

Harmonic Excitation of Bolted Joints

Thesis submitted in accordance with the requirements of the University of Liverpool for
the degree of Doctor in Philosophy by Matthew Oldfield.

February 2004

Abstract

Bolted joints provide one of the most common means of joining two structural components together. The joints themselves use friction to transmit force, torque and motion across a common interface from one component to another. In many cases a pretensioned bolt, running through a common hole at the joint interface, provides the clamping force.

The friction force at a joint interface is highly nonlinear. This makes the analysis of dynamic systems with joints unrealistic with conventional linear techniques. It has also been shown that the contact pressure at a joint interface is not necessarily uniform. A variable contact pressure results in a variable limiting friction load. Where the contact pressure can be shown to be smallest on an interface, local microslip can take place whilst the joint maintains its sticking contact elsewhere. Microslip is responsible for the dissipation of energy from within bolted joints that otherwise maintain their integrity.

The level of energy dissipation caused by microslip can be significantly larger than that provided by other dissipative mechanisms within a structure. This provides an incentive to be able to describe and predict the energy losses and overall joint behaviour accurately. Difficulties arise when considering 3-Dimensional contact, changing contact conditions during dynamic loading and the nonlinear nature of friction phenomena.

To investigate microslip behaviour in bolted joints a detailed finite element model of an isolated lap joint interface was constructed. The joint interface was subjected to a variety of preloads and applied torque. Output from the joint is in the form of hysteresis loops that reveal information about the energy dissipated and overall joint stiffness during a loading cycle.

Representative models are presented that reduce the complexity of the joint, yet still maintain the defining characteristics of the hysteretic behaviour. The first representative model uses Jenkins elements that match the physical response of the joint

at a number of discrete points during the loading cycle. Good agreement between the finite element model and the Jenkins element model is illustrated. The Jenkins element model is also capable of predicting the response of the finite element model when different magnitudes of preload and applied torque are applied.

The second representative model is the Bouc-Wen representation of hysteresis. This model offers significant gains in efficiency when approximating the smooth transition from a fully sticking interface to the onset of joint failure. All of the hysteresis can be described using just four parameters, and matching with the finite element model is demonstrated.

To demonstrate microslip behaviour physically an individual joint was experimentally analysed. A cantilever beam with a single lap joint near the clamped end is resonated to generate the dynamic joint hysteresis. The joint behaviour is monitored by local time domain measurements at a number of different preloads and excitation amplitudes. Microslip is demonstrated in the joint when the preload is reduced from a maximum “rigid” clamping value. Notably at low preloads the spectral content of the response reveals a large contribution from the superharmonics of the excitation frequency. Both the Jenkins element model and the Bouc-Wen model are successfully matched to the hysteresis output of the experimental joint.

Acknowledgements

The work presented here would not have been possible without the support of the Engineering and Physical Sciences Research Council (grant GR/R08520/01).

Foremost I would like to thank my supervisors Dr. Huajiang Ouyang and Professor John Mottershead. In particular I would like to thank them for their encouragement when the work undertaken was not running smoothly and their sound guidance at all times. My thanks also goes to Simon James for being a constant source of advice and practical assistance. Within the Department of Mechanical Engineering I would like to pass on my gratitude to Roy Coates for repeatedly salvaging my computer, and everyone in the Dynamics and Control Research Group who has been generous with their time and thoughts over the course of the last three years.

Finally I would like to thank my family and friends for their support and ability to make me smile.

<u>Chapter</u>	<u>Page</u>
Abstract	i
Acknowledgements	iii
Contents	iv
Chapter 1: Introduction	1
1.1) Structural joints	1
1.2) The Pressure Cone Method	3
1.3) Modelling Issues and Difficulties Associated with Bolted Joints	6
1.4) Methodology and Objectives	11
1.5) Report Outline	14
Chapter 2: Literature Survey	17
2.1) Literature Survey Summary	17
2.2) Introductory Analytical Papers	20
2.3) Further Development of the Analytical Approach	21
2.4) The Harmonic Balance Method	24
2.5) Surface Analysis and Tribological Research	26
2.6) General Investigations of Friction and Dynamic Systems	29
2.7) Microslip	32
2.8) Hysteresis Models	33
2.9) Investigations Utilising the Finite Element Method	36
2.10) Representation of Structural Joints	37
2.11) Studies of Bolted Joints	38
2.12) Joint Parameter Identification	42
2.13) Frames and Built-Up Structures	47

Chapter 3: Finite Element Analysis of an Isolated Joint	50
3.1) Finite Element Method Summary	50
3.2) Benefits of a Commercial Finite Element Package	53
3.3) The Finite Element Model	55
3.4) Implicit and Explicit Contact Algorithms	62
3.5) Pretension Algorithm	68
3.6) Finite Element Output Calculations	69
3.7) Verification Tests	72
3.8) Constant Preload and Variable Torque Magnitude	80
3.9) Variable Preload and Constant Torque	88
3.10) Variable Friction Coefficient	94
3.11) Uniformly Distributed Preload	100
3.12) Spectral Density of Finite Element Output	105
Chapter 4: The Jenkins Element Method	110
4.1) Jenkins Element Summary	110
4.2) Jenkins Element Model	112
4.3) Jenkins Element Parameter Extraction	116
4.4) Evenly Distributed Jenkins Elements	119
4.5) Optimally Spaced Jenkins Elements	125
4.6) Modelling Variable Preload Simulations	130
4.7) Prediction of Finite Element Hysteresis Using Jenkins Elements	135
4.8) Free Vibrations of Jenkins Elements	141

Chapter 5: The Bouc-Wen Method	146
5.1) Bouc-Wen Method Summary	146
5.2) The Bouc-Wen Model	148
5.3) Bouc-Wen Parameter Identification	156
5.4) Equivalent Bouc-Wen Models	159
5.5) Hysteresis Prediction Using The Bouc-Wen Model	168
Chapter 6: Experimental Analysis	172
6.1) Experimental Work Summary	172
6.2) Experimental Setup	174
6.3) Experimental Calculations	181
6.4) Tightest Joint Condition	188
6.5) Smallest Reduction of Joint Preload	193
6.6) Penultimate Reduction in Joint Preload	198
6.7) Loosest joint Condition	205
6.8) Matching Jenkins Elements to Experimental Results	209
6.9) Matching the Bouc-Wen Model to Experimental results	214
Chapter 7: Conclusions	218
7.1) Conclusions	218
7.2) Future Work	223
7.3) Summary of Contribution	226
References	228

Chapter 1. Introduction

1.1) Structural Joints

Joints are used in all manner of built up structures. The joints are used to connect two or more members at a single location to transmit force and torque at its interface. Two of the most common methods for forming these joints are through rivets and bolted connections. Bolted joints offer a significant advantage over riveted joints in that the components can be assembled and disassembled readily for transport. The magnitude of the clamping force can also be more carefully regulated in a bolted joint.

Many bolted connections use the preload in the bolt to force the components that make up the joint into contact with one another. The contact and interaction at the joint interface is what allows the transmission of loads from one member to another in a controlled manner. These types of connection are used in applications on scales ranging from small domestic appliances, through transportation to large civil structures.




Figure 1.1.1. Bolted joints used to connect components in: A simple lap joint, a multiple bolt configuration and a 3-Dimensional truss structure.

How a particular joint is configured plays a major role in the behaviour of the joint locally but also potentially on a global scale in a structure as a whole. In some instances the two clamped components are brought together by inserting the bolt into a threaded shaft in one of the components to be clamped. More generally, the bolt is inserted through a common hole in all of the components that are being clamped, and the preload is generated by tightening a nut on the opposite side of the connection. Tightening the nut results in the bolt being stretched and the tension in the bolt shaft generated between the bolt head and nut itself causes the clamped components to be held together.

Figure 1.1.1. shows the ways in which bolts can be used with clearance holes to join a variety of members together. In almost all structures containing joints of these kinds some kind of dynamic excitation of the joint will take place. In the case of rigid joints the behaviour of the structure is easy to predict. However, the frictional interface present in bolted joints makes their response to dynamic behaviour potentially complex. Friction is a highly nonlinear phenomenon, and in a flexible bolted joint the contact conditions are difficult to determine and subject to change over the period of excitation. The 3-Dimensional nature of the contact problem, nonlinear behaviour of friction and changing contact conditions make bolted joints highly challenging components to model effectively.

3.9) Variable Preload Constant Torque

The variable torque analysis has shown how important the contact pressure distribution is in determining the microslip response of the upper block. As well as the model geometry, the size of the bolt preload is the most important determining factor of the contact pressure at the joint interface. In the following tests, ABAQUS/Standard was used to apply different preloads to the bolt component, before the results were exported to ABAQUS/Explicit where a 240Nm torque was applied at 20Hz. The preloads were multiples of 19kN as this was the value used in the control test. All the preloads were increased as a 240Nm torque nearly caused failure of the joint, and decreasing the preload from 19kN meant that the joint could make the transition from microslip to macroslip (and failure). The magnitude of the preload is potentially a variable in physical situations as even a torque wrench can only provide a guide to the clamping force that has been applied. Different class joints with the same bolt geometry can also be legitimately used to apply preloads several times larger than that used in the control test.

Figure 3.9.1 shows that the contact pressure distribution remains virtually identical to the control preload when it is increased to larger values. The magnitude of the contact pressure increases in proportion to the preload at all the nodal points where measurements are taken. Only at the maximum preload of 85.5kN preload does this relationship break down, and the deviation is still only small.

The control simulation of a 19kN bolt preload and 240Nm torque is illustrated in Figure 3.9.2 and shows the full range of phenomena at the joint interface. Initially the joint is stuck and no energy is dissipated, as the torque is increased from the point of velocity reversal a small amount of microslip is encountered. Further increases in torque result in larger areas of the contact interface sliding and progressively more energy dissipated through friction. More regions of microslip appear as the torque increases and contact stiffness decreases until virtually the entire interface is in sliding contact before velocity reversal. At this point almost the entire stiffness seen in the hysteresis loop is provided by the torsional-spring-like effect of the bolt-nut component.

1.2) The Pressure Cone Method

Generally energy is lost from joints through frictional energy dissipation. The amount of energy dissipated is dependent on the amount of sliding that takes place at the joint's interface. Coulomb's friction law states that sliding will only take place when:

$$F_t \geq \mu F_n \quad (1.2.1)$$

In Equation 1.2.1. F_t represents the tangential force, μ the coefficient of friction and F_n the normal force. What the pressure cone method shows us is that the normal force is not constant over the entire joint interface. Instead the contact pressure distribution between two clamped components can be approximated by a maximum value around the bolt hole, and then decreasing linearly to the 0 if the contact interface is large enough. The rate of decay of the contact pressure is characterised by the pressure cone angle ϕ .




Figure 1.2.1. *The contact pressure distribution between two components clamped by a preloaded bolt in a symmetrical joint.*

The Coulomb friction model is often used to characterise frictional systems on a macroscopic level. However, it can be combined with the contact pressure distribution of Figure 1.2.1. to explain the phenomenon of microslip. Microslip occurs when small regions of the contact interface slide independently of other sticking regions. Microslip exists when an applied load is large enough to overcome the local limiting friction

expressed in Equation 1.2.1. as the normal force is smaller in a particular region. When microslip has progressed across the whole interface, the two components slide against each other en masse. This phenomena is known as gross slip or macroslip, and in the case of a bolted joint generally constitutes failure of the interface.

The pressure cone method is generally used for calculations involving the stiffness of the bolt. However it also serves as a good approximate tool for illustrating the variable pressure distribution at a contact interface. From this initial, relatively crude model, it was possible to approximate the magnitude of the contact pressure that could be expected at the contact interface when various complications such as stress concentrations and surface condition parameters were ignored. This value could then be compared with preliminary results from the finite element analysis to indicate whether the finite element method was behaving as expected.

Further estimates can also be made from the pressure cone model as to the amount of torque required to create gross slip at the contact interface. As the pressure cone model assumes that the pressure distribution drops to zero at the edge of the area nominally in contact then microslip begins to take place at that point instantaneously.

Assuming that the contact pressure distribution is constant around the bolt axis, the area of the joint interface can be represented as follows:

$$A = \int (2\pi r) dr \quad (1.2.2)$$

Here A is the area of a circular element at a radius r from the bolt axis.

The normal force acting on the interface over the area A can be represented by Equation 1.2.3. (below). In this case $p(r)$ is a radially dependent contact pressure distribution described by the pressure cone distribution of Figure 1.2.1. To quantify the pressure distribution, the normal force is equated to the preload applied to the bolt during tightening.

$$F_n = \int (2\pi r p(r)) dr \quad (1.2.3)$$

By analogy to the tangential force required to overcome the limiting friction in Equation 1.2.1. the tangential force can be established in the case of the pressure cone method by multiplying the normal force by the friction coefficient. It is possible that the friction coefficient may also be a function of r , but in this explanation and throughout the work carried out here it is assumed constant.

$$F_t = \int (\mu(r) 2\pi r p(r)) dr \quad (1.2.4)$$

To convert the limiting tangential friction force into an equivalent limiting torque T_l , the following development of Equation 1.2.4 is used:

$$T_l = \int (\mu(r) 2\pi r^2 p(r)) dr \quad (1.2.5)$$

This expression will give an approximation of the torque required to cause failure of the joint if integrated over the extent of the entire contact interface. However an approximation of the onset of microslip can also be obtained by integrating between the limits of a smaller region where the contact pressure was not at its maximum value.

1.3) Modelling Issues and Difficulties Associated with Bolted Joints

It has already been suggested that structures connected by rigid joints are relatively straightforward to analyse. Extracting behavioural parameters from these structures can be achieved readily, and the linear parameters, particularly the stiffness and damping characteristics, can be used to predict the response of a linear, rigidly jointed, structure to dynamic excitation. If the bolted joints between structural members are shown to be nonlinear, then extracting the defining behavioural parameters, and predicting the response of the structure to a given input excitation becomes far more involved.

Through the process of microslip many bolted joints that maintain their integrity cannot be considered linear or rigid. This makes overall analysis of the structure difficult for two reasons. Firstly the stiffness provided by the joint is nonlinear. Secondly the amount of damping in the structure is considerably greater than that predicted by material damping, and environmental effects such as viscous air damping. When analysed the damping also proves to be highly nonlinear. This is because frictional energy dissipation is a function of the sign of the relative velocity between two surfaces. Even if the direction of motion between two surfaces changes gradually, at the point of velocity reversal an abrupt alteration in the characteristic of the friction force takes place. As a consequence of these nonlinearities conventional linear techniques become redundant as analytical tools.

In applications of built up structures in outer space the amount of damping provided by the environment is negligible. This, coupled with design specifications for very lightweight structures requires additional sources of damping to be provided. Active and passive joints are one such source of this additional damping without adding significant mass to a structure.

The problem of incorporating these nonlinear joint characteristics into models of built up structures raises itself initially in modelling the actual stiffness and damping

characteristics for reproduction. The modelling process can only take place with prior identification of the nonlinear contribution from the joint and this in itself is problematic. Virtually all of the driving behaviour for energy dissipation from bolted joints comes from the conditions at the joint interface. This is a very local area that has a significant impact on the behaviour of the structure in global terms. Measurement of the pressure distribution of the joint is difficult as doing so requires a degree of intrusion into the interface itself. Similar problems exist when attempting to establish the local displacements of one surface relative to the other. These local parameters then have to be incorporated into models with components several orders of magnitude larger.

One particularly popular method of extracting the dynamic behaviour of structures is through use of the finite element method. Here the problem of scales becomes very evident. To model the defining behaviour of a bolted joint satisfactorily a small scale, highly detailed model of the contact interface needs to be produced. Such a model includes 3-Dimensional, dynamic contact. This kind of analysis places some of the heaviest demands on processing power. In bolted joints, the problems of dynamic contact are accentuated by the need to constantly reanalyse the normal loading conditions at the interface before further calculations of the dynamic response can be carried out. In certain situations the normal force and tangential motion of the interface are coupled to each other. This type of analysis can take many hours to run even without the complicating factor of attaching the joint to other members. When combined with several structural members, all with joints modelled in the required detail, the result is a computation time that is hugely impractical. Some way of representing the bolted joint behaviour without the detailed modelling of the interface is required to provide an efficient solution to this problem.

One method of providing a solution to the modelling problems outlined above is to develop a linear approximation of the joint behaviour. Such an approximation would maintain the fundamental properties of the joint's dynamic response yet would be represented in a linear manner that would facilitate the prediction of a structure's global response to dynamic excitation. To provide an acceptable linearisation of the problem it

is necessary to first identify the behaviour of the joint through its response to dynamic excitation. From this behaviour it is then possible to reproduce the characteristics that define the nonlinear response but in a linear system.

As bolted joints have a significant impact on the amount of energy dissipated from a system, one of the best ways of characterising energy loss is through hysteresis loops. These are so effective as a single hysteresis loop will reveal information about the rate that energy is dissipated, the amount of energy lost through friction, the stiffness characteristics of the joint and the range of displacement for a given excitation force. Hysteresis loops for a bolted joint are constructed by measuring the displacement of the two clamped surfaces relative to each other and then comparing with the restoring force provided by the contact interface. The exact nature of this restoring force is uncertain, and also dependent to a degree on the scale that the joint is being represented.

It is possible to represent the nonlinear equation of motion for a bolted joint as follows:

$$m\ddot{x} + f(x, \dot{x}, t) = F_{ap} \quad (1.3.1)$$

In Equation 1.3.1. m represents the inertia of the system, F_{ap} the external load, t represents time and x displacement. Overdots are used to signify differentiation with respect to time. The defining hysteresis loop is obtained by plotting the restoring force f against displacement x . At this stage the exact form of this relationship is unknown.

The simplest dynamic system with dissipative behaviour that can be modelled is that with linear stiffness and viscous damping. An equation of motion for the system shown in Figure 1.3.1. can be fully defined by Equation 1.3.2. below:

$$m\ddot{x} + c\dot{x} + kx = F_{ap} \quad (1.3.2)$$




Figure 1.3.1. *A simple 1-degree-of-freedom system with linear viscous damping.*

In the case of the linear system above, the hysteresis loop can be fully identified:

$$f(x, \dot{x}, t) = kx + c\dot{x} \quad (1.3.3)$$

A restoring force of the type defined in Equation 1.3.3. is shown in Figure 1.3.2. By matching the defining qualities of the unknown model with the linear model presented with it, a suitable approximation of the bolted joint will have been achieved. This joint could then be represented by a single linear damping and stiffness term in all subsequent analyses. Future analyses on the joint or built up structure could then be performed by linear methods. The advantages offered by these methods are that they are readily interpretable, thoroughly tested and computationally efficient.




Figure 1.3.2. *A hysteresis loop representing a linear system with stiffness and viscous damping.*

1.4) Methodology and Objectives

It has previously been discussed how bolted joints provide a significant proportion of the damping that exists in built up structures. The prevalence of this type of joint has led to their dynamic behaviour to be of importance in both the modelling and understanding of systems. However, there are inherent difficulties associated with the dynamic analysis of bolted joints in both their identification and representation. These themes lead to the formation of five objectives that were the impetus behind the work carried out here:

- 1) Establish the state of the art in the field of research into frictional energy dissipation, modelling of systems containing friction and dynamic analysis of bolted joints.
- 2) Obtain a full understanding of how microslip dissipates energy under the dynamic excitation of bolted joints.
- 3) Find reduced order models capable of replicating the hysteresis obtained when microslip was the predominant means of energy dissipation. These reduced order models should facilitate the conversion of the nonlinear frictional mechanisms into a single linear spring-damper first order model.
- 4) Verify the predictions about the microslip energy dissipation mechanisms through experimentation.
- 5) Show that the reduced order models are capable of representing the experimentally obtained hysteresis loops in all of their main characteristics.

These objectives have been met through a variety of steps covered in this research. The state of the art in the field of mechanics surrounding the dynamic behaviour of bolted joints was established by conducting a literature survey. This literature survey revealed the extent of the research that had been carried out already and some of the limitations of the work carried out by others to this date.

The mechanism of microslip was investigated by constructing a detailed finite element model of a bolted joint. This detailed model isolated the joint, and used a precise

algorithm for determining the amount of sliding across the joint's interface. Hysteresis loops were generated for the contact interface in the finite element model. These loops illustrated clearly the extent of slip that existed in the joint when it was harmonically excited. By comparing hysteresis loops obtained at different preloading conditions and applied loads a broad picture offering evidence of the slip mechanisms taking place was obtained. All of the hysteresis loops obtained from this detailed finite element were used as the basis for further models that were devised.

The literature survey highlighted several different hysteretic models that had the potential to replicate the findings of the finite element analysis. Not all of these models were based in the field of structural dynamics. Material plasticity also offered a good model for hysteresis. By using analogous parameters between the loss of elastic stiffness in a yielding material, and the loss of stiffness in a joint as microslip progressed, a suitable model was identified. Jenkins elements have been used extensively to model frictional systems and were implemented in a second order system with identical model parameters to the finite element model. A further model, known as the Bouc-Wen model of hysteresis, was also used to replicate the behaviour exhibited in the finite element analysis. The efficiency of both of these reduced order models was tested for a number of different loading conditions.

Once an understanding of microslip had been provided by the work using the finite element method an experimental rig was designed and manufactured to offer a physical demonstration of microslip phenomena. The rig was design in such a way that very local measurements of the joint interface could be performed in the time domain. Measurements of this type were not common in many other investigations of bolted joints. Due to the limitations of not being able to monitor the behaviour on the contact interface directly data was collected as close to the mating surfaces as physically possible.

The experimental results were converted into hysteresis loops to compare the behaviour with that established by the finite element model. Good qualitative agreement

existed between both the finite element hysteresis and the experimentally obtained hysteresis. This allowed the two reduced order hysteretic models to be matched with the experimental output. Confirmation was therefore obtained that the reduced order models could be used to approximate the physical behaviour of bolted joints. Close agreement between the experimental and finite element findings also confirmed a means of predicting joint behaviour without the need to generate many complicated and costly experimental models. By creating a finite element model of a hypothetical isolated joint, reduced parameter equivalent models could be created to replicate computed hysteresis behaviour. The reduced order models could then be substituted into a multi-jointed finite element structure to offer significant savings in computational effort and cost when conducting modelling and dynamic analysis.

1.5) Report Outline

The research carried out into the dynamic behaviour of bolted joints is presented in the following report. The work was broken down into the sections outlined below to provide an overview of the process from its inception to completion.

Chapter 2 details the literature survey that was carried out. In the literature survey a retrospective view of the investigations into bolted joints is taken. However, the related fields of dynamic systems with friction and modelling frictional systems are also investigated to provide both context and relevant background information. Much of the work that has been carried out concerning dynamic systems with friction is associated with managing the nonlinearity of the equations of motion. Many analytical papers dealing with this subject matter are presented. Studies have also been produced in which the friction force is characterised under different physical conditions. These characterisations take the form of both analytical/numerical studies and experimental investigations. The main findings are presented in this chapter. Of particular relevance in this chapter are the papers specifically regarding the dynamic behaviour of bolted joints. Again the research presented was conducted in both the analytical and experimental domain. All of the papers that are considered are put into context by the presence of previous reviews conducted in the relevant fields. Research concerned with developing the analytical tools such as the harmonic balance method and finite element analysis, to facilitate the investigation of bolted joints is also reviewed to provide context for the techniques used herein.

Chapter 3 concerns the finite element work that was carried out into an isolated joint. The built-in algorithms that facilitate complicated 3-Dimensional contact analyses give an indication of the motivation for using a commercial finite element package. Further preliminary methods that were used to validate the finite element model are presented. The validated model of an isolated joint is then used to investigate the effect of varying the magnitude and type of preload and the impact this had on the response to harmonic excitation. Other parameters investigated include the coefficient of friction at

the joint interface and the magnitude of the applied torque. All of the results are presented in the form of hysteresis loops. Investigations of the time domain response of one clamped component relative to another are presented along with the power spectra response of some representative behaviour.

Chapter 4 presents the first reduced order model designed to represent the hysteretic response of a bolted joint to dynamic excitation. The Jenkins elements model is introduced and the method of extracting the model parameters from hysteresis loops is demonstrated. The efficiency of Jenkins elements in representing hysteresis loops characterised as bilinear as well as those exhibiting a large degree of microslip is investigated. Additionally, the Jenkins element system was subjected to initial displacements. The resulting free vibrations are discussed as they demonstrate different characteristics not clearly demonstrated in the case of harmonically forced oscillations.

Chapter 5 presents and develops another reduced-order hysteresis model. The Bouc-Wen model is capable of representing both a hysteretic restoring force and material plasticity. The impact of each of the four defining parameters of the model are discussed and illustrated. Consideration is also given to extracting these parameters from hysteresis loops. Bouc-Wen hysteresis loops are matched to a selection of finite element loops and discussion of these results is given. The ability of a single Bouc-Wen model to recreate the hysteresis obtained at different input conditions is also investigated.

Chapter 6 is a discussion of the experimental aspect of the work that was carried out. Initially a description of the experimental rig and the methods used to extract the joint hysteresis are presented. Hysteresis loops obtained from the joint at different clamping forces and amplitudes of oscillation are illustrated and related to the microslip phenomena taking place at the joint interface. Of particular significance, the power spectrums of the response for some of these tests are shown. The impact of super harmonics of the excitation frequency is illustrated as being of particular significance. To complete the investigation, Jenkins element loops and Bouc-Wen models are matched to the experimental loops in the same manner as the finite element results.

Chapter.7 details the main conclusions of this investigation. The limitations and principal findings are documented. A short summary of proposed further work in relation to that already carried out is presented. To conclude the report a summary of the novel aspects of the research that has been carried out is delivered.

Chapter 2. Literature Survey

2.1) Literature Survey Summary

Before an investigation could be carried out into the energy dissipated from bolted joints, the state of the art in the field had to be established. To achieve this a literature survey was conducted to establish the work that had already been performed on the subject of bolted joints. Research also undertaken in the related field of frictional energy dissipation along with the analytical, numerical and experimental approaches taken in its investigation also needed to be determined.

The practical applications of joint friction were shown to extend from transportation, to civil structures and particularly space structures where the atmospheric and material damping is very low. In civil structures the purpose built frictional isolators were shown to dissipate the energy transferred to a frame during seismic activity. Space frames utilised active and passive joints to dissipate the motion imparted on the structure during deployment and positioning. Frictional joints were illustrated as an ideal source of damping as they added no mass to the overall system and could be tuned to specific requirements.

Many studies showed that the criteria used to define the response of the structure directly related to the joint configuration. The issue of the discrepancy in scales used to characterise joint friction in relation to overall structural dynamics was highlighted as being a particularly difficult issue to overcome. Experimental and analytical studies revealed that the amount of preload that minimised the joint's motion was generally not the same as the amount that maximised the energy dissipation. In these studies it was also shown that the resonant frequency of a system was also dependent on the bolt preload. Looser joints contributed to softer systems, with lower resonant frequencies as a result of the overall of contact stiffness being directly related to the magnitude of bolt preload.

Many analytical and numerical tools were suggested for tackling the problem of frictional systems and joints in particular. These ranged from the finite element method to interpretation of Nyquist plots. Among the different techniques proposed, the most common was the harmonic balance method. However, this method had one significant drawback. Governing equations could only be derived using sticking contact or sliding contact. Further harmonics were required to describe the complicated sticking and sliding that could be a feature of microslip contact.

The amount of damping supplied by a frictional system was in some cases characterised by a single representative stiffness and damping coefficient. These were extracted from both experimental and numerical studies with a particularly popular method being to use the decay rate of free vibrations. In frictional systems the decay rate was shown to be characterised by a linear decay. This was in contrast to the exponential decay rate present in predominantly viscously damped systems.

The system properties of bolted joints were shown to be dependent on many different physical parameters. These have been studied on a microscopic level by a number of authors both numerically and experimentally. Numerical studies have concentrated on the characterisation of contact surfaces and the asperity distribution that is so critical to the frictional contact that is formed. Experiments have shown that the friction coefficient at a surface was generally more complex than the simple signum function known as the Coulomb friction law. Many researchers have found that the friction coefficient between two metallic surfaces was a function of the time they were at rest and the relative velocity between them. Most commonly a friction coefficient that decreases with velocity was presented.

Much research has also been conducted into more general systems with a restoring force containing a friction component. In particular the presence of friction induced vibrations and the stick slip motion possible in 1-Dimension or more have been investigated. The main findings in this type of study were analytical and showed that whilst sticking was an inevitable consequence of 1-Dimension harmonic motion, it did

not necessarily exist when the motion at an interface was 2-Dimensional. Virtually all of these types of studies, while providing relevant background into the techniques used to deal with frictional problems, considered contact on a macroscopic level. Interfaces were either considered to be sticking or sliding. While helpful to characterise the nonlinearity of the friction that exists in bolted joints, microslip could not be reduced to this overall sticking and sliding contact. By its very definition microslip exists as a regime *between* sticking and macroscopic sliding contact.

The hysteresis that microslip provides in a bolted joint was also studied by several authors. Notably hysteretic models not directly associated with structural mechanics were used to characterise microslip. Material plasticity was a popular model used where analogous parameters were established for the degree of plasticity in a material and the amount of microslip that existed at a contact interface. The loss of elasticity in a joint due to microslip and the onset of material plasticity with hysteretic cycling were popular themes, and provided efficient models for replicating the behaviour of bolted joints.

2.2) Introductory Analytical Papers

Den Hartog (1931) provided probably the earliest paper on oscillations with damping provided by friction. It is an analytical study with some experimental evidence of the response to forced vibrations. Initially solutions are provided for second order systems subjected to harmonic excitation with just dry friction. The method is also developed to a system with viscous damping and Coulomb damping. Equivalent viscous damping was also investigated as the means of replacing other forms of damping (Coulomb, hysteretic) at this time. Using the most basic form of Coulomb friction, Den Hartog found non-stop solutions when combining Coulomb friction with viscous damping. The steady state is defined as having the same frequency as the excitation and upward and downward portions of the cycle following the same law. Sudden changes of curvature can be seen at the peaks as curvature of the displacement profile is equivalent to the second derivative in the time domain. The acceleration is also proportional to the restoring force, which is in turn partially defined by a signum function, causing the sudden changes in gradient. Significantly, with the assumptions applied and non-negligible friction an infinite response at resonance is possible. Away from these particular resonance conditions an increase of friction force reduced the response gain. Stopped motion could not be expressed readily in explicit form yet experiment verified multi-stop motion under certain conditions. With viscous and Coulomb damping non-stop motion is described analytically, but any numbers of stops per cycle have to be described by piecewise or numerical integration techniques. A fundamental paper providing closed form solutions to models with basic frictional constraints was contributed by Hundal (1979). The model is similar to Den Hartog's but with base excitation. Both stick-slip and non-stop motion are analysed using a non-dimensional equation of motion. For given values of viscous damping coefficient bounding curves are shown where the friction force becomes large enough to cause one stop per cycle. The potential for more than one stop per cycle is also stated.

2.3) Further Developments of the Analytical Approach

Feeny, Guran, Hinrichs and Popp (1998) recently reviewed stick-slip friction. The paper quotes a value of 5% of gross domestic product lost as a consequence of friction and wear to indicate its relevance as a field of research. Friction as a mechanism is still not widely enough known for its effect to be predicted for an arbitrary set of conditions, and very basic models (Coulomb) are still popular for modelling the phenomenon. A thorough list of parameters that can affect friction is given. These range from:

- Surface conditions.
- Environmental conditions.
- Sticking behaviour.
- Sliding and relative velocity dependent parameters.

Normal motions are indicated as both a possible outcome and driving factor of frictional systems. General chaos and non-smooth behaviour is discussed as even when friction provides the only nonlinearity chaos can still exist in a system. Logarithmic damping is given as a useful tool for estimating the amount of Coulomb damping in a system.

Swayze and Akay (1994) used a Coulomb model to study friction-induced vibrations akin to a carbon brush running along a surface. For coefficients of friction higher than a critical value the origin was found to no longer be an equilibrium point and an offset was created. Critical values of several parameters in the system such as the torsional stiffness of a spring connected to the block and dimensions of the block could stabilise the friction-induced vibrations that were seen in the analysis. Bengisu and Akay (1994) investigated vibrations, and in particular “locking in” of modes between friction induced vibrations and resonance of the system. The models used were multi-degree-of-freedom systems with a friction force applied to one of the degrees of freedom only. The system response was found to be dependent on the slope of the friction force. Phase

space, phase plane responses and root-loci were utilised to assess stability and limit cycles-where the phase space trajectories depended on the number of bifurcations present. Ding, Leung and Cooper (2001) looked in detail at the combination of a system with both hysteretic and Van Der Pol damping. The most notable phenomena identified were bifurcations and again mode locking. It was found that by increasing the hysteretic parameters the sub-harmonic resonance could be reduced. Two frequencies were present in the response, the frequency of excitation and the nonlinear resonant frequency.

Another area of analytical investigation is that of hysteretic damping. Maia, Silva and Ribeiro (1998) looked at the general behaviour of damped systems. Frequency dependent damping forces that are given in terms of fractional time derivatives have been investigated. Setting the power of the fractional derivative to 0 gave hysteretic damping and Viscous damping was achieved by setting the power to 1. Pratap, Mukherjee and Moon (1994a) considered free oscillations of a hysteretic damper in which a hardening parameter is used to define the stiffness of the second part of the bilinear hysteresis loop (the non-dimensional equations for which are solved in a piecewise manner). The hardening parameter has a critical value above which limit cycles are observable. This method allowed the finite number of elasto-plastic oscillations to be calculated before an offset elastic oscillation takes place. Bending moment and angle of rotation were used to define the hysteretic response of a joint at the midpoint of two rigid beams. Pratap, Mukherjee and Moon (1994b) identified the same offset when the system was subjected to periodic impulse loading. Period-doubling bifurcations were not present, but period-two orbits existed when the impulse loading was applied at specific time increments. These orbits were purely elastic and not a result of the bifurcation of the period-one elasto-plastic orbit.

With harmonic excitation stick-slip will always be prevalent if sliding exists. Griffin and Menq (1991) found that in 2-Dimensional motion this was not necessarily the case and a situation could exist where continuous sliding takes place. If subjected to circular motion continuous slip will take place if the aspect ratio of a general case ellipse is large enough. For a periodic response the exact solution shows that the path of a point

followed a closed loop. The phase angle between the motion of the mass and plate (between which the physical contact occurs) was significant as it showed the degree to which the damping displays viscous-like behaviour. The authors showed that in comparison to a 1-Dimensional system 2-Dimensional motion is more effective at reducing the peak amplitudes of vibration. This is attributed to the continual sliding versus stick-slip motion. It was also found that the reduction in amplitude of the vibration is far more significant than the difference in optimum slip load between 1-Dimensional and 2-Dimensional cases. Menq, Chidamparam and Griffin (1994) found this optimum normal force to lie between sticking and continual sliding situations. However, optimisation of one mode in a particular direction may not result in optimal behaviour of a second mode. Sticking or sliding was very much dependent on the magnitude of the friction force rather than the components in either direction. (Again the problems of using a single harmonic in the harmonic balance method are highlighted as being crude in certain specific situations.)

Ferri's (1995) review paper accentuates the positive applications and analytical aspects of friction damping particularly in turbo-machinery blade attachments and seismic isolation. The ability to transmit only limited forces across a slipping interface is seen as most significant in these applications. There is also a considerable section devoted to frictional joints and their positive applications within built up and truss structures. Displacement dependent friction forces are highlighted in particular. For non-experimental studies Ferri advocates exact solutions, numerical integration of system equations and the harmonic balance method as important solution tools. Analytical tools such as these are utilised in studies over three areas:

- Turbo-machinery.
- Transportation (particularly railroad).
- Building systems-notably earthquake isolation.

Ren and Beards (1994a) produced an analytical paper that could deal with local *nonlinear* elements. However, the problem of time integration and finite element

analysis, for steady state response, in terms of computing cost was highlighted. Ren and Beards (1994b) analysed nonlinear friction elements and microslip systems by a method known as the receptance-based perturbative multi-harmonic balance method. This method was capable of producing frequency domain results without needing costly time domain procedures. The algorithm obtained parameters using receptance data and was able to predict the response of other inputs with different excitation. The approach allows responses, for conditions similar to those already found, to be obtained with only few initial calculations. Ferri and Heck (1998) also used singular perturbation theory, but as a means of getting round the numerical stiffness of modelling both high and low frequency resonances of systems containing Coulomb elements. In one simulation two resonant peaks were exhibited in the time integration. None of the reduced order models were able to capture this second peak. However, all the reduced order models significantly decreased the computation time in comparison with the time integration procedure. The first-order-approximation model, while delivering the best results of all the reduced order models required relatively little increase in computational time to complete its analyses when compared with the other two computational models.

2.4) Harmonic Balance Method

Pierre, Ferri and Dowell (1985) used the incremental harmonic balance method to solve the equation of motion of a one or two degree-of-freedom model in the frequency domain. The prime benefit of the incremental harmonic balance method was its efficiency at generating parameter studies in the frequency domain. While the method was good at predicting amplitude and phase of response to harmonic input, at low frequencies (much below natural frequency) it was not good at predicting the multiple steps per cycle of the response. Consequently more harmonics are required to accurately describe the frequency response in this range which is computationally more demanding. Like many harmonic balance studies the equation of motion assumed only sliding contact whereas the time integration was capable of separating regions of stick and slip.

Ferri and Bindemann (1992) developed the harmonic balance method to three types of boundary conditions for a beam clamped at one end. Their findings showed transverse friction yielding damping that is independent of the vibration amplitude of the beam and capable of an unbounded response analogous to Den Hartog's work. The second boundary condition, dependent on the beam's in-plane displacement, yielded hardening spring like behaviour and under certain conditions viscous-like damping. The third model was dependent on the flexural displacement of the beam and gave hydraulic damping. Again only one harmonic was used in the harmonic balance accounting solely for sliding contact. Both the likelihood and criteria for sticking to take place were analysed. Time integration analyses also showed beating that took a very long time to die away before a steady state was obtained. By modifying the beam geometry it was shown how to tailor damping from friction and the importance of modelling the joint correctly. Longitudinally dependent vibration was compared to a self-locking thread under certain circumstances. This locking needs to be avoided for maximum energy dissipation as sliding needs to take place.

Whiteman and Ferri (1997) utilised time integration using modal approximations for its solution. More modes showed the existence of sub-harmonics, generated by the stick-slip condition, not identified by harmonic balance. The normal force for friction was proportional to the transverse displacement of the beam and provided linear structural damping-where the damping coefficient was largely unaffected by changes in the amplitude of the response. If the frequency range goes beyond that of a second mode, more than one mode should be used in the harmonic balance to represent the modal dependence of the frictional force behaviour.

Menq, Griffin and Bielak (1986) also used the harmonic balance method to analyse a single degree-of-freedom system. In the authors' model a friction damper was used which had a normal force was displacement dependent. Allowing lift off and open contact expanded the model from other studies. First order approximations of Fourier series were used to describe the system and a first order harmonic balance to approximate the solution. The conclusion was drawn that there existed an optimum normal force

between loose and rigidly clamped conditions. A coupling term between motion and friction force was also shown to affect both the frequency of maximum response and the frequency peak amplitude when they were calculated as a function of the normal force. Bindemann and Ferri (1995) carried out time integration of a simulated sleeve joint. Clearances were left in the model and a retaining spring was also included. Five modes were used in the calculations with the logarithmic decrement to establish the damping capacity of the joint. Clearances and a stiff retaining spring both reduced the amount of sliding contact and hence the damping ratio. Depending on the restraining spring stiffness the natural frequencies illustrated a hardening spring-like behaviour and viscous-like damping. A Good agreement between the two methods was achieved-except at low frequencies.

2.5) Surface Analysis and Tribological Research

In a bid to understand joints on a more local level it is necessary to investigate a range of papers concerned with microscopic behaviour and surface interactions. Oden and Martins (1985) provided a comprehensive and very widely cited review of friction models and dynamic friction phenomena. The authors defined dynamic friction forces as being affected by:

- Constitution of the interface.
- Time scales and frequency of contact.
- Response of the interface to normal forces.
- Inertia and thermal effects.
- Roughness of the contacting surfaces.
- History of loading.
- Wear and the general failure of the interface materials.
- The presence or absence of lubricants.

The review is divided into three categories; quasi-static friction, dynamic sliding friction and wear and ploughing. Much attention is given to the surface profiles and their

response to normal loads both static and oscillating. Gaps in the available tribological data hindered using microscopic data to define dynamic friction models. This made the coefficient of friction unreliable but to an unquantifiable degree. In the resultant model a differing static and dynamic friction coefficient was used. Oscillations in finite element results and numerical analyses are put down to the drop from the static to dynamic coefficient of friction.

Woo and Thomas (1979) presented a review of surface measuring methods for surfaces in and out of contact. The main finding from many different studies was to show that area of contact varies with load raised to the power 0.8. The methods used to study surfaces actually in contact are:

- Thermal and electrical conduction.
- Neutrographic method.
- Paints and radioactive traces.
- Replica methods.

The neutrographic methods refer to a type of radiograph to be taken of the contact interface. A grease that absorbs neutrons is put between the two surfaces. When the surfaces are then forced together the absorbent grease is squeezed away from the areas of contact. An exposure can then be taken of the contact interface using the contrast between the regions where contact occurred and regions where grease remains and neutrons are absorbed. Although uncertainties and inaccuracies are inherent with all of these methods qualitative and relative results may be obtainable. Tabor (1981) combined ideas of surface modelling and experimental analysis to form a review of friction from a tribological perspective. The review investigated three different aspects of contact. These were true surface contact area, the strength of the bond formed when contact is achieved and the material modifications that take place during sliding. In the review friction force was given as the product of contact area and the shear strength of the material. The shear strength was also described as being a function of the pressure that was applied to it.

Doege, Kaminsky and Bagaviev (1999) also investigated the microscopic behaviour of surfaces and their influence on friction and how microscopic changes in local resolution of the surface can impact on friction. Finite element programs use macro scale mechanics, but these can be determined by micro scale alterations. Consequently the authors matched the evolution of the porosity in a finite element study to the change in microscopic roughness measured experimentally. Barber and Ciavarella (2000) viewed contact as significant as the only other way of applying a force is through body forces and fluid pressure. They too coupled a microscopic level of detail to a macroscale approach to the outcome. This enabled the discovery that friction induced vibrations can exist even with a uniform coefficient of friction. Hess and Wagh (1994) were also concerned with normal vibrations in contact simulations and the chaos associated with them. Two models of interface stiffness were used. The first was an analytical rough surface description, and the second an empirical power law stiffness description. Unusually instead of the Coulomb friction law the adhesion friction law was implemented to allow decoupling of the normal and friction forces. The normal contact dynamics did affect the friction force yet the friction force would not affect the normal contact dynamics. It was also found that tiny normal displacement had a significant impact on the friction force.

