

CLOSED-LOOP FORCE CONTROL FOR HAPTIC SIMULATION OF VIRTUAL ENVIRONMENTS

Craig R. Carignan *

Department of Aerospace Engineering
University of Maryland
College Park, MD 20742

Kevin R. Cleary †

Department of Radiology
Georgetown University Medical Center
Washington, DC 20007

Abstract

Force feedback control is investigated for improving the quality of the haptic feedback in virtual reality applications. Advanced control design can increase the transparency of the haptic device at the haptic interface thereby increasing the realism of the simulation. Force feedback also enables the implementation of admittance control approaches heretofore considered the domain of large robotic platforms. The quality of the haptic interface is first quantified, and the standard open-loop impedance approach to haptic control is reviewed. Dual-input control schemes with sensory force feedback are then introduced, and the resulting quality of the haptic interface is derived. A four-bar linkage example is used to illustrate the improvement in fidelity realized with the various approaches. The tradeoffs encountered in moving to force feedback controllers for haptic applications are also discussed.

Z_{hCL}	closed-loop impedance matrix, N/m
Z_d	desired impedance matrix, N/m
Z_e	impedance error matrix, N/m
ΔZ	relative impedance error (scalar)
K_F	force gain matrix
K_d	desired stiffness gain matrix, N/m
B_d	desired damping gain matrix, N/m/s
C	impedance compensator gain matrix, m/N
jD	joint servo compensator matrix, N-m
D	Cartesian servo compensator matrix, N/m
K_P	joint stiffness gain matrix, N-m
K_D	joint damping gain matrix, N-m/s
J	Jacobian matrix, m
M	inertia matrix, kg-m ²
G	gravity torque vector, N-m

Nomenclature

θ	joint angle vector, rad
θ_c	commanded joint angle vector, rad
x	position vector, m
x_d	desired position vector, m
x_c	commanded position vector, m
τ	actuator torque input vector, N-m
τ_{fwd}	feedforward torque vector, N-m
τ_d	desired torque input vector, N-m
τ_c	commanded torque vector, N-m
F	force output vector, N
F_d	desired force output vector, N
jZ_h	joint impedance matrix, N-m
${}^j\hat{Z}_h$	modeled joint impedance matrix, N-m
Z_h	impedance matrix, N/m

1 Introduction

The development of haptic devices for virtual reality applications is already rivaling the development of industrial manipulators which took place two decades earlier. This rapid pace largely owes its existence to the advent of powerful computer workstations which allow the rendering of virtual environments in real-time. These robotic mechanisms form the kinesthetic counterpart to the visual feedback from the workstation and are a direct descendant of force-reflecting hand controllers used in teleoperation. Desktop haptic devices are already finding widespread use as force “displays” for computer games, nanoscale manipulation, and surgical simulators.

At the core of this development is creating a “realistic” force-position interface with the virtual environment. The quality of this interface can be measured in terms of impedance “accuracy”, which refers to how close the impedance matches that of the virtual environment, and impedance “resolution” or “fidelity”, which refers to the level of impedance discrimination that can be detected at the haptic interface. For high output force applications such as driving simulators and smart exercise ma-

*Research Associate, Space Systems Laboratory

†Research Assistant Professor, Division of Imaging Science and Information Systems

chines, the former is usually more important. For the dexterous procedures involved in surgical simulation, the latter is more likely to be important.

One of the principal barriers to achieving high impedance accuracy and precision are the dynamics of the haptic device itself. The natural dynamics detract from the realism of the haptic feedback since they are perceived by the human operator as part of the simulated environment. Improvements in robotic design over the last several years have done much to reduce the natural dynamics of haptic devices over the previous generation of master arms. These advancements are largely due to more efficient drive trains (cable, harmonic drives) and higher strength-to-weight ratio materials.

However, as the demand for force production continues to rise, it will become increasingly difficult to achieve further reductions by strictly physical means; larger actuators, drive mechanisms, and linkages all lead to more inertia and friction in the haptic device. The trend toward higher output haptic devices is already apparent as traditionally robotic mechanisms are now being employed as haptic devices. Some examples of this include manipulators used as driving and flight simulators [Lee et al., 1998; Clover et al., 1997], motorized “smart” exercise machines [Book and Ruis, 1981; Li and Horowitz, 1996], and astronaut extravehicular (EVA) training using parallel mechanisms [Swaim et al., 1995] and manipulators [Carignan and Akin, 1997b].

Higher fidelity may also imply that the dynamics of the haptic device need to be reduced relative to the impedance of the virtual environment. Thus, even if the force output of current haptic devices is adequate, the natural dynamics might prohibit the finer force discrimination required for simulating more dexterous tasks. A notable example of this is the abrupt change in resistance encountered by a needle or scalpel as it penetrates a layer of tissue during a surgical procedure [Baumann and Clavel, 1998; Popa and Singh, 1998; Cleary et al., 1998].

Further reductions in the haptic device dynamics will likely come from active control. This compensation can take the form of model feedforward or force feedback from a force-torque sensor mounted at the haptic interface. While gravity and friction feedforward are often used to reduce these effects, modification of the inertial dynamics requires force or acceleration feedback. In addition, feedback control allows for a more robust design which is important when physical properties of the haptic device are changing with time. In this paper, the impact of sensory force feedback on the haptic interface will be investigated in the context of several control methodologies.

What this paper is not about is stability. While it is recognized that stable behavior is a prerequisite to implementing any form of haptic control, addressing stability issues as well as haptic quality goes beyond the scope of this paper. The controllers introduced in this report have all been implemented on either haptic or robotic systems so that they are known to have at least a range of stable behavior. The reader is referred to work by Lawrence [1988], Newman [1990], and Adams and Hannaford

[1999] for discussions on stability issues, particularly in the context of impedance control.

This paper begins with a brief history of force-controlled haptic devices and then proceeds to discuss design criteria for haptic interfaces. This is followed by an overview of the open-loop impedance approach to haptic control. The next section discusses advanced control schemes which implement sensory force feedback in the control law. A measure of the impedance tracking accuracy is then developed, and an example is used to illustrate the fidelity realized by the two approaches. Results from the example are then discussed and some conclusions are drawn regarding force feedback implementations.

2 Previous Work

The origin of haptic devices actually dates back to the master arms used in the 1950s for remote handling of radioactive materials [Goertz, 1954]. Initially just passive replicas of the slave arm, master arms were manipulated by the operator to command the slave arm performing the task. Later on, master arms were motorized so that they could impart feedback forces on the operator that were proportional to the forces being felt by the slave arm operating in the remote environment. This increased “telepresence” was found to decrease task times, particularly in environments that were unpredictable or changing [Das et al., 1992; Zhai and Milgram, 1993].

The computational requirements for master/slave (M/S) systems were fairly minimal since the control could be accomplished joint-to-joint. As computers became more powerful, complex kinematics computations could be performed in real-time enabling the development of master arms which bore little resemblance to the slave arms they were controlling. The JPL Force Reflecting Hand Controller (FRHC) was one of the first devices to emerge, and it had better force resolution than most master arms built beforehand [Bejczy and Salisbury, 1983].

It was only a matter of time before researchers began to realize that this new breed of hand controllers could be used for simulating virtual environments as well as reproducing the forces sensed in a “real” environment. In this context, FRHC’s belong to a class of robotic mechanisms usually referred to as “haptic devices”. It has long been recognized that the dynamics of the haptic device interfere with the realism of haptic feedback, and the drive to achieve low dynamics or “transparency” has always been regarded as a desirable feature in hand controller design [Hannaford, 1989; Lawrence, 1992]. By employing new materials and drive components, the mass and friction in these devices have been reduced dramatically making the current generation of haptic devices adequate for many applications.

Force-torque sensors have rarely been implemented in haptic devices due to the added mass and cost penalties. However, the advantages of incorporating them have been known from the start. An example of a simple linear haptic device with force sensing is shown in Figure 1 [Paines, 1987]. “Canute” was

driven by a cable assembly and incorporated strain gauges, an accelerometer, and a position sensor for state feedback. Derived rate information from an encoder was used to modify the apparent viscosity of the joystick. In addition, an accelerometer allowed the user to program different “feels” for the joystick, and even negative values of mass were tested.

