

Haptic Interface Transparency Achieved Through Viscous Coupling*

Abdenbi Mohand-Ousaid, Guillaume Millet, Stéphane Régnier,
Sinan Haliyo, and Vincent Hayward

UPMC Univ Paris 06, UMR 7222, Institut des Systèmes Intelligents et de Robotique, Paris, France.

E-mail: vincent.hayward@isir.upmc.fr

Abstract

Electromagnetic drives are subjected to an inherent inertia-torque tradeoff that fundamentally limits transparency: the higher the torque, the higher the inertia. We describe a dual-stage design that is not subjected to this tradeoff and that is able to approach perfect transparency for human users. It comprises a large, proximal motor and a small, distal motor to reproduce the transients. The two stages are coupled by a viscous clutch based on eddy-currents that, without contact, accurately transforms slip velocity into torque. Such a system can, in general, be controlled to achieve a variety of objectives. Here, we show that an advanced, discrete-time, RST polynomial pole-placement controller can achieve near-perfect transparency. Experimental validation evaluated the human ability to detect small haptic details when using this drive and compared it to when using a conventional, single-motor interface.

KEY WORDS—Electromagnetic drives, discrete-domain pole placement, haptic interfaces

1 Introduction

The ideal haptic device has no mass, infinite bandwidth and can supply unlimited force or torque. In haptic device design, the actuator, or ‘prime mover’, is often a limiting element, followed by mechanical transmissions and structural elements which all restrict the range of signals that can be transmitted to a user (Morrell and Salisbury, 1997; Hayward and Astley, 1996).

Given a particular actuator, no matter how well transmissions and linkages are designed, these elements can only degrade key aspects of performance that include the dynamic range, i.e. the ratio of the largest to the smallest specifiable force, the stiffness linking the prime mover to the manipulandum, or the end-point inertia. In order to cope with these practi-

cal limitations, human sensorimotor performance can be used as a factor to bound the desired performance.

One aspect of human performance that is relevant to device design is the smallest detectable force, since, ideally, this quantity should match the smallest force than can be commanded by the device. At the upper end of the range, the largest force should match the motoric output of the hand. Ideally, the dynamic range of the device should be achievable over the whole frequency range available to touch, that is, from DC to, say, 1 kHz.

While actuator saturation (short and long term) is the factor that determines the upper limit of the range, the lower limit merits discussion. When a manipulandum interacts with a hand, it is subject to several forces.¹ Neglecting the effect of internal elastic forces, that is, if the device operates below its first resonant mode, these forces are:

- a, the actuating force, typically a Laplace force developed in the motor windings;
- b, forces due to induced currents which oppose motion;
- c, forces due to mechanical losses: viscosity, friction;
- d, inertial forces to which moving parts are subjected;
- e, the force applied by the hand which include at least an inertial component, a viscous component, and an elastic component, due to the movement of tissues.

According to Newton’s second law, all these forces must balance. In the above list, it is well known to haptic device designers that all these factors play a significant role in the total result. To illustrate the tradeoffs involved, direct-drive, friction-free floater designs pay the price for nearly absent losses by large end-point inertias due to necessity to actuate against gravity (Berkelman and Hollis, 2000).

*The final, definitive version of this paper published by SAGE Publications Ltd, All rights reserved. ©V. Hayward

¹for brevity, no distinction is made between forces and torques.

We therefore desire a transducer able to command a desired force at the interface between a manipulator and the soft load represented by a finger that can minimize all parasitic forces. Informally, an ideally transparent device should be such that $\mathbf{a} \approx \mathbf{e}$ and that \mathbf{b} , \mathbf{c} , and \mathbf{d} in the above list can be neglected. An estimate, developed further later, shows that the end-point inertia should be on the order of grams, not hundreds of grams, which is the standard in today's devices. We also desire that this force be accurately produced, not only at DC, but over a wide frequency range, up to one kilohertz.

In this article, we first develop target performance figures that result from the above considerations, and propose a dual-stage drive architecture that has the potential to meet these targets and we describe the design of its key elements. We then describe a control design performed in the sampled-data domain in order to maximize accuracy. Lastly, we tested human performance gained when using this device. To this end, we asked participants to perform a foreground-background signal detection task and compare the results with those obtained with a single-stage design.

2 Haptic Transparency

Achieving haptic transparency depends on reducing the parasitic forces under the smallest human-detectable force under all desired operating conditions. Seen from the viewpoint of the hand, we desire the device to have a mechanical impedance (the ratio of force over displacement and derivatives) that is significantly smaller than that of the fingertips since then, the deformation of the tissues would not be the result of parasitic forces, but predominantly from commanded forces.

