Series Elastic Actuators

Gill A. Pratt and Matthew M. Williamson
MIT Artificial Intelligence Laboratory and Laboratory for Computer Science

Abstract

It is traditional to make the interface between an actuator and its load as stiff as possible. Despite this tradition, reducing interface stiffness offers a number of advantages, including greater shock tolerance, lower reflected inertia, more accurate and stable force control, less inadvertant damage to the environment, and the capacity for energy storage. As a trade-off, reducing interface stiffness also lowers zero motion force bandwidth[9][10]. In this paper, we propose that for natural tasks, zero motion force bandwidth isn't everything, and incorporating series elasticity as a purposeful element within an actuator is a good idea. We use the term elasticity instead of compliance to indicate the presence of a passive mechanical spring in the actuator. After a discussion of the trade-offs inherent in series elastic actuators, we present a control system for their use under general force or impedance control. We conclude with test results from a revolute series-elastic actuator meant for the arms of the MIT humanoid robot Cog[5] and for a small planetary rover[2].[1]

I. Introduction

When it comes to the mechanical interface between motors and loads, "the stiffer the better" is a traditional premise of good design[19]. Several authors have previously studied methods for controlling unavoidably flexible structures (such as those expected in space[4]), and the role of interface compliance in stabilizing force control during contact transitions[23]. But with the exception of systems where energy-storage is paramount (such as the legs of a hopping robot[18]), and some passive hand mechanisms[21][11], few have suggested that elasticity should be incorporated into general purpose robotic actuators. This seems strange, particularly for robots executing natural tasks, because elasticity is used for a wide variety of purposes in animals[1].

The robotic "stiffer is better" rule of thumb first arose because increased stiffness improves the precision, stability, and bandwidth of position-control. When either open-loop positioning or colocated feedback are used, increased interface stiffness decreases end-point position error under load disturbances. In non-colocated feedback systems (where the position sensor is located at the load side of the interface), increased stiffness both lowers necessary corrections in response to load variations and raises the resonant frequency of the motor inertia and interface compliance. As a result, the bandwidth of the position control feedback loop may be raised without compromising stability[7][8].

But the use of a stiff interface is not without cost, even in position controlled systems. Most electric motors have poor torque density and thus can obtain high power density only at high speed[15]. To accelerate or support heavy loads, gear reduction become necessary. Unfortunately, gears introduce friction, backlash, torque ripple, and noise. The use of N:1 gearing also causes an N^2 increase in reflected inertia so that shock loads cause much higher forces on the gear teeth. For light weight actuators, it is more often the load rating of the gearing that limits peak torque rather than the motor, and failures of gearing due to shock are not uncommon. Finally, the increased reflected intertia and high backdrive friction of gear trains can cause damage to the environment when unexpected contact occurs.

The disadvantages of using gears are so severe that direct drive is a feasible alternative for industrial robots[3]. But for mobile robots, the power density and force density of direct drive is still too low[2].

II. Benefits of Series Elasticity

Series elasticity can give back to an actuator many of the qualities that are lost when gears are introduced. The basic configuration of a series elastic actuator is shown below:

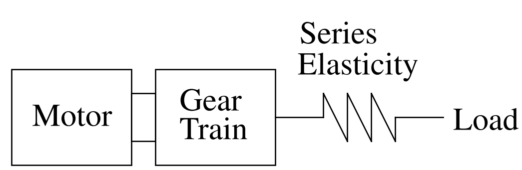

Figure 1: Block Diagram of Series-Elastic Actuator

One effect of the series elasticity is to low-pass filter shock loads, thereby greatly reducing peak gear forces. Although the same low-pass filter that spreads out a shock impulse back driving the actuator also low-pass filters the actuator's output, we believe this is a place for an engineering trade-off, not a one-sided minimization. The proper amount of interface elasticity can substantially increases shock tolerance while maintaining adequate small motion bandwidth.

1. This work was supported by JPL contract # 959333, for which we are most grateful.

Series elasticity also turns the force control problem into a position control problem, greatly improving force accuracy. In a series elastic actuator, output force is proportional to the position difference across the series elasticity multiplied by its spring constant. Because position is more easy to control accurately through a gear train than force, the force errors usually caused by gearing are reduced by lowering interface stiffness.

Increased series elasticity also makes stable force control more easy to achieve. Contrary to the case in position control, it is when the frequency of the interface resonances are **lowered** that matters are improved. The motor's force-feedback loop can operate well at low frequencies, creating an effective zero-rate spring in series with the interface. Because motor and load inertias cannot resonate with zero-rate springs, the system as a whole becomes stable.

