

24-491/492 Departmental Research Honors Spring Design and Optimization

Christopher Atwood
Carnegie Institute of Technology
Carnegie Mellon University
Pittsburgh, PA 15213

May 4, 2005

1 Goal and Scope of Project

As described in CMU-RI-TR-04-24, two extremes of actuator design for running robots are high-bandwidth actuators and tuned natural dynamics. Theoretically, a high-bandwidth motor would allow any dynamic response to be simulated and outputted, albeit with low energy efficiency. But the physical world limits motor response characteristics, which limits this type of actuator to use in low-frequency situations. On the other extreme, a carefully tuned mechanical system can be very efficient and can operate at arbitrary frequencies, but has a narrow operating range.

The Actuator with Mechanically Adjustable Series Compliance (AMASC) takes a middle ground. There is a powerful motor that can respond to low-frequency inputs to the system, and an adjustable spring that responds to high-frequency inputs. During a running gait, the springs store most of the energy required to maintain the gait; the motor adds just enough to compensate for losses in the system.

The AMASC uses fiberglass bar springs that are moment loaded at each end and that are mounted on the thighs of the robot. Limitations of this configuration include an uncontrolled degree of freedom (the springs can vibrate). But more importantly, springs thick enough to hold energy for a running gait would be strained to failure. A design using longer and thicker springs could increase energy storage to a functional level, but would increase the size of the robot past its already larger-than-life size.

It was recognized that for the Next-Generation Actuator with Mechanically Adjustable Series Compliance (NG-AMASC) a different configuration of springs would have to be designed. The aim of the redesign is to remove the limitations (limited energy storage, extraneous degrees of freedom, increasing of sprung mass) of the AMASC spring system while keeping the springs as inexpensive and light as possible. This document outlines the processes and results of the investigation into possible designs of leg springs for the NG-AMASC.

2 Function of Springs

As stated previously, the springs on each leg of the AMASC and NG-AMASC store impulse energy from each leg impact with the ground for use in the next stride. It has been decreed that the absolute maximum load these spring will experience is the 30 kilogram robot falling from a height of 1 meter. Thus, to survive this loading situation each leg must be able to store 300 joules.

This loading situation is an extreme, and will not be encountered in normal operation. While surviving this rare loading is important, perhaps more important is how the springs function in normal gaits. The loading during that normal gait is currently unknown, but a simulation of the robot and leg actuators will eventually provide the desired information. Results of the simulation will greatly influence spring design, but for now all we can design to is the maximum load because it is the only data point available.

2.1 Why Opposing?

Each leg has two opposing springs to allow the leg stiffness to be changed without altering knee or leg position. This mimics the opposing muscle pairs found in nearly every muscle group in mammals.

2.2 Why Variable Stiffness?

Variable leg stiffness is required to maintain a consistent gait when running on varying ground surfaces. A robot without variable leg stiffness moving from hard to soft ground would see uncontrolled gait change, and if the surface change is significant the robot might fall. A robot equipped with opposing springs, allowing for variable leg stiffness, is capable on variable floor surfaces. Additionally, variable leg stiffness can be used to control velocity and other gait parameters.

2.3 Why Non-Linear?

During impact, the two springs on each leg move from some pretension position by some delta. The springs move differentially, thus one spring is being loaded and is storing energy, while the other is being unloaded and is releasing energy. If the springs were linear there would be some net energy storage, but the amount of energy storage would be the same at all pretension positions. To have variable stiffness, and this variable energy storage, the spring systems must be non-linear.