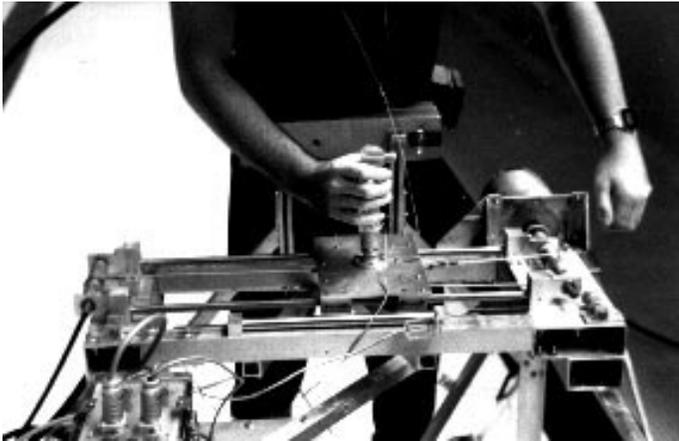


Figure 1. Canute translational hand controller undergoing testing in neutral buoyancy.

The six degree of freedom Cartesian “manipulator” developed by Yokoi et al. [1994] had a force sensor in the wrist to read user force inputs. The master arm had three orthogonal Cartesian axes in the base driven by linear actuators and three orthogonal revolute axes to drive the handle. Feedback from the virtual environment was used to minimize the input forces during free motion and to oppose the user during contact with virtual objects. The closed-loop bandwidth was estimated to be only 4 Hz, and investigators reported instability when the handle was grasped tightly by the user.

A portable haptic mechanism employing force sensing is the GLAD-IN-ART arm exoskeleton [Bergamasco et al., 1994]. This haptic device was designed to be worn over a human arm for the purpose of a more fully immersive simulation. Torques ranged from 20 N-m at the shoulder joints to 2 N-m at the forearm. Gravity feedforward compensation was used to reduce the gravitational loads due to the 10 kg mass of the device. The exoskeleton incorporated a six-axis force-torque sensor below the handle grasped by the operator which was used to close a force loop at the wrist.

Haptic devices with force-sensing capability have also been used as mass simulators in space applications. At the Virtual Reality Laboratory at the Johnson Space Center, EVA crew have trained using fully immersive virtual reality hardware, and have dynamically interacted with massive payloads through the use of a CharlotteTM parallel-cable manipulator system produced by McDonnell-Douglas Corporation [Swaim et al., 1995]. The

Charlotte robot is supported by eight cables connected to the corners of its box structure which are reeled in and out via a capstan/motor system allowing movement within the large workspace area encompassed by the cable assembly.

The Ranger Dexterous Manipulator (DXM) is an example of a serial manipulator that has been used as a mass simulator [Carignan and Akin, 1997b]. By mounting a force-torque sensor at the end of the manipulator, researchers were able to simulate free-floating masses of 125 kg to 1000 kg. In addition, a force “accommodation” mode was also tested for damping values as low as 250 N/m/s. Further reductions might have been realized, but the maximum impedance loop rate attained was 40 Hz due to hardware limitations.

The mechanisms described above are all examples of haptic devices which incorporate force sensors in the hardware. Commonly used haptic devices such as the PHANTOMTM and Impulse EngineTM do not incorporate force sensors in their design. In the next section, the standard open-loop approach to haptic control is examined first, and then closed-loop force control approaches are introduced.

3 Haptic Control Design

Recent developments in haptic interfaces and 3D computer graphics provide the basic hardware and software platforms to incorporate these devices into working systems. For example, a surgical simulator can be constructed using a PHANTOMTM haptic interface, a computer graphics workstation, and an environmental model. While several of these simulators have been built, haptic fidelity is still an issue. Most of the simulators constructed so far concentrate on modeling the environment and use relatively simple control laws for the haptic feedback.

A block diagram of a typical haptic system for surgical simulation is shown in Figure 2 [Cleary et al., 1998]. The simulator consists of four blocks: the physician (human operator), the haptic device, the anatomical model of the patient (environment), and the visual interface (computed tomography (CT) image). The physician interacts with the simulator by positioning the haptic device, which exerts forces back that are similar to the actual patient. The anatomical model computes the desired Cartesian forces and torques on the tool based upon the physical properties of the environment and the position of the tool. The haptic controller calculates the motor torques for the haptic device based on the sensed position and/or force and the desired interaction forces and torques coming from the anatomical model. The motor torques are then sent to the haptic device, and any updates to the model are sent to the visual interface, typically a 3D graphical representation of the environment.

The haptic control and visual computations may be done on different computing platforms, depending on the system architecture. In a typical implementation, the haptic device must be updated at a fairly high rate (on the order of 1000 Hz) to ensure stability and a responsive interface. On the other hand, the visual

interface can be updated at a much lower rate, perhaps as low as 30 Hz, due to the relatively low threshold of visual perception.

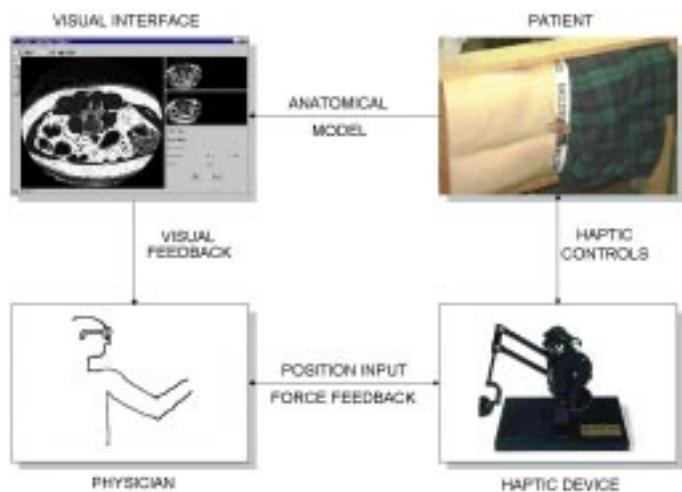


Figure 2. System block diagram for the Georgetown University Spinal Biopsy Simulator.

The haptic device can be regarded as a transmission between the human hand and the motors which generate the forces in the simulation. In between these are the tool interface, the linkages, and the transmissions to the motor shafts. The motors are driven by the haptic controller which is programmed to replicate the desired stiffness, viscous, and mass properties of the environment as well as constraint surfaces or interfaces that might be present.

The PHANToMTM, shown in Figure 3, is a three degree of freedom device consisting of a four-bar parallel linkage attached to a vertical rotational axis. The parallel linkage allows the three motors to be mounted at the base giving an apparent mass at the tool tip of about 100 grams. Cable drives are used for torque amplification to reduce the effect of friction usually found in geared mechanisms. A continuous force of 1.5 N can be generated at the tool tip and peak forces of 10 N can be applied.

Haptic engineers typically employ three criteria when designing haptic mechanisms [Massie and Salisbury, 1994]: (1) *free space must feel free*, (2) *solid virtual objects must feel stiff*, and (3) *virtual constraints must not be easily saturated*. The first criterion implies that the natural dynamics of the haptic device should not distract the user from the environment being simulated (PHANToMTM is < 0.2 N). This implies that the apparent mass and friction of the apparatus should be as low as possible. The second criterion says that the haptic device must be capable of producing a stiffness convincing enough to believe that contact with an immovable object has taken place. This stiffness is usually taken to be a minimum of 20 N/cm. The third criterion implies that the haptic mechanism must be capable of produc-

ing enough force so that virtual objects feel solid. For fingertip contact, a force of 10 N is rarely exceeded. However, for a grasp situation, the force can be much higher. The above values as well as criteria were all derived from Massie and Salisbury [1994].

The first criterion regarding fidelity can be satisfied either through passive design or by active control. The second requirement on stiffness is satisfied by making the mechanism stiff and using a very high bandwidth controller. The latter implies that the control law be computed at a very fast rate. The saturation requirement is a function of the peak torque outputs of the motors. The higher the force produced at the haptic interface, the larger the torques that must be supplied by the motors.

Haptic devices such as the PHANToMTM were designed from the outset to have low dynamic properties and thus easily satisfy the first design criterion. However, in many applications, the operator exerts far more than 1.5 N of continuous force thus violating the third criteria and rendering low output devices such as the PHANToMTM inadequate for simulation and training. For example, during a lumbar puncture insertion, surgeons typically apply in excess of 10 N of force [Popa and Singh, 1998]. It seems likely that in future applications, higher output devices with potentially larger masses and drive mechanisms will be required.

In some haptic devices, certain components of the natural dynamics can be compensated for through mechanical design. The gravitational moments, for example, can be “removed” through counterbalancing. However, there is a penalty associated with the higher inertia from the counterweight. Other passive measures include decoupling the inertia of the apparatus by proper choice of link lengths and masses. While this simplifies the dynamics, it does not reduce the principal inertias. These measures are also restricted to certain classes of mechanisms usually employing parallel linkages [Huissoon and Wang, 1991].



Figure 3. PHANToMTM haptic interface (Image courtesy of SensAble Technologies, Inc.)