Bengisu and Akay (1997) derived a pair of alternative expressions to give the dry friction force on account of the surface profile in two dimensions. The model made advances in that it did not just assume contact at the tips of asperities of the surface. Both surfaces were considered to be rough and were derived from statistical models of the surface profiles. One expression for dry friction was based upon the volumetric deformation of the asperities while the other relied upon the total tangential contact area. It was found that the friction force was dependent on normal and tangential motions with both adhesive and mechanical forces required to calculate the energy dissipated. Bengisu and Akay (1999) then combined work on stick-slip vibrations and determination of friction force from surface profiles. Stick-slip oscillations were found to exist, but only in a narrow band of platform velocities. Outside of these ranges oscillations could be sustained, but they were not stick-slip oscillations. Damping could prevent stick-slip

oscillations-which only occurred in the presence of both adhesion and deformation components of frictional forces.

2.6) General Investigations of Friction and Dynamic Systems

Berger (2002) reviewed dynamic friction models and found four phenomena significant when modelling friction in a dynamic context:

- Reducing coefficient of friction with increasing velocity.
- Increasing sticking friction as dwell time increases.
- A phase lag between the relative velocity and friction force.
- Microslip between two surfaces at an interface.

Many friction models are reviewed on both microscopic and macroscopic levels. Most significant are the conclusions regarding the problems of scale in dynamic friction modelling. Very small-scale effects at joints have an impact on a structure many orders of magnitude larger in scale. Joint damping was given as related on scales from 10^{-6} to 10^0 metres. Relationships between the order of the friction model and various performance criteria such as suitability for the harmonic balance method, computational efficiency and the physical appeal of the model were collated.

Karnopp (1985) devised a friction model that is often used in control applications. Linear models, it is claimed are unsuitable for describing chatter, final sticking and limit-cycles around a final position. The proposed model reduced the need for very short time steps and was based on a small, but finite region around zero velocity where all velocities were assumed to be zero. When inside the zero velocity region the friction force was determined by other forces in the system until a breakaway force was reached. The zero velocity region could be quite coarsely adjusted while still maintaining adequate performance. Tan and Rogers (1998) developed a friction model similar to Karnopp's for a multi-degree-of-freedom situation. The model is described as a force-balance model and was compared with numerical solutions found using modal superpositioning or direct

time integration of a four-degree-of-freedom model. Numerical chatter (occurring around points of zero velocity) was eliminated when using the new friction model. Agreement was found to be good in tests over a short time span only.

Andreaus and Casini (2001) considered a sprung mass on a moving base. Two friction laws were used including a Coulomb friction law and one modelling something analogous to the Stribeck effect. A critical velocity was found for the belt speed above which stick-slip did not take place and the oscillator response was independent of the base speed. Markedly different responses were found for the two friction laws. Hysteresis loops show that the amount of energy dissipated is not the same and that the period of the velocity dependent friction law can be twice that of the basic Coulomb friction law.

Brockley and Ko (1970) produced a paper considering friction-induced vibrations using the pin and disc-style of experimental setup. The shape of the phase plane plot allowed the shape of friction force versus velocity curve to be inferred. Quasi-harmonic friction-induced vibrations could exist if the friction-velocity relationship contained a small hump. Karius (1985): performed experimental validation on a non-joint model. Stop motion was dependent on frequency ratios, the magnitude of the excitation, adhesion and its duration and the size and nature of the dry friction present. The investigation showed that instabilities existed and how the bifurcation points could be calculated. Feeny and Moon (1994) conducted a numerical and experimental study of a mass-spring system with dry friction and forced oscillations. A state variable friction law was found to represent the experimental results (in Poincaré maps) that a traditional and smooth Coulomb friction law could not. The simpler Coulomb laws, although not matching up so well in the 3-Dimensional phase space, were said to be adequate for many applications.

Ibrahim, Zielke and Popp (1999) used free vibrations to establish model parameters when considering metal on metal contact. The friction force was found to be periodic with a superimposed random fluctuation. A higher frequency response of the normal force was also found at a multiple of the excitation frequency. The friction

coefficient varying depending on velocity and normal force was established. Also the friction force was found to deviate in a nonlinear way with relative velocity that was associated with hysteresis. Liang and Feeny (1998) used experiment to investigate compliance in a system including friction and associated sticking and sliding phenomena. The concept of transition speed, where the speed of the mass when the transition from sticking to sliding is made, was introduced as an analytical tool. The experimental rig allowed hysteresis plots to be produced for both displacement and velocity against friction force. Parameters were then extracted to obtain stiffness and transition velocity. Both micro and macro scale stick-slip phenomena were exhibited with the former occurring when the friction force reached its maximum value in transition. Increasing the damping factor of the contact could eliminate this microscale stick-slip.

Gaul and Nitsche (2001) provided a good review of both friction and sticking and sliding phenomena and placed particular emphasis on mechanical joints. A figure of 50 years is given as the amount of time over which joint friction has been studied. The authors conveniently divided friction models into two distinct factions-phenomenological and constitutive models. The phenomenological models were subdivided into static and dynamic models and are based upon experimental observations. Extrapolating the microscopic behaviour that takes place at the joint interface formed the constitutive models. "Elasto Slip Model" was classed as a static model, but is another name for Jenkins element formulation used in dynamic studies such as here. The LuGre model is based on a model developed by Canudas de Wit et al. (1995) and envisaged as contact between elastic bristles. The Valanis (1980) model of plasticity was also ventured as a useful modelling tool. Among the issues raised with constitutive models were microslip and microscale deformation, and fractal representation of surface profiles. Several solution techniques were put forward including the harmonic balance method, time integration procedures and exact solution techniques for certain physical situations. Hysteretic systems were also reviewed and weight is given to studies of the optimal amount of energy that is dissipated. Among the applications put forward in this paper for harnessing the dissipative behaviour in joints were semi-active joints-developed by the authors and frame structures containing many bolted joints.

Karamis and Selçuk (1993) also studied friction in respect to mechanical joints. The effect of surface finish and tightening torque on the behavior of bolted joints was investigated. Despite the different roughness and preloads, the coefficient of friction lay in a narrow band of between 0.12 and 0.33. Smooth surfaces appeared to generally offer the stiffest joints when looking at torsional angle and torsional moment. However, the friction present was dependent on the complicated interplay between a variety of factors and parameters.

2.7) Microslip

In joints that are bolted and not pinned it is microslip and not macroslip that is the main mechanism of energy dissipation. Metherell and Diller (1968) provided a basic introduction to the phenomena exhibited by lap joints. Their model may be considered an over simplification as the pressure distribution was uniform, but many of the most important characteristics were visible. Microslip was shown to progress from the outer edges of the lap inwards, and the instantaneous energy dissipation was shown to be non-sinusoidal despite a harmonic input force. Hysteresis loops were also generated and the “effective” stiffness decreases as the input amplitude increased. Historical effects being responsible for hysteresis were described with regions of constant stiffness in various hysteresis loops highlighted. In linear systems the energy dissipation was said to be proportional to the square of the amplitude while in the case investigated it was proportional to the cube of the amplitude. Groper (1985) defined microslip and macroslip with particular respect to bolted joints. Most important factors defining the energy dissipation were established as the bolt preload and coefficient of friction that varies according to clamping pressure. The method incorporated a coefficient of friction that was dependent on normal load and consequently radial distance from the bolt axis. Groper advocated that High Strength Friction Grip Bolted Joints are used in conjunction with clearance holes and the occurrence of macroslip constituted failure of the joint. A function was also given for defining equivalent viscous damping obtained from hysteresis loops.

Masuko and Ito (1974) investigated microslip hysteresis through quasi-static experimental tests. A large discrepancy was found between the initial loading cycle and subsequent loading cycles where the displacements were almost elastic for high preloads. At lower preloads the flow pressure of the material in question was found to have a strong influence on the type of hysteresis loops observed. Shin, Iverson and Kim (1991) performed modal analysis on a multi-bolted structure with varying preload. When the bolt torque was reduced below 40% of the maximum torque a decrease of the natural frequencies and response amplitudes of the system resulted. Microslip was given as the most likely reason for this and was predicted from the mode shapes. The half power point was used to estimate the damping factor for which various trends were observed dependent on the band of frequency being examined.

2.8) Hysteresis Models

Many hysteresis models are capable of modelling microslip. By using any model that is simply more than bilinear, with a region of sticking contact and a region of gross slippage there is the capability of modelling microslip. It is also possible to make an analogy between elastic and sliding contact and elastic and plastic material behaviour. This has been a common approach adopted by many authors. Iwan (1966) produced an early paper on Multiple Jenkins elements (although not named as such). Jenkins elements connected in parallel were presented for what the author reported as the first time. A thorough description of restoring forces over a loading cycle, and the significance of the displacement prior to velocity reversal was presented by the author. The equations of motion were non-dimensionalised and a parameter was used to define the sharpness of the corners of the loop. Parameter extraction from static tests and gradients of obtained hysteresis loops was shown. It was also found that a softening effect was generated in the frequency domain and an unbounded response was predicted for bilinear and Jenkins element systems with a finite input force in both cases.

Canudas de Wit, Olsson, Åström and Lischinsky (1995) produced one of the most popular, well-cited and versatile mechanical friction models. While generated with the purpose of control in mind, the model is capable of modelling hysteretic behaviour admirably. All of the following phenomena were incorporated into the proposed model:

- The Stribeck effect.
- Hysteresis.
- Microslip.
- Varying breakaway force.

The model was based on the deflection of elastic bristles that can be related to the surface characteristics through a statistical model. The hysteresis of the model was provided through microslip, and the breakaway force is dependent on the rate of change of the applied force. Microslip was referred to as the Dahl model and the Stribeck effect is similar to a velocity-dependent coefficient of friction. Many of the other characteristics were explored experimentally. The varying breakaway force was found to be dependent upon the rate of increase of tangential force but independent of the dwell time. Although defined by only five parameters the model itself was illustrated as both powerful and versatile.

Ramberg and Osgood (1943) introduced a popular method for modelling stress-strain hysteresis that has been utilised by many authors. Its main advantage was being able to non-dimensionalise the force displacement relationship to create a three-parameter model. The model was not quite as flexible in shape as the Bouc-Wen model (1976), but non-integer powers of n could be used to define the rate of onset of plastic deformation. The extraction of parameters was illustrated to be a straight forward process when the results were plotted in log-log format. Shi and Atluri (1992) illustrated the damping provided by joints in space structures in the absence of other forms of energy dissipation. A bilinear Coulomb model was used to describe slipping at joints as one form of hysteretic damping as well as utilising the Ramberg-Osgood model to represent nonlinear stiffness of joints where friction damping existed. The models when applied to multi-

beam structures in two and three dimensions showed a noticeable reduction of amplitude when compared with a rigid jointed equivalent. A limitation of the bilinear model is that it would reduce to elastic response when the amplitudes were small enough whereas the "Ramberg-Osgood" joint always maintained a degree of energy dissipation.

Lenz and Gaul (1995) utilised a joint as a vibration reducer. Longitudinal and torsional vibrations are considered in the microslip to macroslip range. The Valanis model associated with material plasticity was used to represent hysteresis of a joint and was capable of modelling both transient and cyclic loads. Parameters were used to control the initial and final stiffness and also the point at which microslip was initiated within the hysteresis model.

Wen (1989) defined hysteresis as being dependent on the hereditary nature of the responses. Inelasticity is what caused hysteresis and nonlinearity in the model presented. If the loading on the hysteretic model was random, the response would be random and harmonic loading should give cyclic response. The high frequency component of the excitation caused drift that could not be preserved by linearization. Wen (1976) developed Bouc's model to analyse the response of hysteretic systems to random vibrations. Described as an extension of the Markov-vector formulation and Galerkin solution procedure the hysteresis model was influenced by four parameters n , α , β and A . Their influence was outlined by a selection of example hysteresis loops and when applied numerical results compared favourably with Monte Carlo solutions to the problem. Wen (1980) further refined the restoring force model to eliminate the need to differentiate between positive and negative powers (values of n). A prime benefit of the model was that it could cope with anything from elasto-plastic systems to bilinear systems. Equations of motion and restoring force were evaluated using a step-by-step predictor-corrector method and in response to white noise. It was also noted that an elasto-plastic system will dissipate far more energy at low excitation levels than a bilinear system. The author also stated that the method can readily be applied to degrading systems and multi-degree-of-freedom systems.

2.9) Investigations Utilising the Finite Element Method

One of the most powerful tools for performing structural analysis is the finite element method. The analysis of bolted joints here relies heavily on the method to illustrate and verify the behaviour found at the contact interface. However the situation of 3-Dimensional contact is demanding both conceptually and practically. As a result many researchers have applied themselves with good effect in overcoming some of the difficulty associated with performing multi-dimensional contact simulations.

Beer (1985) derived a contact element of zero thickness that could be used in interactions between shell and solid elements in 2 or 3 dimensions. The joints under investigation were assumed to be in initial contact and if the shear strength of the joint was exceeded then irreversible slip took place. Although devised with the main application in rock joints and fractures the element is flexible and probably adaptable to different requirements. Ju and Rowlands (1999) also created contact elements. These were standard elements with a modified formulation for use in situations where there is contact. The advantage of the method was that it could be implemented into current nonlinear finite element codes, in static or dynamic situations with large sliding, and many target surfaces associated with a single contact node. Of prime benefit to computational cost is the maintenance of a symmetric stiffness matrix even in sliding contact.

Gautham, Sarma and Ganesan (1991) produced a 2-Dimensional finite element contact algorithm. The paper provides a nice introduction to the issues of 2-Dimensional open, sticking and sliding contacts and similarities with some of the master/slave contact algorithm used by ABAQUS can be seen. A development of previous contact methods in finite elements, by the implementation of 8-noded quadrilateral elements, was presented by Pascoe and Mottershead (1988). The dynamic coefficient of friction was included explicitly allowing use of a direct, non-iterative solution. Shape functions of the elements were used in determining the forces in the contact zone, with a master/slave type algorithm determining the displacement constraints. Large displacement and heavily

deformed geometry was accounted for by modification of normal and tangent planes accordingly on the elements. The method is analogous to that used in ABAQUS and showed good agreement with theoretical predictions when the mesh refinement was adequate. Mottershead, Pascoe and English (1992) provided a further development of contact algorithms where node on node contact was not required in the finite element analysis. An explicit-like algorithm utilised shape functions to determine the formation of contact and displacement constraints. These constraints could be applied using either the penalty method or Lagrange multipliers. Contact forces in the model were determined by subtracting the vector of applied external loads from the vector of element stress resultants.

2.10) Representation of Structural Joints

Although popular, finite elements are not the only way in which bolted joints have been modelled. Analytical and scale replicas have been used to study joints in conjunction with and independently of the finite element method. Zadoks and Kokatam (2001) investigated the axial stiffness of a bolt. Both quasi-static and dynamic analyses were used to establish the amount of stretch in the joint. Coulomb friction models were used with master/slave algorithms provided by a commercial finite element package. The computational costs were found a significant problem, as were impulses when establishing contact conditions. The deformations of the bolt head were also found to be significant in affecting the calculations of the bolt's axial stiffness. Hwang and Stallings (1994) were more concerned with the flange than bolt behaviour. However of note was an alternative method of applying a preload in a finite element simulation. The bolts were represented by rod elements and the preload simulated by using thermal strains in the model on the rod elements.

Little (1967) outlined pressure cone calculations that represent the contact pressure at a joint interface. As well as being required to generate the contact that transmits moments and forces, Little notes that a preload is need to reduce the effects of cyclic loading. The slope of the clamping pressure cone is said to depend on the geometry

of the joint and tables are included to estimate the pressure distribution dependent on the ratio of the diameters and angle of the pressure cone. A pressure cone of angle 15° - 20° was suggested for preliminary design requirements. Osgood (1979) implicitly suggested a half cone angle of 30° , however, in a study of the design criteria of bolted joints.

2.11) Studies of Bolted Joints

Lee, Ko and Lee (2000) also used a detailed mesh in their simulation involving bolted joints. This approach gave a good indication of issues such as aspect ratios involved in meshing stage of analysis. The aim of the analysis was to show how a condensed stiffness matrix can be used to generate results that are in good agreement with the intuition that modal frequencies are lowered and mode shapes slightly altered. Experimental testing backed this up, but no indication of the savings obtained by the method were given. Beards (1983) gave structural specifications as the main reason for the lack of added damping and suggested joint damping may be the best solution to this problem. The clamping force was given as the main device for controlling the amount of damping in a structure. Relative motion is related to contact pressure and broken down into three categories. These categories were macroslip, microslip and embedding of surfaces and plastic deformation of asperities. The important relationship of energy dissipation being the product of shear force and amount of slip is stated. Also put forward was that the maximum dissipation occurred between rigid and pinned conditions. However it is also stated that the energy dissipated was not sensitive over a wide range of values near the optimum level. Again the bias is towards practical solutions to joint problems. A further reported finding was that the optimum-slip magnitudes were 50% of the undamped displacement amplitude.

Brown and Hickson (1953) used Fosterite to perform a photo-elastic stress analysis on a bolt thread when stressed in a standard nut. Their finding was that the maximum stress at the root of the thread occurred at the point where the stud first joins the nut and another "tension nut" is shown to reduce the stress concentrations in the thread. Maruyama and Yoshimoto (1975) combined analytical, finite element and

experimental work to calculate the spring constant of bolted parts in a bolted assembly. They found, using finite element analysis and point matching method, that the pressure distribution between a bolt head and interface with the clamped component is not uniform yet did not greatly influence the spring constant of the clamped component.

Bradley, Lardner and Mikic (1971) used thermal contact resistance to find consistent trends in contact stiffness according to the geometry of the components in the joint. The cone angle was given and the pressure distribution was integrated over the area in contact and matched with the applied load. The largest discrepancy between finite element results and practical results was found to be around the bolt-hole. Ito, Toyoda and Nagata (1979) considered the interface pressure distribution of a single bolt in a flange assembly. The effect of bolt load, flange material, joint thickness and surface finish all had an impact on the results. The final results were established as plots of pressure against radial distance or pressure-cone half-angle. Ultrasonics were used to qualitatively obtain the pressure distribution before matching it against the bolt loading to establish a quantitative contact pressure distribution (a process similar to that used by Bradley et al.). Differing combinations of factors lead to a pressure-cone half-angle that ranged from about 25°-75°.

Another popular application of friction joints is in the damping of turbine blade vibrations. Menq and Griffin (1985) devised a computationally efficient algorithm that used the response of a beam without friction damping (found using finite elements) to generate receptances to be used with the nonlinear friction force. The nonlinear part used only a few degrees of freedom associated with the friction interface. This could be implemented as an efficient post processor module into commercial finite element software.

Of the work that has been reviewed thus far the vast majority of the cases investigated involved tangential or lateral vibrations. There exists a body of work concerned with the axial vibrations of joints. Rashquinha and Hess (1997) used a lumped parameter model of a fastener and the Karnopp model of friction to describe sticking and

sliding behaviour of the nonlinear joint model. The fifth-order Runge-Kutta procedure was used to solve in state space the nonlinear equations of motion of the bolted joint attached to a cantilever beam. The clamping force was considered after 50 cycles of oscillation of the beam and the optimum positioning of the fastener was considered with this information as a guide. Hess and Sudhirkashyap (1997) used both experimentation and the same lumped parameter model to investigate loosening. Simulations showed that for certain parameters in the model such as contact stiffness, clamped mass and input amplitude there existed regions where no change, loosening and tightening occurred. Increasing the coefficient of friction diminished the amount of loosening that occurred. Basava and Hess (1998) considered the impact of bolt preload and the transient response. Momentary separation between components was identified as the root cause of loosening and tightening. Simulations did however predict that after a certain amount of time the clamping force settled down to a steady value. Esmailzadeh, Chorashi and Ohadi (1996) performed impulse axial loading on a bolted joint's flange. Again an optimal value of preload was found that minimised the bolts deformation and stress in response to the loading. A compromise was needed between the level of preload that optimised the dynamic response of the joint against the level of preload required to maintain the integrity of the joint in a pressure vessel.

Iourtchenko and Dimentberg (2002) considered the response of a bolt to random excitation. A principal benefit of this method of random excitation identification was that it could be carried out in service rather than on a test specimen. From a white noise input viscous damping, Coulomb damping and a combination of the two were identified to within 10 percent of their known values. The data was collected in its entirety over one experimental sweep of duration 6500 times the period of the system's natural period. El-Zhary (1986) analysed a bolt's axial behaviour when subjected to harmonic excitation. Damping caused by the interface was described by a viscous dashpot in what was a very simple two degree-of-freedom model. The joint parameters were investigated for both minimum motion transmissibility and minimum force transmissibility. The author also described how the joint contacts need to remain closed which was good for optimum energy dissipation. Esteban and Rogers (2000) were interested in the high frequency

response of a lap joint containing two bolts. A wave propagation model was derived due to the high frequency response and experimental tests were conducted from the first mode to the 23rd mode. Looser bolts were found to dissipate more energy. The non-destructive testing used was an advantage, but high frequency results could be of limited practical benefit as the lower frequency response will often dominate in practical situations.

Beards (1992) reviewed work on bolted joints' dissipative properties and quotes a value of up to 90% of total damping provided by joints in built up structures. The significance of joint damping due to the high stresses and low inherent damping of structures, coupled with low noise and vibration levels that are desired, was established. 3-Dimensional friction damping is still seen as a subject area that needs further development and investigation. The negative effects of relative movement at joint interfaces such as wear were given prominence. Frequency response data from structures containing joints was seen as a good way to extract parameters. Beards and Neroutsopoulos (1980) investigated the impact of electro discharge machining of joints. The main advantage of the process was that wear on the joint was reduced and simultaneously the damping capacity of the joint could be doubled. The only negative property of the electro discharge machined joint was a loss of dynamic stiffness over the first 1000-10000 cycles. The damping capacity of the treated joints was consistently better than those of a ground joint despite their inferior surface finish as debris in the untreated joints was highlighted as a significant factor in reducing the damping by acting as a solid lubricant.

A numerical and experimental approach was taken by Lee and Feng (2003) to consider the response of a beam with variable clamping force at the joint. As is common the beam could represent anything from a clamped "rigid" joint to a pin joint. A combination of simple 2-beam finite element model with boundary conditions defined according to sticking and slipping was used and a fourth-order Runge-Kutta method was used to integrate the response in the time domain. Sliding was achieved when a critical moment of the joint was surpassed. For fully clamped and pin joints the resonant

frequencies agree with the analytical approach. Slipping was shown to occur at resonant frequencies and over a wider range of frequencies as the input force was increased. When external forcing dominated, the frequency response approximated a linear pinned system. Experimental confirmation of the fact that the joint force can be tuned to give a minimum output at resonance was also provided although the minimising clamping force may have been too low to be practicable.

Gaul and Nitsche (2000) took the approach of using an active joint to increase the amount of damping in a two-beam system. The active joint was created with a view for use in large space structures. It used the Canudas de Wit friction law as the Stribeck effect was visible in the simulation results. The amount of stored elastic energy returned to the system was the quantity that was minimised. The active joint, controllable by feedback, maintained the frictional energy dissipation by preventing the joint from locking up in certain situations. Gaul and Nitsche (2001) developed the work further with an active joint at the interface between two beams. The aim of the joint was to maximise the instantaneous energy dissipation rate and minimise the amount of stored frictional energy in the system. Assumed modes were used to derive the equations of motion of the beam system. A closed loop system including the active joint was then created and implemented on an experimental rig with initial conditions of a fixed displacement. The free response with an active joint coupling the beams exhibited vastly improved energy dissipation over a passive joint with a constant normal force.

2.12) Joint Parameter Identification

A wide body of research exists that seeks to identify the parameters that define a joint's dynamic behaviour either within a structure or in isolation. Ren and Beards (1995) used substructure synthesis to predict the behaviour of a system containing a linear joint. Several different weighting procedures were outlined and used to facilitate the extraction of stiffness and damping parameters from data taken at a variety of positions and frequencies. Criteria and formulae were also given for eliminating those frequency response measurements that were likely to cause erroneous results. The results were

validated when using an existing substructure synthesis method and comparing the frequency response function with and without the inclusion of the joint's influence. Ren & Beards (1998) also isolated the problem of stiff joints within a system. Stiff joints in themselves proved hard to identify, but also made the identification of other joints difficult. Weakly nonlinear joints were viewed as essentially linear. Effective joints were those deemed relatively flexible and other joints were assumed to be rigid in the methodology used. A combination of methods was utilised where the first step was to couple substructures through nominally rigid joints, and the second was to use an identification procedure to establish the properties of the remaining “effective” joints. It was found that results were significantly improved when stiff joints were assumed rigid, coupled and not used in the identification algorithm. The authors also highlighted the problem of using frequency response functions taken at substructure natural frequencies as providing a large source of error.

Ren, Lim & Lim (1998). Claimed to be first paper to attempt to give a general-purpose technique for the identification of nonlinear joint properties. The nonlinear structure was broken down into linear components acted upon by a nonlinear force from the nonlinear component. A multi-harmonic balance method was used to represent the periodic response of the structure. Harmonics were balanced at multiple frequencies and weighted accordingly. Two significant problems that caused ill-conditioned receptance matrices were given as resonance and frequencies that excite other harmonics of the system. The Frequencies used in the identification did not excite other resonant frequencies up to their eighth harmonics. A comparison was made between actual joint force and identified joint force (from measurements away from joint) for an impact joint. Good qualitative agreement was found. All the parameters were identified using Newton-Raphson method and minimisation of a cost function for a joint with an amplitude dependent friction force. Good quality hysteresis loops were obtained where clear regions of stick and slip were visible.

Tsai and Chou (1988) also used substructure synthesis to find the stiffness and damping parameters for two types of bolted joint and were able to predict the behaviour

of other assembled structures if the joint condition was known. The joint properties were found to be frequency dependent. With this method there was no need to generate mathematical models of the stiffness and damping matrices if prototypes of the structure existed. The only measurements needed were the receptances or inertances of the substructures. Two joint types were identified including a lap joint. The identified lap joint properties, when inserted into the synthesis equations, gave a superior match with the experimentally obtained results. The main drawback of the method was that a prototype of the substructures, and knowledge of the joint properties obtained from an equivalent joint in another structure, are required for it to work effectively.

The popular concept of optimum preload is approached by Ito and Masuko (1975) and the damping capacity and change in natural frequency of the beam were also considered. The logarithmic damping decrement was the tool used to establish the damping performance of the joint. It was found that joint properties did not alter much above a certain preload, but as it was lowered significant changes in the natural frequency were observed. The amplitude of vibration was also shown to have an effect on the damping capacity of the joint. Tang and Dowell (1986) experimented on a pinned joint attached to two metal blocks providing friction. Separate dynamic tests were performed to establish the coefficient of friction. The half power method was used to establish an equivalent linear damping ratio. In the time domain the logarithmic decrement was used to also estimate the damping of the first two modes. Good agreement was found between theoretical calculations (stick-slip in the time domain and theoretical mobility for the case of slip) and experimental results for the first mode. Ferri and Heck (1992) also emphasised the use of decay envelopes in systems containing friction. The system used to numerically test the response of these variable normal force joints was the same two-beam, pin-jointed, system used by Gaul and Nitsche (2001). Under the right circumstances the friction damper could represent a viscous damper. Evidence to support this came from the shape of the decay envelopes- exponential for viscous damping and linear decay for predominantly frictional damping.

Crawley and Aubert (1986) used force-state mapping to identify structural elements. Sensitive experimental apparatus created a 3-dimensional space where the axes were velocity, displacement and the input force minus the inertia force of the element. Characteristic profiles were shown for different types of joint properties e.g. cubic and normal stiffness, coulomb friction and impact elements. Experimental joints were also identified in this manner and hysteresis plots could be generated easily from the force-state map. A suggestion of using obtained data in look-up tables during numerical time integration procedures was also made.

Tomlinson and Hibbert (1979) were able to obtain a good match between experimental results and theoretical predictions for a vector response of displacement plotted on an Argand diagram. These plots were used for validation of power dissipation method. A simplified representation of a friction force obtained from the harmonic balance method was used without taking into account problems of sticking and slipping. The power dissipated when resonating a normal mode of the system allowed the size of the friction force together with a separate hysteretic damping constant to be evaluated. For a single degree of freedom system the method worked well, but for more than one degree of freedom a prior knowledge of the location of the frictional element was required. Quasi-static measurements were also used to determine the friction force in the experimental set up. Gaul (1983) analysed a joint's main behavioural characteristics through wave transmission. The joint model was formulated through equivalent linearisation thus maintaining nonlinear features but simplifying the governing nonlinear differential equations. In this paper gas pumping was introduced as a potential loss contributor. Experimental tests showed a variable friction coefficient that was approximated by Coulomb model. Maintaining the area of the hysteresis loops through energy balancing allowed identification of the linearising parameters.

Wang and Sas (1990) identified the stiffness and damping parameters in a multi-degree-of-freedom joint from the modal parameters. Rayleigh's quotient being stationary in the neighbourhood of an Eigenvector allowed an iterative procedure to be used to extract parameters even when only an approximate mode shape could be given. An

example of simple few-degree-of-freedom model is given where stiffnesses and damping capacities were established for each degree of freedom. For more complicated situations finite element analysis was used to generate the approximate Eigenvectors used in the iterative procedures and stiffness and damping was found between a pair of degrees of freedom. Planar loading was assumed meaning only four terms needed to be established, including cross stiffness terms proportional to the elastic modulus. The lap joint could also be identified whilst still attached to its constituent beams.

Lin and Chen (2002) identified contact stiffness by obtaining and minimising the least squares difference between Eigen modes obtained experimentally (using samples of the interface) and finite elements. An optimisation algorithm was then used to obtain the 6 degree-of-freedom contact stiffness coefficients that were applied to a full-scale finite element model. Fairly good agreement was found in mode shapes and modal frequencies between finite element prediction and contact behaviour of full-scale item. The major drawback of the method is the requirement that an accurate sample of the contact interface needed to be produced in isolation of the full model. Meyer, Weiland and Link (2001) used a combination of harmonic balance (to linearise weak nonlinearities), dynamic condensation and least squares fitting of frequency domain response to identify nonlinear joint properties. Coulomb friction had not been established, only powers of damping and stiffness. The experimental model did not show the level of fidelity achieved when validating using a numerical model, but did show a much-improved agreement once the parameters had been updated.

Ma, Bergman and Vakakis (2001) identified a joint through both modal analysis and measurement of the joint force conducted using non-invasive laser vibrometry. Experimentally wandering, or drift, of hysteresis loops was observed and associated with microimpacts of the joints. The algorithm used was based on a comparison made with an equivalent linear beam where the first four modes only were used in the identification. Numerical simulations where linear and nonlinear springs represented the joint were used to validate this methodology. Experimentally obtained hysteresis loops at various joint

torques were also obtained to qualitatively understand the behaviour of the joint. Softening behaviour was found in the joint as the preload on the bolt was decreased.

2.13) Frames and Built-Up Structures

The damping and dissipation that occurs in bolted joints has led to much research being undertaken into their impact on multi-jointed structures and as a controllable damping input. Ungar (1972) observed that random vibration response was dominated by behaviour at near resonant frequencies and that prediction of damping in practical structures was difficult and required complicated measurement. The author also noted built-up structures were found to display much higher damping than monolithic structures and the cause of this came from the joints. In low frequency response inertial effects could be ignored, and maintained wavelengths that were greater than all the relevant structural dimensions. Air pumping in joints was also raised as a dissipative mechanism. Microslip was illustrated as a consequence of both pressure variation and/or material elasticity. The relationship of energy dissipated with respect to μp was also said to be different in rigid, gross sliding mechanisms to those with microslip. The author also observed that closely spaced bolts behave like uniform pressure loading on the joint interface.

Beards and Woowat (1985) presented an investigation of a rectangular frame, rigidly clamped at the bottom, with two friction joints at the top. The clamping force was varied in the frictional joints from an effectively pinned condition to a rigid condition. The frame was then resonated in its second mode for a variety of sizes of clamping force. The magnitude of the frame's response was shown to be dependent on both the size of the input force and the clamping force. Minimum frame response was achieved at a clamping force that was too low to be practical although significant reductions of response were achieved with clamping forces that were not much below the rigidly clamped condition. An accompanying theoretical study was also carried out with each element of the theoretical model being combined using dynamic stiffness coupling. Good

agreement between the theoretical and practical methods was found with only the pinned and clamped models considered.

Rotational slip was used as a damping mechanism whilst having a joint that retains its translational stiffness in the work of Beards and Williams (1977). A rigid, welded, rectangular frame was used with a bolted beam across one of its diagonals. The resonance frequency increased from free to locked condition of the bolted joint, but the mobility of this frequency decreased to a minimum and then increased again towards the locked natural frequency. However, the optimum joint force, measured by static friction capacity, was not consistent for all modes. Ren and Beards (1994c) investigated the use of pseudo joints in a structure that also had to carry static loads. At tuned clamping forces the peaks associated with the lower resonant frequencies could be eliminated. The pseudo joint's optimum clamping force was quite insensitive even with small modifications to the structure. A consistent relationship between the clamping force and excitation force was found allowing scaling of tests conducted with low levels of excitation to full-scale input forces.

Moon and Li (1990) investigated whether a chaotic response was generated from a harmonic input to a 3-Dimensional, pin-jointed space structure. From hammer tests nonlinear joints were found to contribute to distorted mode shapes and natural frequencies below those predicted using finite element analysis. If the input was reduced to a very low level the broadband response disappeared and clear peaks would appear for different modal frequencies and their combinations. The broadband response was contributed by small gaps in the pin joints and a tensioned cable along the length of the structure closed these gaps. Numerical studies without friction effects bore up the concept that it was gaps in the joints that cause the "chaotic" response. Bohlen and Gaul (1987) improved prediction of frame response through joint damping identification from an isolated section. Joint models were coupled with individual components making up two frame structures. Their model could also represent flexure moments as gap elements were included. However, both gaps and damping and the damping attributed to them were claimed to be negligible in the microslip range. Microslip was however found to reduce

natural frequencies. Hysteresis loops had curved corners at larger vibration amplitudes that were idealised with a single Coulomb element. The relative amplification of the input at resonant frequencies was also reduced as the magnitude of the input increased. This could be attributed to the increased amount of sliding, and therefore damping, caused by larger inputs to the joint.

Gaul and Lenz (1997) investigated a multi-bayed space structure similar to that investigated by Moon and Li (1990). A limited, 2-Dimensional, finite element model of the bolt region with a uniform preload instead of a bolt preload was used to investigate hysteresis generated by microslip. The Valanis model of plasticity was used to simulate a specific hysteresis loop from one of the longitudinal experimental tests. Parameters from this particular loop were extrapolated for other input forces with good success. The Valanis model was also incorporated as a joint substructure into a 2-Dimensional model of a space frame with very low material damping. A significant increase in the rate at which the response decayed was shown. Hysteresis loops over time show that initially macroslip was present in the space frame joints and then as energy is dissipated the response became smaller with only microslip being exhibited. Gaul and Lenz (1998) continued their investigation using a hybrid multibody system to simulate the assembled structure. Savings in computation time of two orders of magnitude were achieved over an equivalent finite element model. Placing a piezoelectric disc underneath the nut to modify the contact pressure at the joint interface created an active joint to control the structural vibrations. When implemented in a closed control loop the active joint significantly reduced the amplitude of the space structure's response to both free and forced vibrations. The control strategy was devised by extracting parameters for the Valanis model from a bolted joint excited in isolation.

Chapter 3. Finite Element Analysis of an Isolated Joint

3.1) Finite Element Method Summary

For the last half century the finite element method has been developed into a widely used and powerful numerical tool. From simple stress-strain problems the finite element method can now be utilised in complicated nonlinear contact analyses. Contact analysis can be developed to 3-Dimensional systems using specially designed algorithms included in most commercial finite element packages. In the context of bolted joints, the finite element method allows the detailed behaviour of the contact interface between two components to be analysed to reveal details about local contact pressure and displacements that cannot be measured directly.

In the analysis carried out, a commercial finite element package was used to analyse a bolted joint in isolation. A commercial package was utilised as it allowed full use of pre-programmed 3-Dimensional contact algorithms. By utilising existing codes when performing the contact analysis a significant saving in time was achieved when compared with the effort that would have been required to perform the coding in house.

The ABAQUS suite of programs was used to perform the finite element analysis as it allowed both static simulations using ABAQUS/Standard and dynamic contact problems to be analysed using ABAQUS/Explicit. A preload was applied in ABAQUS/Standard in the form of a bolt preload or a uniform pressure applied to one of the components used in the joint. These variable preloads generated different contact pressure distributions at the joint interface. Once the contact pressure had been created between the two surfaces of the joint interface a sinusoidally varying torque was applied to one of the components in ABAQUS/Explicit. The magnitude of the preload, torque and coefficient of friction were all varied to investigate the microslip behaviour that existed before the joint failed.

Hysteresis loops created for the dynamic part of the simulation revealed that microslip existed to a large degree in the joint with an applied bolt preload. When enough torque was applied the contact stiffness of the joint gradually decreased as greater regions made the transition from sticking contact to sliding contact. At velocity reversals the stiffness of the joint immediately made the transition from its currently reduced value back to the contact stiffness associated with a totally stuck contact interface. This process was established as a steady state cycle that occurred at the same frequency as the applied torque.

A uniformly distributed load applied to one of the components that made up the joint produced a pressure distribution at the contact interface that was largely constant. This was in sharp contrast to the decline from a maximum value around the bolt-hole to near zero pressure at the edge of the interface when a bolt preload was applied. The result of the nearly uniform contact pressure was a bilinear hysteresis loop until the joint failed and macroslip took place.

The most basic Coulomb friction law was applied to the contact algorithm that was used. However, a proportional relationship did not exist between the magnitude of the contact pressure, and hence normal force, and the coefficient of friction. The relationship largely held at friction coefficients smaller than one. Above this value increasing the bolt preload did not produce a comparable result to an equivalent proportional increase of the friction coefficient.

Despite the finite element model being constructed in a way to maintain as much rotational symmetry as possible, the hysteresis loops were offset towards a positive angular displacement. This offset existed to a greater or lesser degree in all of the results and is believed to be a result of an anomaly that occurred in the transition from the implicit static solution to the dynamic explicit solver. The contact algorithm was slightly different between the two ABAQUS programs. When the statically deformed mesh was transferred to ABAQUS/Explicit, a small loss in preload and contact pressure took place

at the interface. The resulting initial contact stiffness was slightly lower than that found in all subsequent loading cycles.

Even though the offset of the first cycle was present in all of the simulations, it did not stop the joint hysteresis from revealing a great deal about the microslip that took place. Microslip existed until the whole contact interface was in sliding contact at which point the output from the finite element simulation became very noisy and the hysteresis was not interpretable. This behaviour was considered to be representative of macroslip. At this transition unconstrained motion was prevented only by the boundary conditions applied to the clamped components and the connections they made with the bolt-head and nut.

3.2) Benefits of a Commercial Finite Element Package

To apply the finite element method to the investigation of bolted joints two options were available. The first was to code an "in-house" program to implement the finite element method and particularly the algorithms required for 3-Dimensional, nonlinear, frictional contact. The other option available was to utilise a commercial finite element package with all the algorithms required for the analysis pre-written and implemented through an input deck.

One obvious advantage of self-coded programs is that the significant cost of software licensing would be circumvented. While an in-house code allowed complete accessibility to all the methods used to implement the finite element algorithms, there was not a program immediately available that provided all the functionality required. A substantial amount of coding would therefore have been necessary, and such an undertaking would undoubtedly have demanded the majority of time invested in this research. The ABAQUS finite element suite of programs facilitated all the different types of analysis that were needed with no additional coding requirements. ABAQUS is already used through a wide range of research and commercial institutions. It is a very versatile product and is readily capable of modelling complicated structural problems with contact in 3-dimensions. The ABAQUS suite has two different components, a "Standard" module that is used for linear and nonlinear static simulations and an "Explicit" component used in linear and nonlinear dynamic simulations.

The use of a commercial finite element package has several advantages over an in-house program. The significant amount of time that would have been required to code an independent program was saved. These man-hours could then be devoted to other necessary areas of research like alternative modelling techniques and experimental validation. Much time can also be saved in the pre-processing and post-processing stages of applying the finite element method when using a commercial software package. A user interface has been predesigned to facilitate the rapid production of models, and ease the interpretation of results. A commercial finite element package also has the benefit of

being tested rigorously as the program evolved through different version releases. ABAQUS version 6.2 and version 6.3 were those licensed to the University of Liverpool and used most comprehensively for the following work. Some primary tests were performed using ABAQUS version 5.8 although these were superseded due to the better compatibility of ABAQUS/Standard and ABAQUS/Explicit modules in later releases.

One of the foremost benefits of a commercial software package is the consistency between all users that is offered. A high level of consistency is achievable between different research institutions either in collaboration or further development of analysis techniques. Other researchers, to ensure a consistency of methodology, could readily replicate and utilise results. Benchmark tests could also be created to ensure that other users of the software package in a different situation are following the same procedures. This would be of particular benefit when considering the relatively complex situation of dynamic, contact simulations.

In-house codes would not offer the same level of transparency to other researchers that a commercial package could. The continuation of the research with different personnel is also eased if experience has already been gained using the same package, substantial documentation is available and a detailed language/program does not have to be learned. ABAQUS provides a thorough background to most of the complicated algorithms (particularly involving contact) that are already in place in the software suite. This ensures there is no shortfall of understanding and transparency despite the software manufacturer's desire to protect its copyright.

In addition to the documentation supplied with the ABAQUS license there was also a vast background of expertise that could be drawn upon in the event of problems. Primarily this was achieved through the manufacturer's technical support facility. There also exist networks of users and contacts with previous experience of the software package along with user groups that exist via the World Wide Web. All of these sources could be approached for advice and support that would not be available when using self-coded finite element algorithms.

3.3) The Finite Element Model

While the blocks that represented an isolated bolted joint were regular in shape, the inclusion of the bolt-hole meant that it would have been extremely ill advised to mesh the blocks with a regular rectangular grid. Another stipulation of the mesh was that when modelling a contact simulation those parts that were in contact should ideally have a matching mesh. Modelling this more sophisticated mesh profile without the use of a graphical pre-processor would have been difficult indeed. The mesh for the finite element model was therefore constructed using Patran as a pre-processor. This meant that the square (in the x-y plane) profile of the two blocks could to be combined with a circular mesh around the bolt head, nut and common hole through the blocks.

