Series Elasticity for Free Free-Space Motion for Free

R. Brent Gillespie* University of Michigan
Dongwon Kim† University of Michigan
Jacob M. Suchoski‡ Stanford University
Bo Yu§ University of Michigan
Jeremy D. Brown¶ University of Michigan

ABSTRACT

Series elastic actuators are used to significant advantage in many robot designs but have not found their way into the design of haptic devices. We use a pneumatic circuit to realize both a flexible power transmission as well as the elastic element in a series elastic actuator. The pneumatic circuit effectively hides the impedance of a high friction, high mass ball-screw actuator, while a low friction, low mass pneumatic cylinder is used at the end-effector. We offer a comparative study of an impedance-control device, admittance-control device, and a device incorporating a series elastic actuator, investigating both the open-loop and the closed-loop impedance displayed to the user. While all hardware and control designs offer an ability to shape the impedance within their operational bandwidths, the series elastic design has the particular advantage of low impedance (a very compliant spring) outside of that bandwidth. Thus, a haptic device featuring series elastic actuation is capable of providing both the low impedances required in free-space and the high impedance required for rendering stiff virtual walls.

Keywords: Series Elastic Actuation, Haptic Device Design

1 INTRODUCTION

The control and hardware design of any haptic device must address two chief challenges: a) render the virtual wall as stiff and crisp as possible and b) render free-space motion as free and light as possible. Whether these two challenges are met satisfactorily can generally be judged with a few quick taps on a virtual object and a wiggle or two in free space. Naturally, interaction with the stiffest possible virtual wall must behave according to expectations: without bounce or chatter if the user imagines, say, contacting the virtual object with a fingertip. Speaking loosely, we describe such a wall as "stable". A significant body of stability analysis and many advances in control and hardware design have led to improvements in the stiffest possible stable virtual walls. The rendering of a particularly "free" free-space motion, on the other hand, has not received the same kind of attention. Generally a certain apparent mass must be tolerated in the feel of any virtual environment rendered through a haptic device, and this apparent mass dominates the response at high frequencies. In particular, the magnitude of the force response to any rapid motion increases steadily without bound as the frequency of that imposed motion increases. This is true whether the device is a so-called *impedance-control* device or an *admittance-control* device.

The goal of a maximally "free" motion in free-space can be quantified as the lowest possible impedance (low to nil force produced in response to a motion imposed by the user), and an apparent mass is undeniably the worst of the canonical mass/damper/spring elements to dominate the impedance at high frequencies. The slope of the mass on a Bode plot or frequency response plot of impedance is the steepest.

Roughly speaking, an impedance-control device has a low impedance when its motors are unpowered, and it uses motor power to increase that impedance when appropriate to render the feel of virtual objects. An admittance-control device is in a sense the dual: it has a high impedance when its motors are unpowered and it energizes its motors to decrease that impedance when appropriate for haptic rendering. Let us address the options for reducing impedance in free-space available to the designer of an impedance control device and an admittance control device separately.

Achieving good free-space behavior using an impedance control device requires astute hardware design to minimize the device inherent impedance. In traditional impedance control devices, there is in fact no feedback loop closed to cancel the device dynamics, other than the loop containing the virtual environment [1,2]. Astute hardware design then involves the selection of a low inertia DC motor and a smooth backlash-free transmission of moderate mechanical advantage. Likewise a handle or end-effector with minimal mass is selected for the user to grasp. The transmission between motor and handle is typically very stiff so as to not introduce dynamics into the control loop or simulation dynamics. In order that the coupled dynamics are those of the user interacting directly with the virtual environment, the dynamics of the haptic device should "disappear" or become "transparent". Compensation for the device dynamics includes model feedforward and force feedback [1] but these are specialty control designs that have not gained significant hold in haptics.

Admittance control devices, on the other hand, do feature a local feedback loop around the device: the device is position-controlled to move so as to *accommodate* the user's imposed force. The imposed force is sensed with a force-sensor placed as close as possible to the point of contact with the user. This force then acts in simulation on the virtual environment in forward-dynamics form to produce the desired position. Because smooth and backlash-free transmissions are still required [1,3], the mass/inertia of the moving parts of the device are generally quite high. Large mechanical advantages can be used, thus achieving a large workspace and high force capacity [3]. Motor power must of course first be used to move the mass/inertia of the device, but with large mechanical advantage transmissions, the required force capacity is readily available. The speed capacity of the motors, however, is also limited and this speed saturation limits the performance at high frequencies.

In this paper, we introduce a new hardware and accompanying control design architecture to the practice of haptic device design—the use of series elastic actuators. The architecture is not new to the field of robotics by any means. It has taken hold in several pockets within the field of robotics, especially in the design of robots intended to interact with humans.

Series elastic actuators (SEAs) have their roots in the desire to achieve smooth force control and minimize impact forces in robots interacting with their environments. Pratt and Williamson [4] noted that "stiffness isn't everything" and purposefully introduced compliance between an actuator and robot end-effector. They coined the construction a "Series Elastic Actuator". The series elastic element provides a sensitive measure of the force of interaction and also significantly reduces the backdrive impedance of the device. The actuator and transmission (including their equivalent mass and non-

*e-mail: brentg@umich.edu
†e-mail:gloryn@umich.edu
‡e-mail:jsucho@stanford.edu
§e-mail:ybo@umich.edu
¶e-mail:jdelaine@umich.edu

IEEE Haptics Symposium 2014
23-26 February, Houston, Tx, USA
978-1-4799-3131-6/14/$31.00 ©2014 IEEE

smooth dynamics) are "hidden" behind the elastic element when viewed from the environment. One advantage of the series elastic actuator is that with a reduced apparent mass and intervening spring, impact forces are attenuated. This provides a safety feature to persons interacting with the robots and by the same taken the robot actuators and transmissions are protected agaist over load damage. The series elastic actuator realizes a lower driving point impedance, in particular at high frequencies. The device appears as a spring coupled to a wall at high frequencies. Baxter [5] is based on series elasticity, and the design by Zinn [6] is intended for interaction with humans, where impacts must be minimized in industrial environments. Series elastic actuators have also been used to advantage in the design of rehabilitation robotics [7–10] and to some extent prosthetic devices [11].