A typical master arm control scheme is depicted in Figure 4. The master arm is comprised of a set of motors/encoders, linkages, drive transmissions, and a gripper. Encoders attached to the motor shafts read the angles of the motors which drive the linkages through a transmission. The joint angles are then sent to a forward kinematics module inside the controller to determine the position of the gripper. The position is then mapped into a desired position command for the slave arm end-effector, x_c , and sent to the slave arm controller. The slave arm controller passes the forces sensed at the manipulator, F_s , back to the master arm controller which then maps this into a force “reflection” command. The force command is mapped using a Jacobian to a set of desired torques to be produced by the motors in the master arm. The signal is then converted into desired currents for the servo amplifiers driving the motors.

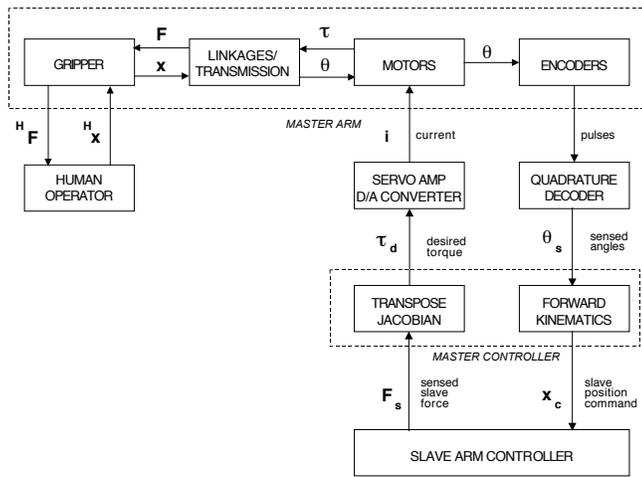


Figure 4. Typical master arm controller used in M/S teleoperation.

This type of force control is called “open-loop” because there is no force feedback from the device to the controller to regulate the force output of the master. The accuracy of the force thus depends largely upon the master arm (or hand controller) having negligible dynamics. Any friction, gravitational, or inertial torques in the device will directly add to the forces felt by the user. In addition, any errors in the geometric modeling or output of the servos will also contribute to inaccuracies perceived by the user.

There are two classes of control schemes available for force reflection: impedance control and admittance control [Baumann and Clavel, 1998]. Impedance controlled systems detect the motion commanded by the operator and control the force applied by the haptic device. Admittance controlled systems detect the force commanded by the operator and control the velocity or dis-

placement of the haptic device. Sometimes, force is used as an additional input to the impedance controller or displacement is used as an additional input to the admittance controller. However, the type of output (force or position) can ordinarily be used to determine the class of controller being used.

In the past, which approach to use often depended on the application being considered. Impedance controllers were generally used when the environment being simulated was highly compliant such as human tissue in surgical simulators. Admittance control was generally used when the environment was unyielding such as flight simulator platforms. However, the use of force and velocity inputs for both classes of control often blurs this distinction. In the remainder of this section, a “pure” impedance controller will be examined as well as two dual-input controllers: impedance control with force feedback [Sakaki et al., 1992; Bergamasco et al., 1994], and admittance control with position feedback [Maples and Becker, 1986; Pelletier and Doyon, 1994].

3.1 Open-Loop Impedance Control

Off-the-shelf haptic devices typically use controllers based on the master arm design discussed in the previous section. A block diagram of a typical haptic controller is shown in Figure 5 [Massie and Salisbury, 1994]. Encoders mounted on the motor shafts read the angles of the motors which are converted to joint angles. The angles are then sent to a forward kinematics module inside the controller to determine the position of the tool tip, x_s . An environmental model uses the position and desired physical properties to generate a desired force response, F_d . The desired force is then mapped using a Jacobian to a set of desired torques to be produced by the motors in the haptic device.

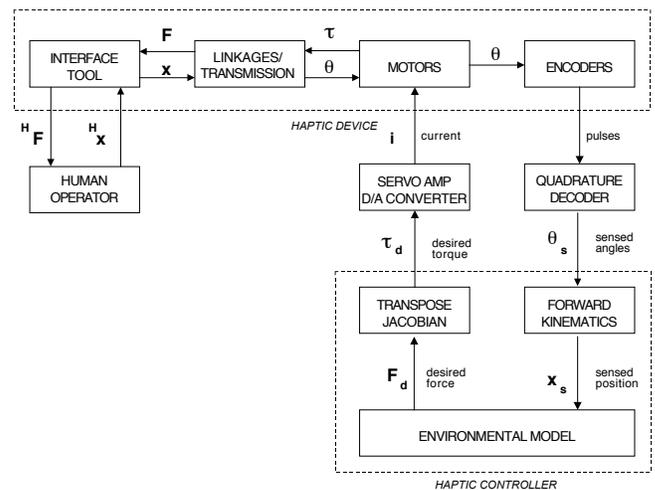


Figure 5. PHANTOMTM haptic interface controller.

This block diagram is almost identical to the master arm controller shown in Figure 4 except the slave arm controller has been replaced by an environmental model. Thus while the master arm is attempting to replicate the forces felt by the slave arm in a real environment, the haptic device is attempting to produce forces generated by interaction with a virtual environment. In addition, the position of the haptic device directly affects the forces generated by the virtual environment, whereas the slave arm intercedes between the master arm and the environment.

The PHANTOMTM controller in Figure 5 can be recast as the control block diagram in Figure 6. In this diagram, the linearized haptic device dynamics are represented by ${}^J Z_h$ which is the relationship between the joint torque inputs and the differential angular output. This will be referred to as the “joint-space impedance” of the haptic device. (Strictly speaking, the impedance is the relationship between velocity and force not position and force.) Since the positions on this linearized block diagram are represented by deviations from some nominal position, x_o , the Cartesian position of the haptic interface and the joint position of the haptic mechanism are related through the Jacobian, J .

Referring to the block diagram in Figure 6, the controller takes the difference between the desired position deflection, $\Delta x_d \equiv x_d - x_o$, and the actual position deflection of the haptic interface, $\Delta x \equiv x - x_o$, and multiplies it by the desired environmental impedance, Z_d , to generate a desired force command, F_d . The commanded force is then mapped into a set of desired torque commands using the Jacobian, J , which are then applied to the haptic device. Note that F is the negative of the force applied by the human operator and represents a physical, not control, input to the haptic device.

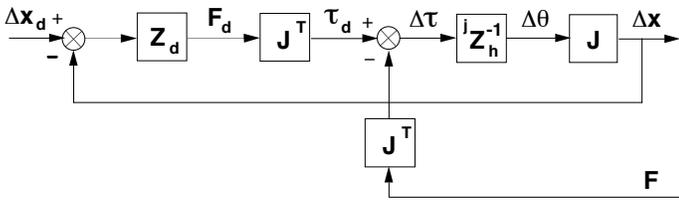


Figure 6. Open-loop impedance controller.

The block diagram in Figure 6 can be used to determine the closed-loop impedance achieved at the haptic interface. The following two relationships can be obtained from the block diagram

$$\Delta x = J^j Z_h^{-1} [J^T F_d - J^T F] \quad (1)$$

$$F_d = Z_d (\Delta x_d - \Delta x) \quad (2)$$

where $\Delta \tau$ denotes the deviation of the total torque applied to the haptic device, τ , from the nominal (gravitational) torque, τ_o . Substituting (2) into (1) and defining the open-loop impedance

of the haptic device, $Z_h \equiv J^{-T} J Z_h J^{-1}$, gives

$$\Delta x = Z_h^{-1} [Z_d (\Delta x_d - \Delta x) - F] \quad (3)$$

Combining terms in Δx and setting $\Delta x_d = 0$ without loss of generality yields

$$[Z_h + Z_d] \Delta x = -F \quad (4)$$

The closed-loop haptic impedance is the relationship between the positional deviation of the haptic interface, Δx , and the force applied by the operator, $-F$,

$$Z_{hCL} = Z_d + Z_h \quad (5)$$

Thus the difference between the desired impedance and the closed-loop impedance is just the natural (open-loop) impedance of the haptic device in Cartesian space, Z_h .

Note that if the desired impedance Z_d contains a mass or acceleration component, then endpoint acceleration feedback is required. Six-axis accelerometers are available, but are generally not as accurate as force-torque sensors. In addition, many of the same problems implicit in force control are also present with explicit acceleration feedback, namely noise and high gain instability. Therefore, if modifying the inertia is a requirement, then it might be best to consider an explicit force signal in the primary feedback path.