The sizes of the components to be modelled by the isolated lap joint were also considered carefully. A joint that was as representative as possible of commonly found structural joints was thought to be best typified by a metric standard bolt of "reasonable size". "Reasonable size" was decided upon by compromising on a bolt that was not so large that it would be used with components too unwieldy for laboratory analysis, yet large enough that it would provide a preload big enough for use in structures large enough for normal measurement techniques. An M10 steel bolt of standard strength, head and nut dimensions was viewed as satisfying these criteria.

It was necessary to produce a model that illustrated a wide range of phenomena, from microslip to macroslip that could be expected to be observed. Most of these phenomena, it was predicted, would stem from the contact pressure interface between the two bolted components. In a relatively tightly clamped bolt some of the contact pressure would be quite large. For microslip to take place there had to be a wide, progressive range of contact pressure at the interface. The pressure cone method (Shigley (1986), Little(1967), Osgood(1979)) outlined in Section 1.2 showed that the pressure distribution for a given preload would be most "spread" if the clamped components were relatively deep. Three conditions now existed to help size the bolted blocks. Firstly the bolt size was known as M10, secondly the contacting faces had to be of a large enough area that

the contact pressure was negligible at radial distance from the bolt axis actually on the interface. The third condition was that the depth of the blocks had to be such that the sticking conditions of a substantial proportion of the interface, caused by the interaction of contact pressure and friction coefficient, could be overcome gradually. This meant that the contact pressure had to be spread over a large proportion of the contact interface. The magnitude of the contact pressure had to be of a small enough magnitude that an applied torque would overcome sticking conditions without significantly deforming the blocks. All of these conditions were satisfied in the pressure cone method (and subsequently proven in the finite element analyses) by blocks that were 80mm square and 30mm deep. The relative sizes of the final finite element model can be seen in Figure 3.3.1.

It should be noted that it is unlikely a beam with such a cross section would be joined to another beam of the same cross section by a single M10 bolt. However a balance had to be found between practically realistic dimensions and dimensions that would best illustrate a broad range of physical phenomena in the joint.

Another area of the model that had to be carefully judged was the overall complexity and detail. It was required that the joint had to represent physically as closely as possible an isolated joint of the most widespread and basic kind. A lap joint was chosen for this purpose as the joint was formed by frictional contact on only one common interface between two flat components. Naturally there would, in a normal joint configuration, be contacting surfaces between washers and the two clamped components and washers and the bolt head and nut. For a fully representative model the washers and associated contact interfaces would have to have been modelled as well. However the ABAQUS solution algorithm involving contact puts very large demands on computing capacity. Again a compromise had to be found between the fidelity of the model and efficiency of the solution. Intuitively it was assumed that very little microslip or macroslip would occur in regions where the contact pressure was very high. The maximum contact pressures could be assumed to exist underneath the bolt head and nut. This, coupled with the relatively small depth of the washers, implied that little change in contact pressure distribution at the contact interface would occur if the washers were not

included in the model. A small area of relatively large contact pressure was assumed to not be responsible for a significant amount of frictional energy dissipation.




Figure 3.3.1 *The components that constituted the finite element model of an isolated lap joint.*

By eliminating the washers from the finite element model, the number of contact surfaces included was reduced from five to three. The benefits were twofold: A drastic improvement in computation time for each simulation was achieved. Also the data regarding frictional energy dissipation was more accessible as it was associated with a single, clearly identifiable, contact interface between the two blocks. To prevent the relative slip between the clamped components and the contacting surfaces of the bolt head or nut the two pairs of contacting surfaces had their motion tied together. The tied constraint meant that any node on the bolt head or nut's surface shared the same motion as the closest point to it on the clamped component's surface. As the meshes matched exactly, this caused nodes on the underside of the bolt head or nut to share the same motion as the node that they were initially in contact with on the upper surface of the top clamped component or the bottom surface of the lower clamped component.

Another necessary simplification of the model and the resulting contact conditions was associated with the bolt thread. It was decided to model the bolt in a clearance hole rather than through a mating thread in the lower component. This is an acceptable configuration of a lap joint. Consequently a bolt-nut configuration (312 elements) could be used which maintained symmetry in the loading conditions and removed the need for

an extensive and complex contact interface between the bolt and lower block. The bolt shaft was modelled as a cylinder of constant cross section, the diameter of which was equal to the bolt's nominal diameter (10mm). The bolt-hole through which the bolt was inserted was also 10mm in diameter. As a result of the connectivity of the elements and the surface definitions in the model no interaction took place between the surface of the bolt and the hole through which it ran. The advantages of this were twofold. Firstly a costly surface interaction was eliminated from the model. Any surface interaction and subsequent energy dissipation from this part of the model would cloud the behaviour of the interface between the two blocks. Secondly an idealistic snug fit between the bolt and shaft meant that the mesh definition, stresses and deformations around the bolt head and nut were less of an uncertain influence on the contact pressure distribution at the joint interface.

The bolt head and nut were designed so that they were axi-symmetric instead of the usual hexagonal profile. This simplification made it much easier to match the meshes of the bolt head and nut to the mesh on the two block components. From the point of view of the physical model, hexagonal profiles would present a problem in that their axial orientation would be indeterminate relative to the blocks and could vary by as much as 60° in any assembly. Also underneath a manufactured nut or bolt head, raised slightly in profile, was a circular ring which actually forms the contact with the clamped components. Even if the whole hexagonal profile did form the contact with the two blocks, it was assumed that the pressure distribution at the interface between the two blocks would still be essentially axi-symmetric. As a consequence the bolt head and nut were modelled as cylinders of diameter equivalent to the width across the flats of standard M10 bolts and nuts. The depth of the bolt head and nut also corresponded to the same standard dimensions (Shigley (1986)).

The reason that the bottom block (612 elements) was marginally larger than the top block (504 elements) was that the contact algorithm in ABAQUS did not converge on a solution during preloading when the surfaces matched exactly and ran to the edges of the components in finite sliding simulations. The skirt around the bottom block prevented

these problems in the case of large relative displacements between the two blocks. However in the contact algorithm used to generate the results used here, the skirt became irrelevant as sliding planes were extended automatically. This became possible due to the "small sliding" algorithm that could be implemented when it was discovered that in the case of microslip the relative displacements at the contact interface were significantly less than those associated with a single element length.

As with all finite element simulations it was necessary to apply boundary conditions that prevented rigid body motion of the model. While relative motion could occur between the top block and the bottom block the model, consisting of the three components, had to be restrained en masse. To make the results as simple to identify as possible, symmetrical boundary conditions about the bolt axis were applied. If the two clamped components were visualised as the local joint at the tips of two beams, the sensible boundary conditions to apply would be to the face of one of the blocks where it joins an imaginary beam component. Loading could then be applied to the other clamped block on the corresponding face where it too joins an imaginary beam. The asymmetry, about the bolt axis, of such a situation would however make interpretation of the results difficult.

The model was therefore made as symmetric as possible with regard to the bolt's axis by constraining the whole of the bottom face of the lower block. All of the nodes were fixed in position using the Encasté boundary conditions-which meant that that the bottom block could not undergo any rigid body motion. Through the tied conditions of the bolt-nut component to the contacting faces of both the bottom and top block, rigid body motions of the two remaining components were correspondingly constrained. Further constraints on the motion of the top block were then enforced by the frictional contact generated between the upper and lower blocks, caused by the preload applied to the bolt-nut component.

All of the constrained surfaces are shown in the exploded view of Figure 3.3.2. From the top the interacting surface pairs are: The underside of the bolt head and top of

the upper component (tied), the bottom of the upper component and top of the lower component, the bottom face of the lower component and the contacting face of the nut (also tied).




Figure 3.3.2 *Exploded view of the surfaces generated on the solid components illustrated in Figure 3.3.1.*

The preload applied to the bolted component corresponded to that attributed to a class 5.8 standard strength bolt as specified in Shigley (1986). A preload of 19kN was applied in the ABAQUS/Standard package to allow a static analysis to be undertaken. In the following section the preloading algorithm is outlined in more detail. There were also some simulations whereby a uniformly distributed load was applied to the top block. This was done to create a uniform contact pressure at the interface between the upper and lower blocks. In these instances the bolt head was untied from the top surface of the upper block to prevent any local interactions around the bolt hole that could have affected the uniformity of the pressure distribution at the contact interface.

The dynamic loading applied to the model was also performed as symmetrically around the bolt axis as possible. To generate the axi-symmetric torque point loads were

applied to the vertices of the upper block once preloading had taken place. As the results gained from the simulations would generally provide qualitative understanding of the contact mechanisms between the two blocks this type of loading, although slightly unrealistic, offered the clearest physical insight.

All dynamic loading tests were performed using ABAQUS/Explicit. The deformed model from the static preloading analysis was exported so that the dynamic loading could be applied. The deformations and contact pressures caused by the particular preload in operation would then determine the response of the blocks to the dynamic torque applied. The torque was applied in a fixed direction, as local deformations were assumed to be small, and using a harmonic variation in magnitude. In all of the simulations an excitation frequency of 125.66rad/s or 20Hz was used. This enabled a controlled response, away from the resonant frequency of blocks to be observed, and enabled a reasonable compromise in the amount of data produced over a given simulation duration. The time increments in the Explicit analysis were predetermined by the model geometry and material properties. A balance therefore had to be struck between computational time, model resonance and resolution of the output data. All the finite element data presented here was extracted at a sampling frequency of 1kHz.

3.4) Implicit and Explicit Contact Algorithms

The contact algorithm used in ABAQUS/Standard was chosen due to its applicability to the modelling of two deformable surfaces. Each surface was created by the joined faces of continuum elements. In particular the algorithm had to effectively reproduce microslip behaviour that occurred as a result of the contact pressure variation across the interface. Precise determination of sticking and sliding conditions for each node at the surface had to be determined, and displacements of a small amplitude relative to the model size also had to be detectable. The contact algorithm offered in ABAQUS/Standard, and the optional modifications that could be applied to it, was able to operate fully in 3-Dimensions. This was critical when considering the level of detail required by the model and also the torsional loading that was applied across the joint. It would have been much more efficient to model tangential loading in either the x or y -direction in 2-Dimensions without representation of the whole contact interface. However, microslip and a contact pressure distribution that could not be assumed axisymmetric meant that a 3-Dimensional, deformable, contact algorithm had to be used.

As has already been reported (Section 3.3) there are three pairs of contacting surfaces within the finite element model of the isolated joint. In ABAQUS/Standard each surface in the pair making up the contact interface had to be distinguished as either a master surface or a slave surface. The bolt-head and nut were defined as slave surfaces and the outer surfaces of the upper and lower blocks, with which these components formed a "contact pair" respectively, given the status of master surface. It did not matter greatly how these two pairs of surfaces were assigned as the tied condition between them made most contact interactions irrelevant. Between the two surfaces of the blocks that made up the central contact interface, the upper block's interface was given slave status, and the lower (slightly larger) block's interface became the master surface. Generally the more flexible component takes on master surface status, and the slightly increased surface area combined with the same depth would suggest that the lower block was more flexible than the upper block. Results during the contact analysis are generated for the slave surface. The skirt used on the lower block purely as a precautionary measure would

have meant that the results would be less representative of the actual interface (defined by the upper surface's dimensions) that were to be measured. Consequently the upper surface was defined as the slave, and the almost identical flexibility of both the clamped components meant that this was acceptable.

ABAQUS/Standard uses a pure master-slave contact algorithm. The main stipulation in this algorithm is that it is only possible for nodes from the master surface to penetrate the slave surface. Constraints are applied within the algorithm to prevent the slave nodes from penetrating the master surface when contact has been established at the interface.

Using a pre-processor facilitates mesh design considerably. Unfortunately, due to round off errors, sometimes nodes that should share the same co-ordinates at the beginning of the analysis are slightly misaligned. To circumvent this problem the nodes at the contact interface were adjusted by moving all of the nodes that were part of both the bottom block and its contacting surface so that any initial overclosures were removed, and any open contacts were closed. This operation, internal to ABAQUS/Standard, was performed before loads and boundary conditions were applied. Each slave surface in the three contact pairs was subjected to this small initial adjustment. Any strains associated with the adjustments of the slave nodes were disregarded as the operation was only to eliminate any numerical discrepancies carried over from pre-processing. It also meant that when the preload was applied to the joint all of the surface nodes were in contact - which saved time during the solution procedure.

As only microslip was being investigated in the joint, the amount of relative displacement between the two blocks was severely limited. The small sliding formulation of contact was therefore used to establish the constraints at the interface. Such a stipulation relied on the fact that the two surfaces would slide no more than an element's length relative to each other. This is reasonable, as such a large amount of relative movement would amount to macroslip and therefore failure of the joint. Each slave node

was therefore defined as acting on the same region of the master surface for the entire duration of the analysis.

The small sliding formulation provided significant computational savings over an equivalent finite sliding model. Each node on the slave surface had a 2-Dimensional tangent plane on the master surface defined for it. The tangent plane was characterised by an anchor point on the master surface and an orientation ascribed to it. This anchor point on the master surface was defined such that its normal vector from the local tangent plane passed through the corresponding slave node. Across the whole of the master surface smoothly varying normals were created before the local tangent planes were established. A limitation of this method was that the local tangent planes offered only a representation of the master surface. This would particularly be the case in heavily distorted meshes or those models that do not start out with a relatively smooth surface. However, contact does not have to be monitored over the extent of a whole master surface for an individual slave node. This provided a drastic saving in computational cost and was perfectly acceptable for the joint interface which exhibits relatively small amount of movement and only minor deformation on application of the preload.

Load transfer was also made more efficient by the small sliding assumption in ABAQUS/Standard. The slave nodes only transferred load to a selection of nodes on the master surface in the immediate locality. The amount of load transferred was determined by the distance of the slave node from the local nodes on the master surface. Over the duration of an analysis the slave node could travel along its local tangent plane. In such an instance the distribution of load from the slave node to the local nodes on the master surface was updated, but the nodes on the master surface associated with each slave node remained constant for the duration. Such a simplification is perfectly adequate for the preloading analysis carried out in ABAQUS/Standard.

The loading normal to the contact interface has been discussed, but the response due to friction was largely seen in tangential motions of the joint. The basic Coulomb friction law specifies these motions as being a function of both the normal loading on the

surface and also the coefficient of friction. In all of the simulations performed in this research, the coefficient of friction (μ) was assumed to be constant and given a value of 0.8. Generally when considering friction in a mathematical context, the limiting value of tangential load required to make the transition from sticking to sliding is defined as:

$$F_l = \mu N \quad (3.4.1)$$

Due to the discretisation process used by the finite element method, and the 3-Dimensional nature of the joint problem, the limiting value is not defined in terms of a limiting load (F_l), but as a limiting shear stress (τ_l) that is a function of pressure (p).

$$\tau_l = \mu p \quad (3.4.2)$$

The value of shear stress that was compared with the limiting value, defined above, was the magnitude of the shear stress in the x and y -directions.

$$\tau_{eq} = \sqrt{\tau_x^2 + \tau_y^2} \quad (3.4.3)$$

If this value, (τ_{eq}), was greater than the limiting value of the shear stress sliding contact would be initiated and the restoring shear stress would be equivalent to τ_l . In all other sticking cases the shear stress would balance that being applied to the contact interface.

To aid computational efficiency and convergence the default friction model allowed some elastic movement of the surfaces relative to one another. Although offering some benefits this penalty friction method did not allow the detection of small amounts of microslip. The small amount of elastic sliding allowed before the limiting shear stress was overcome masked the effects of microslip. Instead a slightly different friction model had to be implemented. Using Lagrange multipliers the sticking and sliding constraints could be enforced precisely. The payoff for the exact solution was a more intensive computational process to achieve convergence. However, the precise results were

required to adequately represent microslip, and the ABAQUS/Standard, static, preloading analysis took a relatively small amount of time to run to completion regardless.

ABAQUS/Standard was used exclusively for the static preloading analysis. The dynamic loading of the joint was performed using ABAQUS/Explicit for reasons of efficiency. If the contact conditions at the interface between the two blocks had to form a convergent solution, as they did in ABAQUS/Standard, the analysis would have taken several orders of magnitude more time than an equivalent simulation using ABAQUS/Explicit. The reason for ABAQUS/Explicit's efficiency in dynamic simulations is that it did not need a convergent solution before attempting the next time step. Each time step was so small that every sample was only a tiny increment from the previous sample. This can be considered to offer a reasonable approximation to the behaviour. Due to this "explicit" solution procedure for dynamic simulations, modifications were made to the ABAQUS/Explicit contact algorithm.

The two main alterations from the implicit solver to the explicit solver was the use of balanced master-slave contact and transition to a kinematic contact algorithm. Balanced contact relies upon the same principals of pure master-slave interaction between surfaces, but applies the master-slave to each surface and vice versa before taking an aggregate of their behaviour.

By combining balanced surface behaviour with a kinematic contact algorithm penetration of the two surfaces was kept to a minimum. The kinematic algorithm worked by initially advancing the model in the simulation with complete disregard to the contact constraints. By doing this it became evident which slave nodes had penetrated into the master surface (an illegal transition within the algorithm). By considering the penetrating node's position, mass and the time increment a force that would resist the penetration was calculated. This force amounted to a value that would have resulted in exact contact of the two surfaces, with no penetration, if it had been applied during the increment. The calculated force was then distributed to the nodes on the master surface along with the mass of any contacting slave nodes-which was combined with the mass associated with

nodes of the master surface. The explicit solver was then capable of forming a corrective acceleration of the master surface nodes. These nodal accelerations, when combined with the amount of penetration and time increment were used to establish corrective accelerations to the slave surface nodes. Once corrective accelerations had been applied to both master and slave nodes a final corrected configuration was achieved where the contact constraints were satisfied for a given increment.

One benefit of the kinematic contact algorithm was that it used the exact Coulomb friction model directly. No penalty sliding was included in the configuration, so microslip could be easily extracted. The frictional constraint worked by establishing the maximum force for the nodal mass, distance slid and time increment to maintain the nodes position on the opposite surface in its configuration. If the critical shear stress (τ_c) was then exceeded its associated load was applied. If the critical/limiting shear stress was not exceeded then the load associated with the actual shear stress generated was applied to both surfaces in opposite directions. Either of the resulting shear stresses were enforced as corrective tangential accelerations to the slave node and the nodes of the contacting master surface facet.

The small sliding algorithm for modelling microslip was also available for use in ABAQUS/Explicit. There were some small changes to the algorithm from that used in ABAQUS/Standard. The most significant alteration was that the anchor point on the master surface was identified as the closest point to the slave node. Previously the *normal* of the anchor point had to pass through the slave node with which it was associated. This allowed calculation of the tangent plane from the average normal of the master surface at the newly defined anchor point. All of the local tangent planes and interacting nodes were calculated at the onset of the analysis. A procedure was used here where the results from the preloading analysis were exported into the explicit harmonic loading analysis. Small discontinuities in the response could therefore occur as a result of the slight changes in contact algorithms used in ABAQUS/Standard and ABAQUS/Explicit.

3.5) Pretension Algorithm

ABAQUS/Standard has the capability to model assembly loads. This capability was utilised to represent the preload that is applied to a bolt to clamp the two joint components together. To define the assembly load a surface was created internally in the bolt across the midpoint of the bolt's shaft. Remote to the rest of the model a controlling node was established and linked to the pretension surface. The controlling node only had a single degree of freedom within the context of the algorithm.

An average normal was calculated facing away from the elements used to create the surface. This normal determined the direction in which the preload or pretension was applied. In the case of the bolt used in the isolated joint, the normal of the pretension section pointed along the bolt axis in the direction from the nut to the bolt-head. The connectivity of the elements was then used to establish which elements lie above and below the pretension section. Those lying below the section, and used to create the surface, were known as the base elements.

To ascribe magnitude to the preload, the remote node associated with the section was given a load. The load acted in the direction of its only degree of freedom and was converted into a self-equilibrating force across the mid-section of the bolt. Equilibrium of the bolt is achieved in the algorithm by adjusting the length of the component at the section until the required force is generated. Nodes of the elements connected to either side of the pre-tension section were moved to generate the change in length that balances the amount of preload required. In the case of the bolted joint, the elements either side of the section were shortened to develop the 19kN (in the control test) preload that was required to clamp the blocks together.

3.6) Finite Element Output Calculations

To understand the mechanism by which the frictional energy was dissipated at the contact interface, hysteresis loops were plotted. The area inside the hysteresis loops revealed the amount of energy dissipated and the gradient and displacement described the overall motion of the joint. ABAQUS/Explicit has an internal function that evaluates the energy dissipated through frictional losses in a model. However, the level of insight of the finite element model could not be replicated in the experimental observations. For this reason a method was used that approximated the behaviour of the joint en masse, but still gave a good indication of the microslip characteristics of the joint.




Figure 3.6.1 Node 1652 on the upper block's contact interface used to collect linear acceleration and displacement data.

To calculate the hysteretic relationship of the joint it was necessary to know both the angle of rotation about the bolt axis and the size of the restoring (hysteretic) torque provided by the interface. The distance travelled by one surface relative to another had to be measured in angular coordinates as the applied load was provided in the form of a torque. A single node (node1652 on Figure 3.6.1) was taken as a reference node. The displacement of this node in the y-direction was assumed negligible when only microslip was taking place. Consequently the angle of rotation of the joint θ was approximated as

the displacement of the node in the x-direction divided by the distance of the node from the bolt's axis. It was assumed that the displacement in the y-direction was negligible due to the largely axially symmetric loading conditions:

$$\theta = \theta_{1652} = \frac{x_{1652}}{r} \quad (3.6.1)$$

In Equation 3.6.1 r was taken as the radial distance (0.04m) from the bolt axis to node 1652. The displacement of node 1652 in all three coordinate directions was a standard output from the ABAQUS/Explicit analysis. Although a simplification, the angular displacement of the interface, was assumed to be constant at all points on the edge of the upper block (the component from which measurements were taken).

What remained was to find the restoring torque that existed at the contact interface to counteract the applied loads. The exact form of this torque could not be assumed as establishing the extent and progression of microslip was the main objective of the work carried out. It was known, however, that the equation of motion of the interface could be represented by:

$$I\ddot{\theta} + f(\theta, \dot{\theta}, t) = T_{ap} \quad (3.6.2)$$

In Equation I represents the moment of inertia of the top block about the bolt axis, T_{ap} represents the applied torque and the restoring torque is represented by f . The angular acceleration of the top block was approximated in the same way as the angular displacement. Part of the output of the ABAQUS/Explicit analysis was the acceleration in the x-direction of node 1652. This allowed the restoring torque of the interface to be approximated as follows with "•" representing differentiation with respect to time:

$$\ddot{\theta} = \ddot{\theta}_{1652} = \frac{\ddot{x}_{1652}}{r} \quad (3.6.3)$$

I was reported by the finite element package as part of the dynamic solution procedure as $1.59744 \times 10^{-3} \text{kg.m}^2$. Therefore:

$$f(\theta, \dot{\theta}, t) = T_{ap} - I\ddot{\theta} \quad (3.6.4)$$

By plotting this quantity against the angular displacement of the upper block the hysteretic relationship of the joint was obtained.

3.7) Verification Tests

When creating a finite element model it is necessary to perform tests that indicate that the actual meshing of the model is not a determining factor in the response to the inputs of the simulation. More particularly it is required to ascertain that the mesh is of a suitable design such that the results of the simulation converge upon a particular, correct solution. Two methods can be used for establishing the convergent solution. Firstly a comparison of the results obtained for different mesh refinements with a known analytical solution can be made. Secondly a solution can be found that is approached as the mesh refinement becomes greater. The model output should then approach this solution, to a predetermined tolerance, through a practical level of mesh refinement. It is almost always preferable to use the least refined mesh that offers suitable convergence, as it is most likely that this mesh design will demand the smallest processing effort. Small computational cost is a good indicator of the efficiency of the method used and also offers a practical benefit in the rate (wall clock) at which a solution is obtained.

As the relationship between contact pressure and the coefficient of friction was fundamental to establishing the slip mechanism at the interface, it was necessary to ascertain that the mesh used gave an acceptable contact pressure distribution when a preload was applied to the bolt-nut component. To perform this check the mesh of the three components in the model was refined in four different configurations. As the contact pressure acted normally to the contact interface the first two tests doubled the density of the mesh in the normal (z) direction.

The distribution of the contact pressure radially about the bolt axis was also significant. Around the bolt-hole itself the distribution of the contact pressure should have been axi-symmetric. Near the corners of the square profiled blocks this need not have held due to stress concentrations and the influence of the geometric discontinuities. The variation of contact pressure with radial distance from the bolt axis should be in agreement with a convergent solution. The contact pressure distribution could also be qualitatively verified with that outlined in Section 1.2. To verify the finite element

simulation model against the pressure cone distribution, and to compare it with an (assumed) convergent solution, two increases in radial density of elements were also applied to the two blocks. Where the mesh overlapped the clamped component the same mesh refinements were applied to the bolt-nut component. One radial refinement consisted of a 50 percent increase in element density, the second refinement doubled the number of elements in the radial direction. Both of these cases were used in conjunction with the normal mesh depth of 4 elements and the doubled mesh depth of 8 elements.

The bolt mesh was also slightly increased from 4 to 6 elements in the cases where the mesh density in the depth of the blocks was doubled. As the initiation of the preload in the bolt was almost entirely a linear strain problem it was anticipated that the mesh density through the depth of the bolt shaft would not have to be very large. The more complicated interactions occurred between the bolt head, nut and the two clamped components where the mesh was more dense already in anticipation.

Initial estimates as to the most efficient mesh design could be approximated from work done by Lee, Ko and Lee (2000), and De Matteis, Mandara and Mazzolani (2000) who modelled joints to a similar level of refinement as that used here. A convergent solution was assumed by the case that combined a doubled mesh density in both the depth and radial directions. It was also possible to compare the total normal force generated at the contact interface with the 19kN bolt preload that was applied. The two forces should be identical-although small numerical effects could be responsible for deviations of the order of a fraction of one percent.

Figure 3.7.1 shows the contact pressure distribution from the nodes where simulation data was collected (node 1652 Figure 3.6.1), through the bolt axis and across the contact interface in the y-direction. For the cases where there were only four elements in the depth of each column the maximum contact pressure is consistently larger than the cases where the mesh is 8 elements deep. The range of the largest and smallest contact pressures at the edge of the bolt-hole is almost a factor of ten. However this large discrepancy occurs in only a very local region formed by the nodes at the edge of the

bolt-hole. The models with a block mesh that is 8 elements deep did not have their maximum contact pressures occurring exactly at the edge of the bolt's clearance hole. In these cases the maximum contact pressures existed at nodes 1 mm from the bolt-hole and were larger, yet of a similar magnitude, to the contact pressure at the edge of the hole.




Figure 3.7.1 Variation of the contact pressure across the joint interface exhibited for a variety of mesh densities subjected to a 19kN bolt preload.

Another main feature of the contact pressure distribution is the fluctuation of the contact pressure in a 3mm radial region from the bolt-hole. Each mesh refinement produced a slightly different variation in contact pressure in this region. All of the different mesh designs featured the contact pressure becoming zero at least once, signalling an open contact between the two surfaces. One possible cause of this open contact is stress concentrations around the bolt-hole and Poisson's ratio effects. The region where the contact was open was very small in terms of radial distance. Due to the

differences between each mesh it is hard to say which is the "convergent" solution, but the most refined mesh would be the likeliest case. The mesh used in the dynamic simulations still exhibited this open contact and so maintained a physical agreement with the more refined convergence tests.

Away from the bolt-hole the pressure distribution shows a gradual decline from a value of between 4 and 10MPa depending on the mesh used. The mesh that is being verified also conforms to this behaviour. This pressure distribution is much more akin to the pressure cone model when viewed in isolation, and is distorted only by the significantly larger contact pressures that can exist around the bolt-hole (Figure 3.7.2). Although the contact pressure does not decrease uniformly in any of the verification tests meshes that were only 4 elements deep showed the largest superimposed fluctuations. The only point at which the decline in contact pressure does not hold is at the very edge of the block where all of the tests show a slight increase in contact pressure to a value of about 1.8MPa. It is the tests with the greatest radial refinement that show the closest agreement with one another in the region of generally declining contact pressure that occurs over most of the contact interface.

All of the mesh refinements, including the least refined mesh that was being verified, generated a normal force at the contact interface of 19kN to a suitable degree of accuracy. The control mesh, despite being the least refined of all the meshes examined, offered the closest agreement of maximum contact pressure between of all the 4-element-deep meshes to those that were 8 elements deep. The control mesh also displayed an open contact and corresponding pressure fluctuations within the 3mm radial region surrounding the bolt-hole. Agreement with the more refined meshes could also be found in the pressure cone-like behaviour away from the locality of the bolt-hole and over the majority of the contact interface. The only notable deviation in behaviour between the control test and more refined meshes was that the control mesh showed a contact pressure that declined right to the edge of the block instead of the small increase shown in the more refined meshes.

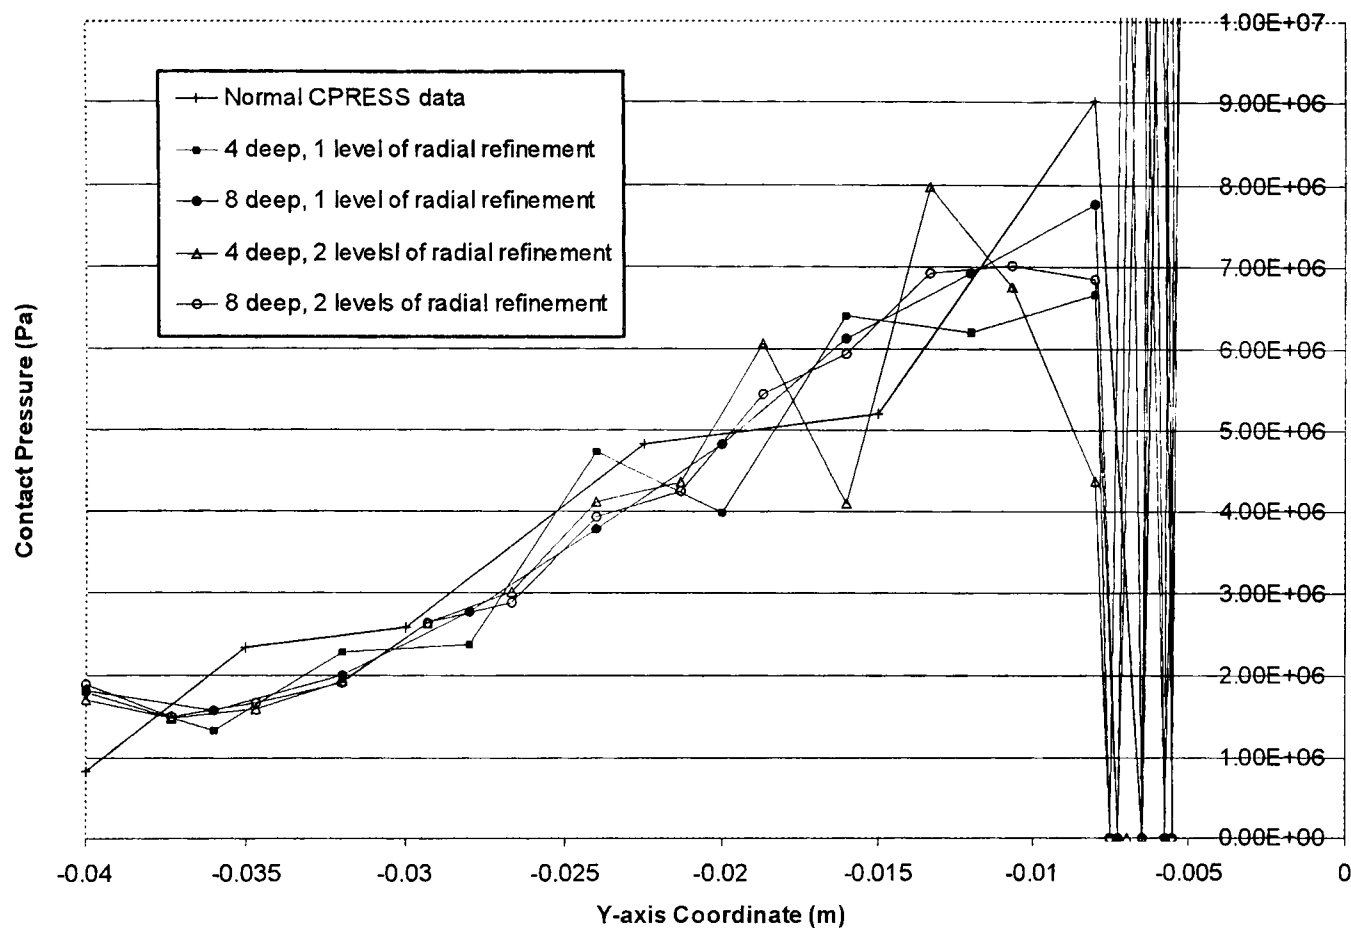


Figure 3.7.2 Contact pressure distribution at the contact interface showing pressure cone-like behaviour over the majority of the contact interface.

In general terms the control finite element model displayed the same contact pressure distribution as the more refined models used to verify it. The control model also had contact-pressure magnitudes of a similar size to the more refined loops. From this perspective the control mesh could be described as suitable. The only areas of inconsistency, close to the bolt hole and at the very edge of the clamped block, were considered unlikely to have a major impact on the overall behaviour of the model. Both were very local effects, and particularly close to the bolt-hole were also very inconsistent across the range of verification tests carried out.

The other verification test that was performed was on the way that the load was applied to the upper block in the dynamic part of the analysis. As the control mesh of Figure 3.3.2 showed acceptable behaviour in the verification tests, this was the mesh that was used in all tests from here onwards. Initially point loads were applied to the upper

corner of the top block to generate a 240Nm torque at 20Hz. The reason behind applying the load at the top corners was that they were as far away from causing local deformations at the contact interface as possible. Results showed that the hysteresis loop generated in this manner was offset (Figure 3.7.3). This was unlikely to be a phenomenon so strongly exhibited by a harmonic torque applied in a purely numerical situation. To try and eliminate this offset experimentation with the load application points was carried out. Firstly the four point loads were applied to the corners of the upper block at each nodal point through the depth of the block.




Figure 3.7.3 *Hysteresis loops generated by a 240Nm torque applied at 20Hz to the upper block through 4 point loads at various depths through the upper block. (Control test represents the same torque applied to the vertices of the upper block).*

As the offset of the block was still too large in all the tests to be considered acceptable, the loads applied to the corners were then distributed from 1 to 3 nodes on the vertex of the upper block. The 3 point loads were applied to adjacent nodes and were tried at the three possible combinations of positions on each vertex. Results for these distributed load tests can be seen in Figure 3.7.4. The reduction of offset in the results suggested that a smaller nodal load, more evenly distributed over the vertex, provide the best results. Offset was minimised when the load was applied to the middle 3 out of 5 nodes on the vertex of the upper block. This suggested that local deformations at the contact interface contributed to the undesirable offset. Minimising the load applied to each node would reduce these local effects and could be obtained through applying the corner loads evenly to every node on the vertex of the model.

An even distribution of nodal forces, with minimised local deformations, contributed to a hysteresis loop that was the most centred of all the tests carried out. By applying the load at the vertices of the block the torque experienced at all points on the contact interface was equal adding useful symmetry to the analysis carried out. Applying the load over the entire block also created a more realistic condition when considering how the joint would be loaded if incorporated to the end of a beam. In such a case the torque would not normally be localised to a small portion of the beam's depth. Figure 3.7.4 shows the control finite element test in comparison with other less well distributed loads. The control test, with the applied vertex load spread over five nodes instead of three, was therefore used for all the finite element simulations in the remainder of this chapter.




Figure 3.7.4 *Hysteresis loops generated with the applied corner load distributed over three nodes, in comparison with the control test of distribution over the five nodes of each vertex.*

3.8) Constant Preload and Variable Torque Magnitude

In this series of tests a fixed bolt preload of 19kN was applied using ABAQUS/Standard. The result was then exported into ABAQUS/Explicit where a sinusoidal torque was applied to the corner vertices of the upper block. The maximum magnitude of the torque was varied from a minimum value of 40Nm to a maximum value of 280Nm. Above the 280Nm torque range the response of the system became highly irregular and is believed to occur as the whole contact interface between the clamped blocks underwent sliding. This was in contrast to the behaviour up to that point when gross slip was assumed not to have occurred and the mechanism for frictional energy dissipation was microslip.




Figure 3.8.1 Contact pressure distribution for a 19kN bolt preload. Measurements taken at nodes along the y-axis, passing through the bolt axis (0.04m), over the full breadth of the top block.

Figure 3.8.1 shows the contact pressure distribution across the interface caused by a 19kN preload. In the extreme locality of the bolt-hole the contact pressure reaches an extreme value before dropping down to zero-signalling an open contact. The most likely cause for this has already been put forward (Section 3.7) as stress concentrations around the bolt-hole. It should be noted that although the contact pressure does not decay to zero at the edges, along the diagonals of the square contact interface the radial distance increases, and it is anticipated that the contact pressure will drop further still at these points.

Even the case where the peak magnitude of the torque was only 40Nm, shows a degree of frictional energy dissipation of about 3.85×10^{-5} (J/cycle). It is only a quarter of the size of the amount of frictional energy dissipated as reported in the ABAQUS/Explicit output. The reason for the discrepancy is likely to be a degree of microslip that takes place in the very local area around the bolt's hole where the contact pressure is very low. As the hysteresis loops are plotted from data obtained at the edge of the block, they are unlikely to display much slippage that takes place near the centre when the edges are still undergoing predominantly elastic contact.

Peak Torque Value (Nm)	Energy Dissipated (J/Cycle)
40	3.85×10^{-5}
80	2.69×10^{-4}
120	5.47×10^{-4}
160	9.53×10^{-4}
200	2.50×10^{-3}
240	9.99×10^{-3}
260	1.54×10^{-2}
280	2.21×10^{-2}

Table 3.8.1 *Frictional energy dissipation values for sinusoidal torques of varying peak magnitude with a constant 19kN bolt preload.*

For the case where the maximum torque applied was 40Nm, there was still some microslip occurring at the extreme edges of the blocks. Figure 3.8.1 shows that the contact pressure does not decrease to exactly zero. However it should be noted that these points are taken from the midpoint of each edge in the x-y plane, and the contact pressure will decrease yet further along the corner-to-corner diagonals as the radial distance from the bolt axis is greater. Consequently these areas of very low contact pressure will undergo sliding at the instigation of a very small amount of torque. Figure 3.8.2 clearly illustrates this small but significant area enclosed by the hysteresis loop despite the fact the frictional energy dissipation is much reduced when compared with the other test cases as shown in Table 3.8.1.

Table 3.8.1 shows that as the maximum torque magnitude is increased, the amount of energy dissipated per cycle also increases. The reason for this is twofold. Firstly increasing the maximum input torque, and maintaining the same frequency of excitation achieve a given magnitude of applied torque more rapidly in each cycle. Areas with a particular level of limiting friction are therefore overcome earlier to initiate local sliding. Consequently the amount of time spent sliding by a given region in each cycle increases along with the energy dissipated per cycle. The second reason for the increased energy dissipation is that areas of the contact interface where the contact pressure is larger start to slide, where before a smaller torque could not overcome the limiting friction. Over the range of maximum torques applied at 20Hz the increase in energy dissipation is not proportional. This is produced by the combined effect of an increased torque causing currently sliding areas to slide further *and* previously elastic contacts to start to slide and dissipate energy.

As these previously stuck areas start to slide the contact stiffness reduces. To accompany the increase in energy dissipation that this causes, there is also a loss of stiffness at the contact interface. This can be seen in overall terms on the hysteresis plots by considering the stiffness from one point of velocity reversal to another. Across the

range of input torques, the average stiffness over a complete cycle decreased as the torque increased.




Figure 3.8.2 *Hysteresis generated by torques of maximum magnitude 40Nm to 200Nm when applied to upper block clamped by a 19kN preload.*

The behaviour of individual hysteresis loops can also reveal interesting information about the amount of the surface that is undergoing sliding contact. In the cases with the smallest maximum torques (40Nm to 160Nm) the loops are largely bilinear (Figure 3.8.2). Each portion of the loop in these cases display very similar stiffness. This suggests that in all of these loops the same areas of the contact interface are undergoing sliding contact. The 200Nm torque case, also exhibits the same regions of constant stiffness as the bilinear loops. When the applied torque approaches its maximum value, shortly before velocity reversal, the stiffness of the hysteresis loop has decreased to a new low. It is inferred from this that a new previously sticking part of the contact

interface has started to undergo sliding and the limiting torque for this region, when considering a 19kN preload, lies between 160 and 200Nm.

Once the maximum torque was increased from 200Nm to 240Nm a new behaviour was exhibited by all of the hysteresis loops. All of the remaining cases (240Nm, 260Nm and 280Nm) display the same behaviour as the 200Nm case over the initial part of the loops after velocity reversals. The bilinear, and then gradually decreasing stiffness, are present in all three cases before the stiffness again decreases to a new low limit. All three blocks demonstrate this same minimum stiffness for a varying duration before the next velocity reversal. The largest, 280Nm torque, leads the region of lowest stiffness for the longest duration (Figure 3.8.3).




Figure 3.8.3 *Hysteresis generated by a torques of maximum magnitude 200Nm to 280Nm when applied to upper block clamped by a 19kN preload.*

A feature of the lowest stiffness limit is the slightly non-smooth profile of the hysteresis loop. The most plausible explanation for this irregularity comes from the fact that when the torque is increased beyond 280Nm the response breaks down completely

and grows dramatically. The nominal area of the hysteresis loops after this stage can become negative and grow larger in magnitude by several orders per cycle. It is inferred that the whole of the contact interface is, after this further increase of torque, undergoing gross slip. The joint in the context used herein is considered to have failed and is not investigated beyond the stable limit of 280Nm applied torque. From this inference of gross slip it is possible to ascribe a phenomenon to the region of constant stiffness immediately before velocity reversal. Only the smallest possible area of the interface between the two blocks will be providing any contact stiffness during the harmonic excitation. This region is likely to be the very local area around the bolt-hole where the contact pressure exhibits its unusually large values.

Despite this very small region of sticking contact the gradient of the hysteresis loop is certainly not negligible. The Encastré boundary conditions applied to the bottom surface of the bottom block, coupled with the tied condition in existence between the bolt-nut component provide a torsional spring-like behaviour from the bolt-nut component.