Paluska and Herr [12] show that series elasticity can be used to increase the power delivered to a mass from an actuator working over a restricted stroke. Shen and Goldfarb [13] show that the passivity of a high-stiffness surface rendered on an impedance-type haptic device can be enhanced with the use of a pneumatic actuator in place of a DC motor. Series elasticity has also been used to achieve actuation with programmable stiffness [14].

In this paper, we highlight the particular advantage of series elasticity for haptic device design. It is not strictly because the series elastic element is a kind of cheap force sensor, nor that it softens impacts. We argue that the utility of series elasticity lies chiefly in its behavior at high frequencies and the associated implications for rendering free-space. The high-frequency behavior of a series elastic element (the inherent dynamics or unpowered impedance) is that of a spring. Its frequency response is flat: the response force to an imposed motion does not increase in amplitude with increasing frequency of the imposed motion. Thus when a user applies the quick and dirty test of wiggling the device end-effector in free-space, the response is simply a stiffness. If the series elastic element is particularly compliant, this stiffness is correspondingly small. Rendering a particularly "free" free-space impedance comes for free!

We employ series elasticity to realize performance gains to meet our particular application requirements. We have built a wearable robot (exoskeleton) that realizes a low weight and high force/torque capacity in addition to a low backdrive impedance. We use a pneumatic circuit to realize not only the elastic element but a flexible transmission and tether as well. This allows us to place a driving cylinder on fixed ground while the tethered driven cylinder is attached to the exoskeleton. The tether is highly flexible, thus accommodating off-axis motion of the user. While the exoskeleton is maximally "transparent" in terms of impedance, it is also transparent in terms of weight and off-axis inertia. The use of pneumatics as a series elastic element appears to be novel in haptics.

In this paper we offer a comparative analysis of impedance control and admittance control devices alongside an analysis of haptic devices featuring series elasticity. We attempt to place all architectures on a level playing field and we analyze not only the inherent (unpowered) impedance, but also the powered impedance relative to the target impedance of the virtual environment to be rendered. We conclude with an empirical evaluation of our device that includes the backdrive impedance and the closed-loop impedance of rendered free-space. The torque output capacity of our device in the frequency domain (operational bandwidth) is also presented.

2 Model Development

2.1 Impedance Control Devices

Many off-the-shelf haptic devices such as the PHANToM from SenseAble Devices [15] and many laboratory devices employ impedance control, essentially as introduced by Hogan in 1985 [16]. Perhaps more descriptively, the most commonly implemented impedance control architecture would be called *open-loop* impedance control, to distinguish it from those that have appeared

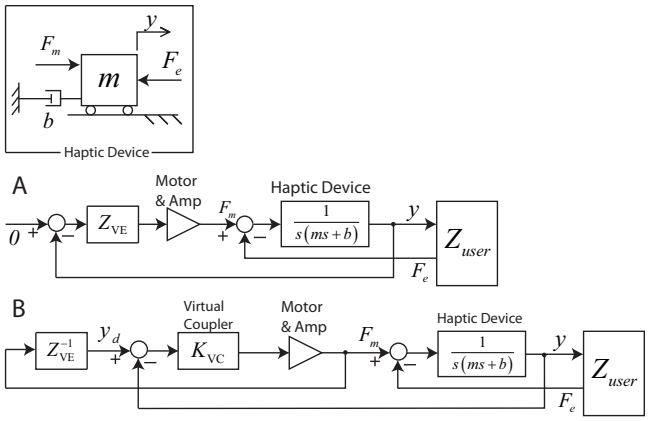

Figure 1: Model of an impedance control haptic device, developed both A) without and B) with a virtual coupler.

featuring feedback or feedforward control loops [1, 2]. Figure 1A shows a block diagram featuring the coupled virtual environment, haptic device, and human user. Note that all elements to the right of the motor/amp are physical while the elements to the left are implemented on a digital computer or embedded microprocessor. The figure inset shows the haptic device schematically, indicating an end-effector of effective mass m to which the user applies a force F_e. The motor applies a force F_m directly to this mass, and m includes the inertia of the motor rotor reflected through the transmission. Transmission dynamics are neglected, as impedance device transmissions are generally designed to be stiff and backlash-free. The model is a simple single-axis model with no distinction between jointspace and taskspace. Naturally the appropriate Jacobian transformations could be used to generalize this model to multi-DOF designs.

The designation of the haptic device as a lumped mass m with viscous damping b to ground with the applied sum of forces $F_m - F_e$ makes it apparent that the loop closed by F_e is not actually a feedback control loop. Rather, it is in the physical domain. The only loop closed around the haptic device is that through the virtual environment dynamics Z_{VE}.