The impedance controller is often augmented by a model feedforward term for the haptic device dynamics, ${}^j \hat{Z}_h$, in the control law. This was the approach used by Hogan [1985] in his development of “impedance control” for robot manipulation and is illustrated by the control block diagram in Figure 7. Using (2) and the modified control input

$$\Delta \tau_{fwd} = {}^j \hat{Z}_h J^{-1} \Delta x \quad (6)$$

$$\Delta x = J^j Z_h^{-1} [\Delta \tau_{fwd} + J^T F_d - J^T F] \quad (7)$$

results in the following closed-loop impedance

$$Z_{hCL} = Z_d + Z_h - \hat{Z}_h \quad (8)$$

The last term in (8), $\hat{Z}_h \equiv J^{-T} J^j \hat{Z}_h J^{-1}$, represents the dynamic model of the haptic device in Cartesian space. The goal of the model feedforward, \hat{Z}_h , is to cancel out the corresponding terms in the dynamics of the haptic device, Z_h .

While model feedforward may represent an improvement in performance, it is fraught with difficulties. One drawback is that the controller is very susceptible to modeling error; if the estimated stiffness, damping, and inertia in the haptic model are incorrect, $\hat{Z}_h \neq Z_h$, then the impedance realized at the output will also be incorrect. Some recent work in adaptive modeling may improve the picture in the future [Liu, 1997], but accurately modeling effects such as friction which is both position and temperature dependent remains elusive. In addition, model feedforward increases the computational load on the controller processor leading to lower loop rates which may compromise the haptic design stiffness criteria.

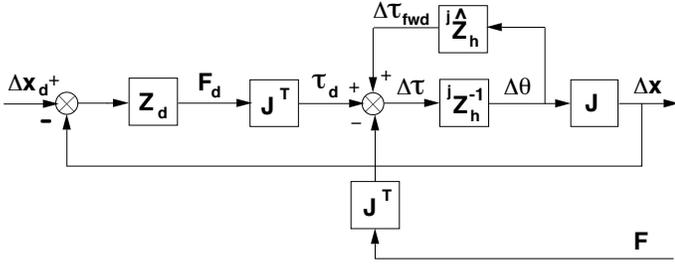


Figure 7. Open-loop impedance controller with model-feedforward.

3.2 Impedance Control with Force Feedback

An impedance controller with force feedback is illustrated in Figure 8. In comparing this controller with Figure 6, there is now a path for the sensed force signal leading back to the haptic controller. The force feedback is used to close the loop on the desired force generated by the controller. An example of a device which uses this type of control is the GLAD-IN-ART Arm Exoskeleton described earlier [Bergamasco et al., 1994].

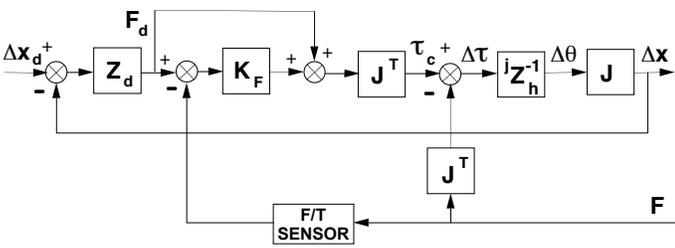


Figure 8. Impedance controller with force feedback.

The dependence of the impedance error on the force gain can be found using the block diagram shown in Figure 8. The following relations are obtained from the diagram

$$\Delta x = J^j Z_h^{-1} (\tau_c - J^T F) \quad (9)$$

$$\tau_c = J^T [F_d + K_F (F_d - F)] \quad (10)$$

$$F_d = Z_d (\Delta x_d - \Delta x) \quad (11)$$

Setting $\Delta x_d = 0$ and substituting (11) into (10) and (10) into (9), the closed-loop impedance can be determined as

$$Z_{hCL} = Z_d + (I + K_F)^{-1} Z_h \quad (12)$$

where P is the dimension of the Cartesian space and I is the $P \times P$ identity matrix. Thus, the errors due to the dynamics of the haptic device are inversely proportional to the force gain, K_F . If the force gain is set to zero, then the controller reverts to the open-loop case in (5). The force gain is typically set as high as possible which, in practice, is bounded by the stability as determined by the sample rate and natural frequencies of the haptic device.

3.3 Admittance Control

While pure admittance controllers (force input only) have been developed for robotic applications [Glosser and Newman, 1994], the approach which seems to have found the most favor is the so-called ‘‘position-based impedance controller’’ developed by Maples and Becker [1986]. To avoid potential confusion with the impedance controller, the terminology ‘‘admittance control with position feedback’’ will be used instead. As illustrated in Figure 9, this type of control consists of two control loops: an outer Cartesian loop for controlling impedance, and an inner joint servo loop for controlling joint position. In the outer loop, the sensed force is used to generate a desired tool position for the haptic device based upon the environmental model embodied by the compensator, C , and desired stiffness, K_d . An inverse kinematics module inside the controller, represented by the block J^{-1} , converts the tool position into a set of commanded joint angles for the haptic device, θ_c . A high gain PD controller at each joint, jD , is then used to servo the joint angle sensed by the encoder to the desired joint angle.

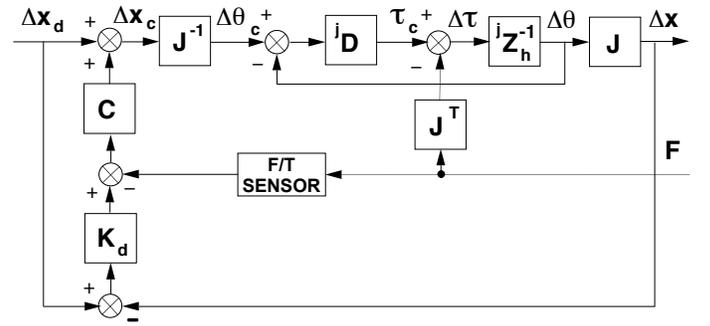


Figure 9. Admittance controller with position feedback.

Using (9) and the following new relations from the block diagram in Figure 9,

$$\tau_c = {}^jD (\Delta \theta_c - \Delta \theta) \quad (13)$$

$$\Delta \theta_c = J^{-1} [\Delta x_d + C (K_d [\Delta x_d - \Delta x] - F)] \quad (14)$$

$$\Delta \theta = J^{-1} \Delta x \quad (15)$$

and performing a considerable amount of algebra, the following closed-loop impedance is derived

$$Z_{hCL} = (DC + I)^{-1} [Z_h + D(CK_d + I)] \quad (16)$$

where $D \equiv J^{-T} {}^jD J^{-1}$ is the inner loop joint PD controller expressed in Cartesian space. As the inner loop compensator gains in jD become larger, $D \rightarrow \infty$ and

$$Z_{hCL} \approx (DC)^{-1} [D(CK_d + I)] \quad (17)$$

which simplifies to

$$Z_{hCL} \approx K_d + C^{-1} \quad (18)$$

If C is chosen to be

$$C = (Z_d - K_d)^{-1} \quad (19)$$

then Z_{hCL} approaches the desired impedance Z_d [Carignan and Smith, 1994]. Typically, an independent joint PD compensator of the form

$$\begin{aligned} {}^jD(s) &= K_D s + K_P \\ K_P &= \begin{bmatrix} \omega^2 M_{11} & 0 \\ 0 & \omega^2 M_{22} \end{bmatrix} \\ K_D &= \begin{bmatrix} 2\zeta\omega M_{11} & 0 \\ 0 & 2\zeta\omega M_{22} \end{bmatrix} \end{aligned} \quad (20)$$

is chosen for the inner loop where ζ is the damping ratio and ω is the control bandwidth in rad/s. In order for this controller to perform as designed, $D \gg C$ and $D \gg Z_h$ must be maintained through high values of the PD gains. K_P and K_D are typically limited by stability considerations which are primarily a function of the sampling rate and flexibility of the drive system.

One of the key features of this form of admittance control is that the computation can be divided into two loops: the inner joint servo loop, which is typically performed by motor controller chips at very high speeds (≥ 500 Hz), and the outer force loop which runs at a much slower rate, perhaps less than 100 Hz. The effects of friction, backlash, and torque ripple are inversely proportional to the inner loop gains which are quite high because of the fast loop rate. The outer loop, which must perform complex inverse kinematics computations, usually runs much slower than the joint servo loop. However, experimental evidence indicates that the outer loop can run 5-10 times slower than the inner loop with little sacrifice in performance [Maples and Becker, 1986; Pelletier and Doyon, 1994; Carignan and Akin, 1997a].

4 Impedance Error Measures

The goal of the haptic controller is to produce a closed-loop impedance at the gripper (tool) interface which matches the desired impedance, Z_d . The difference between the desired impedance and the closed-loop impedance, $Z_e \equiv Z_d - Z_{hCL}$, represents the impedance ‘‘error’’ at the haptic interface. In the multidimensional case, it will be necessary to develop a measure of the ‘‘size’’ of the error matrix. Here, a scalar performance measure of the impedance error will be adopted from a measure developed by Lawrence [1990] in assessing the effects of computational delays on impedance accuracy.

