CHAPTER 3

Design of a variable stiffness spring

3.1 Introduction

In this chapter the design of the variable stiffness spring will be addressed. The design starts with the concept phase where some concepts will be addressed individually and followed by the detailed design of the selected concept.

The objective of this chapter is to design, build and test a variable stiffness spring with low damping and to use smart materials to obtain the stiffness change. The reason for the low damping is due to the fact that it is foreseen to change the damping of the isolator. To be able to give the isolator a reasonable damping change, it is necessary that the primary damping of the spring is not too high.

A usable variable stiffness spring would typically require a change in stiffness of at least 40% (the maximum stiffness must be 1.4 times the minimum stiffness), but a larger stiffness change is preferable. The structural damping constant of the spring must be lower than 0.15. A last factor that must be considered is the response time of the stiffness change. If it takes too long for the spring to change its stiffness, the spring would be impractical. Furthermore it is foreseen to implement an optimisation algorithm at the end to control the spring. A response time of minutes or hours would cause the optimisation to take too long to be of any practical value. Therefore a response time below 30 seconds was chosen as goal, but a faster response time - one less than 10 seconds would be preferable.

3.2 Concepts

Three concepts were investigated. In each case the concept was tested with an experiment to validate the theory. Each concept will be addressed separately.

3.2.1 Magnetorheological elastomer

Magnetorheological materials are materials that react to a magnetic field. These materials are relatively new and were patented in 1992 by the Lord Corporation (Carlson & Weiss, 1992) who are still about the only company specialising in these materials. The common form of this class of materials is magnetorheological fluid, which is commonly used in adaptive dampers (Dyke et al., 1996). The viscosity of a magnetorheological fluid changes in a magnetic field making this fluid ideal for adaptive dampers.

Magnetorheological elastomers are elastomeric materials which stiffness changes under the influence of a magnetic field. The Lord Corporation did the first development of a theoretical model of this material in 1996 (Jolly et al., 1996a; Jolly et al., 1996b) and was followed by the Ford Motor Company in 1999 (Ginder et al., 1999; Ginder et al., 2000) who developed a variable stiffness suspension bushing for their cars.

In its basic form, a magnetorheological elastomer is a normal elastomer like natural rubber or polyurethane that has a certain amount of magnetic iron particles in it. It was shown that a volume fraction of 0.27 yields the best results. These particles are mixed into the elastomer prior to curing. The elastomer is then cured in the presence of a magnetic field of 0.5 Tesla. This causes the iron particles to line up in chains inside the elastomer. During curing these particles is then locked into the elastomer matrix. When the cured elastomer is deformed, these particles move out of the chains. If this happens while the elastomer is in the presence of a magnetic field, the magnetic field restricts the particles from moving out of the chains and causes the modulus of the elastomer to increase. Jolly et al. achieved an increase of up to 39%. The damping of the elastomer also decreased as the modulus increased.

Complex theoretical models were developed that used the magnetic force between two adjacent particles to calculate modulus change. FEM models were also used to predict the modulus changes in the material.

Magnetorheological elastomers seemed to be a viable concept out of the literature study. The next step was to do an experimental test to investigate whether the effect could be reproduced. Polyurethane was used as the elastomer for the tests. Rubber has less damping than polyurethane and made it more suitable for the tests, but polyurethane is much easier to produce than natural rubber. If the tests with polyurethane were successful, the final spring could be made from rubber for lower damping.

The initial design of the spring was taken from the design of the suspension bushing developed by Ford Motor Company depicted in Figure 3. 1.

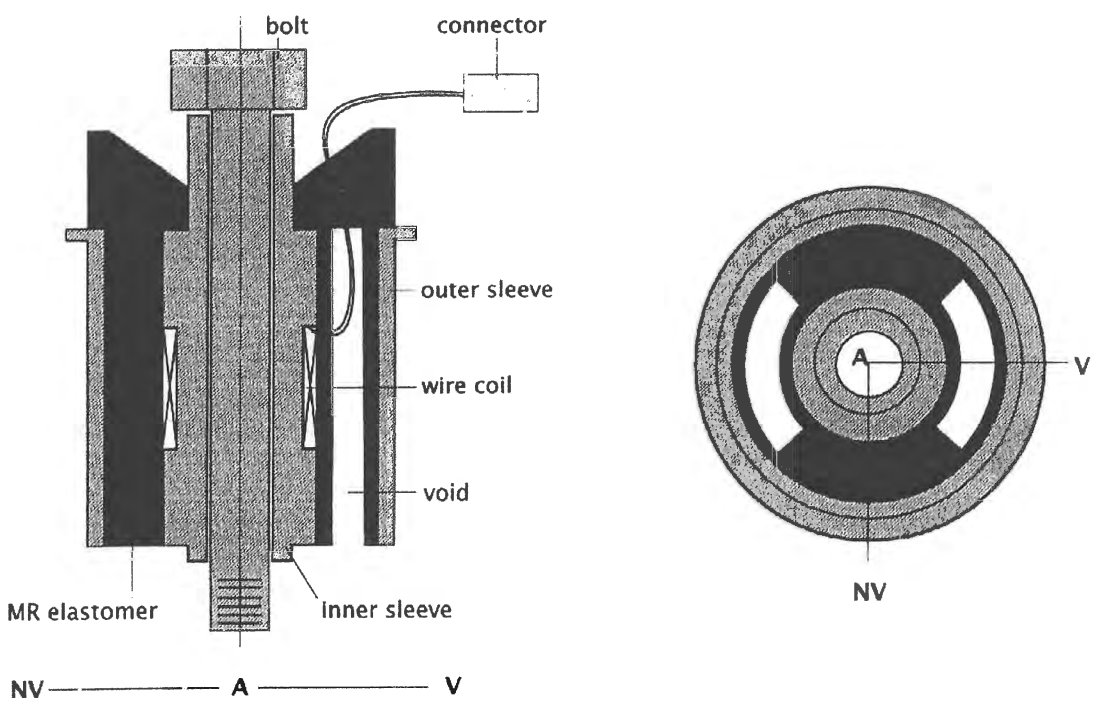

Figure 3. 1 Variable stiffness suspension bushing using MR elastomer (Ginder et al., 2000)

The bushing had an outside and inside metal sleeve to conduct the magnetic field and an electric coil around the inside sleeve to generate the magnetic field. A similar spring was designed, but on a larger scale due to the fact that a large spring was needed for the isolator and not a tiny spring like the bushing. The design is given in Figure 3. 2.

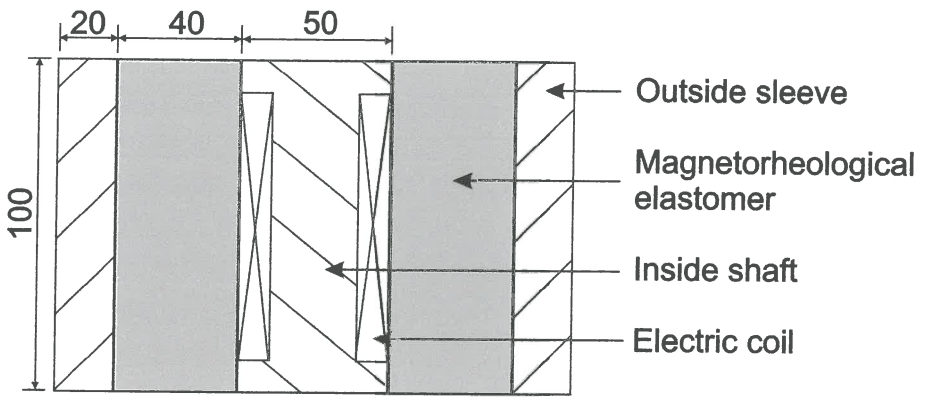

Figure 3. 2 **Design of MR elastomer spring**

The spring was manufactured and tested in a servo-hydraulic actuator by fixing the outside sleeve and moving the inside shaft axially at a constant frequency and amplitude. The experimental spring is depicted in Figure 3. 3.

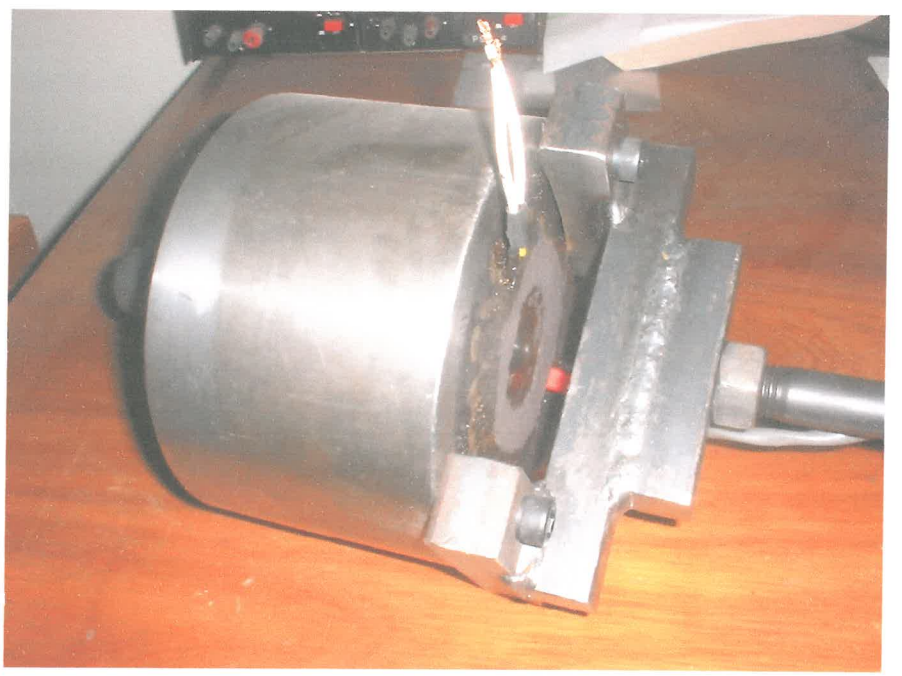

Figure 3. 3 **Experimental MR elastomer spring**

The results in Figure 3. 4 were obtained for an excitation frequency of 16 Hz and different currents through the coil.

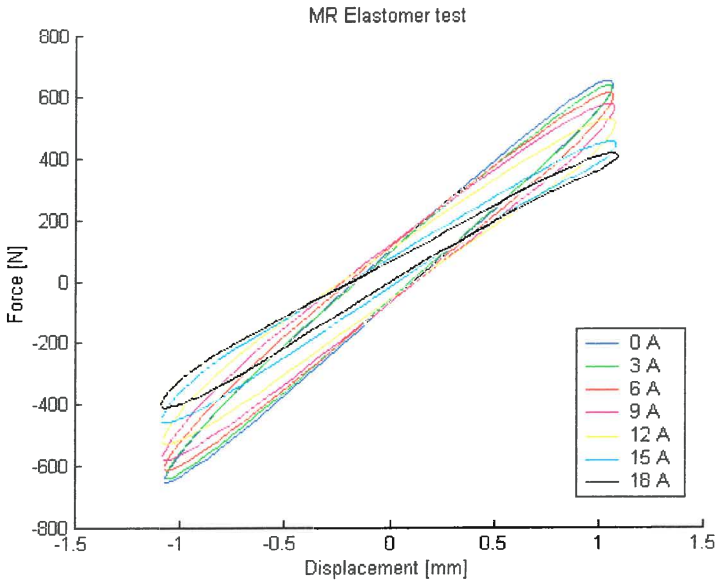

Figure 3. 4 **Hysteresis loops from MR elastomer spring for different currents**

The method used to determine the stiffness and damping from the experimental results can be seen in Appendix A. The corresponding stiffness and damping values are plotted in Figure 3. 5.

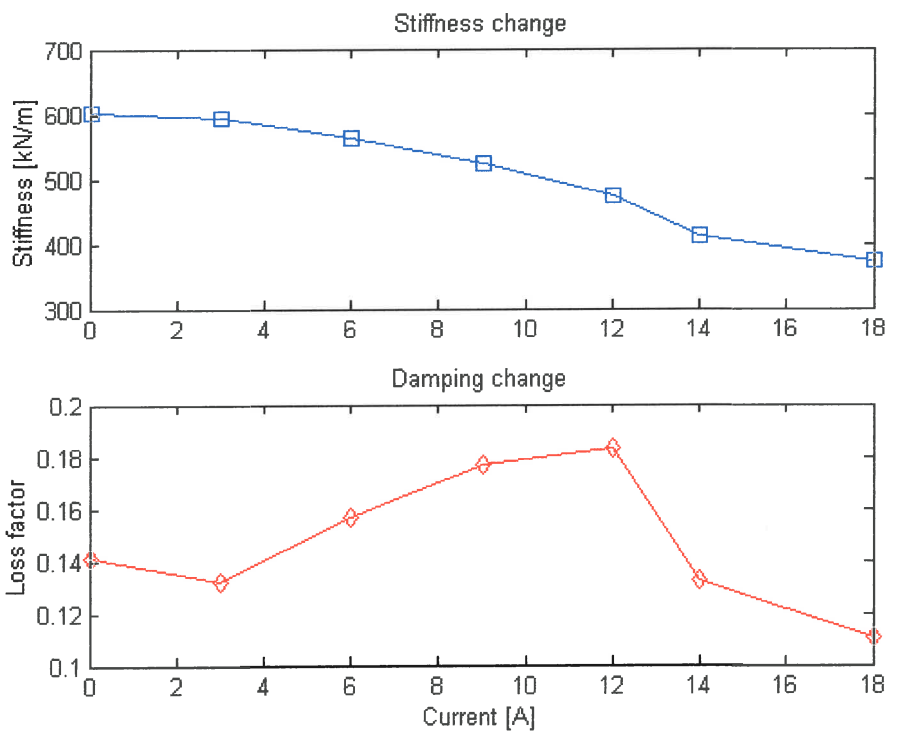

Figure 3. 5 **Stiffness and damping as function of current**

The resulting stiffness change was 60% and the damping was almost in the required range. One problem was that the stiffness of the elastomer decreased as the current or the magnetic field was increased. This is exactly the opposite of what was predicted from the literature study and theory. Another problem was that the polyurethane and the inside shaft did not bond properly and had some amount of slip. Although satisfied with the amount of stiffness change, the discrepancy with the theory resulted in a second test that included a few modifications.

The first modification was that grooves were made in the inside shaft so that the polyurethane will attach better to the shaft. Secondly a magnetic sensor was installed in the polyurethane to be able to measure the magnetic field inside the rubber. The last modification was that a significantly larger coil was inserted around the shaft to decrease the current that was needed for the tests.

The same test was done on the second spring, but the test was repeated directly after each other to test for consistency. The results are given in Figure 3. 6.