In the following, we submit each model to a similar set of analyses to produce two transfer functions in particular: 1) the open-loop impedance Z_{OL} to describe what the user feels when the device is unpowered ($F_m = 0$), and 2) the closed-loop impedance Z_{CL} to describe what is rendered to the user when the device is powered. The closed-loop impedance Z_{CL} can be compared to the virtual environment impedance Z_{VE}, the target impedance to be rendered. With Z_{VE} and Z_{CL} in hand, transparency [17] or distortion [18] can easily be evaluated. To simplify the nomenclature, we define impedance Z and admittance Z^{-1} as relationships between force and displacement father than force and velocity.

Reading out the block diagram in Figure 1, we have:

$$Y = \frac{1}{s(ms+b)}(F_m - F_e), \tag{1}$$

and

$$F_m = -Z_{\text{VE}}Y. \tag{2}$$

The open-loop or unpowered impedance Z_{OL} follows immediately from Eq. (1) by setting $F_m = 0$

$$Z_{\text{OL}} \triangleq \left.\frac{F_e}{Y}\right|_{F_m=0} = -s(ms+b). \tag{3}$$

and the closed-loop impedance Z_{CL} is found by substituting Eq. (2) into (1).

$$Z_{\text{CL}} \triangleq \frac{F_e}{Y} = -\left(Z_h + Z_{\text{VE}}\right), \tag{4}$$

where $Z_h = s(ms + b)$ is the inherent impedance of the haptic device.

We see from this simple analysis that the rendered impedance Z_{CL} has the inherent or open-loop haptic device impedance Z_h added to or 'riding on top of' the virtual environment impedance Z_{VE}. Note further that a pole-zero map of the haptic device impedance Z_h contains a zero at the origin and a real zero in the left-half plane. If the damping b is small (which is usually the case by hardware design) then the two zeros are essentially both at the origin, which describes a double integrator or mass. With increasing frequency of the motion $y(t)$ imposed by the user, the magnitude of the force response increases at a rate of 40 dB/decade on a Bode plot. At high frequencies, the effective mass of the haptic device dominates, as any attempt to compensate for these effects in Z_{VE} will be filtered by the device admittance Z_h^{-1}. Since there is no inner feedback loop in traditional impedance control for haptic devices, the control bandwidth is determined by Z_h^{-1}.

Impedance control is often augmented with a model feedforward term in the control law to cancel the haptic device dynamics Z_h. This approach was part of impedance control originally proposed by Hogan [16]. The use of a force sensor enables a force feedback loop to be closed, which also provides a means to suppress the dynamics Z_h in the rendered dynamics [1]. But both of these methods are rather uncommon in impedance control haptic devices. Instead, careful attention is paid to the design of haptic device hardware to minimize the effective mass m and damping b.

An important variation on the block diagram in Figure 1A is shown in Figure 1B. A so-called *virtual coupler* is used to enable the virtual environment to be simulated in its forward-dynamics form. The forward dynamics correspond to an admittance Z_{VE}^{-1}. The virtual coupler has a physical equivalent of either a spring or parallel arrangement of a spring and damper and simply enables the "causality" of the coupled virtual environment to be reversed (whether the virtual environment has a motion or force imposed on it).

2.2 Admittance Control Devices

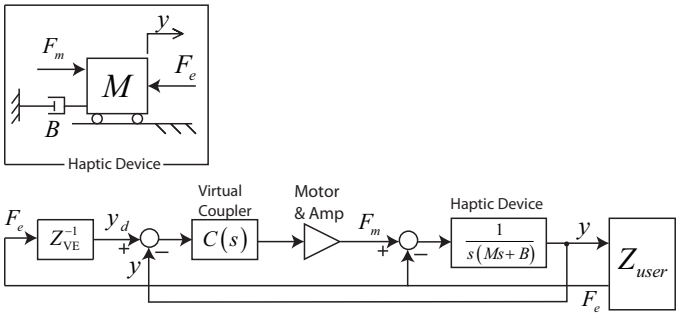

Figure 2: Admittance Control Model

Figure 2 shows schematic and block diagrams highlighting the typical controller used for an admittance-control haptic device. Again the device is modeled as a mass with damping to ground, but since the motor is given a large mechanical advantage over the end-effector, and the end-effector itself has large inertia, the equivalent mass is denoted with a capital M, where M is much greater than m introduced above for the impedance control device. A force sensor is used to sense the force of interaction with the user F_e and this force is fed into the forward dynamics model of the virtual environment Z_{VE}^{-1}. A high gain position controller $C(s)$ is used to drive the haptic device position y to follow the output y_d of Z_{VE}^{-1}. Thus the haptic device is driven to accommodate the user applied force F_e as would the virtual environment Z_{VE}^{-1}. The verb *accommodate*, which is what the controller $C(s)$ attempts to enforce, does describe the rendering of an admittance. And consistent with the rendering of an admittance, the user Z_{user} is modeled as an impedance (responding to imposed motion y with force F_e). But note that even for the impedance-control device in Figure 1, the user was modeled as an impedance and the device was rendering an admittance. The name *impedance control* refers nominally to the causality of the signal-domain portion of the block diagram, behind the motor. In impedance control, the haptic device is in effect a coupler that inverts the impedance Z_{VE} and presents an admittance Z_{CL}^{-1} to the user. While no loop is closed around the haptic device Z_h^{-1} to suppress its dynamics in the nominal form of impedance control, the picture is different in admittance control: device dynamics are suppressed in admittance control by the operation of a feedback loop.