Some target values were suggested by Millet et al. (2009) to specify the characteristics of a haptic interface that can operate at the limits of human performance, see also (Baud-Bovy and Gatti, 2010). These numbers, reported in Table 1, are arrived at by considering that when the hand actively interacts with the handle, parasitic forces should be at threshold, thus not interfering with the Laplace force produced by the motors.

It is worth noting that the dynamic range from the lowest to the highest specifiable force covers four orders of magnitude and that the achievable dynamic range of a high-quality DC motor is only about two orders of magnitude. It is clear that a single standard electric motor, the actuator of choice when building impedance-based haptic devices, cannot meet our objectives.

Table 1: Target performance figures with a 70 mm diameter handle (Millet et al., 2009).

Quantity	Value	Equiv. angular Value
Mass ^a	$0.5 \cdot 10^{-2}$ kg	$6.1 \cdot 10^{-6}$ kg·m ²
Viscosity ^b	$0.5 \cdot 10^{-1}$ N·s/m	$6.1 \cdot 10^{-5}$ N·m·s/rad
Friction ^c	$1.0 \cdot 10^{-3}$ N	$3.5 \cdot 10^{-5}$ N·m
Max. force ^d	5.0 N	$1.8 \cdot 10^{-1}$ N·m
Rise-time ^e	$1.0 \cdot 10^{-3}$ s	

^a from largest acceleration and lowest statically detectable force

^b from largest velocity and lowest statically detectable force

^c from lowest dynamically detectable force (reversal transients)

^d obtained from common sense

^e obtained from vibrotactile performance

2.1 Related Approaches

Perhaps the oldest motivation for dual-motor designs is the desire to modulate the intrinsic dynamics of an actuated joint to imitate the antagonist action of skeletal muscles. Sugano et al. (1992); English and Russell (1999); Stramigioli and Duindam (2001) describe efforts aimed at modulating elasticity with electric motors, and Boulet (1990) describes the modulation of viscosity with hydraulic motors. These ideas have evolved to be included in the design of robots that can interact with people, see (Bicchi et al., 2008).

In the symbolism of Fig. 1, a rectangular box represents an inertia, two parallel lines represent a controllable source of force, typically a Laplace force, a coil represents a spring, a piston represents a damper, a small circle represents the load, and a mechanical connection represents a common velocity on each side. Lack of space prevents us from commenting on the various control options afforded by each design arrangement.




Figure 1: Several ways one or two motors can be coupled to a load. See text for discussion. Combination **j** is adopted for the device described in the present article.

The first case, **a**, is important because it models most actual haptic devices, since the preferred method of construction calls for grounding the actuators. There is an elastic element between the motor and the handle. If elasticity is built-in by design, we have the so-called “series elastic actuators” intended to be employed in closed-loop (Pratt and Williamson, 1995; Lauria et al., 2008). In haptics, one would rather make the transmission as stiff as possible to raise the system’s natural frequency.

The next configuration, **b**, corresponds to the standard, variously called coarse-fine or mini-macro manipulator design (Sharon and Hardt, 1984; Sharon et al., 1988; Salcudean and An, 1989; Khatib, 1990; Hollis et al., 1991). Here, a large motor ‘carries’ a smaller one, but both share the same load since they are on the same load path. Such a scheme has been advocated for haptic devices (Stocco and Salcudean, 1996; Wall and Harwin, 2001). With proper control, this configuration can reduce the apparent inertia of the whole system but the small motor must be able to bear the whole brunt of the load. Addressing this limitation naturally leads to an arrangement, **c**, described by Morrell and Salisbury (1997), and applied to making human-friendly robots and force feedback devices (Zinn et al., 2004, 2008).

The large motor supplies the largest portion of the force and the small motor “fills-in” during transients. With proper control, the user is exposed to the inertia of the small motor but not that of the large. A variant design is proposed in (Conti et al., 2007), where the large motor is replaced by a passive brake.

More generally speaking, the haptic designer could use all three types of mechanical coupling: elastic, dissipative, and inertial, suggesting configuration **d**. In fact, it is by inertial coupling that portable phones, gaming pads or more advanced transducers operate (Yao and Hayward, 2010).

Option **e** can achieve an effect similar to that of arrangement **c** but introduces another tradeoff if the displacement of the distal motor is limited, since inertial coupling establishes a quadratic dependency between this displacement and the lowest frequency at which it can operate.