The stability criterion translates to a minimum load inertia most easily provided by the unavoidable mass of the robot structure. Natural environments do not contain negative masses, so stability is guaranteed under contact with any object.

Finally, series elasticity provides for the possibility of energy storage. In legged locomotion, such energy storage can significantly increase efficiency[1]. By incorporating elasticity into the actuator package, efficiency benefits can be had despite the elasticity being hidden from the higher level control system. In other words, unlike methods that try to account for link elasticity at a systems level[20], the high level control system thinks it is controlling independent force actuators when in fact those actuators have internal springs that provide the aforementioned benefits.

III. Performance Limits

Series elasticity creates the need for elastic deformation whenever force is modulated. This extra motion may add either constructively or destructively to the movement of the load. In other words, depending on the relative amplitude and phase of the load's force and motion waveforms, it is possible for the interface elasticity to either increase or decrease bandwidth. If impedance control is used and the desired impedance is close to the mechanical compliance of the interface elasticity, less motor motion is generally required than in the case of a stiff interface, and bandwidth is increased.

Ignoring output inertia, a series-elastic actuator can be modeled as follows:

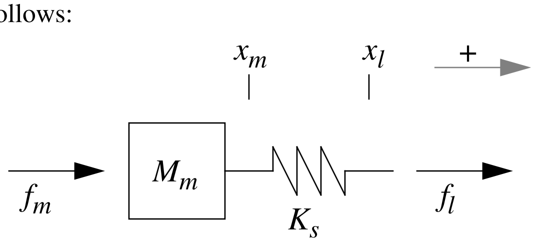

Figure 2: Model of a series-elastic actuator

with the following frequency-domain system diagram:

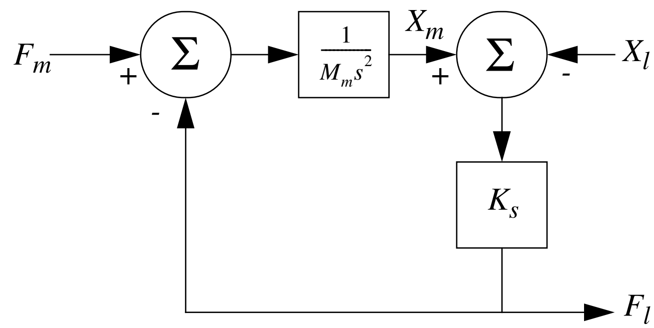

Figure 3: Frequency Domain System Diagram

and the following variable definitions:

f_m, F_m	Magnetic Force Applied to Motor Rotor
f_l, F_l	Force Applied to Load
x_m, X_m	Position of Motor
x_l, X_l	Position of Load
M_m	Motor Mass
K_s	Elasticity Spring Rate

Figure 4: System Variables

From the diagram above we can derive the following equations:

$$F_l = K_s(X_m - X_l) \tag{1}$$

$$X_m = \frac{F_m - F_l}{M_m s^2} \tag{2}$$

Setting $s = j\omega$ and solving for F_m, in terms of F_l and X_l we have:

$$F_m = F_l - \frac{M_m}{K_s}\omega^2 F_l - M_m \omega^2 X_l \tag{3}$$

As can be seen above, the motor force has three components. The first, F_l , is the force applied through the elasticity to the load. The second, $-\frac{M_m}{K_s}\omega^2 F_l$, is the force required to accelerate the motor's mass in order to change the deformation of the elasticity. The third, $-M_m\omega^2 X_l$, is the force required to accelerate the motor's mass so as to track motion of the load. Note that of these three terms, only the middle one is unique to the series elastic actuator.

Ignoring velocity saturation, we can compute performance by imposing a limit on the magnitude of F_m , i.e. $|F_m| < F_{max}$. For most motors, this translates into a bound on the maximum motor current. It is helpful to draw a vector diagram showing the magnitude and phase relationship of F_l and X_l, and the resulting F_m:

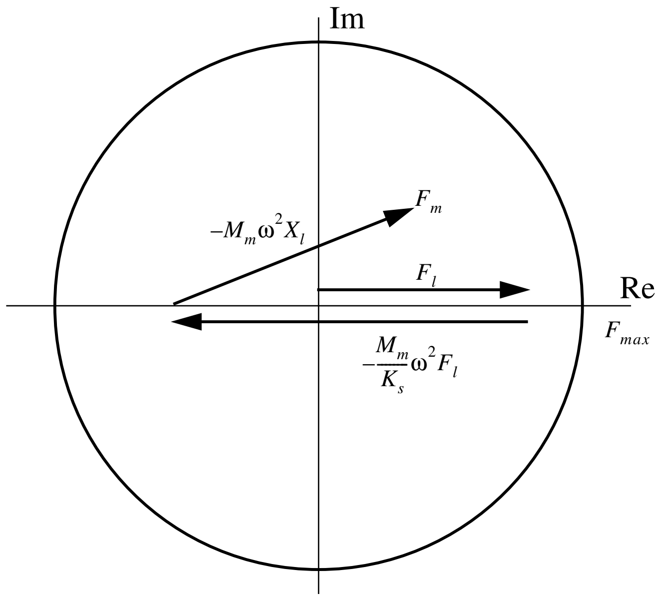

Figure 5: Phase Diagram of Necessary Motor Force

Here we have arbitrarily aligned F_l with the real axis. To satisfy $|F_m| < F_{max}$, end point F_m must land inside the circle of radius F_{max}. Note that the series elasticity term $-\frac{M_m}{K_s}\omega^2 F_l$ opposes the F_l vector. For all frequencies below $\sqrt{2\frac{K_s}{M_m}}$, the series elasticity will bring the starting point of the $-M_m\omega^2 X_l$ vector closer to the circle's center and thus allow for a greater range of possible motion amplitudes and phases than would be possible with a stiff interface. If impedance control[12] is used, the $-M_m\omega^2 X_l$ term of the vector sum will point to the right when simulating positive rate springs, and thus the inclusion of series elasticity will improve actuator performance. At frequencies less than $\sqrt{2\frac{K_s}{M_m}}$, the maximum force amplitude of damping impedances, such as are used in damping control[22], is also increased, as the vertical $-M_m\omega^2 X_l$ vector starts from a place closer to the F_{max} circle's center.

It is informative to view the limits on actuator performance as a function of the mechanical impedance $Z = \frac{F_l}{X_l}$ seen by the load. The equation below gives the load force as a function of desired impedance and motor force.

$$F_l = \frac{F_m Z}{\left(1 - \frac{M_m}{K_s}\omega^2\right)Z - M_m\omega^2} \qquad (4)$$

This function has a zero at $Z = 0$ and, if $\omega \neq \sqrt{\frac{K_s}{M_m}}$, it has a pole at $Z = \frac{K_s M_m \omega^2}{K_s - M_m\omega^2}$, which is the natural impedance seen by the load with the motor unpowered. The magnitude of this function is plotted below using parameter values from the experimental setup discussed later. Plots are shown for frequencies below, at, and above resonance.