To determine the open-loop impedance Z_{OL} for the admittance controlled device, we begin by reading out the following two expressions from the block diagram.

$$Y = \frac{1}{s(Ms + B)}(F_m - F_e), \tag{5}$$

and

$$F_m = C(s)(Z_{\text{VE}}^{-1} F_e - Y). \tag{6}$$

The first of these already yields the open-loop impedance for the admittance control device.

$$Z_{\text{OL}} \triangleq \left.\frac{F_e}{Y}\right|_{F_m=0} = -s(Ms + B), \tag{7}$$

which is the same as for the impedance control device (see Eq. (3)), except for the larger mass M. A substitution from Eq. (6) into Eq. (5) yields the closed loop impedance:

$$Z_{\text{CL}} \triangleq \frac{F_e}{Y} = \frac{Z_h + C(s)}{C(s)Z_{\text{VE}}^{-1} - 1}. \tag{8}$$

In this expression we see that as the gains in the position controller $C(s)$ are tuned up, the impedance Z_{CL} approaches the desired impedance Z_{VE}.

2.3 Series Elastic Actuator Devices—Force Control

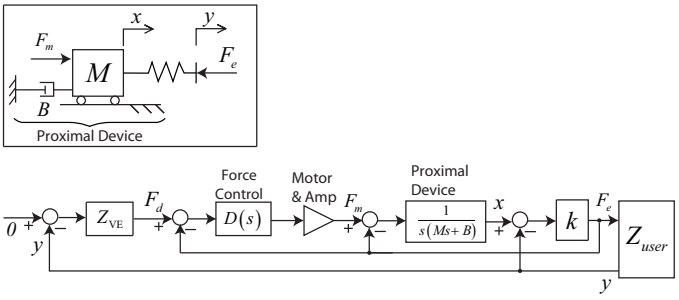

Figure 3: Series Elastic Actuator Control Model—Force Control

We now introduce the first of two control architectures for a haptic device that features a series elastic element. We call the device a *SEA Device*, for Series Elastic Actuator, in accordance with the naming by Williamson and Pratt [19], who introduced series

elasticity in robotics. This first control architecture we call force-controlled.

The inset in Figure 3 shows a spring of stiffness k intervening between a mass/damper configured just as in the admittance control architecture discussed above. This spring is the series elastic element: the product of its stiffness k and its compression $x - y$ expresses the force F_e experienced by the user:

$$F_e = k(X - Y). \tag{9}$$

This force is likewise imposed on what we have named the *proximal* portion of the device, the mass M connected to ground through the viscous damper B. Note that we have modeled the distal part of the spring, including the end-effector that the user grasps as having no mass. This is a simplifying assumption that could be re-visited in a more extended model and analysis. We assume for the moment that this mass is small relative to the mass of the actuator. Its effect on the response plots to be presented below could easily be added; that distal mass would appear to dominate the behavior at high frequencies. However, so long as this mass is small relative to the mass of the actuator, the frequency at which its effect would appear would be high.

The control shown in Figure 3 is based on feeding back an estimate for the force F_e, which may be produced with measures of x and y and knowledge of the stiffness k. Thus the series elastic element may be called a kind of cheap (and rather compliant) force sensor. A controller $D(s)$ applies the force F_m through a motor acting on mass M so as to cause F_e to track the signal F_d produced by the virtual environment impedance Z_{VE}. Note that the force F_e can be interpreted as a disturbance to the action of the controller $D(s)$. The virtual environment Z_{VE} responds to the sensed user motion y as in Figure 1A. A virtual coupler can be used to invert the causality of the virtual environment as in Figure 1B if desired.

To derive expressions for the open and closed-loop impedance in the case of the force-controlled SEA device, we begin by reading out the signal F_e in the block diagram.

$$F_e = k\left(\frac{1}{s\,(Ms+B)}\,(F_m - F_e) - Y\right), \tag{10}$$

from which the open-loop impedance is immediately available:

$$Z_{\text{OL}} \triangleq \left.\frac{F_e}{Y}\right|_{F_m=0} = -\frac{ks(Ms+B)}{Ms^2+Bs+k}. \tag{11}$$

The closed-loop impedance requires reading F_m from the block diagram

$$F_m = -D(s)(Z_{\text{VE}}Y + F_e), \tag{12}$$

which can be combined with Eq. (10) to produce the closed-loop impedance

$$Z_{\text{CL}} \triangleq \frac{F_e}{Y} = -\frac{k\,(D\,(s)\,Z_{\text{VE}} + Z_h)}{Z_h + kD(s) + k}. \tag{13}$$

In this expression we see that high gain control in the force controller $D(s)$ leads to the impedance Z_{CL} approaching the desired impedance Z_{VE}, at least at the low frequency range. At high frequencies, Z_{CL} tends to converge to the stiffness k, given that the effect of the mass M is dominant on the high order term.

2.4 Series Elastic Actuator Devices—Position Control

The second control architecture for the haptic device based on series elastic actuators features an inner position-control loop, as shown in Figure 4A. The position x of the mass M is fed back and compared to a desired position signal x_d produced by the force controller $D(s)$. A new position controller $C(s)$ is tasked with minimizing the error $x_d - x$. Note that the force F_e in the series elastic

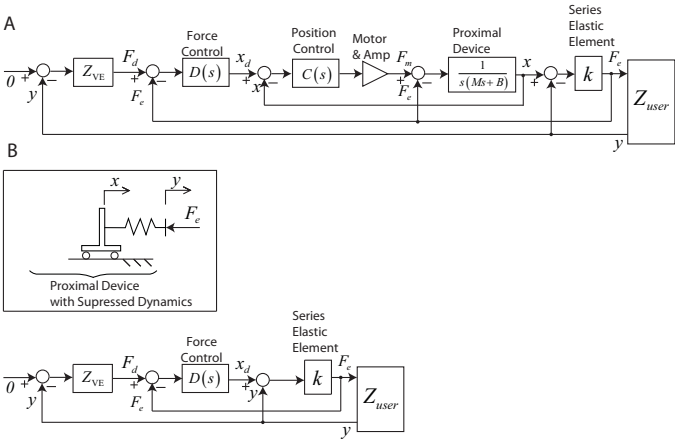

Figure 4: Series Elastic Actuator Control Model—Position Control

element is a disturbance to be rejected by the controller $C(s)$. The outer force-control loop with controller $D(s)$ attempts at the same time to cause F_e to track F_d.