As a result of the behaviour illustrated by all the hysteresis loops it is possible to describe the process of microslip from its onset to the failure of the joint at the point of macroslip. Initially after velocity reversal a negligible amount of the block is in sliding contact. This is the same after every velocity reversal no matter how large the maximum torque becomes. This region of constant "sticking" stiffness can be seen on each hysteresis loops as a small but significant region of greatest stiffness that is consistent among all the cases illustrated in Figures. 3.8.2 and 3.8.3 As the amount of torque is increased a clear transition is reached when more of the contact interface starts to slide and some contact stiffness is lost. After the torque required to exceed the limiting friction of this region was applied, a fairly smooth transition is made to the point where virtually the whole contact interface is sliding. This represents another region of constant stiffness across a number of simulations and exists before the velocity is reversed. Upon velocity reversal the whole process is repeated, starting with maximum stiffness and progressing to macroslip and failure, but in the opposite direction.




Figure 3.8.4 *Time domain response of the upper block to an 80Nm and 240Nm maximum applied torque.*

The softening behaviour is clear from the time domain response of Figure 3.8.4. While the restoring torque is almost entirely sinusoidal, the angular displacement is most definitely not. The reason for this is the decreasing stiffness of the contact interface as the torque is increased to its maximum magnitude. Immediately after velocity reversal it can be seen that the system is stiffer than the average value represented by the almost purely sinusoidal restoring torque. As the cycle progresses, the stiffness of the system reduces until it is at a minimum value (below that of a pure sinusoid) immediately before velocity reversal. In contrast the time series plot for 80Nm maximum torque shows relatively little softening behaviour when compared with the larger applied torque. While still oscillating at the excitation frequency the case of 80Nm torque on average is much stiffer than the cases with larger applied torque. The corresponding hysteresis loop shows bilinear behaviour with two distinct regions of constant stiffness. This behaviour can be seen in

the time domain response where there is a very distinct "corner" between velocity reversals. The corner represents the abrupt transition from sticking stiffness to the second stiffness associated with a bilinear response.

Another factor that is apparent from the time series plots of Figure 3.8.4 and from the hysteresis plots of Figures 3.8.2 and 3.8.3 is that the inertial effects in the system are small. The reason for this is firstly that the moment of inertia about the bolt axis of the top block is very small indeed $0.00159744\text{kg}\cdot\text{m}^2$ and secondly the frequency of excitation is low relative to the natural frequency of the system. The only modifying effect is the rotary inertia which has been shown to be small and therefore the restoring torque is almost identical to the forcing torque.




Figure 3.9.1 *Contact pressure distribution created by bolt preloads of various magnitude. Measurements taken at nodes along the y-axis, passing through the bolt axis (0.04(m)).*

Each of the hysteresis loops generated using 28.5kN, 38kN, 57kN, 76kN and 85.5kN preload in conjunction with a 240Nm applied torque has a very definite bilinear profile. Individual loops differ in the amount of the loop that is defined by regions of a given stiffness. However the two regions clearly have the same stiffness in each case as illustrated by fact that the two portions of the hysteresis loops are parallel in all instances.

All five hysteresis loops share the same initial stiffness as the loop defined with a 19kN preload. This can be considered the initial sticking stiffness of the contact interface where only a tiny proportion of it is undergoing microslip. Very little energy is dissipated, and the response can be considered virtually elastic. What is notable is that the duration of the loading cycle spent in this regime increases with preload. It can be

inferred that by increasing the preload, and hence contact pressure, it takes longer for the first significant departure from this regime to occur i.e. a greater amount of torque is required to cause the onset of significant microslip within each half cycle. The second thing that can be inferred is that the contact stiffness of the interface in the model is not affected by the preload. In tribological terms this may be debatable, but in the case modelled on a macroscopic scale here it holds true.




Figure 3.9.2 Steady state hysteresis created by a 240Nm sinusoidal torque applied to the upper block at bolt preloads from 19kN to 38kN.

Another interesting feature of using 28.5kN to 85.5kN preload hysteresis loops is that all five cases demonstrate the same stiffness in the region where microslip takes place. As well as showing that the contact stiffness does not change with preload, this also shows that there is a plateau of limiting friction at higher preloads that the 240Nm torque cannot overcome. The time spent in the sliding range below this plateau is greatest for the 28.5kN preload case and smallest in the 85.5kN preload case (Figure 3.9.3). That the stiffness in the microslip range below the plateau is much greater than the stiffness solely provided by the torsion spring-like effect of the nut bolt component reveals that a

substantial area of the contact interface undergoes elastic deformation only, at these preloads, with a 240Nm maximum torque.




Figure 3.9.3 Steady state hysteresis created by a 240Nm sinusoidal torque applied to the upper block at bolt preloads from 57kN to 85.5kN.

When contrasting the shapes of the hysteresis loop generated using a 19kN preload with those of the remaining loops it appears that definite regions of fairly constant contact pressure must exist on the contact interface. The first piece of evidence for this claim comes from the observation that there is a very abrupt change in the contact stiffness of the interface when a certain level of limiting friction is overcome by the input torque. It is also apparent that the bilinearity that exists in all of the cases with the larger preload, but hardly at all when a 19kN preload is specified. A region limiting friction exists that can be overcome comfortably by the 240Nm torque even as it is multiplied as the contact pressure is increased. However the next region of contact pressure on the interface, which is overcome in the case of the 19kN preload is not overcome for any of the larger preloads. The suggestion is that there is an area with relatively low contact

pressure, but that the next area of limiting friction on the contact interface is at a level that becomes unobtainable once the preload is increased beyond 19kN. This area is also not encountered gradually, as the bilinear hysteresis loops show a very consistent stiffness and abrupt transition between regions in each segment.

The bilinear hysteresis loops also illustrate a very significant phenomenon that must be considered when using the bolt in a dynamically excited structure. With a constant torque of 240Nm applied, different preloads are required to minimise the response of the structure and maximise the energy dissipated by the joint. The 19kN preload gives the maximum energy dissipation of 9.986×10^{-3} (J/cycle) (Table 3.9.1). However by undergoing microslip almost until the point of macroslip, the joint stiffness softens to the greatest extent and shows the greatest degree of rotation. As the preload increased the joint displayed an average hardening effect over a cycle. With the increased stiffness the amplitude of the response of the joint decreased. In fact the stiffest joint is found with the highest preload, yet the cases of two smaller preloads demonstrate almost the same amount of energy dissipation per cycle.

Joint Preload (kN)	Energy Dissipated (J/Cycle)
19	9.986×10^{-3}
28.5	2.148×10^{-3}
38	2.197×10^{-3}
57	2.429×10^{-3}
76	2.405×10^{-3}
85.5	2.216×10^{-3}

Table 3.9.1 *Energy dissipation from a steady state cycle when a 240Nm torque is applied to a joint with a variable preload.*

Altering the preload has a smaller relative impact on the amount of energy dissipated than the same degree of alteration to the applied torque. This can be

attributed to the fact that the energy dissipated is a function of both the limiting torque and the distance travelled. An increasing preload causes an increase of the limiting torque. In turn however, it causes a decrease of the distance travelled as the joint is stiffer due to the greater application of torque required to overcome the regions of increased contact pressure.

In a design scenario involving passive joint damping for example, it would be necessary to either find the best compromise between energy dissipation and allowable motion or to optimise one of the parameters. The average stiffening of the joint as the preload increases has been shown to alter the resonant frequency of the joint. The preload causing maximum energy dissipation has been shown to occur at a level in between a fully clamped, rigid connection and nominally pinned joint. However, even if the joint is resonated at the fundamental frequency associated with each varied preload, the above finding has been shown to hold by Beards (1985).

3.10) Variable Friction Coefficient

One of the most interesting tests that could be performed on the finite element model was to examine the relationship:

$$\tau_l = \mu p \quad (3.10.1)$$

In Equation 3.10.1 τ_l represents the limiting value of shear stress, μ the friction coefficient and p the contact pressure.

If this simple relationship holds for a bolted joint then an increase in preload by a given factor, when accompanied by a decrease in the friction coefficient by the same factor should, give the same value of limiting friction and consequently the same hysteresis loop. This is acceptable if the stiffness of the contact interface does not alter with varying friction coefficient or preload. Sections 3.8 and 3.9 suggests that this is the case and so the relationship in Equation 3.10.1 can be tested.

Figure 3.10.1 shows that there is a negligible change in the contact pressure distribution when the coefficient of friction is increased or decreased. Only the maximum value around the bolt-hole shows any visible change, and even that is relatively insignificant. Along with the fact that the variation only exists in a very localised region it can be said that the contact pressure distribution is essentially independent of friction coefficient. This means that the equation can be verified by varying the pressure distribution and/or friction coefficient in the knowledge that no coupling existing between them.

A series of simulations was carried out whereby a comparison could be drawn into the effect of varying the friction coefficient in the finite element model and whether this was comparable to a similar alteration in preload. It would have been possible to just maintain a constant limiting value of friction by simultaneously altering both the friction coefficient and the preload. However this strategy was rejected, as by combining the

effect of two variables, the impact of either the friction coefficient or the bolt preload as individual entities may have been masked. The control test of a 19kN preload, 240(Nm) torque and friction coefficient of 0.8 is shown in Figure 3.10.2 It is accompanied by simulations where the coefficient of friction has been increased to 1.6, 2.4 and 3.2. These values are investigated to see whether they give the same results as increasing the bolt preload to 38kN, 57kN and 76kN. The friction coefficients given above are abnormally high suggesting, for example, a large amount of adhesion in the joint. However, as the finite element simulation is purely a numerical model, the physical representation of the behaviour is nominal only and the governing equations are not dependent on physical realism.




Figure 3.10.1 Contact pressure distribution for 19kN bolt preload and varying friction coefficient. Measurements taken at nodes along the y-axis, passing through the bolt axis (0.04m).

There is another simulation not shown in the figure where the friction coefficient is reduced to a value of 0.64. This value comes from a scaling of the 240Nm maximum applied torque relative to the torque that causes gross slip in the joint interface (about 300Nm). The friction coefficient has been inversely scaled by the same factor that, when applied as torque, caused macroslip. As macroslip did occur in this simulation, it suggests the relationship of Equation 3.10.1 does hold better for friction coefficients less than 1.




Figure 3.10.2 Steady state hysteresis response for a 240Nm magnitude, sinusoidal input torque to a system with a 19kN bolt preload.

Whilst doubling the friction coefficient produces almost the same amount of hysteresis as doubling the preload to 38kN (Table 3.10.1), tripling and quadrupling the friction coefficient results in a much smaller amount of energy dissipated per cycle than the same function performed on the preload. A comparison of the hysteresis loops (Figures. 3.9.2 & 3.10.2) shows that they are significantly different even in the case

where the friction coefficient was doubled and the amount of energy dissipated was similar.

The most obvious difference between each linked pair of results in the preload and friction coefficient cases is that the average stiffness is different. Increasing the friction coefficient results in an average stiffness that is consistently larger than that shown when increasing the preload. The reason that the stiffness is larger is that the transition to the second, less stiff, region of bilinearity is delayed. In each cycle a much larger amount of time is spent in almost entirely elastic deformation of the contact interface. Table 3.10.1 shows that as a consequence of this, doubling and tripling the coefficient of friction results in a much smaller amount of energy dissipated than performing the same operation on the preload.

It is therefore apparent that the response to a different coefficient of friction possesses certain properties that are not present when adjusting the preload. To assume the simple Coulomb friction law (Equation 3.10.1) holds in all cases is not possible. The indication is that the basic linear relationship between the parameters is more likely to hold when lower coefficients of friction are used i.e. less than 1. Evidence is provided for this by the fact that gross slip takes place in the joint when using the predicted friction coefficient of 0.64. The fact that such a large value of friction coefficient (0.8) is used in the control analysis makes the problems encountered here more acute. Significant increase of the coefficient introduces numerical problems while a relatively small decrease of 20% results in macroslip at the joint interface. In most physical situations, friction coefficients of 0.3 to 0.5 are more common. Such cases would allow much more range for increasing the coefficient of friction before it became greater than 1 and implied a degree of adhesion that may not exist in reality.

Despite the irregularities discussed above, and when viewed in isolation, increasing the friction coefficient does show some predictable trends. Most notably the energy dissipated per cycle decreases as the friction coefficient is increased. This stems from the fact that the duration of even the smallest amount of microslip decreases as the

friction coefficient, and therefore, limiting friction level increases. As the limiting friction is increased a larger friction force is required to overcome it and this value is reached at a progressively later stage in each sinusoidal loading cycle. Correspondingly the length of the hysteresis loop with the initial stiffness immediately after velocity reversal increases with every progression of the friction coefficient.

Friction Coefficient	Energy Dissipated (J/Cycle)	Preload (kN)	Energy Dissipated (J/Cycle)
0.64	N/A	N/A	N/A
0.8	9.986×10^{-3}	19	9.986×10^{-3}
1.6	2.177×10^{-3}	38	2.197×10^{-3}
2.4	6.354×10^{-4}	57	2.429×10^{-3}
3.2	3.594×10^{-4}	76	2.405×10^{-3}

Table 3.10.1 Comparison of energy dissipation when a 240Nm harmonic torque is applied with a varying friction coefficient or a correspondingly scaled level of preload.

For friction coefficients above a value somewhere between 0.8 and 1.6, the hysteresis loop becomes bilinear. The hysteresis loop for a friction coefficient of 0.8 shows the full range of stiffness and consequently microslip. Above a value of 1.6 only the initial stiffness and a slightly reduced stiffness is exhibited in any of the loops. This suggests only a very small area of the contact interface is undergoing microslip and that the contact pressure in this area is very low. It must be very small as, even when significant increases in limiting friction are achieved by altering the friction coefficient, a value of torque between 0Nm and 240Nm is still capable of overcoming this threshold in all cases. However in the case of a friction coefficient of 3.2 the maximum 240Nm torque is only just capable of producing a force at the contact interface capable of exceeding the friction force in even this (assumed) small area. The duration of sliding in this region does decrease in all of the loops, but most significantly the bilinear hysteresis loops, as the friction coefficient is increased.

Another very significant phenomenon displayed when the friction coefficient becomes very large is the drift of the hysteresis loops. While in certain cases drift does have a physical justification (Section 4.8), here it is likely to be a result of the changing contact algorithms at the beginning of the dynamic loading stage. Most apparently the friction coefficient of 2.4 is offset so that the steady state vibrations are positive angles when the restoring force becomes zero in its cycle (regardless of positive or negative velocity). Only the initial loading cycle beginning from the origin can be claimed to be centred. After the initial velocity reversal the first cycle also largely follows the path of the drifted steady state cycles. As the applied torque is harmonic and symmetric this suggests an inconsistency in the contact algorithm. Only the initial slope appears to be affected as all the subsequent loops settle into some sort of steady state cycle. This points to the importing process of results from ABAQUS/Standard, to ABAQUS/Explicit with its slightly modified contact algorithm, as the most likely source of these unusual results.

3.11) Uniformly Distributed Preload

The contact pressure distribution at the joint interface can be seen to have a very significant impact on the way that a joint behaves. In the case of a bolt preload the areas of low stiffness caused microslip to be initiated almost immediately and to progress gradually as the torque was increased over a cycle. In a multiple bolted joint where the bolts are close together a far more uniform loading is created. The uniformly distributed load applied to the top surface of the upper block caused a similar, nearly uniform contact pressure distribution at the joint interface that did not approach zero at any point. A comparison of the pressure distribution caused by a 19kN bolt preload and a 19kN uniformly distributed load is shown in Figure 3.11.1.




Figure 3.11.1 Contact pressure comparison between bolt preload and uniformly distributed load cases. Measurements taken at nodes along the y-axis, passing through the bolt axis (0.04m).

Consequently a certain magnitude of torque has to be applied to a uniformly loaded upper block before any microslip was visible at all. Over a substantial range of torque (up to at least 240Nm) the amount of microslip in the system remains negligible. The area of the hysteresis loops is very small and the energy dissipated from a structure with this kind of pressure distribution and applied loading will be of the same order, or less than that provided by material damping etc.

When comparing the amount of energy dissipated by 160Nm and 240Nm of torque in the case of 19kN bolt preload and 19kN uniformly distributed load, it can be seen that the amount of energy dissipated in the case of the uniformly distributed load is significantly decreased (Table 3.11.1).

Maximum Applied Torque (Nm)	Energy Dissipated (J/Cycle)
160	5.689×10^{-5}
240	2.2587×10^{-4}
320	1.6092×10^{-3}
360	4.5518×10^{-3}
400	8.0195×10^{-3}

Table 3.11.1 *Comparison of the steady state frictional energy dissipated between the two blocks when subjected to a 19kN uniformly distributed load and torques of sinusoidally varying magnitudes.*

Instead of showing a gradual progression from a sticking through microslip to a state of almost gross slip, the uniformly distributed load only allows a large region of stuck contact and another region of reduced stiffness where microslip is present (Figure 3.11.2). The bilinearity of the hysteresis loop is analogous to the cases where the bolt preload was increased and several comparable hysteresis loops with the same pair of stiffnesses were visible. The main distinction between the cases of bilinear hysteresis found with a bolt preload and the bilinear hysteresis found with a uniformly distributed load is that the uniformly distributed load bilinearity is the extent of the shape of the

hysteresis loop. Increasing the torque to a level required to cause gross slip at the contact interface does not result in the further softening of the contact stiffness seen in the bolt preload case. The maximum torque applied in the case of the uniformly distributed load of 400Nm, and to a lesser degree the 360Nm torque, show the gradual increase in the amount of noise in the system as the condition of gross sliding is approached. That this condition is approached with only a very small decrease in stiffness from that exhibited when microslip is initiated is indicative of only a small variation of contact pressure at the joint interface. With an entirely uniform contact pressure distribution on the contact interface microslip would not exist at all. Instead all parts of the contact interface would start to slide at the same time resulting in an abrupt transition from sticking stiffness to a noisy response associated with macroslip and joint failure.




Figure 3.11.2 *Steady state hysteresis generated by torques of varying maximum amplitude applied to a block at 125.66rad/s. The preload condition is a 19kN uniformly distributed load applied to the upper block to which the torque is applied.*

Figure 3.11.2 does illustrate the behaviour of a contact interface with a largely uniform contact pressure distribution well. Initially the amount of hysteresis in the system is very small and up to a torque of 240Nm the response can be considered to be almost entirely elastic. As the magnitude of the maximum torque was increased, a transition was reached whereby a small amount of microslip was present in the contact interface. The onset of microslip led to a significant increase in the amount of frictional energy that was dissipated and also a reduction in contact stiffness. Both of these behaviours can be seen in the cases where the maximum applied torque was 320Nm and 360Nm.

The final, 400Nm, torque case illustrated the final two phenomena of note. Firstly the amount of microslip did not alter greatly before the onset of a noisy response indicating, as described above, that gross slip was starting to occur. Significantly this happened after only a small region of microslip had been experienced and indicates that the pressure distribution was largely uniform, as shown in Figure 3.11.1. The transition from sticking contact to gross slip therefore took place over a narrow range of applied torque.

The second response that is notable in the 400Nm simulation is the offset of the hysteresis loop in the direction of the positive angular displacement. This offset is a known characteristic of hysteretic vibrations and is described as drift in the Bouc-Wen hysteresis model (Wen (1980)). Many of the hysteresis loops generated through the simulations carried out in this section demonstrate offset, but few do so prominently. Offset never occurs when the response is purely elastic, and is generally more prominent when there is a large degree of microslip tending towards gross slip. One explanation for the offset comes from the fact that the hysteresis loops are not exactly symmetrical in the case of increasing and decreasing angular velocity. Figure 3.11.2 shows that there is a longer duration where no friction is dissipated i.e. sticking contact when the velocity was negative compared to the parts of the cycle where the velocity was positive. The implication is that the amount of microslip was greater, and the stiffness of the interface less for longer when the velocity is positive. As a consequence a greater angular

displacement was experienced in this portion of the hysteresis loop. Again one potential source of this offset is thought to be the discontinuity in the contact algorithm experienced at the transition from implicit to explicit analysis techniques. Such a discontinuity is less likely to show up in a simulation where very little sliding contact takes place at initiation of dynamic loading and then in subsequent cycles.

3.12) Spectral Density of Finite Element Output

An analytical tool for further understanding the response of the joint in the finite element simulations was to investigate the frequency content of both the hysteretic restoring torque and also the angular displacement. The softening behaviour of the angular displacement as the joint lost contact stiffness through progressive microslip will have more than one frequency present in its output. If the response was purely viscous damping, and linear, then the hysteresis loops would be elliptical and the frequency density could be expected to have just a single peak. The peak would occur at the frequency of vibration which in turn would be equivalent to the excitation frequency of 20Hz.

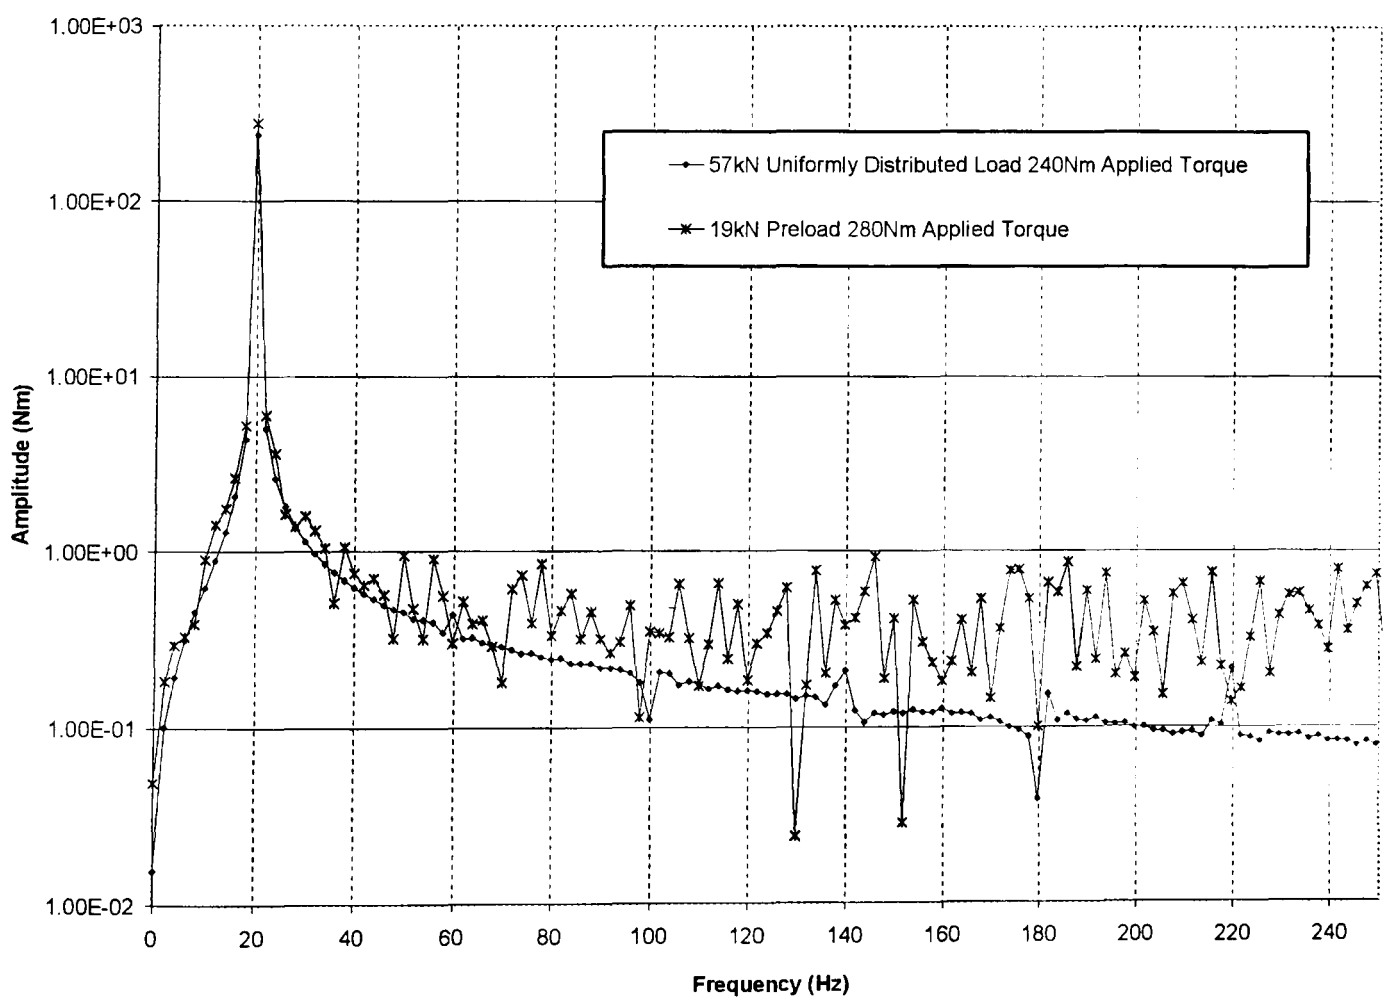


Figure 3.12.1 *Power spectrum density for the hysteretic restoring torque in the cases of an elastic contact interface and extensive microslip.*

Figure 3.12.1 shows that regardless of the conditions at the joint interface the hysteretic restoring torque is almost entirely harmonic. This condition is brought about by the fact that the input force is sinusoidal, and the inertial effects on the joint are also small. The forcing signal dominates the response and only a single frequency is observed that corresponds to the forcing frequency of the applied torque. Each case in Figure 3.12.1 represents a totally different kind of joint response. The 57kN uniformly distributed load resulted in a linear joint with a purely elastic response. No frictional energy was dissipated at any time from the joint interface and as material damping was disregarded the whole simulation can be considered linear conservative. The single peak at the excitation frequency is strong at a value just below the applied torque value. Away from the peak associated with the excitation frequency there is no signal of note at any frequency except a small response at the seventh harmonic. At this particular harmonic the strength of the signal can most likely be attributed to small fluctuations that occurred in the contact conditions at the stage when they were initially exported from ABAQUS/Standard to ABAQUS/Explicit.

The signal for 280Nm applied torque also has only a single peak associated with the frequency of excitation. Extensive macroslip lead to a relatively noisy response as points of velocity reversal were approached where the contact stiffness had decreased considerably. This is demonstrated in the far greater fluctuations in the response away from the excitation frequency. Notably none of the associated superharmonics have any clearly distinguished peaks suggesting that noise associated with approaching macroslip cause the greater fluctuations in the signal.

Next to be investigated were two responses where the hysteresis loop displayed bilinear characteristics for their different loading configurations. These were compared with the linear response obtained with a 57kN uniformly distributed preload. As discussed above the restoring torque does not offer clearly identifiable behaviour to identify the nonlinear response of the joint. The time series response on the other hand shows that the angular displacement does not display a sinusoidal response to match the

applied torque. Instead the angular response shows softening behaviour as the contact stiffness changes throughout the duration of a cycle.

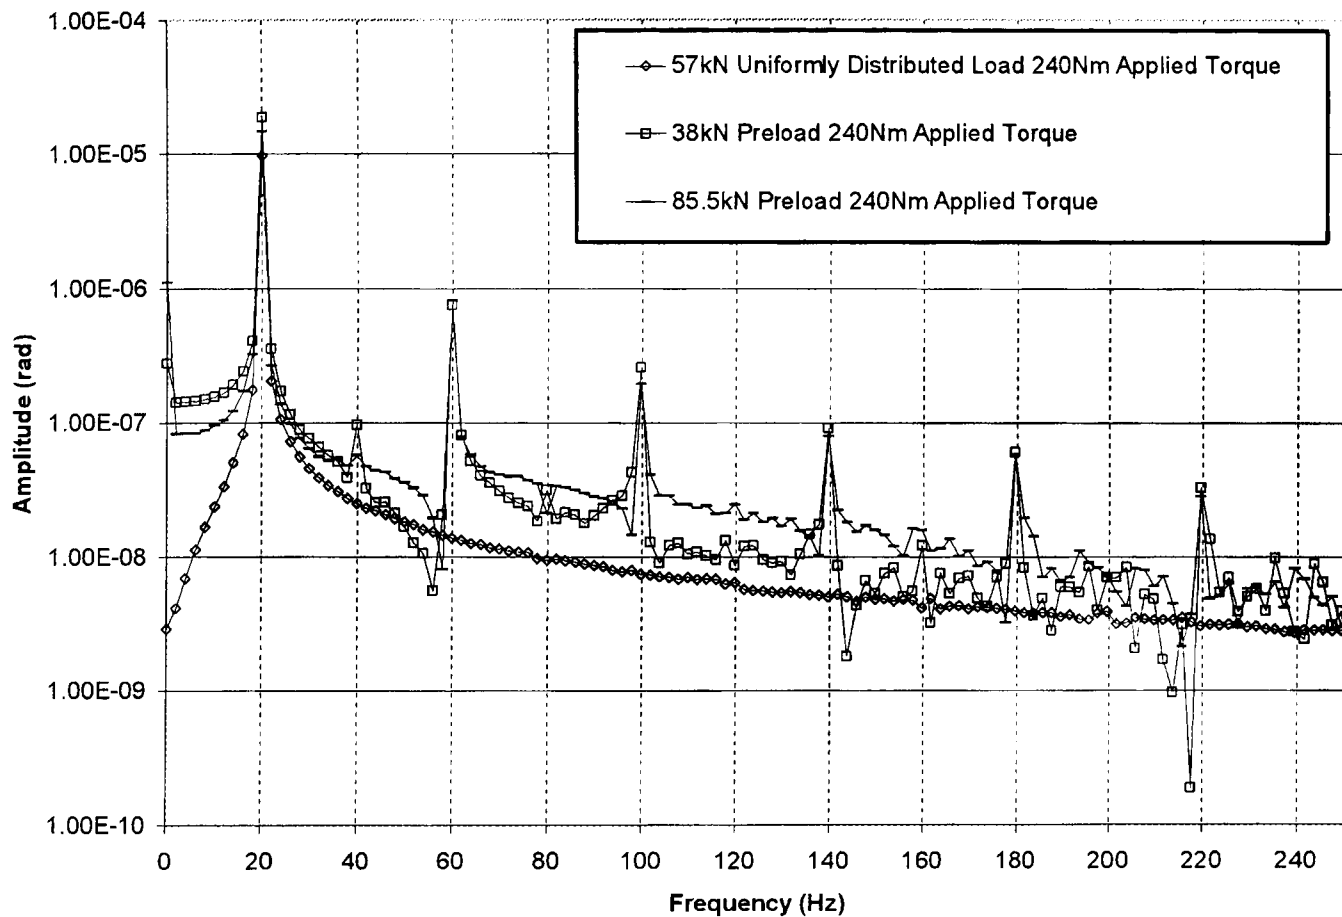


Figure 3.12.2 Power spectrum density for linear joint response and two cases characterised by bilinear hysteresis loops.

Figure. 3.12.2 predictably shows that the linear joint produced a frequency content that was almost entirely the result of the excitation frequency. Once the joint started to become nonlinear, even to a small degree suggested by bilinearity, the power spectrum changed dramatically. Even though the peak associated with the excitation frequency was present it was accompanied by clear peaks at its associated harmonics. The strength of these peaks was greatest in the case of a 38kN bolt preload. When compared with the larger (85.5kN) preload in question, the amount of sliding in the simulation was also greatest. In both cases the signal quality deteriorates towards the higher frequencies, but clear trends are identifiable. Firstly both cases show clear peaks at the odd superharmonics of the excitation frequency. When taking into account the

logarithmic scale of the graph, the superharmonics are even more visible up to 220Hz. In the case of the 38kN preload the 2nd and 4th and 8th superharmonics also show up as distinct peaks. These even harmonics are of a much smaller amplitude than the odd harmonics (even up to the 11th) except in the case of the 2nd harmonic. It is also apparent that the background signal between harmonics is greatest in the case where microslip is lessened by the larger bolt preload. Both signals that demonstrate microslip also have a peak at 0Hz. This offers immediate evidence for offset in the hysteresis loop. It also shows that offset is a function of the sliding that takes place in the simulation as the linear response offers no such peak at 0Hz.




Figure 3.12.3 Power spectrum densities from simulations where the hysteresis loops were generated from a condition of a bolt preload of 19kN.

In Figure 3.12.3 all of the cases demonstrated losses due to friction. The hysteresis loops also showed offset to varying degrees that are indicated by the signal

strength at 0Hz. Even though the preload was maintained at 19kN as the torque was increased, the hysteresis changed from bilinear (80Nm applied torque), to verging on macroslip (280Nm applied torque). The 200Nm case demonstrated an intermediate case where more macroslip was evident than in the bilinear cases, but the joint had not softened to the point associated with macroslip.

The bilinear hysteresis loop was consistent with similar loops obtained at higher preloads. The odd superharmonics are present and show a decreasing amplitude as the frequency increased. As the amount of dissipation increased in the joint, the magnitude of the superharmonics increased as well. At the higher superharmonics, in the cases of 200Nm and 280Nm torque, the amplitude of the peaks levels off to a more consistent level. The peaks of the signal associated with the 280Nm applied torque and larger angular displacement are of a consistently larger magnitude than seen at the smaller applied torques. This behaviour is consistent across the entire frequency range in question.

The most striking difference between the simulation conducted with the largest applied torques is that at the lower torque even harmonics show up quite distinctly. In the case of the 280Nm applied torque, the only even harmonic that is present to a significant degree is the 12th.

Chapter 4. The Jenkins Element Method

4.1) Jenkins Element Summary

Jenkins elements are a means of representing the presence of Coulomb friction in a dynamic system. A single Coulomb element connected in series with a linear spring makes up a single Jenkins element. The Coulomb element is considered to be in its stuck state and behaves as a rigid component until the compressive or tensile force in the attached spring becomes greater than a limiting value. This limiting value is equivalent to the product of the normal force applied to the element and the coefficient of friction. When the limiting value is reached, the Coulomb element makes its transition to a sliding state and offers a constant resistance to the direction of motion, equivalent in magnitude to the limiting value. When several Jenkins elements were placed in parallel, the gradual loss of stiffness, characteristic of the onset and progression of microslip, could be simulated. To demonstrate this ability, various second order systems with Jenkins elements providing the restoring force were created. These models were matched to the hysteresis loops generated using the finite element model. Resistance parameters of the Coulomb elements and the associated spring stiffness were extracted by identifying discrete points on the hysteresis loops. Between these discretisation points, the hysteresis generated by the Jenkins elements was assumed to be joined by regions of constant stiffness.

Different discretisation strategies for the finite element hysteresis loop were evaluated. When evenly distributed points of angular displacement were used, the level of agreement in terms of the energy dissipated per cycle by the two models was only adequate. Increasing the number of Jenkins elements and therefore discretisation points generally improved agreement. An alternative discretisation procedure, whereby the hysteresis loop was broken down visually into a number of different regions of approximately constant stiffness, showed improved prediction in the amount of energy dissipated. A model consisting of four Jenkins elements gave an energy dissipation per cycle to within 5 percent of the value obtained from the reference finite element study.

This model was capable of replicating the finite joint hysteresis behaviour but with several orders of magnitude fewer degrees of freedom.

The predictive power of the Jenkins element model was also evaluated. The Jenkins element models that best replicated the finite element simulations at relatively small and large applied torques were used to predict the response of the finite element joint by applying different torques to the reduced order model. A model consisting of four Jenkins elements and a permanent spring, extracted from the finite element joint with a large amplitude of applied torque offered good predictive ability. Both increasing and decreasing the applied torque in the reduced order model provided a credible prediction of the hysteresis behaviour of the finite element model at an equivalent level of applied torque.

The physical relevance of the Jenkins element parameters was demonstrated by considering the change in the response of the reduced order model when the normal force applied to the Jenkins element was altered. When the normal forces were altered in the same ratio as the bolt preload, the Jenkins element model was capable of predicting the response of the finite element model with a reasonable degree of accuracy. The sum of the stiffnesses of the Jenkins elements was shown to be equivalent to the contact stiffness of the totally stuck interface. The stiffness of the permanent spring could also be related to the residual stiffness in the finite element model, provided by the bolt-nut component's model constraints, when the contact interface was in a state of virtual macroslip.

When subjected to an initial displacement and left to vibrate freely, the energy dissipated led to a linear decay in the amplitude of oscillation. As a result fewer Coulomb elements made the transition from sticking to sliding during oscillations and the average stiffness of the joint increased. Eventually a steady-state, conservative cycle was created when the amplitude of oscillation of the model became so small that no Coulomb elements dissipated frictional energy. The equilibrium point of the conservative vibration was offset in the angular displacement dimension as a result of the time dependent contact conditions.

4.2) Jenkins Element Model

A single Jenkins element consists of a spring and a Coulomb friction element connected in series. Figure 4.2.1 shows a schematic representation of an isolated Jenkins element.



Figure 4.2.1 *A linear spring and Coulomb friction element connected in series to create a Jenkins element.*

In all of the analyses carried out here, the spring component is assumed to be linear. The Coulomb component is established with the most basic formulation of the Coulomb friction law. A limiting value of friction is established (F_1) which determines the sticking or sliding status of the Coulomb element. When the force applied to the Coulomb element via the attached spring is less than the limiting value then the Coulomb element behaves like a rigid component and provides an equal and opposite resisting load. As the applied load increases, the reaction in the Coulomb element also increases until the limit load is exceeded.

The limiting load of the Coulomb element is defined as:

$$F_1 = \mu N \quad (4.2.1)$$

Where the N is normal load and μ is the coefficient of friction. Once the load applied to the element by the spring exceeds the limit load, the Coulomb element offers a constant resistance of magnitude F_1 , in the opposite sense to the relative velocity:

$$F = -\mu N \operatorname{sgn}(\dot{x}) \quad (4.2.2)$$

In some friction models, the Coulomb elements have different static and kinetic coefficients of friction. All of the analysis provided herein assumes a constant friction coefficient.

The force provided by a single Jenkins element (Figure 4.2.1) can therefore be described as:

$$F = \begin{cases} k(x - x_{rev}) + F_1 \operatorname{sgn}(\dot{x}_{rev}) & |k(x - x_{rev}) + F_1 \operatorname{sgn}(\dot{x}_{rev})| < F_1 \\ F_1 \operatorname{sgn}(\dot{x}) & \text{else} \end{cases} \quad (4.2.3)$$

Where the subscript *rev* refers to the value immediately prior to velocity reversal. That the response of the Jenkins element is dependent on the displacement and velocity immediately before velocity reversal is the hereditary factor that makes the system hysteretic. The only exception to this rule is in the initial loading cycle when starting from rest and zero displacement when the parameter x_{rev} is equal to 0.

In a system where the restoring force is provided by a single Jenkins element, the response is elastic-perfectly plastic. This is a particular kind of bilinear hysteresis loop. A more general condition is achieved if a second linear spring is connected in parallel with the Jenkins element. The response remains bilinear, but when yielding of the Coulomb element occurs, the system retains some stiffness. Such a case can be seen in the case of the finite element simulations where some torsional stiffness remained (as a result of the constraints applied to the bolt-nut component) even when the whole surface had yielded in macroslip.

The bilinear Jenkins element model described above is adequate for representing hysteresis loops that are similarly bilinear. However, the finite element analyses carried out have shown that macroslip causes a more gradual loss of stiffness in the joint until the onset of macroslip. Jenkins elements can be used to model this, and have been by Gaul (1997). By connecting a series of Jenkins elements in parallel, and tuning their stiffness and limit loads, the onset of sliding can be made more progressive. If each element slips

after a different amount of relative sliding, then the loss of stiffness over a range of displacement will be gradual. To replicate the microslip exhibited by many of the finite element simulations, a model was devised where the restoring torque was provided by several Jenkins connected in parallel. A permanent spring connected in parallel with the Jenkins elements was also used to provide a means of modelling the residual stiffness when the whole contact interface was on the point of sliding.

The formulation of the restoring torque for n Jenkins elements and a permanent spring connected in parallel is given below:

$$F = k_0 x + \begin{cases} k_i (x - x_{\text{rev}}) + F_{li} \text{sgn}(\dot{x}_{\text{rev}}) & |k_i (x - x_{\text{rev}}) + F_{li} \text{sgn}(\dot{x}_{\text{rev}})| < F_{li} \\ F_{li} \text{sgn}(\dot{x}) & \text{else} \end{cases} \quad (4.2.4)$$

In Equation 4.2.4 k_i is the stiffness of the spring and F_{li} is the limiting load of the Coulomb element in Jenkins element i .

To make the finite element model and the Jenkins element approximation consistent, the coefficients of friction in both cases were 0.8. It was also possible to match the applied load between the two models. The frequency of excitation in the finite element model was 20Hz and so this was the frequency at which the load was applied in the Jenkins element simulations.

Although the Jenkins element model is defined in terms of force and displacement, the finite element hysteresis was described in terms of torque and angular displacement. A simple substitution, for reasons of analogy, of torque for force and angular displacement for linear displacement, removed this inconsistency. In the Jenkins element model the applied torque was the same magnitude as the applied torque in the finite element simulations. The Jenkins element model also replaced mass with the moment of inertia of the upper block about the bolt axis. Once these conversions were completed the response of the equivalent Jenkins element model was calculated using a

numerical time integration procedure. A fourth-order Runge-Kutta algorithm was performed using Matlab with system parameters either as defined by the finite element model or extracted from the finite element hysteresis results. A schematic diagram of the model that was integrated in the time domain is shown in Figure 4.2.2




Figure 4.2.2 The Jenkins element model used to replicate the behaviour of the finite element joint model. (F_{in} =input loading and m =inertia of the upper block).

4.3) Jenkins Element Parameter Extraction

To match the Jenkins element model of Figure 4.2.2 to the finite element hysteresis loop the correct stiffness and Coulomb element parameters had to be identified. All of the other model parameters such as the inertia and applied torque were replicated directly from the finite element analysis. The unknown parameters could be identified from the finite element hysteresis loops through simultaneous solving of the restoring force equation (Equation 4.2.4) at different displacements.

