous relationship is available for this trade-off.

10. Uncertainty Propagation through Radial Basis Function Networks

10.1 Introduction

This chapter examines how uncertainty can propagate through radial basis function (RBF) neural networks for a regression application. As detailed in section 6.3, the radial basis function network is the second main type of neural network in general use, along with the multilayer perceptron network used in chapter 9.

The system being modelled here is the same pre-sliding friction process examined in chapter 9, where the data was obtained from a specially constructed experimental rig. This data takes the form of two SISO time histories with 24000 samples in total, and examines the dependence of the friction force on the measured displacement whilst within a pre-sliding friction contact regime. The data set inputs are therefore displacement values, and the aim is to predict the friction force at each time instant. The experimental data was once again equally divided into three separate sets, labelled as the training, validation and test sets.

As with chapter 9, an interval-valued RBF network is to be created by setting some of the network parameters as interval values. The effect of introducing this uncertainty into the network is to give a range of possible outputs, so the predictions from the model will take the form of interval ranges. Besides examining the effects of uncertainty within RBF networks, this chapter also provides an opportunity to compare the performances of MLP and RBF networks for the same data set.

An initial crisp RBF network solution was required here prior to interval network training taking place, and this was found with the help of the NETLAB neural network toolbox for MATLAB (Nabney, 2001). Unlike chapter 9, there was no prior information available for the optimum architecture of the radial basis function network, so this also formed part of the training process. Due to the hysteretic relationship between displacement and friction force within the data set, the RBF network used autoregressive input nodes.

Different arrangements of input lags, output lags and number of basis functions were systematically tested for the RBF network, with the optimum network parameters found for each individual arrangement. As detailed in section 6.3, the training process for each network architecture involved two separate stages. The first stage was to determine the optimum basis function parameters, and this was done by using the Expectation Maximisation algorithm.

The Expectation Maximisation algorithm created a Gaussian mixture model for the set of autoregressive input nodes used by that particular network architecture, thus determining the centre vector for each different basis function node. A diagonal covariance matrix was used for the basis nodes, so that each node used only a single width parameter σ . The largest squared Euclidean distance between any two basis centre vector elements was calculated, and all of the basis function width parameters were set to this value.

The second training stage then found the optimum network output layer weights by using the pseudoinverse approach. As autoregressive nodes were being used by the network, this pseudoinverse was therefore taken using the network one-step-ahead outputs. Radial basis function networks use an additional bias weight for each network output, so for this single output network application, the number of output weights equalled the number of basis nodes plus one.

After a systematic network training process, the optimum network input arrangement was found to be one input lag, one output lag and an additional input for the current time data input value. The optimum number of basis function nodes was found to be 11, giving a total of 56 adjustable parameters for the crisp RBF network. The training phase was done using the training set one-step-ahead outputs, and a subsequent long range prediction on the previously unseen validation set gave an error of 10.2% MSE.

10.2 Interval Network Training

The RBF network used here differs from the models in chapters 7, 8 and 9 in that it has significantly more adjustable parameters. In relation to computational expense it is not possible to train a network with 56 interval parameters. A single long range prediction using only the two-level vertex method would mean 2^{56} hypercube sample points, which is

not feasible with the available computational power

Based on this, uncertainty could only be introduced into some of the network parameters here, with the remainder staying as crisp values. It was decided to set the network output layer weights as interval values, giving a total of 12 uncertain parameters in the network. For the interval-valued MLP network in chapter 9, the 17 uncertain parameters present there meant that only the two-level vertex method could be used to calculate the interval network output bounds. As the RBF network here has less uncertain parameters than that, available computational time and power allow the reduced transformation method to be used with 6 α -cut levels. The three-level vertex method and modified reduced transformation method are both still too computationally expensive however.

As in previous chapters, the interval output weights of the network were given in the centre and radius form of equation (7.4). The centre of each weight was held fixed at its original crisp RBF solution value, so that only the output weight radii values needed to be trained here. The prediction error was once again calculated by the interval MSE function of equations (7.5) and (7.6)

The interval network training was done using two separate approaches – the proportional weight expansion approach and differential evolution weight radii training. Proportional weight expansion consisted of setting each output layer weight radius as some proportion ρ of its respective central value, and then successively increasing this proportion in regular increments. For each different expansion increment, the long range MSE of the interval RBF network was calculated on the test data set.

The second interval network training approach used the differential evolution algorithm as an optimisation technique for the output weight radii values. Like the previous differential evolution training approaches in chapters 7 and 9, an additional width penalty term was used in the objective function of the algorithm. The cost for a given population solution k was thus given by:

$$Cost_k = MSE_{INTERVAL,k} + \beta \sum_{j=1}^M \left| \frac{w_{jr}}{w_{jc}} \right| \quad (10.1)$$

where M is the number of output layer weights. As in previous chapters, the width penalty coefficient β allows the user to specify the emphasis placed on minimising the network output bound widths.

For a number of different β increments the radii values of the interval output layer weights were trained by the differential evolution algorithm. Time constraints were more severe for this piece of work than that of chapters 7 and 9, so only five separate algorithm runs were undertaken for each β value. This still took into account the stochastic nature of the differential evolution algorithm. Each training run used interval arithmetic to calculate the network one-step-ahead outputs on the training data set, and the resulting solution was then used in a long range prediction on the validation set. From each set of five runs, the solution with the lowest validation set interval MSE was judged to be the best solution, and then used in a long range prediction on the test data set.

10.3 Results and Discussion

As stated in section 10.1, this chapter and the previous one allow an interesting opportunity to compare the performances of MLP and RBF networks for the same data set. The initial crisp MLP network in chapter 9 gave a long range prediction error of 4.1% MSE on the test data set; whereas the initial crisp RBF network in this chapter had an inferior test set MSE of 10.2%. The 17 adjustable parameters of the crisp MLP network were also significantly less than the 56 adjustable parameters of the corresponding RBF network. Of the two neural network types, the multilayer perceptron can therefore be judged as the more appropriate for the pre-sliding friction system studied here.

The training time for the crisp RBF network was significantly shorter however, even considering that the optimal architecture of the crisp MLP network was known in advance. The 4.1% crisp MLP solution came after hundreds of separate differential evolution training runs, and the solutions from these runs were very different to one another. By contrast, for each specific architecture of the crisp RBF network, the Expectation Maximisation algorithm tended to converge to the same Gaussian mixture model for the input data, and thus the same basis function parameters.

For proportional expansion of the network output layer weights, Figure 10.1 shows the

relationship between the expansion proportion ρ and the long range test set MSE of the network. An expansion proportion of zero simply corresponds to the original crisp RBF solution.




Figure 10.1: Proportional expansion long-range MSE values

As with the proportional expansion approach in chapters 7 and 9, there is a monotonically decreasing relationship between ρ and the network prediction error. Increasing the proportional width of the output layer weights corresponds to introducing more uncertainty into the network, and thus there is a wider range of possible network outputs. Although uncertainty has only been introduced into some of the network weights here, the relationship between expansion proportion and network error is still the same as previously seen in chapters 7 and 9.

Figure 10.2 shows a long range test set prediction from the interval RBF network of 5% MSE, obtained by the proportional expansion approach. Figure 10.3 shows an alternative proportional expansion test set prediction, this time with 1% MSE. It can be seen that the output bounds in Figure 10.3 are significantly wider. As shown in Figure 10.1, to decrease the network long range prediction error from 5% MSE to 1% MSE requires a much greater

expansion of the output layer weights, and thus the prediction bounds of the network are much wider.




Figure 10.2: 5% long-range MSE solution obtained from proportional expansion




Figure 10.3: 1% long-range MSE solution obtained from proportional expansion

For differential evolution training of the network output weight radii, Figure 10.4 shows the relationship between width penalty coefficient β and the test set long range prediction error. It can be seen that there is a non-decreasing relationship between the network prediction error and β . Increasing the value of β used in the differential evolution output radii training increases the penalty placed on the width of the output weights, and thus forces the network output bounds to become narrower. These output bounds then contain less of the target data, and so the prediction error of the network increases. As the value of β tends to high values, the prediction error eventually reaches that of the original crisp RBF solution. It can also be seen that the form of the curve in Figure 10.4 is similar to the corresponding differential evolution MLP error curve in Figure 9.6 of the previous chapter.




Figure 10.4: Differential-evolution-trained long-range MSE values

Figure 10.5 shows a long range test set RBF prediction of 5% MSE, with output layer weight radii obtained by the differential evolution algorithm. A direct comparison between this and Figure 10.2 shows that there is very little difference between the proportional expansion and differential evolution approaches here, in relation to the prediction bound widths. Using all 8000 samples in Figures 10.2 and 10.5 the average sample bound width was found to be 0.2% lower for the proportional expansion approach than for differential evolution.




Figure 10.5: 5% long-range MSE solution obtained from differential evolution training

The differential evolution algorithm is therefore marginally worse than the proportional expansion approach for minimising the prediction bound widths here, even though the proportional expansion approach is a much less computationally expensive technique. The initial assumption was made here, as with chapters 7 and 9, that optimising the interval arithmetic one-step-ahead outputs of the network with differential evolution would also optimise the long range predictions using a possibilistic sampling method. The tight interval bounds in Figure 10.5 show that this assumption was valid, but the differential evolution approach does not offer superior results to justify its significantly greater computational expense, so the proportional expansion approach is the more effective of the two training approaches used here.

The use of an interval-valued RBF model allows a trade-off between minimising the prediction error of the network and minimising the widths of the predicted output bounds, as with the interval models used in previous chapters. Due to the continuous monotonic relationships shown in Figures 10.1 and 10.4, in each of the two different interval RBF training approaches, the respective uncertainty parameters ρ and β allow this trade-off to be implemented effectively. As interval-valued RBF networks such as this one could

conceivably be used for any time series regression application, the ability to implement the aforementioned trade-off is an important one. Different network applications may require different levels of precision on the predicted outputs, and this is a choice to be made by the network user.

This is once again an example of the opportunity concept in Ben-Haim's information gap theory. The desired windfall reward is given by a particular value of long range prediction error from the interval network, in this case either 5% or 1% MSE. The level of uncertainty is given by either ρ or β , depending on the interval network training method that is used. The opportunity is thus given by the value of ρ or β that allows the desired windfall reward to be achieved.

11. Interval-Valued Regression Models of Bifurcating

Nonlinear Systems

11.1 Introduction

This chapter examines the effects of bifurcation on an interval-valued discrete-time Duffing oscillator. It extends on the work of a paper by Worden et al (2005) in which bifurcation of a conventional crisp-valued Duffing oscillator was studied.

Bifurcation is an important issue for nonlinear systems, and its effects can be unpredictable. The Duffing oscillator has been chosen here because it is one of the very simplest of nonlinear systems, and therefore suitable as a starting point for representing bifurcation with interval-valued models.

In this work a harmonically-driven Duffing oscillator is studied. It can be shown that within a particular range of forcing frequency, the system bifurcates to create two different stable limit cycles of oscillation. The choice of limit cycle settled upon by the crisp-valued Duffing oscillator is dependent on the initial conditions given to the system. Outside the bifurcation frequency range there exists only a sole stable limit cycle. The Duffing oscillator coefficients initially take crisp values obtained from the background paper, and a NARX model was then used to represent the discrete-time system. It is shown that by introducing uncertainty into the system NARX coefficients within a region local to their initial crisp values, an interval-valued model can be obtained with predicted output bounds that incorporate both of the two possible limit cycles. This interval-valued model therefore indicates the behaviour that the Duffing oscillator is capable of, in relation to the studied bifurcation.

11.2 The Crisp System

The Duffing oscillator is one of the simplest and most well-known of nonlinear systems. It was also studied in chapter 8 of this thesis, although that was for the case of a white noise excitation. The harmonically-driven Duffing oscillator of this current chapter is given by:

$$m\ddot{y} + c\dot{y} + ky + k_3y^3 = X \cos(\omega t) \quad (11.1)$$

In the background paper by Worden et al, the continuous-time Duffing oscillator was integrated with the MATLAB ODE solver *ode45*. In this work, however, the system was represented as a discrete-time system with a fixed sampling frequency. This discrete-time system was given by:

$$m\ddot{y}_i + c\dot{y}_i + ky_i + k_3y_i^3 = X \cos \omega t_i \quad (11.2)$$

with $m = 1\text{kg}$, $c = 20 \frac{Ns}{m}$, $k = 10^4 \frac{N}{m}$, $k_3 = 5 \times 10^9 \frac{N}{m^3}$ and $X = 10$. Note that with the exception of X , these are the same system property values as were used for the Duffing oscillator in chapter 8. A range of different ω values were tested. The following two relations were also used:

$$\ddot{y}_i \approx \frac{y_{i+1} - 2y_i + y_{i-1}}{\Delta t^2} \quad (11.3)$$

$$\dot{y}_i \approx \frac{y_i - y_{i-1}}{\Delta t} \quad (11.4)$$

Using the abbreviation $x_i = X \cos \omega t_i$ and substituting equations (11.3) and (11.4) into equation (11.2), the following NARX model can be obtained:

$$y_i = \left(\frac{2m - \Delta t c - \Delta t^2 k}{m} \right) y_{i-1} + \left(\frac{\Delta t c - m}{m} \right) y_{i-2} + \left(\frac{\Delta t^2}{m} \right) x_{i-1} - \left(\frac{\Delta t^2 k_3}{m} \right) y_{i-1}^3 \quad (11.5)$$

The system response is given by:

$$y(t_i) = Y \cos(\omega t_i - \phi) \quad (11.6)$$

In the background paper, a harmonic balance approach was used to find the response amplitude for the case $k_3 \neq 0$. This gave the following cubic equation in Y^2 :

$$X^2 = Y^2 \left\{ \left[-m\omega^2 + k + \frac{3}{4}k_3 Y^2 \right]^2 + c^2 \omega^2 \right\} \quad (11.7)$$

Figure 11.1 shows the positive real roots of equation (11.7) for a range of forcing frequency ω . It can be seen that within a particular frequency range there are three possible response amplitudes, although the middle one of these three is unstable. This is an example of bifurcation. Outside this frequency range, however, there only exists a single response amplitude solution for each ω value.

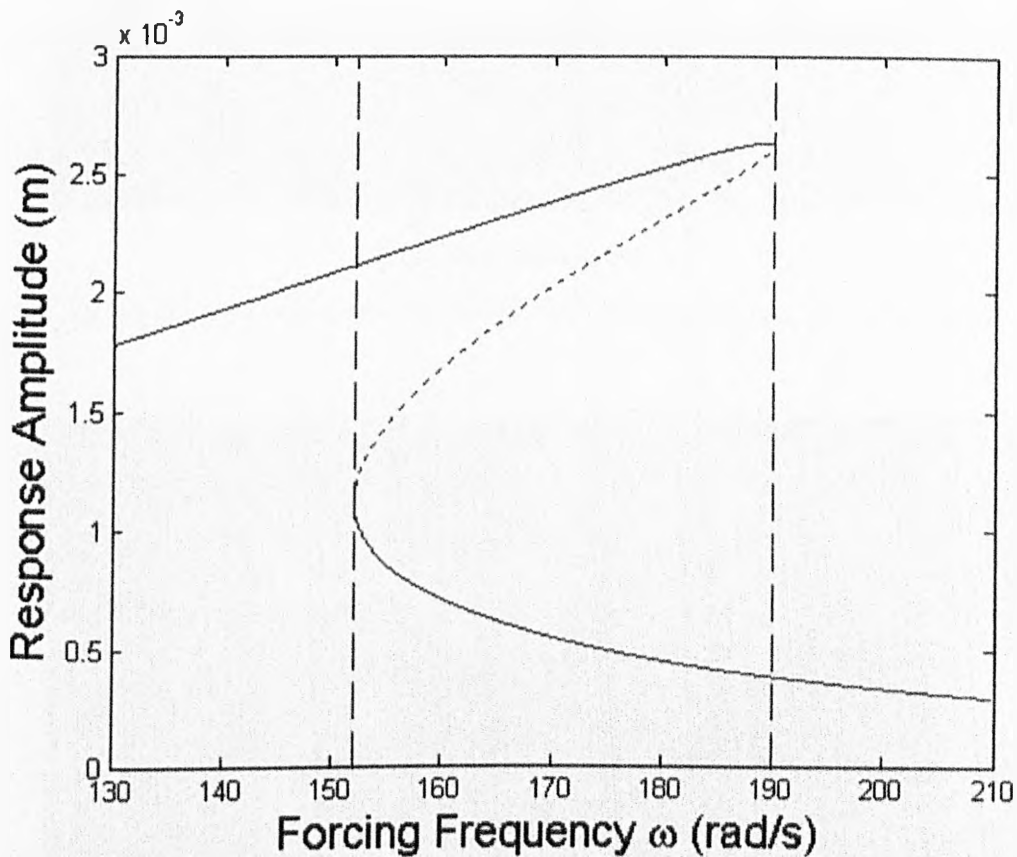


Figure 11.1: Multiple steady-state amplitude solutions in Duffing oscillator FRF

Within the bifurcation frequency range there are thus two steady-state response amplitudes for each value of ω , corresponding to two distinct limit cycles. The limit cycle settled upon is dependent on the initial conditions provided to the system. It was found that this bifurcation frequency range lies between 152 rad/s and 190 rad/s. Figures 11.2 and 11.3 examine the relationship between system initial conditions and steady-state limit cycle reached.




Figure 11.2: Steady-state behaviour vs initial conditions, $\omega = 155 \text{ rad/s}$




Figure 11.3: Steady-state behaviour vs initial conditions, $\omega = 185 \text{ rad/s}$

Figure 11.2 is for a forcing frequency of 155 rad/s, so is at the lower end of the bifurcation frequency range. For a range of initial system conditions, the plot indicates

the limit cycle that the system will settle upon (referred to as “basins of attraction”), with the two limit cycles themselves superimposed onto the plot in blue and red. It can be seen from the plot that all tested initial condition values settle on either one limit cycle or the other. For the range of initial conditions used in Figure 11.2, a sampling frequency of 1000Hz was sufficient to represent the discrete-time system. It was observed, however, that for larger initial condition values, a higher sampling frequency was needed.

Figure 11.3 also examines the obtained steady-state limit cycle in relation to the system initial conditions, but this time for a forcing frequency of 185 rad/s, i.e. at the upper end of the bifurcation frequency range. It can be seen that for Figure 11.2, the higher-amplitude limit cycle dominates the plot, whereas for Figure 11.3 it is the lower-amplitude cycle that dominates. From both Figures 11.2 and 11.3, it can be seen that initial conditions lying on a particular limit cycle do not automatically mean that the system will converge to that cycle at steady-state.

The changing limit cycle basins of attraction observed in Figures 11.2 and 11.3 can be explained by consulting Figure 11.1. Below 152 rad/s only the higher-amplitude limit cycle exists, and above 190 rad/s only the lower-amplitude cycle exists. Between these two frequency limits, therefore, there must be a transition between the two limit cycles. Figures 11.2 and 11.3 suggest that this will be a gradual transition, involving the basins of attraction shrinking for the higher-amplitude limit cycle and expanding for the lower-amplitude cycle.

11.3 An Interval-Valued Model

Within this work the crisp-valued coefficients of the NARX model in equation (11.5) are to be replaced by interval-valued coefficients. As shown in chapter 2, interval arithmetic is unsuitable for long-range predictions of any interval-valued autoregressive model because of its excessive conservatism. For interval-valued NARX coefficients in equation (11.5), the uncertain output bounds can be estimated through the use of the modified reduced transformation method of section 3.6.

As shown in Figures 11.2 and 11.3, the limit cycle settled upon for a forcing frequency within the bifurcation range is dependent on the system initial conditions. The aim of

this work is to show that by using interval-valued coefficients for a discrete-time NARX Duffing oscillator, an output prediction can be obtained that contains both possible steady-state limit cycles.

The NARX regressor coefficients in equation (11.5) can be rewritten as interval numbers. As in previous chapters, these interval coefficients are written in the “centre and radius” form of equation (7.4). Uncertainty can then be introduced into these coefficients through the use of the proportional expansion approach, where the expansion proportion is given by ρ . The NARX coefficients of equation (11.5) are therefore being “expanded” around their original crisp values.

Three different forcing frequencies were used here: 160 rad/s, 170 rad/s and 180 rad/s. For each of these three, the steady state behaviour of the system was examined over a range of different initial conditions, as in Figures 11.2 and 11.3. A different set of initial conditions was chosen for each of these three different forcing frequencies which would lie at a boundary between the higher-amplitude and lower-amplitude limit cycle basins of attraction. This meant that for these initial conditions, a very small change in initial condition would completely change the steady-state behaviour.

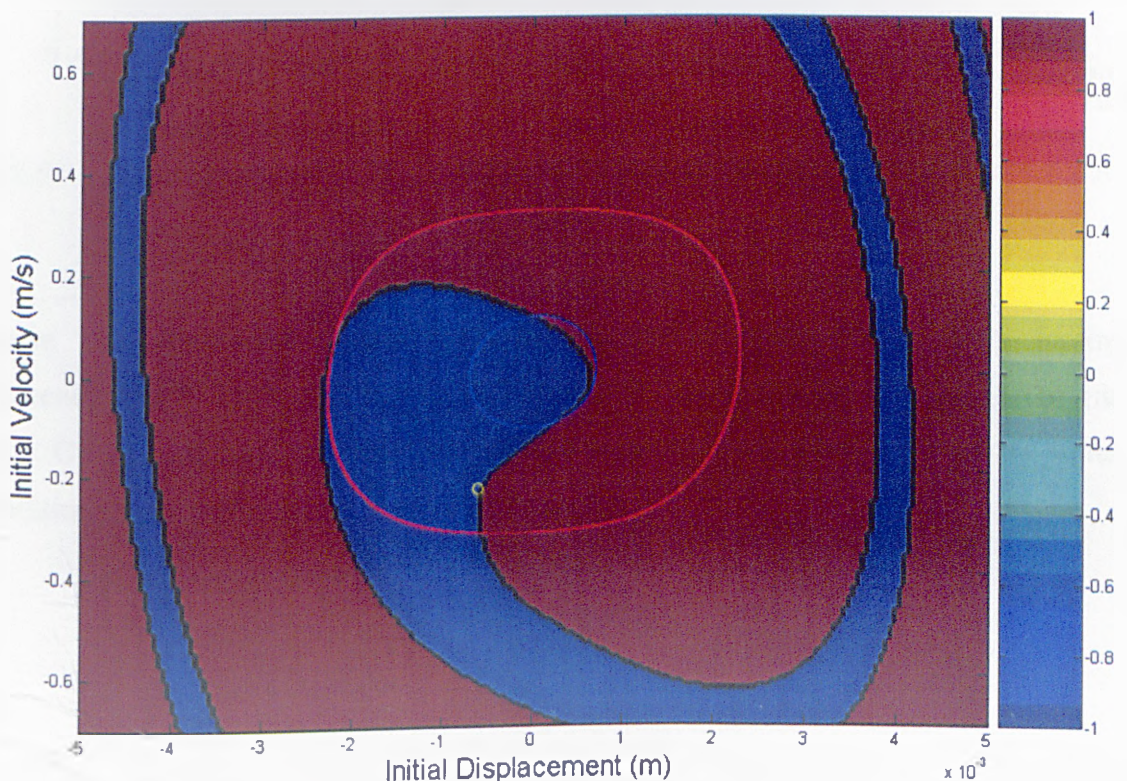


Figure 11.4: Steady-state behaviour vs initial conditions, $\omega = 160$ rad/s

Figure 11.4 shows the steady-state limit cycle vs initial conditions behaviour for a forcing frequency of 160 rad/s. The two limit cycles are superimposed onto the plot, as with Figures 11.2 and 11.3. Also added is the specific initial condition point used for the subsequent interval-valued analysis, shown by the small yellow circle.

Figure 11.5 shows two displacement time-histories of the system in Figure 11.4, with respective initial conditions very close to the highlighted point in Figure 11.4, and on either side of the basins of attraction boundary. 24000 time samples were used in total, and the first 800 samples are shown in Figure 11.5.




Figure 11.5: Displacement time-histories for near-identical initial conditions, $\omega = 160$ rad/s

Figure 11.6 shows the steady-state behaviour vs initial conditions for a forcing frequency of 170 rad/s, and Figure 11.7 shows the same plot for a frequency of 180 rad/s. On both of these two plots are marked small yellow circles to denote the initial condition sets used for the interval-valued analysis.