Constant values of the quadratic formed by the inverse impedance error, $x^T Z_e^{-1} x \equiv c^P$, describe ellipsoids in P -dimensional Cartesian space. This is illustrated for the two-dimensional case in Figure 10. The greater the ‘‘size’’ of Z_e , the larger the ellipse. The semimajor axes of the ellipse are given by $c\sqrt{\lambda_i}$, where λ_i are the eigenvalues of Z_e .

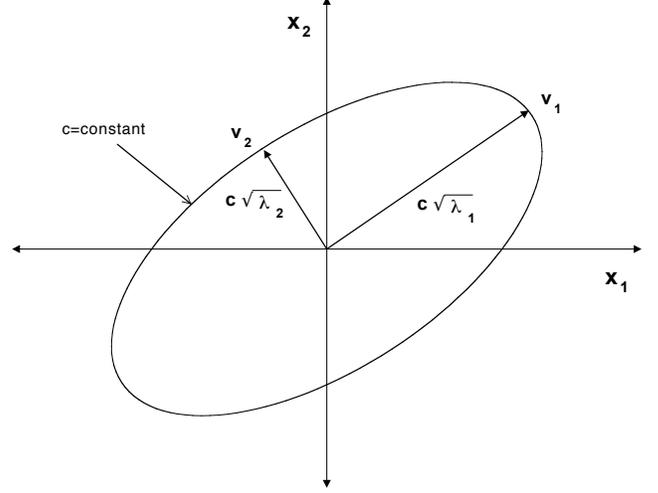


Figure 10. Impedance error ellipse in Cartesian-space.

One possible measure of the size of Z_e is based on the volume of the ellipsoid

$$\|Z_e\| = c \left(\prod_{i=1}^P \sqrt{\lambda_i} \right)^{\frac{1}{P}} \quad (21)$$

where the P^{th} root has been used to retain the dimensions of Z_e . Another possible measure is the average value of the semimajor axes

$$\|Z_e\| = \frac{c}{P} \sum_{i=1}^P \sqrt{\lambda_i} \quad (22)$$

This gives a more conservative measure of the size of Z_e than the volume which would collapse to zero if the error in one of the directions became zero. It was found that the impedance errors in the case of a desired impedance which is isotropic were rather insensitive to the choice of $\|Z_e\|$ as long as the units of Z_e were retained. The more conservative measure (22) was therefore adopted in this work.

The error would be more meaningful if it were somehow normalized by the desired impedance, Z_d . If the same measure is used for Z_d as for Z_e , then a normalized measure of the impedance error is

$$\Delta Z = \frac{\sum_{i=1}^P \sqrt{\lambda_i(Z_e)}}{\sum_{i=1}^P \sqrt{\lambda_i(Z_d)}} \quad (23)$$

Thus (23) represents a scalar measure of the ‘‘relative’’ impedance error. A more conservative measure could be adopted which normalizes $\|Z_e\|$ by the desired impedance along the semiminor axis of Z_d [Lawrence, 1990]. However, (23) represents a balanced approach which weighs the impedance error over all the principal directions.

5 Example

The four-bar planar haptic device shown in Figure 11 will be used to illustrate the differences in impedance accuracy obtained using the various control approaches examined in this paper. First, the design of the haptic device is described. A surgical simulation task is then proposed to illustrate the application of the controllers. Finally, the sensitivities to various parameters are investigated using impedance analysis.

5.1 Haptic Device Design

The geometry of the haptic device is based on the 4-bar linkage design in the CADV1 force reflecting interface [Huissoon and Wang, 1991; Ching and Wang, 1997]. The links in this example are constructed of 2.5 cm square titanium alloy tubes with a 0.5 cm wall thickness. The motors have a continuous output torque of 2 N-m and are mounted at the base to reduce the mass of the links. The rotor inertia is $I_r = 1.5 \times 10^{-4}$ kg-m², and the motor viscous friction is $B_r = 3.5 \times 10^{-4}$ N-m/rad/s. The output of the motors is geared 5:1 to produce a force at the gripper of approximately 16.7 N at full extension. The link parameters are summarized in Table 1.

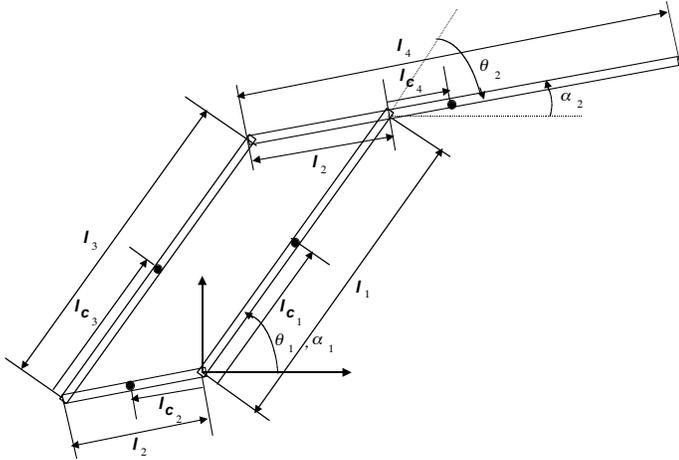


Figure 11. Parallel haptic device used in example.

Since this is a parallel drive mechanism, the dynamics are formulated in the actuator coordinates, α , rather than joint coordinates, θ . However, the analysis used previously for joint space applies equally well here. The position of the tool tip in actuator coordinates is given by

$$x(\alpha) = \begin{bmatrix} l_1 c_1 + (l_4 - l_2) c_2 \\ l_1 s_1 + (l_4 - l_2) s_2 \end{bmatrix} \quad (24)$$

Table 1. Link parameters for parallel haptic device.

link #	length l_i, cm	c. m. l_{c_i}, cm	mass m_i, kg
1	0.300	0.150	0.5352
2	0.150	0.075	0.2676
3	0.300	0.150	0.5352
4	0.450	0.075	0.8028

where $c_i \equiv \cos(\alpha_i)$ and $s_i \equiv \sin(\alpha_i)$. The Jacobian, defined by $J \equiv \frac{\partial x}{\partial \alpha}$, is given by

$$J(\alpha) = \begin{bmatrix} -l_1 s_1 & -(l_4 - l_2) s_2 \\ l_1 c_1 & (l_4 - l_2) c_2 \end{bmatrix} \quad (25)$$

The dynamics, excluding the Coriolis and centripetal force terms, are given by [Asada and Youcef-Toumi, 1987]

$$\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \ddot{\alpha} + \begin{bmatrix} v_1 & 0 \\ 0 & v_2 \end{bmatrix} \dot{\alpha} + \begin{bmatrix} G_1 \\ G_2 \end{bmatrix} = \begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} \quad (26)$$

$$\begin{aligned} \text{where } M_{11} &= I_1 + I_3 + m_4 l_1^2 + I_{m_1} \\ M_{22} &= I_2 + I_4 + m_3 l_2^2 + I_{m_2} \\ M_{12} &= M_{21} = (m_4 l_1 l_{c_4} - m_3 l_2 l_{c_3}) \cos \alpha_2 \\ G_1 &= g(m_1 l_{c_1} + m_3 l_{c_3} + m_4 l_1) \cos \alpha_1 \\ G_2 &= g(m_2 l_{c_2} - m_3 l_2 + m_4 l_{c_4}) \cos \alpha_2 \end{aligned}$$

and v_i is the viscous friction coefficient for actuator i . The inertia of link i about its joint axis, I_i , is found using the parallel axis theorem to be $I_i = \frac{1}{12} m_i l_i^2 + m_i l_{c_i}^2$, and the motor inertia at the output is given by $I_{m_i} = \eta_i^2 I_{r_i}$. The apparent mass at the tool tip when links 1 and 3 are vertical is 1.20 kg in the x-direction and 0.40 kg in the y-direction. The largest apparent mass is therefore an order of magnitude greater than that for the PHANTOMTM.

5.2 Task Description

Suppose that the haptic device is used to simulate a lumbar puncture operation for the purpose of administering spinal or epidural anesthesia as described in Bostrom et al. [1993]. In this procedure, the needle is inserted into the CSF fluid area while the surgeon correlates visual feedback from the simulation with tactile feedback from the needle. As shown in Figure 12, the needle must first penetrate the skin and then an area of dense ligament tissue before reaching the fluid filled sac. Once the needle penetrates the skin, it is held in place by these dense ligaments.