Under the assumption that $C(s)$ is a high gain controller that can impose a position x_d on the proximal device dynamics Z_h^{-1} over a significant bandwidth, we may reduce the block diagram in Figure 4A to that in 4B. An inset schematic is shown with the dynamics of the proximal device suppressed. The proximal device is now simply a position source with imposed position x_d. The force controller $D(s)$ is still in place to cause the force F_e experienced by the user to track the force F_d produced by the virtual environment Z_{VE} in response to the user-imposed position y.

Reading the signal F_e from the block diagram,

$$F_e = k(X_d - Y), \tag{14}$$

where we have assumed that $x = x_d$ due to the action of the high-gain position controller $C(s)$. The open-loop impedance of the position-controlled SEA device may be determined under the assumption that the position of the mass M is fixed at $x = 0$ by commanding $x_d = 0$. We have assumed that the proximal device impedance Z_h has been fully compensated so that it functions as a motion source.

$$Z_{\text{OL}} \triangleq \left.\frac{F_e}{Y}\right|_{x_d=0} = -k. \tag{15}$$

Substituting an expression for the signal x_d into Eq. (14)

$$F_e = -k\,(D(s)(Z_{\text{VE}}Y + F_e) + Y)\,, \tag{16}$$

the closed-loop impedance can be determined

$$Z_{\text{CL}} \triangleq \frac{F_e}{Y} = -\frac{k(D(s)Z_{\text{VE}} + 1)}{kD(s) + 1}, \tag{17}$$

This expression shows that how the rendered impedance Z_{CL} approaches the desired or target impedance Z_{VE} as the gains in the force controller $D(s)$ are tuned up.

2.5 Comparative Analysis

Figure 5 shows the Bode plot of open loop impedance Z_{OL} for an Impedance Control Device, an Admittance Control Device, and a Series Elastic Device Design. To produce these particular curves, the values in Table 1 were used. The mass $m = 100$g for the Impedance Control Device is taken from the PHANToM [15] while the mass M for the Admittance Control and Series Elastic Device

is taken from Table 2. The damping values were set arbitrarily or to achieve overdamped behavior.

The frequency response curves in Figure 5 describe the *open-loop* behavior of each device (without control). For the Impedance Control and Admittance Control Devices, the mass-damper models of the hardware dictate that a magnitude response with slope of 20dB/decade will give way to a slope of 40dB/decade as frequency increases. The high frequency behavior of both the Impedance Control and Admittance Control Devices can be described as a mass (magnitude slope of 40dB/decade), with the Admittance Control device carrying the larger mass ($M > m$). At low frequencies the open-loop impedance of the Series Elastic Device matches that of the Admittance Control Device (higher than the Impedance Control Device). At high frequencies, however, the impedance magnitude of the Series Elastic Device is dominated by the stiffness of the spring k, producing a flat response. The flat impedance response of the Series Elastic Device at high frequencies inevitably leads to a high frequency impedance that is lower than either the Impedance Control or Admittance Control Devices. This is the chief benefit of purposely including a physical spring between the motor and end-effector.

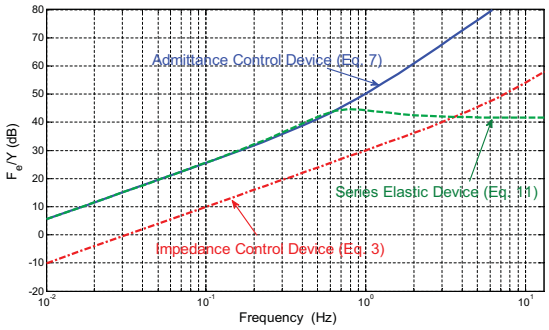

Figure 5: Bode plot showing open-loop impedance Z_{OL} for an Impedance Control Device, an Admittance Control Device, and a Series Elastic Device Design. Parameter values were drawn from Table 1.

Table 1: Parameter Values

Parameter	symbol	value
impedance device effective mass	m	0.1 kg
admittance device effective mass	M	6.56 kg
series elastic device effective mass	M	6.56 kg
impedance device damping	b	5 Ns/m
admittance device damping	B	30 Ns/m
series elastic device damping	B	30 Ns/m
series elasticity	k	120 N/m

3 Apparatus and Target Application

We are developing powered exoskeletons to be worn across joints of an amputee's residual limb. The exoskeleton power is used to provide haptic feedback regarding the interaction forces that develop between a powered prosthetic hand and a grasped or touched object. In particular, the interaction force is displayed as an extension or flexion moment across the elbow. Thus the exoskeleton is a force reflecting "master" while the powered prosthetic terminal device is the "slave" of a teleoperator. The wearer teleoperates a prosthetic hand from his residual limb and the interface is very much like that of a body-powered prosthesis. In related research, we are comparing the ease of use of body-powered and myoelectric prosthetic devices [20,21]. Note that myoelectric devices lack force reflection. Figure 6b shows a prototype exoskeleton that employs a the driven cylinder. The driving cylinder is shown in Figure 6c.