Dissipative forces can also be used to achieve desirable couplings. The counterpart of case **a**, case **f**, is described by Chew et al. (2004) for robotic applications. It is not appealing for haptics since coupled to elastic load like a hand, this coupling will introduce an attenuation of 6 dB/octave in addition to the 12 dB/octave attenuation due to the motor inertia.

We could then think of yet another option, **g**, to ‘take up’ the response in the high-frequency, but similarly to **b**, this configuration would solve this problem by creating another. Option **h** is worth mentioning

since it can be used to generate distortion-free haptic information up to 100 Hz but does not operate at DC (Campion et al., 2008). Arrangement **i** would suffer from the same difficulties as **e**.

This enumeration naturally leads us to the last design, **j**, that could compete with the elastic coupling of configuration **c**. A comparison between the two options is carried out in the next section.

2.2 Discussion

Elastic couplings have the property of storing energy. This property is an advantage in applications where energy storage and restitution is a desirable property, such as in a walking machine, but in haptics, energy storage is inconvenient, even hazardous.

An elastically coupled drive system must be wound-up to output a torque. If the output torque needs to be brought to zero abruptly, for instance when breaking contact with a virtual wall or when letting go of the handle during interaction, the limited speed of coupler unwinding forces a delay in the regulation of the output torque since this torque is two integrals away from the displacement of the motor. Moreover, the stored elastic energy could be uncontrollably converted to output kinetic energy, causing a hazard.

With viscous coupling the proximal stage must spin, also storing energy but in kinetic form. This energy cannot be transferred at the output transiently since the coupling naturally dissipates energy in itself, rather than in the user, without any need for active control. In contrast with the elastic coupling case, the energy storing quantity—velocity—is only one integral away from the output torque.

The basic servo mechanism that must govern the proximal stage is velocity control, which is easier to synthesize than a position servo since the system has one fewer order. The proximal stage can react more quickly than in the case of an elastic coupling, reducing the reliance on the distal stage to produce an accurate output torque.

There are other advantages of the viscous coupling over the elastic coupling, chiefly among them is accuracy. Dual-stage designs produce accurate output torques because, in essence, they are force-feedback systems. In the elastically coupled case, the feedback control system must regulate the spring deflection. If the elastic element is not of stringent quality, there could be hysteresis and high-frequency modes. More generally speaking, viscous coupling enjoys the general advantage of precluding the occurrence of structural dynamics present in any mechanical system having elastic elements.

With viscous coupling the onus is also on the coupler to provide an accurate relationship between dis-

placement and torque. Due to the absence of contact, the main source of the injurious properties of mechanical systems, couplers based on eddy-current induction create resistive torques that are *exactly* proportional to the relative slip velocity. At very high velocities, the viscous coefficient decays, see the discussion in Section 3.2.1.

3 Prototype Drive

We now describe the system depicted in Fig. 2. The objective was to prototype a haptic interface that is such that, seen from the handle, the system's apparent mechanical impedance is small, but the torque that it can deliver is large. The result is a very high-performance system that closely meets the goals listed in Table 1 and thus is close to achieve perfect transparency.




Figure 2: Prototype realized according to strategy **j**.

3.1 Design Parameters

The quantities that participate in the performance are:

- J_1 : the inertia connected to the proximal shaft that includes the proximal rotor and the clutch inductor;
- J_2 : the inertia related to the distal axis that includes the rotor of the distal motor, the clutch armature, and the handle;
- b : the coefficient of viscosity of the clutch.

Rapid changes in torque benefit from a short time constant for the proximal stage. This stage can be represented by a first-order system with a time constant equal to the ratio of its inertia to the viscosity of the clutch. Reducing inertia and increasing viscosity improves the capacity of the proximal stage to effect rapid torque changes without the help of the distal stage.

Achieving precise output torque requires the perturbing torque caused by friction to be minimized. In open loop, the inertias of both stages should be minimized. In closed loop, however, the inertia of the proximal axis can be hidden, since the energy needed to back-drive the system is supplied by an external source—the electronic amplifier, and not by the load.

The control is based on the estimation of the relative speed of the two shafts and can suffer from quantization effects. For a given torque transmitted by the clutch and a given quantization resolution, a higher viscosity decreases relative speed and thus the signal/noise ratio. These conflicting requirements require the use of high-resolution encoders or of a state observer.