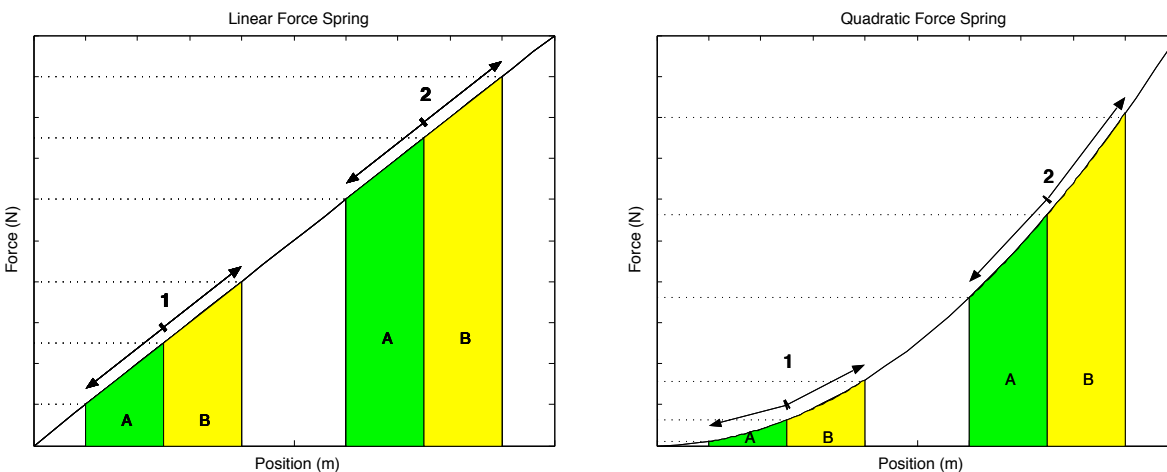


Figure 1: Spring Force Diagrams

Spring Type	Pretension Position	Spring A (J)	Spring B (J)	B - A (J)
Linear	1	1.1667	2.1667	1
	2	4.5	5.5	1
Quadratic	1	0.2167	0.7167	0.5
	2	3.05	4.55	1.5

Table 1: Spring Energy Storage Values

Figure 1 shows the energy storage of linear and quadratic spring systems at two different pretension positions, numeric values are shown in Table 1. Since the springs A and B are opposing, the system's net energy storage is the difference between them. The table shows that pretension position has no effect on a linear spring system. To get a variable system stiffness and variable system energy storage, the spring system used must be non-linear.

It is worth noting that the physical springs can be linear. Spiral pulleys connected to the leg springs can give the spring system an arbitrary force function. This investigation assumed that the spring systems would have quadratic force function, but any relation can be achieved by the spiral pulleys.

3 General Spring Optimization

All materials have some strain where they fail and break. Obviously, loading a spring on the NG-AMASC to that point would be disastrous. But there is a narrower constraint that must be considered, the elastic limit of the material. If the spring is deformed past the elastic region and into the plastic region it is forever deflected by some amount, limiting its energy storage from then on. Thus, to maintain consistent performance over the spring's life, it must only operate in its elastic region.

Some materials, like Carbon Fiber and Fiberglass, go directly from their elastic region to failure. For these materials it was decided that the elastic limit was 75% of that failure point. Titanium's failure point is around 8% strain, but most of that is in the plastic region, so the elastic limit of titanium was assumed to be 1%. These values give a factor of safety of about 1.3, which allows for some error in calculations or construction without compromising performance.

As described above, during impact the springs bend from some initial pretension position by some delta to a final position. An additional factor of safety was applied by assuming that the impact delta during this maximum loading would bend the spring by only 40% of the elastic limit from the initial pretension of 50%. Thus, one spring bends from 50% to 10%, and the other from 50% to 90%.

3.1 Physical Property Values

These physical property values were used in all calculations.

Material	Density	Modulus of Elasticity	Elastic Strain Limit
Fiberglass	1800 kg/m^3	$3.79 \cdot 10^{10} Pa$	1.5 %
Carbon Fiber	1750 kg/m^3	$2.34 \cdot 10^{11} Pa$	1.31 %
Titanium Beta C	4280 kg/m^3	$1.04 \cdot 10^{11} Pa$	1 %

3.2 Deriving Spring Rate

During impact one spring is being loaded, and one is being unloaded. During maximum loading, the sum of these energy changes must equal the energy storage requirement or the actuator will fail. And since the springs are non-linear the amount of energy stored is constantly changing. To determine the energy storage amount we must integrate the spring force equation from its initial pretension position to its final position.