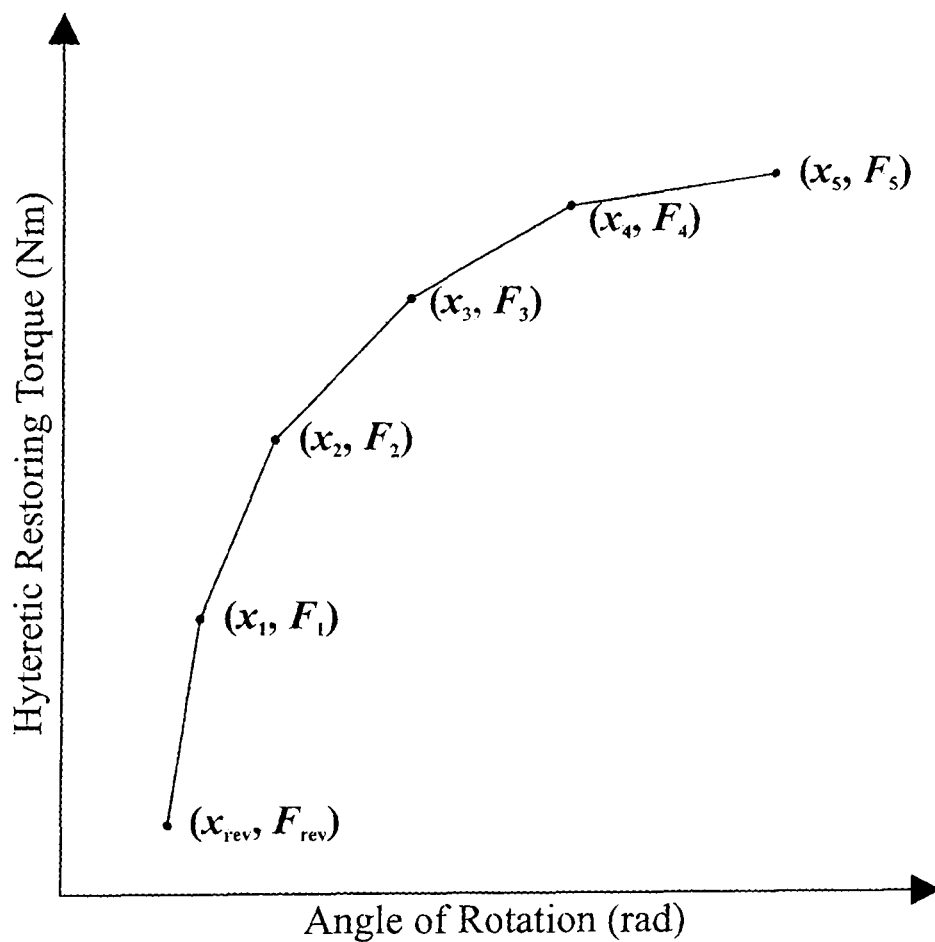


Figure 4.3.1 *The discretisation of a finite element hysteresis loop for Jenkins element parameter extraction.*

The finite element hysteresis loops had to be discretised before any parameter identification could take place. This was done by locating points on the finite element hysteresis loops and joining them with regions of constant stiffness. For a Jenkins element model with n elements the hysteresis loop was discretised into $n+1$ segments of

constant stiffness between velocity reversals. This is illustrated for the case of a 4 Jenkins element model in Figure 4.3.1. The first segment represents the stiffest system where none of the Jenkins elements were sliding. Between (x_4, F_4) and (x_5, F_5) the only stiffness contribution was provided by the permanent spring k_0 . Over each sequential segment one Jenkins element made the transition from sticking to sliding contact in the Coulomb element.

The stiffness of the springs was easy to identify from the gradient of the hysteresis loop between pairs of discretisation points. Between (x_{rev}, F_{rev}) and (x_1, F_1) the stiffness k_{tot} was given by:

$$k_{tot} = k_0 + k_1 + k_2 + k_3 + k_4 \quad (4.3.1)$$

Between (x_2, F_2) and (x_3, F_3) the total stiffness was given by:

$$k_{tot} = k_0 + k_1 + k_2 \quad (4.3.2)$$

As k_1 could be identified directly between (x_4, F_4) and (x_5, F_5) , all of the other stiffnesses could then be extracted sequentially.

Identification of the Coulomb elements' normal force was a slightly more involved process. When performing the parameter identification, a resistance was found that corresponded to the limiting friction force for each Coulomb element. This was converted into a normal force by dividing by the coefficient of friction.

At the point (x_1, F_1) , the point had been reached where one Coulomb element made the transition from sticking to sliding contact. Either equation of motion (sticking or sliding) for the Coulomb element in question described this point equally well. As a convention, it was taken that the Coulomb element was still stuck. As the angular displacement and restoring torque at this point were known, the restoring torque (F_1) could be described as follows:

$$F_1 = k_0 x_1 + (k_1 + k_2 + k_3 + k_4)(x_1 - x_{rev}) - R_1 - R_2 - R_3 - R_4 \quad (4.3.3)$$

Equation 4.3.3 is simply an expansion of Equation 4.2.4 with R_1 to R_4 used to represent the magnitude of the resistance force provided by the Coulomb elements when in sliding contact. Each discretisation point was described in such a manner. Equations 4.3.4 show the describing function at (x_3, F_3) and (x_5, F_5) respectively:

$$F_3 = k_0 x_3 + (k_1 + k_2)(x_3 - x_{rev}) - R_1 - R_2 + R_3 + R_4 \quad (4.3.4)$$

$$F_5 = k_0 x_5 + R_1 + R_2 + R_3 + R_4$$

As all of the stiffnesses (k_0 to 4) were known along with F_{rev} to 5 and x_{rev} to 5 it was possible to solve for R_1 , R_2 , R_3 and R_4 simultaneously. All of the parameters of the Jenkins element model could then be inserted into the time integration procedure for verification.

4.4) Evenly Distributed Jenkins Elements

A strategy had to be devised for discretising the hysteresis loop to obtain the parameters of the Jenkins elements matched the finite element hysteresis behaviour. The finite element simulation that was used for the comparative work was the control loop of generated by a 19kN preload and 240Nm maximum applied torque. This loop was used because of its representation of a full range of microslip behaviour from full sticking to the point of macroslip.

For general-purpose analysis a consistent algorithm for identifying the discrete data points on the finite element hysteresis loop had to be settled upon. This algorithm would ideally be suited to any hysteresis loop that was generated using finite elements or experimental observations.

It was also desired to test the agreement of the Jenkins element method for differing numbers of elements. The main reason for having an equivalent Jenkins element system was that a drastically reduced order model could represent the behaviour of the contact interface. Judging from the nature of the hysteresis loops obtained, it was not anticipated that the number of Jenkins elements required would exceed 10. This already showed that the number of degrees of freedom of the joint had been reduced by several orders of magnitude when compared with the finite element model. Further efficiency could be achieved by minimising the number of Jenkins elements in the representative system. Not only would this make the representative solution more efficient in isolation, but the bolted joint being investigated could be used many times over in a frame structure *and* feature more global degrees of freedom than just the angular rotation about the bolt axis.

Even in a finite element investigation of a frame structure with many bolted joints where tuned Jenkins elements replaced the detailed model of the joint interface, computational savings through an efficient model were desirable. A finite element model would not simply have to integrate a Jenkins element system in the time domain, but

would also need to combine it with all of the other constraints and element calculations inherent to the method. The effect of the extra calculations of the finite element method, further degrees of freedom and the potential for multiple joints meant that every joint needs to be refined to as high a degree as possible.

The most repeatable approach for testing the efficiency of the Jenkins element model was a regular discretisation of the finite element hysteresis loop from one point of velocity reversal to another. The hysteretic restoring torques that accompanied these angular displacements were obtained graphically from the finite element simulation. To use an automated extraction process for the data points required a hysteresis loop that was more consistent in the steady state oscillations, and was less noisy as gross slip was approached.

Figure 4.4.1 shows the agreement of the Jenkins element model when evenly spaced points were used for parameterisation. Five Simulations were carried each with an extra Jenkins element added to the model and subjected to the same input conditions as the finite element simulation that was to be replicated. Very poor agreement is shown by the single Jenkins element, as it is only capable of representing bilinear hysteresis with the permanent spring. The average stiffness of the loop was in good agreement with the finite element simulation. However this was almost a default situation as the extreme torques and displacements had been used as parameterisation points in the study. The energy dissipated, or frictional damping, of the reduced order model was much lower than that given by the finite element hysteresis. An obvious explanation for this is that the parameterisation is conducted by discretisation of points lying on the finite element hysteresis loop. As the Jenkins elements offer a linear stiffness, when not in the sliding state, the path linking discretisation points is a straight line. In the finite element simulation the gradual microslip exhibited causes the transition in stiffness between discrete data points on the hysteresis loop to be nonlinear and lie outside the Jenkins element approximation.




Figure 4.4.1 *Illustration of the degree of fidelity that exists between the control finite element simulation's hysteresis output and that generated by the reduced order model containing a variable number of Jenkins elements evenly spaced between velocity reversals.*

The relatively coarse, single-Jenkins-element, model cannot represent other significant behavioural properties of the finite element simulation. Immediately before and after velocity reversals the stiffness of the Jenkins element model is too inaccurate. After velocity reversal the stiffness is approximated to a value much lower than that required to properly represent the totally stuck contact stiffness. In the control finite element simulation gross slip is approached just before velocity reversal. The stiffness associated with the torsional spring behaviour of the bolt-nut component is approximated by the stiffness just before velocity reversal. The Jenkins element model overestimates this stiffness considerably due to the way that the hysteresis loop is approximated.

When considering the positive velocity component of the finite element hysteresis loop, even the addition of a single additional Jenkins element offers a significant improvement in the agreement between the loops. Further Jenkins elements increase the agreement between the hysteresis loops from the two different sources. All of the Jenkins element loops show good agreement in the average stiffness between velocity reversals, yet as more Jenkins elements are added (from 2 upwards) the similarity in the stiffness represented by the contact interface just before and just after velocity reversal shows an improvement. Up to four Jenkins elements every additional component added to the reduced order model furthers the agreement between both types of simulation. This is illustrated best by a visual comparison between the loops and also by comparing the amount of frictional energy dissipated in a given cycle (Table 4.4.1).

Number of Jenkins Elements	Energy Dissipated (J/Cycle)	Percentage Difference from Finite Element Result
1	7.3386×10^{-3}	-26.51%
2	8.7878×10^{-3}	-12.00%
3	8.8129×10^{-3}	-11.75%
4	9.1738×10^{-3}	-8.13%
5	9.0269×10^{-3}	-9.60%

Table 4.4.1 Comparison of the energy dissipated in simulations with evenly spaced stick-slip transitions in the Jenkins elements when compared with finite element simulation of 19kN preload and 240Nm torque.

Notably the addition of a fifth Jenkins element actually hinders the agreement between the finite element simulation and the numerical time integration. The difference between the loops with four and five Jenkins elements is caused by the way that the finite element loop is discretised. In the case of four Jenkins elements the points where changes in stiffness occurred, caused by the change in state of a Coulomb element, coincided better with the points where different regions of the contact interface also made the transition from sticking to sliding. The five-element model did not coincide so well with

these points as regions of constant stiffness in the finite element hysteresis loop were poorly approximated by the stiffness of the discretised regions on either side.

Another failing of the five-element model was that it had a redundant element in the part of the hysteresis loop before velocity reversal. In this region a relatively large range of angular displacement was represented by a single stiffness. While the four-element model covered the region with a single Jenkins element, the five-element model divided it into two even though no change in the contact state was evident in the finite element hysteresis. The additional obsolete element was still included in the time integration algorithm, and consequently represented dead weight in terms of computational cost. The four-element model was therefore a more efficient and accurate means of representing the finite element hysteresis than its more refined, five-element, counterpart.

In small displacement simulations the representation of the totally sticking contact stiffness become more significant as it characterises a greater proportion of each cycle. The range of displacement that actually represents this contact stiffness with a 19kN preload is small. To use the discretisation scheme suggested here, and represent that region accurately would require in the region of forty elements. Not only would the computational cost of such a model be extremely high even for a short duration of oscillations, but most of the elements would be obsolete as they would lie on regions of constant stiffness on the hysteresis loops. This is the major limitation of the evenly spaced method of obtaining parameters for the Jenkins element model. Particularly in the case of bilinear hysteresis loops the reduced order model should be able to accurately represent the finite element hysteresis with just a single appropriate Jenkins element and permanent spring.

Despite the limitations expressed above, the four or five element model visually corresponds well with the positive velocity portion of the hysteresis loop from which the parameters are extracted. The discrepancy in the amount of energy dissipated and frictional damping is worse than this matching of the profiles suggests. Figure 4.4.1

shows that in fact the finite element loop is slightly different in the negative velocity part when compared with the positive velocity part. One reason for this is that the hysteresis loop is slightly asymmetrical. There is also a small amount of offset in the finite element hysteresis loop that cannot be replicated by Jenkins elements without a superficial correcting term. Both of these features of the finite element loop make the accurate extraction of the finite element parameters problematic.

The reasons for the behaviour of the finite element simulation are unclear. As the loading and boundary conditions are equivalent in both the positive and negative velocity responses it is thought that the asymmetry is caused by inconsistencies in the contact algorithm when transferring from an implicit to an explicit analysis. This is particularly the case in the initial loading cycle. Potentially there are ways that the impact of these problems can be reduced. The most obvious is to extract the parameters for the positive and negative-velocity portions of the cycle and then average the results. Another potential solution would be to artificially centre the finite element simulation before parameter extraction was performed. Artificial centring of the hysteresis loop would maintain all of the main characteristics of the joint including the energy dissipated and all defining stiffnesses. Neither of these "patches" to the problem are ideal, and they do show that the Jenkins element method is limited when trying to represent drift in purely sinusoidal vibrations. However as the amount of offset is quite small when compared with total angular displacement the impact of drift is also correspondingly so.

4.5) Optimally Spaced Jenkins Elements

The contact pressure at the finite element joint interface did not decay linearly from the maximum at the bolt-hole and the minimum at the edge. It followed from this that the hysteresis loop did not show a regular reduction of stiffness from one point of velocity reversal. Partly due to the geometry of the contacting blocks and preloading, and partly due to the discretisation of the finite element mesh, there were parts of the hysteresis loop that showed an almost constant stiffness over a finite range of displacement. The transition from one region of constant stiffness being abrupt meant that instead of being approximated by equal regimes of constant stiffness, the hysteresis loops could, in certain circumstances, be discretised according to these regions of constant stiffness.

In the case of the bilinear hysteresis loops it was not necessary to discretise the loops into more than two regions of constant stiffness. However the case investigated as a control analysis (19kN bolt preload, 240Nm applied torque) showed a range of stiffnesses from one velocity reversal point to another. The initial stiffness corresponded to the fully stuck contact interface. The final stiffness was that contributed at the contact interface just before gross sliding was initiated and is attributed to the torsional stiffness provided by the bolt-nut component. In between these regions of constant stiffness at least three other reasonably discrete regions can be seen (Figure 4.5.1) and do not occur at regular relative angular displacements of the interface.

Previously data points for obtaining the Jenkins element parameters were obtained by matching the hysteretic restoring torque on the control finite element loop at evenly distributed angular displacements from one velocity reversal to another. It has been shown that this method consistently under-predicts the amount of energy dissipated in a cycle. The hysteresis loop traced out by the Jenkins elements always lies inside the finite element loop, except at the discrete data points where the two loops touch. A strategy for an improved fit of the Jenkins element loop with the finite element loop was to use visually-optimised data points to establish the parameters of the Jenkins elements. These

optimised data points would be able to take advantage of the regions of constant stiffness, and in many cases could be located at the points of abrupt changes in stiffness shown on the hysteresis loops.

Although a superior match between the two hysteresis loops was anticipated, this match was still reliant on a sufficient number of Jenkins elements being used to offer a suitable resolution in the results. By optimising the position of the discrete data points it was possible to match the constant stiffness offered by the Jenkins elements to the regions of constant stiffness in the finite element model. It should be noticed that the average stiffness of the control finite element loop and the Jenkins element hysteresis loop could be matched exactly due to the velocity reversal points being used in the discretisation procedure whether the data points were evenly spaced or optimised.

Figure 4.5.1 shows the hysteresis loops obtained with visually optimised data points. The parameters were obtained using the same method of obtaining the stiffness of each section directly from the finite element loop, and the resistance force by simultaneously solving the equations describing each segment of the hysteresis loop. It is obvious again that all of the hysteresis loops generated using Jenkins elements fall inside, or along the profile of the finite element loop. The aim of optimisation is that if enough Jenkins elements are used then it will be possible to match the two loops by virtue of the finite number of constant stiffness regions of the finite element loop.

It is immediately apparent from Figure 4.5.1 that the location of the optimal discretisation point for the 1-Jenkins element model is almost the same as that of the evenly spaced Jenkins element model. Little benefit can be seen of the more flexible approach to the placement of the parameterisation of the Jenkins elements in this particular case. The range of stiffness described in the finite element simulation is too large for a model this crude to be effective.




Figure 4.5.1 Control finite element loop (19kN bolt preload and 240Nm applied torque) matched by a second order system with the restoring force described by a varying number of Jenkins elements. Jenkins element parameters are obtained from visually-optimised discretisation.

The poor agreement of the 1 Jenkins element model can be seen in the comparison of energy dissipated by the two models (Table 4.5.1). Even the addition of just one additional Jenkins element to the model resulted in a dramatic improvement (over 13%) in agreement of the energy dissipated per cycle. The main reason for this improvement in agreement is that there is a far superior approximation of the stiffness shown over the final two thirds of the region between velocity reversals. The approximation of the initial contact stiffness where the whole interface is in a sticking state is still relatively poor. However the displacement range where this approximation is made is smaller than when using a single Jenkins element-which does give some improvement in agreement of the energy dissipated per cycle.

Number of Jenkins Elements	Energy Dissipated (J/Cycle)	Percentage Difference from Finite Element Result	(Percentage Difference of Evenly Spaced Jenkins Elements)
1	7.7350×10^{-3}	-22.5%	-26.51%
2	9.0454×10^{-3}	-9.42%	-12.00%
3	9.4203×10^{-3}	-5.66%	-11.75%
4	9.6408×10^{-3}	-3.46%	-8.13%

Table 4.5.1 Comparison of the energy dissipation per cycle when differing numbers of optimised Jenkins elements are used to represent the joint hysteresis of a finite element simulation with 19kN preload and 240Nm torque.

A major contributory factor to the poor agreement of the systems with two or less Jenkins elements was their inability to efficiently represent the small region of total sticking that occurred at the beginning of a cycle. This region was only a small fraction of the relative angular displacement between velocity reversal yet was responsible for a large underestimation of the amount of energy dissipated over approximately the first third of the cycle. The addition of a third Jenkins element to the model resulted in the initial, totally sticking contact, being properly represented. Benefits of this refinement over the two previous approximations were twofold. Firstly the amount of energy dissipated showed a significant improvement over the two-Jenkins-element model. The level of agreement was certainly approaching the region where it would be acceptable for a large range of engineering applications ($\approx 5\%$). It also offered better agreement than the best offered by up to five evenly spaced Jenkins elements.

The second benefit of this model was that it provided a very good representation of the fundamental features of the finite element model along with the improved agreement in energy dissipation. The initial stiffness, final stiffness, overall relative displacement, average stiffness, range of restoring torque as well as the energy dissipated per cycle are all effectively represented. This is achieved with only the 7 parameters

(Including the permanent spring) required to describe the whole response to dynamic loading.

The addition of a final Jenkins element, taking the total up to four resulted in a further improvement in the representative behaviour. However the improvement could no longer be said to be definitely worth the increase in computational cost of its implementation. Less than a 3% improvement in energy dissipation was achieved and the physical model was only modified slightly in the region shortly before velocity reversal. A large proportion of the error that was still present could also be as a result of the slight asymmetry of the finite element hysteresis that could not be represented by the Jenkins element simulations. Another potential source of error between the energy dissipated by the Jenkins element model and the finite element simulation could also be the noisy response that occurred just before velocity reversal-as reported in Chapter 3. This noisy region made parameterisation slightly approximate.

The computational cost of the virtually equivalent system with four Jenkins elements, like previous Jenkins element analyses, was at least two orders of magnitude less than the finite element simulation. The finite element simulation took about two hours to complete two cycles in comparison with less than a minute for the Jenkins element model. It should also be noted that the computational impact of adding a single Jenkins element to the time integration procedure was negligible. In a multi-jointed structure where contact may occur at several joints and between several degrees of freedom this kind of saving may become cumulative and significant. The use of three Jenkins elements over four could be advocated as it offers the best compromise between model complexity, computational cost and fidelity to the finite element simulation. What is also clear is that the parameters obtained from the optimally discretised hysteresis loop offer a much more efficient and accurate representation than those obtained through evenly spaced discretisation.

4.6) Modelling Variable Preload Simulations

A series of the finite elements simulations showed that when the bolt preload was increased above the control value of 19kN and the applied torque was maintained at 240Nm, hysteresis loops were generated that were essentially bilinear (Section 3.9). It is evident from the work carried out so far that a bilinear hysteresis can be represented by a system containing just one Jenkins element and one permanent spring. The aim was therefore to match this simple Jenkins element system to the bilinear hysteresis loops obtained using finite element analysis. In all the simulations that were converted into Jenkins element cases the parameters of the Jenkins elements were established by discretising the finite element hysteresis loops at the point at which the initial contact stiffness changed into the microslip stiffness. Evenly spacing the discretisation points on the hysteresis loops would not take advantage of the bilinearity that was to be exploited.

The overall behaviour in the series of tests was quite acceptable in all cases except that where the preload was the lowest i.e. 28.5kN. In this case the hysteresis loop is almost identical to the loop obtained when using a 38kN preload. This is the case to a degree in the finite element study. However, Figure 4.6.1 shows there was a small difference in their behaviour that does not show up clearly when replacing the finite element analysis with a second order Jenkins element simulation. One particular reason for the disagreement between the 28.5kN preload case and the bilinear model used to represent it is that a single Jenkins element combined with a permanent spring did not represent the overall hysteresis accurately enough. The finite element study reveals that despite appearing to be essentially bilinear the response was best represented by three segments. This resulted in a significant under-prediction of the amount of energy dissipated by the model (Table 4.6.1).




Figure 4.6.1 *Hysteresis loops generated using Jenkins elements whose parameters were extracted from the finite element hysteresis loops obtained with varying preloads and 240Nm applied torque.*

Table 4.6.1 Shows that the agreement in all the solutions is relatively poor considering that with preloads of 38kN and above the hysteresis loops are very bilinear. The main reason for this poor agreement is the difficulty in extracting accurate parameters for the Jenkins element model when the finite element model displays relatively large amounts of drift. Figure 4.6.2 shows that in such a case agreement in stiffness at all parts of the loop and a distinct point from which to extract the model parameters could not guarantee good agreement. The size of the offset is significant in causing the disagreement, but most notably so when it is large relative to the overall range of angular displacements. Quite clearly Table 4.6.1 Shows the best agreement offered by the 57kN preloaded case and this is the simulation in this particular series of tests that demonstrated the smallest relative offset-as can be seen in Figure 3.9.3.

Bolt Preload (kN)	Energy Dissipated by Jenkins Elements (J/Cycle)	Energy Dissipated in Finite Element Simulation (J/Cycle)	Percentage Difference from Finite Element Result
28.5	1.782×10^{-3}	2.1281×10^{-3}	-16.26%
38	1.8500×10^{-3}	2.197×10^{-3}	-15.79%
57	2.2537×10^{-3}	2.429×10^{-3}	-7.22%
76	2.1312×10^{-3}	2.405×10^{-3}	-11.38%
85.5	1.8279×10^{-3}	2.2159×10^{-3}	-17.51%

Table 4.6.1 Comparison of the energy dissipated between the finite element simulations of varying preload and 240Nm torque with optimised simulations generated by a single Jenkins element and permanent spring.

One possible way of getting around the problem of the offset hysteresis loops is to extract parameters from both the positive and negative velocity sections of the loop and then average the parameters obtained in both cases. However, the problem does not seem to lie in the case of the method of using Jenkins elements or the means used to extract them. Well-centred loops can be matched very well as previous sections have shown and the 57kN preload case also illustrates. As has already been discussed the reasons for the offset, in simulations where the finite element simulation is not exhibiting gross slip, are thought to lie within the finite element contact algorithm. None of the finite element cases with varying preload show any other behaviour than microslip and so drift is highly unlikely to be exhibited without some sort of numerical anomaly in the contact algorithm. It should be recognised that in certain simulations and experimental tests, where contact is less determinate and drift may occur legitimately, these simulations have shown a limitation of the Jenkins element method of representation.




Figure 4.6.2 *Hysteresis loops generated by Jenkins elements and the finite element simulation showing offset of the finite element result and poor agreement between both cases.*

Despite these limitations there were some interesting phenomena displayed in the parameter extraction and behaviour of the attempts to model the effect of varying preload. It has been discussed (Section. 3.9) how each of the finite element cases for 38kN preload and above represented the same portion of the contact interface in sticking and/or sliding contact. The parameter extraction for all the Jenkins element cases used to represent these finite element simulations backed up this finding. In all the cases the stiffness in the initial segment, where only sticking contact was assumed to take place, was between 20.69×10^6 and 21.61×10^6 Nm/rad. Even the small deviation displayed in these results could be justified by measurement errors in the extraction process. The second segment where a small amount of microslip existed, showed a stiffness range from between 9.31×10^6 to 10.11×10^6 Nm/rad-which could be accounted for by the same

measurement errors. This value was also well below the stiffness of $2.45 \times 10^6 \text{Nm/rad}$ that existed just before gross slip. It was therefore possible to conclude that in all the cases the same region of microslip was being represented, and gross slip was not occurring. Drift should therefore be minimal in response to a harmonic input torque of constant magnitude.

Figure 3.9.1 associated with the contact pressure at the joint interface reveals that a proportional increase in contact pressure accompanies a change in bolt preload. The normal force extracted for each preload case from 38kN upwards showed an increase that was nearly of the same proportion as that displayed by the preload. The exception to this rule was the case of a 28.5kN preload and the shortcomings of using bilinear representation in this particular case have already been discussed. One reasonable conclusion from the above findings is that if the pressure distribution at the joint interface is known to be of a particular type it should be possible to use Jenkins elements to predict the change in response to a varying preload. This finding is only valid if the contact pressure distribution contributes to a hysteresis loop with a constant number of segments (i.e. regions of constant stiffness) and microslip is the only dissipative process that exists at the contact interface. If the preload and normal forces were decreased then the model must have enough Jenkins elements so that it can represent new regions of microslip that may appear on the contact interface. Bilinearity is a useful assumption, but will only hold if the preload is increased and the response either remains bilinear or becomes totally elastic.

4.7) Prediction of Finite Element Hysteresis Using Jenkins Elements

The behaviour of the Jenkins elements has so far been shown to replicate all of the main features of the finite element simulations with the only exception being drift of the hysteresis loop. Each individual simulation performed by the finite element method has been represented by a single hysteresis loop where the parameters for an equivalent Jenkins element model have been extracted. The case of varying preload has shown that a satisfactory relationship can be drawn between the normal force applied to the Coulomb element and the amount of preload applied to the bolt in the finite element simulation.

This section shows that if parameters for the Jenkins elements are extracted over a large enough range of angular displacements then, at a given preload, the response of the system to inputs of varying magnitude can be predicted from one set of parameters. The limitation of extrapolating the results from a small applied torque to a model where the torque was significantly larger is also shown.

A model where the parameters were extracted from a finite element simulation with 19kN preload and maximum applied torque of 80Nm was used. The response of this simulation was essentially bilinear meaning that only a single Jenkins element could be used to describe the response of an initially sticking interface and a small amount of microslip.

Even without a comparison of the energy dissipated it can be seen in Figure 4.7.1 that the bilinear Jenkins element model defined at small torques does not represent the results achieved at higher torques well. At larger torques than those used to define the bilinear result an increased and progressive degree of microslip is visible. The one-Jenkins-element model that is so efficient at modelling the small torque response is incapable of representing the softening behaviour and increased energy dissipation found when more microslip occurs. Even the range of angular displacement is much smaller than that required to represent the response to 240Nm of applied torque. This is easily explained by the bilinear loop's increased stiffness that is only exhibited by the large

torque result at the beginning of the cycle between velocity reversals. As the torque was increased the system softens and so more relative displacement for a given applied torque occurred than was seen in the extrapolated loop.




Figure 4.7.1 *Demonstration of the limitation of extrapolating small torque, bilinear Jenkins element simulations to represent large torque, large microslip finite element results.*

Although a bilinear loop was capable of predicting the response at torques below those used in parameter extraction, a different loop was required to describe the whole range of phenomena that the joint was capable of exhibiting. The control loop of a 19kN preload with parameters extracted from a simulation where the maximum applied torque magnitude was 240Nm was used. At this applied torque before velocity reversal, when the torque was at its largest value, the contact interface was on the point of gross slippage. No more microslip behaviour could be exhibited, and the interface over the

course of a cycle displayed both the maximum possible and minimum possible contact stiffness. Initially the response was calculated where the maximum torque magnitude was reduced from 240Nm to 80Nm and 200Nm. These values were chosen on account of the fact that the finite element simulations associated with them displayed bilinearity and the onset of further microslip after bilinearity respectively. In all the tests the parameters used were those extracted for four Jenkins elements when the discretisation of the hysteresis loop was visually optimised.

Once the tests for reduced maximum torque had been performed the torque was increased to values of 260 and 280Nm. This represented the upper values of torque that could be applied before the finite element response broke down and gross slip was displayed by the contact interface. In the range of applied torque between 240Nm and 280Nm no further decrease was evident in the stiffness before velocity reversal. The results from the reduction and increase in applied torque in the Jenkins element simulations are shown in Figure 4.7.2. A comparison of the energy dissipated by the time integration analysis and the finite element response is given in Table 4.7.1.

Visually the results in Figure 4.7.2 are good. The significant behaviours of each hysteresis loop have been well represented by the Jenkins element model. By increasing the torque to 260Nm and 280Nm the good agreement is certainly predictable. Each Jenkins element makes the transition from sticking to sliding due to the relative angular displacement after a velocity reversal. These limiting angular displacements have already been discovered for torques up to 240Nm, assuming regions of constant stiffness, by the way the finite element response was discretised. After the final Jenkins element made the transition to sliding in the final segment of the range between velocity reversals the only restoring force was due to the permanent spring in the model. Analysis of the finite element response in Chapter 3 showed that the stiffness before velocity reversal, where only a small region was still in sticking contact, did not change in the region of applied torque of 240Nm to 280Nm. Therefore no additional Jenkins elements are required to describe the region over which the results are extrapolated. As the previous part of the extrapolated cycle is common to all the cycles which cause sliding contact across all the

Jenkins elements the extrapolated zone is simply an extension of the permanent spring and easy to predict.




Figure 4.7.2 Figure showing the ability, of a second order model with 4 Jenkins elements and a permanent spring, to predict the response of a finite element simulation when a different magnitude of torque was applied.

The significance of relative displacement can be seen when considering the loops with smaller torques than 240Nm. Again, as the state of the Coulomb element is dependent on relative angular displacement after velocity reversals, the hysteresis loops associated with smaller torque should display regions of constant stiffness over the same range of angular displacements. The point at which velocity reversal occurs is dependent on the point at which the applied torque is reversed and so the profile that extends from velocity reversal to velocity reversal for higher torque follows the same "path" but is cut off at an earlier point. As long as the same Jenkins element parameters are used to describe the high torque simulations as the low ones, the smaller torque hysteresis loops

should be identical up to the point of velocity reversal if the starting angular displacements are co-ordinated.

Input Torque (Nm)	Energy Dissipated by Jenkins Element Simulation (J/Cycle)	Energy Dissipated in Finite Element Simulation (J/Cycle)	Percentage Difference from Finite Element Result
80	1.6564×10^{-4}	2.6900×10^{-4}	38.42%
200	2.1337×10^{-3}	2.5006×10^{-3}	14.67%
260	1.4945×10^{-2}	1.5402×10^{-2}	2.97%
280	2.0244×10^{-2}	2.2104×10^{-2}	8.84%

Table 4.7.1 Comparison of the hysteresis loops generated when the Jenkins element parameters obtained for the case of a visually optimised 19kN preload and 240Nm applied torque are extrapolated for different input torques.

Table 4.7.1 does show that there is a considerable deterioration when going from a large torque of 240Nm to a smaller torque of either 200Nm but particularly 80Nm. The reason for this can be attributed to the difficulty in accurately extracting the parameters for relatively small displacements from a hysteresis loop that represents everything from total sliding to a point only slightly before the onset of gross slip. This effect can be seen from the fact that from Figure 4.7.2 a good visual agreement is found for the smaller torques. However, if the scale of the graph was altered to show only the comparison of 80Nm torque, then the visual agreement would be shown to be far less satisfactory. When the parameters were extracted directly from a simulation with 80Nm of applied torque the agreement was much better, as can be seen in Figure 4.7.1.

Response quality from the Jenkins elements is largely due to the resolution of the finite element simulation or experimental test. The extrapolation process will not work adequately unless a suitable number of visually-optimised, parameter-extraction points are used over the whole range of stiffness exhibited by the contact interface. The best

solution for identifying parameters over a wide range of torque is to apply both a large torque to the point of macroslip and also a small torque where a microslip response dominates. The results can then be combined. Parameters extracted from the small amplitude response can be used to describe the beginning of the cycle and those from a state close to macroslip used to describe the end of a cycle when the response amplitude becomes large enough. The resolution of the tests performed in different ranges will then allow accurate parameters to be extracted for the entire cycle.

4.8) Free Vibrations of Jenkins Elements

Much useful information about the behaviour of Jenkins elements in a second order system has already been established. In all of the examples used so far the simulations have been compared with the response of a finite element model of a joint to forced harmonic vibrations. The response of Jenkins elements to free vibrations also illustrates their behaviour, but also gives some useful information as to their suitability as a tool for modelling the behaviour of the high order, finite element, isolated joint.

A series of simulations were performed on a second order model with 4 Jenkins elements and a permanent spring where the parameters for these components were the same as those used in the visually-optimised case. The reason why this configuration of model was chosen was that it offered the best agreement with the control finite element simulation in all the key areas; average stiffness, stiffness immediately before and after velocity reversals and energy dissipated over the course of a cycle. Each simulation consisted of giving an initial displacement before releasing the system to exhibit free vibrations. A range of initial displacements was used ranging from 1×10^{-5} to 5×10^{-5} rad. This range of displacements was chosen to give a variety of responses from that associated with small displacements, to a maximum initial amplitude approximately equivalent to that required to cause gross slip in the finite element model. The oscillations were allowed to continue until a steady state was adjudged to have been reached.

It can be seen in Figure 4.8.1 how the vibration amplitude dies away over relatively small number of cycles. This happens because energy is being dissipated from the system through frictional sliding. During forced oscillations more energy was put into the system through frictional sliding. In the free response this source of energy input does not exist, and the response being frictionally damped reduces to a steady state level. As has been found previously the decay envelope of the system response is linear. This is a feature of friction damped systems as opposed to those with viscous or proportional damping. The linear decay is most evident in the case of an initial displacements of 1×10^{-5} and 3×10^{-5} rad.




Figure 4.8.1 Free vibration response of 4-Jenkins-element model to initial displacement.

Within the model itself, as the energy is dissipated from the system, the response amplitude decreases, and less Jenkins elements in each cycle make the transition from sticking to sliding. Eventually when the response has been damped to a certain level, no Jenkins elements have enough displacement to cause the transition from sticking to sliding. The only restoring force in the system is provided by the permanent spring and the four springs attached to the Jenkins elements. When this occurs the Jenkins element system has become conservative and causes the steady state vibrations seen at the latter end of the time span of the analysis. The frequency of these vibrations is equivalent to that of the fully stuck contact interface. This can be found as the root of the totally stuck contact stiffness ($22.5 \times 10^6 \text{ Nm}$) divided by the inertia $0.0016 \text{ (m}^4\text{)}$ and is $119 \times 10^3 \text{ (rad/s)}$ or $18.9 \times 10^3 \text{ Hz}$. The amplitude of the conservative vibrations is also significant as it corresponds in all of the cases to the range of angular displacement, after a velocity

reversal, before the first Jenkins element makes the transition from sticking to sliding (1.5×10^{-6} rad).




Figure 4.8.2 Comparison of the hysteretic restoring force for the 4-Jenkins-element system subjected to initial displacements and free vibrations.

When comparing the hysteretic restoring torque the response also dies away to a constant steady state oscillation (Figure 4.8.2). The phase difference between the simulations with different initial displacements is caused by the fact that when the amplitudes are large the contact interface is less stiff and so oscillates at a lower frequency. As the amplitudes get smaller, due to frictional dissipation, the contact becomes stiffer on average over a cycle. This causes the resonant frequency to increase and the period of the oscillation to decrease until the steady state is reached.

Almost all of the free vibration tests were in the bilinear stiffness region after only a single cycle. Only in the situation where the initial displacement was 3×10^{-5} rad did a

cycle take place after the initial oscillation where the amplitudes were still large enough to cause two Jenkins elements to make the transition from sticking to sliding. The bilinearity of the oscillations can be seen in the comparison of hysteresis loops of Figure 4.8.3 All of the hysteresis loops show the response settling into an elastic oscillation where the energy dissipated becomes zero. The loops also show how less energy is dissipated per cycle as the oscillations become smaller until the steady state is reached.




Figure 4.8.3 Comparison of the hysteresis demonstrated by the free vibration of the second order system with initial displacement (circled).

Instead of nesting, the loops, as is the case with forced oscillations of varying magnitude, are offset to varying degrees. The restoring torque is cyclic about zero, so the offset is entirely a result of the angular displacement. As the amount of energy dissipated is dependent on the amplitude of vibration, different simulations reach the point where steady state oscillations occur at different points in their cycles. When that point is reached the steady state oscillations of amplitude 1.5×10^{-6} rad start regardless of whether

the point of onset is symmetrical about the origin or not. Other parts of the cycle are offset for the same reason, and are indicative of a hysteretic response i.e. the point at which enough energy has been dissipated to prevent a certain Jenkins element making the transition to sliding ever again.

In the case of uniform pressure distribution, it has been shown that for large parts of the cycle the response can be elastic. However in the case of preload, energy dissipation commences almost as soon as any torque is applied to one of the clamped components. Consequently the situation above where conservative oscillations are set up would not apply to the case where a 19kN bolt preload is in evidence. Frictional dissipation would exist almost for as long as there was movement between the two surfaces. In the case of the uniform loading however, these elastic oscillations could occur quite reasonably as small amplitude vibrations would not change the state of the entirely sticking contact interface. In reality the oscillations would die out eventually, but an extra damping component would need to be included in the representative model to take account of this. This shows that Jenkins elements alone may not be suitable for modelling the small amplitude, or free response of interfaces where the contact pressure decays to very small values.

This particular demonstration of free vibration also helps to confirm that the source of offset in the seemingly symmetrical finite element simulations is likely to be due to a change (even small) in the conditions at the contact interface. Certainly this is the case when moving from an implicit solver to an explicit solver. The symptoms demonstrated in the free vibration cases above correspond to those displayed in the finite element simulations. These symptoms of changes in contact condition are a symmetric restoring force, but an offset angular displacement, and a hysteresis loop that is then consistent about this offset location. Again, all are consistent with the offset behaviour encountered in the finite element simulations with forced oscillations.

Chapter 5. The Bouc-Wen Method

5.1) Summary of the Bouc-Wen Method

The Bouc-Wen model is capable of describing hysteretic restoring forces and material behaviour with particular relevance to random excitation. The hysteresis described by the Bouc-Wen model can also be applied to a harmonically excited system as well as those driven by a random input. The principal benefit offered by the Bouc-Wen model was that it was capable of representing a smooth transition from an initial elastic restoring force to a contact interface where the stiffness had been reduced to nothing.

Computationally this smooth transition was expensive to approximate using Jenkins elements. The large number of elements required to provide a gradual loss of stiffness of a system lead to costly time integration solutions. At every time increment it was also necessary to establish whether a given Coulomb element had made the transition from sticking to sliding or vice versa. In contrast the Bouc-Wen model was capable of representing the gradual loss of contact stiffness with just the 4-parameters constituting the describing function.

The impact of each of the 4 parameters on the form that the restoring torque took has been shown. Although the totally-sticking stiffness was the only parameter that was directly identifiable from the describing function, the remaining model parameters could be scaled to provide a fixed value of the yield torque when all stiffness of the contact interface was lost. The model parameters could also be used to tune a given hysteretic response to demonstrate smooth softening behaviour or even hardening of the contact stiffness before the later onset of softening due to microslip. In this flexibility of hysteretic profile and the efficiency of providing a gradual loss of contact stiffness, the Bouc-Wen model had its main benefits.

To examine the effectiveness of the Bouc-Wen model in describing bolted-joint hysteresis, the results from the finite element simulations were matched using numerical

integration of the Bouc-Wen equation. The simulations that were analysed demonstrated a wide variety of phenomena that existed in the joint from small microslip bilinearity to a loss of contact stiffness associated with the imminent onset of macroslip. A procedure to extract the parameters of the Bouc-Wen model proved effective in identifying the initial contact stiffness immediately after velocity reversals. The average stiffness of the contact interface over a complete loading cycle was well estimated in all cases by this parameter extraction procedure.

In cases where the microslip was not extensive on the contact interface the Bouc-Wen model accurately predicted the energy loss due to joint hysteresis. When the amount of microslip became large, the predicted energy dissipation by the Bouc-Wen model was found to be on the whole larger than that demonstrated during finite element tests. The overall profile of the hysteresis loop did not relate closely to the loops generated by the finite element simulations. The stiffness of the loops at equivalent angular displacements were generally unreliable with the exception at the critical point before immediately before velocity reversal.

Further tests were also conducted to establish the effectiveness of a set of Bouc-Wen model parameters in determining the response at conditions different to those for which they were extracted. The method was shown to have potential when predicting the response to lower magnitudes of applied torque, but was highly unreliable when extrapolating results obtained at these lower torques to greater magnitudes. The limitation of the Bouc-Wen model was its generally poor agreement with the reference hysteresis loop except at velocity reversals. Better parameter extraction would reduce the impact of these issues, and still allow the computationally efficient describing function to be utilised in describing the transition from small-scale microslip to macroslip.

5.2) The Bouc-Wen Model

The Bouc-Wen model derives its name from the original presentation by Bouc (1967) and further developments and generalisation of the model by Wen (1976, 1980 & 1989). Motivation for the model came for the need to model a smooth restoring force with hysteretic properties. In particular the model had to be capable of representing the restorative force in situations where the excitation force was random as opposed to periodic.

The model is described by Wen (1980) as follows:

$$Q(\theta, \dot{\theta}) = g(\theta, \dot{\theta}) + z(\theta, \dot{\theta}, t) \quad (5.2.1)$$

In Wen's formulation θ is replaced by x , but here the angular displacement of the joint is being investigated rather than the lateral stiffness. Component g of the hysteresis force is a linear or nonlinear function of the instantaneous angular displacement and velocity. The component z is a state variable, function of the time history of θ and therefore the hysteretic component. In this study it is assumed that the hysteresis of the joint under investigation can be represented solely by the hysteretic term z .