Figure 11.6: Steady-state behaviour vs initial conditions, $\omega = 170$ rad/s




Figure 11.7: Steady-state behaviour vs initial conditions, $\omega = 180$ rad/s

11.4 Results and Discussion




Figure 11.8: Interval coefficient expansion, $\rho = 1 \times 10^{-5}$, $\omega = 160$ rad/s




Figure 11.9: Interval coefficient expansion, $\rho = 1 \times 10^{-4}$, $\omega = 160$ rad/s

Figures 11.8 and 11.9 both show displacement predictions from interval-valued models of equation (11.5), with a forcing frequency of 160 rad/s. The proportional expansion approach was used to generate the interval NARX coefficients, and the system initial conditions were those indicated by the small yellow circle on Figure 11.4. The interval bounds were estimated by the transformation method with 20 hypercube levels

(excluding the crisp central point), as detailed in section 11.3. Figure 11.8 used an expansion proportion of $\rho = 1 \times 10^{-5}$, and Figure 11.9 used $\rho = 1 \times 10^{-4}$. Both of these two figures show only the first 800 of the 24000 time samples.

Figure 11.10 shows two displacement time-histories for a forcing frequency of 170 rad/s. The two sets of initial conditions for these two curves are taken at near-identical points very close to the highlighted yellow circle on Figure 11.6, where these two points were deliberately chosen to lie on opposite sides of the boundary for the steady-state limit cycle basins of attraction.




Figure 11.10: Displacement time-histories for near-identical initial conditions, $\omega = 170$ rad/s

Figures 11.11 and 11.12 show displacement predictions from interval-valued models of equation (11.5), with a forcing frequency of 170 rad/s. These two predictions are for initial conditions given by the small yellow circle on Figure 11.6, and the proportional expansion approach was used for the interval NARX coefficients. Figure 11.11 is for an expansion proportion $\rho = 1 \times 10^{-5}$ and Figure 11.12 is for $\rho = 1 \times 10^{-4}$.




Figure 11.11: Interval coefficient expansion, $\rho = 1 \times 10^{-5}$, $\omega = 170$ rad/s




Figure 11.12: Interval coefficient expansion, $\rho = 1 \times 10^{-4}$, $\omega = 170$ rad/s

Figure 11.13 shows two displacement time-histories for a forcing frequency of 180 rad/s. The two curves have almost identical initial conditions, both close to the small yellow circle on Figure 11.7. These two curves were chosen to lie on opposite sides of a basins of attraction boundary for the steady-state limit cycles.




Figure 11.13: Displacement time-histories for near-identical initial conditions, $\omega = 180$ rad/s

Figures 11.14 and 11.15 show displacement predictions from interval-valued models of equation (11.5), with a forcing frequency of 180 rad/s. The initial conditions for these two interval plots are given by the small yellow circle on Figure 11.7, and once again the proportional expansion approach was used for the interval-valued coefficients. Figure 11.14 is for a coefficient expansion proportion $\rho = 1 \times 10^{-5}$ and Figure 11.15 is for $\rho = 1 \times 10^{-4}$.




Figure 11.14: Interval coefficient expansion, $\rho = 1 \times 10^{-5}$, $\omega = 180$ rad/s




Figure 11.15: Interval coefficient expansion, $\rho = 1 \times 10^{-4}$, $\omega = 180$ rad/s

Figures 11.2, 11.3, 11.4, 11.6 and 11.7 relate the steady-state behaviour of the system to the initial condition values that are used. Each of these figures is for a different value of ω . It can be seen that the “basins of attraction” of the two limit cycles change significantly as ω passes through the bifurcation frequency range. At the lower end of this frequency range it is the higher-amplitude limit cycle that dominates the plot, as can be seen from Figure 11.2. Towards the upper end of the bifurcation frequency range, however, it is the lower-amplitude limit cycle that has the largest basins of attraction, as shown in Figure 11.3. Figures 11.4, 11.6 and 11.7 show that there is a gradual transition between those two extreme cases, with each figure showing a different intermediate ω value.

For the interval-valued analysis, three different values of forcing frequency were considered individually. This was to show that the technique is valid over the whole bifurcation range. In each of these three ω cases, an initial condition was chosen to lie very close to a basins of attraction boundary for the two steady-state limit cycles.

The reasoning behind this was that the introduction of a relatively small amount of uncertainty into the NARX model coefficients would create an interval range prediction that would cross both zones of attraction. The interval bounds should therefore contain both of the limit cycles but yet remain relatively narrow, as the level of coefficient expansion is still small.

Figures 11.5, 11.10 and 11.13 show the effects of slight differences in system initial conditions for the three different forcing frequencies. For each of these three plots it can be seen that the two displacement histories remain near-identical for a small handful of oscillations, before diverging to the two alternative steady-state limit cycles. The sensitive dependence of the system on initial conditions can clearly be seen from these three figures. The figures also each demonstrate the levels of their two respective amplitude cycles for the subsequent interval analysis.

Figures 11.8 and 11.9 show interval coefficient expansions of the model for a forcing frequency of 160 rad/s. The initial condition was chosen to lie just inside the basin of attraction for the lower limit cycle, at the yellow point indicated on Figure 11.4. Figure 11.8 shows that for a coefficient expansion proportion of $\rho = 1 \times 10^{-5}$, the interval model only includes the lower-amplitude limit cycle after the transients have decayed. This level of ρ is therefore insufficient for detecting both limit cycles.

Figure 11.9 then shows that if the expansion proportion ρ is increased to 1×10^{-4} , the prediction bounds from the interval model include both limit cycles. It can be seen that once the steady state is reached, these bounds contain the two alternative limit cycles with the minimum necessary width. This interval-valued NARX model is therefore effective for illustrating distinct cases of steady-state behaviour. It can also be seen from Figure 11.9 that the largest range of the interval bounds is in the region of sample numbers 300-500. This is the region where alternative displacement histories have already begun to diverge from one another, but where steady-state behaviour has not yet been reached.

The forcing frequency was then changed to 170 rad/s and a different set of system initial conditions was chosen. As with the 160 rad/s case, the set of system initial conditions was chosen to lie very close to a basins of attraction boundary for the steady-state limit cycles, and just inside the lower-amplitude cycle region.

The same behaviour was observed for the interval-valued models as with $\omega=160$ rad/s. Figure 11.11 shows that a small coefficient expansion proportion of $\rho = 1 \times 10^{-5}$ is insufficient to represent both steady-state limit cycles, so only includes the lower-amplitude cycle. The larger expansion of $\rho = 1 \times 10^{-4}$ in Figure 11.12, however, is

enough for the interval prediction bounds to contain both limit cycles. As with Figure 11.9, the interval bounds contain the two steady-state limit cycles without any excessive prediction width, and the largest range of system behaviour is before the transients have decayed.

Adjusting the forcing frequency ω to 180 rad/s still led to the same behaviour of the interval-valued models. Figure 11.13 shows a pair of displacement time-histories for near-identical initial conditions. These two initial conditions both lay very close to the small yellow circle indicated on Figure 11.7, and were chosen to lie on opposite sides of the basins of attraction boundary. As with Figures 11.5 and 11.10, it can be seen that these two displacement plots are almost identical for the first few oscillations of the system, before diverging to the two different steady-state limit cycles.

Figures 11.14 and 11.15 show displacement plots for interval-valued models with $\omega=180$ rad/s. Both of these two plots used an initial condition set that lay just inside a basin of attraction for the lower-amplitude limit cycle, as shown by the yellow circle on Figure 11.7. Figure 11.14 shows that for an interval coefficient expansion with expansion proportion $\rho = 1 \times 10^{-5}$, the interval bounds only contain the lower-amplitude limit cycle after the steady state has been reached. Figure 11.15 then shows that if ρ is increased to 1×10^{-4} , the interval bounds are wide enough to include both steady-state limit cycles with the minimum required width. As with Figures 11.9 and 11.12, the largest range of outputs in Figure 11.15 is found in the transient region.

This work has therefore used interval-valued models to indicate the potential behaviour of a harmonically-driven Duffing oscillator capable of bifurcation. The interval-valued models were all used in regions where the system showed a high sensitivity to initial conditions. Despite the existence of two distinct steady-state behaviours, these two behaviours were both predicted simultaneously through the deliberate introduction of uncertainty into the NARX model coefficients. These interval NARX coefficients were “proportionally expanded” around their original crisp values, so that the studied system coefficient ranges were local to these original values. The expansion parameter ρ allowed the level of the uncertainty introduced into the model to be controlled.

The system studied here was only a simple nonlinear system with a bifurcation whose effects were relatively minor. This interval-valued technique could also potentially be

used, however, for indicating the possible behaviour of a more complex and severe bifurcation of a nonlinear system.

12. The Dynamic Effects of Uncertainties within Automotive

Spot Welds

12.1 Introduction

The work in this chapter examined the effects of uncertainties on the dynamic response of an automotive spot-welded component. The work was conducted during a nine-month industrial placement at Centro Ricerche Fiat in Turin, Italy, as part of a European Union research and training network entitled MADUSE (Modelling Product Variability and Data Uncertainty in Structural Dynamics Engineering).

The technique of resistance spot welding is ubiquitous within the automotive industry. A typical car can contain thousands of spot-welded joints, so their effect on the dynamic performance of the car is significant. Spot welding is much cheaper and less time-consuming than using continuous welds to join together panels in a car. Spot welds have a large effect on the overall stiffness of the car, and also play a significant role in its vibration damping characteristics.

The work here used a possibilistic uncertainty modelling approach to represent the response of the spot-welded component, where this component was based on an actual automotive component with some modifications. The original part was not symmetrical, although the design was modified so that symmetry was present. Previous empirical experience has shown that symmetry can amplify certain dynamic effects, so this made the modelling process more of a challenge. Three different cross-sectional designs were initially generated, in order to suggest possible modifications for reducing the subsequent manufacturing costs of the component.

The spot-welded component was represented both by a computational approach, in the form of finite element modelling, and by an experimental approach, using manufactured test specimens. The intention was to generate a finite element model of the system response, taking into account any relevant parameter uncertainties, and then obtain results from this computational model. These results could then be validated by modal testing of the manufactured specimens, with the use of a suitable experimental setup. A good

correlation between the finite element and experimental modal testing results would verify that the computational model was accurate.

By analysing any discrepancies between the two sets of results, any weaknesses or deficiencies of the finite element model could be deduced. This would be an effective method for suggesting improvements in the work, especially if it were to be repeated at some future time.

12.2 Resistance Spot Welding

The technique of producing automotive spot welds is generally referred to as resistance spot welding, since it utilises the resistance of a metal sheet interface to a high electrical current. The process uses two electrodes to join the two metal sheets in question, and these electrodes can be assembled in the form of a hand-held tool to be used by a production operative, or can be utilised by an automated production-line robot.

The two metal sheets are joined at a small ellipsoidal area or “spot”, and the welding process has four distinct steps, the first of which is called the squeeze cycle. The two sheets to be welded are placed together, and the two electrodes are added to their respective outer surfaces. A load is applied to the electrodes in order to press the two sheets together.

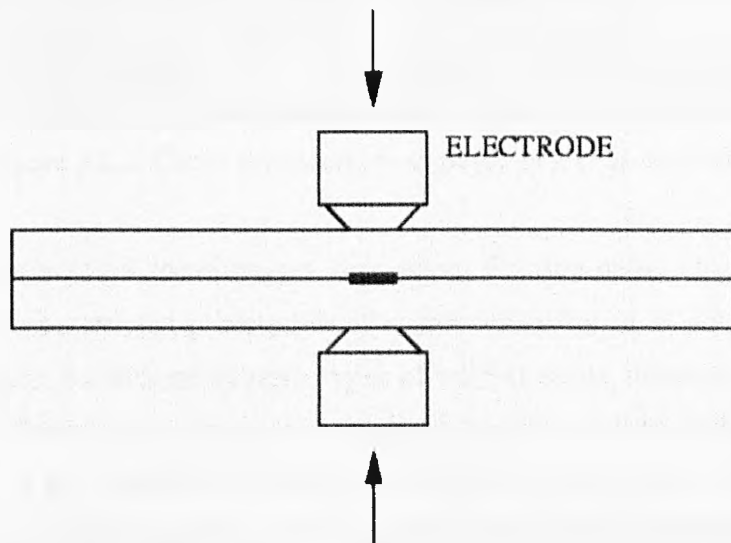


Figure 12.1: The “weld cycle” for a resistance spot weld

The second step is called the weld cycle. A high electrical current is passed through the

electrodes, and resistance to the current at the interface between the two sheets causes the material at this point to become plastic. The metal at the interface between the electrodes fuses together to form a small ball joining the two sheets, known as the spot weld nugget. This step is illustrated graphically in Figure 12.1.

The third step is called the hold cycle. The electrical current through the surface electrodes is turned off, but the load from the electrodes to the sheet surfaces is maintained. The pressure arising from this load allows the molten nugget to cool and solidify. The fourth step is called the off cycle, and simply consists of the electrodes being removed from the metal sheet surfaces.




Figure 12.2: Cross-sectional photograph of a spot-welded joint

The end result is a small metal nugget connecting the two metal sheets together. Figure 12.2 shows a cross-sectional photograph of a spot-welded joint, albeit a poor quality one with weld porosity. As with many other types of welded joints, immediately adjacent to the joint is a heat affected zone, where the material properties will be different from those of the base metal. In the region close to the spot weld there will be no physical link between the two metal sheets, but these sheets are likely to be frequently in contact with each other. During vibrations of the system, the friction losses at these contact interfaces will cause a large amount of energy to be dissipated. In a typical car which may contain thousands of spot-welded joints, this energy dissipation is the principal source of vibration damping prior

to any damping treatment being added.

12.3 The Spot-Welded Component

Three different initial designs were originally considered for the spot-welded component here. This was to suggest possible reductions in the cost of manufacturing the design at a later date, for use as test specimens. All three of these designs had a constant cross section, an axial length of 500mm and a sheet thickness of 1mm. The designs show the sheet midplanes.




Figure 12.3: First design for the spot-welded component

A cross section of the first design is shown in Figure 12.3. This design has the dimensions

of the original automotive part on which the component is based, but adjusted to give symmetry about the vertical axis.

The second design replaced the two interior angles of the specimen with a single angled section, as shown in Figure 12.4. This would require a single ply operation for each angle, as opposed to two, so was thought to save on manufacturing costs. Although both of the designs in Figures 12.3 and 12.4 are symmetrical about the vertical axis, neither of them have horizontal symmetry. The second design was adjusted to have this horizontal symmetry, giving the third design shown in Figure 12.5.




Figure 12.4: Second component design




Figure 12.5: Third component design

After consultation with workshop operatives, it was found that replacing the two interior component angles with a single angled section offered no real saving in manufacturing

costs, so the second and third component designs were scrapped. It was decided that the specimen needed to have vertical and horizontal symmetry, however, so the first design was adjusted to give the design shown in Figure 12.6. This design was thus decided upon as the final design of the spot-welded component, and was used as the nominal model for the work here.



Figure 12.6: Final design for the spot-welded component

Initially it was planned to use 24 spot welded joints within the design in Figure 12.6. These were to be arranged into two groups of 12 spot welds, and each arranged symmetrically on the 17mm vertical faces at the top and bottom of the design. When arranged in regular increments along the axial component length, the spacing between each spot weld would have been approximately 40mm.

Preliminary FE testing was conducted with Nastran for the use of these 24 spot welds. A number of different finite element weld types were used, including RBE2 rigid elements, linear solid elements, and the new Nastran CWELD element. Parameters such as the FE

mesh density and spot weld diameter were adjusted, as well as small changes to the number of spot welds used. For each of these adjustments the finite element model was solved, and modal analysis was conducted on the result. Ignoring the first six rigid body modes of the model, the frequencies of the following ten modes were recorded in each case.

These initial changes to the model did not lead to a great variation in the measured modal frequencies. The conclusion was reached that the nominal 24 spot welds being used was too many, and therefore needed to be significantly reduced. After consideration, it was chosen to use only 10 spot-welded joints in the model, to be equally divided amongst the top and bottom edges of the design. Figure 12.7 shows a side view of the selected design in Figure 12.6, and highlights the location of these 10 spot welds in yellow.




Figure 12.7: Side view of the spot-welded component

12.4 Spot Weld Uncertainties

The component design introduced in section 12.3 was viewed as being the nominal design for the work here. The task was then to consider how uncertainties could be introduced into this design, and how these uncertainties could be represented. There were a large number of possibilities for uncertainty within the spot-welded component. The options initially considered here were:

- Spot weld diameter
- Differences in spot weld positions
- Number of missing spot welds
- Spot weld stiffness

- Contact area around a spot weld
- Residual pressure distribution around a spot weld
- Operational loads on the component
- Component dimensions
- Sheet thickness
- Surface quality of the sheets prior to spot welding

As with any engineering modelling exercise, computational time and effort was an issue here, so it simply wasn't possible to examine the effects of all these uncertainties. For a variety of reasons, many of the uncertainty sources listed above were viewed as less important or difficult to model, and therefore were not selected.

It was decided to consider the effects of just three uncertainty sources, these being:

- Sheet thickness
- Differences in spot weld positions
- Number of missing spot welds

The latter part of this work was to involve experimental testing of manufactured specimens. Whereas computational finite element modelling does not involve any physical components, this manufacture of experimental specimens is subject to production costs and lead time. The quantity of experimental specimens to be manufactured therefore needed to be within practical limits, whilst simultaneously giving a suitable distribution for the spot weld uncertain parameters.

It was thus decided to use three levels for the sheet thickness, three levels for the difference in spot weld positions, and two levels for the number of missing spot welds. This would give 18 possibilities in total, including the nominal design, and hence 18 different manufactured specimens. For the sheet thickness, a uniform thickness was assumed for each design. In reality, any manufactured component is likely to exhibit distributions in sheet thickness over the entire component, but this uniform thickness simplification was necessary for the modelling to take place. The three levels of sheet thickness used were 0.9mm, 1.0mm and 1.1mm.

The levels for the differences in spot weld position uncertainty can be illustrated by Figure 12.8, which shows an enlarged view of two spot welds in Figure 12.7 with the central section of the component removed. Note that compared to previous diagrams, Figure 12.8 is no longer to scale. The three points A, B and C show the three possible levels for this parameter, where B corresponds to the nominal design. As with the sheet thickness, a simplifying assumption was also made for modelling the differences in spot weld positions. Allowing the ten spot welds to be displaced independently of each other would give far too many possibilities to be adequately represented, so it was assumed that the component spot welds would be displaced together.




Figure 12.8: The three levels for the difference in spot weld positions uncertainty

For situation A, the model spot welds were all displaced outwards from the component central axis by a vertical distance of 2.5mm, i.e. the 5 spot welds at the top edge of Figure 12.7 were displaced upwards, and the 5 lower spot welds were displaced downwards. For situation C the opposite was true, so that all ten spot welds were displaced inwards, moving together towards the central axis of the component. For situation B, all spot welds remained

at their original positions given by the nominal model.

As with the two previous uncertain parameters, a modelling simplification was also made for the missing spot welds parameter. The first level for this parameter was to have all ten spot welds present, as in the nominal model in Figure 12.7. The second level was for two missing spot welds, but the assumption was that it would always be the same two spot welds that were absent. This is shown in Figure 12.9, where positions 2 and 8 give the locations of these two missing spot welds.




Figure 12.9: Position indices for the ten component spot welds

In a simplified design such as this one, with only a small number of spot welds, it is unlikely that any of these spot welds would be omitted during production. For the case of a full automobile however, where there are thousands of spot weld locations, it is likely that a very small number of these spot welds would be accidentally missed by a production operative, therefore this study of missing spot welds is justified. Table 12.1 shows the entire arrangement of uncertain parameters for all 18 manufactured specimen designs.

Test Number	Sheet Thickness	Spot Weld Positions	Missing Spot Weld Positions
1	1.1mm	C	-
2	1.1mm	C	2 & 8
3	1mm	C	-
4	1mm	C	2 & 8
5	0.9mm	C	-
6	0.9mm	C	2 & 8
7	1.1mm	B	-
8	1.1mm	B	2 & 8
9	1mm	B	-
10	1mm	B	2 & 8
11	0.9mm	B	-
12	0.9mm	B	2 & 8
13	1.1mm	A	-
14	1.1mm	A	2 & 8
15	1mm	A	-
16	1mm	A	2 & 8
17	0.9mm	A	-
18	0.9mm	A	2 & 8

Table 12.1: Arrangement of uncertainties for the 18 manufactured test specimens

12.5 The Finite Element Model

Altair Hypermesh 6.0 was used as a pre-processor here, along with MSC Nastran as the finite element solver. Hypermesh was also used in the post-processing phase for viewing modal analysis results, although the majority of the post-processing operations were done by Veiprod, a specialist piece of dynamic analysis software created by Centro Ricerche Fiat.

The spot-welded component was represented by CQUAD4 shell elements with an average element size of 2.5mm. This corresponded to 28430 nodes and 27860 elements in the FE model. For representing the spot welded joints, the available options were to use RBE2 Nastran rigid bar elements, linear solid elements, or the new Nastran CWELD element. The latter is a point-to-point weld that is able to link the centre of one mesh element to the centre of another, so does not rely on the finite element meshes being congruent at the weld faces. CWELD elements allow the diameter and material type of the weld to be specified, and the weld is modelled elastically. This is in contrast to the RBE2 element, which is simply a rigid connection between two nodes. The linear solid (brick) weld element also treats the welds elastically, but does not represent the true cross-sectional shape of a spot

weld nugget. For these reasons, it was decided to use the CWELD element for the spot welds here.

Whereas the number of manufactured specimens is limited by material and production costs, these constraints do not apply for finite element models. The only things to consider are the computational time and available processing power. As the model here is relatively small by finite element standards, neither of these two considerations should be an issue here. This enables far more tests to be done in the finite element modelling than for the experimental testing.

Based on this, it was decided to conduct a more exhaustive finite element investigation of missing spot welds than for the experimental specimens. Up to two spot welds were to be omitted from the finite element model. As already stated, the spot-welded component is symmetrical in all three principal orthogonal directions, and this symmetry can be used to best advantage.

For all ten spot welds present in the model, there is obviously only one possible arrangement for achieving this. For one missing spot weld, there initially seems to be ten possibilities for this, i.e. any weld from index 1-10 on Figure 12.9 not present. However, due to the symmetry of the model, weld number 1 missing is the same as weld numbers 5, 6 or 10 missing. Similarly, weld number 2 missing is the same as weld numbers 4, 7 or 9 missing, and weld number 3 missing is the same as weld number 8 missing. Thus there are only truly three possibilities for one missing spot weld – index position 1, 2 or 3.

By a similar procedure, there are 15 different possibilities for two missing model spot welds, listed in Table 12.2. With one possibility for zero missing spot welds and three possibilities for one missing spot weld, there are thus 19 possibilities for up to two missing model spot welds. When used in conjunction with the same three sheet thickness levels and three spot weld position levels of section 12.4, there are 171 different finite element models to be used here. Although this number may seem excessive, each model can be run in under five minutes, so the testing schedule is well within practical limits.

Test Number	Missing Index Positions	
1	1	2
2	1	3
3	1	4
4	1	5
5	1	6
6	1	7
7	1	8
8	1	9
9	1	10
10	2	3
11	2	4
12	2	7
13	2	8
14	2	9
15	3	8

Table 12.2: Potential arrangements of any two missing spot welds (with symmetry)

12.6. Modal Analysis of the Spot Welded Component

Modal analysis is the study of the resonant vibration for a given system, where resonance can be described as excessive, sustained oscillatory motion. Modal analysis is conducted in the frequency domain, so involves searching for the natural frequencies of the system in question. It can be done either computationally, with finite element packages such as Nastran, or experimentally through the use of accelerometers and impact hammers or electrodynamic shakers. Modal analysis is a very well-known tool within structural dynamics, having been in use for decades, and is used extensively by industry.

12.6.1 The Modal Assurance Criterion

The modal assurance criterion (MAC) is a well-known and long established modal analysis technique, and was proposed by Allemang and Brown (1982). It is used to compare two different estimates of a modal vector for consistency.