It was assumed that the helical spring has a maximum deflection of 360° because this is the largest deflection before the tension cable would wrap on top of itself. As stated previously, during maximum loading the spring travels from a pretension of 50% to either 90% or 10%. This means the helical spring travels from 180° to 324° or 36°, depending on which spring it is. The bar spring was assumed to have a maximum deflection of 120° because this seems a reasonable bend amount to have in the torso of the robot. And thus during impact, the bar deflects by 40% to 12° or 108°.

This series of calculations and their results is shown below.

$$E_{required} = m \cdot g \cdot h = \Delta E_A + \Delta E_B$$

$$SpringForce(x) = K \cdot x^2$$

To determine the amount of energy stored or released in a spring we integrate the spring's force equation from the spring's initial pre-tension position by the impact delta. This calculation applied to springs A and B is below.

$$\Delta E_A = \int_{x_3}^{x_3 - \Delta x} SpringForce(x) \cdot dx = \frac{K}{3} \cdot \left((x_3 - \Delta x)^3 - x_3^3 \right)$$

$$\Delta E_B = \int_{x_3}^{x_3+\Delta x} SpringForce(x) \cdot dx = \frac{K}{3} \cdot \left((x_3 + \Delta x)^3 - x_3^3 \right)$$

Substitute ΔE_A and ΔE_B into the $E_{required}$ equation, and solve the resulting equation for K, the effective spring constant of the joint.

$$K = \frac{3 \cdot E_{required}}{(x_3 + \Delta x)^3 + (x_3 - \Delta x)^3 - 2 \cdot x_3^3}$$

$$K = \frac{3 \cdot E_{required}}{6 \cdot x_3 \cdot \Delta x^2} \quad (1)$$

where:

Variable:	Description:
m	Mass of robot
h	Jump height of robot
ΔE_A	Energy released from spring A
ΔE_B	Energy stored in spring B
K	Spring constant
x_3	Mechanical pretension position
Δx	Mechanical travel during impact

Results of calculation:

Spring Type	Pretension Position	Impact Delta	Spring Rate
Helical	180°	144°	7.41 N/m
Bar	60°	48°	200 N/m

Notice that to store the required amount of energy the bar spring needs a much larger spring rate. This is due to its much smaller range of travel, 120° versus 360°, and the cubic relationship between deflection and energy storage. The much larger spring rate is feasible to build because the bar spring does not need the extreme flexibility the helical design requires. The helical design will have a longer and thinner flexure, while the bar design will have a shorter and thicker one. The mass of the two spring designs should be the same if the same material is used.

3.3 Deriving Wall Thickness

For every material there is some deflection at which failure occurs. In beam bending the outermost layer undergoes the most deflection, and thus the thickness of the beam is directly proportional to the material's maximum stress percentage. Other factors in maximum thickness are flexure length and bend angle. As the flexure gets longer or as bend angle decreases the thickness can increase.

$$y = \frac{\epsilon \cdot L_f}{\theta_{max}} \quad (2)$$

where:

Variable:	Description:
ϵ	% material max strain
θ_{max}	Spring bend angle
L_f	Length of spring flexure
y	Distance from max strain to zero strain

3.4 Deriving Width

To calculate the rest of the physical parameters of the spring we can use the spring energy storage equation.

$$E_{stored} = \frac{E \cdot b \cdot h^3 \cdot \theta^2}{24L_h} \quad (3)$$

Note that E_{stored} is the maximum of ΔE_A and ΔE_B and is larger than $E_{required}$. The wall thickness found previously is h and the equation is solved for b .

$$b = \frac{24 \cdot E_{stored} \cdot L_f}{E \cdot h^3 \cdot \theta^2} \quad (4)$$

where:

Variable:	Description:
E	Modulus of elasticity
b	Width of flexure
h	Depth of flexure
L_f	Length of flexure
θ	Angular deflection of spring

4 Torsionally Loaded Helical Spring

Unlike a normal helical spring that is vertically loaded and compressed, a helical spring in the NG-AMASC would be torsionally loaded and would twist. This loading is analogous to moment loading the ends of a flexure strip. Of course as the helix angle increases, the moment loading goes off-axis and becomes a torsional loading. This twists the strip and causes non-uniform deflection and shear stress in the flexure strip. If the spring were made of a laminate material like Fiberglass or Carbon Fiber this has the potential to delaminate the layers and cause the spring to fail prematurely.

In torsional loading, a helical spring has a distinct advantage over a solid tube spring in the NG-AMASC. It goes back to the tradeoff between stiffness and deflection. A solid tube has high stiffness and low maximum deflection when loaded torsionally. Conversely, a helical spring has low stiffness and large maximum deflection, due to the longer flexure length. In the NG-AMASC the latter is greatly preferred because the longer travels mean that stresses in the cables and parts of the actuator will be lower. Additionally, the longer travels keep the reduction ratios lower, which simplifies the design and lowers reflected inertia.

4.1 Wall Thickness

In a helical spring the flexure length can be calculated from the length of the spring and the angle of the winding as expressed in Equation 5

$$L_f = \frac{L_{helix}}{\sin(\theta_h)} \quad (5)$$

This equation is plugged into Equation 2 and spring thickness can be determined. The results of this calculation, assuming wall thickness is twice y and with an arbitrarily chosen spring length of $0.4m$ and helix angle of 15° , are in the following table. It shows that allowable wall thickness increases as maximum strain increases.

Material	Maximum Strain	Wall Thickness
Fiberglass	1.5%	0.74 cm
Carbon Fiber	1.32%	0.65 cm
Titanium Beta C	1%	0.49 cm




Figure 2: CAD Render of Helical Spring

4.2 Width and Weight

As shown previously, each spring material has a different maximum strain, which means that each has a different maximum wall thickness. The amount of energy stored with each spring is the same, thus the width of each spring is different. Results of Equation 4 and overall weight calculations are shown in the table below. Weight was determined by multiplying volume of the spring by the density of the material. It does not account for the non-spring material at each end of the helix that is used for attachment to the actuator. The proportions of these weights will be the same no matter what spring geometry is chosen, thus Carbon Fiber will always be lightest and will therefore be the most desirable spring material.

Material	Width	Weight
Fiberglass	2.82 <i>cm</i>	0.58 <i>kg</i>
Carbon Fiber	0.67 <i>cm</i>	0.12 <i>kg</i>
Titanium Beta C	3.47 <i>cm</i>	1.13 <i>kg</i>

Once the ballpark dimensions of a helical spring were determined, we started looking into sources to supply these springs. We located a few companies that could wrap Fiberglass or Carbon Fiber strands into a near-net helical spring shape. This seemed promising because machining a helix from a solid composite tube would be difficult. Unfortunately, these manufacturers had no past experience with composite springs and wanted much larger orders than we felt comfortable placing for a prototype spring. So we turned to the Titanium spring idea. We located a manufacturer who has designed and fabricated many titanium springs, both helical and non-helical. What prevented us from pursuing on this idea was the estimated \$300 per spring cost and the large weight of the Titanium design. Adding 2.5*kg* to each leg would nearly triple their mass and thus would significantly increase the actuator’s motor size. As an alternative to leg mounted Titanium springs we next considered using bar springs mounted external to the NG-AMASC. Since two NG-AMASCs are going to be incorporated into a biped robot, as the leg length actuators, there are additional options for mounting that did not exist in the AMASC.