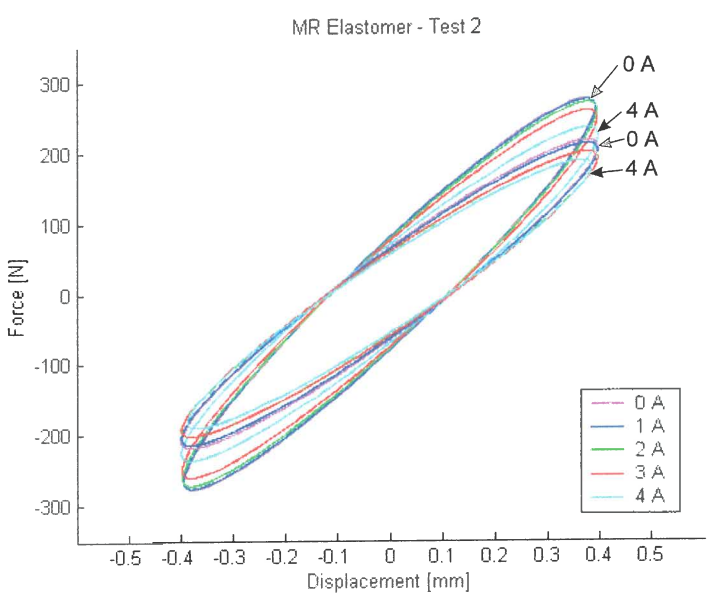

Figure 3. 6 Hysteresis loops for second experimental test

The stiffness again decreased as the current was increased as in the first test. An interesting observation was that when the test was repeated the second time, the

stiffness was lower than the first time. Even the stiffness of the 0A test was lower than the 4A test of the first run. It was concluded that the stiffness change was not due to the magnetorheological effect, but due to an increase in temperature of the polyurethane. Due to the large currents that were sent through the coil that was situated inside the polyurethane, the coil heated up and heated the polyurethane. That is the reason why the stiffness in the second test was lower than in the first test.

Further tests did not result in any indication of the magnetorheological stiffening of the spring. One main reason for this was the magnitude of the magnetic field. The magnetic sensor installed in the polyurethane indicated that the magnetic field in the polyurethane was not even 1% of the desired value. To be able to create a magnetic field of 0.5 – 1 Tesla in an elastomer, even in a very small specimen, without heating the elastomer will require an external magnetic source (electromagnet) or a very small and well-designed specimen. This concept was just not viable on the scale that was needed for this project.

3.2.2 Heating of elastomer

Although the results from the first concept did not show the magnetorheological effect, it showed that a reasonable stiffness change could be obtained by heating the polyurethane. Therefore a logical second concept was to investigate the heating of an elastomer as a variable stiffness spring.

The biggest problem foreseen with this concept was response time. To heat and especially cool a solid piece of polyurethane is not an easy task. Although forced convection is the obvious way to approach such a problem, it was not preferable in this case due to the complexity that comes with it. Natural convection was the first choice in this case and the design tried to increase the response time without using forced convection.

Basically the same design of the spring was used as in the magnetorheological elastomer. The changes made were an inside sleeve instead of the solid shaft and

much thinner steel parts to conduct heat faster. A soft polyurethane was used as the elastomer for this spring. A heating coil was moulded inside the polyurethane together with a thermocouple to be able to measure the temperature of the polyurethane. The design is depicted in Figure 3. 7.

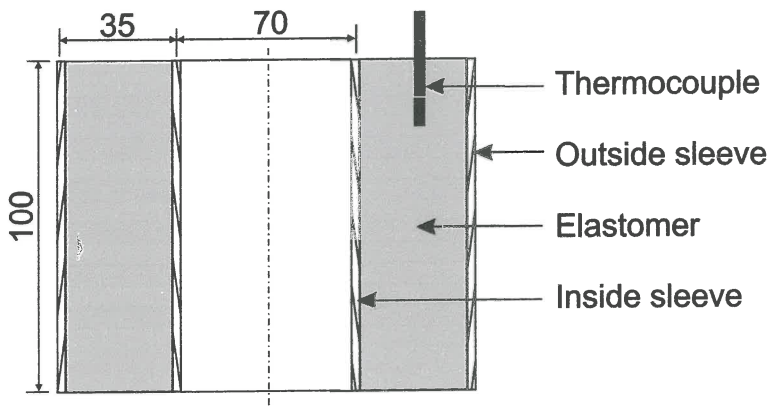

Figure 3. 7 **Design of elastomeric spring**

The temperature was controlled with a 0 – 200 °C temperature controller using a Pt 100 thermocouple. The response of the spring was measured at 5 °C intervals from 30 °C to 80 °C. The response was measured at different frequencies to be able to see the frequency effect to stiffness and damping. Figure 3. 8 show the results.

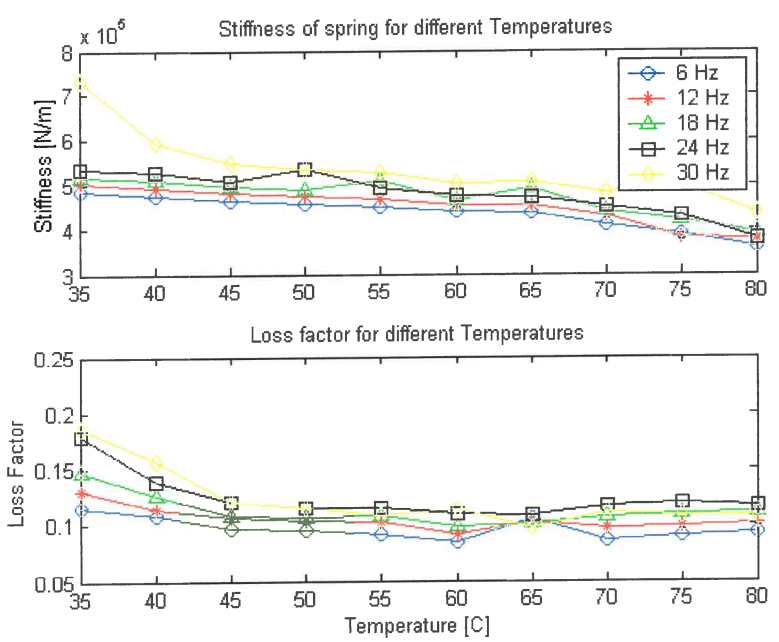

Figure 3. 8 **Stiffness and damping as a function of temperature**

The decrease in stiffness and damping with an increase in temperature can clearly be seen. The discrepancies in the linearity of the graphs were due to a faulty servo-hydraulic controller and should not be taken into account. Overall the graphs indicate a quite linear decrease in stiffness with an increase in temperature. Furthermore it can be seen that the higher frequencies resulted in higher stiffness and damping.

Although the damping was in an acceptable range, the stiffness change of 34 % was not quite sufficient for the application. Furthermore, the response time was definitely too long to even be considered as an option. This concept was discarded due to these reasons.

3.2.3 Compound leaf spring

The third concept originated out of a publication (Walsh & Lamancusa, 1992) that claimed to have achieved a stiffness change of 6200% (62 times) of which 4500% (45 times) was linear with respect to the separation of the leaf springs. Walsh & Lamancusa used very thin leaf springs to get a very low minimum stiffness that is probably the reason for the impressive figures. This result has been regarded by other authors (Brennan, 2000) as impractical and not possible to implement on a practical problem. It is possibly the case, but it shows that the concept has quite a lot of potential. Even if only 10% of the reported stiffness change can be achieved, it will still be a very significant stiffness change in the application of a LIVE isolator. Another very attractive property of this concept is the fact that steel springs are used instead of elastomeric springs that were the case in the previous two concepts. Steel springs are known to have a very low loss factor and therefore it would be possible to decrease the damping of the spring significantly.

The first test was a basic experiment with two straight leaf springs separated with a lead screw. The experiment is depicted in Figure 3. 9.

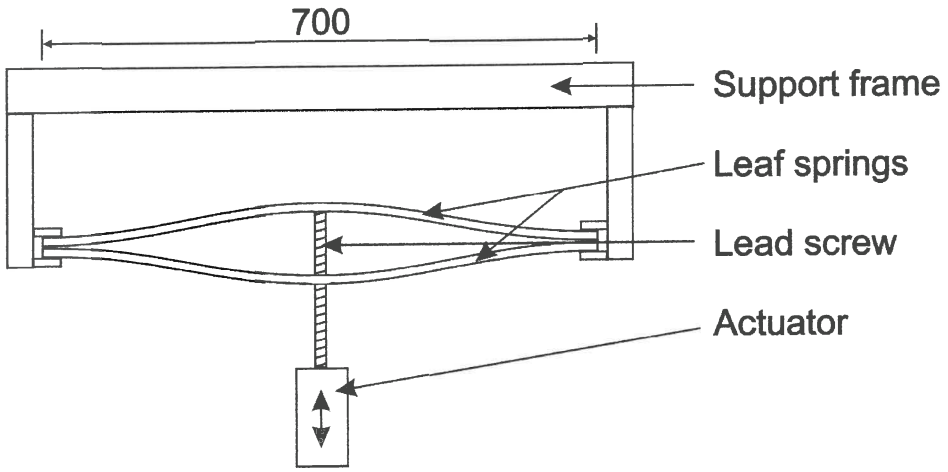

Figure 3. 9 Compound leaf spring design

This was a quite big spring and not suitable for implementation on an isolator. This test was therefore only a preliminary test to validate the concept and to see what stiffness change is practically possible for this concept.

The results are given in Figure 3. 10.

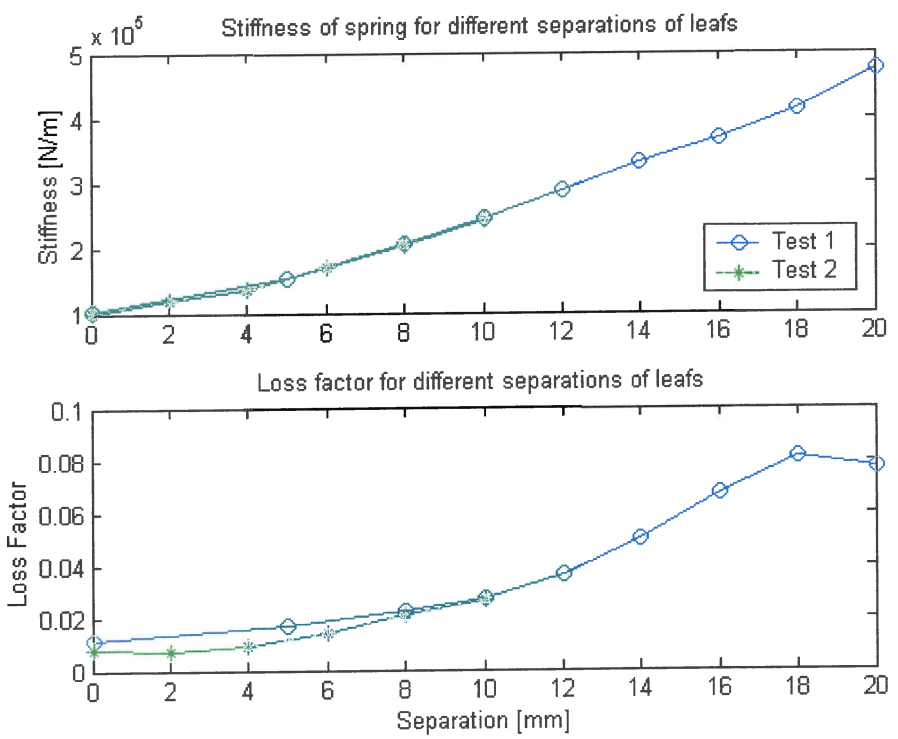

Figure 3. 10 Stiffness and damping as a function of separation

Two prominent features of the results are the reasonable stiffness change (4.5 times) and the very low damping of the spring. Both these two observations were expected beforehand and therefore the test was successful in proving the concept and giving some estimated figures regarding stiffness change and loss factor. From the graphs it can be seen that two test runs were performed and that the result were basically the same. This proves the repeatability of the tests.

Finite element modelling provides a method of analysing structures without setting up an experimental test. Therefore it would be ideal to have an accurate finite element model of the springs to be able to calculate the stiffness for different designs and configurations. A FEM model was built of one of the springs that were used in the test to compare with the experimental results.

The finite element package from MSC software was used and included MSC Patran and MSC Marc, a pre-processor and a non-linear solver. The leaf spring was meshed with brick elements and deformed in the way the spring deformed during the test by clamping the ends of the spring and displacing the centre. The reaction forces were extracted and the stiffness calculated at different displacements.

Two simulations were done. The first modelled the ends to be clamped in all directions except in the x direction (horizontally in-line with the spring). The second simulation modelled the ends clamped in all directions. The reason for this is that the boundary condition will have a significant influence on the stiffness and that in the practical test, the ends were not completely fixed and not completely free, so it was suspected to be somewhere in between. The results are given in Figure 3. 11.

The method of obtaining separation force from the force on a single spring is described in Appendix B.

As expected the experimental results were between the results from the two simulations. One interesting observation is the amount of difference in stiffness change between the clamped and free simulations. This is important to understand for future design of such springs and need to be explained.

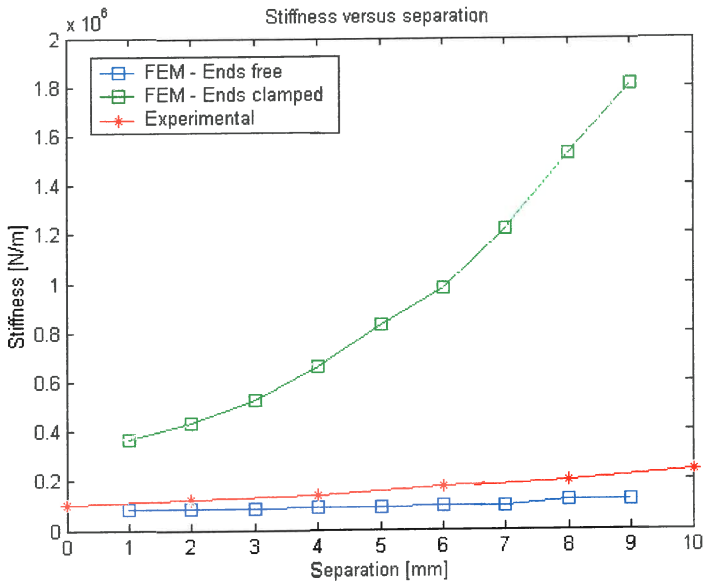

Figure 3. 11 Comparison between stiffness calculated by FEM and measured

The reason why the stiffness of the compound spring changes is because the separate springs are non-linear. Linear springs will not have an increase in stiffness as the displacement increases. Therefore the more non-linear the springs are, the more the stiffness change will be. The following two mechanisms are responsible for the non-linearity of the springs:

- Change in geometry
- Elongation of springs

The change in geometry is the most obvious. As the two leafs are separated, the change in geometry cause the spring to stiffen. It can be compared to an increase in the second moment of area of a beam. If the height of the beam is increased, the second moment of area and therefore the stiffness of the beam increases. The boundary condition will also not have a significant influence on this mechanism. The second mechanism can be explained with Figure 3. 12.

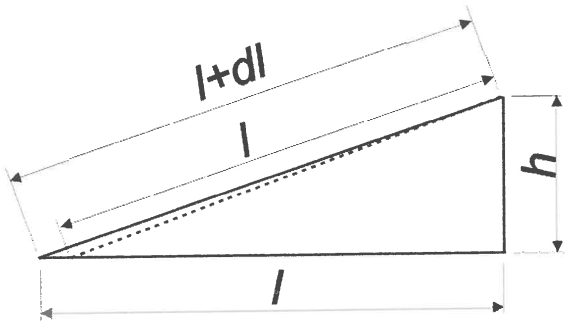

Figure 3. 12 Explanation of elongation of spring

If the centre of the spring is displaced upwards by *h*, one of the following two situations can occur:

- If the ends are free to move horizontally, the spring will deform as indicated by the dotted line. The centre of the spring will move upwards and the ends will move towards the centre. This cause the spring itself to keep a constant length and only bending of the spring is needed and no elongation.
- If the ends are clamped, the springs will have to stretch to be able to be displaced at the centre as illustrated by the solid line in the sketch. The tension generated in the springs is the same as tightening a guitar string. This causes the spring to be a lot stiffer and more non-linear. This mechanism is the main reason for the difference in stiffness of the two simulations and is an important aspect of these types of springs. To be able to get a significant stiffness change, the boundary conditions of the experimental spring must be as close to clamped as possible.