The modal assurance criterion is generally used in conjunction with the modal scale factor (MSF), so this must also be explained in any outline of the technique. For a test with n output degrees of freedom (DOFs) and m input DOFs, the system can be written in Laplace notation as:

$$X(s) = H(s)F(s) \quad (12.1)$$

where $X(s)$ is the displacements vector, $H(s)$ is the $(n \times m)$ transfer function matrix and $F(s)$ is the forcing vector.

The transfer function matrix can then be written as the partial fraction expansion:

$$H(s) = \sum_{r=1}^n \frac{[A(r)]}{s - p(r)} + \frac{[A^*(r)]}{s - p^*(r)} \quad (12.2)$$

where the $(n \times m)$ matrix $A[r]$ is known as the residue matrix.

This residue matrix can be rewritten as:

$$[A(r)] = k(r)U_{OUT}(r)U_{IN}(r)^T \quad (12.3)$$

where $k(r)$ is a scaling constant, $U_{OUT}(r)$ is the modal coefficients vector for the system output DOFs and $U_{IN}(r)$ is the modal coefficients vector for the system input DOFs.

Any two rows or columns c and d of the residue matrix $A(r)$ for a given mode are thus proportional to each other. The constant of proportionality is given by the modal scale factor:

$$MSF(c, d, r) = \frac{U(c, r)}{U(d, r)} \quad (12.4)$$

The modal scale factor is a complex scalar, and can be rewritten as:

$$MSF(c, d, r) = \frac{A(c, r)A^*(d, r)}{A(d, r)A^*(c, r)} \quad (12.5)$$

where $A(c, r)$ and $A(d, r)$ are two rows or columns of the residue matrix $A(r)$.

The modal scale factor can thus be thought of as the ratio between two modal vector estimates c and d , where d is the reference vector. Modal vector c can be viewed as having two parts – one part which is correlated with reference vector d , and the other part which is not.

The modal assurance criterion is a measure of the consistency between modal vector estimates c and d . It is given by:

$$MAC(c, d) = \frac{|A(c, r)A^*(d, r)|^2}{A(c, r)A^*(c, r)A(d, r)A^*(d, r)} \quad (12.6)$$

The value of the MAC will always be a real number between 0 and 1 inclusive. Two modal vector estimates that are completely consistent with each other will have a MAC value of unity, and two estimates with absolutely no consistency between them will have a zero MAC value.

The MAC can be used in several different ways. One is to compare two different modal vector estimates from the same source, for example, from different rows or different columns of a frequency response function matrix. A second approach is to compare two modal vector estimates from different sources, so a finite element modal vector could be compared for similarity with an experimental modal vector. A third application is to compare a set of modal vectors for orthogonality, commonly known as the “auto-MAC”.

The auto-MAC compares a set of modal vectors against themselves. For p measured modes within the frequency range of interest, this gives a $(p \times p)$ comparison matrix, where each element is an individual MAC value between 0 and 1. The leading diagonal of this matrix is made up of unity values, since the MAC of a modal vector with itself will always give a value of one. For a completely orthogonal set of modal vectors, the auto-MAC will give the identity matrix, so that all off-diagonal terms are zero. In reality this ideal situation is likely to be unattainable, but the user often settles for low off-diagonal terms, indicating near-orthogonality of the modal vectors.

12.6.2 The Frequency Domain Assurance Criterion

One of the deficiencies of the MAC is that it is only designed to be used with modal vectors. To obtain modal vector estimates often requires significant post-processing of results after the modal testing phase has been completed. In certain situations, modal vector information may not be available.

Frequency response functions from modal testing are much simpler to obtain. Based on this, the frequency domain assurance criterion (FDAC) was proposed by Pascual et al (1997). This technique allows two different frequency response matrices to be compared for similarity, on a frequency-by-frequency basis. The end result is a matrix that enables any frequency increment in one frequency response matrix to be compared with any frequency increment in the second frequency response matrix. Regions of similarity between the two frequency response matrices can then be observed.

If $[H_a(\omega)]$ is the frequency response matrix for an analytical model and $[H_x(\omega)]$ for an experimental test, then the FDAC can be defined as:

$$FDAC(\omega_a, \omega_x, j) = \frac{|\{H_a(\omega_a)\}_j^T \{H_x(\omega_x)\}_j|^2}{(\{H_a(\omega_a)\}_j^T \{H_a(\omega_a)\}_j) (\{H_x(\omega_x)\}_j^T \{H_x(\omega_x)\}_j)} \quad (12.7)$$

where ω_a and ω_x are specific analytical and experimental frequencies respectively, and j is for a specific column of the frequency response matrix.

As with the modal assurance criterion, the FDAC returns a value between 0 and 1, where 1 indicates complete consistency and 0 signifies absolutely no consistency between the two frequency response matrices at the two given frequencies.

One of the advantages of the FDAC is that it uses entire columns of the frequency response matrix, so considers all response DOFs together. It may be that particular modes are observable at certain output DOFs but not at others, so this feature takes this into account.

12.6.3 Modal Testing

The aim of this work was to examine the effects of the selected spot weld uncertainties on the dynamic response of the component, and thus modal analysis was used here. The frequency range of interest was 0 – 1024 Hz.

The modal testing of the component was to be done by a standard hammer testing process, with the response to be measured by accelerometers. Ten accelerometers were to be used, as this was deemed in advance to be a sufficient quantity. Using Nastran, three different finite element versions of the nominal specimen design were tested. Two of these models included additional concentrated masses at ten arbitrarily selected accelerometer locations, and the third model was unchanged. One of the two concentrated mass models featured masses of 2.5g each, to represent uniaxial accelerometers, and the other used masses of 12.5g for triaxial accelerometers. Figure 12.10 shows an inertance plot taken at the same point for these three models, for a free-free loading arrangement.

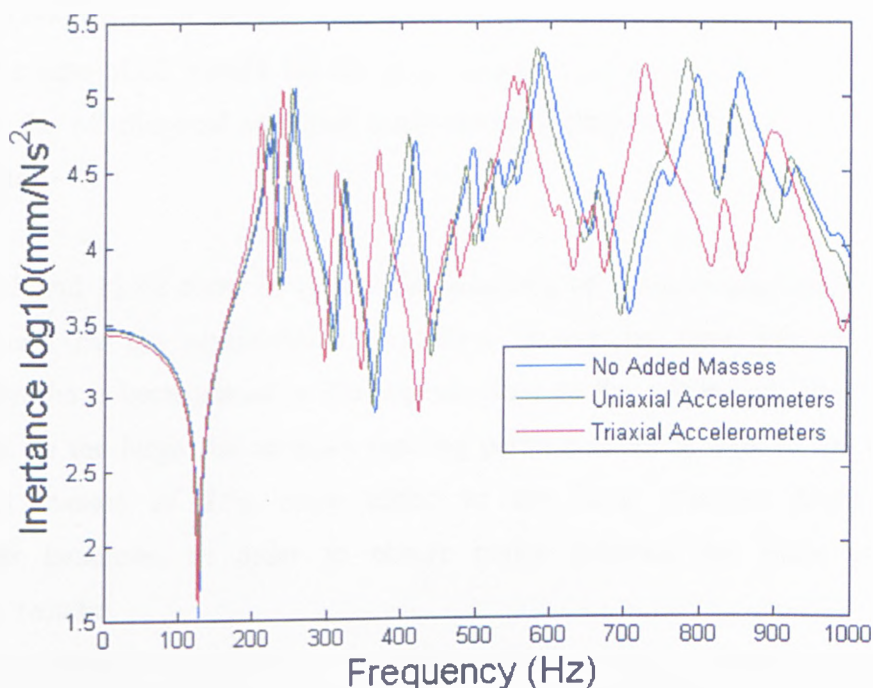


Figure 12.10: Effects of added accelerometers on the finite element model response

It can be seen that the addition of ten triaxial accelerometers noticeably changes the response of this nominal FE model, whereas ten uniaxial accelerometers have a much smaller effect. For this reason, it was decided to use uniaxial accelerometers in the modal

testing here, and measure the response in each orthogonal direction separately.

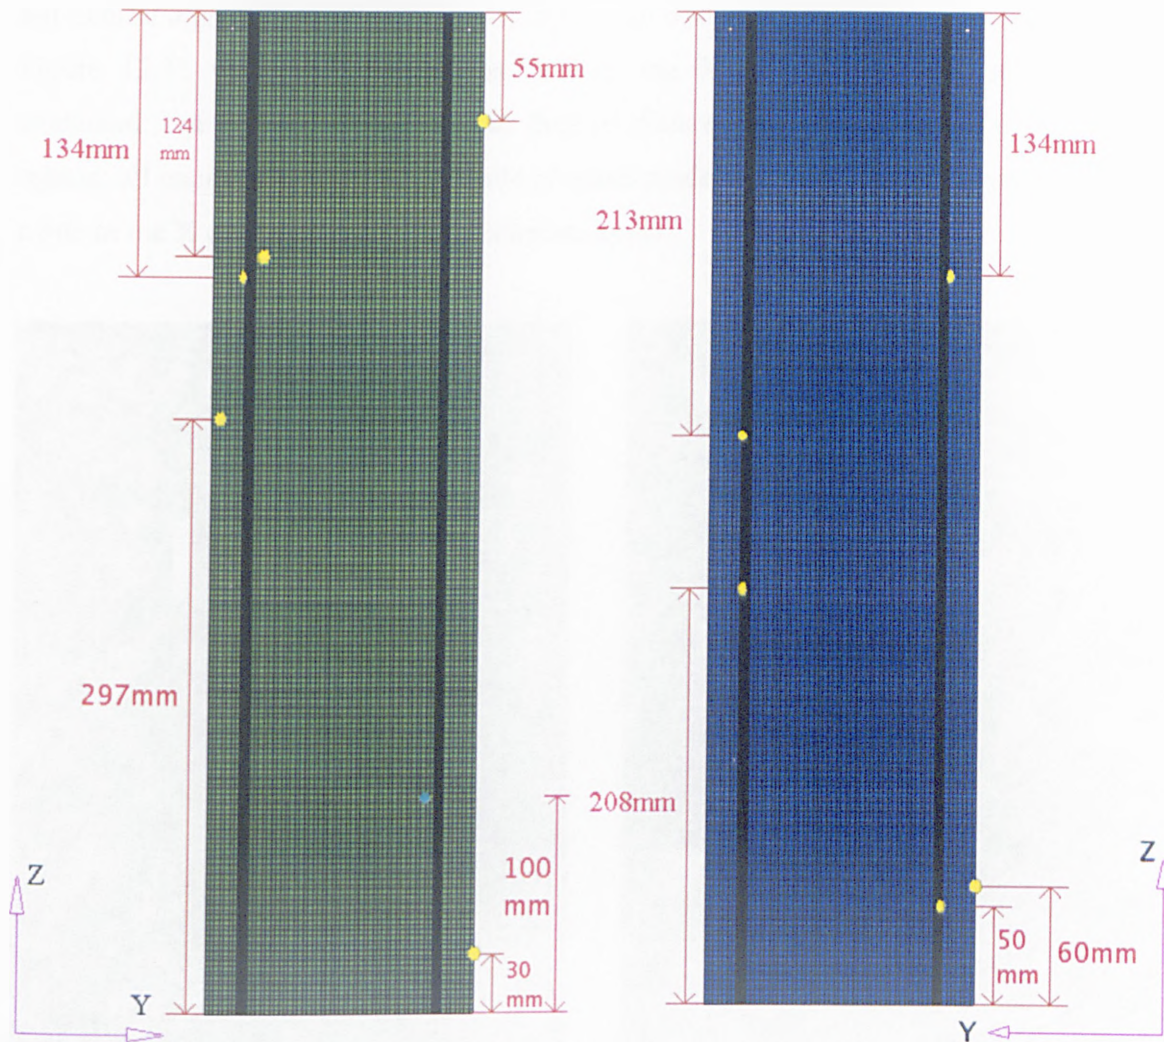
The ten accelerometer locations in Figure 12.10 were chosen completely arbitrarily, and for the sole purpose of examining the effects of the added accelerometer masses onto the specimen during testing. Prior to any modal testing taking place, the final locations of the accelerometers needed to be chosen. It was decided to use an optimisation procedure for this task, in conjunction with the auto-MAC. For the work here, the optimum ten accelerometer locations will be a set of points that are able to detect each mode in the frequency range of interest with no ambiguity. The quality of each accelerometer position set can be assessed by two measures – the sum of the off-diagonal auto-MAC matrix terms and the maximum off-diagonal term.

An optimisation software package called Mode Frontier was used to implement this procedure. An alternative finite element model of the nominal specimen was generated with an increased element size of 15mm and solved in a single Nastran solution run, with the response at each model node stored in an output file. The Mode Frontier software initially generated 33,000 different random combinations of the ten accelerometer positions, and calculated the auto-MAC matrix for the responses at these points. For each different test arrangement the off-diagonal sum and maximum off-diagonal term of the MAC matrix were calculated.

Figures 12.11 and 12.12 show in yellow the locations of the optimum ten accelerometer positions found by the optimisation procedure. It can be seen that all ten of the accelerometers have been placed in the regions close to the component spot welds, with none located on the large flat sections running parallel to the Z axis of the component. Concentrated masses of 2.5g were added to the finite element models at these accelerometer locations, in order to ensure parity between the finite element and experimental results.

In Figures 12.11 and 12.12, the two halves of the component are shown in the orientation in which they are to be suspended during the subsequent experimental modal testing. Two suspension holes were drilled in each manufactured specimen in order to facilitate the modal testing process, and elements corresponding to these holes were removed from the finite element models. These missing elements can be observed by close examination of the

top of Figure 12.12. The blue point shown 100mm from the base of Figure 12.11 corresponds to the hammer excitation point, used in both the experimental and finite element modal testing.



Figures 12.11 and 12.12: Optimal accelerometer locations on the two halves of the component

Preliminary finite element testing of the nominal model also showed that the Z direction response (i.e. along the axial length of the specimen) was much less significant than in the X and Y directions. Frequency response plots taken at several different points showed that the Z direction responses were relatively featureless, with specimen modes giving only small rises in the FRF amplitude. Based on this, and also taking into account time constraints, only the X and Y direction responses of the spot welded components were considered here, for both the finite element models and manufactured test specimens.

The experimental modal testing was done using a HP-Unix workstation with a spectral analyser and software by LMS International. Figures 12.13 and 12.14 show photographs of the test setup. For the nominal manufactured specimen, two different excitation points were initially tested. The accelerometer responses from these two different excitation points did not exhibit a great deal of difference however, so only the single excitation point shown in Figure 12.11 was used. Excitations in both the X and Y directions were originally examined, but did not exhibit a great deal of difference in the measured FRFs. For this reason, all excitations in both the finite element models and manufactured specimens were made in the X direction at the sole excitation point.




Figure 12.13: Experimental testing setup




Figure 12.14: Experimental testing setup

Initially the test specimens were suspended by a thick elastic cord with small hooks at each end to fit through the drilled suspension holes. After a number of specimens had been tested, it was suspected that the elastic cord was too stiff to capably represent a free-free test arrangement. The thick elastic cord was thus replaced by a much thinner cord and instead of using hooks; the ends of the cord were knotted through the suspension holes.

Figure 12.15 compares the frequency response function at the same accelerometer point and response direction for these two different suspension types. Towards the higher end of

the frequency range it can be seen that the responses are significantly different. This suggests that this region of the original thick cord response exhibits modes of the suspension hooks, as opposed to modes of the actual specimen. The earlier modal testing results were thus disposed of, and all manufactured specimens were suspended by the thinner elastic cord with no attached hooks.






Figure 12.15: Comparison of suspension types for experimental modal testing

12.7 Results

Figures 12.16 and 12.17 examine the differences in measured response at each of the ten different accelerometer points, for the case of the finite element model. Both of these figures are for the nominal model design. These two figures show all ten accelerometer FRFs in the X and Y response directions respectively. Figures 12.18 and 12.19 also examine the differences in measured response at each accelerometer location, but this time for the nominal manufactured test specimen. Figure 12.18 shows all ten accelerometer FRFs in the X response direction, and Figure 12.19 shows the same ten FRFs, but for the Y response direction.




Figures 12.16 and 12.17: X and Y direction finite element responses for all ten accelerometers



Figures 12.18 and 12.19: X and Y direction experimental responses for all ten accelerometers

In all finite element and experimental tests conducted here, the nominal design was always the design with all ten welds present, a sheet thickness of 1.0mm and spot welds located in their central positions. Figure 12.20 compares the X direction FRF for both the nominal finite element model and nominal manufactured specimen at the same accelerometer point. It can be seen that the correlation between these two curves is reasonably good, but there are certain regions of difference. The general convention for finite element modelling is to undertake model updating, where parameters of the finite element model are continually modified to improve the correlation of the computational model results with measured experimental data. In this case however there was simply not time available to undertake any model updating. It may be that a suitably thorough model updating procedure could

improve the correlation between the two responses in Figure 12.20. Figure 12.21 then shows the same comparison between computational and experimental results of the nominal specimen design, but for the Y direction response at the same accelerometer location.



Figures 12.20 and 12.21: Comparison of nominal model responses for finite element and experimental results

The work here examined the effects of the three studied uncertainties by two different measures – the range of FRF amplitudes and the range of nominal modal frequencies. The range of FRF amplitudes approach was undertaken graphically through the use of an FRF envelope. For each different combination of uncertainties, this corresponded to a number of different finite element models or manufactured specimens. FRFs were taken from this collection of models/specimens at the same response point, and a convex hull was created to contain all of these curves within the minimum possible area. The width of the FRF envelope indicates the degree of uncertainty on the measured outputs.

For the second approach, the nominal finite element model had 16 modes within the frequency range of interest (0-1024Hz), and the nominal manufactured specimen had 22 modes in the same range. For both the finite element and experimental testing regimes, it was wished to see if all of these nominal modes were present for a given arrangement of the uncertain parameters, and if so, at what modal frequencies.

Initially this was attempted by visual inspection of the measured FRFs. For the finite element models, the X and Y direction FRFs at a particular response point were compared

to those of the nominal model. For certain modes it was very easy to see which nominal modes they corresponded to, but this was less true for other modes. Certain models exhibited more modes than the nominal model, and it was often not possible to tell which modes had appeared anew.

Within a short while it became obvious that the changes in modal frequencies could not be done by visual FRF inspection alone. This method was a very subjective and unreliable approach, so it was chosen to use the FDAC instead. The FDAC generally requires the selection of a particular complete column from the frequency response matrix $[H(\omega)]$. Here, however, the use of a single excitation DOF meant that the frequency response matrix only had a single column, so the need for this choice was eliminated.

Observing all FRFs of the nominal models simultaneously, it could be seen that the frequency of each mode varied slightly over the different response DOFs. For both the nominal finite element model and manufactured specimen the modal frequencies were therefore recorded relative to one specific response location. Certain modes could be seen in the X direction FRF but not the Y direction FRF, and vice versa.

For each model/specimen other than the nominal ones, the FRF responses in both the X and Y directions were observed simultaneously at the same specific response location. The frequencies were recorded for every point of the plot which looked like it could be a mode, however unlikely; often resulting in many more modes than shown for the nominal models. The entire FDAC matrix was calculated for all response DOFs of the sample model against those of the nominal design.

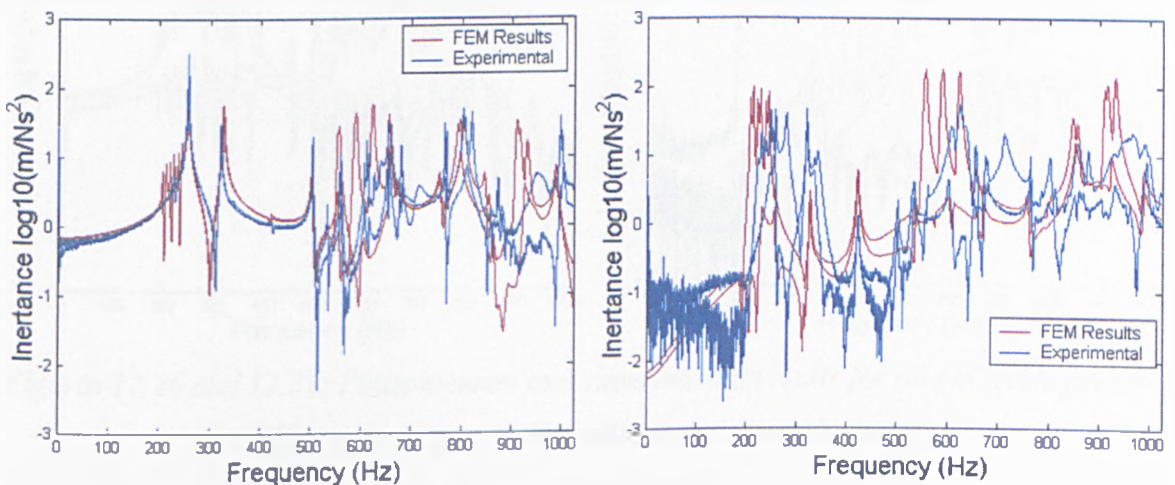
The appearance of the nominal modes in the model was decided solely based upon the FDAC index at that point. For the situation where there were several candidates for the nominal mode, the one was chosen with the highest index. A measured mode was only considered to be a candidate for a nominal mode if its respective FDAC index was at least 0.5, so that weak mode correlations were ignored. If a suitable match could not be found, that nominal mode was judged not to have appeared in the relevant model/specimen. Within the numerous tables of modal frequency results that follow, mode correlations with an FDAC index less than 0.5 are shown in these tables, but these correlations were not used in subsequent calculations of the ranges of the nominal mode frequencies.

For the case of possible missing spot welds from the component, this could be divided into four distinct situations – all welds present, any single weld missing, welds 2 and 8 missing, any two welds missing. These four situations were therefore considered separately.

12.7.1 All Welds Present

Initially all ten spot welds were used in the testing process, so that the only uncertainties present were the three different spot weld positions and three different sheet thicknesses. This corresponded to nine different model designs in total, which were available as both finite element models and manufactured test specimens.


Figures 12.22 – 12.27 show the effects of uncertainties on the measured FRF amplitudes, but for the case of no missing spot welds, i.e. all ten spot welds present. In each of these figures the respective uncertain amplitude ranges from the finite element models and experimental specimens are compared directly, and each figure is for the same response point. Figures 12.22 and 12.23 respectively show the X and Y direction FRFs for differences in spot weld positions only. The specimen thicknesses were thus held fixed at their nominal values of 1.0mm, which corresponded to using three different specimen designs.



Figures 12.22 and 12.23: Finite element and experimental results for all ten welds present, differences in weld positions only

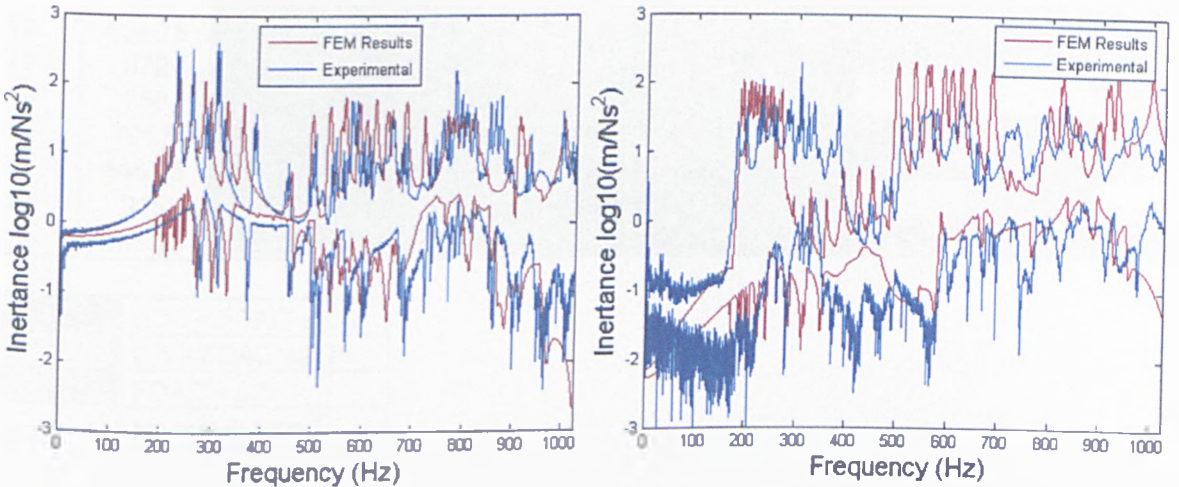
Figures 12.24 and 12.25 respectively show the X and Y direction FRFs for differences in specimen thickness only. The spot weld positions were held fixed at their central locations

in Figure 12.8, so this also corresponded to three different specimen designs.