5 Vertically Loaded Bar Spring

A vertically loaded bar spring has one major advantage over a torsionally loaded helical spring: it is a much simpler design. We originally swayed away from this configuration of springs because of its much larger size, which is a significant concern if the springs are mounting on the leg. However, if the springs were mounted in the torso of the robot they could be almost arbitrarily large. The major obstacles to this type of mounting was finding a way to have hip rotation not move the leg-length springs. Clever use of cable differentials allowed us to decouple these two motions, opening up torso mounted bar springs as a viable spring option.

Mounting springs in the torso of the robot allows them to be much larger than leg mounted springs for the same overall robot size. Also, physical attachment of the spring is simpler because no four-bar linkage is required (used in the helical design to prevent rotation but allow linear movement of the spring’s non-loaded end). And, low-quantity manufacturing costs for a bar spring are much lower than for a helical spring.

But, uniform loading is critical when trying to store the maximum amount of energy in a spring. If there are stress concentrations the spring will fail at those points before maximum loading is reached. A torsionally loaded helical spring has uniform loading intrinsically, due to the moment loading, but a vertically loaded bar spring does not have this characteristic. To prevent stress concentrations in a bar spring it must bend in a perfect circle. Thus, any bar spring must have a circular support underneath, which unfortunately adds weight to the robot. In the design of the NG-AMASC this support was designed into the main structure of each leg, which serves other purposes, and thus weight gain was minimal.




Figure 3: CAD Render of Bar Spring

5.1 Wall Thickness

The wall thickness calculations using Equation 2, assuming wall thickness is twice y and with an arbitrarily chosen spring length of $0.6m$, are in the following table. The wall thicknesses for these bar springs are greater than the helical design because the torso-mounted bar spring allows for a longer flexure.

Material	Maximum Strain	Wall Thickness
Fiberglass	1.5%	0.86 <i>cm</i>
Carbon Fiber	1.32%	0.76 <i>cm</i>
Titanium Beta C	1%	0.57 <i>cm</i>

5.2 Width and Weight

As before a Carbon Fiber spring gives the best energy storage per unit weight. Note that although the dimensions of each spring have changed from the helical design, the weights have not. This makes sense because an object's energy storage capacity is directly proportional to its mass.

When the time comes to build the NG-AMASC, machining down a stock bar to the exact size recorded below would be expensive and useless. Instead, commercially available bar stock sizes will be run through the calculations backward, and the size slightly larger than necessary will be used.

Material	Width	Weight
Fiberglass	6.24 <i>cm</i>	0.58 <i>kg</i>
Carbon Fiber	1.48 <i>cm</i>	0.12 <i>kg</i>
Titanium Beta C	7.67 <i>cm</i>	1.13 <i>kg</i>

6 Experiments

To test that the theoretical spring calculations were valid, experiments were designed and performed on off-the-shelf Fiberglass bar springs. The springs tested have the required energy storage for use on the NG-AMASC but require too great a bend angle and are too wide for easy integration. This explains the differences in dimensions between the test springs ($90cm$ by $5cm$ by $0.75cm$) and theoretical springs designed for the NG-AMASC.

6.1 Testing the Spring Bending Math Model

One concern was that effects of the changing angle of attack in the system would be difficult to model. In the NG-AMASC, and thus in the experiment, the tension cable always pulls from the same direction. As the bar spring bends, the angle between the cable and bar changes, which causes the effective moment on the bar spring to change. As the bar's deflection increases, more of the tension in the cable goes toward stretching the bar and less goes toward bending the bar. An approximate mathematical model of this situation was created which, along with cantilevered beam equations, was used to create the simulation line in Figure 5.




Figure 4: Photo of First Experiment

Legend	
A	Test spring mount point
B	Test spring
C	Steel cable
D	Known coil spring
E	Linear actuator
F	Fixture Frame

In the first experiment (see Figure 4 photo), we loaded a fiberglass bar spring using a linear actuator. A coil spring of known spring rate conveyed the load from the actuator to the test fiberglass spring. At each data point, the coil spring's deflection and the fiberglass spring's position were measured. Accuracy of this measurement was estimated to be within $0.8mm$. As loadings increased and the limit of the coil spring was reached, a larger and stiffer coil spring was substituted in its place. A total of three coil springs were used, with spring rates ranging from $125kg/m$ to $900kg/m$.