As said previously, the springs are non-linear with respect to the deflection of the spring. If two springs are combined in a spring assembly to form a variable stiffness spring, this spring can however be assumed to be linear around its neutral position if the spring is displaced by a small amount. That is the reason why a linear spring model can be used for such a spring although it consists of non-linear springs.

One last aspect of the FEM simulation that was investigated was the effect of the number of elements used in the model. In any FEM model, the question of "How many elements are enough?" will probably be asked at some stage. With modern

computers and software, it is really not more difficult or more time consuming to mesh a model with 10 or 10 million elements. The problem comes in at the solving of the model. The more elements the model has, the larger the matrices become and the longer the calculations take. It is not uncommon at all these days to have models that can take up to a few days to solve.

The problem is that only up to a certain point will more elements give you a more accurate answer. The accuracy of the results converges quite quickly to the actual answer with an increase in elements after which a further increase in elements will only increase the processing time dramatically and will not yield a significantly better answer.

Therefore the simulation of the leaf spring was run for different number of elements and the results combined in Figure 3. 13.

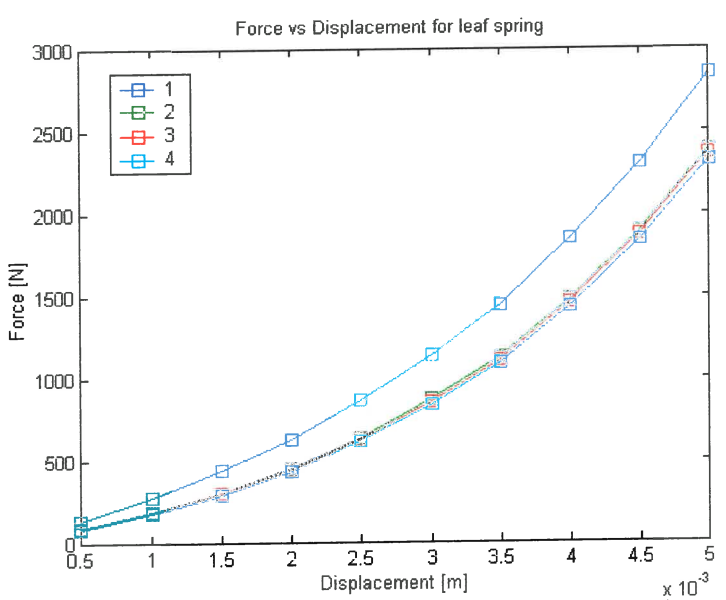

Figure 3. 13 Comparison between different number of elements in FEM model

The first test entailed the least number of elements and the fourth test the most. It can be seen that in the first test too few elements were used and that the results are not accurate (stiffness was over predicted). From the second test onwards, the results are

more accurate every time, but by such small a margin that it is not worth the extra processing time for that small amount of extra accuracy (test 2 took about 2 minutes of processing time and test 4 about 16 hours).

Out of this preliminary work the following conclusions could be made:

- The concept works and a significant stiffness change can be achieved with relatively low damping.
- Future designs need a way to practically clamp the ends of the springs for maximum stiffness change and to be able to compare it to the FEM model.
- The spring needs to be smaller and more compact to be used in an isolator.

This concept was chosen out of the three that were investigated and with these considerations the design process could progress into the final design of the spring.

3.3 Design of spring

3.3.1 Design concepts

The spring that needed to be designed must have the following properties as seen from the previous section:

- It must be smaller and more compact to be to implemented practically
- The ends of the spring must be as close to clamped as possible
- The design must produce repeatable test results.

The current design of two long straight leaf springs was not very well suited to be implemented in an isolator and has major problems as far as the clamping of the ends is concerned due to the fact that a very thick and stiff structure would be needed to achieve clamping. Therefore an entirely new design was needed.

LIVE isolators are usually circular in section and therefore a circular spring would probably be a suitable design to go with the isolator. For a circular spring, an inside

and outside ring will be needed that would move axially with respect to each other with elements between them that will deform and supply the stiffness. Figure 3. 14 show the simplest form of the design.

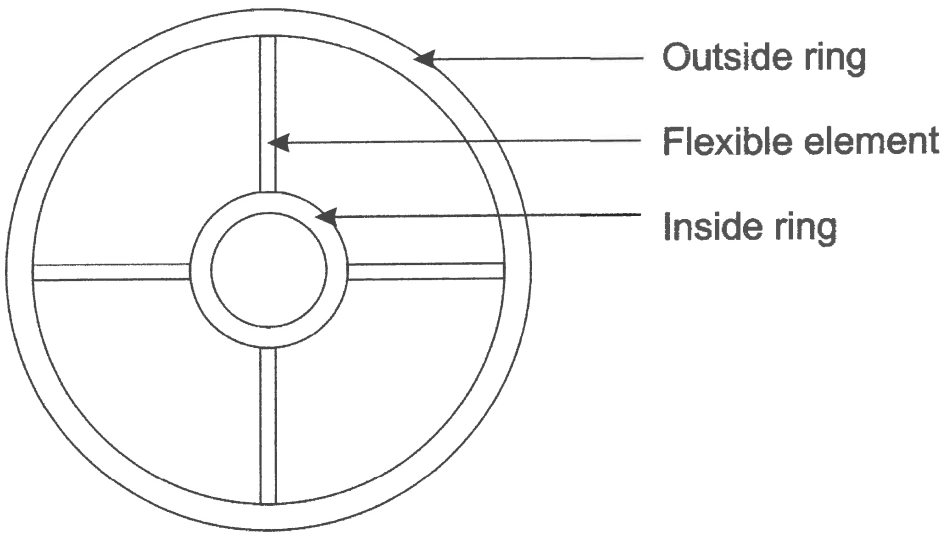

Figure 3. 14 Basic design of circular spring

Apart from the fact that the circular design will be nicely incorporated with the isolator, it has a few other advantages:

- It has a high lateral stability in all directions. This is important for an isolator.
- The outside ring already gives a clamped effect for the flexible elements due to the fact that a circle is a very strong shape. Furthermore, it will be possible to clamp the ring itself to produce a solid base that will result in a nearly perfectly clamped spring.

This design basically satisfies all the requirements. Within this design, the flexible elements can still be changed to suit the stiffness requirements needed. The simplest way to insert the flexible elements is to make the entire spring out of one piece from a thin spring steel sheet. This may sound impossible to manufacture, but with modern laser cutting techniques, it is actually much easier to manufacture in this way than by any other method.

To determine a suitable design, four designs were made and evaluated against each other with the FEM package. The four designs are depicted in Figure 3. 15.

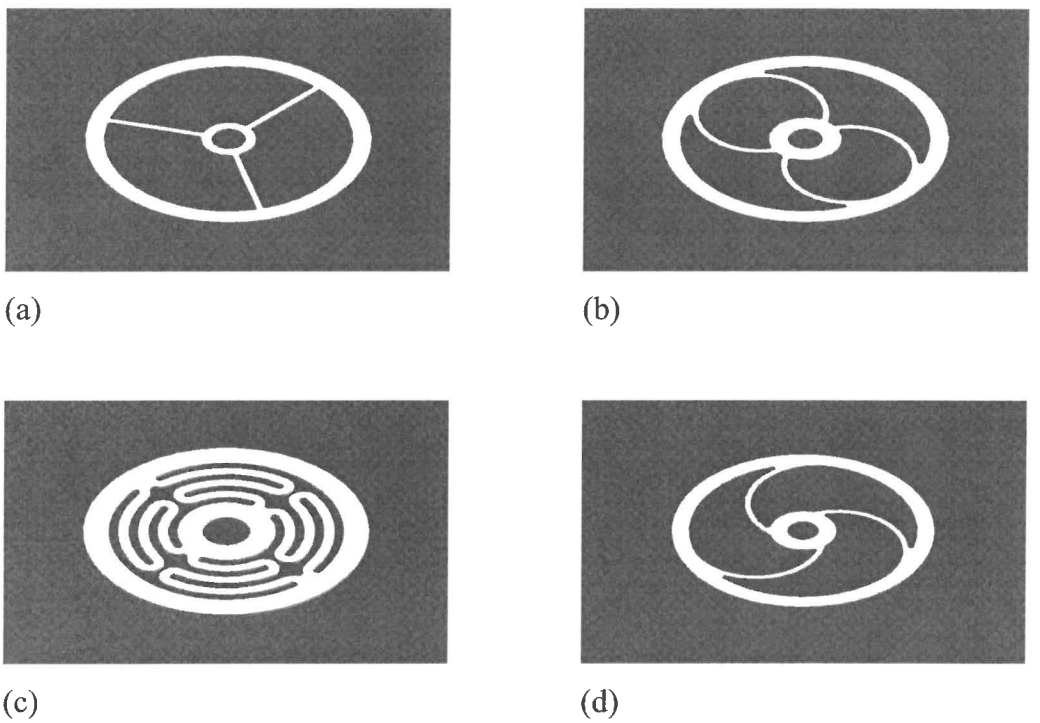

Figure 3. 15 Four designs of circular springs

The first design is the normal design with straight spokes. In the other three designs it was tried to obtain a less stiff spring but without increasing the size of the spring. All of these designs were simulated with MSC Patran and MSC Marc and the result compared to each other as shown in Figure 3. 16.

Out of the result it can be seen that the first design with the straight spokes yielded the best increase in stiffness. This is again due to the fact that the spokes have to stretch for the inside to move upwards, where the spokes in the other designs are able to bend or twist and do not have to stretch. It is noticeable that the third design did not yield the lowest stiffness. That however is due to the thicker spokes of the design compared to the last design that yielded the lowest stiffness. The first design was chosen due to the largest stiffness change. Although a less stiff spring is preferable in an isolator, the percentage stiffness change is the primary criteria for the application.

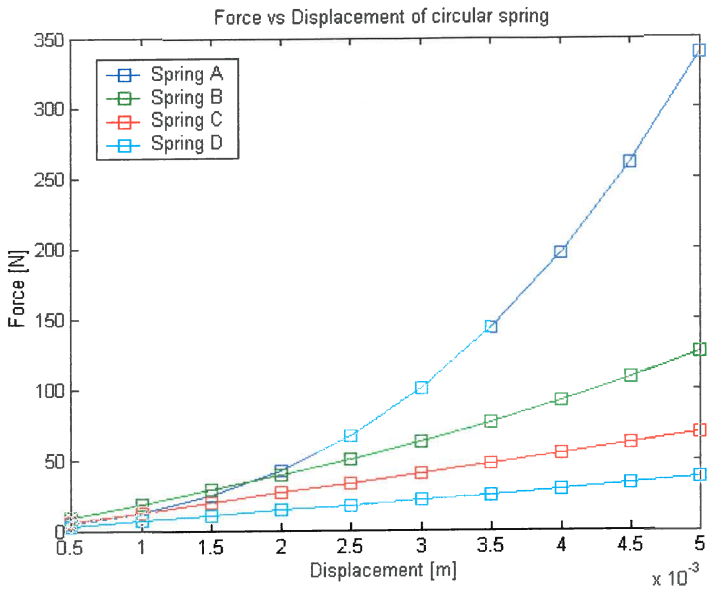

Figure 3. 16 **Stiffness of the four circular springs**

3.3.2 Final design

Three spokes were chosen in the first design, because it gave to lowest stiffness while still retaining stability and lateral stiffness in all directions that would not be the case with 2 spokes. To obtain the final design, many simulations were done with the FEM model to try and minimise the stresses in the spring while obtaining favourable stiffness values. The dimensions were kept as small as possible to keep the design compact. The design is depicted in Figure 3. 17.

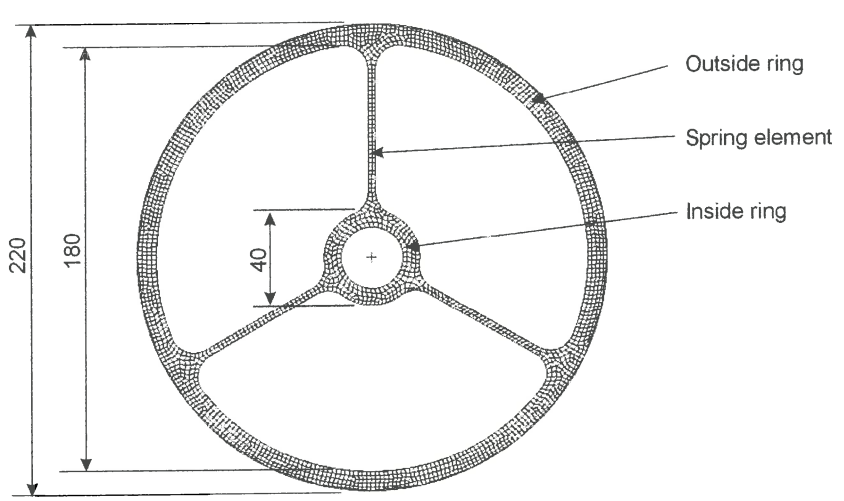

Figure 3. 17 **Mesh of circular spring**

In the FEM model the spring was meshed with about 4000 brick elements. The outside was clamped and the inside displaced 5mm upwards. 10 steps were used in the output so that data for every 0.5mm displacement were obtained. The displaced structure at 5mm displacement with corresponding stresses is given in Figure 3. 18.

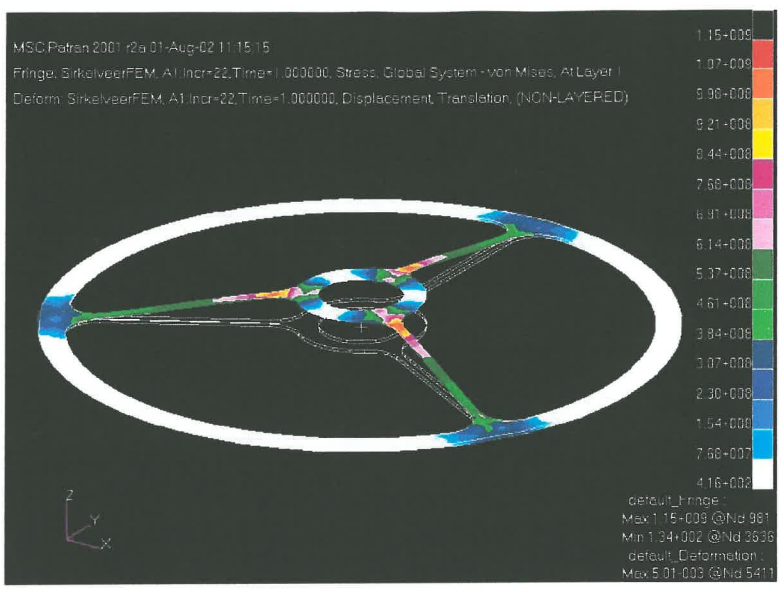

Figure 3. 18 Stresses in circular spring at 5mm displacement

The nodal forces and displacements were extracted from the results for all 10 steps. The vertical nodal reaction forces for the inside ring were added together to obtain a single force needed to displace the spring. From this the force vs. displacement graph in Figure 3. 19 was plotted.