Figure 6a shows a design rendering featuring both the exoskeleton with driven cylinder attached and the powering device with the driving cylinder. The driving cylinder is actuated through a 5mm pitch ball screw by a Maxon RE65 motor. A carriage riding on two linear slides couples the motion of the ballscrew nut and the piston of an Aircylinder RLF10A-DAP-NA00 1.5in bore double acting cylinder (Norgren®, Brookville OH). The driven cylinder is an Airpel E16D5.0U 16mm bore antistiction double acting cylinder (Airpot Corp., Norwalk CT). The driven cylinder is mounted on a lightweight exoskeleton made out of ABS plastic. The cylinders are connected on both ends of the double action. We selected a high performance cylinder for the slave device but a comparatively low performance cylinder for the master, since the friction (and other dynamic effects) of the driving cylinder are suppressed through feedback control and hidden from the user by virtue of the series elasticity, whereas the dynamic effects in the driven cylinder are not suppressed. Note, the current design may be susceptible to leaks across the piston seals causing drift of the force reading, though we have not noticed this issue. Such susceptibility could be accounted for in future design iterations by implementing a pressure sensor in the pneumatic line instead of force estimation through position sensing.

The compliance of the device is affected by the volume of air in the pneumatic circuit. For example, increasing the length of the tether would reduce the stiffness and likewise the high frequency response. Increasing the bore size and stroke length of either cylinder would have the same effect. Increasing the driving cylinder bore size would also increase the mechanical advantage of the transmission, increase the available workspace for the driving cylinder, and increase overall force capacity of the device. Because it would reduce the elastic element stiffness, however, a given force output would require more displacement in the driving cylinder and would require more capacity (increase velocity and force saturation) from the driving motor.

To prepare to experimentally validate some of the transfer function models presented above, we require some parameter values from our physical apparatus. From our force-control series elastic model in Figure 3 we require values for the three mechanical elements M, B, and k. The stiffness k of the series elastic element and the damping b_a will be identified using system identification experiments presented below, but the mass M is essentially available from data sheets. Note that the mass M is an equivalent mass representing the lumped mass of our cylinder rod, motor carriage, and other items that translate, and the moment of inertia of our motor and ball screw and other items that rotate each divided by the ball screw pitch squared. The values for each of these parameters was taken from corresponding data sheets, and can be found in Table 2. The equivalent mass M_{eq} is then 6.56 kg.

4 Experimental Results

The apparatus described above has replaced a previous design in the service of ongoing experiments in our labs concerning sensory feedback for upper-limb prosthetic devices. The previous design was based on a traditional impedance control device featuring a capstan drive and DC motors [22]. To characterize the performance of the new apparatus and validate some of the analytical models discussed above, we conducted a battery of system identification experiments.

In an initial experiment, we measured the force/displacement relationship at the driven cylinder while holding the driving cylinder fixed. Backdriving the driven cylinder in both directions produced

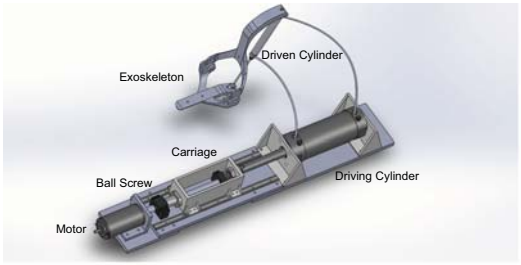

(a)

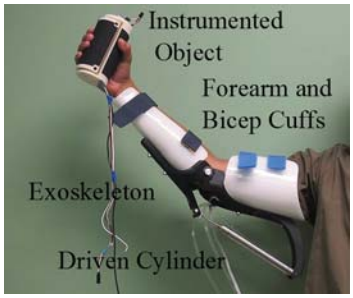

(b)

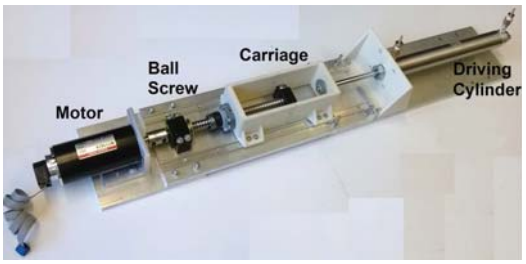

(c)

Figure 6: Prototype Prosthetic Interface Featuring Force Feedback. An exoskeleton worn across the elbow (part of a transradial amputee's residual limb) is powered through a pneumatic circuit. The *driven* cylinder spans the elbow axis and is powered through flexible tubing by a *driving* cylinder that sits atop a table in the present prototype. The driving cylinder is powered in turn through a ball screw and DC motor according to a force signal acquired in the instrumented object. The prosthetic hand (not shown) can be driven by position, myoelectric, or electroencephalogram (EEG) signals.

hysteresis in the force/displacement relationship, from which we estimated damping in the driven cylinder at 23.5 Ns/m. Note damping in the driven cylinder is not captured in the models developed above; the damping b_a represents dissipative effects acting on the *driving* cylinder. We used a line fit to the upper and lower portions of the hysteresis loop to estimate stiffness at 120 N/m. Systematic deviation from a straight line was not apparent in our data. We did not measure damping in the driving cylinder directly, though its value can be estimated indirectly through the frequency response experiments below.