As can be seen in Figure 5, the mathematical model and experimental results match well. I hypothesize that the significant error in both elevations with the large test spring is primarily due to non-linearities in the test spring and not to problems in the math model. This conclusion was reached because of the shape of the large spring data points. The initial flat range is due to the bar spring bending and the test spring not stretching, which only happens when the test spring is being loaded less than its pretension amount. The spring specification stated a pretension of $20kg$, which is significantly lower than the measured plateaus. This means the spring specification is incorrect and thus the spring's linearity is also under question. Later testing was performed with a more accurate measurement method which confirms the correctness of the simulation model.

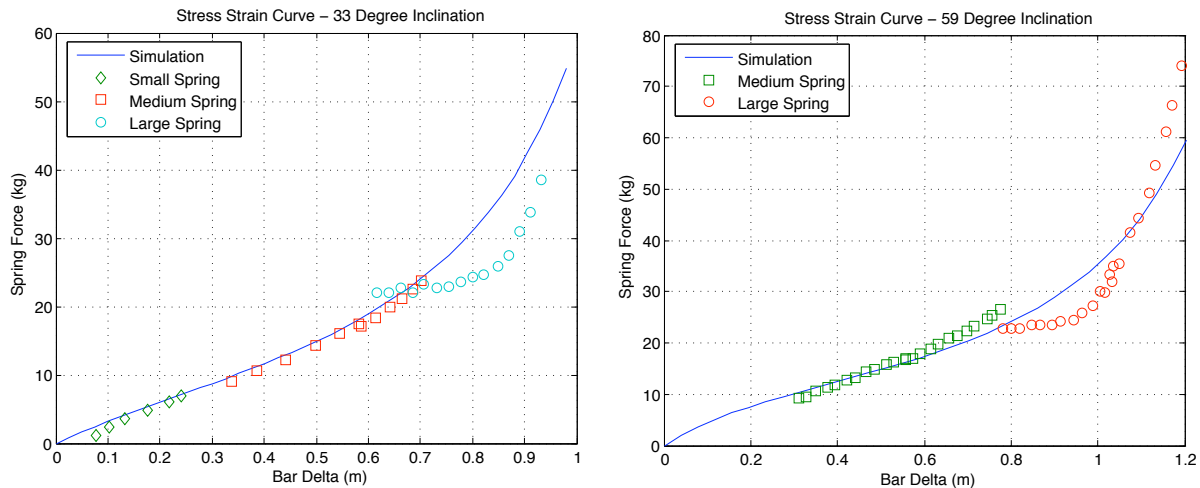


Figure 5: Experimental Results

6.2 Testing Spring Energy Capacity

To completely verify the math model of a bending spring, a second experiment was performed. The goals here were to overcome the two main limitations of the first experiment.

The first problem to correct was the presumed non-linearities of the known coil springs. To compensate for this we could have put each coil spring through a thorough deflection versus load test and created a model of the varying spring rate, or we could substitute the coil springs for a high accuracy load cell. The latter method was chosen. A properly sized load cell would not only solve the inaccuracy problem, it would be able to take measurements throughout the spring's deflection, which would remove the time-consuming process of switching between different coil springs.