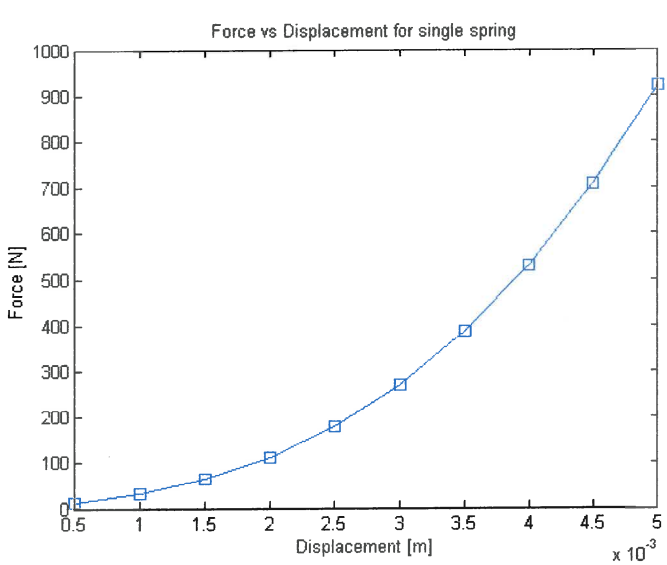

Figure 3. 19 FEM analysis results

3.3.3 Validation of FEM model

It is good practise to validate FEM results with an experimental test. For this purpose the designed spring was manufactured from hardened spring steel (EN42D). Laser cutting was used to manufacture the complex geometry of the spring. A frame was manufactured to clamp the spring on the outside and to mount the spring in order to do the test. A load cell was used to measure the applied force and a strain gauge displacement transducer was used to measure the displacement of the centre of the spring. The setup is depicted Figure 3. 20.

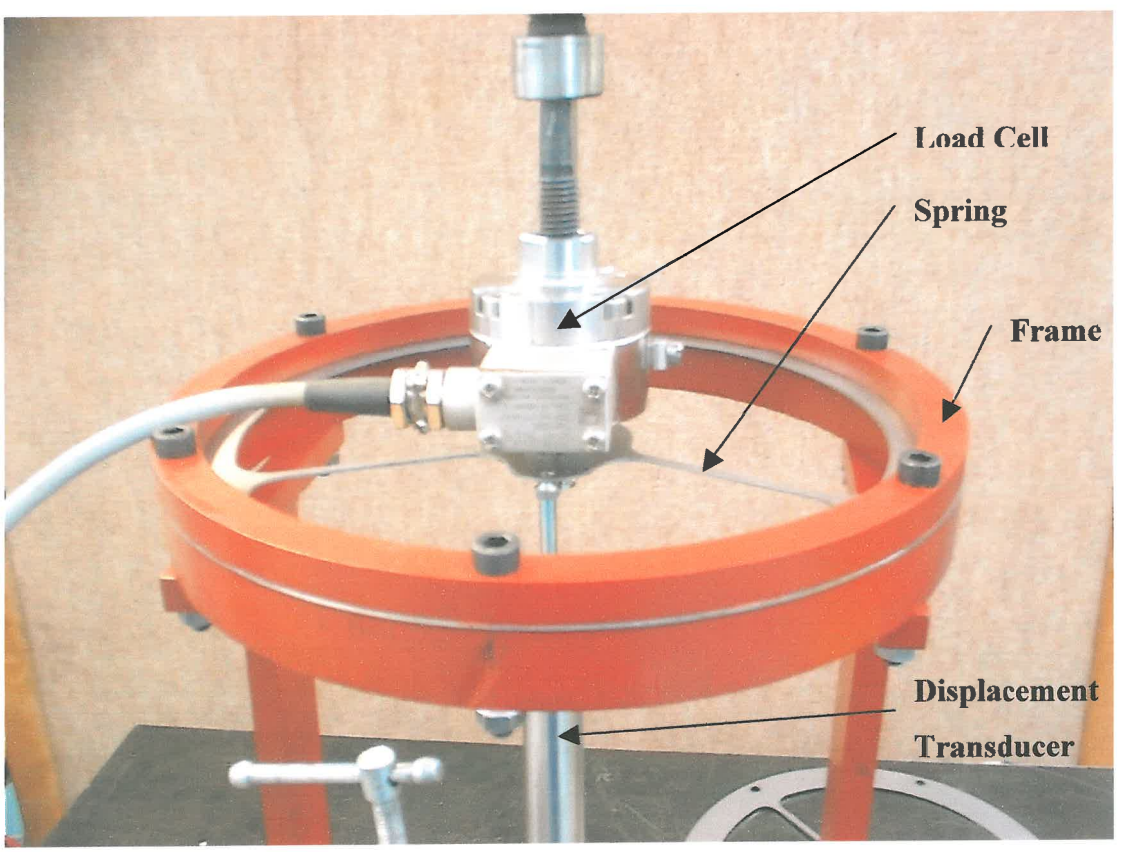

Figure 3. 20 Experimental setup

The spring was displaced 5mm and the displacement and force recorded every 0.5mm. After the first test it was seen that the spring slipped within the frame with

displacements larger than 4mm. That caused the change in gradient of the graph. The test was repeated with the spring tightened better, but the slipping occurred again.

The results from the FEM analysis and the practical test were compared to each other by plotting them on the same graph in Figure 3. 21.

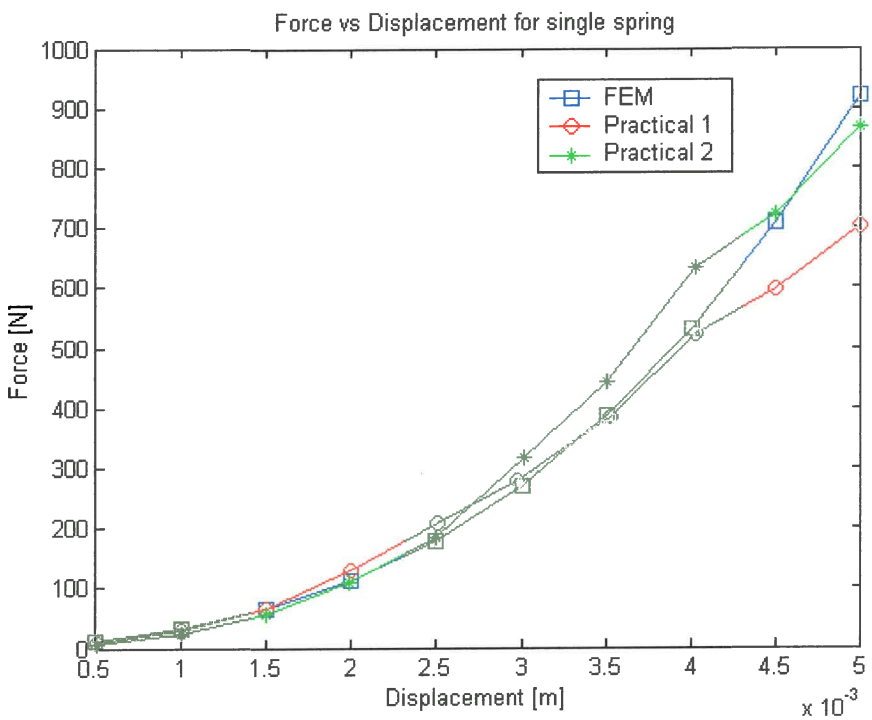

Figure 3. 21 Comparison between FEM and experimental results

It can be seen that there is a very good correlation between FEM and the experimental results with displacements up to 4mm. The slipping of the spring caused the difference above 4mm displacement and could be eliminated by better fixing the spring to the frame.

The spring has been designed, manufactured and compared to the FEM model. The last part of the spring assembly that has still not been addressed is the separation mechanism that will separate the two springs and control the stiffness of the spring. This will be addressed in the following section.

3.3.4 Separation mechanism

As stated in chapter one, it is preferable to use smart materials to accomplish the variation in stiffness of the variable stiffness spring. The current design consists of steel springs and do not comply with this aim. Therefore the separation mechanism will be the area where smart materials will be implemented. The separation mechanism must have the following properties:

- Must have a displacement of up to 10 mm.
- Must be able to deliver a force of up to 900 N at 10mm displacement (Figure 3. 21).
- Must be a simple mechanism with as few moving parts as possible. A solid-state type of actuator is preferable.
- The response time must be below 30 seconds and even less if possible.

These properties are guidelines for the optimal performance of the spring. It is therefore possible to use a mechanism that does not comply with all of the properties, but it will result in a spring with probably less stiffness change than possible.

3.3.4.1 Actuators

To minimise the moving parts in the separation mechanism, the actuator must preferably act directly upon the springs without an amplification device. The two major problems regarding smart actuators and the specifications are the displacement and force. The two most common smart actuators namely piezoelectric actuators and shape memory alloy actuators will not be able to produce displacements and forces of the required magnitude. Piezoelectric actuators have minute displacements (a few microns), but do have quite a large force. The problem is that the amount of displacement amplification needed will decrease the force to probably far below 1 N. Shape memory alloys on the other hand can in some cases give 10 mm displacement by means of bending, but result in very small forces. The amount of actuators needed to produce 900 N will just not be practical.

Other smart actuators that may possibly work are bi-metal strips and wax actuators. Bi-metal strips will probably have the same problem as shape memory alloys due the bending mechanism it uses. Wax actuators are a range of very compact actuators with quite a large force and displacement and were therefore well suited for the application.

3.3.4.2 Wax actuators

Wax actuators are not commonly sold commercially as actuators and the only company that could be found that do specialise in these types of actuators is Starsys. They manufacture these actuators (they call them paraffin actuators) for space applications and have a whole range offering different stroke length, force and accuracy.

Maximum values are 1.25 inches (32 mm) stroke, 500 lbf (2224 N) force and 0.3 micron accuracy. All of these figures can obviously not be obtained simultaneously, but gives an indication of the capabilities of these actuators. The 0.3 micron accuracy is obtained with a closed loop displacement control system. It is possible for them to design an actuator for a specific application regarding temperature range, force, stroke, internal spring and heating element. A quotation from them revealed that such a custom made actuator would be very expensive. One problem with the actuators is the response time, which is given in the region of 200 seconds. This will be a problem, but it will probably be possible to improve this with better heating and cooling.

A much more common and less expensive form of wax actuators is found in the thermostats of cars. The wax actuator in the thermostat is heated by the car's cooling water through forced convection. When the temperature of the water rises above a certain temperature, the wax actuator extends to open a valve that lets the water flow through the radiator. These thermostats can be obtained from any motor spares outlet for about R40. In the thermostat assembly there is usually a wax actuator, a spring to enable the actuator to retract and a housing to fit the thermostat to the vehicle.

Figure 3. 22 Thermostat and wax actuator

Taking the wax actuator out of the housing revealed a very small and compact actuator depicted in Figure 3. 22. The actuator consists of a copper cup filled with wax, a steel shaft and a rubber seal to prevent the wax inside the actuator from leaking (Figure 3. 23). The actuator works with a basic hydraulic principle in the sense that when the wax is heated, it changes from a solid to a liquid resulting in a volume increase or a pressure increase inside the copper cup. If the shaft is free to move, the shaft will extend. If the shaft is clamped, the pressure will rise inside the actuator and will increase the force on the shaft. Therefore it is important to note that the displacement of the actuator is dependant not only on the temperature, but also on the force produced on the shaft.

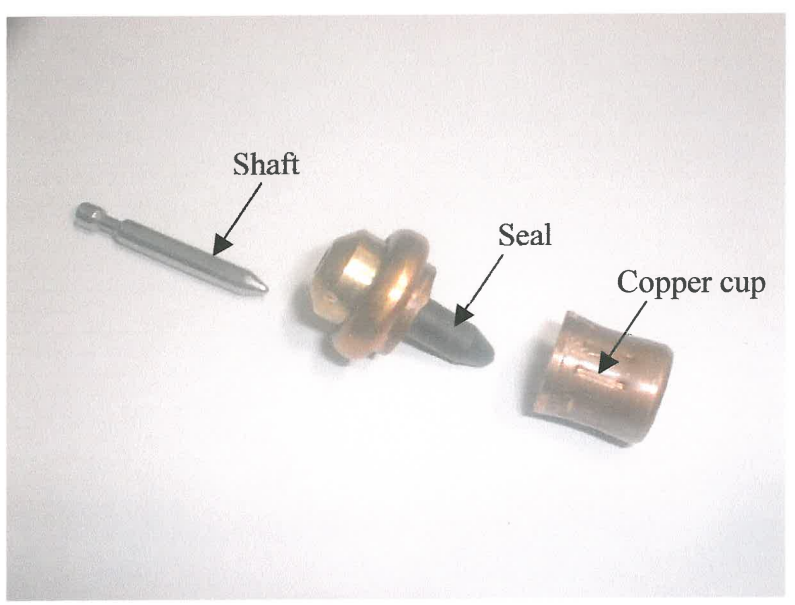

Figure 3. 23 Parts of a wax actuator

A study has been made on the heat transfer of such an actuator (Ogontz, 1998). It uses the normal heat flow equations to model such an actuator. The company, Ogontz, which sell valves actuated by wax actuators, indicates that these actuators have some level of hysteresis between heating and cooling. They indicated this with the graph in Figure 3. 24.

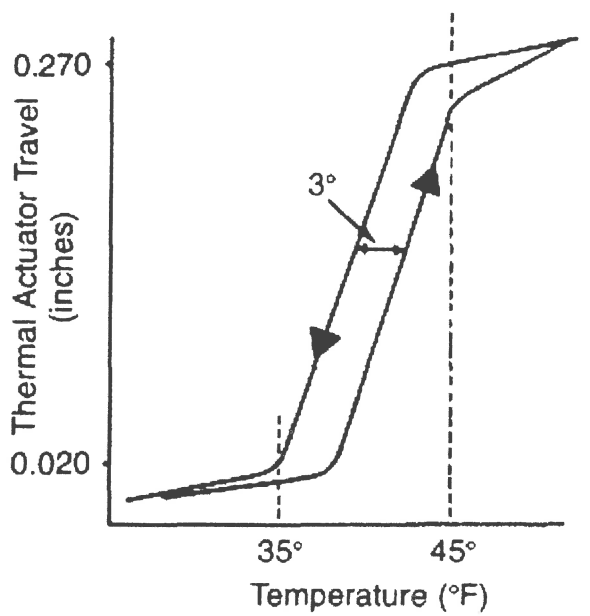

Figure 3. 24 Hysteresis loop of a wax actuator (Ogontz, 1998)

Note the low temperatures and small temperature range over which the actuator operates. This is another indication of the flexibility of these actuators.

The availability of the thermostats made it practically possible to use wax actuators and together with the excellent force and displacement characteristics, it was chosen as the actuators to be used. Because no information about the actuators themselves was available, it was necessary to characterise the actuator and determine its force and displacement characteristics.

The first test was to determine the maximum force and displacement of the actuator. For the maximum displacement, a dial gauge was used to measure the displacement while the actuator was heated. The experiment is depicted in Figure 3. 25.