Using average values of 3.0 N/cm and 1.5 N/cm/s for the stiffness and viscosity at the interface of the needle with the ligament [Popa and Singh, 1998], the desired impedance objective can be written as

$$Z_d(s) = B_d s + K_d \quad (27)$$

$$K_p = \begin{bmatrix} 300 & 0 \\ 0 & 300 \end{bmatrix} N/m$$

$$K_v = \begin{bmatrix} 150 & 0 \\ 0 & 150 \end{bmatrix} N/m/s$$

Realistically, the stiffness in the direction perpendicular to the needle travel is probably much higher, but isotropic properties will be assumed here for simplicity.

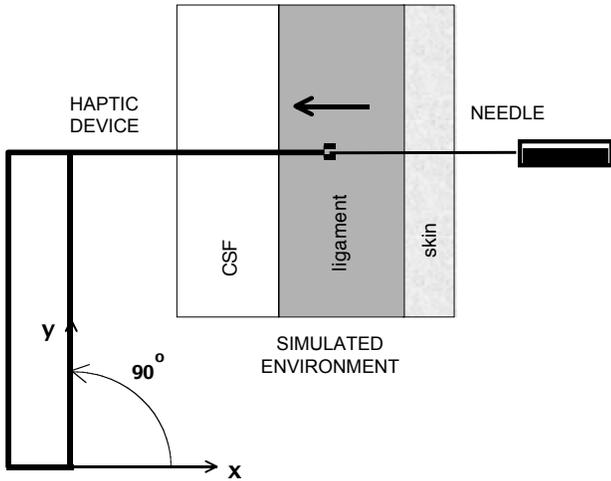


Figure 12. Simulation of lumbar puncture operation.

5.3 Impedance Analysis

The cases which will be compared are the impedance controller with and without force feedback and the admittance controller with position feedback. In all cases, perfect gravity model compensation is assumed. In addition, a nominal velocity of zero is assumed so that the linearized Coriolis and centripetal torques are zero. The joint PD servo for the admittance controller was critically damped.

5.3.1 Force Gain The relative impedance errors for three different levels of force gain in the impedance controller are shown in Figure 13. The error clearly decreases as the level of force feedback increases. In the quasistatic range, even the impedance controller without force feedback is superior to the admittance controller shown here at a bandwidth of 50 Hz. However, the performance of the open-loop impedance controller crosses over the admittance controller at about 0.2 Hz when the inertial forces begin to dominate. Note that even small changes in relative error can be large in absolute terms; a 1% change in relative error represents an absolute error of 3 N/m near dc. Considering that Srinivasan and Chen [1993] recommend a force res-

olution at the haptic interface of 0.01 N, even a deflection as low as 1 cm could lead to a perceivable error in this example.

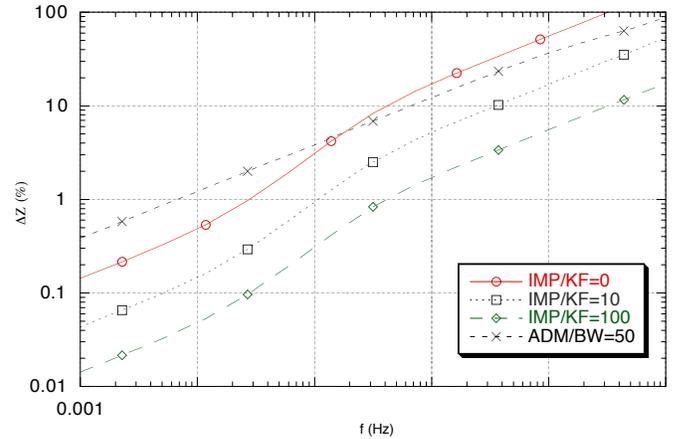


Figure 13. Relative impedance errors (ΔZ) as function of input frequency (f) in nominal case for different levels of force gain [admittance (ADM), impedance (IMP), force gain (KF), bandwidth (BW)].

5.3.2 Joint Servo Control Bandwidth The impedance error as a function of the input frequency and control bandwidth can be represented as a surface in three-dimensional space as shown in Figure 14. The control bandwidth is plotted up to a value of 100 Hz which represents the performance limit of most state-of-the-art haptic devices [Burdea, 1996]. The plot in Figure 13 represents a cross section of this plot at a control bandwidth of 50 Hz. The impedance controller with and without force-feedback is shown for comparison.

The error for the admittance controller is reduced significantly by an increase in bandwidth. The performance crossover boundary between the controllers is the line formed by the intersection of the admittance and impedance control surfaces. The crossover point for a 50 Hz bandwidth is approximately 0.2 Hz and the crossover for the 100 Hz bandwidth is 0.02 Hz. Thus, the performance crossover decreases by a decade with only a doubling of the control bandwidth.

5.3.3 Friction Feedforward Figure 15 shows the impedance errors for the open-loop impedance controller with and without friction compensation. The 50 Hz bandwidth admittance controller is also shown for comparison. The friction compensation clearly decreases the error for the impedance controller in the quasistatic range, but the advantage begins to diminish rapidly at about 0.1 Hz when the inertial effects begin to take over.

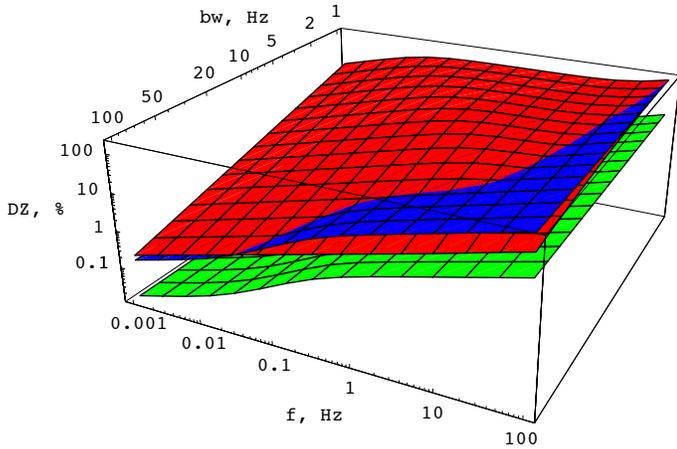


Figure 14. Relative impedance error (DZ) as a function of the inner loop bandwidth (bw) and input frequency (f) [admittance (red), impedance (blue), impedance/ $KF=100$ (green)].

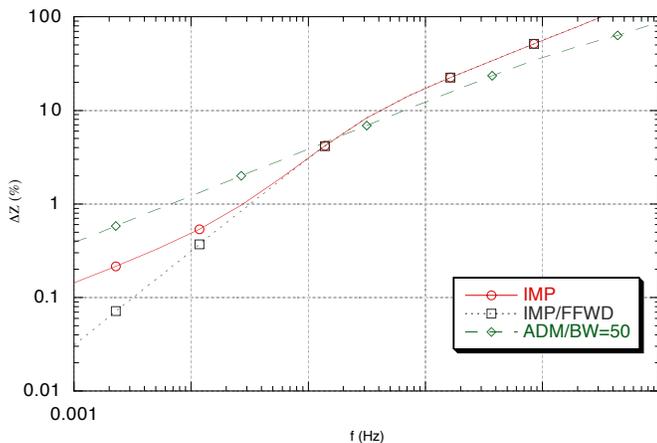


Figure 15. Effect of friction compensation on impedance error (ΔZ) versus input frequency (f) for the open-loop impedance controller [admittance (ADM), impedance (IMP), friction feedforward (FFWD), bandwidth (BW)].

5.3.4 Inertia The effect of inertia was investigated for values both higher and lower than the nominal values. Low values were obtained by decreasing the link cross sectional area from 4 cm^2 to 1 cm^2 and using direct drive. The apparent mass at the haptic interface was reduced to 0.29 kg in the x -direction and 0.09 kg in the y -direction. High values were obtained by increasing the cross section to 16 cm^2 and increasing the gear ratio to $10 : 1$. The apparent mass increased to 4.8 kg in the x -direction and 1.6 kg in the y -direction.

The resulting impedance errors for high and low values of the inertia are illustrated in Figure 16. As expected, the effect of reducing the inertia was to decrease the error in the impedance controller since the error is proportional to the natural dynamics.

However, the effect is to worsen the performance of the admittance controller. This is because the PD loop gains decrease with lower inertia in order to maintain the same bandwidth for the inner PD loop.