In all cases, a design generally benefits from minimizing the inertia of the distal stage, J_2 , as well as minimizing the time constant of the proximal stage, J_1/b .

3.2 Analysis

3.2.1 Design Criteria Approximations

The quantities J_2 and J_1/b are affected by the mechanical and magnetic characteristics of the coupler. An eddy-current coupler has more parameters to optimize than a brake. There are two rotors: a permanent-magnet inductor and a disc armature. A multiple-loop magnetic circuit was designed according to the guidelines found in (Gosline and Hayward, 2008) and from the results of the prototype described in (Millet et al., 2009). Figures 3a and 3b illustrate the adopted geometry. As seen in Fig. 3b, a multiplicity of closed magnetic circuits create a periodic magnetic flux that crosses the annular armature mostly orthogonally. As a result, relative movement induces eddy-current loops as represented in Fig. 3c.

The analysis of the two optimization criteria, J_2 and J_1/b , provided insight regarding the critical design parameters. The geometric parameters are defined in Fig. 3b and 3c. The first criterion, J_2 , is the sum of the inertias of the armature, the rotor of the distal motor, and the handle. Provided that the width of the armature w_a is much smaller than its mean radius r , to good approximation,

$$J_2 \approx 2\pi\rho_a t_a w_a r^3 + J_{r2}, \quad (1)$$

where ρ_a is the mass density of the armature, t_a is its thickness, and J_{r2} is the inertia of the distal rotor plus that of the handle. The mean radius of the clutch cubically scales the inertia of the armature, so it should be kept small. In our prototype, the handle accounts for about two thirds of J_2 and the armature for one third.

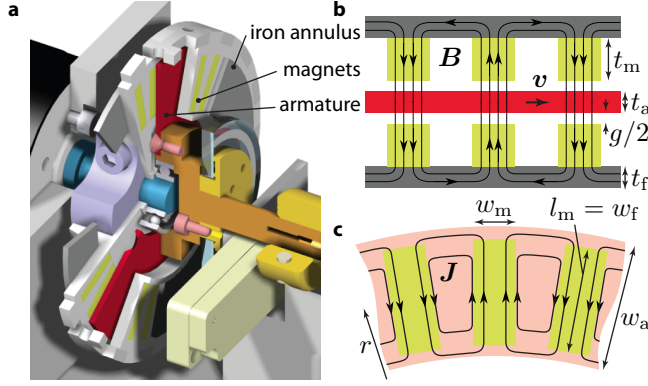


Figure 3: Eddy-current viscous coupler. **a**, Cutaway view of the coupler. **b**, Magnetic flux density streamlines generated by the magnets. **c**, Current density streamlines in the mid-plane of the armature during relative movement.

The second criteria to minimize is the time constant of the proximal stage J_1/b where $2n$ magnets, two iron annuli, the rotor of the proximal motor, and the magnet holder participate in the moment of inertia. Considering that the width of the annuli, w_f , and the length of the magnets, l_m , are much smaller than the mean radius, r , the inertia is approximated by

$$J_1 \approx 2nl_m w_m t_m r^2 \rho_m + 4\pi w_f t_f r^3 \rho_f + J_{r1} \quad (2)$$

where ρ_m and ρ_f represent the mass densities of the magnets and of the annuli, t_m and t_f their thicknesses, respectively, and J_{r1} is the inertia of the proximal rotor plus that of the magnet support.

The modeling of the magnetic field and induction in complex geometries is difficult because it involves several nonlinear phenomena (Balakrishnan et al., 1997; Heald, 1988; Srivastava and Kumar, 2009). Nevertheless, several simplifying assumptions can be made to find a closed-form expression that estimates the viscosity of the clutch. When the armature rotates relatively to the magnets, the magnetic field, \mathbf{B} , induces an electrical field, $\mathbf{E} = \mathbf{v} \times \mathbf{B}$, orthogonal to \mathbf{B} and to the linear velocity \mathbf{v} . Following the reasoning of Wouterse (1991), an approach to estimate the current density induced in the armature is to suppose that the return path around the magnet is of negligible resistance. The current density induced under magnets is then $\mathbf{J} = \sigma_a \mathbf{v} \times \mathbf{B}$, where σ_a is the armature's conductivity. Assuming that \mathbf{B} is uniform through the volume of the armature covered by magnets and is sufficiently greater than the magnetic field induced by the eddy currents, the power dissipated by the eddy currents can be computed by integrating the currents over the volume of the armature that is