In a separate experiment, we measured and compared the open-loop and closed-loop impedance of the system. To measure the open-loop impedance $\left(\frac{F_e}{Y}|_{F_m=0}\right)$, we turned the motor "off." We then imposed a sinusoidal displacement Y on the driven cylinder by hand with frequencies in the range $0.8-5$ Hz and measured the resulting force F_e. Basically a human user produced a backdrive motion that was rich in frequencies while force and displacement were measured for about three minutes. To measure the closed-loop impedance $\left(\frac{F_e}{Y}\right)$, we placed the motor and driving cylinder under proportional-derivative control with setting $Z_{\rm VE}$ to 0. We again imposed a sinusoidal displacement Y on the driven cylinder as in the open-loop case, and measured Y and the resulting force F_e. The data from both the open-loop and closed-loop tests were processed in Matlab to produce an estimate of the open-loop and closed-loop impedance frequency response, respectively, as shown in Figure 7.

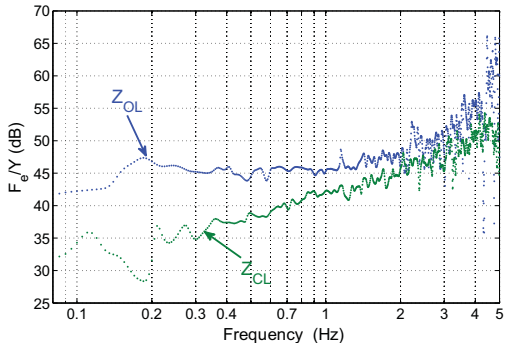

Figure 7: Experimental Frequency Response Estimate: open-loop control impedance and closed-loop control impedance; input was a broad spectrum position y, while the output was the measured force F_e.

From these data, we find that the open-loop impedance is consistent with our model in Eq. (11). As frequency increases, the open-loop impedance (the transfer function from displacement to force) converges to a value corresponding closely to our previously identified stiffness $k = 120$ N/m. Since the device dynamics involving M and B has a high impedance (the device is unlikely to move), we can see the same phenomenon that the impedance magnitude is "flat" even at low frequencies as in Eq. (15) in which case the device is fixed at 0 with a high gain controller $C(s)$. However, our apparatus diverges at high frequencies mainly due to a small mass in the distal portion of the device not accounted for by our model. We measured the mass in the distal portion of the device as 34 grams.

For the closed-loop controller, the impedance is lower than that of our open loop impedance for frequencies up to $\sim$ 2 Hz. This lower impedance is consistent with our model in Eq. (13) assuming $Z_{\rm VE} = 0$ (moving in free space). After $\sim$ 2 Hz the closed-loop impedance approaches that of the open-loop impedance because in effect the frequency of movement is too fast for the motor driving the proximal device to respond, effectively turning it 'off.'

In another experiment, we conducted an analysis in which we commanded sinusoidal motions to the driving cylinder X and monitored the force F_e produced. Based on the stiffness, we can scale X to the motor force F_m. Thus for a series of sinusoidal frequencies in the range $0.02-4$ Hz, we characterized the frequency response of our device $\left(\frac{F_e}{F_m}\right)$. We found our bandwidth to be 2 Hz as shown in Figure 8.

5 Conclusion

We have built a haptic device featuring a pneumatic transmission that plays multiple roles. It transmits power and realizes a significant mechanical advantage, it is the elastic element in a series elastic actuator, and it realizes a flexible power linkage that allows us to minimize the weight of a wearable robot. The present instantiation of our design features a very compliant elastic element,

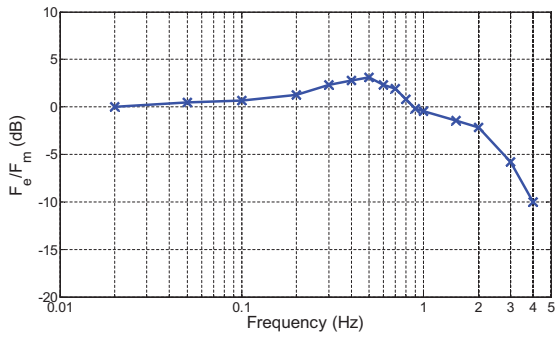

Figure 8: Experimental Frequency Response Estimate: input is a sinusoidal position of the driving cylinder (scaled as desired force F_m) at particular frequencies and output is the force produced by the driven cylinder.

whose value has not been tuned for performance. While we recognize there exists a tradeoff between the benefits of low compliance and the cost of curtailed bandwidth, our current design does not seem to lie up against that tradeoff curve. We do believe we are out from under the design tradeoff that plagues common haptic device designs, whether impedance-control or admittance control devices, namely a tradeoff between a maximally free free-space behavior and a compromised capacity for responding with large force.

Table 2: Hardware Parameter Values

Parameter	Symbol	Value	Unit
Motor rotor inertia	I_{rotor}	1.38×10^{-4}	kgm^2
Ball screw inertia	I_{screw}	2.6×10^{-5}	kgm^2
Carriage mass	m_{carriage}	0.254	kg
Piston and rod mass, driving cylinder	m_{driving}	0.2	kg
Piston and rod mass, driven cylinder	m_{driven}	0.034	kg
Driving cylinder stroke	S_{driving}	0.254	m
Driven cylinder stroke	S_{driven}	0.127	m
Driving cylinder bore	A_{driving}	0.0011	m^2
Driven cylinder bore	A_{driven}	0.0002	m^2
Proximal device equivalent mass	M_{eq}	6.56	kg

ACKNOWLEDGEMENTS

This work was supported in part by grant IIS-1065027 from the National Science Foundation.