Figures 12.24 and 12.25: Finite element and experimental results for all ten welds present, differences in sheet thickness only

Figures 12.26 and 12.27 then show the X and Y direction FRFs for differences in specimen thickness and spot weld positions. This corresponded to nine different specimen designs.



Figures 12.26 and 12.27: Finite element and experimental results for all ten welds present, differences in spot weld positions and sheet thickness

For the ranges of the nominal modal frequencies, the finite element and experimental testing results are presented separately.

12.7.1.1 Finite Element Modelling

The finite element modelling was undertaken first, with nine different models in total. As detailed previously, the FDAC was used to deduce the frequencies of the 16 nominal modes in each different model. Table 12.3 shows the entire set of finite element results for the case of all ten spot welds present. The colour coding represents the quality of the correlation between the nominal model and studied model for that particular mode, with respect to the FDAC index. A high FDAC value signifies with a high degree of confidence that this is the same mode as seen in the nominal model.

Mode Number	Nominal	0.9	0.9	0.9	1.0	1.0	1.1	1.1	1.1
		Central	Inwards	Outwards	Inwards	Outwards	Central	Inwards	Outwards
1	221.25	199.75	214.25	185.25	237.25	206.25	242.5	260	225
2	227.75	205.75	220.25	191.25	244	213	249.75	267.25	232.25
3	256	230.5	230.75	230.25	256.25	257.25	281	281.5	280.5
4	325.75	293	293.25	293	325.75	327	358.25	358.25	358.25
5	422.25	392	392	392	421.75	423	450.25	450	450.5
6	500.75	455.25	455	455	500.5	500.75	544.5	544.25	543.75
7	534.75	486	486.75	545.25	535.75	533.75	582.5	583.75	583.25
8	552.75	498.75	498.5	-	552.25	601.75	606	605.5	658.75
9	589	531.25	561.5	497.25	622.75	555.75	646.25	683.25	604.5
10	654.75	589.5	589.5	589.5	654.5	658	719.75	719.75	719.75
11	672	606	605.25	606.5	671.25	672.5	738	737	739
12	756	682	682	682	756	756.25	-	-	832.5
13	796.5	774.75	782.25	765.5	804.25	792.25	814	822.5	803.25
14	846.75	798.5	796	799.75	845.75	854.5	895.5	896	891
15	915	823.25	819.25	824	910.5	930.75	-	999	-
16	988.25	891.25	891.25	891.5	988.25	988.5	945.5	945.25	946.75

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
	No match found

Table 12.3: Entire finite element modal frequency results for all ten welds present

Table 12.4 examines the effects of difference in spot weld positions only, so that the sheet thickness is held at the nominal value of 1.0mm. This corresponded to three different finite element models, and the FDAC was once again used. A range was calculated for each of the sixteen nominal modes when subject to this position uncertainty.

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	206.25	-6.78	237.25	7.23
2	227.75	213	-6.48	244	7.14
3	256	256	0	257.25	0.49
4	325.75	325.75	0	327	0.38
5	422.25	421.75	-0.12	423	0.18
6	500.75	500.5	-0.05	500.75	0
7	534.75	533.75	-0.19	535.75	0.19
8	552.75	552.25	-0.09	601.75	8.86
9	589	555.75	-5.65	622.75	5.73
10	654.75	654.5	-0.04	658	0.50
11	672	671.25	-0.11	672.5	0.07
12	756	756	0	756.25	0.03
13	796.5	792.25	-0.53	804.25	0.97
14	846.75	845.75	-0.12	854.5	0.92
15	915	910.5	-0.49	930.75	1.72
16	988.25	988.25	0	988.5	0.03

Table 12.4: Finite element modal frequency variations for all ten welds present, differences in spot weld positions only

Table 12.5 examines the effects of differences in sheet thickness only, so that the spot weld positions were held fixed at their central values in Figure 12.8. This once again corresponded to three different finite element models.

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	199.75	-9.72	242.5	9.60
2	227.75	205.75	-9.66	249.75	9.66
3	256	230.5	-9.96	281	9.77
4	325.75	293	-10.05	358.25	9.98
5	422.25	392	-7.16	450.25	6.63
6	500.75	455.25	-9.09	544.5	8.74
7	534.75	486	-9.12	582.5	8.93
8	552.75	498.75	-9.77	606	9.63
9	589	531.25	-9.80	646.25	9.72
10	654.75	589.5	-9.97	719.75	9.93
11	672	606	-9.82	738	9.82
12	756	682	-9.79	756	0
13	796.5	774.75	-2.73	814	2.20
14	846.75	798.5	-5.70	895.5	5.76
15	915	823.25	-10.03	915	0
16	988.25	891.25	-9.82	988.25	0

Table 12.5: Finite element modal frequency variations for all ten welds present, differences in sheet thickness only

Table 12.6 then examines the combined effects of both differences in sheet thickness and differences in spot weld positions, corresponding to nine different finite element models.

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	185.25	-16.27	260	17.51
2	227.75	191.25	-16.03	267.25	17.34
3	256	230.25	-10.06	281.5	9.96
4	325.75	293	-10.05	358.25	9.98
5	422.25	392	-7.16	450.5	6.69
6	500.75	455	-9.14	544.5	8.74
7	534.75	486	-9.12	583.75	9.16
8	552.75	498.5	-9.81	658.75	19.18
9	589	497.25	-15.58	683.25	16.00
10	654.75	589.5	-9.97	719.75	9.93
11	672	605.25	-9.93	739	9.97
12	756	682	-9.79	832.5	10.12
13	796.5	765.5	-3.89	822.5	3.26
14	846.75	796	-5.99	896	5.82
15	915	819.25	-10.46	999	9.18
16	988.25	891.25	-9.82	988.5	0.03

Table 12.6: Finite element modal frequency variations for all ten welds present, differences in weld positions and differences in sheet thickness

12.7.1.2 Experimental Testing of Manufactured Specimens

Mode Number	Nominal	Specimen Number							
		1	3	5	7	11	13	15	17
1	229	333	252	-	269	197	257.25	222	181.75
2	248.75	-	279	248.25	299.25	220.5	288.75	237.5	197.25
3	256.5	304.5	256.5	230	309.75	226.25	308.5	256	226.75
4	321	380.25	321.5	289.25	383.25	283.75	383.75	321	285.75
5	335.75	364.75	344.75	333.75	346.5	329.75	340.25	334.5	330
6	422.25	483.75	421.75	391.5	619.5	386.75	561.75	422.5	389
7	502.5	593.25	502	451.75	599.5	445	595.75	503	449.25
8	532.25	725.25	679.75	-	726	488.75	723.5	618	477
9	560.25	668.25	557.5	497.75	678.25	491.75	674.5	560.25	496.75
10	605.25	732	606.75	-	733	536.25	740.25	614	542.25
11	615.75	833	766.25	-	755	818.25	733.25	607.75	817.75
12	626.25	923	674	604.75	741.25	574.75	-	-	507.25
13	639.25	817.25	-	625.25	-	548.5	730.5	618	517
14	654.25	873.25	651.5	584.75	793.5	760.75	793.75	653.75	581.25
15	674.25	-	-	601.75	821	589.75	825	677	-
16	678	781	656.25	582.75	824	674	826.75	682	604.5
17	685.5	815.5	-	892.5	829.75	567.75	819.25	879	681
18	766.25	-	610.25	-	933.5	-	933	769.25	-
19	806.75	-	820.25	783.75	862	766.75	851.75	803.25	757
20	871	925.5	679.75	774	733	596.5	-	874.5	602
21	880.25	998.75	929.75	933.5	771.75	619.75	778.25	899	585
22	996.75	783.5	770.25	686	743.5	875	823.25	1000	884.75

0.75 < FDAC ≤ 1.0
0.5 ≤ FDAC ≤ 0.75
FDAC < 0.5
-
No match found

Table 12.7: Entire experimental modal frequency results for all ten welds present

For the nominal manufactured test specimen there were 22 modes within the frequency range of interest (0 – 1024Hz). Table 12.7 shows the complete set of results from the nine manufactured specimens with all ten spot welds present.

Table 12.8 examines the effects of differences in spot weld positions only, corresponding to three manufactured specimens.

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	229	222	-3.06	252	10.04
2	248.75	237.5	-4.52	279	12.16
3	256.5	256	-0.19	256.5	0
4	321	321	0	321.5	0.16
5	335.75	334.5	-0.37	344.75	2.68
6	422.25	421.75	-0.12	422.5	0.06
7	502.5	502	-0.10	503	0.10
8	532.25	532.25	0	679.75	27.71
9	560.25	557.5	-0.49	560.25	0
10	605.25	605.25	0	614	1.45
11	615.75	607.75	-1.30	766.25	24.44
12	626.25	626.25	0	674	7.62
13	639.25	618	-3.32	639.25	0
14	654.25	651.5	-0.42	654.25	0
15	674.25	674.25	0	677	0.41
16	678	656.25	-3.21	682	0.59
17	685.5	685.5	0	879	28.23
18	766.25	610.25	-20.36	769.25	0.39
19	806.75	803.25	-0.43	820.25	1.67
20	871	679.75	-21.96	874.5	0.40
21	880.25	880.25	0	929.75	5.62
22	996.75	770.25	-22.72	1000	0.33

Table 12.8: Experimental modal frequency variations for all ten welds present, differences in weld positions only

Table 12.9 examines the effects of differences in sheet thickness only, also corresponding to three manufactured specimens.

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	229	197	-13.97	269	17.47
2	248.75	220.5	-11.36	299.25	20.30
3	256.5	226.25	-11.79	309.75	20.76
4	321	283.75	-11.60	383.25	19.39
5	335.75	329.75	-1.79	346.5	3.20
6	422.25	386.75	-8.41	619.5	46.71
7	502.5	445	-11.44	599.5	19.30
8	532.25	488.75	-8.17	726	36.40
9	560.25	491.75	-12.23	678.25	21.06
10	605.25	536.25	-11.40	733	21.11
11	615.75	615.75	0	818.25	32.89
12	626.25	574.75	-8.22	741.25	18.36
13	639.25	548.5	-14.20	639.25	0
14	654.25	654.25	0	793.5	21.28
15	674.25	589.75	-12.53	821	21.76
16	678	674	-0.59	824	21.53
17	685.5	567.75	-17.18	829.75	21.04
18	766.25	766.25	0	933.5	21.83
19	806.75	766.75	-4.96	862	6.85
20	871	596.5	-31.52	871	0
21	880.25	619.75	-29.59	880.25	0
22	996.75	743.5	-25.41	996.75	0

Table 12.9: Experimental modal frequency variations for all ten welds present, differences in sheet thickness only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	229	181.75	-20.63	333	45.41
2	248.75	197.25	-20.70	299.25	20.30
3	256.5	226.25	-11.79	309.75	20.76
4	321	283.75	-11.60	383.75	19.55
5	335.75	329.75	-1.79	364.75	8.64
6	422.25	386.75	-8.41	619.5	46.71
7	502.5	445	-11.44	599.5	19.30
8	532.25	477	-10.38	726	36.40
9	560.25	491.75	-12.23	678.25	21.06
10	605.25	536.25	-11.40	740.25	22.30
11	615.75	607.75	-1.30	833	35.28
12	626.25	507.25	-19.00	923	47.39
13	639.25	517	-19.12	817.25	27.85
14	654.25	581.25	-11.16	873.25	33.47
15	674.25	589.75	-12.53	825	22.36
16	678	582.75	-14.05	826.75	21.94
17	685.5	567.75	-17.18	892.5	30.20
18	766.25	610.25	-20.36	933.5	21.83
19	806.75	757	-6.17	862	6.85
20	871	596.5	-31.52	925.5	6.26
21	880.25	585	-33.54	933.5	6.05
22	996.75	686	-31.18	1000	0.33

Table 12.10: Experimental modal frequency variations for all ten welds present, differences in weld positions and differences in sheet thickness

Table 12.10 examines the combined effects of differences in spot weld positions and differences in sheet thickness, corresponding to nine different manufactured specimens.

12.7.2 Any Single Weld Missing

The effects of one missing spot weld were then examined. As the manufactured specimens considered only zero or two missing spot welds, all results given in this section were obtained from finite element modelling only. From consulting Figure 12.9 and using the previously mentioned symmetry of the model, any one missing spot weld corresponds to any one of the spot welds at positions 1, 2 and 3 not being present.

Nominal	Position 1 Missing								
	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	201.75	-	170.75	223.5	240	188.75	245	263.25	227.5
227.75	-	-	170.75	-	-	207.5	222.5	459	205.25
256	213.75	214.25	213.25	233.5	234.25	233	252.5	253.5	252
325.75	288	288	288	318.25	318.25	318.5	349	349	349.25
422.25	-	368.75	365	427.75	391	425	454.5	411.5	452.5
500.75	514	544.5	-	559	603.5	603.25	547.25	561	543.5
534.75	486.25	487.5	558.5	535.25	537	618.25	583.5	662	583.5
552.75	530	560.25	495.75	587.5	621.25	549.75	664.75	681.75	603
589	553	523.75	484.75	602.75	580.5	535.25	661.25	585.25	589
654.75	588.75	588.75	588.5	654	654	653.75	719	719	718.75
672	605.25	605	606	671.25	670.75	672.25	735.75	736.75	736
756	808	770.25	803.5	758.5	758.5	758.5	832.75	833	832.75
796.5	633.75	664.5	653.75	713.5	713.75	708	763.25	769.25	756.75
846.75	795.5	794.5	795	842.5	843.5	839.25	890.25	892	886
915	821.5	818.5	839.25	914	909	915.25	945.75	1000	946.25
988.25	891.25	891.5	888.75	987.5	988.25	984.5	995.25	946.25	993.5

	$0.75 < \text{FDAC} \leq 1.0$
	$0.5 \leq \text{FDAC} \leq 0.75$
	$\text{FDAC} < 0.5$
	No match found

Table 12.11: Entire finite element modal frequency results for weld position 1 missing

Tables 12.11 – 12.13 give the complete set of modal frequency results for any one missing spot weld. Taking into account the three different spot weld position levels and the three sheet thickness levels, these three tables are comprised of 27 different finite element model results. The responses used to calculate the uncertain FRF amplitudes in Figures 12.28 -

12.35 were all taken at the same response point; this being a different point to that used in section 12.7.1.

Nominal	Position 2 Missing								
	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	193.25	207	179.25	214	229.25	198.5	234.5	251.25	217.75
227.75	204.5	219	190	226.5	242.5	204.5	241.75	265.75	224.25
256	230	230.25	229.75	255.25	255.5	254.75	280.25	280.75	279.75
325.75	292.25	292.25	292.25	324.5	324.5	324.5	356.75	356.75	356.75
422.25	391.75	391.75	390.75	421.75	421.5	422	450	449.75	450
500.75	453.75	453.75	453.25	498.75	498.5	498	541.75	541.75	548.75
534.75	487.25	487	544.75	536.75	572	550.75	582.75	582.75	658.25
552.75	498.25	497.25	466.25	551.75	550.75	515.25	604.75	603.25	563.25
589	531	561.25	496.75	588.75	622.25	537.25	645.75	682.75	604
654.75	589	589	589	654.25	654.25	654.25	719.25	719	719.25
672	605.75	605.25	606.25	672	671.25	672.5	738	737.25	738.5
756	701.5	681.5	681.25	-	755.5	375	831.5	-	833.75
796.5	762	771.75	749	776.75	791	768.75	797.75	809.25	786
846.75	799.75	798.25	799.25	849.75	854	845.5	897.25	906	890.25
915	822.5	818.25	824.75	914.5	909.75	916	946.5	947	-
988.25	891.25	754.75	891.25	988.25	816.75	988.25	838.75	865.5	946.25

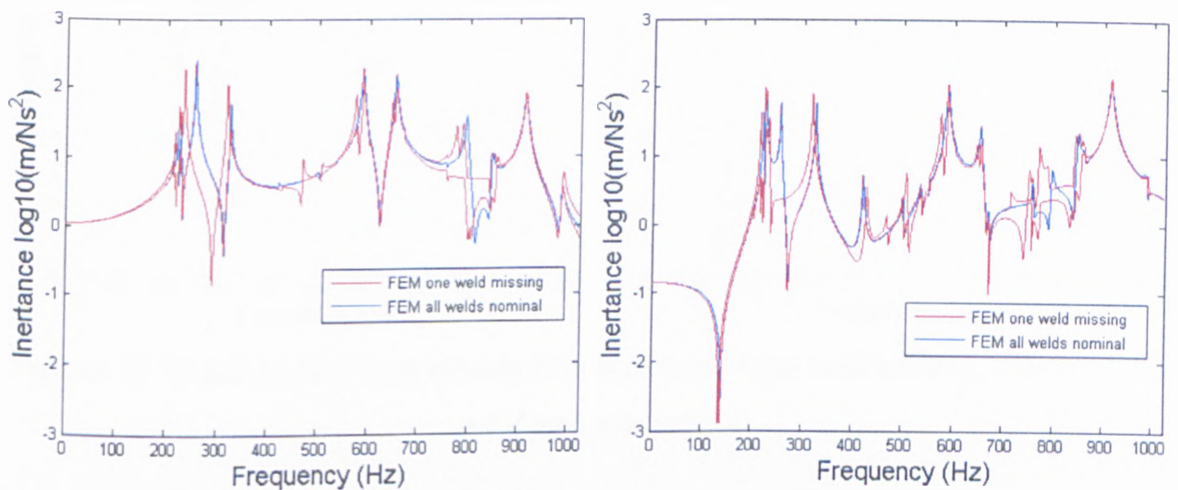
Table 12.12: Entire finite element modal frequency results for weld position 2 missing

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
	No match found

Nominal	Position 3 Missing								
	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	194.75	208.75	180.5	215.5	231	200	236.25	253.25	219.25
227.75	205.75	220.25	191.25	227.75	244	211.75	249.75	267.25	232.75
256	230.25	230.75	230	255.75	256	255.25	280.75	281.25	280.25
325.75	292.25	292.25	292.25	324.5	324.5	324.5	356.75	356.75	356.75
422.25	390.75	391	389.75	420.75	420.5	420.75	448.75	448.25	448.75
500.75	455	454.5	454.75	500.25	500	500	544	543.5	543.5
534.75	485.75	486.75	484.75	534.5	535.75	537	582.5	583.75	588.75
552.75	498.25	497.75	545	551.75	551.25	515	604.5	604	563
589	531.5	562	497.25	589.25	623	551.25	646.75	683.75	604.75
654.75	589.5	589.25	589.5	654.75	654.5	654.75	719.75	719.5	719.75
672	606	605.25	606.5	672	671.25	672.75	737.75	736.75	739.25
756	833.5	681.75	779.25	-	755.75	861	-	712	711
796.5	766.25	775.25	754.25	856.25	792	774.25	799	811.75	791
846.75	805.25	805.75	798.25	789.5	955.5	843	818.75	857.5	886
915	823	819	824.75	915	910.5	916.25	946	-	945.75
988.25	711	745.5	891.25	771	808.75	988	832	839.75	832

Table 12.13: Entire finite element modal frequency results for weld position 3 missing

Firstly the sole effects of any one missing spot weld were examined, so the specimen thickness and spot weld positions were held fixed at their respective nominal values of 1.0mm and the central position. This corresponded to three different specimen designs. Figures 12.28 and 12.29 show the X and Y direction uncertain FRFs for these tests. Added to these two figures are the responses for the nominal model with all ten welds present, so that the effect of one missing spot weld can be compared to the original case. Table 12.14 gives the range of finite element modal frequencies for the sole case of any one missing spot weld.




Figures 12.28 and 12.29: Finite element results for any single weld missing only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	214	-3.28	223.5	1.02
2	227.75	226.5	-0.55	227.75	0
3	256	233.5	-8.79	255.75	-0.10
4	325.75	318.25	-2.30	324.5	-0.38
5	422.25	420.75	-0.36	421.75	-0.12
6	500.75	498.75	-0.40	500.25	-0.10
7	534.75	534.5	-0.05	536.75	0.37
8	552.75	551.75	-0.18	587.5	6.29
9	589	588.75	-0.04	602.75	2.33
10	654.75	654	-0.11	654.75	0
11	672	671.25	-0.11	672	0
12	756	-	-	-	-
13	796.5	776.75	-2.48	776.75	-2.48
14	846.75	842.5	-0.50	849.75	0.35
15	915	914	-0.11	915	0
16	988.25	987.5	-0.08	988.25	0

Table 12.14: Finite element modal frequency variations for any single weld missing only

The effects of any one missing spot weld were then examined in conjunction with differences in weld positions. The specimen thicknesses were held fixed at their nominal values, so this corresponded to nine different specimen designs. Figures 12.30 and 12.31 show the X and Y direction uncertain FRFs for these tests, and Table 12.15 gives the range of nominal modal frequencies.




Figures 12.30 and 12.31: Finite element results for any single weld missing, differences in weld positions only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	188.75	-14.69	240	8.47
2	227.75	204.5	-10.21	244	7.14
3	256	233	-8.98	256	0
4	325.75	318.25	-2.30	324.5	-0.38
5	422.25	420.5	-0.41	422	-0.06
6	500.75	498	-0.55	500.25	-0.10
7	534.75	534.5	-0.05	618.25	15.61
8	552.75	515	-6.83	621.25	12.39
9	589	535.25	-9.13	623	5.77
10	654.75	653.75	-0.15	654.75	0
11	672	670.75	-0.19	672.75	0.11
12	756	755.5	-0.07	755.5	-0.07
13	796.5	768.75	-3.48	792	-0.56
14	846.75	839.25	-0.89	854	0.86
15	915	909	-0.66	916.25	0.14
16	988.25	816.75	-17.35	988.25	0

Table 12.15: Finite element modal frequency variations for any single weld missing, differences in weld positions only

Following this, the effects of any one missing spot weld were examined in conjunction with

differences in sheet thicknesses. This time the spot weld positions were held fixed at their central positions, and this once again corresponded to nine different specimen designs. Figures 12.32 and 12.33 show the X and Y direction uncertain FRFs for these tests, and Table 12.16 gives the range of nominal modal frequencies.




Figures 12.32 and 12.33: Finite element results for any single weld missing, differences in sheet thickness only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	193.25	-12.66	245	10.73
2	227.75	204.5	-10.21	249.75	9.66
3	256	213.75	-16.50	280.75	9.67
4	325.75	288	-11.59	356.75	9.52
5	422.25	390.75	-7.46	450	6.57
6	500.75	453.75	-9.39	544	8.64
7	534.75	485.75	-9.16	583.5	9.12
8	552.75	498.25	-9.86	644.75	16.64
9	589	531	-9.85	661.25	12.27
10	654.75	588.75	-10.08	719.75	9.93
11	672	605.25	-9.93	738	9.82
12	756	701.5	-7.21	701.5	-7.21
13	796.5	762	-4.33	799	0.31
14	846.75	795.5	-6.05	897.25	5.96
15	915	821.5	-10.22	915	0
16	988.25	891.25	-9.82	988.25	0

Table 12.16: Finite element modal frequency variations for any single weld missing, differences in sheet thickness only

The combined effects of any one missing spot weld, differences in specimen thickness and differences in spot weld positions were then examined. In total this corresponded to 27

different finite element models. Figures 12.34 and 12.35 show the X and Y direction FRFs for these tests, and Table 12.17 gives the range of nominal modal frequencies.



Figures 12.34 and 12.35: Finite element results for any single weld missing, differences in weld positions and sheet thickness

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	170.75	-22.82	263.25	18.98
2	227.75	190	-16.58	232.75	2.20
3	256	213.25	-16.70	281.25	9.86
4	325.75	288	-11.59	356.75	9.52
5	422.25	389.75	-7.70	450	6.57
6	500.75	453.25	-9.49	544	8.64
7	534.75	485.75	-9.16	662	23.80
8	552.75	466.25	-15.65	681.75	23.34
9	589	484.75	-17.70	683.75	16.09
10	654.75	588.5	-10.12	719.75	9.93
11	672	605	-9.97	738.5	9.90
12	756	681.5	-9.85	755.5	-0.07
13	796.5	749	-5.96	811.75	1.91
14	846.75	794.5	-6.17	906	7.00
15	915	818.25	-10.57	1000	9.29
16	988.25	754.75	-23.63	988.25	0


Table 12.17: Finite element modal frequency variations for any single weld missing, differences in weld positions and sheet thickness

12.7.3 Welds 2 and 8 Missing


The manufactured specimens only considered the case of two specific missing spot welds, given by weld index positions 2 and 8 on Figure 12.9. A direct comparison between finite element and experimental results for two missing welds could therefore only be made for

these particular two welds, not for any two welds.