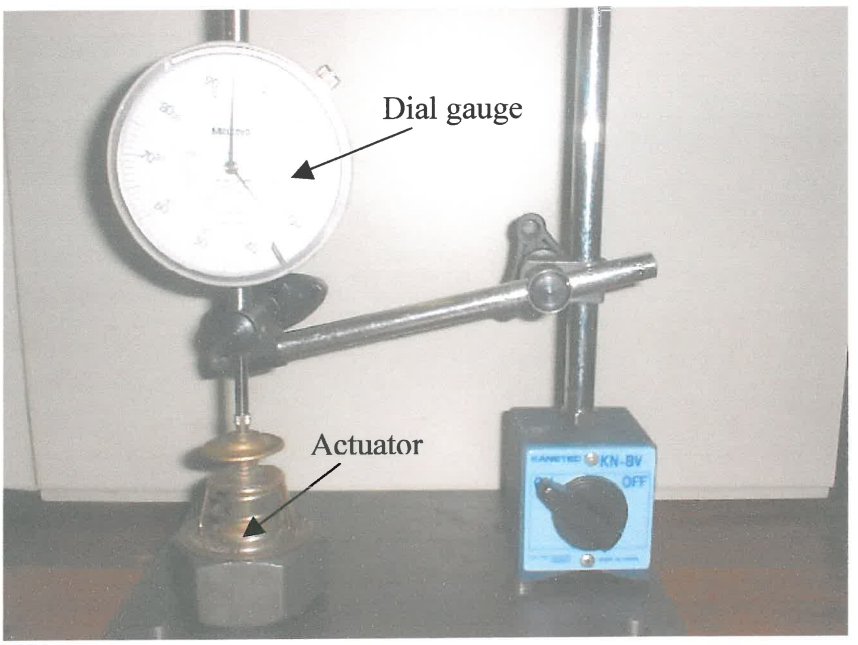

Figure 3. 25 Measurement of maximum displacement of wax actuator

The actuator was heated with hot air until it was extended to its maximum and the displacement documented. To measure the force of the actuator, the actuator was clamped with a load cell. Because of the hydraulic effect of the actuator, it was suspected that the maximum force would be dependant on the position of the actuator, so the maximum force was measured for different positions of the actuator. The force measurement experiment is depicted in Figure 3. 26.

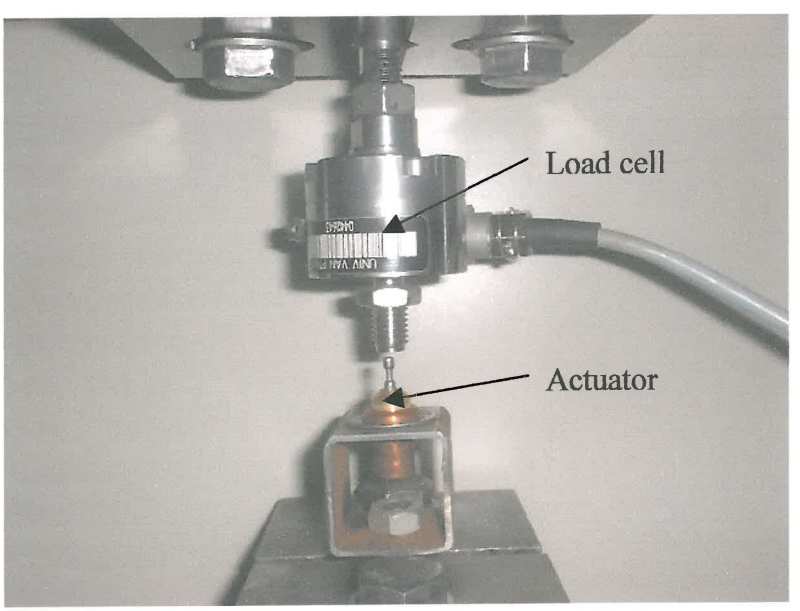

Figure 3. 26 Measurement of maximum force of wax actuator

The following results were obtained:

Table 3. 1 Maximum force of wax actuator

Displacement	Force [N]
0 mm	500 N
5 mm	430 N
10 mm	370 N

The last readings were at maximum displacement determined in previous tests. Further displacement caused the shaft end to be pushed out of the actuator. Considering that the shaft is a mere 3.8 mm in diameter, the pressure inside the copper cup must be 44 MPa and even higher in the walls of the cup for a force of 500 N on the shaft. The temperature of the actuator was also measured with a thermocouple to get an idea of the temperature range of these actuators. The actuator started to open at about 85 °C and the maximum force was obtained at about 170 °C. Therefore the temperature range of these actuators is quite large and it would probably be quite easy to control them.

The force values of the actuator are not enough to separate the springs 10 mm. As stated before, this will result in less stiffness change than ultimately possible, but because we do not want to use an amplification device, this actuator was accepted. It is actually possible to determine the amount that this actuator will separate the springs, by plotting the force vs displacement graphs of the actuator and the springs on the same graph as shown in Figure 3. 27. The displacement axis of the springs graph must only be scaled to display separation distance between the springs and not the displacement of the one spring.

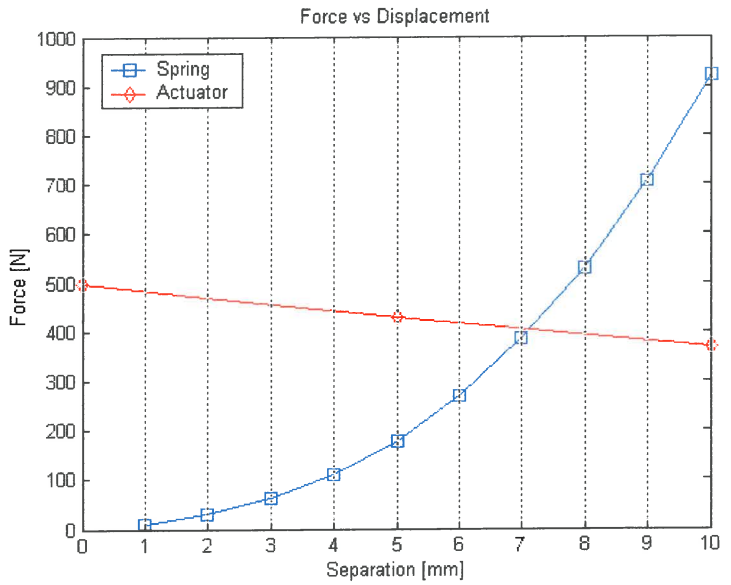

Figure 3. 27 Comparison between spring and actuator force

It can be seen that the actuator would be able to separate the springs about 7.2 mm. This will not result in the maximum possible stiffness change, but will be sufficient for the application in a LIVE isolator.

3.3.4.3 Heating mechanism

The wax actuator is controlled by controlling the temperature of the wax and through it the amount of phase transition in the wax from solid to liquid. Because the wax must be heated, it will be the best to insert a heating coil inside the wax, but that would require the manufacturing of an actuator, which was not part of this project. As far as external heaters are concerned, there are firstly normal heat conducting elements that consist of a heating element attached to the outside of the copper cup and the heat is conducted through the copper to the wax. Secondly, there are forced convection heaters that will make use of a hot fluid or gas that flows around the actuator to heat it.

The conduction is usually slower than forced convection, but the problem with forced convection is that a large volume of fluid or gas has to be heated that will result in a lot of losses. Reaction time is important, so forced convection will be used for

heating the actuator. To use water as medium to transmit the heat will slow the process down due to the large thermal capacitance of water. Air is a much better medium to use because it can be heated and cooled down more quickly than water.

The air must be heated well above 170 °C to be able to quickly heat the actuator to the maximum desired temperature. The easiest way to produce very hot air is with a hot air gun that can be bought commercially in a hardware store. A Bosch PHG 2 Hot air gun (Figure 3. 28) was used that was rated at 1.5 kW and that can produce air of up to 500 °C. The cover of the gun was removed and only the core-heating module used together with its integrated DC motor fan.

Figure 3. 28 Hot air gun

The wax actuator was mounted in the hole where the hot air exits, to form a closed controllable environment for the heating of the actuator. The fan was powered with 20V DC Power Supply, to supply the airflow necessary to heat or cool the actuator.

To be able to control the temperature of the actuator, the power dissipation of the heating coil in the hot air gun needs to be linearly controllable from zero to maximum power. This presented a challenge due to the large currents and voltages needed to drive the coil. The current of the coil can be calculated as follow:

$$I = \frac{P}{V} = \frac{1500W}{220V} = 6.818A \tag{3.1}$$

Therefore a controller is needed that can be controlled with a 0-10V DC signal from a computer and that can deliver about 7 A at 220 V or as close as possible to it. The controller can be AC or DC due to the fact that a coil is driven and that the type of current doesn't matter.

The first concept that was investigated was an AC drive that used pulse width modulation to control the power of the AC signal. The principle work as follows:

The basic idea is to control the part of the sine wave that is transmitted. A triac (switch) is used to switch the sine wave off after a certain time on every half cycle. Therefore only part of every half cycle will be transmitted and less power will be given. Figure 3. 29 show the concept in graphical form.

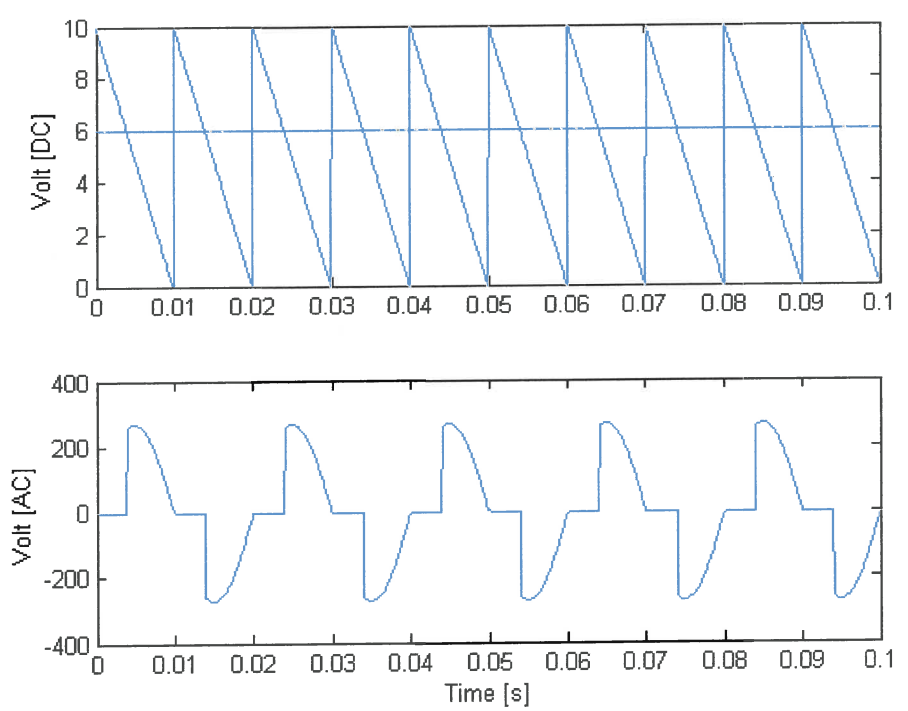

Figure 3. 29 Pulse width modulation principle

The top graph has a ramp signal in blue and a reference signal in green. The reference signal is the 0-10 V signal that is sent from the computer. The ramp signal is generated by the electronics by detecting the zero crossing of the sine wave. At every zero crossing (half cycle) a capacitor is charged to 10 V. The capacitor then discharges to 0 V in exactly the time of a half cycle of the sine wave after which is charged again to 10 V at the zero crossing. In this way the saw tooth signal is generated. The switching then works on comparing the reference and ramp signals to each other to see which one is larger. When the ramp signal is larger than the reference signal, the sine wave is switched off and when the reference signal is large than the ramp signal, it is switched on. In the bottom figure the chopped sine wave can be seen. Now it is possible to control the power delivered by changing the reference signal between 0 and 10 V.

This principle works well although it is not linear due to the fact that the power of a sine wave is not linear over time, but for the application it will work very well. The problem was that the electronic components to do all of this were not available. A complete unit will cost around R3000 to R4000, so another option was considered.

The other option was to look at DC signals. A normal DC motor controller was a perfect device to use for this application. It was capable of delivering the required current, but usually worked at about 180V instead of 220V. Such controllers were available in the laboratory and a simple test revealed that the coil could indeed be controlled with such a controller. After small modifications it was also possible to implement the control via a 0-10V DC signal. A Saftronics Quadro drive was used to control the heating coil.

Although convinced that the hardware was all in place, it was still not possible to drive the DC motor controller with the 0 – 10V DC signal from the computer. After investigation and research, it was discovered that the switching power supply of the computer caused the computer's ground to be different from the DC motor controller's ground. The difference was measured as a 110 V AC difference and that caused the problem. Therefore it was needed to isolate the grounds from each other.

For this an isolation amplifier was used. An isolation amplifier transfers the potential difference between the signal and ground, but isolates a difference of up to 1500 V between the grounds. A Burr-Brown ISO 124 isolation amplifier was used and its schematic can be seen in Figure 3. 30.

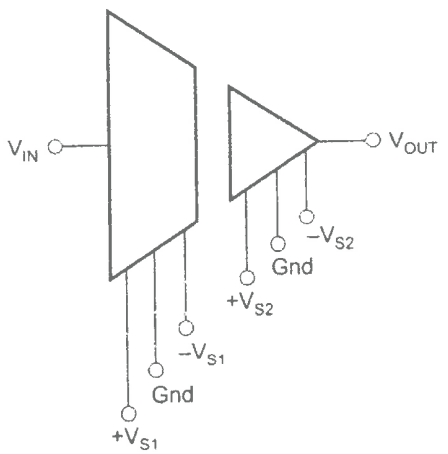

Figure 3. 30 Schematic of Burr-Brown ISO 124 isolation amplifier

This isolation amplifier was built into a box with connector terminals for the power supplies and signals and installed as depicted in Figure 3. 31.

Figure 3. 31 Isolation amplifier with power supplies

The four boxes at the bottom are four DC power supplies to supply the dual ρ12V power for the isolation amplifier. The box on top is the isolation amplifier with its connectors.

The whole heating mechanism worked well with the DC motor controller controlling the voltage over the coil of the hot air gun to control the temperature of the actuator. The DC motor controller is controlled with the computer through a digital to analog card. The whole system is schematically represented in Figure 3. 32.

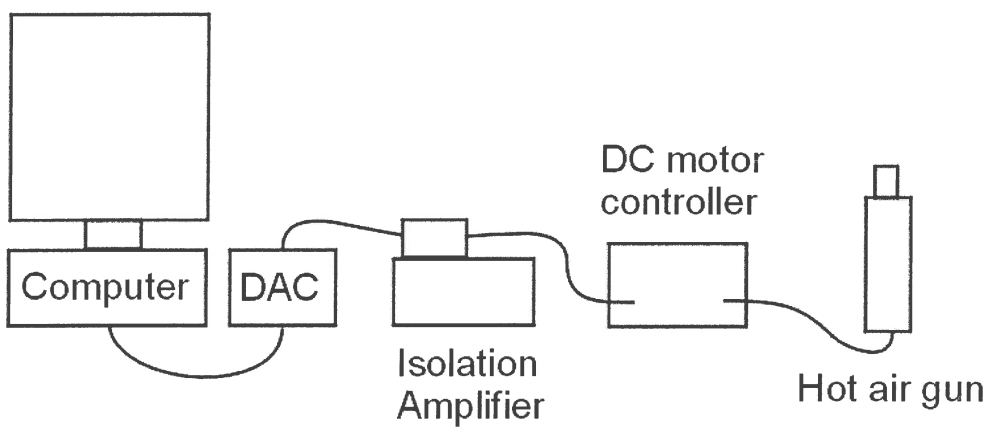

Figure 3. 32 Schematic of whole heating system

3.3.4.4 Displacement measurement

The last aspect to be addressed before control of the actuator can be obtained is the mechanism to measure the displacement of the actuator for the closed loop displacement control system. A very small and compact transducer is needed that has a stroke of at least 10 mm. A LVDT is an obvious choice, but it was found that a LVDT with 10mm linear stroke is quite large and would not fit in the allowable space. Another type of sensor was a laser displacement transducer that was a non-contact sensor, which is quite small and compact. The problem was that the sensor needed to be 80 mm from the surface that needed to be measured.

The last type of sensor considered was a linear potentiometer. Potentiometers are commonly used in volume controls and dimming devices and are usually found as rotational devices. Sliding potentiometers are also available and they are very small and compact with sufficient stroke and durability for the application. A 15mm linear sliding potentiometer (RS 855-377) was used and fitted to the springs to measure the displacement between the springs.