Wen related θ to z using the following nonlinear equation:

$$\dot{z} = -\alpha|\dot{\theta}|z|z|^{n-1} - \beta\dot{\theta}|z|^n + A\dot{\theta} \quad (5.2.2)$$

This formulation provides a hysteretic restoring torque that is totally smooth between velocity reversals. To achieve a similar degree of smoothness the number of Jenkins elements required would be extensive. Each Jenkins element is defined by two parameters namely the stiffness and resistance of the Coulomb element. The Coulomb element also requires a coefficient of friction that, if not uniform, can count as a further parameter to be defined in the model. As Chapter 4 has shown each additional Jenkins

element adds to the computational cost of the model. One advantage of the Bouc-Wen model is that a totally smooth hysteresis loop can be defined by just 4 parameters: α , β , n and A .

Another way of viewing the Bouc-Wen restoring torque is shown below. The signs used are dependent on the power of n used to represent the hysteresis, and the part of the loop is being described in terms of the sign of θ and z :

$$\frac{dz}{d\theta} = A \pm (\alpha \pm \beta)z^n \quad (5.2.3)$$

Table 5.2.1 shows the form that each branch of the hysteresis loop takes depending on whether the parameter n is odd or even. It is possible to integrate each branch of the loop in closed form, but not as a whole cycle.

	Odd n	Even n
$\dot{\theta} > 0, z > 0$	$A - (\alpha + \beta)z^n$	$A - (\alpha + \beta)z^n$
$\dot{\theta} < 0, z > 0$	$A + (\alpha - \beta)z^n$	$A + (\alpha - \beta)z^n$
$\dot{\theta} < 0, z < 0$	$A + (\alpha + \beta)z^n$	$A - (\alpha + \beta)z^n$
$\dot{\theta} > 0, z < 0$	$A - (\alpha - \beta)z^n$	$A + (\alpha - \beta)z^n$

Table 5.2.1 *Combination of Bouc-Wen parameters depending on the branch of the hysteresis loops and value of n .*

The Bouc-Wen model is capable of describing everything from elastic-perfectly plastic behaviour to a torque that demonstrates an immediate onset of plasticity. Each of the four parameters is used to define the overall response of the restoring torque and can be identified and tuned to match specific characteristics of a wide variety of hysteresis loops. The initial stiffness of the restoring torque, is represented by the parameter A . In

physical terms this is equivalent to the contact stiffness of the interface before any microslip has taken place. By considering the defining Equation 5.2.2 it can be seen that when z (the restoring torque) is zero the equation reduces to:

$$\dot{z} = A\dot{\theta} \quad (5.2.4)$$

If the assumption is made that the restoring torque is zero when the angle of rotation between the two components is also zero, as is found at the onset of the movement between the two surfaces, then the above relationship holds. A then defines the maximum torsional stiffness possible in the joint. Any subsequent path of the hysteresis loop will lie inside an imaginary line equivalent to the stiffness A starting at a point of velocity reversal.

A second defining relationship can be derived from Equation 5.2.3 The joint can be considered perfectly plastic when the gradient of the hysteresis loop is reduced to zero. At this point the yield load or torque (z_y) will have been reached:

$$\frac{dz}{d\theta} = A \pm (\alpha \pm \beta)z_y^n = 0 \quad (5.2.5)$$

In the case of the initial loading slope (starting from $\theta=0$, $z=0$ and continuing until the first velocity reversal), whether n is odd or even, Equation 5.2.5 reduces to give:

$$z_y = \left(\frac{A}{(\alpha + \beta)} \right)^{\frac{1}{n}} \quad (5.2.6)$$

Consequently it is possible to get an idea of the scale of the hysteresis loop. It is governed by the magnitude of $(\alpha+\beta)$ relative to A . What follows is a series of example hysteresis loops created using the Bouc-Wen model showing the influence of varying the 4 parameters that define its profile.




Figure 5.2.1 Hysteresis loops with a yield torque of 0.5Nm and initial stiffness of 1Nm/rad calculated using increasing powers of n .

Initially the impact of a variable value of n is demonstrated. In all of the following examples all of the other parameters are maintained at fixed values or to give a constant yield torque (z_y). All of the examples were obtained from the numerical integration of Equation 5.2.3 using a second order central difference method. Figure 5.2.1 illustrates the case where n was increased from 1 to 9. The initial stiffness was maintained at a constant value and α and β were equated to each other and scaled to give a constant yield torque of 0.5Nm. Although the initial stiffness stays constant at 1Nm/rad the amount of the hysteresis cycle that maintains this value varies according to n . With low values of n , the hysteresis loop loses its initial stiffness A quickly. As n is increased the stiffness A is maintained for increasing proportions of the cycle and the transition to the yield torque becomes more rapid. When n tends to infinity, the hysteresis loop tends to the properties associated with an elastic-perfectly plastic system. A similar “ n ” parameter was used by

Ramberg and Osgood (1943) in their less flexible model used to describe hysteresis attributed to material plasticity.




Figure 5.2.2 Comparison of the effect of the α and β parameters in the Bouc-Wen model with constant yield torque and initial stiffness.

A significant characteristic of the hysteresis associated with bolted joint used in the finite element analysis is the softening behaviour of the joint interface. Some nonlinear dynamic systems demonstrate hardening behaviour, and this is equally possible in the event of variable normal forces in a joint. The Bouc-Wen model is capable of modelling both hardening and softening conditions. Tailoring of the hysteresis loop to meet the hardening and softening conditions is performed by altering the relative size of the parameters α and β . Figure 5.2.2 shows that with a small value of α relative to β the hysteresis loops show hardening behaviour leading to softening characteristics later in the cycle. Assuming a constant value of $\alpha + \beta$, as the size of α increases relative to β the

hardening behaviour disappears and the hysteresis loop demonstrates softening from one peak torque value to the other.

Figure 5.2.3 shows the impact of changing the initial stiffness of the hysteresis loop. All of the other parameters have been manipulated to maintain the same shape characteristic and yield torque of the hysteresis loop. Increasing A results in an increased gradient, and therefore stiffness, after the velocity reversal points in the hysteresis loops. As a value of $n=3$ was used, the stiffness is shown to be largely maintained until after the restoring torque has changed sign during each half cycle.




Figure 5.2.3 Hysteresis loops generated using the Bouc-Wen model with constant yield torque, $n=3$, $\alpha=\beta$ and varying initial stiffness parameter A .

Also of note in Figure 5.2.3 is the fact that the hysteresis loop with the smallest value of A does not reach its yield torque (z_y). The consequence of this is that the stiffness

of the hysteresis loop at velocity reversal is greater than zero. Upon velocity reversal the stiffness is instantly back to the value A and traces an identical path in the opposite direction. This demonstrates the hysteretic property of the Bouc-Wen model admirably. The shape of the hysteresis curve is defined by the current velocity (sign) and angular displacement. It is also partially defined by the angular displacement relative to the point at which velocity reversal occurred. This characteristic is not altered by the amount of yielding that has taken place since the previous velocity reversal. Even when the smallest initial stiffness is used, and the yielding force is not met over the angular displacement range in question, the same initial stiffness A is in evidence after each velocity reversal.

The final model parameter that is illustrated is the yield torque z_y . Whilst not one of the four parameters (A , n , α and β) used to define a particular Bouc-Wen model, it is a crucial quantity in terms of model behaviour. Figure 5.2.4 shows how the size of the yield torque influences the hysteresis loops that exist when oscillating between angular displacements of ± 1 rad. As the yield torque increases the range of displacement is not large enough for a totally plastic-macroslip response to be in evidence. Progressively more stiffness is retained in the joint at each velocity reversal as the transition to fully yielded behaviour has developed to a smaller degree.

However, the yield torque is only an implied parameter of a given Bouc-Wen model, and is more difficult to isolate from a hysteresis loop generated by another method (finite element, Jenkins elements or experimentally). Measuring the displacement where the hysteresis loop becomes flat is much harder than identifying the initial stiffness A . It is also defined by the ratio of A to $(\alpha + \beta)$ and so a constant yield torque z_y can exist for a range of these three parameters. Each of the other four parameters used in the Bouc-Wen equation (Equation 5.2.2) can be isolated explicitly and are therefore more integral to reproducing hysteresis curves than the yield torque.




Figure 5.2.4 *Hysteresis loops with varying yield torque, generated using the Bouc-Wen model with $n=3$, $A=1$ and $\alpha=\beta$.*

5.3) Bouc-Wen Parameter Identification

Unlike the Jenkins element model, it was not possible to characterise all of the parameters of the Bouc-Wen model directly from the hysteresis loop. Instead, the describing differential equations had to be partially solved before the data from the hysteresis loop (either generated through finite element analysis or experiment) could be used to finalise the model.

One parameter, A could be identified directly from the hysteresis loop that was being modelled. As it represents the initial stiffness of the loop, A was extracted by measuring the gradient at the onset of vibrations or after every velocity reversal. Assuming a steady state had been reached meant that all of the Bouc-Wen parameters including the stiffness after velocity reversals remained constant over time.

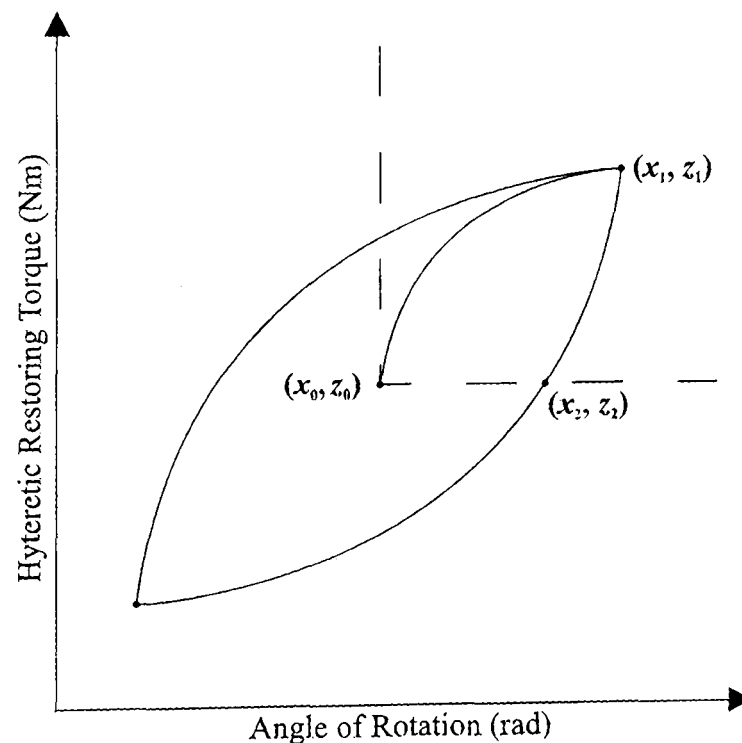


Figure 5.3.1 *Defining points of a hysteresis loop used to extract the Bouc-Wen parameters α and β .*

The parameter extraction process relied upon an assumption that the value of n in Equation 5.2.3 was no bigger than 2. When n was equal to 1 or 2 it was possible to solve

the governing differential equations in the first two branches. From these solutions the parameters α and β were established. Further powers of n did exist as analytical solutions, but it was beyond the scope of the work carried out here to develop them. Instead, α and β at higher powers of n could be scaled by assuming constant A and yield torque z_y in Equation 5.2.6.

The procedure for identifying the parameters first required the development of analytical solutions for Equation 5.2.3 in the branch of the hysteresis loop extending from (x_0, z_0) to (x_1, z_1) . In this branch, the differential equation linking x and z was given by the first row in Table 5.2.1 Full solution of the differential equations required the boundary conditions to be applied. For the first branch these boundary conditions were taken as (x_0, z_0) . The solutions to the differential equations for $n=1,2$ with these boundary conditions were found to be:

$$z = \frac{A[1 - e^{-(\alpha+\beta)x}]}{(\alpha + \beta)} \quad \text{for } n = 1. \quad (5.3.1 \text{ a \& b})$$

$$z = \sqrt{\frac{A}{(\alpha + \beta)} \frac{e^{2\sqrt{(\alpha+\beta)x}} - 1}{e^{2\sqrt{(\alpha+\beta)x}} + 1}} \quad \text{for } n = 2.$$

Both equations are clearly functions of $(\alpha+\beta)$, x , and A , yet by inserting the values of x and z from (x_1, z_1) in Figure 5.3.1 they reduced to functions of $(\alpha+\beta)$ only.

In the second branch of the hysteresis loops, defined from (x_1, z_1) to (x_2, z_2) , the solutions of the governing differential equations were:

$$z = \frac{[A + (\alpha - \beta)z_1]e^{(\alpha - \beta)(x - x_1)} - A}{(\alpha - \beta)} \quad \text{for } n = 1.$$

(5.3.2 a & b)

$$z = \frac{\sqrt{\frac{A}{(\beta - \alpha)}} \left(\frac{\sqrt{A} + \sqrt{(\beta - \alpha)}z_1}{\sqrt{A} - \sqrt{(\beta - \alpha)}z_1} \right) e^{2\sqrt{A(\beta - \alpha)}(x - x_1)} - 1}{\left(\frac{\sqrt{A} + \sqrt{(\beta - \alpha)}z_1}{\sqrt{A} - \sqrt{(\beta - \alpha)}z_1} \right) e^{2\sqrt{A(\beta - \alpha)}(x - x_1)} + 1} \quad \text{for } n = 2.$$

When $n=1$, and A , x_1 and z_1 were all known, the restoring torque (z) became a function of only $(\alpha - \beta)$. Under similar conditions when $n=2$ the restoring torque became a function of $(\beta - \alpha)$. The values of x and z for the hysteresis loop being modelled were identified at (x_2, z_2) . For $n=1, 2$ this left defining equations for each branch where the only unknowns in the first branch were $(\alpha + \beta)$, and $(\alpha - \beta)$ and $(\beta - \alpha)$ respectively in the second branch. For a given value of n these pairs of equations could then be solved simultaneously to establish both α and β independently.

5.4) Equivalent Bouc-Wen Models

To study the effectiveness of the Bouc-Wen model for representing the hysteresis found in bolted joints, a selection of finite element simulations were replicated. These simulations were selected to illustrate the full range of microslip behaviour. In all cases the finite element joint was subjected to a 19kN preload before harmonic torques of 3 different magnitudes were applied. The first torque that was considered had a maximum magnitude of 80Nm and was chosen as it demonstrated the bilinear stiffness associated with only a small region of microslip at the contact interface. The remaining applied torques were of magnitude 240Nm and 280Nm. In the case of 240Nm applied torque the amount of microslip had progressed until the most of the interface had made the transition from sticking to sliding. Similarly the 280Nm applied torque resulted in the contact interface being on the point of macroslip, with the smallest possible residual stiffness before velocity reversal.

For the case of the 80Nm applied torque, the peak displacement and restoring torque were matched with the finite element analysis and the parameters were extracted by the method outlined in Section 5.3. The initial stiffness was varied according to the different values extracted from the finite element simulation. Due to the change in contact conditions, brought about by the use of an explicit solver after the initial implicit static solution, the stiffness of the initial loading curve in the finite element simulation was different to that found after every subsequent velocity reversal. From the initial loading cycle, the maximum stiffness of the contact interface (identified as parameter A in Equation 5.2.4) was taken as $15.1 \times 10^6 \text{Nm/rad}$. In all of the other cycles across all of the finite element tests the totally sticking contact stiffness was found to be between $20 \times 10^6 \text{Nm/rad}$ and $22 \times 10^6 \text{Nm/rad}$ dependent on measurement inconsistencies. An average value of $21 \times 10^6 \text{Nm/rad}$ was therefore used to identify the parameters in a subsequent test. For completeness, a value of $1.73 \times 10^6 \text{Nm/rad}$ was also used for A as it lay between the previous upper and lower estimates. All of the tests were carried out with parameter $n=1$. One additional test was also performed at $n=2$ using the largest value of stiffness.

Finite Element Torque (Nm)	Initial Stiffness Parameter A (Nm/rad)	n	Energy Dissipated (J/Cycle)	Energy Dissipated in Finite Element Test (J/Cycle)	Percentage Difference
80	15.1×10^6	1	2.4883×10^{-4}	2.69×10^{-4}	-7.50%
80	17.3×10^6	1	2.6121×10^{-4}	2.69×10^{-4}	-2.90%
80	21×10^6	1	2.7510×10^{-4}	2.69×10^{-4}	2.27%
80	21×10^6	2	3.3924×10^{-4}	2.69×10^{-4}	26.17%

Table 5.4.1 Comparison of the energy dissipated by the Bouc-Wen model when matched using different parameters to the finite element model with 80Nm maximum applied torque.

It can be seen (Table 5.4.1) that in the case of the two largest initial stiffnesses, when $n=1$, that the agreement in energy dissipated between the two hysteresis models was very good indeed. This can be considered in conjunction with the defining model parameters. From the identification process it was already known that the peak displacement at velocity reversals offered good agreement. The stiffness A was also identified directly from the finite element hysteresis loop so offered good correlation between the two models. Figure 5.4.1 shows how well further defining physical characteristics are replicated by the Bouc-Wen model.

All of the Bouc-Wen simulations offer reasonable agreement with the stiffness of the finite element loop before velocity reversal with the exception of the case where $n=2$. There is also reasonable agreement of the restoring torque in the steady state cycles, although consistently in all of the Bouc-Wen models it was slightly overestimated. The cause of the overestimation was a slight change in the response of the Bouc-Wen model between the first cycle and all following cycles. The points of velocity reversal were required for parameter extraction and were therefore replicated exactly in the initial loading cycle. At the point of the second velocity reversal the response of the Bouc-Wen model tended to the steady state shown in Figure 5.4.1. These subsequent velocity

reversals were identified internally by the Bouc-Wen model and not constrained by the points extracted directly from the finite element hysteresis loop in the way that the initial loading cycle was.




Figure 5.4.1 Comparison of the hysteresis generated by the finite element model with maximum input torque of 80 Nm with simulations using the matched Bouc-Wen model.

Notably Figure 5.4.1 shows that the overall shape of the hysteresis loop did not agree terribly well with the Bouc-Wen models. The finite element model lost its initial stiffness at an abrupt point whereas the Bouc-Wen model illustrates a much more smooth transition from the initial stiffness to the contact stiffness before velocity reversals. The discrepancy in shape was not matched by a discrepancy in energy dissipated by the corresponding hysteresis loops as the Bouc-Wen loop wove in and out of the finite element loop. The aggregate effect is that parts of the cycle where correspondingly more energy

was dissipated than the finite element analysis are counteracted by parts of the loop where less energy was dissipated.

Overall the Bouc-Wen model in the 80Nm case offers good agreement in all of the main areas that matter. A good match was found in all but one case of the energy dissipated, the points of velocity reversal and the stiffness exhibited before and after steady state velocity reversals. In comparison with the Jenkins element method however, the behaviour between velocity reversal points was not so representative. The regions of microslip that were identifiable by abrupt changes in stiffness were only approximated by the smooth behaviour of the Bouc-Wen model. The time domain response would also be altered as the stiffness demonstrated by the Bouc-Wen model between velocity reversals varied differently to the finite element model. In contrast the matched Jenkins element model was represented by just 3 model parameters to the Bouc-Wen model's 4. With fewer parameters, the Jenkins element model was able to offer a better representation of the bilinearity of the finite element model, and could therefore be considered a more efficient solution.

To accurately represent the response of the control finite element model (19kN preload and 240Nm maximum applied torque) 4 Jenkins elements were required. This constituted a 9-parameter model. More Jenkins elements were required as the amount of microslip in the model was greater and the transition in stiffness between velocity reversals could be viewed as more smooth. The Bouc-Wen model used to replicate the control hysteresis was still characterised by only 4 parameters. The parameter extraction method only yielded two solutions for α and β . Both existed when the initial stiffness A was taken as the highest value of $21 \times 10^6 \text{Nm/rad}$. One solution was obtained for $n=1$ and a second was found for $n=2$. The full steady state response was again established between the displacement limits of the finite element hysteresis loop.

Finite Element Torque (Nm)	Initial Stiffness Parameter A (Nm/rad)	n	Energy Dissipated (J/Cycle)	Energy Dissipated in Finite Element Test (J/Cycle)	Percentage Difference
240	21×10^6	1	1.141×10^{-2}	0.999×10^{-2}	14.21%
240	21×10^6	2	1.312×10^{-2}	0.999×10^{-2}	31.33%

Table 5.4.2 Comparison of the energy dissipated by the Bouc-Wen model when matched using different parameters to the "control" finite element model with 240Nm maximum applied torque.

Despite offering a superior approximation of smooth microslip response the amount of energy dissipated by the Bouc-Wen model over-predicted that of the finite element model (Table 5.4.2). The reason for this can be seen in the comparison of the hysteresis loops shown in Figure 5.4.2.

Again the Bouc-Wen model traced a path that lay at times inside and outside the reference finite element loop. For the most accurate agreement between the two different methods of hysteresis generation, the parameter $n=1$ showed that there was also good correlation of the stiffness immediately preceding velocity reversal. In the case of the Bouc-Wen model when $n=2$ the stiffness was greater for longer than the corresponding model for $n=1$. However, the loss of stiffness was more marked, and could be related to a situation where a rapid transition was made to a contact interface with more sliding regions than predicted by the finite element analysis. The overall impact of the rapid transition to a large degree of microslip outweighed the case where by microslip started earlier in the cycle, but had progressed more slowly to a smaller area when the velocity was reversed. Further evidence of this was given by comparing the contact stiffness at velocity reversal, which was almost nothing in the case of $n=2$.




Figure 5.4.2 Comparison of the hysteresis generated by the Bouc-Wen model with the "control" finite element simulation.

It can be seen in Figure 5.4.2 that the Bouc-Wen model that best matched the defining physical characteristics of the finite element response also provided the best approximation of the energy dissipated. A large range of microslip was represented by a model that was described by half the number of parameters required by the equivalent Jenkins element model. The result of the simplified model was that the Bouc-Wen simulation was far less demanding of processing power and ran to completion in significantly shorter time than the Jenkins element model. The payoff for improved efficiency was that the physical behaviour of the finite element joint was only approximated by the smooth profile of the Bouc-Wen hysteresis. The Jenkins element method, on the other hand, offered more direct physical insight to the contact interface's behaviour with clearly defined stiffness at all points on the hysteresis loop. In contrast the Bouc-Wen model was only capable of directly matching the initial, fully stuck, contact

stiffness and the stiffness prior to velocity reversal when the parameters were extracted well enough.

The final case that was tested was the case where 280Nm was applied to the finite element joint model, and almost the entire joint interface was in sliding contact. The full range of contact stiffness was in evidence in the finite element hysteresis loop ranging from fully stuck to the stiffness offered by the torsional spring-like behaviour of the bolt-nut component. Four different sets of parameters were extracted to define the Bouc-Wen model. Two cases were carried out for an initial stiffness A of 15.3×10^6 Nm/rad with $n=1$ and $n=2$ respectively. Similarly two different values of n were used at the higher stiffness parameter value of 21×10^6 Nm/rad.

Finite Element Torque (Nm)	Initial Stiffness Parameter A (Nm/rad)	n	Energy Dissipated (J/Cycle)	Energy Dissipated in Finite Element Test (J/Cycle)	Percentage Difference
280	15.3×10^6	1	2.8842×10^{-2}	2.21×10^{-2}	30.5%
280	15.3×10^6	2	3.1395×10^{-2}	2.21×10^{-2}	42.1%
280	21×10^6	1	3.3207×10^{-2}	2.21×10^{-2}	50.3%
280	21×10^6	2	3.51×10^{-2}	2.21×10^{-2}	58.8%

Table 5.4.3 Comparison of the energy dissipated by the Bouc-Wen model when matched using different parameters to the finite element model with 280Nm maximum applied torque.

In all of the Bouc-Wen models established, the amount of energy dissipated was over-predicted. Unlike the previous two cases where the higher initial contact stiffness resulted in the best estimation of the energy dissipated per cycle, the case with least error was found with a stiffness parameter of $A=15.3 \times 10^6$ Nm/rad and $n=1$ (Table 5.4.3). In each of the three loading situations, the least error was consistently found to occur when $n=1$. The comparison between the finite element simulation and the Bouc-Wen hysteresis

shown in Table 5.4.3 illustrates that the parameters were not good enough to generate an acceptable solution in terms of the amount of energy dissipated.

Despite giving a poor approximation of the amount of energy dissipated, the Bouc-Wen model was still capable of replicating many of the important features of the finite element hysteresis model. The initial stiffness and points at which velocity reversal occurred were represented exactly for the initial loading cycle, and offered perfectly acceptable agreement during steady state oscillations. Figure 5.4.3 shows that the stiffness before velocity reversal was under-predicted in all cases. Again the analogy made (for the case of 240Nm applied torque) regarding the increased amount of sliding contact in the model, and subsequent increased energy dissipation that this reduced stiffness suggests, could be made. Another similarity between the Bouc-Wen models extracted at 240Nm and 280Nm was that the finite element simulations became less stiff far earlier in the loading cycle than the equivalent Bouc-Wen models. Only the Bouc-Wen model with $A=15.3 \times 10^6$ and $n=1$ showed a significant amount of stiffness at velocity reversal. The implication was that the parameters extracted for the remaining cases resulted in a yield torque z_y (Equation 5.2.6) very close to the 280Nm applied torque.

Despite the inaccuracies of the stiffness before velocity reversal and the overall amount of energy dissipated, the Bouc-Wen model offered a further significant benefit over the Jenkins element beyond the computational efficiency. Free vibration tests have shown that when the amplitude of vibration became small enough, the Jenkins element model of a joint becomes conservative.

This behaviour would not be exhibited by the Bouc-Wen model and as a result offered a significant improvement in describing free vibrations. The Bouc-Wen model only has a stiffness A immediately after velocity reversals. The moment the relative displacement of the surface became non-zero, the Bouc-Wen model began dissipating energy resulting in a system that is non-conservative at all times during free vibrations. Such a system would occur in all real joints, but also in any simulation where the contact

pressure at an interface decreased to zero between the two components. Such an interface would result in the immediate onset of microslip at velocity reversal and a non-conservative system. The existence of such interfaces (including the joint modelled in the finite element analysis) illustrates the practical benefit of the Bouc-Wen model. This benefit remains, with the computational efficiency in approximating large microslip situations, even if preliminary parameter extraction does not yield hysteresis loops with acceptable agreement of energy dissipation.




Figure 5.4.3 *Finite element hysteresis loop obtained with a 280Nm applied torque in comparison with matched Bouc-Wen models.*

5.5) Hysteresis Prediction Using the Bouc-Wen Model

It was shown that when using the parameters extracted for one Jenkins element simulation, it was possible to predict the response to a reduced input torque for example. The ability for one set of Jenkins element parameters to be used in several scenarios was one of the model's most valuable assets. To establish whether the same property existed for a given set of Bouc-Wen parameters, two tests were performed. In the first test the parameters extracted for the best matching loop at 280Nm of applied torque ($A=15.3 \times 10^6 \text{ Nm/rad}$, $n=1$, $\alpha=19760$ and $\beta=30428$) were used to predict the response over the range of angular displacement associated with 240Nm and 80Nm of applied torque. In the second test the best matching loop at 80Nm of applied torque ($A=21 \times 10^6 \text{ Nm/rad}$, $n=1$, $\alpha=55819$ and $\beta=94162$) was extrapolated over the range of angular displacement of the two larger input torques.

Tables 5.5.1a and b. show that the agreement in the energy dissipated in all of the cases was relatively poor. The only exception was when the results from the 80Nm analysis were extrapolated to the 280Nm range of displacement. The agreement in this case could be put down to coincidence rather than the applicability of the procedure carried out.

Torque (Nm)	Actual Energy Dissipated (J/Cycle)	Predicted Energy Dissipation (J/Cycle)	Percentage Error
240	0.999×10^{-2}	1.2463×10^{-2}	24.75
280	2.21×10^{-2}	2.3558×10^{-2}	6.60

Table 5.5.1a Comparison of the predicted energy dissipated with the actual energy dissipated when model parameters obtained for 80Nm input torque were calculated for larger values.

Torque (Nm)	Actual Energy Dissipated (J/Cycle)	Predicted Energy Dissipation (J/Cycle)	Percentage Error
80	2.69×10^{-4}	9.2225×10^{-5}	-65.71
240	0.999×10^{-2}	1.4138×10^{-2}	41.52

Table 5.5.1b Comparison of the predicted energy dissipated with the actual energy dissipated when model parameters obtained for 280Nm input torque were calculated for smaller values.

Figure 5.5.1 shows that the superior agreement achieved when extrapolating from small values of torque to high values was not representative of superior agreement in all of the other critical hysteresis characteristics. The yield force predicted by the Bouc-Wen parameters for the 80Nm model create a yield torque of 140Nm. At this point all stiffness in the model disappeared and any further increases in displacement were not accompanied by an increase in the magnitude of the restoring torque. The range of angular displacement that both encompassed were enough to cause yielding for the extrapolated parameters which created two errors in the extrapolated results. The first error was the zero contact stiffness before velocity reversal. In all the finite element analyses, every loop had a significant stiffness at velocity reversal and this was a significant behavioural characteristic. The second discrepancy caused by the yielding resulted from the fact that no increase in restoring torque occurred above 140Nm as the displacement increased. This led to both extrapolated hysteresis loops having a much lower average stiffness over the entire cycle than the finite element simulation and the Bouc-Wen models with separately identified parameters.

The fact that both hysteresis loops offered a credible agreement in the amount of energy dissipated could largely be put down to chance. This is clearly illustrated in Figure 5.5.1. Nearly an equivalent area was occupied by the extrapolated results outside the finite element hysteresis loop and the area of the finite element hysteresis loop above and below the restoring torque of 140Nm and -140Nm respectively.




Figure 5.5.1 A comparison of the extrapolated hysteresis loops with their equivalent finite element simulation results for 240Nm and 280Nm maximum applied torque.

When the results were extrapolated to a lower torque from a model established at 280Nm, the yield torque was no longer an issue. Poor agreement in the amount of energy dissipated still existed for both the range of displacements equivalent to 80Nm and 240Nm. The reason for this poor agreement was further evidence of one of the limitations of the Bouc-Wen model.

Figure 5.5.2 shows the comparison between the "exact" hysteresis loops and those obtained using the Bouc-Wen parameters established at 280Nm of applied torque. Although a reasonable approximation of contact stiffness was given at velocity reversal, the torque at which it occurred was over-predicted in both cases. The Bouc-Wen method provided only a smooth approximation of the overall hysteresis loop for which its parameters were extracted. Unlike the Jenkins element method that fixed the

approximated model to the finite element hysteresis loop at discrete points, the Bouc-Wen model was only constrained at the points of velocity reversal. This led to the approximation, which offered the best fit of the finite element hysteresis as a whole, offering a poor fit particularly when the range of displacement was reduced.

Care would therefore be required when putting a single Bouc-Wen model into a structure where the magnitude of vibration was not predetermined. For a Bouc-Wen model to be used in this type of situation the parameters would need to offer a far better agreement over the whole cycle rather than the critical points around velocity reversals.




Figure 5.5.2 A comparison of the hysteresis generated by the Bouc-Wen parameters extracted for 280Nm torque when used over angular displacement ranges for 80Nm and 240Nm.

Chapter 6. Experimental Analysis

6.1) Experimental work Summary

In a bolted joint that maintains its integrity the main means of dissipating energy is through microslip. The angular displacements associated with microslip are very small, yet essential for characterising the energy dissipated through joint hysteresis. Consequently the aim of the experimental work carried out in this chapter was to generate microslip in a simple structural joint and then measure both the angular displacement and restoring torque of the joint as locally as possible.

The joint that was used for the experimental investigation was a lap joint mounted near the fixed end of a composite cantilever component. To generate the motion at the joint location, the beam was excited at its first resonant frequency in the same plane as the contact interface. The motion of the joint and its surroundings was measured using accelerometers whilst the bending moment at the joint location was calculated from measurements taken by strain gauges on either side of the bolt. Time domain data from all the transducers allowed the relative angle of rotation of the top free component to be established against the bottom clamped component. The restoring torque at the joint itself was taken to be the bending moment at the bolt axis minus the inertial effects of the free component. Four different levels of preload were applied to the M10 bolt that held the joint together. At each level of preload different amplitudes of beam excitation were investigated.

At the tightest joint condition no microslip was in evidence from the experimental hysteresis. Instead the energy dissipated across the joint could be characterised as viscous by the elliptical profile of the hysteresis loop. This result was consistent for all amplitudes of vibration. As the bolt preload was decreased the amount of microslip detectable from the joint hysteresis increased. For a fixed preload microslip became more prevalent as the amplitude of vibration increased. Overall the joint behaved as a softening

spring with amplitude dependent stiffness. With certain exceptions, within each cycle the trend was for decreasing stiffness as the amplitude of oscillation increased.

When the preload on the joint was at its lowest two settings, the power spectrum densities revealed an increasing influence of the odd superharmonics. This manifested itself in hysteresis loops that appeared to show some hardening behaviour over the initial part of the cycle, before softening as the velocity reversal point was approached. The increased damping and subsequent reduced amplitudes of vibration at the lower preloads meant that extracting the joint behaviour from the background noise became increasingly difficult. A particular problem was that the natural frequency of the beam (47.7Hz) was close to the 50Hz power supply frequency.

The experimentally obtained hysteresis loops were also matched using the Jenkins element model and the Bouc-Wen model. Both were capable of representing the main behaviour of the joints in terms of the average stiffness, the totally stuck contact stiffness and points of velocity reversal. It was the Jenkins element model that best approximated the overall amount of energy dissipated by one cycle of experimental vibrations. However, the potential of the Bouc-Wen model lay in its ability to reproduce the hardening, then softening, behaviour associated with the increased superharmonic content at lower preloads. Until the parameters that identified the Bouc-Wen loops were more representative of the actual experimental hysteresis this potentially superior matching was not harnessed fully.

6.2) Experimental Setup

The aim of the experimental rig was to demonstrate as explicitly as possible microslip behaviour in a bolted joint. The type of joint used had to be as simple as possible to isolate the frictional dissipation to very localised area. For this purpose a lap joint was used that had just a single contact interface between the two bolted components. Each component when joined rigidly made up an equivalent cantilevered member with a joint located on its length. The components used to clamp the fixed end of the composite cantilever were bolted in place by joints larger and stiffer than those used in the lap joint under investigation. The larger clamping force and geometry of the components used to restrain the fixed end of the composite cantilever prevented motion between any parts of the experimental rig except in the unrestrained length of the composite cantilever. An M10, mild-steel, zinc-coated bolt was used to generate the clamping force at the joint interface so the experimental rig corresponded to the finite element model as closely as possible. To ensure complete consistency a beam with the same cross section as that used in the finite element study would be required. Such a cross section (80mm×30mm) would have meant a tremendous amount of energy was required to generate the required motion. Instead a beam with a smaller cross section was used as the self-weight would have less impact on the joint, the resonant frequency would be lower due to increased flexibility and it would be more representative of practical structural members. The solid mild steel blank used was machined down to give a section of 30mm×16mm. As a composite beam structure the wall to tip length was 693mm

To generate a full range of microslip in the joint it was necessary to generate as much movement as possible at the region where the frictional interface was located. In a cantilever beam the largest bending moment exists at the wall where the beam end is clamped. In the experimental rig the torque applied to the joint was supplied by the bending moment of the beam. Although the range of motion in the first mode of excitation is greatest at the free end of a cantilever, the relative motion between the two contact interfaces would be greatest where the torque was at its largest value-at the wall

of the cantilever. The joint interface was therefore placed as close to the wall as possible while still allowing room for all the measurement devices.

Excitation of the beam was supplied by a single shaker (Ling Dynamic Systems V406 Shaker). The shaker was attached to the bolted joint via a stinger that was screwed into both the shaker and the free beam component. To generate the maximum amount of motion in the beam and therefore applied torque to the joint, the whole beam ensemble was shaken at its first bending natural frequency in the direction of the width of the joint.

The power supply that generated the response of the shaker was provided by a signal generator (Muirhead D-880-A 2-Phase L.F Decade Oscillator) passed through a power amplifier (Ling Dynamic Systems PO 100 Power Oscillator). A (Blackstar 3225MP) multimeter was used to monitor the amount of current going into the shaker at all times. It was connected in series between the power amplifier and shaker and had an upper limit of 10A capacity before the circuit was automatically broken.

An accelerometer (PCB Piezoelectronics M353B16 Shear Accelerometer) attached with wax mounted 23mm away from the free tip of the beam showed the movement of the beam when the bolt was tightened to its maximum level. The signal from the accelerometer was passed through a PCB Piezoelectronics (Model 482A06 I.C.P) power supply to both a frequency counter and an oscilloscope (Farnell Dt12-5 Oscilloscope). By tuning the frequency output from the decade oscillator and monitoring the amplitude of the signal using the oscilloscope, the frequency for maximum tip acceleration was found to be 47.7Hz using the frequency counter. This was assumed to be the resonant frequency of the rigidly clamped system.

The basic behaviour of the beam and subsequent calculations had been defined using a number of assumptions:

- The local mass of the joint area was not considered as significant in altering the behaviour of the beam.

- When the joint was rigidly clamped no frictional sliding took place at the interface and the two beam components could be considered as a single cantilevered structure.
- There was no acceleration, velocity or displacement of the beam at the clamped end.
- The beam's cross section and material properties were assumed constant along the whole length.

Figure 6.2.1 shows the localised joint region in some detail. The design of the localised joint area offered the best opportunity of measuring microslip. Both surfaces of the interface were milled so to achieve a consistent surface roughness that would allow a reasonable amount of actual contact that a more coarsely machined surface would not. The nominal contact surface was 30mm×30mm and large enough to allow a reasonable amount of variation in contact pressure at the interface. The interface itself lay as closely as possible to the centroid of the main beam structure to prevent twisting during oscillations. It is also notable that there are thicker areas around the joint which were filleted down to the main cross section of 30mm×16mm. These thicker areas were to allow the positioning of measuring devices and to prevent the cracking of the beam due to stress concentrations and reduction in cross sectional area.




Figure 6.2.1 *Localised joint detail of the experimental rig showing filleting and sectional view illustrating the clearance hole and tapped accelerometer locations.*

Around the joint area there were tapped holes to accept four accelerometers of the same specification as the one used near the tip of the beam. Two of the accelerometers were attached to the clamped component of the cantilever and two are fixed to the free component. Each accelerometer was attached by screwing it into the tapped hole, but with a small amount of wax on the base to prevent self-loosening during dynamic excitation of the rig. The accelerometers around the joint did not have a large enough signal to be detected during normal use of the rig. All four had their signal passed through a pre-amplifier (PCB Model 442B104 4 Channel I.C.P Sensor Signal Conditioner) boosting it by a factor of 100.

Also in the locality of the bolt, two pairs of strain gauges were connected in a half bridge configuration to measure the bending strain either side of the bolted joint along the beams length. The positions of the strain gauges, shaker attachment and accelerometers are shown in Figure 6.2.2 The output from the strain gauges was passed through 2 (Fylde) strain indicators. All of the data from the measuring devices was passed through the respective signal processing devices and via a PCI-DAS 1000 Connection unit into a personal computer for direct processing in Microsoft Excel.

Two calibration tests had to be performed before the experimental results could be taken to derive the joint's dissipative behaviour. The first calibration was performed on the strain gauges to establish a relationship between the gauge output and the bending moment at the gauge location. A second calibration test was also performed on the data collection device that linked the output from the experimental rig to the personal computer.

To calibrate the strain gauges a series of static loading tests were performed on the rigidly clamped joint. A series of known masses were hung from the beam tip which allowed bending moments to be calculated at the strain gauge locations. The strength of the signal received by the personal computer could then be easily related to the size of the bending moment that was then known to exist at each gauge. A linear relationship was found to exist between the signal supplied to the personal computer and the size of the

bending moment in the beam for each gauge. This scaling factor was consistent in the full range of displacement that was to be measured during the experimental investigation.




Figure 6.2.2 Composite beam with all transducer locations marked: A-accelerometers, S-strain gauges and B-shaker location.

The data collection device worked by sampling each input channel in turn. As each channel was read sequentially a time delay existed between the readings. This time delay was found to be a function of the data collector's sampling frequency. Such a delay was particularly significant as all of the results were measured in the time domain.

The exact nature of the relationship between the time lag for sequential channel readings and the sampling frequency (f_s) was established using a sinusoidal input applied to all channels simultaneously. As the sampling frequency increased the time delay between channels became shorter. Over a range of sampling frequencies between 1-2kHz a constant ratio was found between the time lag per channel and the period of the sampling frequency. This ratio was converted to a phase angle (ϕ_s) of 1.047rad/channel.

For all of the tests herein a sampling frequency of 1.5kHz was used as this gave an acceptable resolution of about 30 data points per cycle of the rig at resonance.

The time lag T_{lag} between each channel was therefore calculated as:

$$T_{lag} = \frac{\phi_s}{f_s \times 2\pi} \quad (6.2.1)$$

Before the time lag was calculated:

$$ch_n = f(t) \quad (6.2.2)$$

In Equation 6.2.2 ch_n is channel n and t represents time. After the time lag was incorporated into the results:

$$ch_n = f(t + nT_{lag}) \quad \text{for } n=1 \text{ to } 6 \quad (6.2.3)$$

All of the data from channels 1 to 6 were now normalised with respect to the base channel (channel 0). However, the data across all of the channels was still not sampled at coincident points in this time domain. To express the data for channels 1 to 6 at the same time values used for channel 0 the results were linearly interpolated on to the base time scale using Matlab.

Now all of the data from the transducers was given at the same points on a normalised time scale. What remained was to establish the bending moment at the joint interface and the relative angular displacement between the two clamped components. Both the torque applied to the joint and the angular displacement had to be established at the same point of the joint interface. The point chosen was at a point the same distance from the clamped end as the bolt axis, on the length of the beam, at the contact interface.

The applied torque to the joint was taken to be equal to the bending moment experienced at the reference point on the interface. This bending moment was taken by taking a cubic interpolation at the reference point using three other known bending moments on the beam. Two of these bending moments were established using the strain gauge measurements. The third known bending moment was taken at the tip of the beam where it had to be 0Nm. This interpolation was completed for the every sample through the duration of the test.

To calculate the angular displacement of the two components was more complicated. In both cases, cubic interpolations were carried out to establish the displacement profile of the upper (free) beam component and the lower (fixed) beam component. In both cases the angular acceleration at the edge of the joint was calculated first. For the free beam component, the three locations where the acceleration was known across the width of the beam were at the bolt axis, 23mm from the free tip of the beam and 61mm along the beam's length from the bolt axis. A cubic curve was then fitted through the three points to give an approximation of the acceleration profile of the "free" beam at every point in time for the duration of the test.