The second major limitation this experiment was to overcome was non-uniform loading. During the first experiment, the test spring was supported only on one end and the rest of the spring floated in air. During light loading this did not matter; the spring deformed in a near-perfect circle because the angle between the cable and the end of the spring was nearly 90° and thus there was little pure tension in the spring. As the loading increased, and the angle between the cable and spring end increased, the moment loading of the bar decreased and the tensional loading increased. This had the effect of straightening out the cable end of the spring and putting most of the bend at the static end. Near the end of the test,

the spring was nearly in a “J” shape, not the expected and desired partial circle. The problem is not that energy cannot be stored if the spring is in this shape; the issue is the extremely tight bend radius at the static end of the spring. If the spring were mounted this way on the NG-AMASC, the stress concentrations caused by the tight bend would cause premature failure. To force the spring into uniform loading in the second experiment, a circular wood form was created for the spring to follow during loading. Figures 6, 7(a), and 7(b), of the spring under low, medium, and high loading, show it progressively wrapping more around this wood form.

The actuation method in the second experiment was also significantly different. The cable attached to the test spring passes around a passive pulley to wrap around a centrally located spindle. On the back side of the test fixture, not pictured, a 20cm lever arm extends from the spindle and rests against a load cell. To actuate the spring, an operator turns a 45cm disk on which the load cell is mounted, and takes a reading after every 60° of rotation. With a spindle diameter of 3.2cm, this means there is a data point every 1.8cm of cable retraction. All of the force being put into the system passes through the load cell, but since it is on a lever 10 times larger than the spindle’s radius, it only sees 1/10th of the force that the spring sees. This allows use of a low rated force (25kg) yet high precision load cell while testing the spring to nearly 100kg.






Figure 6: Second Experiment: Low Loading

Legend	
A	Circular wood form
B	Test spring
C	Nylon cable
D	Passive Pulley
E	Spindle



(a) Medium Loading



(b) High Loading

Figure 7: Photos of Second Experiment

The numeric results of this second experiment, shown in Figure 8, match the math model even closer than results in the first experiment. The loading in this experiment got up to about 100kg, twice the first

experiment, and even there the model matches closely. The higher loading was achieved because of the greater elevation angle of the mounted spring (96°), and the wood form constraining the spring to uniform loading. The experimental best fit line in the first graph of Figure 8 is a third order polynomial that was fitted using least squares regression. Both it and the simulation data were numerically integrated to determine spring energy storage, shown in the second graph of Figure 8.

The one issue with the experimental data is its waviness. Throughout cable retraction, the force recorded by the load cell plateaus or drops for a data point or two. This happens cyclically and matches the rotation of the spindle. The abnormality could be caused by a bearing with a bad ball, which when loaded at a certain rotation would be rougher than at other rotations. Another source could be weight imbalance of the large rear disk, caused by the block of metal to which the load cell is mounted. These two sources of error should be constant throughout cable retraction, which matches with the magnitude consistent waviness observed. Another possibility is that the waviness is the result of how the cable winds onto the spindle: at each revolution the cable has to shift over, which could cause a momentary spike in the load recorded because between those data points the actual cable retraction would be greater than the nominal retraction. But this type of error should be proportional to the magnitude of the force at that retraction, which is not the case with the waviness observed here.

Other sources of error include an observed drift in the zero-force output of the load cell and the use of a non-round cable. The zero-force output of the load cell was measured to be $35.9mV$ before the experiment and $38mV$ after. This shift of 5.5% was assumed to have happened linearly throughout the experiment, and thus adjustments were made to load cell readings to compensate for this linear shift. If the shift occurred all at once, the results reported here could be 5.5% off. Additionally, the cable used in the experiment was slightly non-circular. The cable was in fact rectangular and had a cross section of $3.63mm$ by $3.23mm$; in calculations the cable diameter was assumed to be the average of these dimensions. Since the braided nylon cable is very flexible, and thus can deform, I do not expect the non-circularity to have effected the experimental results much.