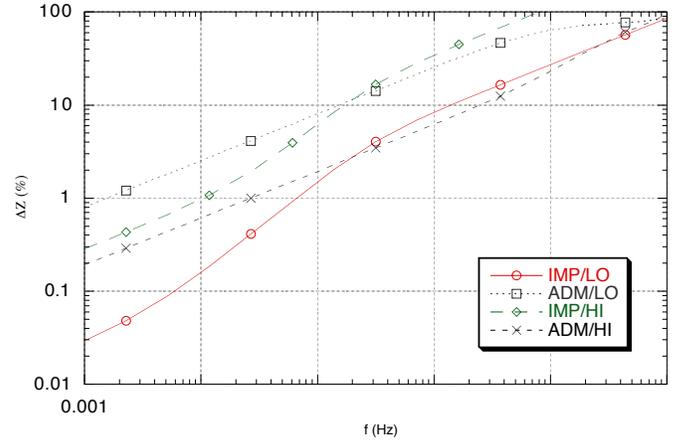


Figure 16. Effect of inertia on the relative impedance error (ΔZ) versus input frequency (f) [50 Hz BW admittance (ADM), open-loop impedance (IMP), high inertia (HI), low inertia (LO)].

5.3.5 Target Impedance The control performance is also strongly dependent upon the desired or “target” impedance. If the target impedance is examined for high and low values, the optimum choice alternates between the two controllers. Figure 17 shows the impedance error when the target impedance is ten times the nominal value. The impedance controller is clearly superior in all cases. When the target impedance is instead reduced by a factor of ten, the admittance controller is superior to the open-loop impedance controller for bandwidths above 10 Hz as shown in Figure 18.

As the desired impedance increases, the effect of the haptic device dynamics in the impedance controller is reduced. This is because the error from (5) becomes smaller relative to the desired impedance. The effect on the admittance controller is the opposite. The impedance compensator design in (16) depends on the inner loop gains in $D(s)$ being much higher than the desired impedance Z_d . Thus, the admittance controller performs better at low target impedances while the impedance controller is superior for high target impedances.

6 Results and Discussion

Some important observations can be made with regard to the results of this example. In the impedance controller, force feedback will always increase the impedance accuracy assuming that stability can be maintained. This is because the impedance error

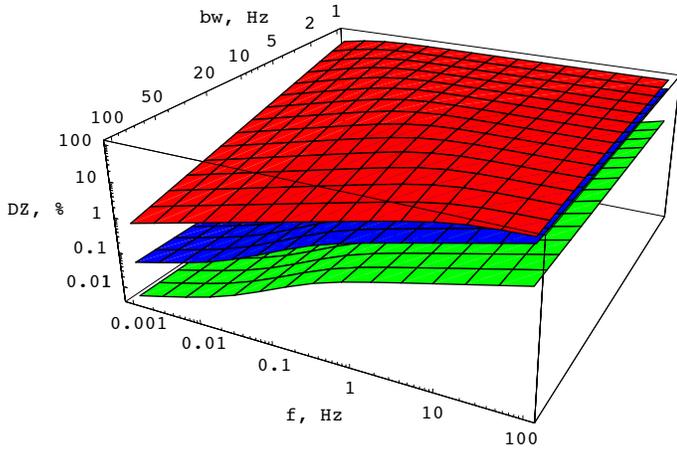


Figure 17. Relative impedance error (DZ) versus input frequency (f) for $K_d = 3000$ N/cm and $B_d = 1500$ N/cm/s [admittance (red), impedance (blue), impedance/KF=100 (green)].

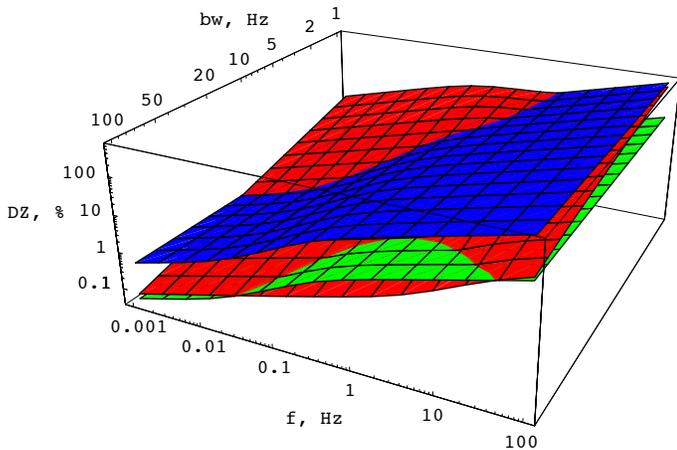


Figure 18. Relative impedance error (DZ) versus input frequency (f) for $K_d = 30$ N/cm and $B_d = 15$ N/cm/s [admittance (red), impedance (blue), impedance/KF=100 (green)].

is inversely proportional to $1 + K_F$. In addition, model feedforward, if exact, will also improve the performance of the controller. However, the range of performance increase for gravity and friction compensation is limited to the quasistatic range.

The caveat on stability for this controller is an important one. While using this form of control to simulate a lumbar puncture procedure, Popa and Singh [1998] noted a significant amount of “noise” in the force output which was attributed to noise in the force sensor. However, it seems more likely that the noise was caused by non-collocation of the force-torque sensor with the joint torque commands which excited unmodeled modes in the arm and the drive train [Eppinger and Seering, 1987]. This is perhaps not as problematic in the admittance controller where

the force signal passes through the impedance compensator prior to being input to the joint servo controller, thus filtering out the high frequency dynamics before they reach the joint level.

The dynamic instability associated with the use of force control in robotics was also cited by An and Hollerbach [1987]. They seem to have been the first to propose a multiloop control architecture which combines wrist and joint force sensing to provide a stable, high-bandwidth force-controlled system. This was accomplished by combining a high bandwidth open-loop joint torque controller at each joint to get stability and fast response with a lower bandwidth outer loop implementing wrist force sensing to improve the accuracy of the force produced at the end-effector. The wrist force signal was modulated by a low-pass filter to avoid exciting the unmodelled dynamics in the arm. This approach is strikingly similar to admittance control with position feedback except that the inner loop is a high gain torque controller rather than a position controller.

Another result from this example is that the performance of the admittance controller improves rapidly with an increase in the control bandwidth of the joint servo loop. As with the force gains in the impedance controller, the PD loop gains are also limited by stability considerations. Because the position feedback is collocated with the torque control, the control bandwidth of this loop is typically much higher (perhaps a decade) than that of a Cartesian loop. This is also why the performance of the admittance controller improves with increasing inertia; larger inertias can support larger inner loop gains thus improving the accuracy of the impedance compensator.

One possible advantage of the admittance controller over the impedance controller is that it relies on high inner loop gains to eradicate the effect of the haptic device dynamics. The impedance controller with model feedforward relies on an accurate model of the dynamics of the haptic device which may fluctuate with operating conditions. The additional computational burden imparted on model-based controllers also lowers the loop rate thereby increasing the discretization errors. The impedance controller with force feedback relies on high force gains to mitigate the effect of the haptic device dynamics. However, the high gain nature of explicit force feedback tends to make these controllers more susceptible to instability.

An important difference between the impedance and the admittance control approaches is the role of the human operator who functions as the “environment” to the haptic device. In the impedance approach, instability is more likely to occur when the operator lets go of the device, imparting zero impedance. The worst case for the admittance controller is when the operator grasps the handle tightly, imparting maximum impedance. This corresponds to the hard contact stability problem in robotics where the coupling of the robot to a stiff environment causes highly undamped poles to occur in the closed-loop system. The compliance of the human hand and arm joints relative to a manipulator render this less likely, although such instability has been reported with a pure admittance controller by Yokoi et al. [1994].

In the event that instability does occur, Adams et al. [1998] propose a virtual coupling network between the haptic display and the virtual environment to guarantee stable interaction.

While not covered in this work, sampling effects are also an important issue in haptic controller design. Colgate and Schenkel [1994] showed that the sample rate for the Cartesian loop imposes upper bounds on the target stiffness and damping values in the impedance controller. There seems to be scant data on the sample rates for the impedance controller when force feedback is present. Bergamasco et al. [1994] suggest rates in the vicinity of 1 KHz although lower values [Popa and Singh, 1998] have been reported. Experimental evidence for the admittance controller indicates low sensitivity of the impedance accuracy to the outer (impedance) loop rate. Rates below 100 Hz have been successfully implemented [Maples and Becker, 1986; Pelletier and Doyon, 1994; Carignan and Akin, 1997a], although it seems likely that the required outer loop rate should increase with the target stiffness, K_d , and decrease with the target damping, C^{-1} .

7 Conclusion

This report began with a review of force-controlled haptic devices and proceeded to quantify their effects at the haptic interface. The impedance error for the open-loop impedance controller was derived first, followed by a model-feedforward implementation. Dual-input impedance and admittance controllers were then examined, and an example was used to illustrate the sensitivities to position bandwidth, force gain, haptic device inertia and friction, and target impedance.