Calibration of the potentiometer against a laboratory quality strain gauge displacement transducer revealed that the potentiometer had a linear behaviour and a calibration factor was obtained. The potentiometer was supplied with 10V power and the output measured with an analog to digital converter on the computer.

The whole separation mechanism was designed and implemented complete with a displacement transducer and all the connections needed to control it with a computer. The control of the spring will be addressed in the next section.

3.3.5 Design of the controller

3.3.5.1 Characterisation of the system

It is important to characterise the open loop system before the loop is closed. The open loop system is the whole spring and actuator assembly and incorporates the isolation amplifier, DC motor controller and the potentiometer as shown in Figure 3. 33. Therefore the transfer function of the open loop system will be the transfer function between the voltage signal that is given to the DC motor controller and the output voltage signal from the potentiometer.

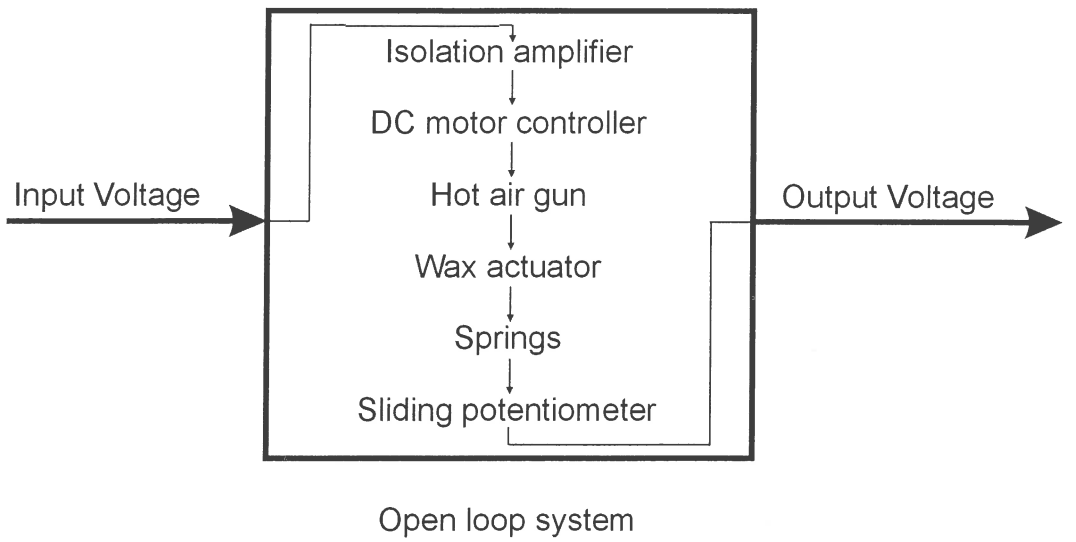

Figure 3. 33 Open loop system

The open loop system is obviously a very complex system, but to simplify the whole matter, the response of the open loop system will be measured and a theoretical model will be fitted to the response. This fitted model can be a first, second or higher order system, although it is preferable to keep it simple.

The first test that had to be done was to determine the range of the system. This is the steady state output of the system for certain input voltages. The system was therefore simply given a constant input voltage and the output was recorded after some time to make sure the system had stabilised. From these tests it was seen that the system starts to respond at 2 V and that maximum displacement of 7.2 mm is reached at an input of 3.4 V. This value would be the limit for the input to the open loop system.

It is difficult to determine whether a first or second order system will model the physical system the best, so both cases will be considered throughout the modelling and compared to each other at the end.

First we will consider a first order system. To characterise a first order system, it is necessary to measure the step response of the system. The step needed to be within the range of the system, so an input step between 2.2 V and 3.4 V was given to the system and the response measured. It was strongly suspected that the upward and

downward response of the system would differ, so both were measured. Figure 3. 34 give the results obtained.

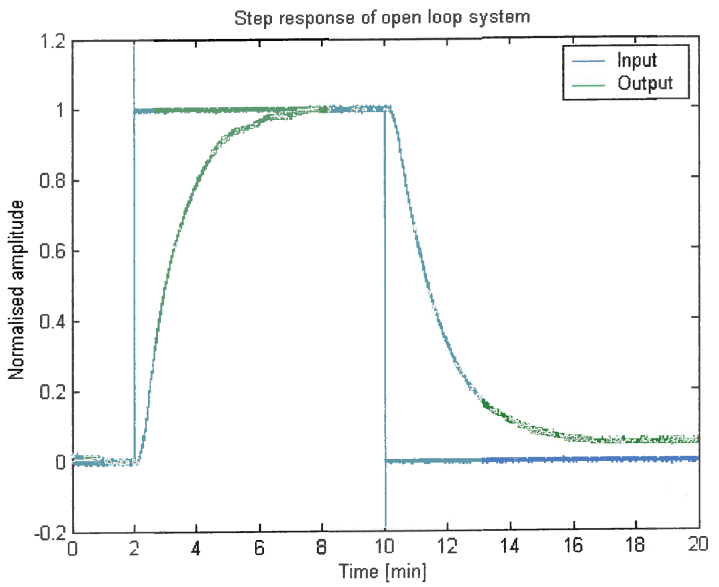

Figure 3. 34 Step response of open loop system

The input and output signals were scaled to fall between 0 and 1. To fit a first order system to such a response, the time constant has to be determined. Lets consider a first order system with the following transfer function (Theron, 2002):

$$G(s) = \frac{a}{s+a} = \frac{1}{\frac{1}{a}s+1} \tag{3.2}$$

The time constant of the system is defined as $\tau=1/a$.

To determine the response of the system with a step input the following can be done:

$$Y(s) = \frac{a}{s+a}\cdot\frac{1}{s} = \frac{a}{s(s+a)} = \frac{1}{s} - \frac{1}{s+a} \tag{3.3}$$

The 1/s is for the step input and the last term was obtained with partial fraction expansion. By performing an inverse Laplace transform, the response can be obtained as:

$$y(t) = 1 - e^{-at} = 1 - e^{\frac{-t}{\tau}} \tag{3.4}$$

After one time constant has elapsed, the response of the system would be the following:

$$y(\tau) = 1 - e^{-1} = 0.6321 \tag{3.5}$$

Therefore, any first order systems response would be 0.6321 after one time constant has elapsed from a step input to the system. If a step response was measured, the time at which the response reached 0.6321 will be the time constant of the system.

Figure 3. 35 give the results if this is implemented on the measured responses.

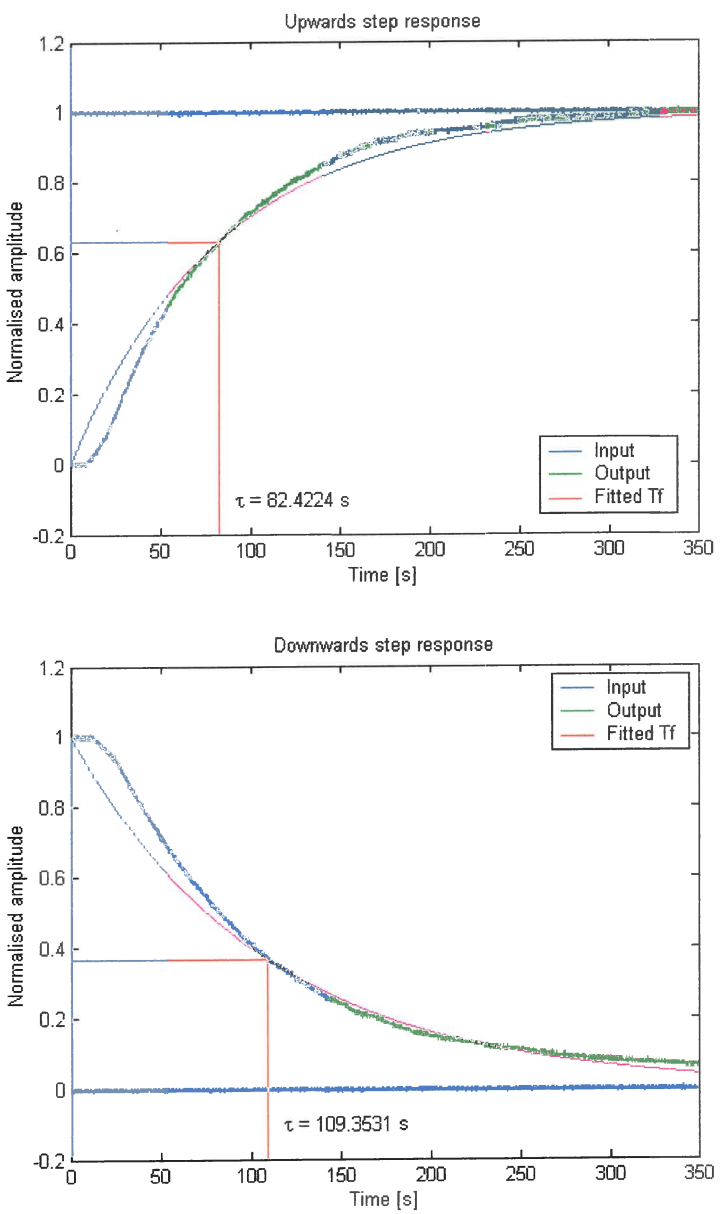

Figure 3. 35 Fitted first order system to upward and downward step response

It can be seen that the downward response is slower than the upward response due to the cooling that is less effective than the heating. The system is therefore non-linear, but to simplify the modelling, a linear system must be fitted to the response. Therefore the average time constant will be taken as 95.888s. The transfer function of the first order system would then be:

$$G(s) = \frac{1}{95.888s + 1} \tag{3.6}$$

The system has one pole at –0.0104.

If a second order system is considered, a step response is not sufficient to determine a theoretical model of the system. Therefore it was chosen to measure a transfer function of the system in the frequency domain. With a system with such a slow response, normal random or chirp excitation techniques are not possible. Therefore sine wave excitation must be used to determine the amplitude and phase of the transfer function at each frequency.

The same type of test was done as with the step response, except that a sine wave instead of a step input was used. The sine wave had an amplitude between 2.6 and 3 V and the period of the sine wave was varied between 100 and 5000 seconds to cover the whole frequency range. At each frequency measured, one sine wave input and output was measured. A theoretical sine wave was then fitted to the output to obtain amplitude and phase values for each frequency. A typical measurement is given in Figure 3. 36.

The red fitted line has been done using the following equation:

$$y(t) = A\sin(\omega t + \phi) \tag{3.7}$$

Where A is the amplitude, ω the circular frequency and ϕ the phase angle with which the output is leading the input. These values for all the frequencies were combined in a Bode diagram. For a second order system, the transfer function can be given as:

$$G(s)=\frac{\omega_n^2}{s^2+2\zeta\omega_n s+\omega_n^2} \tag{3.8}$$

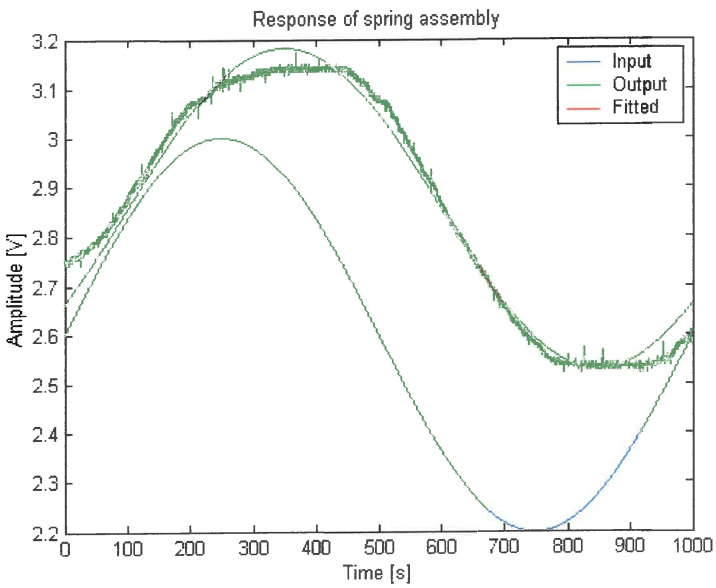

Figure 3. 36 Sine wave measurement

From this, the equations for the amplitude and phase graphs on the Bode diagram can be derived as:

$$\left|G(j\omega)\right|=\frac{1}{\sqrt{\left[1-\left(\frac{\omega}{\omega_n}\right)^2\right]^2+\left[2\zeta\frac{\omega}{\omega_n}\right]^2}} \tag{3.9}$$

$$\phi=-\tan^{-1}\left[\frac{2\zeta\frac{\omega}{\omega_n}}{1-\left(\frac{\omega}{\omega_n}\right)^2}\right] \tag{3.10}$$

Either one of these two equations can be used to fit a second order system to the measured bode diagrams. The result is given in Figure 3. 37.

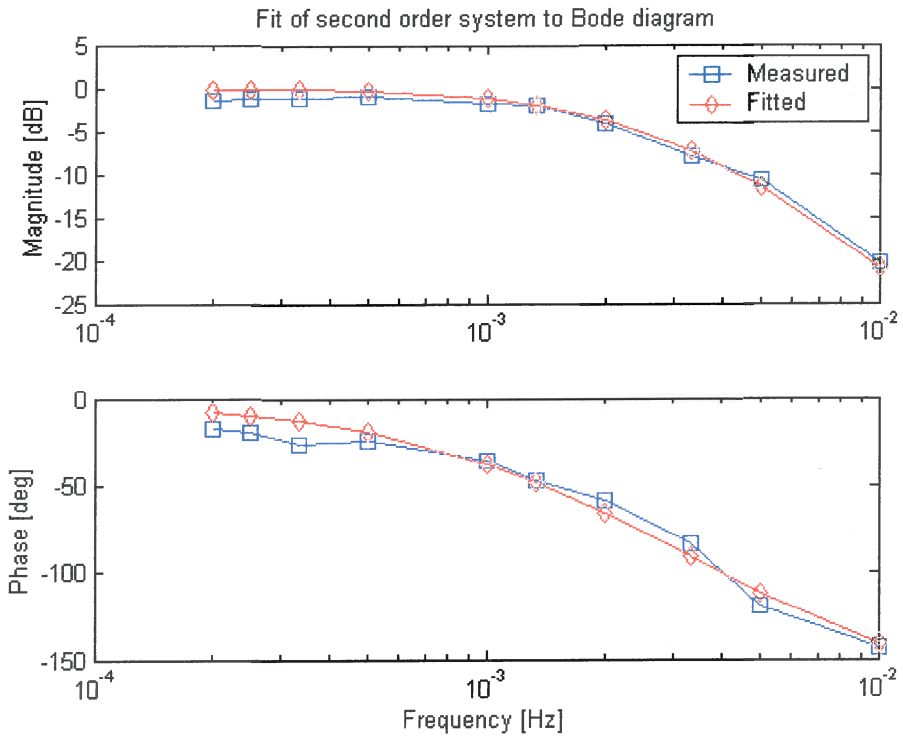

Figure 3. 37 Bode diagram

The fit on the phase graph yielded better results and were used. The fitted second order system had a natural circular frequency of ω_n=0.02055 rad/s and a damping ratio ζ=1.133. This means that it is an over damped second order system. The transfer function of the system is:

$$G(s) = \frac{P}{s^2 + 0.046566s + 0.0004223} \tag{3.11}$$

The poles of this system lie on -0.034228 and -0.012338.