REFERENCES

[1] C. R. Carignan and K. R. Cleary, "Closed-loop force control for haptic simulation of virtual environments," *Haptics-e*, vol. 1, no. 2, pp. 1–14, 2000.

[2] M. Ueberle and M. Buss, "Control of kinesthetic haptic interfaces," in *Proc. IEEE/RSJ Int. Conf. on Intellig. Rob. and Syst., Workshop on Touch and Haptics*, vol. 22, 2004.

[3] R. Q. Van der Linde, P. Lammertse, E. Frederiksen, and B. Ruiter, "The hapticmaster, a new high-performance haptic interface," in *Proc. Eurohaptics*, 2002, pp. 1–5.

[4] G. A. Pratt, M. M. Williamson, P. Dillworth, J. Pratt, and A. Wright, "Stiffness isn't everything," pp. 253–262, 1997.

[5] C. Fitzgerald, "Developing baxter," in *Technologies for Practical Robot Applications (TePRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1–6.

[6] M. Zinn, O. Khatib, B. Roth, and J. K. Salisbury, "Playing it safe [human-friendly robots]," *Robotics & Automation Magazine, IEEE*, vol. 11, no. 2, pp. 12–21, 2004.

[7] J. F. Veneman, R. Ekkelenkamp, R. Kruidhof, F. C. van der Helm, and H. van der Kooij, "A series elastic-and bowden-cable-based actuation system for use as torque actuator in exoskeleton-type robots," *The international journal of robotics research*, vol. 25, no. 3, pp. 261–281, 2006.

[8] H. Vallery, J. Veneman, E. Van Asseldonk, R. Ekkelenkamp, M. Buss, and H. Van Der Kooij, "Compliant actuation of rehabilitation robots," *Robotics & Automation Magazine, IEEE*, vol. 15, no. 3, pp. 60–69, 2008.

[9] K. Kong, J. Bae, and M. Tomizuka, "A compact rotary series elastic actuator for human assistive systems," *Mechatronics, IEEE/ASME Transactions on*, vol. 17, no. 2, pp. 288–297, 2012.

[10] H. Yu, S. Huang, N. V. Thakor, G. Chen, S.-L. Toh, M. Sta Cruz, Y. Ghorbel, and C. Zhu, "A novel compact compliant actuator design for rehabilitation robots," in *Rehabilitation Robotics (ICORR), 2013 IEEE International Conference on*. IEEE, 2013, pp. 1–6.

[11] K. B. Fite, T. J. Withrow, X. Shen, K. W. Wait, J. E. Mitchell, and M. Goldfarb, "A gas-actuated anthropomorphic prosthesis for transhumeral amputees," *Robotics, IEEE Transactions on*, vol. 24, no. 1, pp. 159–169, 2008.

[12] D. Paluska and H. Herr, "The effect of series elasticity on actuator power and work output: Implications for robotic and prosthetic joint design," *Robotics and Autonomous Systems*, vol. 54, no. 8, pp. 667–673, 2006.

[13] X. Shen and M. Goldfarb, "On the enhanced passivity of pneumatically actuated impedance-type haptic interfaces," *Robotics, IEEE Transactions on*, vol. 22, no. 3, pp. 470–480, 2006.

[14] ——, "Simultaneous force and stiffness control of a pneumatic actuator," *Transactions-American Society Of Mechanical Engineers Journal Of Dynamic Systems Measurement And Control*, vol. 129, no. 4, p. 425, 2007.

[15] T. H. Massie and J. K. Salisbury, "The phantom haptic interface: A device for probing virtual objects," in *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, vol. 55, no. 1. IOS Press, 1994, pp. 295–300.

[16] N. Hogan, "Impedance control: An approach to manipulation: Part illapplications," *Journal of dynamic systems, measurement, and control*, vol. 107, no. 2, p. 17, 1985.

[17] D. A. Lawrence, "Stability and transparency in bilateral teleoperation," *Robotics and Automation, IEEE Transactions on*, vol. 9, no. 5, pp. 624–637, 1993.

[18] P. G. Griffiths, R. B. Gillespie, and J. S. Freudenberg, "A fundamental linear systems conflict between performance and passivity in haptic rendering," *Robotics, IEEE Transactions on*, vol. 27, no. 1, pp. 75–88, 2011.

[19] G. A. Pratt and M. M. Williamson, "Series elastic actuators," in *Intelligent Robots and Systems 95.'Human Robot Interaction and Cooperative Robots', Proceedings. 1995 IEEE/RSJ International Conference on*, vol. 1. IEEE, 1995, pp. 399–406.

[20] J. D. Brown, R. B. Gillespie, D. Gardner, and E. A. Gansallo, "Co-location of force and action improves identification of force-displacement features," in *Haptics Symposium (HAPTICS), 2012 IEEE*. IEEE, 2012, pp. 187–193.

[21] J. D. Brown, A. Paek, M. Syed, M. K. O'Malley, P. A. Shewokis, J. L. Contreras-Vidal, A. J. Davis, and R. B. Gillespie, "Understanding the role of haptic feedback in a teleoperated/prosthetic grasp and lift task," in *World Haptics Conference (WHC), 2013*. IEEE, 2013, pp. 271–276.

[22] B. Gillespie, J. Baker, M. O'Malley, P. Shewokis, and J. L. Contreras-Vidal, "Functionally biarticular control for smart prosthetics," in *EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint*. IEEE, 2009, pp. 627–628.