Figure 12.36 examines the effects of removing welds 2 and 8 on the finite element response at a particular accelerometer in the X response direction. Figure 12.37 makes the same comparison, but in the Y response direction.



Figures 12.36 and 12.37: Effects on the finite element response of removing welds 2 and 8





Figures 12.38 and 12.39: Effects on the experimental response of removing welds 2 and 8



Figures 12.38 and 12.39 then examine the effects on the experimental response of removing welds 2 and 8. These two figures are again for the X and Y response directions respectively, and were calculated at the same response point as for Figures 12.36 and 12.37.

Figures 12.40 – 12.47 then compare the finite element and experimental uncertain FRF

amplitude results for the case of weld positions 2 and 8 absent. All of the FRFs shown in these figures are for the same response point, this being a different response point to those used in sections 12.7.1, 12.7.2 and Figures 12.36 – 12.39. Figures 12.40 and 12.41 directly compare the nominal finite element and nominal experimental FRFs in the X and Y directions respectively. These two figures can be compared to the nominal comparisons made in Figures 12.20 and 12.21, where all ten component spot welds were present.




Figures 12.40 and 12.41: Comparison of finite element and experimental responses for welds 2 and 8 missing



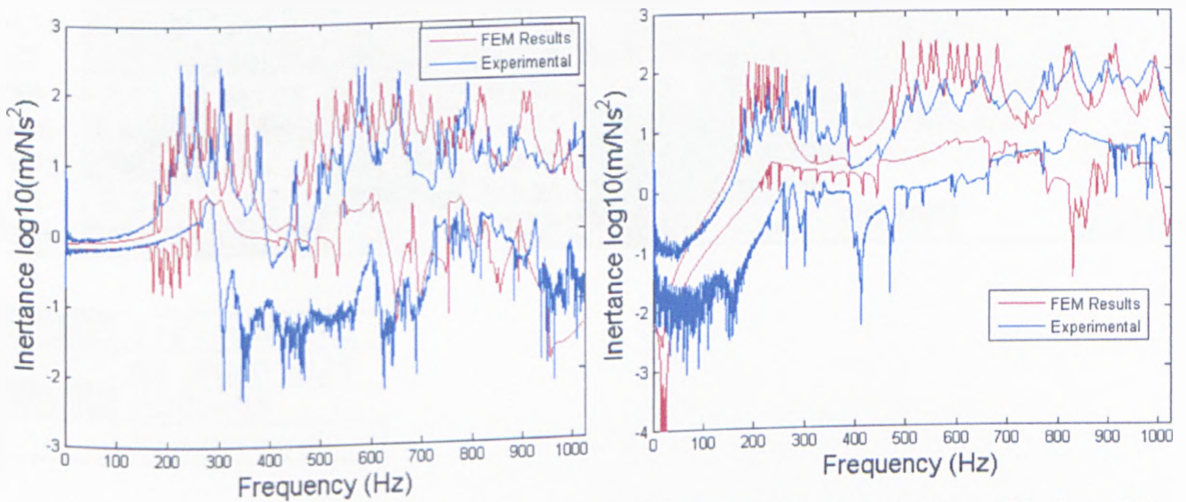
Figures 12.42 and 12.43: Finite element and experimental results for welds 2 and 8 missing, differences in weld positions only

Figures 12.42 and 12.43 then compare the X and Y direction uncertain FRF amplitudes for the combined effects of the two missing spot welds and differences in spot weld positions. The specimen thicknesses were held fixed at their nominal values of 1.0mm. This

corresponded to three different specimen designs. Figure 12.44 and 12.45 examine the combined effects of the two missing spot welds and differences in sheet thicknesses, so now the spot weld positions were held fixed at their nominal values. This once again corresponded to three different specimen designs.



Figures 12.44 and 12.45: Finite element and experimental responses for welds 2 and 8 missing, differences in sheet thickness only



Figures 12.46 and 12.47: Finite element and experimental responses for welds 2 and 8 missing, differences in weld positions and sheet thickness

Figures 12.46 and 12.47 then compare the finite element and experimental uncertain FRFs for the combined effects of the two particular missing spot welds, differences in weld positions and differences in specimen thicknesses, corresponding to nine different models. These two figures show the X and Y direction responses respectively.

For the ranges of the nominal modal frequencies, the finite element and experimental results for welds 2 and 8 missing are given separately.

12.7.3.1 Finite Element Modelling

Table 12.18 gives the complete set of finite element modal frequency results for welds 2 and 8 missing. Table 12.19 shows the range of nominal modal frequencies for differences in spot weld positions only.

Nominal	Positions 2 and 8 Missing								
	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	188.75	202.25	175	209	224	194	229	245.25	212.5
227.75	204.25	218.5	189.75	226	242	210.25	247.75	265.25	230.5
256	229.75	230.25	229.5	255	255.25	254.5	279.75	280.25	279.25
325.75	291.5	291.5	291.5	323.5	323.5	323.5	355.25	355.25	355.25
422.25	390.75	390.75	389.5	420.5	420.5	420.5	448.5	448	448.5
500.75	453.5	453.25	453.25	498.25	498	487.75	541.25	540.75	540.5
534.75	486.25	486.75	473.25	535.75	536	523.75	584	582.75	573.5
552.75	390.75	496.5	544.25	525.75	549.5	603.5	604	601.75	606.5
589	529.5	559.5	495.5	587	620.25	549	644	680.75	602.25
654.75	588.75	588.5	588.75	653.5	653.5	653.75	718.25	718	722.5
672	605.75	605.25	606	671.75	671.25	672.25	737.75	737	738.25
756	707.25	681.75	681	755.75	755.5	882.25	982	945	718.25
796.5	754.75	744	740	766.75	782.75	758.75	788.75	800.25	765.25
846.75	737.25	767	800	793.25	818	843.25	849.5	866	885.75
915	825.25	817	824.75	915.25	903.75	915.75	-	-	-
988.25	719	902.5	891.5	988.75	810.5	988	947	-	947

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
-	No match found

Table 12.18: Entire finite element modal frequency results for welds 2 and 8 missing

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	194	-12.32	224	1.24
2	227.75	210.25	-7.68	242	6.26
3	256	254.5	-0.59	255.25	-0.29
4	325.75	323.5	-0.69	323.5	-0.69
5	422.25	420.5	-0.41	420.5	-0.41
6	500.75	497.75	-0.60	498.25	-0.50
7	534.75	535.75	0.19	536	0.23
8	552.75	525.75	-4.88	549.5	-0.59
9	589	549	-6.79	620.25	5.31
10	654.75	653.5	-0.19	653.75	-0.15
11	672	671.25	-0.11	672.25	0.04
12	756	755.5	-0.07	755.5	-0.07
13	796.5	758.75	-4.74	782.75	-1.73
14	846.75	793.25	-6.32	843.25	-0.41
15	915	903.75	-1.23	915.75	0.08
16	988.25	988	-0.03	988	-0.03

Table 12.19: Finite element modal frequency variations for welds 2 and 8 missing, differences in weld positions only

Table 12.20 gives the nominal modal frequency ranges for differences in sheet thickness only. Table 12.21 examines the combined effects of differences in spot weld positions and differences in sheet thickness on the nominal modal frequencies.

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	188.75	-14.69	229	3.50
2	227.75	204.25	-10.32	247.75	8.78
3	256	229.75	-10.25	279.75	9.28
4	325.75	291.5	-10.51	355.25	9.06
5	422.25	390.75	-7.46	448.5	6.22
6	500.75	453.5	-9.44	541.25	8.09
7	534.75	486.25	-9.07	584	9.21
8	552.75	390.75	-29.31	604	9.27
9	589	529.5	-10.10	644	9.34
10	654.75	588.75	-10.08	718.25	9.70
11	672	605.75	-9.86	737.75	9.78
12	756	707.25	-6.45	707.25	-6.45
13	796.5	754.75	-5.24	786.75	-1.22
14	846.75	793.25	-6.32	849.5	0.32
15	915	825.25	-9.81	915.25	0.03
16	988.25	-	-	-	-

Table 12.20: Finite element modal frequency variations for welds 2 and 8 missing, differences in sheet thickness only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	175	-20.90	245.25	10.85
2	227.75	189.75	-16.68	265.25	16.47
3	256	229.5	-10.35	280.25	9.47
4	325.75	291.5	-10.51	355.25	9.06
5	422.25	389.5	-7.76	448.5	6.22
6	500.75	453.25	-9.49	541.25	8.09
7	534.75	486.25	-9.07	584	9.21
8	552.75	390.75	-29.31	604	9.27
9	589	495.5	-15.87	680.75	15.58
10	654.75	588.5	-10.12	722.5	10.35
11	672	605.25	-9.93	738.25	9.86
12	756	681.75	-9.82	755.5	-0.07
13	796.5	740	-7.09	800.25	0.47
14	846.75	767	-9.42	885.75	4.61
15	915	817	-10.71	915.75	0.08
16	988.25	891.5	-9.79	988	-0.03

Table 12.21: Finite element modal frequency variations for welds 2 and 8 missing, differences in weld positions and sheet thickness

12.7.3.2 Experimental Testing of Manufactured Specimens

Nominal	Specimen Number								
	2	4	6	8	10	12	14	16	18
229	272.75	242	251.75	268.5	225.75	188.75	346	203.25	181.75
248.75	315	278.25	230.25	305.75	250.25	220.75	298	229	203.75
256.5	306.75	256.75	229.25	309.75	256	226	304	256	226.25
321	382.25	321.75	289	382.25	319.5	283	370.25	319.5	286
335.75	349.75	338.75	333.25	346.5	334.25	328.75	-	332	476.25
422.25	424.5	420.75	390.25	577.75	475.25	387.5	618	513	389
502.5	594.75	501.5	452	595	502	446	589.25	502.25	448.75
532.25	-	532.75	335	726.5	532.25	520.25	667.25	534.5	-
560.25	668.5	558	497	670.25	558	491.75	816.75	611.25	495.5
605.25	722.75	608	540	-	603	537	777.5	588.5	544.25
615.75	-	-	681	-	-	561.75	800.75	-	843
626.25	-	682	-	-	612.75	577.25	787.75	570.25	523
639.25	785	-	886.75	772.5	628	765.5	812.75	560.75	508
654.25	790.25	655.25	780	792.25	656.25	571.5	771	653.5	575.25
674.25	-	-	778	818.5	675	753.5	-	679	605.25
678	815	-	604.25	819.5	677.5	572	816.75	682.5	607
685.5	-	723.5	897.25	824.5	633.75	598.5	821	880	834.25
766.25	924	-	543.25	786.75	765.25	-	-	770.75	684.5
806.75	878.25	818.75	588.25	854.5	801.25	596.5	834.25	778.5	750.5
871	-	-	1022	-	-	1005.5	-	658.5	777.5
880.25	815	771.25	858.25	828.5	900.75	827.25	921.25	723.25	843
996.75	733.5	-	685	928.5	-	573.75	733	1000	1022.25

Table 12.22: Entire experimental modal frequency results for welds 2 and 8 missing

Table 12.22 gives the complete set of experimental modal frequency results for the case of welds 2 and 8 missing. Table 12.23 gives the nominal modal frequency ranges for differences in spot weld positions only, and Table 12.24 examines the effects of differences in sheet thickness only.

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	229	203.25	-11.24	242	5.68
2	248.75	229	-7.94	278.25	11.86
3	256.5	256	-0.19	256.75	0.10
4	321	319.5	-0.47	321.75	0.23
5	335.75	332	-1.12	338.75	0.89
6	422.25	420.75	-0.36	475.25	12.55
7	502.5	501.5	-0.20	502.25	-0.05
8	532.25	532.25	0.00	534.5	0.42
9	560.25	558	-0.40	611.25	9.10
10	605.25	589.5	-2.60	608	0.45
11	615.75	-	-	-	-
12	626.25	570.25	-8.94	682	8.90
13	639.25	560.75	-12.28	628	-1.76
14	654.25	653.5	-0.11	656.25	0.31
15	674.25	675	0.11	679	0.70
16	678	677.5	-0.07	682.5	0.66
17	685.5	633.75	-7.55	880	28.37
18	766.25	765.25	-0.13	770.75	0.59
19	806.75	779.5	-3.38	819.75	1.61
20	871	658.5	-24.40	658.5	-24.40
21	880.25	723.25	-17.84	900.75	2.33
22	996.75	1000	0.33	1000	0.33

Table 12.23: Experimental modal frequency variations for welds 2 and 8 missing, differences in weld positions only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	229	188.75	-17.58	268.5	17.25
2	248.75	220.75	-11.26	305.75	22.91
3	256.5	226	-11.89	309.75	20.76
4	321	283	-11.84	382.25	19.08
5	335.75	328.75	-2.08	346.5	3.20
6	422.25	387.5	-8.23	475.25	12.55
7	502.5	446	-11.24	595	18.41
8	532.25	520.25	-2.25	726.5	36.50
9	560.25	491.75	-12.23	670.25	19.63
10	605.25	537	-11.28	603	-0.37
11	615.75	561.75	-8.77	561.75	-8.77
12	626.25	577.25	-7.82	612.75	-2.16
13	639.25	628	-1.76	772.5	20.84
14	654.25	571.5	-12.65	792.25	21.09
15	674.25	675	0.11	818.5	21.39
16	678	672	-0.88	819.5	20.87
17	685.5	589.5	-14.00	824.5	20.28
18	766.25	765.25	-0.13	786.75	2.68
19	806.75	596.5	-26.06	854.5	5.92
20	871	1005.5	15.44	1005.5	15.44
21	880.25	827.25	-6.02	900.75	2.33
22	996.75	573.75	-42.44	928.5	-6.85

Table 12.24: Experimental modal frequency variations for welds 2 and 8 missing, differences in sheet thickness only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	229	181.75	-20.63	346	51.09
2	248.75	203.75	-18.09	315	26.63
3	256.5	226	-11.89	309.75	20.76
4	321	283	-11.84	382.25	19.08
5	335.75	328.75	-2.08	476.25	41.85
6	422.25	387.5	-8.23	618	46.36
7	502.5	446	-11.24	595	18.41
8	532.25	335	-37.06	726.5	36.50
9	560.25	491.75	-12.23	816.75	45.78
10	605.25	537	-11.28	777.5	28.46
11	615.75	561.75	-8.77	843	36.91
12	626.25	523	-16.49	787.75	25.79
13	639.25	508	-20.53	886.75	38.72
14	654.25	571.5	-12.65	792.25	21.09
15	674.25	605.25	-10.23	818.5	21.39
16	678	604.25	-10.88	819.5	20.87
17	685.5	589.5	-14.00	897.25	30.89
18	766.25	543.25	-29.10	786.75	2.68
19	806.75	588.25	-27.08	879.25	8.99
20	871	658.5	-24.40	1022	17.34
21	880.25	723.25	-17.84	921.25	4.66
22	996.75	573.75	-42.44	1022.25	2.56


Table 12.25: Experimental modal frequency variations for welds 2 and 8 missing, differences in weld positions and sheet thickness

Table 12.25 gives the combined effects of differences in spot weld positions and differences in sheet thickness on the nominal modal frequencies.

12.7.4 Any Two Welds Missing

The symmetry of the model design could once again be used here for the case of any two missing spot welds. As given in Table 12.2, there are 15 different arrangements for the case of any two missing spot welds. The results for any two missing welds were obtained from finite element modelling only, since the manufactured specimens only omitted the same two weld positions. With the three different spot weld position arrangements and the three different levels of sheet thickness, this gave 135 different finite element models in total. There is insufficient space to show the complete set of modal frequency results from these 135 models here, but they are given in Appendix B.

The uncertain FRFs given in Figures 12.48 – 12.55 were all taken at the same response point, this being a different point to the four locations used previously. Figures 12.48 and 12.49 show the X and Y direction uncertain FRFs for any two missing spot welds, but with the same specimen thickness and spot weld positions, i.e. their nominal values. This therefore corresponded to 15 different finite element models. The nominal modal frequency ranges for any two missing welds only are given in Table 12.26.




Figures 12.48 and 12.49: Finite element responses for any two welds missing only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	198.75	-10.17	224	1.24
2	227.75	195.75	-14.05	227.75	0
3	256	216	-15.63	255.5	-0.20
4	325.75	315.25	-3.22	323.75	-0.61
5	422.25	419.5	-0.65	421.5	-0.18
6	500.75	486.75	-2.80	503	0.45
7	534.75	524.5	-1.92	612.5	14.54
8	552.75	516.25	-6.60	588.5	6.47
9	589	535.75	-9.04	611.25	3.78
10	654.75	602.5	-7.98	654.75	0
11	672	669.25	-0.41	672.25	0.04
12	756	685.25	-9.36	757.25	0.17
13	796.5	689.25	-13.47	768.75	-3.48
14	846.75	781.75	-7.68	852.25	0.65
15	915	862.5	-5.74	953	4.15
16	988.25	749	-24.21	988.75	0.05

Table 12.26: Finite element modal frequency variations for any two welds missing only

Figures 12.50 and 12.51 are X and Y direction uncertain FRFs for the combined effects of any two missing spot welds and differences in spot weld positions, but with the nominal specimen thickness. This corresponded to 45 different finite element models. The corresponding range of nominal modal frequencies for these 45 models is given in Table 12.27.




Figures 12.50 and 12.51: Finite element responses for any two welds missing, differences in weld positions only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	183.75	-16.95	240	8.47
2	227.75	180	-20.97	244	7.14
3	256	215.5	-15.82	255.75	-0.10
4	325.75	315.25	-3.22	323.75	-0.61
5	422.25	419.25	-0.71	421.5	-0.18
6	500.75	485.25	-3.10	503	0.45
7	534.75	511.75	-4.30	620.25	15.99
8	552.75	516.25	-6.60	642	16.15
9	589	522	-11.38	646.25	9.72
10	654.75	602.5	-7.98	657.5	0.42
11	672	669	-0.45	686.25	2.12
12	756	657.75	-13.00	757.5	0.20
13	796.5	678.75	-14.78	784.25	-1.54
14	846.75	781.75	-7.68	852.25	0.65
15	915	812.75	-11.17	953	4.15
16	988.25	749	-24.21	989	0.08

Table 12.27: Finite element modal frequency variations for any two welds missing, differences in weld positions only

The combined effects of any two missing spot welds and differences in sheet thickness were then examined, so the spot welds were held fixed at their central positions. This once again corresponded to 45 different finite element models. Figures 12.52 and 12.53 show the uncertain FRFs for these models in the respective X and Y response directions, and Table 12.28 gives the range of nominal modal frequencies.




Figures 12.52 and 12.53: Finite element responses for any two welds missing, differences in sheet thickness only

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	179.75	-18.76	245.5	10.96
2	227.75	177.5	-22.06	249.75	9.66
3	256	200	-21.88	280.5	9.57
4	325.75	277.5	-14.81	355.5	9.13
5	422.25	389.75	-7.70	449.5	6.45
6	500.75	446	-10.93	546	9.04
7	534.75	478.25	-10.57	671.75	25.62
8	552.75	390.75	-29.31	681.5	23.29
9	589	496.25	-15.75	670.5	13.84
10	654.75	583.25	-10.92	720	9.97
11	672	603.5	-10.19	737.75	9.78
12	756	595.25	-21.26	831	9.92
13	796.5	623.75	-21.69	787	-1.19
14	846.75	781.75	-7.68	898.75	6.14
15	915	775.5	-15.25	953.75	4.23
16	988.25	681.5	-31.04	988.75	0.05

Table 12.28: Finite element modal frequency variations for any two welds missing, differences in sheet thickness only

Finally the combined effects of any two missing spot welds, differences in specimen thickness and differences in spot weld positions were examined. This corresponded to a huge 135 different finite element models. Figures 12.54 and 12.55 show the X and Y direction uncertain FRFs for these tests, and Table 12.29 gives the range of nominal modal frequencies.



Figures 12.54 and 12.55: Finite element responses for any two welds missing, differences in weld positions and sheet thickness

Mode Number	Nominal	Lower Bound	% Difference	Upper Bound	% Difference
1	221.25	166.5	24.75	263	18.87
2	227.75	150.75	33.81	267.25	17.34
3	256	199.5	22.07	281	9.77
4	325.75	277.5	14.81	355.75	9.21
5	422.25	358.25	15.16	449.75	6.51
6	500.75	401.75	19.77	546	9.04
7	534.75	436.25	18.42	676	26.41
8	552.75	390.75	29.31	708.25	28.13
9	589	418.25	28.99	683.75	16.09
10	654.75	522.5	20.20	722.5	10.35
11	672	539.25	19.75	797.75	18.71
12	756	591.5	21.76	831.25	9.95
13	796.5	620.25	22.13	993	24.67
14	846.75	757.5	10.54	898.75	6.14
15	915	730.25	20.19	1013	10.71
16	988.25	681.5	31.04	993.25	0.51

Table 12.29: Finite element modal frequency variations for any two welds missing, differences in weld positions and sheet thickness

12.8 Discussion and Conclusions

In this chapter an extensive study was conducted of the effects of uncertainties within a spot-welded component. This component itself was a basic “double omega” beam with symmetry in all three orthogonal directions, so was a somewhat idealised structure. The spot-welded joints introduced nonlinearity and complexity into the design, however, to give a more substantial problem.

In contrast with the previous five thesis chapters, the uncertainty approach used in this chapter was much more straightforward, consisting of a simple design-of-experiments approach. The advantage of this chapter though is the range of tests that were conducted, both numerical and experimental. 171 different finite element models were used here and 18 different manufactured specimens, so that all the selected uncertainties were able to be examined thoroughly. The experimental tests conducted here were also in contrast to the simulated problems of the previous five chapters, and allowed a direct comparison between numerical and experimental results, something that had not been possible in the previous chapters.

This chapter presented a sizeable body of results, beginning with Figures 12.16 – 12.19. These four figures examined all ten responses simultaneously, for both the X and Y

directions, and for the finite element and experimental tests. All four of these figures show that, in general, the specimen modes can be seen at all of the response locations.

Figures 12.20 and 12.21 show that the agreement between the nominal finite element model and nominal manufactured specimen responses is reasonably good, but could be improved. Unfortunately there was not time available for any model updating, as this could have led to an improved agreement between the two.

The nominal manufactured specimen exhibited many more modes than that of the nominal finite element model, with 22 compared to 16. These two sets of modes were compared to one another with the FDAC, in order to see which 16 of the experimental modes corresponded to the finite element modes, and which could be ignored. The correlation between modes with the FDAC was very poor, however, with no pair of modes giving an index greater than 0.5. For this reason, an unequal number of modes had to be used for the nominal finite element and experimental models.

In relation to the presence of welds in the component, four different sets of circumstances were examined; all welds present, any single weld missing, welds 2 and 8 missing, and any two welds missing. Finite element models were tested for all four of these situations, but only two of the situations were verified by experimental results – all welds present and welds 2 and 8 missing.

Figures 12.22 – 12.27 examined the range of FRF amplitudes for different combinations of uncertainties with all ten welds present. In Figure 12.22 the agreement between the finite element and experimental uncertain ranges is good, especially below 500Hz. In Figure 12.23 the match between these two ranges is not quite as good, but still of a reasonable quality. In Figure 12.22 the actual ranges of the finite element and experimental predictions are similar, whereas for Figure 12.23 the experimental range appears to be the wider of the two.

For Figures 12.24 and 12.25 the uncertain FRF ranges are wider than in Figures 12.22 and 12.23, suggesting that differences in sheet thickness is a more significant uncertainty than differences in spot weld position. As with Figures 12.22 and 12.23, the finite element and experimental X direction responses in Figure 12.24 match better than the Y direction

responses in Figure 12.25.