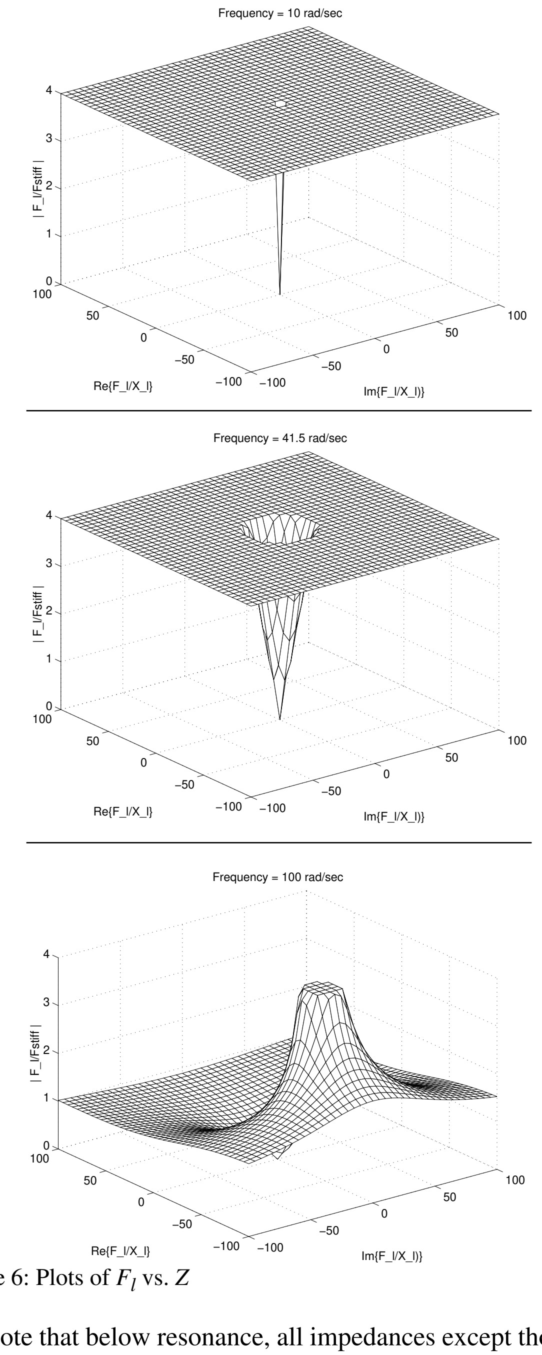

Figure 6: Plots of F_l vs. Z

Note that below resonance, all impedances except those near $Z = 0$ may be generated at the full motor force. At low frequencies it is particularly easy to generate impedances that simulate the motor's mass M_m. This is not apparent above because F_l has been artificially limited (as might occur due to the spring's yield point or maximum gearbox torque, for example). Generating significant force near $Z = 0$ entails large motions and is thus difficult regardless of the elasticity of the interface. This difficulty increases with frequency.

At resonance, the pole disappears and the ability to generate force is a simple linear function of the complex distance from $Z = 0$.

Above resonance, the pole re-appears, moving asymtoptically towards $Z = -K_s$ as frequency increases. Thus, at high frequencies it is particularly easy to generate impedances close to that of the series elasticity.

To compare the performance of the series elastic actuator with a stiff actuator, one can take the ratio of possible output force as a function of output impedance and frequency. This is:

$$\frac{F_{l_{elastic}}}{F_{l_{stiff}}} = \frac{Z - M_m\omega^2}{\left(1 - \frac{M_m}{K_s}\omega^2\right)Z - M_m\omega^2} \tag{5}$$

The magnitude of this function is plotted below using parameters from the experimental setup compared to a system of identical parameters, but with a stiff interface:

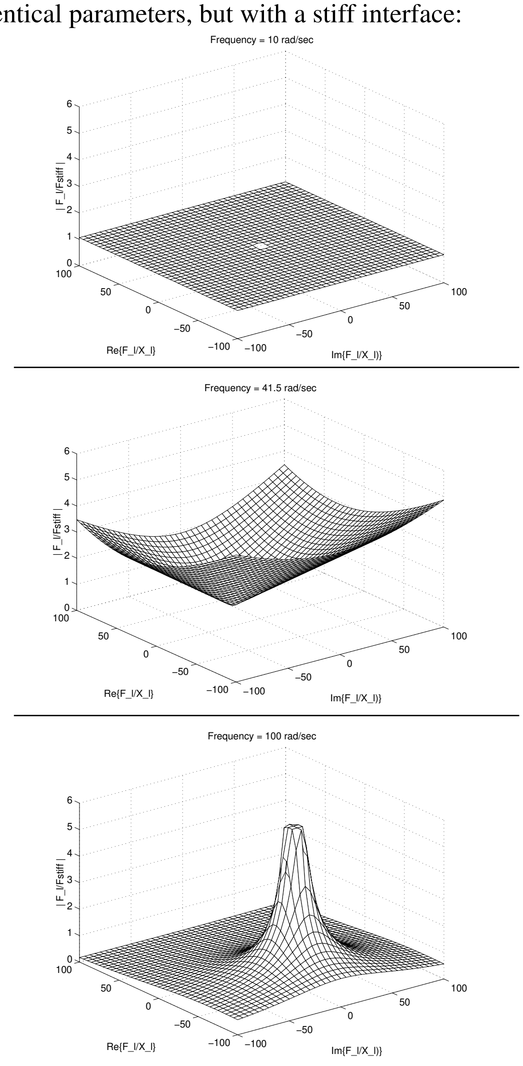

Figure 7: Plot of F_l / F_{stiff} vs. Z

These curves are similar to the previous ones except that the zero has moved from $Z = 0$ to $Z = M_m\omega^2$, indicating that the series-elastic actuator cannot match the unlimited performance of the stiff interface actuator at generating impedances close to that of the motor's mass (because the stiff interface design needs no motor power at all). The series-elastic actuator can generate such impedances, but they require some motor power.