6.3) Experimental Calculations

The raw data collected using the personal computer gave very specific parameters for carefully determined points on the two components that made up the cantilever beam. Each of the five accelerometers gave accelerations at the specific location where they were mounted. Similarly the strain gauges only revealed the bending moment at their position on the composite structure. All of this data was sampled at 1.5kHz over 1000 samples.

One of the problems with the output from the experiment was that the data was not necessarily centred about the middle of the sampling range of voltage input to the data collector. Before any further manipulations were performed, the output from each transducer was centred about the middle of the digital range over which it was observed. The duration over which samples were measure was 0.666s and the frequency of oscillation of the beam was 47.7Hz, over 30 cycles of data were collected from each transducer. Taking an average of the readings over the entire period during which data was collected approximated the midpoint of the sampled data from each transducer. Each transducer's output was then measured relative to this artificial zero that had been established. As only microslip was taking place, it was assumed that the amount of offset in the simulation would be strictly limited and the results would therefore be well centred.

Once the digital data was sampled it could be converted into a voltage output from each transducer. After the transducer voltage output was established it was possible to find either the bending moment being measured at each strain gauge or the acceleration at each accelerometer. Each accelerometer had a slightly different scaling factor for converting the voltage output into a value of acceleration and these had to be applied to each transducer in turn. The accelerations from the four accelerometers around the joint also had to be divided by a factor of 100 to take into account their signal amplification. Similarly the strain gauges had different scaling factors for converting the voltage output

into the bending moment at the gauge location. These were established during static calibration tests (Section 6.2).

After scaling of the voltage output each transducer had readings of physical significance for the point where they were mounted on the composite beam. The strain gauges gave values of bending moment, and the accelerometers revealed local acceleration. None of these physical parameters in isolation revealed anything about the energy dissipated at the joint between the two components. When combined the data obtained could reveal a great deal about the local deflection of the joint.

To create a hysteresis loop of the joint, the approximate behaviour of the interface had to be known. Like the finite element studies, the relative rotation of the two joint components could only be reduced to the relative angular displacement measured at the edge of the interface. Unfortunately it was not possible to measure the movement of the joint components right next to the joint interface. Instead measurements were taken from the middle of the depth of the two clamped components and it was assumed that plane sections remained plane throughout. If the relative angular displacement at the edge of the joint interface was known along with the applied torque then good approximations of the joint hysteresis could be established.

Before any of the transducer outputs could be combined to reveal the angular displacement from the joint interface, the time domain data had to be normalised to remove the natural phase shift caused by the sequential sampling of the data collector Equation 6.2.1. To achieve this the time domain response of the first channel data was assumed to have the base time range. All of the other results were then attributed new time co-ordinates that took into account the phase lag between the channels relative to the channel 0 output.

The phase lag was 1.047rad for a sampling frequency of 1500Hz. This corresponded to a time lag between each channel of 0.11063ms. Although sampling times had been normalised to the base time scale associated with channel 0 it was still required

to obtain readings at coincident times for all of the transducers. The output from each of the transducers was therefore linearly interpolated (using Matlab) at the time increments used for channel 0. An approximation was therefore obtained for all of the transducer's outputs at identical points in time.

Having a series of physical parameters at discrete times over a 0.666s interval facilitated the understanding of the joint hysteresis only after further manipulations. The joint itself covered a finite area, yet it was not possible to know the behaviour over the entire interface. Instead an approximation similar to the one used to analyse the finite element results was used.

Plotting the restoring torque against the angular displacement of the joint could generate each hysteresis loop. An approximation of the angular displacement of the joint was taken by measuring the displacement of the joint at the edge of the interface. This measurement approximation is shown in Figure 6.3.1. The point used was on the outside of the interface at the same length from the clamped end as the bolt axis. It was impossible to directly measure the angular displacement at this location, but possible to obtain an approximation from the profile of the two components that made up the composite beam.




Figure 6.3.1 *The joint interface with angular displacement approximated by θ and the position at which the measurement took place labelled A.*

Both components that made up the beam had 3 points where the acceleration was known. For the top "free" component, the two accelerometers around the joint were used in combination with the accelerometer at the beam's tip. The lower "fixed" beam component had the two accelerometers around the joint, but it was also assumed that the acceleration, velocity and displacement at the point where the cantilever was clamped were zero. Once the relevant scaling factors had been applied for each transducer, and the data had been normalised in the time domain, accelerations were known at these six locations.

Using Matlab's cubic interpolation procedure a curve was produced to approximate the linear acceleration profile of the two components. The approximation extended from the clamped point to the bolt axis on the fixed component, and from the bolt axis to the accelerometer mounted close to the tip of the free component. For both components the acceleration profile was very finely discretised near the bolt axis. Using the fine discretisation at the bolt axis, an approximation was taken as to the gradient of the linear acceleration profile at the bolt itself. In the case of both the fixed and free components it was assumed that the gradient offered a good approximation of the angular acceleration at the interface at A in Figure 6.3.1. Again this relied on the assumption that throughout the length of the beam all plane sections were remaining plane.

The actual deflection (y) profile of a cantilever beam in its first bending mode is defined by the following relationship (Blevins (1984)):

$$y = \cosh\left(\lambda \frac{l}{L}\right) - \cos\left(\lambda \frac{l}{L}\right) - \sigma_1 \left(\sinh\left(\lambda \frac{l}{L}\right) - \sin\left(\lambda \frac{l}{L}\right) \right) \quad (6.3.1)$$

Where λ and σ are constants associated with the mode of vibration in question, l represents the length from the clamped end and L the total length of the beam. The theoretical acceleration at a point l on a cantilever's length is given below in Equation 6.3.2.

$$\frac{d^2y}{dt^2} = \lambda^2 \cosh\left(\lambda \frac{l}{L}\right) + \lambda^2 \cos\left(\lambda \frac{l}{L}\right) - \sigma_1 \lambda^2 \left(\sinh\left(\lambda \frac{l}{L}\right) + \sin\left(\lambda \frac{l}{L}\right) \right) \quad (6.3.2)$$

The theoretical profile of both the linear displacement and acceleration could not be assumed for the composite beam being investigated. Due to the presence of the joint, the boundary conditions used to derive the above relationships were altered for each component. For the fixed component, the tip was not free, but attached to the free component. In contrast the free component did have a free end at the tip, but was not rigidly clamped at its fixed end. The presence of the built up area around the bolted joint also meant that there might have been small deviations due to concentrations of mass along the length of the beam. By using a cubic approximation of the acceleration profile, a reasonable representation of the theoretical variation of acceleration with length was provided. However, any inconsistencies in the beam profile due to the factors given above were accounted for in the interpolated relationship. It should also be noted that the proximity of the pair of accelerometers on each component around the bolted joint region meant that the interpolation was more tightly constrained in this area.

The relative angular acceleration between the two components at the joint interface was established using the simple relationship:

$$\ddot{\theta}_{rel} = \ddot{\theta}_{free} - \ddot{\theta}_{fix} \quad (6.3.3)$$

The relative angular displacement was found by integrating the relative angular acceleration twice in the time domain. The integration was performed using the trapezium rule provided in Matlab. Before integrating the angular acceleration signal and the angular velocity, the signals were passed through a filter. Again an internal Matlab command was used to create a high-pass Butterworth filter that removed the lowest frequency content of the signal. It was the low frequency content (much below the excitation frequency of the beam) that resulted in offset of the signal. This offset was present in the experimentally obtained angular accelerations and also in the angular

velocity produced by the first integration in the time domain. The filter was designed in such a way that a full strength signal was achieved above 30(Hz), but from 0(Hz) to 30(Hz) the signal strength increased smoothly from zero to its full strength. After filtration it was possible to relate the relative angular displacement θ_{rel} at the joint to the angular acceleration as follows:

$$\theta_{rel} = \iint \ddot{\theta}_{rel} dt dt \quad (6.3.4)$$

Due to the assumptions of zero offset for microslip simulations, and the filtration of the frequency content that exists around 0(Hz), the constant of integration in the above integration was assumed to be zero.

A similar interpolation procedure was used to find the applied torque created by the bending of the composite beam. The bending moments were known at the two strain gauge locations either side of the joint region. One further point was required to establish the third cubic interpolation point. It was known that the bending moment in the beam was zero at the free tip. By applying a cubic interpolation to these three known values of bending moment it was possible to approximate the bending moment and therefore torque at the bolt axis location. However by taking measurements either side of the interpolation point and quite close to the bolt itself it was anticipated that errors and uncertainties were minimised.

A similar procedure to that used in the finite element simulations was carried out to identify the value of the restoring torque within the joint. It was taken that the joint restoring torque T_{res} was equivalent to the applied T_{ap} torque minus inertial effects:

$$T_{res} = T_{ap} - I\ddot{\theta}_{rel} \quad (6.3.5)$$

The moment of inertia of the beam was calculated with a number of assumptions included. Firstly the inertial effects of the short, fixed components were taken to be

negligible. Secondly the inertial effects of the joint region itself were considered to be small. The free beam component was also assumed to be a uniform beam of cross section 30mm×16mm. This did not take into the different profile around the joint where the filleted profile added some extra mass. However, all of the accelerations of the free component were assumed to be relatively small, as a consequence of the low frequency of excitation and small range of displacements. As a result any simplifications of the representation of the moment of inertia would have only a slight impact on the size of the restoring torque.

The moment of inertia of the free beam component about the bolt axis was calculated using the parallel axis theorem. Initially the moment of inertia was calculated for the component about its own centroid I_G . This was then corrected to represent the moment of inertia about the bolt axis. The correction was achieved by adding the component mass m multiplied by the length from the bolt axis to the centre of mass of the component l_g squared.:

$$I = I_G + ml_g^2 \quad (6.3.6)$$

The mass of the component was calculated using a standard density for mild steel of 7800kg/m³. With this value the moment of inertia of the free beam component about the bolt axis was 0.2022kg.m². When combined with the small relative acceleration values any impact on the joint hysteresis could be seen to be small.

At every point in the time domain values of relative angular acceleration, relative angular displacement and applied torque were known at the specific location on the joint's interface. This allowed the establishment of hysteretic relationships for the interface between the two components making up the composite beam. From this behaviour hysteresis plots were obtained to illustrate the mechanism by which energy was dissipated by the lap joint.

6.4) Tightest Joint Condition

The experimental results could be broken down in a series of tests performed at four different preloads. At each preload the amplitude of oscillation was varied to establish the frictional damping's dependence on the preload and other factors. For each preloading case the bolt was loosened by approximately $1/8^{\text{th}}$ of a turn. This caused a decrease in preload incrementally from an initial tightly clamped situation. The tightly clamped joint condition was reached when the bolt could be tightened no further by hand when using a pair of wrenches. The loosening process was stopped once the bolt became free enough for the head and nut to slide against the beam surfaces, and when there was not enough force to prevent the self-weight of the beam pulling the lap joint contact area apart.

These conditions were only approximate and reliant on qualitative judgement to establish when they were reached. Without the use of a load cell it was impossible to quantify the bolt preload in these conditions. Such a load cell could however disturb the distribution of pressure within the joint and cause an asymmetry of the loading conditions. The qualitative approach taken did allow a wide range of joint phenomena to be investigated and a good understanding of the dissipative processes to be established. Even though the pinned condition investigated by other authors (Moon and Li (1990)) was not generated by the loading conditions, such a joint would be deemed to have failed when only microslip was permitted as a dissipative mechanism. In the experimental analysis carried out here the hand-tight to rigidly clamped conditions demonstrated the most likely service requirements of a bolted joint and prevented failure due to macroslip. This corresponds directly to the cases modelled using the finite element method, Jenkins element model and Bouc-Wen model of hysteresis-thus allowing a direct comparison of results obtained in each case to be made.

The nominally rigid bolted connection shows the significance of microslip in dissipation of energy from the jointed beam. With the bolt tightened as much as possible, the preload, and therefore contact pressure, were at their peak values. In this situation, the

amount of microslip at the bolt interface was at a minimum. If the contact pressure had a finite magnitude at all points on the surface then the joint may never undergo microslip unless the applied torque was of a certain magnitude. The joint would have a particular stiffness until microslip was initiated. At this point energy would be dissipated, and the stiffness of the contact interface would decrease. However if the limiting value of torque was not reached at any point in the harmonic oscillations then the response of the contact interface would remain elastic in the context of the representative models that have been used thus far.

In the case of the finite element model the interface area was relatively large compared to the depth which caused the contact pressure to decay to zero at certain points. The dimensions of the joint used in the experimental analysis meant that this was not necessarily the case, and that all areas of the interface would experience some magnitude of contact pressure. Any assumptions made about the symmetry and regularity of the contact pressure also had to take into account the asymmetric loading caused by the self-weight of the beam and surface irregularities within the limits machining process. However it was assumed that these effects would be relatively small, and that the pressure distribution was largely axi-symmetric about the bolt axis.

Figure 6.4.1 shows the effect of increasing the amplitude of oscillations of the beam when it was in the rigidly clamped condition. It is an interesting result as the hysteresis loops do not exhibit any of the characteristic features associated with purely friction damping. While it is clear that energy is being dissipated from the beam, the mechanism for this dissipation does not appear to be caused by friction. The loops generated all have an elliptical profile characteristic of viscous damping (Thomson (1998)).




Figure 6.4.1 *Hysteresis loops obtained from the experimental joint when varying amplitudes of vibration were applied at maximum bolt preload.*

At the largest two largest amplitudes of vibration, the response begins to flatten just slightly. This is indicative of microslip and suggests at torques of 50Nm and above the joint starts to slide in very localised areas. Until this level of torque is applied to the joint the energy dissipated is likely to be through material damping and atmospheric damping. A small contribution may have also be made by magnetic effects as the joint became slightly magnetised over the duration of its service. The relative magnitude of the viscous-like damping effects is much larger than any dissipation due to friction at the maximum preload. It is therefore possible to characterise the joint as exhibiting linear damping. Softening spring behaviour is not present when considering the average stiffness over a cycle of oscillations. Table 6.4.1 shows that the average stiffness of the hysteresis loops between points of maximum angular displacement remain largely consistent with a maximum variation of only 4.3% between the largest and smallest stiffnesses. Sharp changes in stiffness at velocity reversals are also missing from the

linear damping shown in Figure 6.4.1. The profile of the hysteresis loop is therefore smooth and does not exhibit regions of constant stiffness associated with microslip in the finite element joint model.




Figure 6.4.2 Power spectrum density plot for the angular displacement of the joint at maximum preload and applied torques of varying maximum amplitude.

The power spectrum shown in Figure 6.4.2 displays the characteristics associated with linear damping in most of the simulations. There is a strong peak at the resonant/excitation of the beam. At the smaller oscillation amplitudes this peak remains the only significant contribution to the angular displacement of the joint. However the two of the three largest input amplitudes also result in a small peak at the third harmonic of oscillation. These features are commensurate with the onset of friction damping. The fact that they occur at the larger input torque values is consistent with the expected behaviour as small regions experience a torque that can overcome the limiting friction values associated with them. At these larger torques a slight distortion of the elliptical

profile can also be seen as the velocity approaches zero. This can be attributed to a small degree of nonlinear softening behaviour at the joint interface. Even at the largest input torque the amount of frictional behaviour can be considered small.

Table 6.4.1 shows that the damping provided by the joint is almost linear as the proportional increase in angular displacement is accompanied by a similar increase in the amount of energy dissipated. The amount of energy does increase at a slightly larger rate than the angular displacement which does suggest a small degree of nonlinear damping. Whilst not of a large enough degree to be visible from the hysteresis, or spectral density plots this is further confirmation that a small amount of dissipative microslip exists in the joint even with the largest applied preload. However, for this damping to become significantly evident then the amplitude of vibration needs to be in the upper range shown here or a reduction in bolt preload must be produced.

Relative Applied Amplitude	Energy Dissipated (J/Cycle)	Maximum Angular Displacement (rad)	Average Cycle Stiffness from Peak to Peak Displacement (Nm/rad)
Smallest	1.2868×10^{-2}	4.61×10^{-4}	6.1171×10^4
Second Smallest	2.3830×10^{-2}	6.70×10^{-4}	6.1194×10^4
Intermediate	3.8961×10^{-2}	7.95×10^{-4}	5.8994×10^4
Second Largest	5.2292×10^{-2}	9.83×10^{-4}	5.8698×10^4
Largest	6.3565×10^{-2}	1.04×10^{-3}	6.1250×10^4

Table 6.4.1 *A comparison of hysteresis loop characteristics found when a range of torques were applied to a joint with maximum preload.*

6.5) Smallest Reduction of Joint Preload

Other than frictional damping all other forms of energy dissipation were assumed to be independent of the bolt preload. Any material damping in the beam would not be affected by a change in the bolt preload. Changes in the joint's behaviour could therefore be attributed directly to the change in contact conditions brought about by the change in bolt preload. Immediately the bolt was loosened from its tightest clamping force its behaviour changed dramatically. A joint with smaller overall contact pressure implies a joint with more microslip for a given applied torque. The finite element studies of Chapter. 3 indicate that the contact pressure changes in proportion to the applied preload. In this case as the preload decreases, the contact pressure can be assumed to decrease and more microslip and therefore energy dissipation may take place at the interface between the two components.

The first sign that the amount of damping had risen when the preload was decreased was seen in the amplitude of oscillation. Much more power had to be input into the shaker to generate an amount of displacement comparable to the cases with the nominally rigid joints. It was, in fact, only possible to generate a *maximum* amplitude in the reduced preload cases that corresponded approximately to the *minimum* amplitude of the rigid joint tests. Figure 6.5.1 shows the hysteresis loops obtained in the cases of a slightly reduced joint preload. It is immediately apparent that the hysteresis loops can no longer be characterised as elliptical. Even the smallest amplitude response shows a profile that is much more similar to the hysteresis loops generated in the finite element studies. The portion of the loop before velocity reversal demonstrates distortion as the interface becomes less stiff. This is attributed to greater amounts of microslip taking place in the joint and is most prevalent at the largest two amplitudes where it appears that the joint has lost almost all of its stiffness and is on the point of failure. While this may be the case, other forms of damping have to be considered. What is actually being shown in Figure 6.5.1 is not only frictional damping, but also the viscous-like damping that was present independently of the joint preload. Therefore the shape of the hysteresis loop cannot be put entirely down to the complete loss of contact stiffness at the interface. A

more telling indication of the increased amount of slip at the large torque cases is the cleanliness of the signal just before velocity reversal. Only the largest amplitude response shows a decline in signal quality before velocity reversal. This is a comparable response with the finite element studies, which showed that when the contact stiffness became small, and macroslip was being approached, then the signal quality was reduced. This would be the expected outcome in the largest amplitude response (if it was to occur at all) of all the cases at this reduced preload.




Figure 6.5.1 *Hysteresis loops generated by varying input amplitudes when the bolt preload was slightly reduced ($1/8^{\text{th}}$ of a turn).*

Another feature of friction damping that was not displayed in the rigidly clamped case was an average softening of the joint as the amplitude of oscillation increased. At this slightly looser preload the gradual softening of the joint can be seen more clearly when considering the peak to peak stiffness in each cycle. Table 6.5.1 shows that at the smallest amplitudes of vibration, where the frictional dissipation is likely to be least, the stiffness of the joint is very close to the value found in the rigidly clamped condition.

Although the relationship is not smooth, the general trend is for a slight decrease in stiffness from the smallest amplitudes to a value about 5% lower when the amplitude of oscillation is largest. The peak displacement points can only be estimated as the contribution of the viscous-like damping makes their precise identification difficult. Pure friction damping yields a sharp change in stiffness with the maximum displacement magnitude occurring at the same time as the maximum torque magnitude. Here this is not the case, but if the point where the maximum restoring torque *or* the maximum angular displacement is used to calculate the average cycle stiffness, then softening behaviour is detectable.

Relative Applied Amplitude	Energy Dissipated (J/Cycle)	Maximum Angular Displacement (rad)	Average Cycle Stiffness from Peak to Peak Displacement (Nm/rad)
Smallest	6.3969×10^{-3}	3.46×10^{-4}	5.9249×10^4
Second Smallest	9.4674×10^{-3}	4.01×10^{-4}	6.1596×10^4
Intermediate	1.3877×10^{-2}	4.91×10^{-4}	5.8900×10^4
Second Largest	2.0269×10^{-2}	5.69×10^{-4}	5.6063×10^4
Largest	2.3420×10^{-2}	6.29×10^{-4}	5.7234×10^4

Table 6.5.1 *A comparison of hysteresis loop characteristics for a range of displacements when applied to a joint with preload reduced by a 1/8th relative turn of the bolt head.*

Table 6.5.1 also illustrates a nonlinear increase in the amount of energy dissipated per cycle relative to the change in angular displacement. The difference in proportionality between the two parameters is not very different from that found in the rigidly damped condition however. Although not conclusive, the amount of energy dissipated at a given displacement does seem to be consistently higher in the cases illustrated here than in the "rigidly" clamped joint. Similar behaviour was identified by Gaul and Lenz (1997) in the gradually increased amount of energy dissipated per cycle as the amount of microslip present in the joint became greater. The cause of the increase in microslip was

progressively larger loads applied to the joint. Only when macroslip was initiated did the energy dissipated per cycle start increasing proportionally to the applied load.




Figure 6.5.2 Power spectrum for the various amplitudes of vibration when the bolt preload is reduced by a single increment.

Another sign of frictional sliding is the appearance of superharmonics in the power spectrum density plot (Figure 6.5.2). All of the angular displacement data, with the exception of that obtained for the smallest amplitude of oscillation, show a significant contribution from at least the first odd superharmonic. This corresponds largely with the first signs of softening of the joint interface (Table 6.5.1). The two largest amplitudes of vibration also demonstrate a peak in the response at the fifth harmonic of the excitation frequency. In all of the cases no further harmonics were distinguishable against the background signal even though they could be expected to exist. This was consistent with the power spectrum response of the finite element results when macroslip took place. It also suggests that at this smaller preload frictional damping was present to a limited

degree in all of the simulations. The only exceptions were those with an output smaller than, or equal to, the minimum amplitude of vibration analysed at this particular preload.

6.6) Penultimate Reduction in Joint Preload

Further reduction of the bolt preload was achieved by turning the bolt head a further 45° relative to the nut. This produced a further significant change in the dissipative behaviour of the joint. Another drop in the relative angular displacement of the top beam member relative to the bottom beam member was seen for a given amount of current supplied to the shaker. When considering the nominally rigid joint the amount of movement of the beam was the limiting factor to the power that could be supplied to the shaker without damage occurring to either it or the beam itself.




Figure 6.6.1 A comparison of the spectral strength of the background signal and the signal from the accelerometer located at the joint, on the free beam component, during an experimental test. Bolt preload is at the level obtained after two loosening and at the second smallest amplitude of vibration.

As the joint preload was reduced the damping produced by the joint increased and the beam displacement reduced drastically. The limiting feature of the experimental setup became the amount of power that could be supplied to the shaker without damage occurring to the Ammeter. Even at this maximum power supply the joint's range of angular displacement was still smaller than the minimum value used when the two components were rigidly clamped together. The response from the accelerometers around the joint became so small that the output from them was affected by the background noise in the system. This was a particular problem with the accelerometers whose output was passed through the signal amplifier. Figure 6.6.1 shows a spectrum of this background signal and compares it with the digital output from one of the accelerometers during an experimental test at this preload. What is immediately apparent is that there are peaks in the background signal at almost exactly the same frequency as the resonant frequency of the beam and its harmonics. The reason for this is that the beam resonated at 47.7Hz and the power supply has a frequency of 50Hz. Superharmonics of the power supply frequency can also be seen in the background signal. In the higher frequency ranges the superharmonics of the frictional vibration of the beam and the electrical power supply became more widely separated. However at the resonant frequency and the 3rd and 5th harmonics the combination of the two signals is problematic when the beam's angular displacement is small enough.

The power spectrum density plot of the displacement data also shows that the contribution from additional super harmonics could no longer be considered negligible. In the cases of the two largest displacement amplitudes the fifth superharmonic is very prominent, but in the smaller amplitude cases it is very hard to distinguish it from the background noise. The presence of the extra harmonics is far more clear from spectral density plots of the acceleration data (Figure 6.6.2). In the case shown, the harmonics of particular significance are the 3rd and 7th. However after integration of the signal only the fundamental and 3rd harmonic have a significant impact on the displacement signal. Time integration of the signal means that after each integration the effect of harmonics above the natural frequency of the beam diminishes. The variation of the signal amplitude at high frequencies (above the 5th superharmonic) in the displacement data shows smaller

variation relative to the fundamental frequency than that seen in the acceleration data. Consequently the relatively large signal at the 6th and 7th harmonics of the excitation frequency do not drastically alter the angular displacement behaviour.




Figure 6.6.2 Comparison of the spectrum density in angular acceleration and angular displacement signals. Results taken for bolt with twice reduced preload and second smallest amplitude of oscillation.

The relative angular acceleration signal also carries significant signals from the 5th and 6th harmonics and at the 3rd, 4th and 9th superharmonics of the power supply frequency. As well as being of diminished importance after integration, the Figure 6.6.2 is plotted using a logarithmic scale which suggests that even the impact on the acceleration is limited. The 3rd superharmonic (150Hz) of the power supply frequency is the only exception as there is still a clear peak in the displacement signal. Figure 6.6.2 also shows the effect of filtering the response well below the resonant frequency of the

beam. It is very prominent in the angular acceleration and yet almost totally eradicated in the angular displacement plot. Importantly, filtration of the signal responsible for the offset after integration does not have any influence on the peak associated with the fundamental frequency of the beam.

The interference of the background noise, and the increasing amount of sliding in existence in the joint can both be held responsible for the lack of cleanliness in the relative angular displacement of the two beams at the further reduced preload. This is particularly the case for the two smallest peak vibration amplitudes. Figure 6.6.3 shows that identifying properties and/or physical behaviour from these two simulations is virtually impossible.

Hysteresis loops obtained at larger peak-to-peak amplitudes of vibration are undoubtedly affected by the same problems associated with small overall amplitudes of vibration. However, from the intermediate to the largest input amplitude the loops exhibit the behaviour associated with frictional sliding more than those obtained with higher preloads. Each of the largest three loops show further flattening of the response as the amplitude increased. Despite the combined effect of viscous damping it is clear that the stiffness of the joint approaches zero at the largest amplitude of vibration. This implies that a point of macroslip is almost encountered despite the smaller overall displacements that occurred as a result of the increased damping relative to previous simulations. The loops are also more consistent in their stiffness from one point of velocity reversal to another. This suggests that the viscous effects at these small amplitudes are beginning to be masked by the frictional damping that is also present.

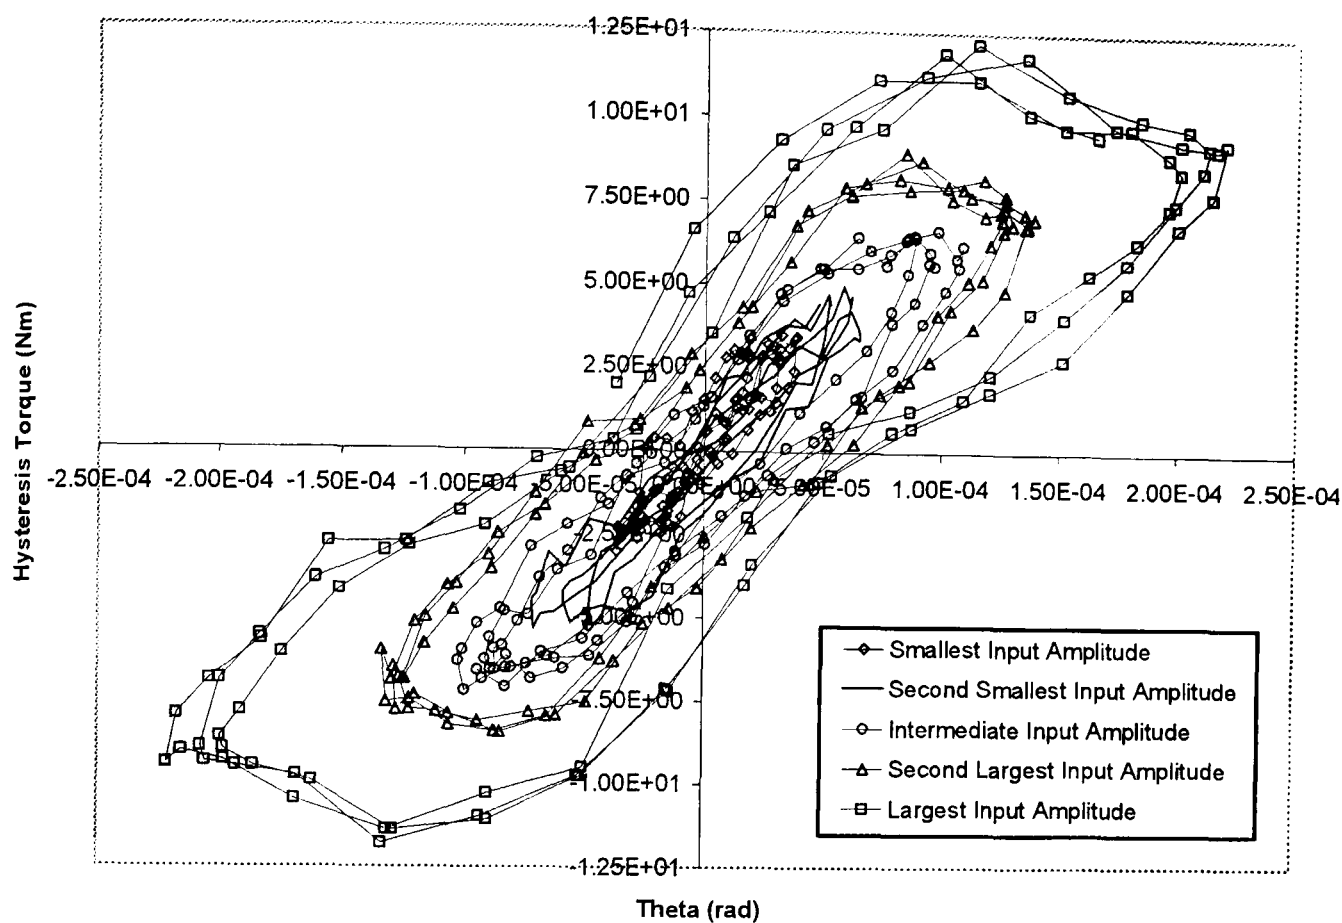


Figure 6.6.3 *Hysteresis loops obtained at various amplitudes of vibration when the bolt preload is reduced from the rigidly clamped condition by approximately 1/4 turn.*

The power spectrum of the output from the loosest joint condition encountered so far demonstrates a disproportionately large influence from the 3rd and 5th harmonics (Figure 6.6.4). These are most evident in the largest oscillations and in particular the test with the greatest peak-to-peak amplitude. In this case there is a definite concavity about the hysteresis loop that is directly attributable to the size of the 3rd harmonic. The two largest loops both show a similar profile with a relatively flat portion before velocity reversal and regions of comparable stiffness in between. Such behaviour is consistent with that seen in the finite element model where the same regions of microslip occurred in a pair of simulations. Consequently regions of similar stiffness were visible as parallel sections in both loops with the main difference being the amount of time spent with the lowest contact stiffness by the loop with the largest peak-to-peak amplitude. The intermediate hysteresis loop also shows a more gradual onset of microslip which is also consistent with the finite element tests presented earlier. Separating the gradual onset of microslip from viscous effects is difficult, yet the increased proportion of damping

provided by friction at this preload indicates that this is the phenomena most likely being observed.




Figure 6.6.4 Power spectrum plot achieved at a preload reduced from the rigidly clamped condition by two loosening the bolt by 1/4 turn.

Table 6.6.1 shows that the joint is much softer on average over a cycle than had been seen with larger preloads. Although the stiffness parameters obtained at the two smallest amplitude ranges are unreliable due to the background noise encountered, an undoubted average softening occurs as the amplitude increases. This, along with the shape of the hysteresis loops and the presence of odd superharmonics is the most significant evidence that microslip is taking place at this particular preload.

Table 6.6.1 also shows that the amount of energy dissipated for a given displacement range is not necessarily more at this particular preload than the tighter joint used previously. The principal suggestion of this is that despite the reduced amplitudes of

vibration, the optimum preload in terms of energy dissipated may lie closer to the condition where the bolt is only loosened by 1/8 turn rather than 1/4 turn. The presence of concavity resulting from the unusually large 3rd harmonic is significant in reducing the amount of energy dissipated in a given loop at the preload investigated here however. This may slightly distort the evidence for the preload that maximises the amount of energy dissipated. Further tests would be required where the input torque is more closely monitored, but the evidence shown so far seems to confirm that different preloads exist for optimising the amount of energy dissipated and the amount of movement in the joint. This is consistent with the findings of Beards and Woowat (1985) and is one of the most important characteristics of bolted joints.

Relative Applied Amplitude	Energy Dissipated (J/Cycle)	Maximum Angular Displacement (rad)	Average Cycle Stiffness from Peak to Peak Displacement (Nm/rad)
Smallest	1.0085×10^{-4}	4.69×10^{-5}	5.3731×10^4
Second Smallest	3.1718×10^{-4}	6.25×10^{-5}	6.9440×10^4
Intermediate	6.4250×10^{-4}	9.00×10^{-5}	7.556×10^4
Second Largest	1.4400×10^{-3}	1.35×10^{-4}	5.2519×10^4
Largest	3.1007×10^{-3}	1.95×10^{-4}	4.2641×10^4

Table 6.6.1 *A comparison of hysteresis loop characteristics found when a range of torques were applied to a joint with two reductions in the applied preload.*

6.7) Loosest Joint Condition

After turning the bolt-head a further approximately 45° relative to the nut the smallest preload was obtained before the joint interface began to separate. Reducing the preload yet further would result in a pinned condition which was not considered a legitimate range of joint loading conditions at the outset of the study. The preload associated with this final loosening of the bolt was therefore regarded as the loosest possible for the type of experiment being performed here. It amounted to a level of torque applied a little greater than that possible with hand tightening.

The lowest level of torque applied to the bolt would result in the smallest preload and in turn the lowest contact pressures at the joint interface. Both the degree of microslip and/or macroslip would be at their largest values for a given range of angular displacement at this preload. Figure 6.7.1 shows again that the maximum angular displacement achievable is at its lowest level for the safe amount of power that can be applied to the shaker. This is as a result of the increased amount of damping in the system relative to all the other preloads experienced thus far.

The cleanliness of the hysteresis loops in all cases has deteriorated from the cases where the clamping force was larger. Again the reasons likely to cause this are the increased amount of slippage in the joint, the large relative impact of background noise on the results and the increased presence of superharmonics. A change in behaviour of the largest amplitude case can be seen from the similar maximum amplitude case at the previous level of preload (shown in Figure 6.6.3). The duration that the cycle remains largely flat before velocity reversal has increased as the bolt preload is reduced. This implies that regions of the joint interface are sliding for a greater duration during each cycle than at higher preloads. Again this result is consistent with the general microslip behaviour illustrated by the finite element simulations. At the two lowest amplitudes of vibration less relevant information of this kind can be obtained as the amount of background noise contaminating the signal becomes too great. Generally though it can be seen that increased amounts of energy are dissipated as the amplitude of the oscillations

increases. What proportion of the energy is dissipated due to viscous-like damping is unclear.

Relative Applied Amplitude	Energy Dissipated (J/Cycle)	Maximum Angular Displacement (rad)	Average Cycle Stiffness from Peak to Peak Displacement (Nm/rad)
Smallest	3.0209×10^{-5}	2.90×10^{-5}	5.5517×10^4
Second Smallest	1.3203×10^{-4}	4.42×10^{-5}	7.5539×10^4
Intermediate	2.8201×10^{-4}	6.53×10^{-5}	5.1302×10^4
Second Largest	9.0801×10^{-4}	1.00×10^{-4}	5.2100×10^4
Largest	2.667×10^{-3}	2.13×10^{-4}	3.1596×10^4

Table 6.7.1 *A comparison of hysteresis loop characteristics for a range of displacements when applied to a joint with preload reduced by 3/8th relative turn of the bolt head.*

The larger amplitude vibrations show the friction characteristics more clearly. Notably the softening characteristic of the interface can be seen as the amplitude of oscillations is increased. These are also illustrated quantitatively in Table 6.7.1. Although the exact points of velocity reversal are very difficult to identify if either the points of maximum hysteresis torque or maximum angular displacement are used the same softening behaviour can be seen. It is almost undoubtedly the presence of some viscous behaviour in the joint that causes the peak amplitude to not coincide with the point of peak hysteretic restoring torque. The only other factor that may be of influence is the inertial effects of the beam and these can be considered small.

Table 6.7.1 demonstrates that the rate of increase of energy dissipated per cycle is far greater than the rate of increase in angular displacement. Coupled with the largest softening effect of the joint, the disproportionate increase in energy dissipated points to a joint that is highly nonlinear. Such characterisation is reasonable from the standpoint of the increased amount of frictional sliding that takes place in the presence of the lowest preload and contact pressure experienced so far. Despite this the amount of energy

dissipated is still not as large for a given range of angular displacement as that found when the joint was only loosened by a single increment from its rigidly clamped state. This offers further confirmation that the joint configuration with most flexibility and hence microslip taking place may not be the one that provides the optimal amount of energy dissipation.




Figure 6.7.1 *Hysteresis obtained with the smallest amount of preload in the bolt at various peak-to-peak amplitudes of vibration.*

The hysteresis shows that the impact of the third harmonic is greater at the lowest preload than in the test utilising a preload one increment larger. Even with lower amplitude vibrations the concavity produced by the large relative magnitude of the third harmonic can be seen. With the larger preload (Figure 6.7.1) only the largest amplitude oscillations displayed concavity to a significant degree. Some of the noise of the cycles may also be further superimposed oscillations caused by the presence of the 5th and 7th harmonics of the frequency of vibration. It is also possible that at this smallest value of

preload the largest amplitudes of vibration also contain a significant input from the 2nd harmonic for the first instance in the series of tests carried out. Figure 6.7.2 shows the different harmonic's impact on the overall response of the top joint member relative to the bottom joint member. The smallest amplitude response, in particular shows the difficulty in extracting a clean signal from the background noise that is present. Even at high frequencies the variation in amplitude of all frequency content is still high, and distinguishing small peaks at the odd harmonics of excitation frequency is not possible.




Figure 6.7.2 *Spectral density plot for the smallest amount of preload in the bolt when the joint was subjected to a variety of amplitudes of input torque.*

6.8) Matching Jenkins Elements to Experimental Results

It has already been shown that the Jenkins element model can be used effectively to match the hysteresis loops generated using an isolated joint and the finite element method. To verify that the same matching process can be carried out on physical joints, the Jenkins element models were fitted to a selection of the hysteresis loops generated experimentally. The experimental loops that were matched were chosen as they displayed a perceivable amount of microslip and demonstrated a relatively clean response. A number of different simulations were found to satisfy these requirements at two different clamping forces. At the first reduction in clamping force from the rigid connection the three largest amplitudes of vibration were matched. Similarly, at the next decrease in the clamping load, the hysteresis loops with the three largest ranges of angular displacement were selected for replication.

At the larger of the two preloads, the impact of a certain amount of viscous damping made identification of clear points of velocity reversal very difficult. Where the hysteresis loop had its smallest gradient was not necessarily equivalent to the peak magnitude in angular displacement or restoring torque. Judgement had to be used to select these reversal points and then the following discretisation points on the experimental hysteresis loops. At the larger preload each hysteresis loop was visually discretised into 4 segments, creating a model with three Jenkins elements and a permanent spring.

The hysteresis output from these Jenkins element models are shown in Figure 6.8.1. Although the slightly elliptical nature of the experimental hysteresis loops was not entirely reproduced by the Jenkins element model, the overall system dynamics were well represented. Points of velocity reversal were well matched ensuring that the average stiffness over a cycle was accurate. The initial, fully sticking, contact stiffness was also found to be relatively consistent for the preload in question. Although varying, the initial stiffness remained a similar order of magnitude with discrepancies attributable to the subjective methods used to discretise the experimental hysteresis loops. All of the

hysteresis loops when discretised also had a large region of similar stiffness that existed between velocity reversals. The angular displacement range of this region was not constant across each cycle, but its presence suggested a similar amount of microslip was taking place. The discretisation process facilitating the use of Jenkins elements helped to highlight this behaviour, which was consistent with that identified in the finite element model.




Figure 6.8.1 *Matching of Jenkins element models to the experimentally obtained hysteresis loops at the first reduction in preload from the rigidly clamped state.*

Good agreement between the experimental results and the Jenkins element hysteresis was also demonstrated in the amount of energy dissipated over a cycle. This would be significant in terms of establishing an equivalent damping factor for the two sets of hysteresis. Such close agreement was possible largely due to the smaller variation in contact stiffness seen for much of the experimental cycle allowing regions of constant stiffness to better approximate the experimental behaviour. Quantitative results showing the effectiveness of the Jenkins element model at this preload are given in Table 6.8.1

Amplitude of Oscillation	Initial Stiffness (Nm/rad)	Jenkins Element Energy Dissipated (J/Cycle)	Experimental Energy Dissipated (J/Cycle)	Percentage Difference
Largest	149×10^3	2.3240×10^{-2}	2.3420×10^{-2}	-0.77%
Second Largest	218×10^3	2.0105×10^{-2}	2.0269×10^{-2}	-0.81%
Intermediate	130×10^3	1.5313×10^{-2}	1.3877×10^{-2}	10.35%

Table 6.8.1 Comparison of the agreement of experimental results and Jenkins element approximations at the first reduction in preload.

Fitting Jenkins elements to the hysteresis loops obtained at the second reduction in preload was more problematic. Due to the large presence of particularly the 3rd harmonic in the experimental output, the experimental hysteresis loops tended towards an s-shaped profile with slight hardening taking place before the softening behaviour associated with microslip became apparent. Jenkins elements were only capable of representing the softening behaviour. This meant that when the experimental hysteresis loop was discretised, the narrower part of the loop lay inside the region of constant stiffness defined by the Jenkins element. The outcome of this was an over estimation of the energy dissipated by the Jenkins element model. However, this was compensated in the lowest stiffness region before velocity reversal. Here, the Jenkins element approximation lay *inside* the experimental hysteresis. The profile of the experimental hysteresis did not show a continuous increase in restoring torque to the point of velocity reversal. Instead the experimental loop dipped slightly which meant an approximation of the response had to be considered just before velocity reversal.