The P in the numerator in equation (3.11) still has to be determined. This factor has to do with the amplitude scaling and time data is needed to determine the value accurately. A step response between 2.2 V and 3 V was measured and the step response of the system described in equation (3.11) was fitted to it by changing the value of P.

The result was a value of P=0.00046831. The response is given in Figure 3. 38.

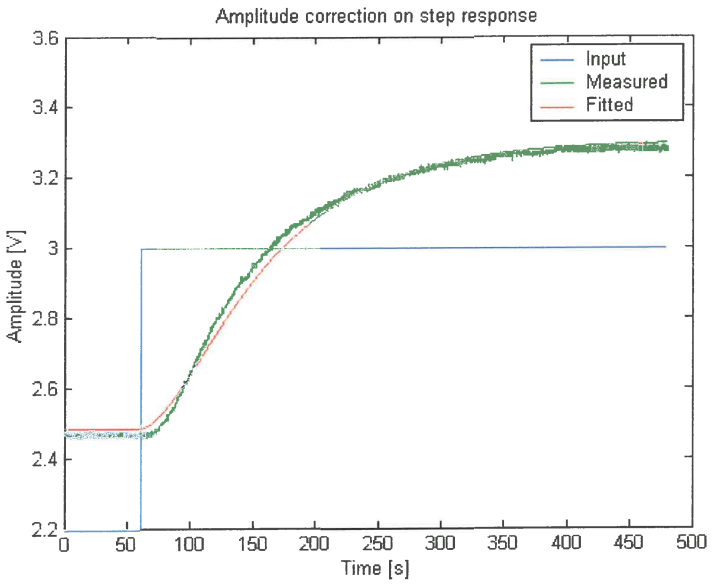

Figure 3. 38 Step response of second order system

Because the amplitudes needed to be fitted here, the graphs were not scaled to 0 – 1 like in the previous tests. Now a complete second order system was fitted to the response of the system and the transfer function of the system is:

$$G(s) = \frac{0.00046831}{s^2 + 0.046566s + 0.0004223} \tag{3.12}$$

3.3.5.2 Closing the loop

It was now possible to close the control loop by using displacement feedback to try and improve the system performance. The displacement signal was obtained from the sliding potentiometer that was installed for this purpose. The control was done in Matlab on the computer and a D/A and an A/D card were used to send and receive the voltage signals.

There are a few constraints on the system that must be pointed out.

- The input to the DC motor controller must be in the range of 0 to 10 V DC. 0 V will yield the maximum amount of cooling and 10 V the maximum amount

of heating. The control system may expect larger or smaller values, but it will be limited by this constraint.

- In a normal control system, an input of 0 V to the system will result in a static system at the current position. With this system it is definitely not the case because 0 V input will lead to maximum cooling and the system will not be stationary. To counter this effect, a constant of 2 V will be added to the output to the system. The reason for 2 V is because it is the value at which the system starts to respond.
- The system is non-linear. This must be remembered especially when the linear theoretical models are compared to the practical results.

The simplest way to close the loop is to use negative displacement feedback. This means that the current displacement is subtracted from the asked displacement, the difference multiplied with a gain input to the system. This is illustrated in Figure 3. 39 in block diagram form.

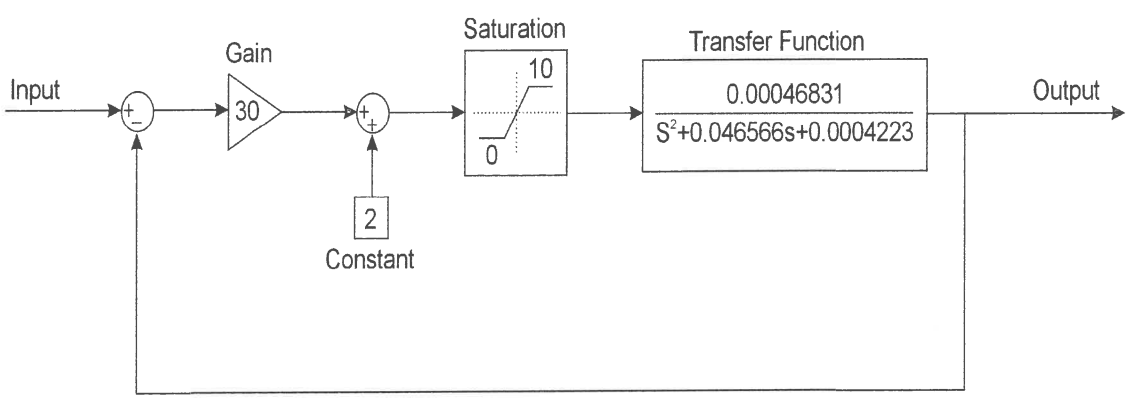

Figure 3. 39 Negative displacement feedback

Such a control program was written in Matlab and implemented on the actual spring. Different gain settings were used to obtain the best possible response. A step response was used to compare the different gain settings to each other. The responses were again normalised to range between 0 and 1 for easier comparison and are presented in Figure 3. 40.

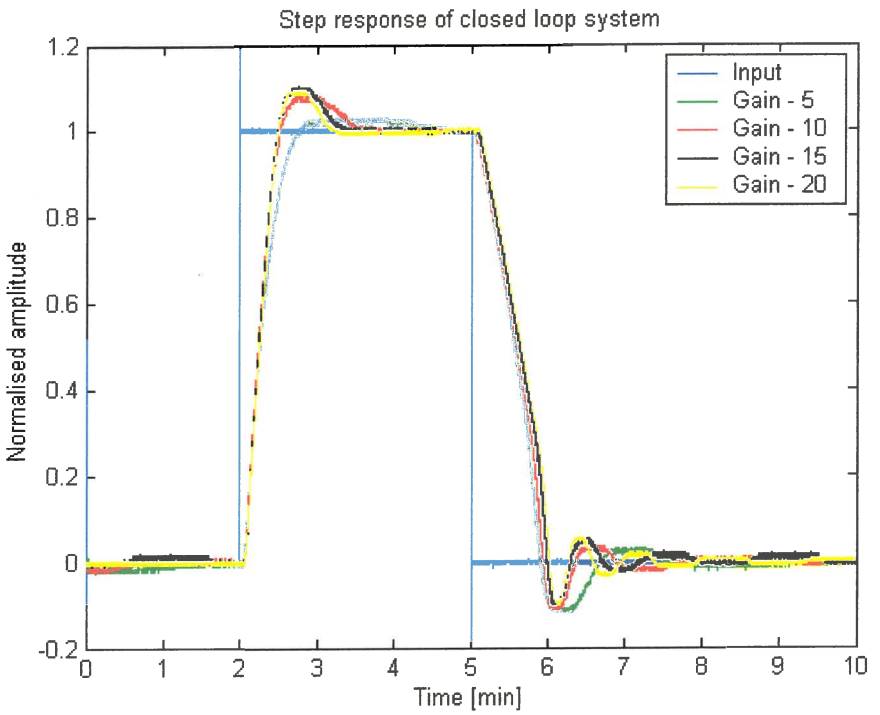

Figure 3. 40 Step response of closed loop system with negative displacement feedback

One of the important observations is the enormous decrease in response time compared to the open loop system. Furthermore, it can be seen that the response is improved by an increase in the gain, but that the overshoot and oscillation also increased. A large gain was needed to minimise the response time, but the overshoot and oscillation were unfavourable. The way to minimise overshoot and oscillation in a control system, is to implement a differentiator.

A differentiator determines the gradient of change of the system and adds that, multiplied by a differentiator gain, to the displacement of the system to account for the dynamics of the system. The effect will be that if the system is moving upwards, the differentiator will cause the measured displacement to appear larger than it actually is to account for the speed at which it is moving upwards. The control system will therefore start to brake the system before it reaches the desired position and therefore minimise the overshoot. The system is given in Figure 3. 41 in block diagram form.

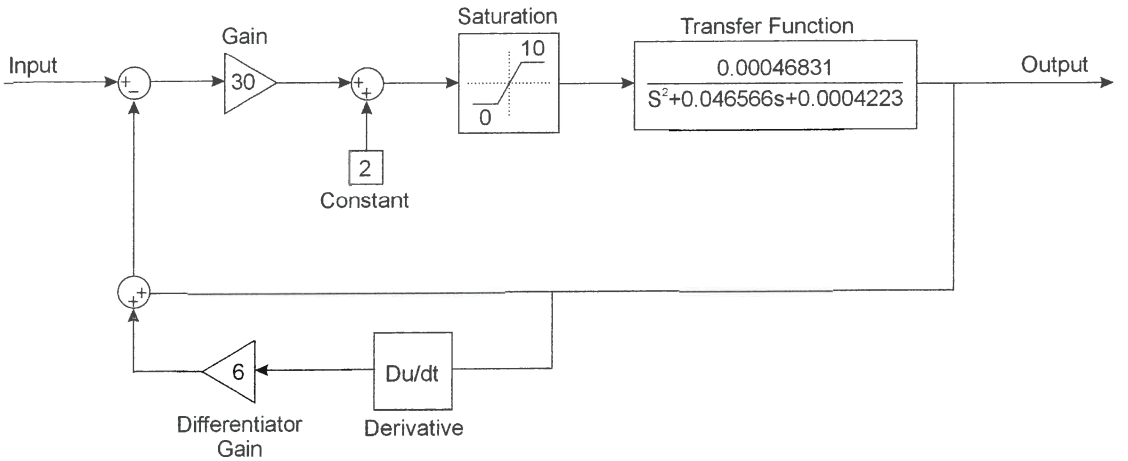

Figure 3. 41 Closed loop system with differentiator

This was also implemented in the Matlab program and the system was tested with this new control system. Different gains were again used to obtain the best results. Three gains were used (20, 30 and 40) and at each gain, the differentiator gain was varied between 2 and 10 and the response measured. Figure 3. 42 give the results.

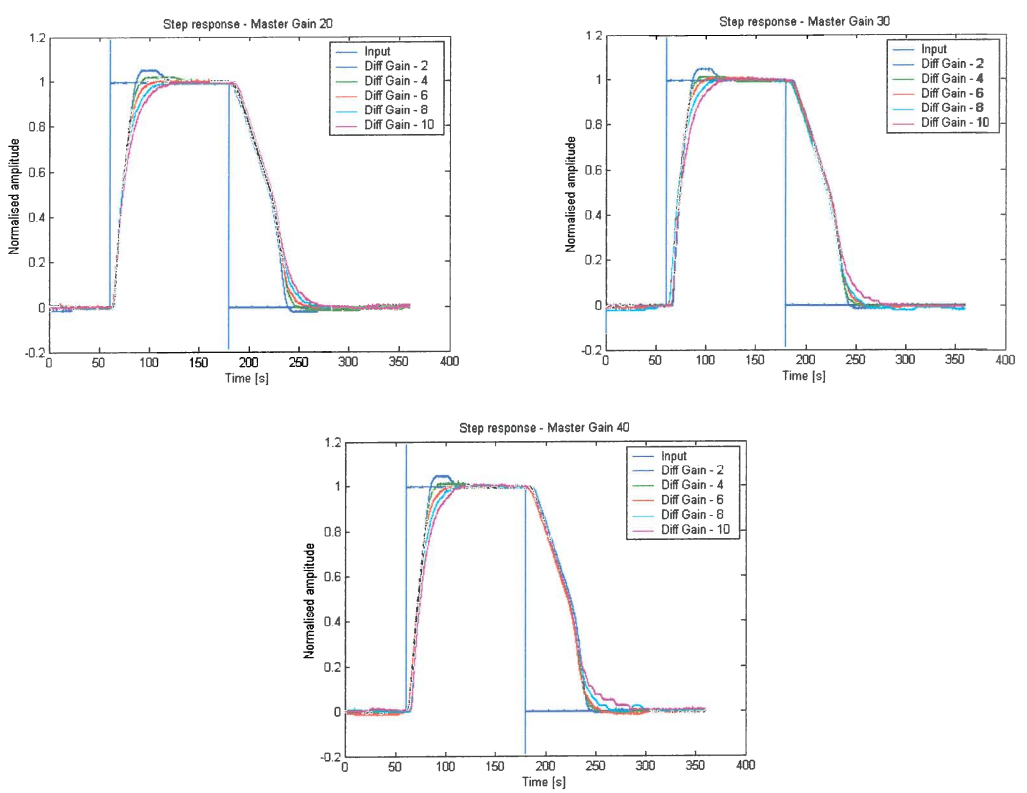

Figure 3. 42 Step response of closed loop system with differentiator

It can be seen that the overshoot and oscillation decreases as the differentiator gain is increased. The problem is now to determine which setting delivers the best response. To be able to determine this, an evaluation criterion must first be chosen. The different criteria available (Dorf & Bishop, 1995) are explained in Figure 3. 43.

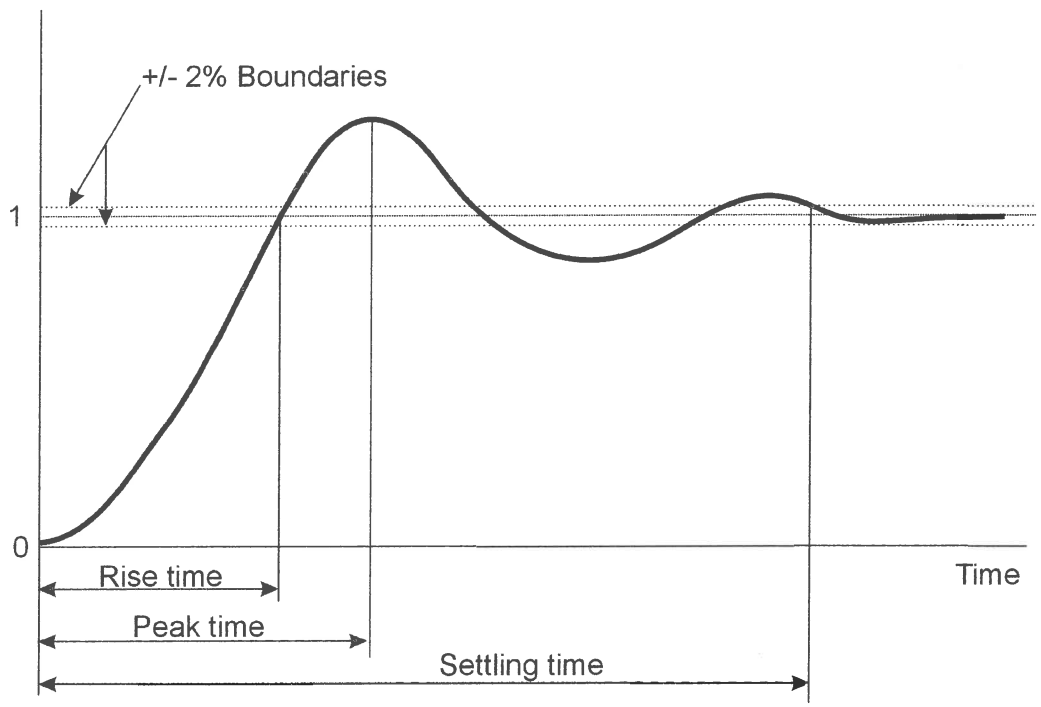

Figure 3. 43 Evaluation criteria

Rise time: The time the system takes to first reach the final value.

Peak time: The time the system takes to reach its maximum value.

Settling time: The time the system takes to settle within a ±2% region around the final value.

If the spring is considered, the stiffness has effectively changed to the new value once the system has reached the desired value and has settled on it. Therefore the settling time would be a valid criterion to evaluate the performance of the system.

The settling time for all the tests were determined for the upward and downward step as shown in Figure 3. 44.