Some strengths and weaknesses of these approaches were wrought out by the analysis. The error for the impedance controller is inversely related to the level of force feedback used. In addition, the relative error increases in proportion to the inertia of the haptic device and decreases in proportion to the target impedance. The error for the admittance controller is reduced by an order of magnitude for a doubling of the joint servo loop bandwidth. The error for this controller grows as the target impedance increases. However, increased joint PD gains achieved through higher inertias in the device reduce the impedance error.

It cannot be concluded from this analysis that one control approach is clearly superior. The impedance controller with force feedback is promising if the stability problems associated with explicit force feedback can be overcome. Admittance control may offer more immunity to this problem because of its filtering effect on the force signal. The disturbance rejection properties of the dual-input admittance controller are attractive, but the need for high joint PD gains contraindicates its use at low values of joint inertia and sample rate. The effect of target impedance is clearly divisive as the admittance controller is superior at low target impedances and the impedance controller is superior at high target impedances. The choice may end up being determined by the environment being simulated as well as the characteristics of the haptic device.

The results in this paper represent only one step toward improving the haptic interface for virtual environments. While further mechanical improvements in haptic technology may yet be realized, it is clear that control design has a significant impact on the quality of the haptic feedback. Whether driven by a need for higher force output or higher fidelity, it seems likely that force feedback controllers will play a vital role in the successful implementation of future haptic interfaces.

Acknowledgements

The authors would like to acknowledge the students of ENAE788Y, "Haptic Controls and Applications", who endured the initial exposure to many of these ideas. We would also like to thank Corinna Lathan of Catholic University whose collaboration on the Spine Biopsy Simulator at Georgetown University helped inspire this work.

REFERENCES

- Adams, R. J. and Hannaford, B. (1999). Stable haptic interaction with virtual environments. *IEEE Trans. on Robotics and Automation*, 15(3):465–474.
- Adams, R. J., Moreyra, M. R., and Hannaford, B. (1998). Stability and performance of haptic displays: Theory and experiments. In *Proc. of the ASME Dynamics Systems and Control Division*, pages 227–234.
- An, C. H. and Hollerbach, J. M. (1987). Dynamic stability issues in force control of manipulators. In *Proc. of the IEEE Conference on Robotics and Automation*, pages 890–896. IEEE.
- Asada, H. and Youcef-Toumi, K. (1987). *Direct-Drive Robots*. MIT Press, Cambridge, Mass.
- Baumann, R. and Clavel, R. (1998). Haptic interface for virtual reality based minimally invasive surgery simulation. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 381–386.
- Bejczy, A. and Salisbury, J. K. (1983). Controlling remote manipulators through kinesthetic coupling. *Computers in Mechanical Engineering*, 2(1):281–290.
- Bergamasco, M., Allotta, B., Bosio, L., Ferreti, L., Parrini, G., Prisco, G. M., Salsedo, F., and Sartini, G. (1994). An arm exoskeleton system for teleoperation and virtual environments applications. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1449–1454.
- Book, W. and Ruis, D. (1981). Control of a robotic exercise machine. In *Proc. of the Joint Automatic Control Conference*, pages WA–2A.
- Bostrom, M., Singh, S., and Wiley, C. (1993). Design of an interactive lumbar puncture simulator with tactile feedback. In *Proc. IEEE Virtual Reality Annual Int. Symposium (VRAIS)*, pages 280–286.
- Burdea, G. C. (1996). *Force And Touch Feedback For Virtual Reality*. John Wiley and Sons, New York.
- Carignan, C. and Akin, D. (1997a). Achieving impedance objec-

- tives in robot teleoperation. In *Proc. of the IEEE Conference on Robotics and Automation*, pages 3487–3492.
- Carignan, C. and Akin, D. (1997b). Actively controlled mockups for EVA training in neutral buoyancy. In *Proc. IEEE Conf. on Systems, Man, and Cybernetics*, pages 2369–2374.
- Carignan, C. and Smith, J. (1994). Manipulator impedance accuracy in position-based impedance control implementations. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1216–1221.
- Ching, M. and Wang, D. (1997). A five-bar-linkage force reflecting interface for a virtual reality system. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 3012–3017.
- Clery, K., Lathan, C., and Carignan, C. (1998). Simulator/planner for CT directed needle biopsy of the spine. In *Surgical-Assist Systems*, pages 218–224. SPIE.
- Clover, C., Luecke, G., Troy, J., and McNeely, W. (1997). Dynamic simulation of virtual mechanisms with haptic feedback using industrial robotics equipment. In *Proc. of the IEEE Conference on Robotics and Automation*, pages 724–730.
- Colgate, J. E. and Schenkel, G. (1994). Passivity of sampled data systems: Application to haptic interfaces. In *Proc. American Control Conf. IEEE*.
- Das, H., Zak, H., Kim, W., Bejczy, A., and Schenker, P. (1992). Operator performance with alternative manual control modes in teleoperation. *Presence-Teleoperators and Virtual Environments*, 1(2):201–218.
- Eppinger, S. D. and Seering, W. P. (1987). Understanding bandwidth limitations in robot force control. In *Proc. of the IEEE Conference on Robotics and Automation*, pages 904–909.
- Glosser, G. D. and Newman, W. S. (1994). The implementation of a natural admittance controller on an industrial robot. In *Proc. of the IEEE Conference on Robotics and Automation*, pages 1209–1215.
- Goertz, R. C. (1954). Manipulator systems development at ANL. In *Proc. 12th Conf. on Remote Systems Technology*, pages 117–136. American Nuclear Society.
- Hannaford, B. (1989). Stability and performance tradeoffs in bilateral telemanipulation. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1764–1767.
- Hogan, N. (1985). Impedance control: An approach to manipulation. *J. of Dynamics Systems, Measurement, and Control*, pages 1–24.
- Huissoon, J. P. and Wang, D. (1991). On the design of a direct drive 5-bar-linkage manipulator. *Robotics*, 9:441–446.
- Lawrence, D. (1988). Impedance control stability properties in common implementations. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1185–1190.
- Lawrence, D. (1990). Measuring the effects of computational delay in manipulator control architectures. University of Cincinnati, Department of Electrical and Computer Engineering.
- Lawrence, D. (1992). Stability and transparency in bilateral teleoperation. In *Proc. IEEE Conf. Decision and Control*.
- Lee, W.-S., Kim, J.-H., and Cho, J.-H. (1998). A driving simulator as a virtual reality tool. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 71–76.
- Li, P. and Horowitz, R. (1996). Control of smart exercise machines. In *IFAC World Congress*.
- Liu, K.-C. (1997). Experimental evaluation of adaptive neuro-controllers for a prototype space robotic manipulator, AIAA 97-0004. In *35th Aerospace Sciences Meeting and Exhibit*.
- Maples, J. and Becker, J. (1986). Experiments in force control of robotic manipulators. In *Proc. of the IEEE Conference on Robotics and Automation*, pages 695–702.
- Massie, T. and Salisbury, J. K. (1994). The PHANTOM haptic interface: A device for probing virtual objects. In *Proc. ASME Winter Annual Meeting: Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*.
- Newman, W. S. (1990). Stability and performance limits of interaction controllers. Technical Report 90-144, Case Western Reserve University, Cleveland, Ohio.
- Paines, J. (1987). Optimization of manual control dynamics for space telemanipulation: Impedance control of a force-reflecting hand controller. Master's thesis, M.I.T., Cambridge, Mass.
- Pelletier, M. and Doyon, M. (1994). On the implementation and performance of impedance control on position controlled robots. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 1228–1233.
- Popa, D. O. and Singh, S. K. (1998). Creating realistic force sensations in a virtual environment: Experimental system, fundamental issues and results. In *Proc. IEEE Int. Conf. on Robotics and Automation*, pages 59–64.
- Sakaki, T., Iwakane, T., and Tachi, S. (1992). Impedance control of a robot manipulator stabilized by PID endpoint motion compensator. In *Proc. Second Int. Symposium on Measurement and Control in Robotics (ISMCR)*.
- Srinivasan, M. and Chen, J. (1993). Human performance in controlling normal forces of contact with rigid objects. In *Advances in Robotics, Mechantronics, and Haptic Interfaces, DSC-49*, pages 119–125. ASME.
- Swaim, P., Thompson, C., and Campbell, P. (1995). The Charlotte intra-vehicular robot. Technical Report N95-23703, NASA.
- Yokoi, H., Yamashita, J., Fukui, Y., and Shimojo, M. (1994). Development of the virtual shape manipulating system. In *Proc. of the 4th Int. Conf. on Artificial Reality and Tele-Existence (ICAT'94)*, pages 43–48.
- Zhai, S. and Milgram, P. (1993). Human performance evaluation of manipulation scheme in virtual environments. In *IEEE Virtual Reality Annual Int. Symposium (VRAIS)*, pages 155–161.