IV. Control

Stable, accurate, force control can be obtained by using the architecture shown below, where F_d is the desired force:

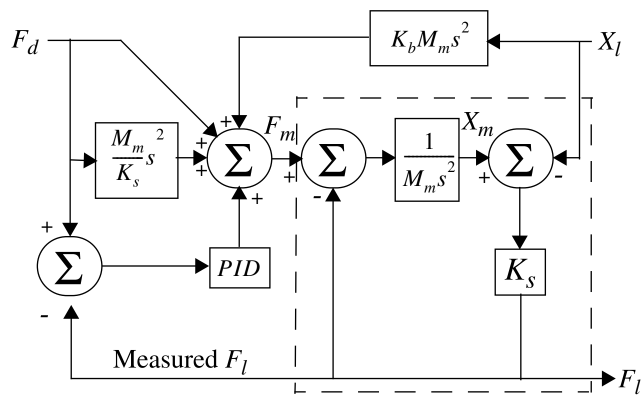

Figure 8: Control Architecture

As shown above, the control architecture contains both feed-forward and feed-back paths. The feed-forward terms attempt to fully compensate for all three terms of equation 3, with the exception of the last (load motion) term, where a back drive gain K_b is set to a little less than 1 so as to prevent feedback inversion and instability.

Feedback to compensate for modeling errors and $K_b < 1$ is accomplished by an ordinary PID loop, operating on force error. This loop has a transfer function of:

$$PID(s) = K_p + \frac{K_d s}{1 + \tau_d s} + \frac{K_i}{1/\tau_i + s} \tag{6}$$

Where the parameters are defined as follows:

K_p	Proportional Gain
K_i	Integral Gain
τ_i	Integral Roll-Off
K_d	Derivative Gain
τ_d	Derivative Roll-off

Figure 9: Feedback Parameters

Stability can be analyzed by looking at the output impedance as a function of frequency $s = j\omega$ with a commanded force $F_d = 0$:

$$Z = \frac{F_l}{X_l} = \frac{K_s(1-K_b)M_m\omega^2}{K_s(PID(j\omega)+1) - M_m s^2} \quad (7)$$

If the imaginary part of this impedance is less than or equal to zero, than the actuator as whole will be passive and thus stable when interacting with any passive load[13][14]. The only imaginary component of the impedance comes from the PID term, which is in the denominator. Thus, for the impedance to have a negative imaginary part, the PID term must have a positive imaginary part, i.e.:

$$\mathrm{imag}\left(K_p + \frac{K_d j\omega}{1+\tau_d j\omega} + \frac{K_i}{1/\tau_i + jw}\right) \geq 0 \quad (8)$$

which is guaranteed for all ω when

$$\tau_i \leq \sqrt{\frac{K_d}{K_i}} \quad (9)$$

i.e., when the integral gain is rolled off below a sufficiently high frequency. A theoretical bode plot of output impedance with zero command force (for the experimental setup's parameters and without considering saturation) is shown below:

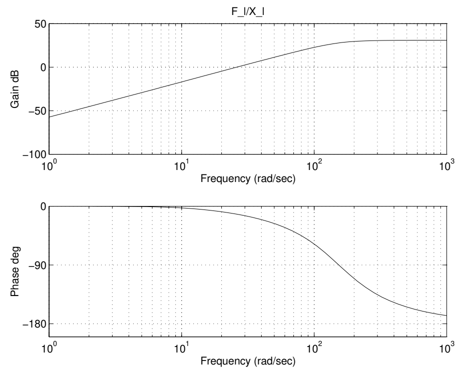

Figure 10: Theoretical Bode Plot of Output Impedance

In a real system with motor saturation, the actuator will take on the natural impedance of the series elasticity at sufficiently high frequencies[10]. Thus, a light load mass may resonate with the series elasticity. To avoid this problem, placing a minimum mass on the load will lower the resonant frequency to where the control loop operates well. At this low frequency, the impedance of the series elasticity disappears from the overall impedance (which is very low), and resonance cannot occur.

Because of saturation limit-cycle concerns or sampling rate limitations (if the control system is implemented digitally), the bandwidth of the PID system may be limited by design. However, at mid-range frequencies, the feed-forward components of the control system still operate unfettered, and performance is limited only by the capabilities of the motor.