In all of the cases at the second reduction in preload, the Jenkins element response was approximated by 2-elements and a permanent spring. The hysteresis output in Figure 6.8.2 shows that a good overall agreement was found between the hysteresis generated by the two different sources. It also highlights the comparable contact interface stiffness that existed between all of the different amplitudes of oscillation at this level of preload.




Figure 6.8.2 *Matching of Jenkins element models to the experimentally obtained hysteresis loops at the second reduction in preload from the rigidly clamped state.*

The acceptable agreement in energy dissipated between the experimental hysteresis and Jenkins element approximation is recorded in Table 6.8.2. The fully sticking contact stiffness in existence immediately after velocity reversal is also shown to lie within the same order of magnitude for each of the experimental cases. This is illustrated by the initial stiffness of the Jenkins elements and demonstrates the consistency of the results obtained experimentally and in the modelling process.

Amplitude of Oscillation	Initial Stiffness (Nm/rad)	Jenkins Element Energy Dissipated (J/Cycle)	Experimental Energy Dissipated (J/Cycle)	Percentage Difference
Largest	101.74×10^3	3.6912×10^{-3}	3.1007×10^{-3}	19%
Second Largest	266.67×10^3	1.9195×10^{-3}	1.4400×10^{-3}	33%
Intermediate	209.00×10^3	0.5831×10^{-3}	0.6425×10^{-3}	-9.25%

Table 6.8.2 Comparison of the agreement between experimental results and Jenkins element approximations, at the second reduction in bolt preload.

6.9) Matching the Bouc-Wen Model to Experimental Results

As with the Jenkins element model, it was necessary to verify that a physical joint's behaviour could be adequately represented by the Bouc-Wen model of hysteresis. The same experimental hysteresis loops were used to compare the two sets of results as were used when matching the Jenkins element hysteresis. The parameters for matching the Bouc-Wen results were extracted by using the same initial stiffness after velocity reversals as the Jenkins element model. This represented an assumed region of totally sticking contact. The same estimated points of velocity reversal were also used in the Bouc-Wen parameter extraction as were used in the Jenkins element model. This enabled α and β to be defined using the same parameter extraction process outlined in Section 5.3. In all of the Bouc-Wen models the parameter n was taken to be either 1 and 2.




Figure 6.9.1 Experimental results approximated by Bouc-Wen models of order $n=1$ and $n=2$ with a preload obtained by a single loosening of the bolt.

The comparison of the results obtained for the larger of the two preloads used are shown in Figure 6.9.1. The elliptical profile of the experimental hysteresis data is not well represented by the Bouc-Wen model. However, the initial stiffness and average stiffness of the Bouc-Wen approximation are well defined in the extraction process and therefore offer good agreement. Although the overall profiles of the hysteresis loops were not ideal, the energy dissipated over a complete cycle was the most significant factor in potentially finding a net damping coefficient for the joint in question. The agreement in this criterion was better than the match in profile between the Bouc-Wen approximations and the experimental results.

Table 6.9.1 shows that it was the Bouc-Wen models with $n=1$ that offered the closest agreement with the experimentally obtained energy dissipation. This was consistent with the finite element comparisons with the Bouc-Wen model. Despite the relatively poor approximation of the experimental hysteresis profile, the Bouc-Wen model offers a reasonable estimation of the overall energy dissipation. Particularly at this higher clamping load, the Bouc-Wen model was not suited to matching the experimental, elliptical profile, and a reasonable approximation of the enclosed area was the best that could be offered.

At the further reduced preload that was matched the Bouc-Wen model had a significant advantage over the Jenkins element model in matching the experimental behaviour. By choosing the appropriate values of α and β , it was possible to generate a Bouc-Wen hysteresis loop that demonstrated both hardening and softening behaviour between points of velocity reversal. This type of behaviour has been described as occurring at the smaller bolt preloads in the experimental analysis. Consequently a superior match of the experimental hysteresis profile was possible when the superharmonics of the excitation frequency became prominent in the experimental response. In all cases, except the smallest angular displacement range, the most accurate model of the amount of energy dissipated per cycle was obtained with parameter $n=1$.

Angular Displacement Range	n	Bouc-Wen Dissipation (J/Cycle)	Experimental Dissipation (J/Cycle)	Percentage Difference
Largest	1	1.7825×10^{-2}	2.3420×10^{-2}	-23.88%
Largest	2	3.8871×10^{-2}	2.3420×10^{-2}	65.97%
Second Largest	1	1.9891×10^{-2}	2.0269×10^{-2}	-1.86%
Second Largest	2	2.3934×10^{-2}	2.0269×10^{-2}	18.08
Intermediate	1	1.2972×10^{-2}	1.3877×10^{-2}	-6.52%
Intermediate	2	1.6026×10^{-2}	1.3877×10^{-2}	15.49%

Table 6.9.1 Comparison of the energy dissipated by an experimental joint and the Bouc-Wen models at a preload reduced from the rigidly clamped condition by one increment.

Angular Displacement Range	n	Bouc-Wen Dissipation (J/Cycle)	Experimental Dissipation (J/Cycle)	Percentage Difference
Largest	1	3.2209×10^{-3}	3.1007×10^{-3}	3.88%
Largest	2	3.6645×10^{-3}	3.1007×10^{-3}	18.18%
Second Largest	1	1.8858×10^{-3}	1.4400×10^{-3}	30.96%
Second Largest	2	1.9811×10^{-3}	1.4400×10^{-3}	37.58%
Intermediate	1	0.4571×10^{-3}	0.6425×10^{-3}	-28.86%
Intermediate	2	0.5676×10^{-3}	0.6425×10^{-3}	-11.66%

Table 6.9.2 Comparison of the energy dissipated by an experimental joint and the Bouc-Wen models at a preload reduced from the rigidly clamped condition by two increments.

Figure 6.9.2 shows that the qualitative agreement in the profile of the hysteresis loops was superior using the Bouc-Wen model than when approximating regions of decreasing stiffness using Jenkins elements (Figure 6.8.2). However, the overall relative

error in the amount of energy dissipated by the Jenkins element model and the Bouc-Wen model existed within the same bounds at the lower magnitude of preload. This suggested that for smaller preloads the Bouc-Wen model was potentially a better method of representing the experimental hysteresis. Such a conclusion became even more significant when the number of parameters that define the model, and the numerical efficiency of the two different methods, were considered. In cases where there were superharmonics present, the Bouc-Wen model was shown to have the potential to match the experimental hysteresis more fully than the Jenkins element approximation. However, until Bouc-Wen model parameters are capable of being identified to give this superior performance, the flexibility that the model offers cannot be used to its full potential.




Figure 6.9.2 *Experimental results approximated by Bouc-Wen models of order $n=1$ and $n=2$ with a preload obtained by twice loosening the bolt from the rigidly clamped condition.*

Chapter 7. Conclusions

7.1) Conclusions

In a structure of components connected by bolted joints, when subjected to dynamic excitation the main source of damping can be frictional energy dissipation. This is particularly the case in space structures where design specifications result in structures with low inherent material damping, and no means of energy dissipation through atmospheric effects.

Joints that maintain their structural integrity can still dissipate energy through frictional sliding at their interfaces. The contact pressure distribution at the joint interface caused by a bolt preload results in the decline of normal force in a radial direction from a maximum value near the bolt-hole. Where the contact pressure is lowest, sliding contact may be initiated in local area whilst regions of higher contact pressure remain in their sticking state. This phenomenon is known as microslip and was investigated with respect to an isolated joint subjected to a harmonically applied torque.

A literature survey revealed that one of the main difficulties in investigating joints with microslip was the difference in scales. Microslip takes place on a very local level and manifests itself in deflections that are orders of magnitude smaller than a representative joint dimension. The impact of microslip can then be displayed in the modified behaviour of a structure several orders of magnitude larger than the joints that connect its components. The difference in scale between the phenomena that causes the structural damping and the size of the structure that experiences the damping is one of the largest problems associated with incorporating joint friction into a model. Further significant factors that make joint friction difficult to analyse and model are the highly nonlinear nature of friction itself and the complexities of representing this nonlinear behaviour in 3-Dimensional situations.

In an ideal situation an equivalent model where the joint friction's effects could be represented by a single linear spring and linear viscous damping component could be created. To achieve this a detailed finite element model was constructed that enabled microslip to be analysed and characterised. Equivalent models representing the hysteretic energy dissipation of the detailed finite element model were then evaluated. Two equivalent models were investigated. The first utilised Jenkins elements to represent the local slipping behaviour associated with microslip. The second, Bouc-Wen, model was a smooth hysteretic representation of the loss of stiffness between velocity reversals of a single joint. Both models demonstrated joint hysteresis using several orders of magnitude fewer degrees of freedom than the detailed finite element joint.

The finite element model showed that the microslip behaviour of the isolated joint was different depending on the preload that was applied to it. A bolt preload generated contact pressure at the joint that declined to a near zero value near the edge of the interface. In this situation, the onset of microslip occurred at a relatively small value of applied torque. Increasing the magnitude of the bolt preload also increased the level of applied torque required to initiate microslip. The microslip that existed propagated over a larger area of the joint interface as the applied torque was increased. This occurred as regions of larger contact pressure made the transition from sticking to sliding contact. When the interface was on the point of macroslip the hysteresis loop that characterised the energy dissipated showed a relatively smooth transition in the loss of joint stiffness between points of velocity reversal. At higher preloads and smaller applied torque, the hysteresis loop showed smaller variation in the contact stiffness and could be characterised as bilinear. A similar bilinearity existed when the contact pressure distribution was more uniform. In these instances only two distinct regions of microslip existed before macroslip was initiated.

Jenkins elements offered a direct physical insight into the hysteretic behaviour of the detailed finite element hysteresis. The stiffness provided by the finite element model when no microslip existed was equivalent to the sum of the stiffness of the permanent spring and each of the Jenkins elements that were connected in parallel. On the point of

macroslip, the stiffness of the reduced order model was provided only by a permanent spring. This residual stiffness was equivalent to the torsional stiffness provided by the bolt nut component in the finite element model.

Visually discretising the finite element hysteresis loop into regions of approximately constant stiffness allowed its accurate reproduction by a model of just four Jenkins elements and a permanent spring. The points of velocity reversal, stiffness immediately before and after velocity reversal and energy dissipated per cycle agreed well with the finite element joint model. It was also possible to predict with reasonable accuracy the hysteretic behaviour of the finite element joint at different applied torques to the case used in the extraction of the Jenkins element parameters. Prediction of this sort was only successful when the full range of macroslip was visible and well matched to the reference finite element loop.

The hysteretic response of the finite element model could be predicted at different levels of bolt preload. When the bolt preload was applied in the finite element model, the contact pressure at points on the interface was directly in proportion to the magnitude of the preload. Altering the resistance, and consequently the normal force applied to the Coulomb elements in the same proportion to the change in bolt preload, generated a good approximation of the finite element behaviour. The Jenkins element model also demonstrated that during free vibrations caused by an initial deflection, the hysteresis loops would be generated about points that were not centred about an angular displacement of 0(rad). This was attributed to changes in contact conditions as greater or fewer elements made the transition from sticking to sliding contact when the amplitude of vibration was not periodic. Once the amplitude of the angular displacement became small enough the model reached a conservative steady state vibration.

The Bouc-Wen model had an advantage over the Jenkins element model in that a smooth transition from total sticking to macroslip was achieved with only four defining parameters. The Bouc-Wen model was also more flexible in identifying hysteresis where the response was typical of a hardening then softening spring. Beyond the initial fully

sticking contact interface stiffness however, the parameters of the Bouc-Wen model did not have an immediate physical significance in the finite element model. By scaling certain parameters with respect to each other, the point where a complete loss of contact stiffness was reached could be identified.

When matched to the finite element model, the Bouc-Wen approximation of hysteresis offered good agreement in the overall amount of energy dissipated and contact stiffness of the fully sticking joint interface. Apart from the points of velocity reversal, the model was not constrained to the finite element hysteresis loop at any other points during the cycle. The behaviour was therefore less representative of the time domain response than that achieved with Jenkins elements.

The Bouc-Wen model achieved its smooth transition from sticking to virtual macroslip with only four parameters used to define it. The model did not require the identification of the sticking or sliding state at individual time increments. This made it computationally more efficient than a Jenkins element model with a smooth hysteretic profile. However, the lack of physical significance of the Bouc-Wen parameters and the overall smoothing of the joint hysteresis made the prediction of the finite element hysteresis at different conditions quite inaccurate. This could only be improved with a more sophisticated method for identifying the parameters used to define the Bouc-Wen model from the reference finite element hysteresis.

To validate the findings of finite element macroslip behaviour and the models used to approximate it, experimental tests were performed on a single lap joint. Measurements enabled local identification of the macroslip behaviour from a joint in a cantilever beam that was sinusoidally excited. Good qualitative agreement was found between the behaviour exhibited by the finite element model and the experimental behaviour. At high bolt preloads, the energy dissipated was viscous in nature and macroslip was not detected at any input amplitude of vibration. Once the bolt preload was reduced macroslip became apparent and the amount of energy dissipated for a given amplitude of response increased. Further reductions in bolt preload resulted in increased

damping and microslip effects dominating the characteristics of the experimentally obtained hysteresis. At the two lowest levels of preload, the magnitude of the response attributed to the third harmonic of the excitation frequency became large. As a result, the experimental hysteresis displayed both hardening then softening behaviour between velocity reversals.

Both Jenkins elements and the Bouc-wen model were shown to match well the defining characteristics of the experimental behaviour. The Jenkins elements provided a better overall agreement in the amount of energy dissipated than the Bouc-Wen model. However, the Bouc-Wen model was capable of reproducing the hardening then softening behaviour observed at lower experimental preloads. This flexibility could not be reproduced by the Jenkins element model, but further improvements in the parameter extraction of the Bouc-Wen model would be required before its full potential is realised.

7.2) Future Work

As with any research that has reached performed, the work that has been carried out to this point has illuminated several different areas that could be investigated further. With the exception of the work performed using Jenkins elements, all areas of the research that have been undertaken could be expanded and developed. A pursual of the work performed so far would be necessary to address both limitations of the methods used to this point and to fulfil the potential of the processes that have been introduced.

With the exception of the Jenkins element method, all of the procedures carried out could be expanded to broaden the knowledge that exists with regard to bolted joints and microslip. The finite element method has been shown to offer considerable insight that was not available through experimental measurement. Principally the advantage of the method lay in its ability to reveal information about the model behaviour actually on the joint interface. This behaviour ranged from the contact pressure to nodal accelerations and displacements. Further potential developments to the finite element joint model are wide ranging. As the contact analysis has been shown to be reliable, more complicated interactions could be incorporated into the model. These range from contact interfaces present but ignored in the current model to producing detailed models of threaded connections. The finite element model could also be developed to include a more sophisticated friction law. Principally the implementation of a velocity and pressure dependent coefficient of friction would address the findings of many other researchers.

The finite element method once validated for the simple lap joint under investigation here could be expanded to model more complicated joint configurations. These include different joint geometry, joints with more than a single contact interface and contact that takes place in a number of different planes simultaneously. By producing finite element models of this kind the cost of building experimental components would be saved and fabrication time would be virtually eliminated.

Of all of the modelling processes the Bouc-wen model was shown to be the most efficient and flexible in representing microslip. Its main limitation manifested itself as a poor approximation of the finite element or experimental results at all points on the hysteresis loop. Before the full potential of the model can be realised further work would be required in extracting the parameters that defined the model. In particular, the identification of parameters that describe fully the hardening and softening behaviour demonstrated in certain experimental cases would be of significant benefit.

The experimental work that has been carried out to this point was vital in confirming qualitatively the microslip mechanism for dissipating energy. One area where there was a limitation in the experimental design was the lack of means of measuring the preload in the bolt. Quantifying this load would greatly enhance the relevance and applicability of results obtained. To achieve this a load cell would provide the data that was required. Tests, using finite element models, would be necessary to confirm that the presence of a load cell did not significantly change the characteristics of the contact interface. A further improvement to the experimental rig would be obtained by resonating the beam at a frequency well away from the 50Hz power supply frequency. An immediate means of achieving this would be to add a mass to the free end of the composite beam.

Contact pressure distributions were an automatic by-product of the finite element contact analysis. An inability to physically measure this was one of the fundamental limitations of the experimental investigation. One solution to establishing the contact pressure distribution in the experimental joint would be the application of pressure sensitive films. A general trend in the contact pressure distribution could therefore be established at different levels of preload to outline the variation over the contact interface and dependence on the magnitude of static preload applied.

The ultimate goal of the research being carried out is to reduce the complicated nonlinear joint behaviour to an equivalent pair of linear parameters that describe adequately the joint stiffness and energy dissipation. Development of a methodology for

converting Bouc-Wen models to a single stiffness and damping coefficient would provide a significant tool for analysing multiple jointed structures.

The behaviour of multiple jointed structures was the main source of inspiration for the work carried out here. There is scope for the experimental analysis of structures with several joints of the type used here. These results could be compared with finite element simulations where the reduced order models replace the detailed joint model. By substituting a detailed joint with an equivalent Bouc-Wen model, and using existing beam and truss elements, the full potential of the different scales of investigation carried out here will have been met. Ultimately the joints in a built up finite element body could be replaced by the equivalent linear damping and stiffness parameters. In a commercial finite element package such as ABAQUS stiffness and damping components are provided as standard. The modal response from such a built up model could then be compared with modal tests performed on an experimental rig.

It is this potential to effectively predict and linearise the local joint behaviour for which the work carried out here provides the impetus. Considerable savings in terms of computation time and experimental test costs would be achieved. This would be combined with development of the analytical techniques initiated here and a deeper understanding of the frictional phenomena that exist in bolted joints.

7.3) Summary of Contribution

Bolted joints have been investigated for a number of years by different researchers (Chapter 2). Investigations have been performed to identify equivalent stiffness and damping parameters for the joint in question. Research had also been performed on the hysteresis generated by microslip in bolted joints

The work presented here combines the detailed analysis of microslip in bolted joints with an approach that progresses towards a single degree-of-freedom linearised model. Advances have been made in the detailed investigation of microslip by using a commercial finite element package to perform the complicated 3-Dimensional contact algorithms. Physical features such as the bolt-hole, bolt preload and the applied torque have been used in dynamic 3-Dimensional contact simulations for what is believed to be the first time.

Jenkins elements have been extensively used to model microslip behaviour in the past. The developments presented here provide a means of finding the most efficient match between the microslip behaviour of a joint and an equivalent Jenkins element model. Prediction of different hysteresis response is also presented for different loading conditions of the joint from which parameters were extracted. It is believed that this represents an advancement of previous models using Jenkins elements.

The Bouc-Wen model has been used in the past to model the hysteretic response to random vibration. From the literature search that has been conducted it is believed that this thesis represents the first time the model has been used to describe bolted-joint hysteresis.

Many experiments have been performed on bolted joints to identify microslip. The measurement technique presented herein is a novel approach to measuring directly the localised response of the joint interface. Resonant input conditions have been used in the past to generate microslip in a bolted joint. These methods have been applied to joints

that were deliberately isolated. The method used here is versatile enough to be used on a component that required no isolation or unusual boundary conditions.

References

- Anderson, J. R. and Ferri, A. A., "Behaviour of a single-degree-of-freedom system with a generalized friction law", *Journal of Sound and Vibration*, Vol. 140, No. 2, pp. 287-304, 1990.
- Andreus, U. and Casini, P., "Dynamics of friction oscillators excited by a moving base and/or driving force", *Journal of Sound and Vibration*, Vol. 245, No. 4, pp. 685-699, 2001.
- Barber, J. R. and Ciavarella, M., "Contact mechanics", *International Journal of Solids and Structures*, Vol. 37, pp. 29-43, 2000.
- Basava, S. and Hess, D. P., "Bolted joint clamping force variation due to axial vibration", *Journal of Sound and Vibration*, Vol. 210, No. 2, pp. 255-265, 1998.
- Beards, C. F., "The damping of structural vibration by controlled interfacial slip in joints", *Journal of Vibration, Acoustics, Stress and Reliability in Design*, Vol. 105, pp. 369-373, 1983.
- Beards, C. F., "Damping in structural joints", *The Shock and Vibration Digest*, Vol. 24, No. 7, pp. 3-7, 1992.
- Beards, C. F. and Neroutsopoulos, A. A., "The control of structural vibration by frictional damping in electro-discharge machined joints", *Journal of Mechanical Design*, Vol. 102, pp. 54-57, 1980.
- Beards, C. F. and Williams, J. L., "The damping of structural vibration by rotational slip in joints", *Journal of Sound and Vibration*, Vol. 53, No. 3, pp. 333-340, 1977.
- Beards, C. F. and Woowat, A., "The control of frame vibration by friction damping in joints", *Journal of Vibration, Acoustics, Stress and Reliability in Design*, Vol. 106, pp. 26-32, 1985.
- Beer, G., "An isoparametric joint/interface element for finite element analysis", *International Journal for Numerical Methods in Engineering*, Vol. 21, pp. 585-600, 1985.

Bengisu, M. T. and Akay, A., "Stability of friction-induced vibrations in multi-degree-of-freedom systems", *Journal of Sound and Vibration*, Vol. 171, No. 4, pp. 557-570, 1994.

Bengisu, M. T. and Akay, A., "Relation of dry-friction to surface roughness", *Journal of Tribology*, Vol. 119, pp. 18-25, 1997.

Bengisu, M. T. and Akay, A., "Stick-slip oscillations: Dynamics of friction and surface roughness", *Journal of the Acoustical Society of America*, Vol. 105, No. 1, pp. 194-205, 1999.

Berger, E. J., "Friction modeling for dynamic system simulation", *Applied Mechanics Reviews*, Vol. 55, No. 6, pp. 535-577, 2002.

Bickford, J. H., "An Introduction to the Design and Behaviour of Bolted Joints", *Marcel Dekker*, New York, 1981.

Bindemann, A. C. and Ferri, A. A., "Large amplitude vibration of a beam restrained by a non-linear sleeve joint", *Journal of Sound and Vibration*, Vol. 184, No. 1, pp. 19-34, 1995.

Bouc, R., "Forced vibration of mechanical systems with hysteresis (Abstract)", *4th Conference on Nonlinear Oscillations*, Prague, pp. 315, 1967.

Bradley, T. L., Lardner, T. J. and Mikic, B. B., "Bolted joint interface pressure for thermal contact resistance", *Journal of Applied Mechanics*, Vol. 38, pp. 542-545, 1971.

Brockley, C. A. and Ko, P. L., "Quasi-harmonic friction-induced vibration", *Journal of Lubrication Technology*, pp. 550-556, 1970.

Brown, A. F. C. and Hickson, V. M., "A photo-elastic study of stresses in screw threads", *Proceedings of the Institution of Mechanical Engineers*, pp. 605-612, 1953.

Canudas de Wit, C., Olsson, H., Åström, K. J. and P., L., "A new model for control of systems with friction", *IEEE Transactions on Automatic Control*, Vol. 40, No. 3, pp. 419-425, 1995.

- Chen, S. H., Cheung, Y. K. and Xing, H. X., "Nonlinear vibration of plane structures by finite element and incremental harmonic balance method", *Nonlinear Dynamics*, Vol. 26, pp. 87-104, 2001.
- Crawley, E. F. and Aubert, A. C., "Identification of nonlinear structural elements by force-state mapping", *ALAA Journal*, Vol. 24, No. 1, pp. 155-162, 1986.
- De Matteis, G., Mandara, A. and Mazzolani, F. M., "T-stub aluminium joints: influence of behavioural parameters", *Computers and Structures*, Vol. 78, pp. 311-327, 2000.
- Den Hartog, J. P., "Forced vibrations with combined Coulomb and viscous friction", *Journal of Applied Mechanics*, Vol. APM-53-9, pp. 107-115, 1931.
- Ding, Q., Leung, A. Y. T. and Cooper, J. E., "Dynamic analysis of a self-excited hysteretic system", *Journal of Sound and Vibration*, Vol. 245, No. 1, pp. 151-164, 2001.
- Doege, E., Kaminsky, C. and Bagaviev, A., "A new concept for the description of surface friction phenomena", *Journal of Materials Processing Technology*, Vol. 94, pp. 189-192, 1999.
- El-Zahry, R. M., "Optimum design of preloaded bolted joint under harmonic excitation", *Journal of Sound and Vibration*, Vol. 108, No. 3, pp. 455-470, 1986.
- Esmailzadeh, E., Chorashi, M. and Ohadi, A. R., "Analysis of preloaded bolted joints under exponentially decaying pressure", *Journal of Pressure Vessel Technology*, Vol. 118, pp. 393-398, 1996.
- Esteban, J. and Rogers, C. A., "Energy dissipation through joints: theory and experiments", *Computers and Structures*, Vol. 75, pp. 347-359, 2000.
- Feeny, B., Ardèshir, G., Hinrichs, N. and Popp, K., "A historical review on dry friction and stick-slip phenomena", *Applied Mechanics Reviews*, Vol. 51, No. 5, pp. 321-341, 1998.

- Feeny, B. and Moon, F. C., "Chaos in a forced dry-friction oscillator: Experiments and numerical modelling", *Journal of Sound and Vibration*, Vol. 170, No. 3, pp. 303-323, 1994.
- Ferri, A. A., "Friction damping and isolation systems", *Journal of Vibration and Acoustics*, Vol. 117/B, pp. 196-206, 1995.
- Ferri, A. A. and Bindemann, A. C., "Damping and vibration of beams with various types of frictional support conditions", *Journal of Vibration and Acoustics*, Vol. 114, pp. 289-296, 1992.
- Ferri, A. A. and Heck, B. S., "Analytical investigation of damping enhancement using active and passive structural joints", *Journal of Guidance, Control and Dynamics*, Vol. 15, No. 5, pp. 1258-1264, 1992.
- Ferri, A. A. and Heck, B. S., "Vibration analysis of dry friction damped turbine blades using singular perturbation theory", *Journal of Vibration and Acoustics*, Vol. 120, pp. 588-595, 1998.
- Galvanetto, U. and Bishop, S. R., "Stick slip vibrations of a two degree-of-freedom geophysical fault model", *International Journal of Mechanical Sciences*, Vol. 36, No. 8, pp. 683-698, 1994.
- Gaul, L., "Wave transmission and energy dissipation at structural and machine joints", *Journal of Vibration, Acoustics, Stress and Reliability in Design*, Vol. 105, pp. 489-496, 1983.
- Gaul, L., "Identification of nonlinear structural joint models and implementation in discretized structure models", *The 1987 ASME Design Technology Conferences-11th Biennial Conference on Mechanical Vibration and Noise*, Boston, Massachusetts, pp. 213-219, 1987.
- Gaul, L. and Lenz, J., "The influence of microslip on the dynamic behaviour of bolted joints", *13th International Modal Analysis Conference*, Nashville, Tennessee, pp. 1995.
- Gaul, L. and Lenz, J., "Nonlinear dynamics of structures assembled by bolted joints", *Acta Mechanica*, Vol. 125, No. 1-4, pp. 169-181, 1997.

- Gaul, L., Lenz, J. and Sachau, D., "Active damping of space structures by contact pressure control in joints", *Mechanical Structures and Machines*, Vol. 26, No. 1, pp. 81-100, 1998.
- Gaul, L. and Nitsche, R., "Friction control for vibration suppression", *Mechanical Systems and Signal Processing*, Vol. 14, No. 2, pp. 151-165, 2000.
- Gaul, L. and Nitsche, R., "The role of friction in mechanical joints", *Applied Mechanics Reviews*, Vol. 54, No. 2, pp. 93-106, 2001.
- Gaul, L. and Nitsche, R., "Damping control in systems assembled by joints", *International Conference on Structural System Identification*, Kassel, Germany, pp. 37-48, 2001.
- Gautham, B. P., Sarma, G. B. and Ganesan, N., "A critical comparison of two methods for solving elastic contact problems with friction", *Computers and Structures*, Vol. 41, No. 1, pp. 93-97, 1991.
- Griffin, J. H. and Menq, C.-H., "Friction damping of circular motion and its implications to vibration control", *Journal of Vibration and Acoustics*, Vol. 113, pp. 225-229, 1991.
- Groper, M., "Microslip and macroslip in bolted joints", *Experimental Mechanics*, pp. 171-174, 1985.
- Guyan, R. J., "Reduction of stiffness and mass matrix", *AIAA Journal*, Vol. 3, No. 2, pp. 380, 1965.
- Hess, D. P. and Sudhirkashyap, S. V., "Dynamic loosening and tightening of a single bolt assembly", *Journal of Vibration and Acoustics*, Vol. 119, pp. 311-316, 1997.
- Hess, D. P. and Wagh, N. J., "Chaotic normal vibrations and friction at mechanical joints with nonlinear elastic properties", *Journal of Vibration and Acoustics*, Vol. 116, pp. 474-479, 1994.
- Hibbit, Karlsson & Sorensen, Inc., "ABAQUS/Users Manuals Version 6.2", Pawtucket, USA, 2001.

- Hong, H.-K. and Liu, C.-S., "Coulomb friction oscillator: Modelling and responses to harmonic loads and base excitations", *Journal of Sound and Vibration*, Vol. 229, No. 5, pp. 1171-1192, 2000.
- Hong, H.-K. and Liu, C.-S., "Non-sticking oscillation formulae for Coulomb friction under harmonic loading", *Journal of Sound and Vibration*, Vol. 244, No. 5, pp. 883-898, 2001.
- Hundal, M. S., "Response of a base excited system with Coulomb and viscous friction", *Journal of Sound and Vibration*, Vol. 64, No. 3, pp. 371-378, 1979.
- Hwang, D. Y. and Stallings, J. M., "Finite element analysis of bolted flange connections", *Computers and Structures*, Vol. 51, No. 5, pp. 521-533, 1994.
- Ibrahim, R. A., Zielke, S. A. and Popp, K., "Characterization of interfacial forces in metal to metal contact under harmonic excitation", *Journal of Sound and Vibration*, Vol. 220, No. 2, pp. 365-377, 1999.
- Inc., M., "Matlab User's Guide", Natick, Massachusetts, 1996.
- Iourtchenko, D. V. and Dimentberg, M. F., "In-service identification of non-linear damping from measured random vibration", *Journal of Sound and Vibration*, Vol. 255, No. 3, pp. 549-554, 2002.
- Ito, Y. and Masuko, M., "Study on the damping capacity of bolted joints-effects of the joint surface condition", *Bulletin of the JSME*, Vol. 18, No. 117, pp. 319-326, 1975.
- Ito, Y., Toyoda, J. and Nagata, S., "Interface pressure distribution in a bolt-flange assembly", *Journal of Mechanical Design*, Vol. 101, pp. 330-337, 1979.
- Iwan, W. D., "A distributed-element model for hysteresis and its steady-state dynamic response", *Journal of Applied Mechanics*, pp. 893-900, 1966.
- Johnson, K. L., "Contact Mechanics", *Cambridge University Press*, Cambridge, UK, 1992.

- Ju, S. H. and Rowlands, R. E., "A three-dimensional frictional contact element whose stiffness matrix is symmetric", *Journal of Applied Mechanics*, Vol. 66, pp. 460-467, 1999.
- Karamis, M. B. and Selcuk, B., "Analysis of the friction behaviour of bolted joints", *Wear*, pp. 73-83, 1993.
- Karius, D., "A contribution on forced periodic oscillations with piecewise linear and constant damping", *Journal of Vibration, Acoustics, Stress and Reliability in Design*, Vol. 107, pp. 383-391, 1985.
- Karnopp, D., "Computer simulation of stick-slip friction in mechanical dynamic systems", *Journal of Dynamic Systems, Measurement and Control*, Vol. 107, pp. 100-103, 1985.
- Laura, P. A. A., Pombo, J. L. and A., S. E., "A note on the vibrations of a clamped-free beam with a mass at the free end", *Journal of Sound and Vibration*, Vol. 37, No. 2, pp. 161-168, 1974.
- Lee, S.-Y., Ko, K.-H. and Lee, J. M., "Analysis of dynamic characteristics of structural joints using stiffness influence coefficients", *KSME International Journal*, Vol. 14, No. 12, pp. 1319-1327, 2000.
- Lee, Y. and Feng, Z. C., "Dynamic responses to sinusoidal excitations of beams with frictional joints", *Communications in Nonlinear Science and Numerical Simulation*, In Press, 2003.
- Leine, R. I., Van Campen, D. H., De Kraker, A. and Van Den Steen, L., "Stick-slip vibrations induced by alternate friction models", *Nonlinear Dynamics*, Vol. 16, pp. 41-54, 1998.
- Liang, J.-W. and Feeny, B. F., "Dynamical friction behaviour of a forced oscillator with a compliant contact", *Journal of Applied Mechanics*, Vol. 65, pp. 250-257, 1998.
- Lin, Y. and W., C., "A method of identifying interface characteristic for machine tools design", *Journal of Sound and Vibration*, Vol. 255, No. 3, pp. 481-487, 2002.

- Little, R. E., "Bolted Joints: How Much Give?" *Machine Design*, Vol. 39, pp. 173-175, 1967.
- Ma, X., Bergman, L. and Vakakis, A. F., "Identification of bolted joints through laser vibrometry", *Journal of Sound and Vibration*, Vol. 246, No. 3, pp. 441-460, 2001.
- Maia, N. M. M., Silva, J. M. M. and Ribeiro, A. M. R., "On a General Model for Damping", *Journal of Sound and Vibration*, Vol. 218, No. 5, pp. 749-767, 1998.
- Maruyama, K. and Yoshimoto, I., "On spring constant of connected parts in bolted joints", *Bulletin of the JSME*, Vol. 18, No. 126, pp. 1472-1480, 1975.
- Masuko, M. and Ito, Y., "Horizontal stiffness and microslip on a bolted joint", *Bulletin of the JSME*, Vol. 17, No. 113, pp. 1494-1501, 1974.
- Menq, C.-H., Bielak, J. and Griffin, J. H., "The influence of microslip on vibratory response, part 1: A new microslip model", *Journal of Sound and Vibration*, Vol. 107, No. 2, pp. 279-293, 1986.
- Menq, C.-H., Chidamparam, P. and Griffin, J. H., "Friction damping of two-dimensional motion and its application in vibration control", *Journal of Sound and Vibration*, Vol. 144, No. 3, pp. 427-447, 1991.
- Menq, C.-H. and Griffin, J. H., "A comparison of transient and steady state finite element analyses of the forced response of a frictionally damped beam", *Journal of Vibration, Acoustics, Stress and Reliability in Design*, Vol. 107, pp. 19-25, 1985.
- Menq, C.-H., Griffin, J. H. and Bielak, J., "The influence of microslip on vibratory response, part 2: A comparison with experimental results", *Journal of Sound and Vibration*, Vol. 107, No. 2, pp. 295-307, 1986.
- Metherell, A. F. and Diller, S. V., "Instantaneous energy dissipation rate in a lap joint-uniform clamping pressure", *Journal of Applied Mechanics*, pp. 123-128, 1968.
- Meyer, S., Weiland, M. and Link, M., "Modelling and updating of local nonlinearities using frequency response residuals", *International Conference on Structural System Identification*, Kassel, Germany, pp. 151-160, 2001.

- Moon, F. C. and Li, G. X., "Experimental study of chaotic vibrations in a pin-jointed space truss structure", *AIAA Journal*, Vol. 28, No. 5, pp. 915-921, 1990.
- Mottershead, J. E., Pascoe, S. K. and English, R. G., "A general finite element approach for contact stress analysis", *International Journal for Numerical Methods in Engineering*, Vol. 33, pp. 765-779, 1992.
- Mottram, J. T. and Shaw, C. T., "Using Finite Elements in Mechanical Design", *McGraw-Hill*, London, 1996.
- Oden, J. T. and Martins, J. A. C., "Models and computational methods for dynamic friction phenomena", *Computer Methods in Applied Mechanics and Engineering*, Vol. 52, pp. 527-634, 1985.
- Osgood, C. C., "Saving Weight in Bolted Joints", *Machine Design*, Vol. 51, No. 24, pp. 128-133, 1979.
- Palm III, W. J., "Introduction to Matlab 6 for Engineers", *McGraw-Hill*, New York, 2001.
- Pascoe, S. K. and Mottershead, J. E., "Linear elastic contact problems using curved elements and including dynamic friction", *International Journal for Numerical Methods in Engineering*, Vol. 26, pp. 1631-1643, 1988.
- Pierre, C., Ferri, A. A. and Dowell, E. H., "Multi-harmonic analysis of dry friction damped systems using an incremental harmonic balance method", *Journal of Applied Mechanics*, Vol. 52, pp. 958-964, 1985.
- Popp, K. and Stelzer, P., "Stick-slip vibrations and chaos", *Philosophical Transactions of the Royal Society of London A*, Vol. 332, pp. 89-105, 1990.
- Pratap, R., Mukerherjee, S. and Moon, F. C., "Dynamic behaviour of a bilinear hysteretic elasto-plastic oscillator, Part 1: Free oscillations", *Journal of Sound and Vibration*, Vol. 172, No. 3, pp. 321-337, 1994a.
- Pratap, R., Mukerherjee, S. and Moon, F. C., "Dynamic behaviour of a bilinear hysteretic elasto-plastic oscillator, Part 2: Oscillations under periodic impulse forcing", *Journal of Sound and Vibration*, Vol. 172, No. 3, pp. 339-358, 1994b.

- Ramberg, W. and Osgood, W. R., "Description of stress-strain curves by three parameters", *NACA Technical Note*, NS02, No. 902, 1943.
- Rao, S. S., "Mechanical Vibrations", *Addison-Wesley*, Reading, Massachusetts, 1995.
- Rashquinha, I. A. and Hess, D. P., "Modelling nonlinear dynamics of bolted assemblies", *Applied Mathematical Modelling*, Vol. 21, pp. 801-810, 1997.
- Ren, Y. and Beards, C. F., "On application of the harmonic balance method to structures containing one non-linear element", *Journal of Sound and Vibration*, Vol. 175, No. 5, pp. 705-710, 1994a.
- Ren, Y. and Beards, C. F., "A new receptance-based perturbative multi-harmonic balance method for the calculation of the steady state response of non-linear systems", *Journal of Sound and Vibration*, Vol. 172, No. 5, pp. 593-604, 1994b.
- Ren, Y. and Beards, C. F., "An experimental study on the dynamic response of a beam structure containing a pseudo joint", *Proceedings of the Institution of Mechanical Engineers*, Vol. 208, Part C, pp. 321-328, 1994c.
- Ren, Y. and Beards, C. F., "Identification of joint properties of a structure using FRF data", *Journal of Sound and Vibration*, Vol. 186, No. 4, pp. 567-587, 1995.
- Ren, Y. and Beards, C. F., "Identification of "effective" linear joints using coupling and joint identification techniques", *Journal of Vibration and Acoustics*, Vol. 120, pp. 331-338, 1998.
- Ren, Y., Lim, T. M. and Lim, M. K., "Identification of properties of nonlinear joints using dynamic test data", *Journal of Vibration and Acoustics*, Vol. 120, pp. 324-330, 1998.
- Shi, G. and Atluri, S. N., "Nonlinear dynamic response of frame-type structures with hysteretic damping at the joints", *AIAA Journal*, Vol. 30, No. 1, pp. 234-240, 1992.
- Shigley, J. E., "Mechanical Engineering Design", *McGraw-Hill*, Singapore, 1986.
- Shin, Y. S., Iverson, J. C. and Kim, K. S., "Experimental studies on damping characteristics of bolted joints for plates and shells", *Journal of Pressure Vessel Technology*, Vol. 113, pp. 402-408, 1991.

- Swayze, J. L. and Akay, A., "Effects of system dynamics on friction-induced oscillation", *Journal of Sound and Vibration*, Vol. 173, No. 5, pp. 599-609, 1994.
- Tabor, D., "Friction- the present state of our understanding", *Journal of Lubrication Technology*, Vol. 103, pp. 169-179, 1981.
- Tan, X. and Rogers, R. J., "Simulation of friction in multi-degree-of freedom vibration systems", *Journal of Dynamic Systems, Measurement and Control*, Vol. 120, pp. 144-146, 1998.
- Tang, D. M. and Dowell, E. H., "Damping in beams and plates due to slipping at the support boundaries, part 2: Numerical and experimental study", *Journal of Sound and Vibration*, Vol. 108, No. 3, pp. 509-522, 1986.
- Thomson, W. T. and Dahleh, M. D., "Theory of Vibration with Applications", *Prentice Hall*, Upper Saddle River, New Jersey, 1998.
- Tomlinson, G. R. and Hibbert, J. H., "Identification of the dynamic characteristics of a structure with Coulomb friction", *Journal of Sound and Vibration*, Vol. 64, No. 2, pp. 233-242, 1979.
- Tsai, J.-S. and Chou, Y.-F., "Identification of dynamic characteristics of a single bolt joint", *Journal of Sound and Vibration*, Vol. 125, No. 3, pp. 487-502, 1988.
- Ungar, E. E., "The status of engineering knowledge concerning the damping of built-up structures", *Journal of Sound and Vibration*, Vol. 26, No. 1, pp. 141-154, 1973.
- Vakakis, A. F. and Ewins, D. J., "Effects of weak nonlinearities on modal analysis", *10th International Modal Analysis Conference*, San Diego, pp. 72-78, 1992.
- Valanis, K. C., "Fundamental consequences of a new intrinsic time measure. Plasticity as a limit of the endochronic theory", *Archives of Mechanics*, Vol. 32, pp. 171-191, 1980.
- Wang, J. and Sas, P., "A method for identifying parameters of mechanical joints", *Journal of Applied Mechanics*, Vol. 57, pp. 337-342, 1990.

- Wen, Y.-K., "Method for random vibration of hysteretic systems", *Journal of Engineering Mechanics Division, Proceedings ASCE*, Vol. 102, EM2, pp. 249-263, 1976.
- Wen, Y.-K., "Equivalent linearization for hysteretic systems under random excitation", *Journal of Applied Mechanics*, Vol. 47, pp. 150-154, 1980.
- Wen, Y.-K., "Methods of random vibration for inelastic structures", *Applied Mechanics Reviews*, Vol. 42, No. 2, pp. 39-52, 1989.
- Whiteman, W. W. and Ferri, A. A., "Multi-mode analysis of beam-like structures subjected to displacement-dependent dry friction damping", *Journal of Sound and Vibration*, Vol. 207, No. 3, pp. 403-418, 1997.
- Woo, K. L. and Thomas, T. R., "Contact of rough surfaces: A review of experimental work", *Wear*, Vol. 58, pp. 331-340, 1980.
- Zadoks, R. I. and Kokatam, D. P. R., "Investigation of the axial stiffness of a bolt using three-dimensional finite element model", *Journal of Sound and Vibration*, Vol. 246, No. 2, pp. 349-373, 2001.