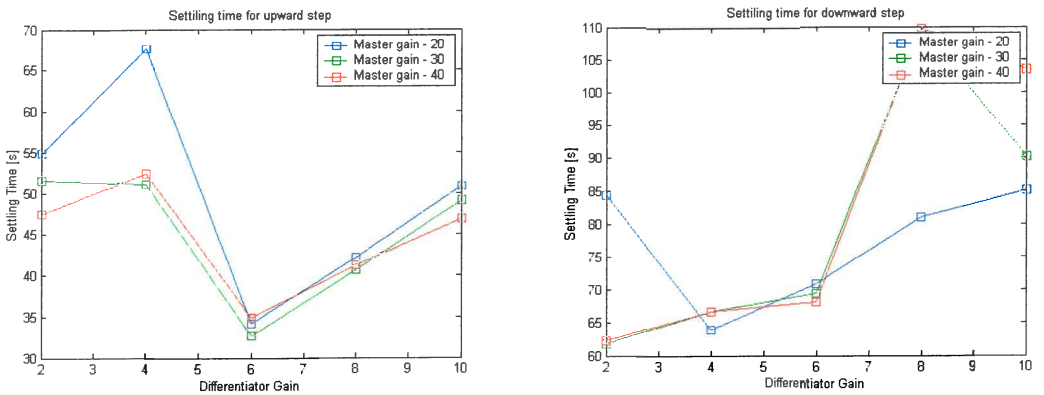

Figure 3. 44 Settling time for closed loop system with differentiator

From the graphs, it can be seen that a combination of a gain of 30 and a differentiator gain of 6 will yield the best performance for the upward step. For the downward step, there are a few discrepancies between the different gain settings. It seemed that a differentiator gain of 2 would yield the best performance. On closer investigation of the time signals, it was however noted that with a differentiator gain of 4 the graphs of the 30 and 40 gain settings just crossed the –2% boundary and increased settling time. It could very well have been that the graphs had not crossed the boundary, which would cause the differentiator gain of 4 to give the best performance as is the case with the gain of 20. Therefore a differentiator gain of 4 was chosen for the downward step. It is possible to implement different gain settings for upward and downward movement on the computer due to the fact that a Matlab program is used to control the settings. The block diagram of the final closed loop system is given in Figure 3. 45.

To validate the theoretical models that were developed, their predicted response must be compared to the measured response of the closed loop system. The previously mentioned constraints on the system complicated the analytical determination of the response of the theoretical systems. It is obviously necessary to include these constraints in the models otherwise the response would certainly not be comparable to the real system. Simulink, a simulation tool of Matlab, provided the tools to numerically implement all the constraints and determine the response of the theoretical models to certain inputs.

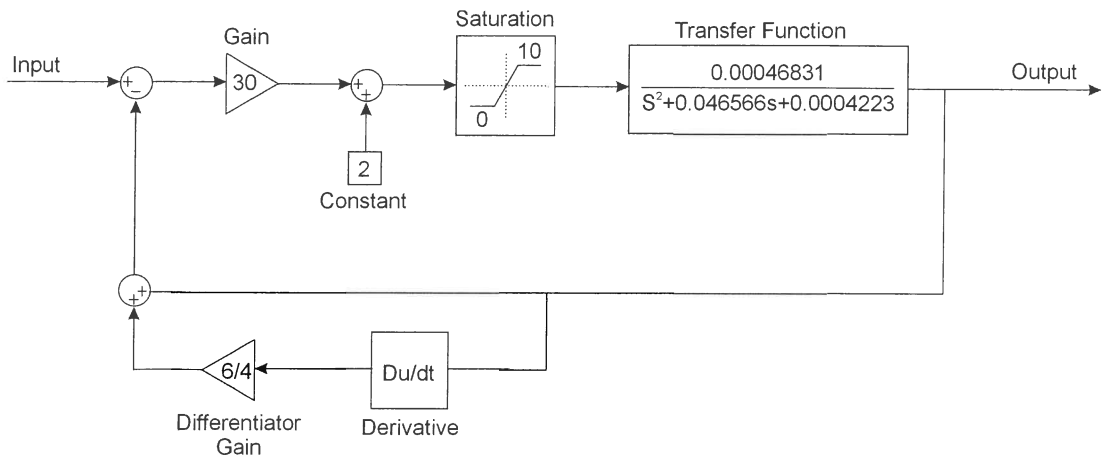

Figure 3. 45 Final closed loop system

A simulink model was developed to incorporate the first and second order system and the constraints to the system. The second order system model is given in Figure 3. 46.

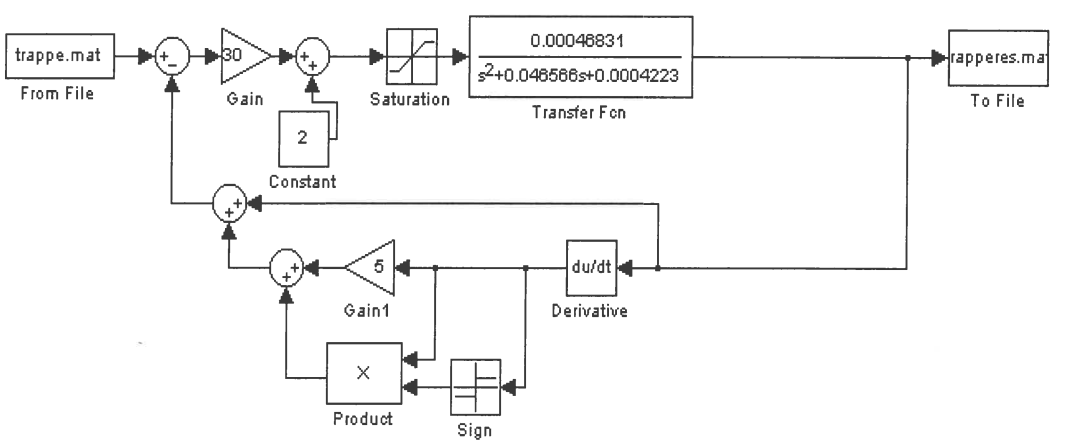

Figure 3. 46 Simulink model of second order closed loop system

To test the system, an input file was created that stepped from 2.2 V (in steps of 0.2 V) to 3.2 V and down again. The response of the actual system with the file as input was measured and the same file was used as input for the simulink model. The responses were plotted on the same graph to be compared in Figure 3. 47.

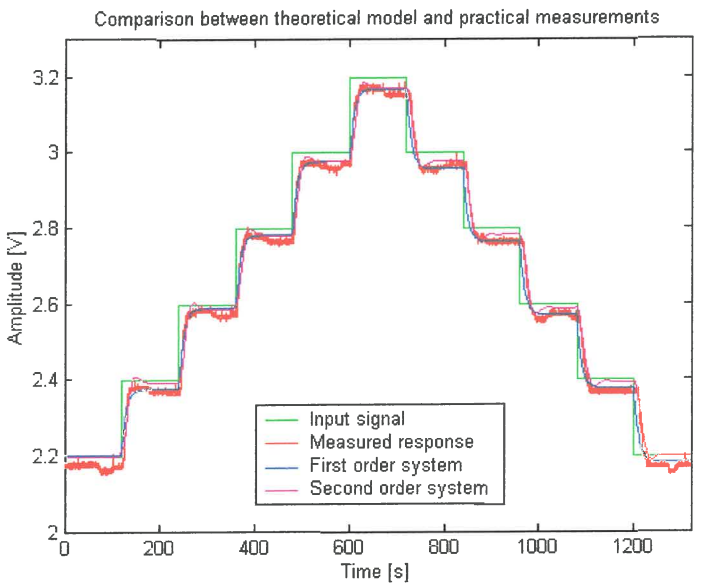

Figure 3. 47 Comparison between theoretical model and actual system response

It can be seen that both the first and second order models gave a quite good correlation to the actual response of the system. There is actually a very easy way to conclude that the second order system is the better fit: The overshoot. A first order system cannot have any overshoot and out of the tests it is clear that the system definitely has overshoot. Therefore the second order system yields the best theoretical model of the physical system.

From this last test it can also be seen that the smaller steps were completed in about 30 seconds. Therefore it was concluded that the time response of the system is adequate for the application. The last spring characteristic necessary to determine whether the spring complies with the specifications, are the stiffness and damping values. This will be reported in the next section.

3.4 Testing of the spring

The whole spring assembly with actuator, hot air gun, displacement transducer and frame was assembled to be able to test the spring on a servo hydraulic actuator. The Solid Edge model of the actuator and spring assembly is depicted in Figure 3. 48.

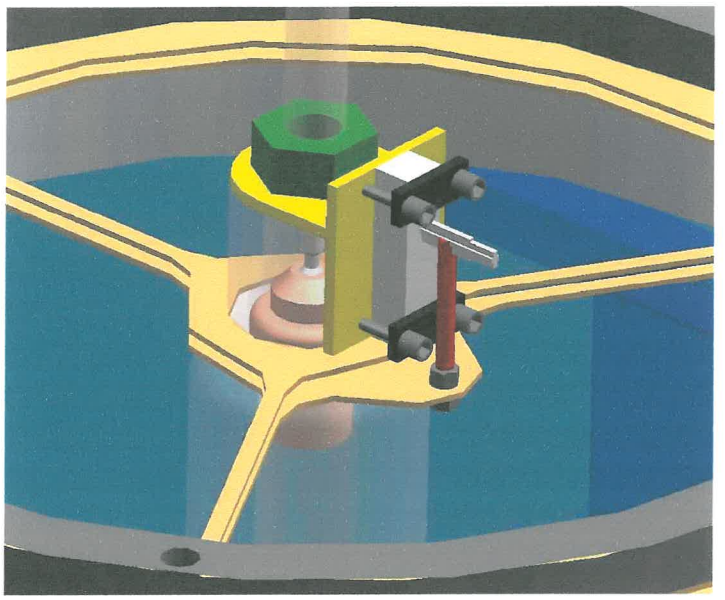

Figure 3. 48 Spring assembly

The actuator is between the two springs with the hot air gun beneath it. A cup-like part was made to fit on top of the actuator to act as a platform for the actuator to push against. The sliding potentiometer was fitted to measure the displacement between the two springs. In real life, the whole assembly is depicted in Figure 3. 49.

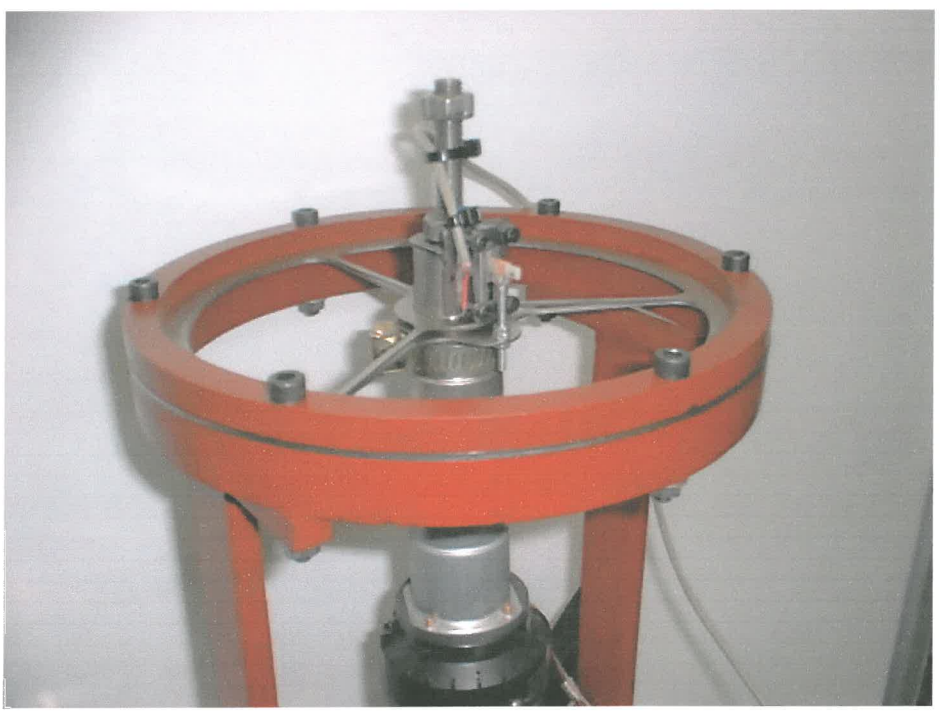

Figure 3. 49 Experimental spring assembly

This assembly was clamped at the top and the bottom was fixed to the servo hydraulic actuator. The force and displacement signals were measured and the stiffness and damping were determined. The results are given in Figure 3. 50.

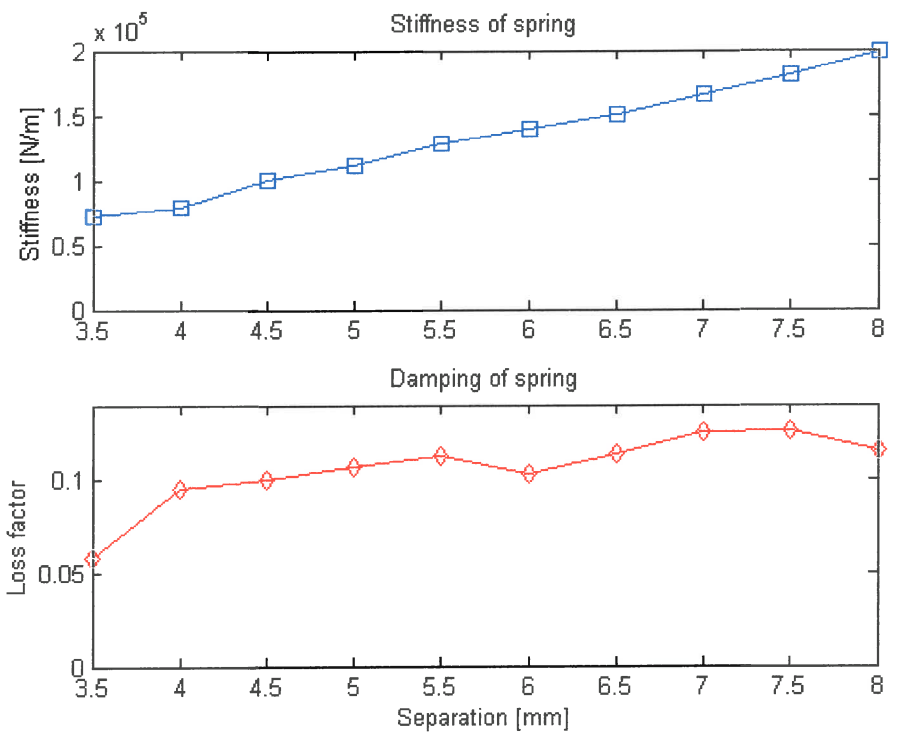

Figure 3. 50 Stiffness and damping of spring assembly

A stiffness change of 2.73 times was obtained and the loss factor ranged between 0.06 and 0.12. These values fulfil the requirements that were set for the spring and therefore the spring can be implemented on a LIVE isolator. The above stiffness and damping values can be used to design the LIVE isolator to obtain a desired frequency response.

3.5 Conclusion

In this chapter the whole development process of the variable stiffness spring was addressed starting at the concept phase with three concepts and then progressing to

the detailed design phase of the chosen concept. A control system was also implemented to control the stiffness of the spring by using a wax actuator. Closed loop displacement and velocity control was used to achieve the desired time response and accuracy of the actuator.

The final variable stiffness spring complied with the requirements set at the beginning and performs even better than needed, especially in light of the amount of stiffness change available. It is concluded that a successful variable stiffness spring was developed that can be implemented on a LIVE isolator or on various other practical applications.