V. Experimental Setup

A series elastic actuator, shown in the photograph below, was constructed to evaluate performance:

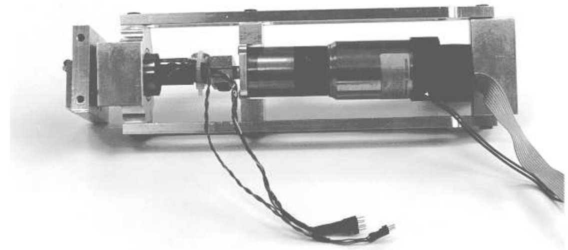

Figure 11: Photograph of Experimental Series-Elastic Actuator

The motor used was a MicroMo 3557K (48V, 25W) with a HEDS5010 encoder. A 66:1 reduction planetary gearbox was used, the output shaft of which was attached to a steel torsion spring, which formed the series elasticity. The actuator output was taken from the other end of the spring. The spring was of a cross-shaped cross-section, which was found to give the best stiffness v. strength characteristics. The inertia of the motor at the output of the gearbox was calculated to be 0.02 kgm^2 and the stiffness of the spring was 34 Nm/rad, making the natural frequency of the system 41 rad/s or about 7Hz.

As shown below, the twist in the spring was measured using strain gauges mounted on the flats of the spring:

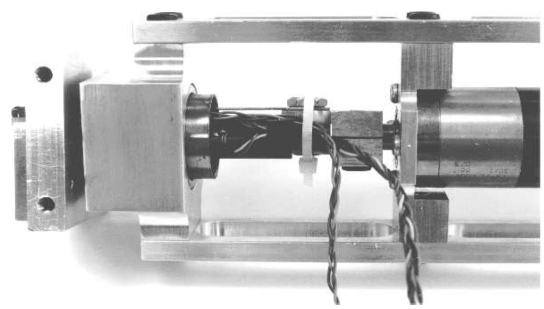

Figure 12: Mounting of Strain Gauges on Series Spring

The strain gauges were connected to an amplifier (AD 1B31) and then to the control computer.

The control loop used on the actuator was similar to that shown in fig. 8, only the $\frac{M_m s^2}{K_s}$ term was not implemented. The control parameters were set as follows:

K_p	12.41
K_i	12.41
τ_i	0.08
K_d	0.124
τ_d	.0015

The derivative of the strain gauge reading was obtained by sampling the output of an analog differentiating circuit, with a roll off of about 100 Hz. The output of the control loop was a desired motor current, but due to hardware constraints in the simple motor amplifier, a feedforward motor model was used to calculate an appropriate PWM duty cycle. The motor model had two components, one for the resistance of the motor windings, and a second corresponding to back emf, using a calculation of motor velocity from the encoder. The sampling rate for the control loop was 1 kHz.

The control loop ran on a Motorola 6811 microcontroller, mounted on a custom board with circuitry for encoders, power amplifiers for the motors, and communications with a Motorola 68332 processor. The 68332 was used to send gains to the 6811, and receive data on the performance of the system. The 68332 runs L[6], a version of Lisp written by Rodney Brooks at MIT.

For the test results included in this paper it was desired to change the angle of the actuator output shaft in a controlled manner. This was implemented by placing a position controlled actuator with a stiff interface so that its output shaft was connected to the series elastic actuator's output. The complete test setup is shown in the photograph below:

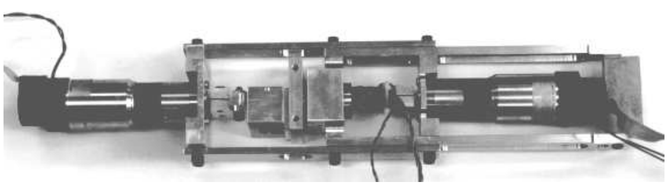

Figure 13: Dual actuator test rig

VI. Results

The first test consisted of commanding both the output force and the output position to move in sine waves of the same frequency, and varying the magnitudes and the phase difference between the two signals, thus setting the impedance of the actuator output. The performance could then be measured by calculating the root mean square force error and normalizing with respect to the commanded force amplitude. In order to reduce the effects of backlash, the force was given a bias, ensuring that the desired force was always positive. There was also a limit on the lowest impedance that could be generated in the system, due to motion limits on the position controlled motor.