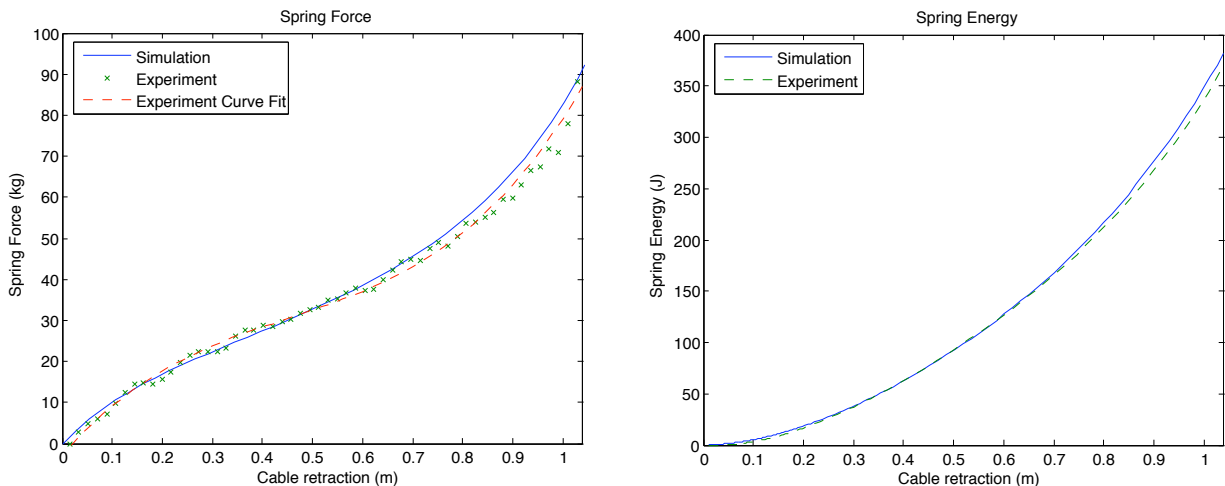


Figure 8: Experimental Results

6.3 Experiment Conclusions

These two experiments have shown that the energy storage and spring force equations developed during the course of this research are valid. The theoretical spring calculations were for a spring system with total storage of $300J$, meaning that the spring storing energy would have to handle $370J$ to offset the spring releasing $70J$. This is significant because the spring in the second test was able to store $360J$, very close to

this theoretical target. Additionally, the weights of the two Fiberglass springs were nearly the same ($0.58kg$ versus $0.6kg$). The matching of energy storage and weight, along with the fact that springs survived extreme bending in the first experiment, shows that the equations and their factors of safety are correct and proper.

The experiments have also shown that the effects of the changing angle between tension cable and spring end can be accurately modeled and thus can be compensated for when designing the NG-AMASC spiral pulleys. As stated previously, the spring force function does not have to be quadratic, or linear, or any particular mathematical function; it just has to be known so that it can be compensated for when designing the spiral pulleys.

Note that although the springs tested in both experiments were Fiberglass, the results can be used to verify math models of both Fiberglass and Carbon Fiber bar springs because of the similarities in their material properties and physical make up.

7 Spring Recommendation

The helical spring opens up some concerns about shear stress delaminating the composite materials. If a helical spring were chosen as the leg spring, titanium would be the preferred material, even though it is heavier. The possibility of delamination and premature spring failure is just too high with the composite material to risk the weight savings. Creation of custom torsional springs is an expensive and time-consuming process: initial quotes put titanium springs at \$300 plus material for each spring. Additionally, these springs would be mounted around the leg members of the robot and would add significant mass to them. Performance of the robot might be compromised.

The vertically loaded bar spring does not have the same delamination issue, so lighter composite materials can be considered. From an energy storage to weight standpoint, carbon fiber is a much better choice than fiberglass. Carbon fiber and fiberglass bar stock is easily procured and is on the order of dollars per foot. This type of spring would be mounted in the torso region of the robot, thus the spring weight would not be on the legs and would not impact motor performance as much. The only drawback to bar springs is their large size. This was a major problem on the AMASC, but the size problem is less significant here because of the torso versus leg placement.

Carbon fiber bar springs placed in the torso of the NG-AMASC provide a lightweight, inexpensive, and easy solution to its energy storage problem.