In Figures 12.26 and 12.27 the uncertain FRFs are noticeably wider than previously, for the case of both the experimental and finite element tests. This implies that a combination of spot weld position and sheet thickness uncertainties has a greater effect on the response than either of the two individually, a result that could be deduced intuitively. Once again the finite element and experimental X direction responses are matched better than the Y direction responses. This may be related to the excitations only being made in the X direction.

Table 12.3 gives the complete set of finite element modal frequency results for the case of all ten spot welds present. It can be seen that for each different combination of sheet thickness and spot weld positions, the correlation between the sample modes and those of the nominal model is good. There are very few poorly correlated modes or modes where a match was not possible, and these tend to be for the same few nominal modes. The correlations in general are very good, with FDAC indices greater than 0.75.

Table 12.4 examines the effects of differences in spot weld positions for the finite element model. It can be seen that some modal frequencies are susceptible to the positions of the spot welds (modes 1, 2 and 9), but others are very unsusceptible to them. This is in contrast to Table 12.5, where virtually all of the nominal modes are susceptible to a change in sheet thickness. These results are in agreement with Figures 12.22 – 12.27, where sheet thickness uncertainties had a greater effect on the FRF responses than weld position uncertainties. Table 12.6 then shows that in every case, the combined effects of differences in spot weld positions and sheet thickness is greater than either of the two individually. This is in agreement with the finite element uncertain FRF plots in Figures 12.26 and 12.27, where the same conclusion can be drawn.

Table 12.7 gives the complete set of modal frequency results for experimental tests with all ten welds present. Although the correlation of the various specimen modes with the set of nominal modes is still good, in comparison with the corresponding finite element case in Table 12.3, the results are slightly poorer. There are many more nominal modes that cannot be matched, even considering that there are more nominal experimental modes than nominal finite element modes. There are also more average quality correlations, where the

FDAC index is between 0.5 and 0.75. This poorer nominal mode correlation for the experimental case may be related to the wider experimental FRF bounds in Figures 12.22 – 12.27, especially the Y direction responses.

As with the finite element results, there are some nominal modes that are susceptible to changes in spot weld positions alone, but other modes that are virtually unaffected by it. Also as with the finite element results, all nominal mode frequencies are affected by a change in sheet thickness, and the combined effects of changes in sheet thickness and spot weld positions is greater than either of their effects individually. This is once again in agreement with the experimental uncertain FRF bounds in Figures 12.22 – 12.27.

The second situation for weld availability was then examined, with any single missing spot weld. This situation was studied by finite element modelling only, and the complete set of modal frequency results is given in Tables 12.11 – 12.13. As explained previously, this corresponded to one of three spot welds absent, and it can be seen from the tables that the quality of correlation with the nominal modes is generally dependent on which weld index is missing.

Of the three missing weld positions, the correlation between specimen modes and nominal modes is best when position index 2 is missing (see Figure 12.9). The modal correlations become worse if the missing weld is at position 3, and are even worse if position number 1 is omitted. The results for position 1 or 3 missing also show a poor correlation is almost always found for particular nominal modes; mode numbers 5, 6, 12 and 13 for position 1 absent; mode number 12 for position 3 missing. This is most likely due to those particular modal vectors being especially dependent on the spot welds that have been removed, i.e. as nodes of vibration. Without these welds present, the modal response is much less significant, and therefore the correlation with the nominal model FRFs becomes worse.

Figures 12.28 and 12.29 examined the sole effects of any one missing spot weld on the finite element responses. In both the X and Y direction responses it can be seen that the range of FRF amplitude is very small, and that the correlation with the nominal model is very good. The conclusion from this is that the individual effect of a single missing weld is small. This is reinforced by modal frequency results in Table 12.14, where the frequency ranges of the nominal modes are small.

Figures 12.30 and 12.31 then suggest that the combined effects of any single missing weld and differences in spot weld position are much more significant than those of any single missing weld alone. This is confirmed by comparing the nominal modal frequency ranges in Table 12.15 with those in Table 12.14. Comparing Figures 12.32 and 12.33 with Figures 12.30 and 12.31, or the modal frequency ranges in Table 12.16 with those in Table 12.15, these results suggest that differences in sheet thickness is more significant than differences in spot weld positions, as had also been found with all ten component spot welds present. Figures 12.34 and 12.35, and Table 12.17 show that, as with all ten welds present, differences in sheet thickness and spot weld positions have a greater effect on the response than either of the two individually.

The third examined situation for weld availability was welds 2 and 8 absent, and this also enabled a comparison between finite element and experimental results. Figures 12.36 and 12.37 examine the effects on the response of removing welds 2 and 8 from the nominal finite element model. Although the amplitude of the FRF plots changes with the removal of these two welds, it can be seen that many of the modal frequencies remain virtually the same. The same can be seen for the removal of welds 2 and 8 from the nominal experimental specimen in Figures 12.38 and 12.39; that the modal frequencies remain much the same.

Figures 12.40 and 12.41 then compared the finite element and experimental responses in the X and Y directions for welds 2 and 8 missing. If these two figures are compared with Figures 12.20 and 12.21 where all ten welds were present, it can be seen that the removal of welds 2 and 8 has made the correlation between the finite element and experimental results worse, especially in the X response direction.

Figures 12.42 and 12.43 compared both the finite element and experimental uncertain FRFs for differences in spot weld positions and the two missing spot welds. These two figures can be compared with Figures 12.22 and 12.23 when all ten spot welds were present, and it can once again be observed that the correlation between finite element and experimental results has become worse. This is again truer in the X response direction than the Y direction.

Figures 12.44 and 12.45 compare the finite element and experimental uncertain FRFs for

the two missing welds and differences in sheet thickness. It can be seen that the range of the uncertain FRFs is greater than in Figures 12.42 and 12.43. This suggests that as previously with all ten welds present and any single weld missing, the sheet thickness uncertainty has a greater effect on the results than the spot weld positions uncertainty. Figures 12.46 and 12.47 reinforce that, as observed previously for other weld availability situations, the combined effects of differences in weld positions and differences in sheet thickness is greater than either of the two individually.

Table 12.18 gave the complete set of finite element modal frequency results for welds 2 and 8 missing. If this table is compared to Table 12.3 where all ten welds are present, it can be seen that the FDAC correlation indices between the measured and nominal modes have become worse, with many more poor correlations. This reinforces the behaviour observed in Figures 12.36 and 12.37, where removing welds 2 and 8 changed the measured FRFs of the nominal finite element model.

The nominal modal frequency ranges in Tables 12.19 and 12.20 show that differences in spot weld positions are less significant on the finite element response than differences in sheet thickness, which is in agreement with Figures 12.42 – 12.45, and has been observed in other results here. Comparing Table 12.21 with Tables 12.19 and 12.20 shows that the combined effects of differences in spot weld positions and differences in sheet thickness is greater than either of the two individually, which is in agreement with Figures 12.46 and 12.47.

Table 12.22 gives the complete set of experimental modal frequencies for welds 2 and 8 missing. A direct comparison with Table 12.7 where all ten welds were present shows that the removal of the two welds has made the FDAC correlation indices with the nominal experimental modes worse. This is as with the finite element modal frequencies in Table 12.19. The difference however is the extent of the nominal mode correlations. If Tables 12.3 and 12.19 are compared with Tables 12.7 and 12.22, it can be seen that the correlation with the nominal modes has deteriorated more for the finite element models than the experimental specimens. In other words, removing welds 2 and 8 from the finite element model seems to have had a greater effect than removing them from the experimental specimens. This can be confirmed by re-examining Figures 12.36 – 12.39. There appears to be a greater change between the nominal and measured FRFs for the finite

element responses in Figures 12.36 and 12.37 than for the experimental responses in Figures 12.38 and 12.39.

Tables 12.23 – 12.25 then restate the conclusions made from the finite element modal frequency ranges in Tables 12.19 – 12.21, i.e. that differences in sheet thickness are more significant than differences in spot weld positions, and that the combined effects of these two uncertainties are greater than either of the two individually.

The fourth and final examined situation for weld availability was then any two welds missing. These results were obtained from finite element modelling only, as the manufactured specimens only dealt with a specific pair of missing spot welds. As given previously in Table 12.2, the symmetry of the component design gave 15 different arrangements for two missing spot welds. Combined with the spot weld position and sheet thickness uncertainties this gave 135 different finite element models, so the majority of the 171 finite element models were with two missing spot welds.

Figures 12.48 and 12.49 gave the sole effects of any two missing spot welds on the FRF responses. If these two figures are compared with the sole effects of a single missing weld in Figures 12.28 and 12.29, it can be seen that the range of FRF amplitudes for any two missing welds is much greater. The uncertain FRFs in Figures 12.48 and 12.49 were calculated from 15 different models, however, compared to the three models used in Figures 12.28 and 12.29. Table 12.26 gives the modal frequency ranges for any two missing welds only. If this table is compared with Table 12.14 for any single missing weld, it can be seen that the modal frequency range for two missing welds is greater, which reinforces the uncertain FRF observations made for Figures 12.48 and 12.49. Thus, removing a greater number of spot welds from the model leads to a greater variation of the model response

Figures 12.50 and 12.51 gave the combined effects of any two missing welds and differences in spot weld positions on the finite element uncertain FRFs. The range of FRF amplitudes is greater here than in Figures 12.48 and 12.49. Table 12.27 also shows that differences in spot weld positions give an additional increase in the modal frequency ranges when compared to Table 12.26.

As with many previous results, Figures 12.52 and 12.53 show that differences in sheet thickness have a greater effect on the FRF responses than differences in spot weld positions. Table 12.28 confirms this for the nominal modal frequency ranges. Figures 12.54 and 12.55 then show that the combined effects of differences in spot weld positions and differences in sheet thickness for any two missing spot welds are greater than either of the two individually. By this point, however, the amount of uncertainty within the models is so great that the uncertain FRF plots in Figures 12.52 – 12.55 give results that are not actually very meaningful, other than to illustrate the amount of uncertainty that is present in the predicted responses. This is also true for the nominal modal frequency ranges in Table 12.29, where the variation of these modal frequencies is now quite high in relative terms.

The technique of estimating FRF amplitude ranges by using a prediction envelope worked reasonably well here. This was a simple technique that could be implemented quickly and efficiently. For a high number of uncertainties, however, i.e. any two welds missing, the amplitude ranges predicted by this approach are very wide, giving vague results where individual modes in the uncertain FRF plots cannot be distinguished.

The use of the FDAC to calculate the modal frequency variations proved to be an effective technique here. A method was needed to detect the appearance of the nominal modes in each different finite element model or manufactured specimen that was tested, and as stated previously, attempting this solely by visual FRF inspection was much too subjective an approach. The majority of finite element models and experimental specimens examined had more modes than those of the nominal cases, so the FDAC enabled the pruning of these new surplus modes.

For the frequent problem where several of the measured sample modes were possible matches for the nominal modes, the FDAC indices of these matches enabled a quick decision to be made on the correct mode to select. All 20 of the response DOFs (X and Y responses at 10 different accelerometer locations) were used in the calculation of the FDAC, so there was no redundant information in the multitude of FRFs that were measured.

As already outlined in section 12.4, simplifications were made for both the spot weld positions and sheet thickness uncertainty. Although these two uncertainties were somewhat

unrealistic for a real-life case, they were necessary for the uncertainty study undertaken here to be practical. It has been suggested in the past that there are more effective modelling representations of finite element spot weld joints than the Nastran CWELD element. Complex arrangements involving several rigid bar connectors can be used for each separate weld, but this leads to more time needed for mesh generation.

For the work here, the 2.5mm displacement in spot weld positions was exactly equal to the mesh element size. One finite element model was initially created with no weld elements, and this was used as a template for all of the finite element models used here. The use of the CWELD element meant that the weld positions could be changed very quickly for each different model, and the sheet thickness uncertainty involved a simple change of the overall material parameters.

This chapter has seen an extensive investigation into the dynamic effects of three selected uncertainties for a spot-welded steel component, using both numerical and experimental tests. The overall conclusions that can be made from the results obtained are:

- Of the three uncertainties studied here, the sheet thickness is the most important in relation to the effects on both the FRF amplitudes and the nominal modal frequencies.
- The spot weld positions uncertainty is less important, and the number of missing welds parameter is the least important of the three.
- Combinations of parameter uncertainties have a greater effect on the model response than using them individually.
- In relation to the FDAC, the finite element FRF correlation with the nominal model is worse for any one missing spot weld than for all welds present
- For the finite element models, the range of the nominal modal frequencies is greater for any two missing welds than for any one missing weld
- The effect of removing welds 2 and 8 on the FDAC FRF correlation with the nominal model is worse for the finite element case than the experimental case.

The work here was undertaken as part of the MADUSE project: a European Union research and training network comprising nine different partner institutions, some of them academic and others industrial.

MADUSE
Contract number MRTN-CT-2003-505164

RATHOISIE UNIVERSITEIT LEUVEN

lms LMS INTERNATIONAL
Empowering Engineering Innovation

INSTITUT FOR MECHANIK

uic

RENAULT

PSA

CENTRO RICERCHE FIAT

isvr

Modelling Product variability and Data Uncertainty in Structural Dynamics Engineering

13. Conclusions and Further Work

13.1 Conclusions

The title of this thesis was chosen to be deliberately broad, in order to incorporate a wide range of work. Uncertainty propagation and quantification is an important topic to consider within structural dynamics, although it also crosses over to a number of other academic disciplines. This thesis has thoroughly examined the effects of uncertainty, both in simulated problems and real-life examples.

As outlined in the thesis introduction, uncertainty can be categorised into two different types: aleatory uncertainty (often called “variability”) and epistemic uncertainty. Aleatory uncertainty is the inherent variation of a quantity; epistemic uncertainty arises from a lack of knowledge of the studied system or process. Both of these two different types can have a significant effect on the behaviour of a system, so need to be considered carefully. The traditional “factor of safety” design approach involving a nominal model is often overly conservative for engineering applications, so it is more beneficial to model the effects of the uncertainties present in the system.

An important feature of this thesis is the use of the information-gap uncertainty theory, and this was outlined in the thesis introduction. The earliest and most popular method of representing uncertainty is the “probabilistic” approach, but information-gap theory takes a very different approach to this. It was proposed for problems with “severe” uncertainty, where there is a large gap between the information that is needed to model the system, and the information that is actually available. Information-gap theory proposes that uncertainty may be either favourable or detrimental, and provides measures for assessing both of these two characteristics. In this thesis it is the favourable aspect that is examined.

After the thesis introduction there followed five “theory” chapters, all of which were necessary for the remainder of the thesis. The first of these was entitled “Interval techniques”, and concentrated on the uncertainty techniques of interval arithmetic, affine arithmetic and quadratic arithmetic. Although the methodologies of all three of these

techniques were outlined (especially interval arithmetic), the main purpose of the chapter was to highlight why these techniques were unsuitable for the subsequent thesis work.

Interval arithmetic is highly conservative and often provides excessively vague predictions for an uncertain system. This is also referred to as “dependency” or “the bound explosion effect”, and is especially true if the uncertain system is autoregressive, which many of the simulated examples here were. Affine and quadratic arithmetic both exhibit this same conservative behaviour to a lesser extent, but at the cost of a significant increase in computational expense.

Continuing this line of thought, chapter 3 then proposed a number of methods that were suitable for modelling the uncertain autoregressive systems to be encountered here. The title “Possibilistic sampling methods” was chosen as an umbrella term to incorporate all of the work presented in the chapter. Possibilistic modelling is an alternative branch of uncertainty quantification, as opposed to the older and more popular probabilistic methods. All of the methods within chapter 3 used a methodology similar to design-of-experiments.

Chapter 3 concluded by introducing the transformation method, a possibilistic uncertain approach used frequently within this thesis. The transformation method was developed as a practical means of conducting fuzzy arithmetic, and is used for simulating systems with uncertain parameters. Three different versions of the transformation method were outlined, where these versions differed in computational expense.

Chapter 4 then gave a guide to evolutionary computation methods, beginning with genetic algorithms. This was in preparation for an introductory section on the differential evolution algorithm; an evolutionary strategy that was used several times in this thesis. Differential evolution is a global optimisation algorithm using the principles of natural selection to find the optimum parameters of a system. Solutions are stored in real-number format and vector differentials are used to mutate existing solutions.

As many of the systems in subsequent chapters were to be represented by autoregressive models, chapter 5 gave an introduction to this type of model. This began with an introduction to the ARX model, then described the differences between one-step-ahead and

long-range predictions. The NARX model (an extension of the linear ARX model to nonlinear systems) was then also described. Chapter 5 concluded with an explanation of the pseudoinverse approach (a simple matrix-inversion training approach for systems that are linear-in-the-parameters) and the recommended process for training autoregressive models from scratch.

Chapter 6 then gave an introduction to artificial neural networks, which were to be used in two of the subsequent thesis chapters. Following a brief section on the theory and origin of neural networks, the two most popular neural network types were presented individually. These were the multilayer perceptron (MLP) and radial basis function (RBF) networks. The operation and training approaches for these two networks were covered, in preparation for the later work.

Once the five “theory” chapters had been completed, the emphasis then shifted onto the new and original work within the thesis. Chapters 7-12 each presented a different example of an uncertain system, where the predictive models used in these chapters tended to successively increase in complexity. Each of these chapters presented its methodology and results individually within that chapter, although, where relevant, the chapters occasionally referenced one another also.

Chapters 7-11 were all computational problems, although chapters 9 and 10 used pre-sliding friction data obtained experimentally. In each of these five chapters a crisp-valued model was initially used to represent the system in question. Uncertainty was then deliberately introduced into the parameters of these crisp-valued models, so that the model outputs became interval ranges. This was analogous to the information-gap concept of “opportunity”, where the favourable aspect of uncertainty is assessed.

Chapters 7, 9 and 10 all used the same methodology to train their respective interval-valued models, although the type of model and the data were different. Chapter 7 used one of the simplest kinds of model – an ARX model for a linear system. Simulated data was generated for this linear system, but this original data was much too simplistic to provide a true, challenging problem to study. Zero-mean Gaussian noise was therefore added to the outputs of this data set, so that, in truth, the system was now no longer linear.

An initial crisp ARX model was found for this noisy data set by using the pseudoinverse approach. Different combinations of input and output lags were systematically tested so that the optimum arrangement of ARX regressor terms could be found. As would also follow in chapters 9 and 10, the uncertain model for this system was trained by two different approaches: proportional expansion and the differential evolution algorithm.

Proportional expansion simply means that each element of a crisp solution vector is expanded by the same proportional value ρ into an interval number centred around its original value. This maintained the relative magnitudes of the interval coefficients in respect to each other, and was computationally simplistic. It was argued, however, that this approach did not take into account the individual contribution of each solution vector element, so did not consider the possibility that the crisp solution contained a number of “bad” elements.

The differential evolution algorithm was therefore used as a second interval training approach in order to train the interval model coefficients/weights independently of one another. It was recognised that minimising the width of the interval model prediction bounds was equally important as minimising the model prediction error, so a “width penalty term” was incorporated into the objective function of the differential evolution training algorithm. This term was preceded by an adjustable coefficient β , so that differing emphasis could be placed on minimising the widths of the model bounds.

For these two different interval training methods, the coefficients ρ and β therefore controlled the amount of uncertainty introduced into the network. The rationale behind using these two different training approaches was that the differential evolution method should give tighter prediction bounds as it optimised the interval model coefficients individually, as opposed to the proportional expansion approach which did not.

As shown early in the thesis in chapter 2, interval arithmetic cannot be used for long range predictions from any kind of uncertain autoregressive method, due to its excessive conservatism. One of the possibilistic sampling methods from chapter 3 therefore had to be used, but this led to an issue of computational expense for the differential evolution algorithm. An autoregressive model should be assessed on the quality of its long-range

predictions, since these are representative of the future use of the model once it has been trained. It therefore makes sense for the differential evolution algorithm to train the interval model on its long-range predictions. If this is done using a possibilistic sampling method, however, it is simply not possible due to computational expense, so a key assumption was made in chapter 7. It was assumed that optimising an uncertain autoregressive model on its interval arithmetic one-step-ahead predictions with the differential evolution algorithm would also optimise the long-range predictions using a possibilistic sampling method. Chapters 7, 9 and 10 all used this assumption, so the results were awaited to see how valid the assumption was.

As chapter 7 was the first chapter of original work, it was viewed as a good opportunity to directly compare many of the possibilistic sampling methods from chapter 3. The long-range prediction errors of the uncertain ARX model were estimated with four of these methods for both the proportional expansion and differential evolution approaches, over a range of ρ and β values respectively.

Figure 7.2 gives the proportional expansion results for the four different possibilistic sampling methods. It can be seen that these four curves are very similar, and all show a continuous monotonic relationship between the long-range prediction error of the interval models and the level of uncertainty introduced. This raises the issue of a “trade-off” between prediction error and solution precision; a theme that resurfaces in subsequent thesis chapters.

Figures 7.4 and 7.6 gave the differential evolution results for the four different possibilistic sampling methods. Unlike the proportional expansion curves in Figure 7.2, some of these curves are significantly different from one another, and not all of them show a monotonic relationship between long-range prediction error and level of uncertainty introduced. This shows that different possibilistic sampling methods can give different results, so the choice of which method to use is important. The reduced and modified reduced transformation method curves do show a continuous monotonic relationship between long-range prediction error and level of uncertainty introduced, so these are the two methods that are recommended for use here.

The central conclusion from chapter 7 was therefore that if some degree of tolerance can be accepted on the predictions from a model, it is possible to reduce the prediction error. The continuous trade-offs in both Figure 7.2 and the transformation method variants of Figure 7.6 are useful here, because this is very much an application-dependent and user-specific problem. Some problems may require a much higher solution precision than others, and a higher error may be tolerable on some problems than others. The use of the coefficients ρ and β allows a user of this method to choose a suitable trade-off level relative to their particular application.

Chapter 8 then examined replacing a nonlinear system with an uncertain linear model. Equivalent linearisation of a nonlinear system is a very well-known control technique, to the point where it is taught within an undergraduate degree syllabus. The approach here differs from that, however, in that an interval-valued linear model is used, as opposed to a conventional crisp linear model. The Duffing oscillator was used here as it is one of the simplest nonlinear systems available.

Nonlinear predictive models are able to give a good prediction over a wide range of system operating points, as was shown in Figure 8.3. A linearisation of the system gives an inferior prediction quality, however, and this prediction quality deteriorates rapidly as the operating point deviates from the linearisation point, illustrated in Figure 8.1. The aim of chapter 8 was therefore to create a model that had the operating range of a nonlinear model, but the computational simplicity of a linear model. This was achievable with an uncertain linear model, given here as an interval-valued ARX model.

As in chapter 7, an initial crisp-valued model was generated for the system. Only the proportional expansion approach was used to generate the interval models here though, so the differential evolution algorithm was not required in chapter 8. The operating point of the interval ARX model was represented separately by two different measures: system input RMS and additional system output noise. 1% long-range MSE was used as the desired prediction error of the model.

Figures 8.4 and 8.5 give the interval model performances in relation to the two specified measures. Figure 8.4 gives the range of system input RMS than can be tolerated in order to

give an interval long-range MSE of up to 1%, with respect to the expansion proportion ρ of the regressor coefficients. The long-range MSE of the initial crisp ARX model at the linearisation point was 1.9%, so for a small value of ρ it is not possible to obtain 1% interval MSE, regardless of the input RMS. The remainder of Figure 8.4 then shows that as more uncertainty is introduced into the ARX model coefficients, a greater range of input RMS can be tolerated. As with Figure 8.4, Figure 8.5 shows that for small values of ρ it isn't possible to obtain 1% interval MSE, so the results curve is not plotted in that region. The remainder of Figure 8.5 then shows that as the ARX regressor coefficients are expanded further, a greater level of additional output noise can be tolerated to still obtain an interval MSE of up to 1%.

Figures 8.4 and 8.5 also exhibit the trade-off behaviour experienced in chapter 7. As more uncertainty is introduced into the ARX model, this model can then tolerate a greater range of operating conditions, but this will inevitably come with the price of wider model prediction bounds. The measures assessed in Figures 8.4 and 8.5 are both continuous with the expansion proportion ρ , which enables a decision to be more easily made on the optimum trade-off level. As with chapter 7, this trade-off is application dependent and user-specific. Although the Duffing oscillator studied in this chapter was only a simple nonlinear system, this technique could also potentially be used for more complex nonlinear systems.