The following figures show the results below, at, and above, resonance:

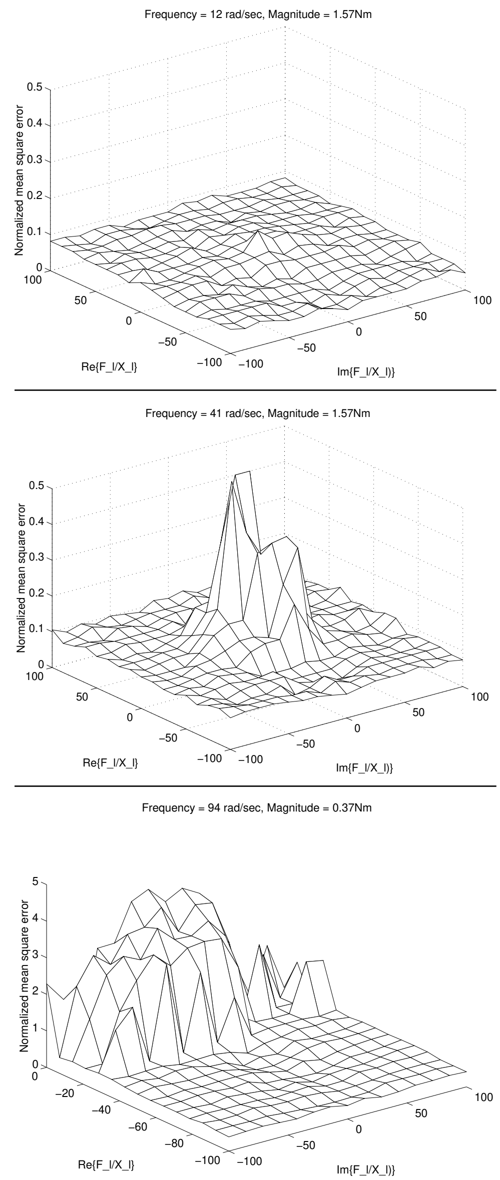

Figure 14: Experimental Results

At low frequencies, the error is pretty constant and low, as suggested by the plots in figure 6. The small spike corresponds to the lowest impedance that could be generated on the test rig: 8.1 Nm/rad. At resonance there are more errors near the zero impedance point, while at larger impedances the error is still small. Above resonance, the magnitudes of the force and the position profiles were reduced because neither motor could provide large movements or torques at this frequency. It can be clearly seen that at high frequency the actuator only performs well when its output impedance has a negative real part, which corresponds to spring-like behavior or damping. The minimum point on the surface is at about -40 Nm/rad, which roughly corresponds to the impedance of the spring (34 Nm/rad). The reason the errors are so big for impedances near zero is that both motors are saturating.

VII. Conclusions

In this paper we have shown that for natural tasks, series elastic actuators can provide many benefits when compared to traditional stiff actuators. These benefits include shock tolerance, lower reflected inertia, more accurate and stable force control, less damage to the environment, and energy storage. Although zero motion force bandwidth is reduced as series elasticity is increased, the force bandwidth for many tasks that involve load motion is improved. This is particularly true for natural tasks that are spring- or damper-like in their output impedance[22].

We have also shown that a simple control system can generate a range of output impedances - not just that of the passive series elasticity, and have demonstrated experimentally that accurate, stable control is possible over a range of generated impedances. Since force is the direct variable of control, energy field methods[16] may also be used with series-elastic actuators.

Several avenues are open for future work:

In cases where the reduction of zero motion bandwidth cannot be tolerated, one can connect the outputs of two differently tuned series-elastic actuators in parallel. The series elasticity in each actuator isolates its motor's inertia from the output, thus making it possible to connect the output of a high frequency actuator (perhaps with lower ratio gearing and a stiffer series elasticity) to the output of a low frequency actuator (with higher ratio gearing and a less stiff series elasticity). In much the same way a hi-fi speaker uses a tweeter and woofer to cover a larger frequency range, two series-elastic actuators can be configured to cover a wider bandwidth. Because frictional elements may differ in the two motors as well, a parallel connection can also improve dynamic range. This approach is currently being investigated at MIT by John Morrell and Ken Salisbury[17].

Another interesting idea is to use variable-rate springs where modulation of the bias point can effect changes in passive stiffness. This type of mechanism has been studied before[21] and a more sophisticated version is currently being investigated at MIT by Ken Salisbury.

In many cases (e.g. tendon manipulators and pneumatic actuators), series elasticity is unavoidable. It is our belief that when such an actuator is to be used in natural tasks, this property should be looked on as a blessing rather than a curse. Mike Binnard at MIT is presently investigating series-elastic control in a small pneumatic robot.

The revolute actuators discussed in this paper have been designed for use in the arms of the MIT humanoid robot COG. We plan to develop minitiarized versions for use in a planetary rover. In yet another project, one of the authors (Pratt) and a number of graduate students are completing a bipedal walking robot with series-elastic tendon actuators.

We believe that series elasticity will endow all of these robots with better mechanical characteristics, particularly for interactions with the natural environment.

VIII. References

[1] Alexander, R. McNeill, "Elastic Mechanisms in Animal Movement", Cambridge University Press, 1988.

[2] Angle, C.M. and Brooks, R.A., "Small Planetary Rovers", IEEE International Workshop on Intelligent Robots and Systems, Tsuchiura, Japan, July 1990 pp. 383-388.

[3] Asada, H. and Kanade, T. "Design of Direct-Drive Mechanical Arms", ASME J. of Vibration, Acoustics, Stress, and Reliability in Design 105(3), pp. 312-316.

[4] Balas, Mark J., "Active Control of Flexible Systems", Proc. of the 1977 Symposium on Dynamics and Control of Large Flexible Spacecraft, June 13-15, 1977, Blacksburg, VA pp 217-236.

[5] Brooks, R.A. and Stein, L.A., "Building Brains for Bodies", to appear in Autonomous Robots, (1:1), 1994.

[6] Brooks, R.A., "The L Manual", MIT AI Lab Internal, 1994.

[7] Cannon, Robert H. Jr., and Rosenthal, Dan E., "Experiments in the Control of Flexible Structures with Noncolocated Sensors and Actuators", J. Guidance, Vol. 7 No. 5 pp. 546-553, Sept.-Oct. 1984.

[8] Cannon, Robert H. Jr., and Schmitz, E., "Initial Experiments on the End-Point Control of a Flexible One-Link Robot", Int. J. of Robotics Research, Vol. 3 No. 3 (1984) pp 62-75.

[9] Eppinger, Steven D., and Seering, Warren P., "Three Dynamic Problems in Robot Force Control", IEEE Intl. Conf. on Robotics and Automation, 1989, pp 392 - 397.

[10] Eppinger, Steven D., and Seering, Warren P., "Understanding Bandwidth Limitation in Robot Force Control", IEEE Int. Conf. on Robotics and Automation, April 1987.

[11] Hashimoto, Minoru, and Imamura, Yuichi, "An Instrumented Compliant Wrist Using a Parallel Mechanism", Japan/USA Symp. on Flexible Automation, V. 1 ASME 1992, pp 741-744.

[12] Hogan, N., "Impedance Control: An Approach to Manipulation: Part I - Theory, Part II - Implementation, Part III - Applications", J. of Dyn. Syst., Measurement Contr., 107:1-24 (1985)

[13] Hogan, Neville, "On the Stability of Manipulators Performing Contact Tasks", IEEE J. of Robotics and Automation, V4 N6, Dec. 1988, pp 677-686.

[14] Hogan, Neville, and Colgate, Ed, "An Analysis of Contact Instability in term of Passive Physical Equivalents", IEEE Intl. Conf. on Robotics and Automation, 1989, pp 404-409.

[15] Hunter, Ian W., Hollerbach, John M., and Ballantyne, John, "A comparative analysis of actuator technologies for robotics", Robotics Review 2, MIT Press, 1991.

[16] Khatib, O., "Real Time Obstacle Avoidance for Manipulators and Mobile Robots", Int. J. of Robotics Res. V5 N1 (1986).

[17] Morrell, John B., and Salisbury, J. K., "Parallel Coupled Actuators for High Performance Force Control: A Micro-Macro Concept", Submitted to IROS 95.

[18] Raibert, M. H., “Legged Robots That Balance.” Cambridge, Mass.: MIT Press (1986).

[19] Salisbury, K., Eberman, B., Levin, M., and Townsend, W., “The Design and Control of an Experimental Whole-Arm Manipulator”, Proc. 5th Int. Symp. on Robotics Research. (1989)

[20] Spong, M. W., “Modeling and Control of Elastic Joint Robots”, Trans. of the ASME, Vol 109, Dec. 1987, pp 310-319.

[21] Sugano, S., Tsuto, S., and Kato, I., “Force Control of the Robot Finger Joint equipped with Mechanical Compliance Adjuster”, Proc. of the 1992 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, 1992, pp 2005-2013.

[22] Whitney, Daniel E., “Force Feedback Control of Manipulator Fine Motions”, J. Dyn. Syst. Measurement Contr. 98:91-97 (1977)

[23] Whitney, Daniel E., “Historical Perspective and State of the Art in Robot Force Control”, Int. J. of Robotics Research, V6, N1, Spring 1987 pp 3-14.