Chapters 9 and 10 then increased the complexity of the computational model used; moving from the ARX models of chapters 7 and 8 to artificial neural networks. The two most popular neural network types were used: the multilayer perceptron in chapter 9 and the radial basis function network in chapter 10. The data set used here was pre-sliding friction data obtained experimentally in a background paper from a specially constructed test rig.

As with chapters 7 and 8, an initial crisp-valued model was found for chapter 9, although the architecture of this MLP network was taken directly from the background paper. This two-layer network had a sufficiently small number of weights for uncertainty to be introduced into all of them. As in chapter 7, the uncertain weights were trained by both the proportional expansion approach and the differential evolution algorithm, where these two methods were respectively governed by the adjustable parameters ρ and β .

The long-range interval MSE values of the network were found over a wide range of ρ and β values, where these results are shown in Figures 9.3 and 9.6. As with Figure 7.2 and the transformation method variants in Figure 7.6, Figures 9.3 and 9.6 show that there is a continuous monotonic relationship between the level of uncertainty introduced into the network (given by ρ or β) and the interval prediction error. This once again allows a trade-off between these two quantities, where the level for this trade-off is application-dependent and user-specific. For the proportional expansion approach, Figures 9.4 and 9.5 show that if an inferior initial crisp solution were used for the MLP network, the bounds would need to be wider in order to obtain the same value of interval MSE. This is a result that could be deduced intuitively.

Unlike chapter 9, the optimum architecture for the radial basis function network in chapter 10 was not originally available, so this had to be deduced as part of the training phase. An initial crisp-valued RBF network was found for the data, although this had a significantly greater number of adjustable parameters than the MLP network in chapter 9, with 56 in total. This was too many for uncertainty to be introduced into all of them, so the decision was therefore made to only set the network output layer weights as interval numbers. As with chapters 7 and 9, these uncertain parameters were trained by both the proportional expansion approach and the differential evolution algorithm.

The uncertain network output weights were trained for a variety of different ρ or β values respectively, and these results are shown in Figures 10.1 and 10.4. There is once again a monotonic continuous relationship between the level of uncertainty introduced into the network and the interval MSE. Introducing uncertainty into the network output weights causes the predicted output bounds to become wider, leading to a trade-off between minimising the network prediction error and minimising the widths of the predicted bounds. As in chapters 7 and 9, the continuous relationships in Figures 10.1 and 10.4 allow this trade-off to be made more easily by the user. The optimum level for this trade-off is application-dependent and user-specific. Figures 10.2 and 10.3 show that to obtain a lower value of interval MSE from the same initial crisp solution with the proportional expansion approach, the proportion of expansion ρ needs to be greater. This is in agreement with the proportional expansion curve in Figure 10.1, and is also a result that could be deduced intuitively.

As the same data set was used in chapters 9 and 10, this gave an interesting opportunity to directly compare the MLP and RBF neural network types for the same problem. In chapter 9 the optimum crisp MLP network had 17 adjustable weight parameters and gave a long-range MSE of 4.1% on the test set. The crisp RBF network in chapter 10 had 56 adjustable weight parameters and gave a long-range MSE of 10.2% on the test set. In the case of this pre-sliding friction data therefore, the MLP network is able to provide a much better prediction with a simpler model structure. Even considering that the training phase for this MLP network was much longer than for the RBF network in chapter 10, the MLP network is still the superior of the two types here.

A direct comparison can also now be made for the proportional expansion approach and the differential evolution approach, based on the results obtained in chapters 7, 9 and 10. It can be seen that the proportional expansion curves in Figures 7.2, 9.3 and 10.1 all have very similar shape, showing an exponential-like decay of the interval model MSE as the expansion proportion ρ is increased. The proportional expansion approach therefore works in much the same way here, regardless of the type of computational model.

In chapter 7, Figures 7.3 and 7.7 directly compare the proportional expansion and differential evolution approaches for the same value of long-range interval MSE. It can be seen that the bounds from the proportional expansion approach are significantly narrower, even though this is a much simpler method for obtaining an interval-valued model. In chapter 9, Figures 9.4 and 9.7 make the same comparison between the two different interval training methods. Here the differential evolution bounds in Figure 9.7 are very slightly narrower, but this is barely visible on these two figures. In chapter 10, Figures 10.2 and 10.5 also compare the two different approaches. In this case the average sample bound width over the whole 8000 samples of the test set is 0.2% lower for the proportional expansion approach than the differential evolution training approach, so the proportional expansion approach produces marginally narrower prediction bounds.

It was hypothesised early on in this thesis that the differential evolution training approach would produce narrower prediction bounds than the proportional expansion approach, due to it training the uncertain quantities independently of one another. Based on the above comparisons between these two methods in chapters 7, 9 and 10, this hypothesis appears

not to be true. Even though the differential evolution training approach first introduced in chapter 7 is much more computationally expensive than the proportional expansion approach, it gives inferior results. This raises the question as to how its use can be justified.

As stated earlier, it was assumed that optimising an uncertain autoregressive model on its interval arithmetic one-step-ahead predictions with the differential evolution algorithm would also optimise the long-range predictions using a possibilistic sampling method. It is likely that this assumption is the problem here, not the idea of training uncertain parameters independently of one another.

One-step-ahead and long-range predictions for an autoregressive model are not the same, and although it could be estimated that minimising the former would also minimise the latter, this is certainly not guaranteed to be true. Interval arithmetic also differs from the possibilistic sampling methods introduced in chapter 3, so assessing outputs on these two methods is also likely to generate different values. This could explain why the differential evolution algorithm does not predict tighter long-range bounds than the simpler proportional expansion approach. It is possible that if the computational power was available to train the uncertain model parameters on their long-range outputs with differential evolution, this minimisation of output bound widths could be achieved. With the computational power available for this thesis, however, this was impractical to attempt.

Chapter 11 then changed direction slightly, with a look at a simple nonlinear system capable of bifurcation. This system was a harmonically-driven symmetric Duffing oscillator modelled in discrete-time, where the coefficients of this system were the same as for the Duffing oscillator used in chapter 8. The steady-state response amplitudes were calculated over a range of different forcing frequencies and it was found that within two particular forcing frequency limits, an additional steady-state limit cycle appeared. The steady-state limit cycle that the system settled upon was dependent on the initial conditions provided to the system.

Unlike the previous four chapters, the aim of chapter 11 was not to model the system, but merely to show the behaviour of which the system was capable. Bifurcation of a nonlinear system can lead to significant changes in the system behaviour, as was the case here.

Figures 11.2 – 11.4 and 11.6 – 11.7 show how the basins of attraction change for the two limit cycles as the forcing frequency is changed. Figures 11.5, 11.10 and 11.13 then show that very small changes in initial conditions can give significant changes in output behaviour; a characteristic that is well known for chaotic nonlinear systems.

Three different forcing frequencies were chosen within the bifurcation range in chapter 11. In each of these three different cases, a different set of initial conditions was chosen that lay very close to a basins of attraction boundary on the phase portrait, so that a very small change in initial conditions would give different steady-state system behaviour. As with chapters 7-10, an initial crisp model was obtained for the system. This took the form of a NARX model, where the regressor coefficient values were given in terms of the specified system parameters.

The proportional expansion approach was used in chapter 11 to transform the initial crisp NARX model into an interval-valued model for the three different forcing frequencies. As this was once again an uncertain autoregressive model, a possibilistic sampling method had to be used to approximate the long-range prediction bounds. Figures 11.8 and 11.9 give uncertain output bounds for two different expansion proportions ρ , in the case of the first forcing frequency $\omega=160$ rad/s. Figure 11.8 shows that a small ρ value was unable to detect both steady-state limit cycles of the system, but the larger ρ value in Figure 11.9 was able to do this.

The uncertain output bounds in Figure 11.9 therefore indicate both steady-state behaviours that the bifurcating discrete-time system is capable of, but without excessively wide output bounds. Figures 11.11 and 11.12 then show similar behaviour for the second forcing frequency of $\omega=170$ rad/s, where the uncertain prediction in Figure 11.12 shows both steady-state behaviours without overly-conservative output bounds.

Figures 11.14 and 11.15 then give uncertain output predictions for the third forcing frequency of $\omega=180$ rad/s. As with the previous two forcing frequency increments, the expansion proportion in Figure 11.14 was insufficiently large to represent both steady-state limit cycles. The larger ρ value in Figure 11.15 gave prediction bounds that contained both steady-state limit cycles for that ω value, but that were not unnecessarily wide to achieve

this task. Comparing the uncertain model predictions in Figures 11.9, 11.12 and 11.15 with the displacement histories in Figures 11.5, 11.10 and 11.13, it can also be seen that the largest ranges of output from the interval models occur in the transient region, which is something that could have been deduced intuitively.

Chapter 11 therefore showed that an interval-valued model could be used to illustrate the possible behaviour of a bifurcating nonlinear system. Bifurcation is often a problem for nonlinear systems, and its effects can be severe. The deliberate introduction of uncertainty into the NARX model here allows the user to better understand how the system is likely to behave when bifurcation is a possibility, so this chapter is another example of the favourable aspects of uncertainty.

Chapter 12 was the largest chapter in this thesis, and presented the findings from nine months of research. This chapter is a contrast to the remainder of the thesis, as it gives many results obtained experimentally from an industrial environment, as opposed to the simpler simulated problems in chapters 7-11. Another feature of chapter 12 is that it presents work conducted in the frequency domain, as opposed to the time domain work in the previous five chapters.

Unlike chapters 7-11, where uncertainty was deliberately introduced into adjustable coefficients of various different computational model types, chapter 12 instead looked at the effects of existing uncertainties. The system in question was an automotive spot-welded component with symmetry in all three orthogonal directions. The effects on the component dynamic response were examined for three different selected uncertainties: component sheet thickness, differences in spot weld positions and varying numbers of missing spot welds. These were considered as uncertainties that the component was likely to experience in a real-life situation. It was chosen to use spot welds in this chapter because they are very widely used in the automotive industry, and they exhibit nonlinear behaviour.

In chapter 12 the system was represented by both computational modelling (in the form of finite element analysis) and experimental modal testing. This enabled a direct comparison between these two approaches. A design-of-experiments approach was used for the uncertainty modelling, to continue the possibilistic uncertainty theme running through this

thesis. The three selected uncertainties were investigated quite intensively, with 18 manufactured specimens and 171 different finite element models used in total.

The ten accelerometer locations for the component were selected using an optimisation routine, and the additional weight from these was included in the finite element models. The component was always excited at the same point and in the same direction, and only two of the three orthogonal response directions were studied, as the third was shown to be much less significant. The uncertainty in the component response was assessed by two different measures: the range of FRF amplitudes and the variation of the nominal modal frequencies. The appearances of the nominal modes in other specimens were judged using the frequency domain assurance criterion (FDAC), a frequency-domain equivalent of the well-known modal assurance criterion (MAC).

The ten welds in the nominal component design were assigned numerical indices corresponding to their locations. Four different situations were examined for the number of missing spot welds parameter: all ten welds present, any single weld missing, welds 2 and 8 missing, any two welds missing. The manufactured experimental specimens only had either all ten welds present or welds 2 and 8 missing, so a finite element-experimental comparison was not possible for all four situations.

After calculating and inspecting all of the results from chapter 12, the following set of conclusions were drawn:

- Of the three uncertainties studied, the sheet thickness is the most important in relation to the effects on both the FRF amplitudes and the nominal modal frequencies.
- The spot weld positions uncertainty is less important, and the number of missing welds parameter is the least important of the three.
- Combinations of parameter uncertainties have a greater effect on the model response than using them individually.
- In relation to the FDAC, the finite element FRF correlation with the nominal model is worse for any one missing spot weld than for all welds present

- For the finite element models, the range of the nominal modal frequencies is greater for any two missing welds than for any one missing weld
- The effect of removing welds 2 and 8 on the FDAC FRF correlation with the nominal model is worse for the finite element case than the experimental case.

Unfortunately it had not been possible for any finite element model updating to be done in this work, due to time constraints. This widely-used modelling procedure could have improved the comparison between the finite element and experimental results, as the comparisons in Figure 12.20 and 12.21 were not ideal, although reasonably good.

Chapter 12 showed that the dynamic effects of three different uncertainties could be adequately modelled with a design-of-experiments uncertainty approach, and that predicted results obtained from computational and experimental procedures were largely in agreement. Although the spot-welded component was a somewhat idealised structure due to its orthogonal symmetry, chapter 12 still showed that uncertainty modelling can be used just as effectively in an industrial environment as for simulated computational problems.

Uncertainty propagation and quantification have therefore been extensively studied within this entire thesis. A number of different computational model types have been studied, along with finite element and experimental modal testing in an industrial environment. The thesis has concentrated on possibilistic uncertainty modelling techniques, as the work in chapters 7-12 has been more suited to this approach. All the uncertainties present here, whether the deliberate uncertainty introductions in chapters 7-11 or the existing uncertainties modelled in chapter 12, took interval forms. No information was available as to the distributions of these uncertainty quantities, so although the more popular probabilistic uncertainty approaches could have been used, possibilistic approaches were a better choice.

The various possibilistic sampling method applications in chapters 7-11 were used here as a way of circumventing the conservatism problems of interval arithmetic, as interval arithmetic is unsuitable for the long-range predictions of any uncertain autoregressive model. In chapters 7-11 all of the different possibilistic sampling methods described in

chapter 3 were used, with the choice of method being dependent on the required computational expense.

Uncertainty was deliberately introduced into all of the computational models in chapters 7-11, and it was shown that the effects of this uncertainty could be favourable. Chapters 7-10 showed that if some degree of tolerance could be accepted by the user on the predicted outputs, it was possible to obtain a lower prediction error when compared to the original crisp models. Chapter 11 showed that introducing uncertainty into a discrete-time representation of a bifurcating nonlinear system could illustrate all potential effects of this (albeit simple) bifurcation.

Another common theme in chapters 7-11 was the concept of a trade-off between the level of uncertainty introduced into the model and the level of precision obtained on the result. As already stated earlier, the idea of deliberately introducing uncertainty into a model to exploit the favourable aspects of that uncertainty originates from information-gap theory, and within this theory there also exists the idea of a “trade-off” for both the opportunity and robustness immunity functions.

As already hinted upon, available computational power and time acted as limiting factors on much of the work done in this thesis. The idea of using differential evolution to independently train the uncertain parameters of interval models was good in theory, but the necessary compromise required for practical computational time meant that the proposed precision benefits did not materialise. Totally independent uncertain parameter training may be realisable in the cases of parallel computing power or the use of a “supercomputer”, but was not feasible for the stand-alone von Neumann-style computers used in this work. Increased time and computational power could also have permitted a greater investigation of spot weld uncertainties for the finite element models in chapter 12, or for a greater number of experimental specimens to have been manufactured and tested.

In summary, this thesis has studied the effects of uncertainty of a number of different systems, and has given several examples of how the information-gap theory can be applied to actual problems. It has shown that uncertainty can often be favourable, and has highlighted the application-dependent trade-off that frequently occurs between the level of

uncertainty present in a model, and the level of solution precision. The thesis has shown the importance of modelling uncertainty within engineering applications, and has presented an alternative to the traditional “factor of safety” design approach.

13.2 Future Work

Chapters 7-11 deliberately introduced uncertainty into the adjustable parameters of various different crisp-valued computational models, and showed the beneficial effects of these uncertainties. This is analogous to the “opportunity” concept in Ben-Haim’s information-gap theory. Opportunity is only one of the two information-gap immunity functions, however. The second of these immunity functions, robustness, has not been considered at all within this thesis.

This is therefore a suggestion for future work. Time could be spent examining both the opportunity and robustness characteristics of introducing uncertainty into computational models. Robustness, however, is more suited to quantifying the effects of existing uncertainties, for example incomplete or noisy training data. This poor data could be represented by interval data, to indicate the severity of the data flaws. After selecting a nominal form for the computational model, uncertainty could then be introduced into the coefficients of this model, and the model could be assessed on both the opportunity and robustness performances simultaneously. This may lead to an optimum uncertainty level; one which best balances out the favourable and detrimental aspects of uncertainty for the model.

Another suggestion is to increase the complexity of the computational models used. From chapters 7-11, the highest degree of nonlinear model complexity reached was the neural networks in chapters 9 and 10. There are more complex nonlinear models available than these, however, either for the number of adaptive parameters used, or the processing operations. The techniques presented in chapters 7-11 could also be used for one of these models. In theory, any nonlinear model with adjustable coefficients/weights could be transformed into an uncertain model by the introduction of uncertainty into these adjustable quantities.

A further suggestion is to study more complex nonlinear systems. Chapter 7 modelled a linear system with added output noise, and chapters 8 and 11 examined the Duffing oscillator, one of the simplest nonlinear systems available. Chapters 9 and 10 represented a hysteretic pre-sliding friction data set, which was the most complex of the systems in chapters 7-11. There are countless nonlinear systems that are more complex than these though, giving a much more challenging problem to attempt.

This call for more complicated nonlinear systems especially applies to chapters 8 and 11. The replacement of a nonlinear model with an uncertain linear model in chapter 8 has great potential for further use, as an uncertain linear model with operational range approaching that of a nonlinear model is a tempting prospect. In order to be credible though, the technique must be used on more difficult nonlinear systems than the Duffing oscillator. Likewise in chapter 11, using an uncertain model to demonstrate the potential behaviour range of a bifurcating nonlinear system is also an attractive proposition. As with chapter 8 though, it needs to be implemented on more substantial bifurcations in the future.

All of the uncertainties introduced in chapters 7-11 were modelled using possibilistic uncertainty quantification techniques. This is only one of the two main branches of uncertainty quantification though, the other being the older and more widely used probabilistic branch. Probabilistic uncertainty quantification was not used at all within this thesis, as possibilistic modelling was considered to be the more suitable of the two approaches. An interesting idea for future work could be to see if the favourable uncertainty behaviour observed here would also occur for uncertainties introduced with probabilistic distributions. The uncertain model outputs would then also have probabilistic distributions, and an alternative interval MSE function would be needed to that of equations (7.5) and (7.6).

For the MADUSE work in chapter 12, there are a number of suggestions that can be made in relation to future work. The use of finite element model updating would have made a difference here, as the correlation between the nominal finite element and experimental responses is not exact. This could be a reason why the removal of 2 welds has a greater effect on the finite element response FRFs than for the experimental responses.

One suggestion for future work in chapter 12 is to eliminate the simplifications used in the spot weld position and sheet thickness uncertainties. If the spot weld were displaced independently of one another, or local regions of sheet thickness variation were used, this could give a study that was more akin to a real-life industrial application.

A second suggestion is to use a more intricate representation of a finite element spot welded joint. Finite element joint models have been proposed in the past that take into account contact friction, residual pressure distributions and other factors. If experimental tests were again used, a model updating phase could specifically target the simulation parameters relating to these weld models.

A third suggestion is to examine uncertainties relating to material properties, i.e. mass, stiffness and damping, as opposed to the uncertain parameters that were chosen here. A fourth choice could be to use attempt an advanced uncertain modelling technique such as the transformation method, although this would require a much greater number of tests to be undertaken.

Another suggestion is to select a small number of modes that are most important in regard to the operation of the studied component, as opposed to the wide frequency range examined here. The effects of uncertainties on these specific modes could be studied, possibly including a figure for the maximum allowable uncertainty on a given parameter, considering its effect on the component as a whole.

A final suggestion is to study a more detailed model than the spot welded component used in chapter 12. This could be a built-up system consisting of several constituent parts, each with their own respective uncertainties, but still capable of being recreated in both a laboratory environment and by an engineering simulation software package.

At the time of writing this thesis there is in excess of two years remaining on the MADUSE project, so it is likely that some of these suggestions for future work arising from chapter 12 are presently being implemented by the current MADUSE researcher at Centro Ricerche Fiat.

References

- Abbas, H.A. 2002. "A Memetic Pareto Evolutionary Approach to Artificial Neural Networks". *Lecture Notes in Artificial Intelligence*. Vol 2256. Springer-Verlag
- Allemang, R.J. & Brown, D.L. 1982. "A Correlation Coefficient for Modal Vector Analysis", *Proceedings of the 1st International Modal Analysis Conference*, pp 110-116
- Bayes, T. 1964. "An essay towards solving a problem in the doctrine of chances", *Philosophical transactions of the Royal Society of London*, 53, pp 370-418
- Ben-Haim, Y. 2001. *Information Gap Decision Theory*. Academic Press.
- Billings, S.A., Chen, S. & Korenberg, M.J. 1989. "Identification of MIMO non-linear systems using a forward-regression orthogonal estimator", *International Journal of Control*, 49 (6), pp.2157-2189.
- Broomhead, D.S. & Lowe, D. 1988. "Multivariable Functional Interpolation and Adaptive Networks". *Complex Systems*, 2, pp 322-355
- Comba, J.L.D & Stolfi, J. 1993. "Affine Arithmetic and its Applications to Computer Graphics", *Proceedings of the 6th Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAP'93)*
- Cristianini, N. & Shawe-Taylor, J. 2000. *An Introduction to Support Vector Machines (and Other Kernel-Based Learning Methods)*. Cambridge University Press
- Dempster, A.P., Laird, N.M. & Rubin, D.B. 1977. "Maximum likelihood from incomplete data via the EM algorithm", *Journal of the Royal Statistical Society*, B 39 (1), pp 1-38

- Dong, W & Shah, H.C., 1987. "Vertex method for computing functions of fuzzy variables", *Fuzzy Sets and Systems*, 21, pp 65-78
- Drago, G.P & Ridella, S. 1998. "Pruning with Interval Arithmetic Perceptron", *Neurocomputing*, 18, pp 229-246
- Dubois, D. & Prade, H. 1978. "Operations on fuzzy numbers", *International Journal of Systems Science*, 9, pp 613-626
- Dubois, D. & Prade, H. 1979. "Fuzzy real algebra: Some results", *Fuzzy Sets and Systems*, 2, 327-348
- Fisher, Sir R.A. 1935. "The Design of Experiments"
- Goldberg, D.E. 1989. "Genetic Algorithms in Search, Optimisation and Machine Learning", *Kluwer Academic Publishers*, Boston, MA
- Hanss, M. 2002. "The Transformation Method for the Simulation and Analysis of Systems with Uncertain Parameters, *Fuzzy Sets and Systems*, 130, pp 277-289
- Holland, J.H. 1975. "*Adaptation in Natural and Artificial Systems*", The University of Michigan Press
- Hornik, K. M., Stinchcombe, M. & White H. 1989. "Multilayer feedforward networks are universal approximators", *Neural Networks*, 2 (5), pp 359 – 366
- Korenberg, M., Billings, S.A., Liu, Y.P. & McIlroy, P.J. 1988. "Orthogonal parameter estimation algorithm for non-linear stochastic systems", *International Journal of Control*, 48 (1), pp 193-210

Leontaritis, I.J. & Billings, S.A. 1985. "Input-output parametric models for non-linear systems. Part I: Deterministic non-linear systems", *International Journal of Control*, 41 (2), pp 303-328

Leontaritis, I.J. & Billings, S.A. 1985. "Input-output parametric models for non-linear systems. Part II: Stochastic non-linear systems", *International Journal of Control*, 41 (2), pp 329-344

Masters, T. & Land, W. 1997. "A New Training Approach for the General Regression Neural Network", *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics, Computational Cybernetics and Simulation*, Vol 3

McCulloch, W.S. & Pitts, W. 1943. "A logical calculus of the ideas immanent in nervous activity", *Bulletin of Mathematical Biophysics*, 5, pp 115-133

Metropolis, N. & Ulam, S. 1949. "The Monte Carlo method", *Journal of the American Statistical Association*, 44, pp 335-341

Metropolis, N., Rosenbluth, A., Rosenbluth, M., Teller, A. & Teller, E. 1953. "Equation of State Calculations by Fast Computing Machines", *Journal of Chemical Physics*, Vol 21, 6, pp 671-680

Montana, D.J. & Davis, L. (1989). *Training Feed-Forward Neural Networks Using Genetic Algorithms*. BBN Systems and Technologies, Cambridge, MA.

Moore, R.E. 1966. "*Interval Analysis*", Prentice-Hall

Nabney, I.I. 2001. "*Netlab: Algorithms for pattern recognition*", Springer-Verlag

Oberkampf, W., Helton, J. & Sentz, K. 2001. "Mathematical representation of uncertainty", *Non-Deterministic Approaches Forum, AIAA*

Parlitz, U., Hornstein, A., Engster, D., Al-Bender, F., Lampaert, V., Tjanjowidodo, T., Wong, C.X., Worden, K. & Manson, G. 2004. "Identification of Pre-sliding Friction Dynamics", *CHAOS* (Focus issue on *Nonlinear dynamics in spatially extended mechanical systems*)

Pascual, R., Golinval, J.C., Razeto, M. 1997. "A Frequency Domain Correlation Technique for Model Correlation and Updating", *Proceedings of the 15th International Modal Analysis Conference*.

Piroddi, L. & Spinelli, W. 2003. "Long Range Prediction: A Case Study", *Proceedings of the 42nd IEEE Conference on Decision and Control*, pp.3984–3989.

Plagianakos, V.P., Magoulas, G.D & Vrahatis, M.N. 2001. "Learning in Multi-Layer Perceptrons Using Global Optimisation Strategies". *Nonlinear Analysis, Theory, Methods and Applications*, Vol 47, 5.

Rosenblatt, F. 1962. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Washington DC: Spartan.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. 1986. Learning Internal Representations by Error Propagation. In D.E. Rumelhart and J.L. McClelland, (Eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, pp 318-362, Chapter 8. Cambridge, MA and London: MIT Press.

Storn, R. & Price, R. 1997. Differential Evolution – a Simple and Efficient Heuristic for Global Optimisation Over Continuous Spaces. *Journal of Global Optimisation*, 11, pp 341 – 359

Tarassenko, L. 1998. *A Guide to Neural Computing Applications*. Arnold.

Worden, K., Manson, G., Lord, T.M., Friswell, M.I. 2005. "Some Observations on Uncertainty Propagation Through a Simple Nonlinear System", *Journal of Sound and Vibration*, awaiting publication.

Zadeh, L.A., 1965. "Fuzzy sets", *Information and Control*, 8, pp 338-353

Appendix A

Given below are details of the academic conference publications given by the author for the work presented in chapters 8-11 of this thesis:

Chetwynd, D., Worden, K. & Manson, G. 2004. "An Application of Interval-Valued Neural Networks to a Regression Problem", *Proceedings of the 29th International Symposium on Noise and Vibration Engineering (ISMA2004)*

Chetwynd, D., Worden, K. & Manson, G. 2005. "On Replacing a Nonlinear System with an Uncertain Linear Model. Part I – The Time Domain", *Proceedings of the 23rd International Modal Analysis Conference (IMAC2005)*

Chetwynd, D., Worden, K., Manson, G. & Pierce, S.G. 2005. "Uncertainty Propagation through Radial Basis Function Networks. Part I – Regression Networks", *Proceedings of the 6th European Conference on Structural Dynamics (EURODYN2005)*

Chetwynd, D., Worden, K., Manson, G. & Pierce, S.G. 2005. "Interval-Valued Regression Models of Bifurcating Nonlinear Systems", *Accepted for presentation at the Ninth International Conference on Recent Advances in Structural Dynamics (RASD2006)*

Appendix B

Presented here are additional results from chapter 12. Given are the complete set of modal frequency results for any two missing spot welds, corresponding to 15 different combinations of two missing weld positions.

Nominal	Positions 1 and 2 Missing								
	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	200.5	213.5	166.5	221.5	234.25	206.25	242	254.5	226
227.75	213.5	221.5	187.25	233.25	243.25	225.5	253.75	265.5	245.25
256	155	155.5	143.25	166.75	170.75	166.5	178.5	180.5	178
325.75	357.25	357.25	243.5	393.75	393.75	394	430.5	430.75	430.5
422.25	391.5	391.75	358.25	421.5	421.25	421.25	449.25	449	449.5
500.75	446	446.5	401.75	486.75	487.5	485.25	524.5	525.25	522.75
534.75	478.25	559.5	496.5	526.5	620.25	533.5	573	579.5	671
552.75	496.75	544.75	441	550	596.5	548.75	602.5	652.75	602
589	529.25	487.25	418.25	535.75	537.5	522	583.5	585.5	573.5
654.75	587.5	587.5	522.5	652.25	652	652.25	714.75	714.25	715
672	605.5	605.25	539.25	669.5	669.5	669.75	737.5	736.75	738.25
756	683.25	683.5	693	914.5	757.75	863.5	830	827.25	830.5
796.5	641.5	641.75	847	689.25	689	688	726	726	724.5
846.75	799	799.75	667	854.5	820.25	843.75	898.75	866.75	885.75
915	819.75	816.25	730.25	953	907	914.75	947	996.75	946.25
988.25	892.25	771.25	943	988.25	989	986	848	922.5	811

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
	- No match found

Positions 1 and 3 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	179.75	192.75	166.5	198.75	213.25	183.75	217	233	201.25
227.75	200.75	-	185.5	221.5	236.75	205.25	242.25	260	225
256	211	212.25	210.25	230.25	231.75	229.25	248.5	250.25	247.25
325.75	288	287.75	288	318.25	318	318.5	349	348.75	349.25
422.25	497	351	394.75	213.25	243.5	424.5	237.5	269	452
500.75	462.5	499.25	458.25	504.75	552.25	500.75	546	550	542
534.75	486.5	487.5	543	535.75	537	599.5	583.75	585.25	656.75
552.75	525	554.5	554	582.25	614.75	612.5	638.75	674.5	670
589	501	518.5	476.25	554	575.25	546.5	607	631.75	599.5
654.75	586	585.25	585.5	649.5	648.5	648.75	711.75	711.25	710.25
672	603.75	603.75	604	669.75	669.5	670	945.25	735	935
756	808	808.5	776.25	863.25	982.75	863	-	948	945
796.5	635.75	643.75	678.5	848.5	861.25	683.5	895.75	751.75	733.75
846.75	796.5	797.25	794.25	755	816.25	838	824.5	867	881.25
915	821.75	816.75	821.25	911.25	907.25	912.5	-	998	984.25
988.25	724.75	757.5	890.25	787.75	758	986	829.75	923.25	832.75

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
-	No match found

Positions 1 and 4 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	183.25	196.5	169.75	202.5	217.25	187.75	221.25	237.5	219
227.75	194.75	208.75	180.5	215	227.25	200	235.75	458.75	205.25
256	213.5	214.25	213	233.5	234.25	232.75	252.25	253.5	251.75
325.75	286.75	286.75	286.5	316.5	316.25	317	346.5	346.25	347
422.25	398.5	367.75	394.75	427.5	390.25	424.75	454.25	411	452.25
500.75	498.75	523.5	480.25	552	580.5	500.25	604.5	637	539.5
534.75	553	486.5	556.25	612.5	535.75	616	671.75	583.75	675.25
552.75	516.25	545.5	500.5	572.25	604.5	553.5	627.75	663.25	605.75
589	504	541.75	487.25	558.25	552	539.25	612.75	871.25	591.25
654.75	583.5	582	583.75	602.5	646.5	647	706.25	708.5	706.75
672	605.5	605.25	606	671.75	671.25	672.25	735.5	735.5	735.5
756	595.25	596.5	682.75	757.25	757.75	757.5	830.5	832	834.75
796.5	856.5	682.75	850.75	710	716.75	763.25	757.75	764.5	750.25
846.75	795.75	797	944	845	824.5	841	890	876.5	884.5
915	953.75	946	821.5	862.5	935	946.5	948	946.25	994.25
988.25	891.25	892	888.75	988	988.75	984.5	995.75	993.5	945.25

Positions 1 and 5 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	187.75	-	174	207.5	222.75	192.5	227	243.75	210.5
227.75	177.5	190.75	164	195.75	210.75	180	213.75	230.25	196.25
256	200.75	201.5	200.75	217.5	218.5	217.5	233.5	234.5	233.25
325.75	277.75	278	277.75	304	304.5	452	479.25	483	476.75
422.25	166.25	426.75	160.5	186.75	203.5	-	207	226	449.25
500.75	403.75	409	400.5	855.75	855.75	858	891.75	463.5	582.25
534.75	508.25	545.25	483.25	563	604.25	555.75	617.75	585.25	608.25
552.75	711.75	487.25	502.75	535.75	536.75	534.25	583.75	662.5	587
589	496.25	515.5	450.5	550.5	571.5	485.5	604.5	627.75	519.25
654.75	450	451.25	323.25	483.75	484	343.75	516.75	516.25	357.25
672	776.75	603.75	467.25	-	669.5	670.5	965	735.75	946.25
756	681.5	681.25	681.5	371.75	755.75	756.25	390.5	421	831.25
796.5	623.75	628.25	620.25	669.75	686.25	678.75	738.75	742.25	735.25
846.75	794.25	794.25	793.25	840	841.5	837.25	883.5	886	878.5
915	815.75	876.25	876.5	907.75	907.25	778	985.75	996.5	1013
988.25	891.25	891.75	891.75	987.75	988.25	985.5	946	946	993.25

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
-	No match found

Positions 1 and 6 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	-	216.75	187	224	240	207.25	245.5	263	227.25
227.75	168	180.25	155.5	185.5	199.25	171.5	202.5	217.75	187.25
256	200	201	199.5	216	217.25	215.5	230.75	232.5	230.25
325.75	285.25	285.25	285.5	316	316	316	346.5	346.5	346.75
422.25	355.25	366.75	406	378.25	416	435	400.5	418.5	482
500.75	460.25	463.25	457.5	503	506	499.5	543.75	547	540
534.75	530	500.75	485.75	602.25	602.5	534.75	670.75	661.5	676
552.75	500	560.75	496	587.75	621.5	547.5	644.75	681.75	603.25
589	527.5	488	474.75	533.25	542.5	614.25	660.75	437	661.3
654.75	589.25	589.25	589.25	654.5	654.5	654.5	719.5	719.5	719.5
672	603.5	603.25	617.5	669.25	669	686.25	734.75	734.5	754.75
756	683	683.25	683	685.25	757.5	757.25	753.25	751.75	734.75
796.5	460.25	528.25	522.25	553.5	563	851.75	574.25	877.25	557.25
846.75	798.75	795	797.75	841.25	842.25	839.25	888.25	889.75	882.25
915	775.5	775.5	775.5	914.75	910	915.5	947	1001	948.25
988.25	891.25	891.75	889.25	988	988.25	985.25	977	978.25	974.75

Positions 1 and 7 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	202	218.5	164	216	244.5	207.25	244.75	270.75	227.25
227.75	190	-	187.25	223.25	239	-	240.75	263.5	198
256	213.25	213.75	212.75	232.75	233.5	232.25	251.5	254.75	251
325.75	287.25	287.25	287.25	317.25	317	317.75	348	349	348.5
422.25	498.75	368.75	394.5	427	391	424.25	453.5	421.25	451.75
500.75	485.75	486.75	651	535	565.75	501	544.5	629.5	541
534.75	491.75	559.25	558.75	545.25	620.25	617.75	598.75	658.25	676.75
552.75	529.25	540.75	495.75	586.75	599.5	549.5	643.75	687	602.75
589	543.25	498.5	473.5	601.75	536	521.5	659.5	982.5	584
654.75	587.25	587	587.5	652	651.75	652.25	716.5	719.75	716.5
672	605.5	605.25	605.5	671.5	671.25	671.25	735	735.25	966.25
756	821.5	776	810.75	-	757.75	863	-	854.25	945.5
796.5	656	661.75	459.25	718.25	716	704.5	759.25	770	752.75
846.75	800.25	794.5	544	845.25	851.5	840.25	890.25	545.25	883.5
915	870.5	818.5	823	852.5	908.75	915	946	882	988
988.25	891.5	892.25	889.75	988	988.75	985.75	995.25	914.75	830.75

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
-	No match found

Positions 1 and 8 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	180.75	194	167.5	199.75	214.25	184.75	218.25	234.25	202.25
227.75	200.25	-	185.25	221.25	-	205.25	242	259.25	224.75
256	213.5	214.5	213.25	233.25	234.25	233	252.25	253.25	251.75
325.75	287.25	287	287.25	317.25	317	317.5	347.5	347.5	348
422.25	776.5	368.5	-	426.25	240.25	423.25	452.75	266.5	450.5
500.75	513.5	487.75	461	602	537	502.5	546	604.75	543
534.75	486.5	554.5	393.5	535.75	614.75	602.75	584	674.5	579.25
552.75	525	536.75	560.25	582	594.25	821	638.5	651	661.5
589	501.5	519	493	611.25	575.5	646.25	670.5	632	599.25
654.75	585.75	584.5	586	650	649	650	713.5	712.75	713.5
672	605.25	604.75	605.5	671.25	670.75	671.5	737.5	736.75	737
756	817.5	591.5	512.88	862.25	657.75	862.75	-	723	936.75
796.5	653.75	664.5	993	713.5	719.5	707.75	762.75	769	756.25
846.75	729.5	853.25	794.25	862.75	809.5	837.5	848.25	867.5	879.25
915	876	808.25	821.5	951	906.75	913.25	895.5	996.75	990.75
988.25	892	816.75	890.75	988.75	881	987	946	947.25	965.25

Positions 1 and 9 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	183.25	196.5	159.75	202.25	217.25	187.5	221.25	237.5	205.25
227.75	195.25	208.75	150.75	215.75	227.75	200	236	458.75	219.25
256	213.5	214	209.75	233.25	234	232.75	252	253.25	251.5
325.75	287	287	285.75	316.75	316.75	317.25	347	347	347.25
422.25	192	207	392.25	427.5	227.75	424.75	454.25	411.25	452.25
500.75	498.75	523.5	424	552	580.5	501.75	604.75	637	541.5
534.75	485.75	545.75	436.25	534.75	604.75	617	582.75	663.25	580
552.75	516.5	542	431	572.5	600.25	553.75	628	658	606.25
589	504	487	463	558.5	536.25	538.25	613	958.75	590
654.75	583.25	583.25	581	646.25	646.25	646.5	708.25	708.25	708.75
672	605.25	604.75	605.25	671	670.75	671.75	737.5	736.75	735.5
756	683.25	880.25	754	757.25	757.75	863	964.25	831.25	831
796.5	658.25	664	926.25	712	718.5	705.5	759	765.75	751.5
846.75	746.75	800.5	790	845	824	841	890.25	890	884.25
915	868.75	767.25	804	923.75	944.25	912.5	825.5	998.25	736.25
988.25	891.5	894.5	901.25	988	988.75	985.25	845.5	845.5	945.5

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
-	No match found

Positions 1 and 10 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	186.75	200.25	173	206.5	221.75	191.25	225.75	243	209.25
227.75	180	193	166.75	198.5	213.25	183	216.5	232.75	199.5
256	201.25	202.25	201	218.25	219.25	217.75	234	235.25	233.75
325.75	277.5	277.5	277.5	303.75	304	304	328.5	329	328.75
422.25	172	435.75	161.75	192	461	183	211.75	485.5	455.75
500.75	485.5	500.5	544.5	524.75	553.25	518	662.25	605.5	555
534.75	544.5	577	466.25	563.5	605	-	584.25	663.5	662
552.75	526	546	553.5	603.5	642	536	681.5	708.25	588
589	497	515	483.5	551.25	571.5	603.5	608.5	627.5	455.75
654.75	372.75	415.25	356	393.25	441.75	378.25	410	466.25	399
672	617	604	615.75	771.25	669.75	428.75	459.5	735.5	455.75
756	683.5	777.25	683.5	757.25	757.5	757.25	831	831	831.25
796.5	636.25	658	641.75	743.5	718.25	705.25	757	763	750.75
846.75	892.5	811	948.75	839	840.25	836.25	885.25	888.25	880.5
915	777.25	795.25	777.25	878.25	873.75	881	942.75	782.5	1010
988.25	891	891.25	889	987.5	988	985.25	990.25	989	989

Positions 2 and 3 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	187	200.5	173.5	207	221.75	192	226.75	243	210.25
227.75	204.25	218.75	189.75	226.25	242.25	210.25	248	265.5	230.5
256	229	229.25	228.75	253.75	254	253.5	278.25	278.75	278
325.75	285.5	285.5	285.75	315.25	315.25	315.5	344	344	344.25
422.25	383	385	377.75	411.25	413.5	405	437.75	440.25	430.5
500.75	472.75	449.25	451.75	494.75	758.5	496	535.75	533.25	538
534.75	491	530.75	542.25	604.25	541.25	511.75	583.75	581.75	658
552.75	543	-	466.5	516.25	523.5	549.25	663.25	575.75	576
589	529.75	559.75	495.5	587.25	620.75	542.75	644.25	681	602.5
654.75	589.25	589	589.25	654.25	654	654.25	719	718.75	719.25
672	606	605.5	606.5	672.25	671.5	673	737.75	737.75	726
756	776.25	776.75	776.25	875.75	-	758.5	832.25	832.5	832
796.5	683.25	688	661.25	715.25	727	705.25	759.25	769	750.5
846.75	795.75	794.25	795	841.5	841.5	839.5	885.5	886.75	880.75
915	820.75	815.75	822.75	912.5	908.25	914.5	998.75	998.5	993.75
988.25	891.25	891.25	891	988	988	894.25	945	519.5	945

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
	- No match found

Positions 2 and 4 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	188.75	202.5	175.25	209.25	224.25	194	229.25	245.5	212.75
227.75	201	215	186.75	222.5	238.25	207	244	261	226.75
256	229.25	229.5	228.75	254.25	254.5	253.75	278.75	279.25	278.25
325.75	291.75	291.75	291.75	323.75	323.75	323.75	355.5	355.75	355.5
422.25	390.5	391.5	390.5	421.5	421.25	421.5	449.5	449.25	449.75
500.75	452.75	452.5	452	497	496.75	496	539.5	539.5	538
534.75	486.5	486.25	475	530	534.5	526	580.5	581.25	576.5
552.75	478.25	495.5	498.5	535.5	548	552	603	599.5	604.75
589	523	551.5	490.25	579.75	611.5	542.75	636	671	595
654.75	585	585	592.25	648	648	657	710.25	710.25	721.5
672	605.5	605.25	606	671.75	671.25	672.25	737.75	737	738.75
756	681.5	681.5	682	755.75	755.5	891.5	800.5	-	997.25
796.5	752	763.25	735.5	768.75	781.75	754	786.75	800	766
846.75	801	799.5	799	851.75	825.5	844.5	902.25	878.25	888.75
915	833.25	839.75	827	917.75	918	916.25	-	931.25	948.25
988.25	891.25	891.75	891	988	806.75	975.5	935.5	946.5	834.75

	Positions 2 and 7 Missing								
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	186	199.5	172.75	206	220.75	191.25	225.75	242	209.75
227.75	204.25	219	182.5	226.25	242.5	202	248	265.5	221.5
256	229.5	229.75	229	254.5	254.75	254	279.25	167.5	278.75
325.75	291.25	291.5	291.5	323.25	323.25	323.5	355	355.25	355.25
422.25	391.5	391.5	390.5	421.5	421.25	421.5	449.5	449.25	449.75
500.75	452.5	452.25	451.75	496.75	496.75	495.75	539.25	539	537.75
534.75	485.25	485.5	-	532.75	534.25	266.75	580.5	581.75	662.25
552.75	476.25	496.25	496.75	588.5	549	550.5	588	682.75	604
589	531	541.25	542.5	601.25	622.25	600.75	645.75	697.5	658.5
654.75	589	589	589.25	654.25	654.25	654.25	719.25	719.25	719.25
672	605.5	605.25	606	671.5	671.25	672	737.75	737	738.5
756	779	717.5	954	863	-	710.25	964.75	959.5	719.25
796.5	749.75	761	733.5	764.25	778.5	751	781.25	795.75	765
846.75	807.25	771.25	801.5	781.75	818.5	846	898.25	864	889.5
915	822.5	818	823.75	914.5	909.75	915.5	946.75	-	999.25
988.25	711.25	890.75	891.75	988.5	755.75	988.25	840.25	947.25	834

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
-	No match found

	Positions 2 and 8 Missing								
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	188.75	202.25	175	209	224	194	229	245.25	212.5
227.75	204.25	218.5	189.75	226	242	210.25	247.75	265.25	230.5
256	229.75	230.25	229.5	255	255.25	254.5	279.75	280.25	279.25
325.75	291.5	291.5	291.5	323.5	323.5	323.5	355.25	355.25	355.25
422.25	390.75	390.75	389.5	420.5	420.5	420.5	448.5	448	448.5
500.75	453.5	453.25	453.25	498.25	498	497.75	541.25	540.75	540.5
534.75	486.25	486.75	473.25	535.75	536	523.75	584	582.75	573.5
552.75	390.75	496.5	544.25	525.75	549.5	603.5	604	601.75	966.5
589	529.5	559.5	495.5	587	620.25	549	644	680.75	602.25
654.75	588.75	588.5	588.75	653.5	653.5	653.75	718.25	718	722.5
672	605.75	605.25	606	671.75	671.25	672.25	737.75	737	738.25
756	707.25	681.75	681	755.75	755.5	862.25	862	945	712.25
796.5	754.75	744	740	766.75	782.75	758.75	786.75	800.25	765.25
846.75	737.25	767	800	793.25	818	843.25	849.5	866	885.75
915	825.25	817	824.75	915.25	903.75	915.75	-	-	-
988.25	716	903.5	891.5	988.75	810.5	988	947	-	947

Positions 2 and 9 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	189.5	203	175.75	209.75	224.75	194.5	229.75	246.25	213.25
227.75	200.75	214.75	186.5	222.25	237.75	206.75	243.75	260.75	226.5
256	229.5	229.75	229	254.5	255	254	279.25	279.75	278.75
325.75	291.25	291.25	291.5	323.25	323.25	323.25	355	355	355.25
422.25	391.5	391.5	390.5	421.5	421.25	421.5	449.5	449.25	449.75
500.75	452.5	452.25	451.75	496.75	496.5	495.75	539.25	539	537.75
534.75	485.75	486	941.25	524.5	562	525.5	574.25	581.5	574.5
552.75	498.25	496.25	498.75	551.5	549	552.25	604	601.25	605.25
589	522.5	551.5	488.75	579.5	611.5	541.75	635.75	671	594
654.75	586.25	586.25	586.25	650	650	657.5	720	713.25	721.75
672	605.5	605.25	605.75	671.5	671	671.75	737.25	797.75	737.25
756	707.5	742	681.5	755.5	755.5	862.5	964.5	703	831.25
796.5	751.25	681.5	735.25	767.25	780.5	753.5	784.5	797.75	765.25
846.75	544.75	780.25	800.75	852.25	825.25	845.5	898.75	874.25	889.25
915	803.75	750	824	922.75	812.75	914.75	946.75	827.5	1000
988.25	891.5	879	892.25	988.5	998.5	988.25	811.5	865.75	834.75

	0.75 < FDAC ≤ 1.0
	0.5 ≤ FDAC ≤ 0.75
	FDAC < 0.5
	- No match found

Positions 3 and 8 Missing									
Nominal	0.9 Central	0.9 Inwards	0.9 Outwards	1.0 Central	1.0 Inwards	1.0 Outwards	1.1 Central	1.1 Inwards	1.1 Outwards
221.25	189.75	203.5	176	210	225.25	195	230.25	246.75	213.75
227.75	205.75	220.25	191.25	227.75	244	211.75	249.75	267.25	232.25
256	230.25	230.5	229.75	255.5	255.75	255	280.5	281	280
325.75	291.5	291.5	291.5	323.5	323.5	323.5	355.25	355.25	355.25
422.25	389.75	390.25	388.5	419.5	419.5	419.25	447.25	446.75	447.25
500.75	454.5	454	454.75	500	499.25	500	543.5	542.75	543.5
534.75	485.75	486.75	475.25	535	535.75	528	582	583.75	576.25
552.75	475.75	497	514.25	527.5	550	603	587.75	602.5	601.25
589	531.5	562	497.25	589.5	623	551.25	646.75	683.75	604.75
654.75	589.5	589.5	589.5	654.75	654.5	654.75	719.75	719.5	719.75
672	605.75	605.25	606.25	671.75	671.25	672.5	737.5	736.75	739.25
756	838.25	881.75	714.75	644	755.5	861.5	706.5	706	945.25
796.5	759.25	772.5	744.75	775.75	784.25	759.25	787	801.25	781.75
846.75	709.75	757.5	799.25	766	800.75	840.75	817.5	849.75	881.25
915	823	819	825	915	910.5	917	946.5	947.5	999
988.25	681.5	744.5	681.75	749	828.75	988.5	856.75	893.75	831