covered by magnets,

$$P_{\text{diss}} = \frac{1}{\sigma_a} \int |\mathbf{J}|^2 dV \approx n\sigma_a t_a l_m w_m r^2 \mathbf{B}_m^2 \omega^2, \quad (3)$$

where n is the number of magnetic loops and \mathbf{B}_m the magnetic field produced by the magnets through the armature. The coefficient of viscosity, b , is then approximated by

$$b \approx n\sigma_a t_a l_m w_m r^2 \mathbf{B}_m^2. \quad (4)$$

That the viscosity coefficient is independent of the velocity follows from the assumption that the magnetic field induced by eddy current can be neglected compared to \mathbf{B}_m , which is borne out at low speeds (Anwar, 2004). For the present prototype, this assumption breaks down at high torques, that is, at speeds higher than 2000 rpm where b would probably noticeably begin to decrease. This effect, however, is of no consequence to our main objective. We can now evaluate the proximal stage time constant

$$\frac{J_1}{b} \approx \frac{2nl_m w_m t_m r^2 \rho_m + 4\pi l_m t_f r^3 \rho_f + J_{r1}}{n\sigma_a l_m w_m t_a r^2 \mathbf{B}_m^2}. \quad (5)$$

From this expression, it follows that stronger magnets, higher conductivity, and a lower mass density in the armature all enhance performance, which is obvious. Finding the optimal number of magnets and their geometry, however, is not trivial.

Similar issues have been discussed for the design of magnetic undulators used in free-electron lasers and synchrotron radiation facilities. Literature from this topic provides some insight (Quimby and Pindroh, 1987; Elleaume et al., 2000), but the tradeoff with inertia is not considered. These considerations justify a multi-physics numerical design optimization approach.

3.2.2 Numerical Simulation and Optimization

A 3D simulation was set up with a multiphysics simulation package (COMSOL AB, Stockholm, Sweden), a finite element analysis simulator suitable to simulate eddy currents induced in a moving conductor. The two symmetries (mid-plane and n angular anti-periodicity) in the problem were used to reduce computation time. By establishing anti-periodic boundary conditions at the mid-planes of two consecutive magnets, the simulation of one of the $2n$ anti-periodic sectors was sufficient.

The mesh resolution was set to 0.3 mm for the regions of the armature shadowed by the magnets and coarser elsewhere. An important criterion for selecting the mesh resolution was to mesh the boundary of

the annulus with element sizes sufficiently small to account for the skin effect. An ungauged potential formulation was used to solve the electromagnetic problem. Typically, convergence was achieved in 15 steps for a relative precision of 10^{-6} .

The skin depth was about 12 mm for a relative speed of 21 rad/s, hence invading the whole armature at higher speeds. This effect is visible in Figure 4 that shows the current density induced in the armature. Density is slightly smaller in the mid-plane of the armature. Figure 4c shows the results of the simulation for a relative speed of 21 rad/s, generating a braking torque of 50 mN·m.




Figure 4: **a**, Current density induced in a 32 magnet pair armature rotating at 21 rad/s. **b**, Zoom on one sector. **c**, Electromagnetic simulation of a half magnetic loop.

The influence of the thickness and width of the armature, t_a and w_a , the number of pairs of magnets n , and the gap, g , between the armature and magnets were studied. Other parameters were fixed. The mean radius was $r = 20$ mm, the magnet dimensions $2 \times 2 \times 10$ mm, their remanent field $B_m = 1.37$ T. The geometry of the iron annuli was constrained by the magnets and the need to prevent magnetic saturation.

Figure 5 shows that increasing the number of magnets increases the total torque until a certain limit where magnets become too close, causing magnetic leakage. The reduction of the inertia of the proximal stage J_1 conflicts with the optimal number of mag-

nets. The best tradeoff for the present design corresponds to a number of magnets pairs between 32 and 36 and an armature thickness of 1.0–1.5 mm.




Figure 5: Viscosity coefficient and time constant of the proximal stage axis according to the armature thickness and the number of magnets. Other parameters were fixed, $w_a = 10$ mm and $g = 1$ mm.

Figure 6 shows that selecting the width of the armature to be slightly longer than the length of magnets significantly increase the viscosity coefficient, while not affecting the inertia of the distal axis significantly. This result agrees with experimental observations found in (Gosline and Hayward, 2008). Lastly, Figure 6 shows that the viscosity coefficient is strongly influenced by the gap which should be kept as small as possible, within the limits of manufacturing feasibility.




Figure 6: Influence of the gap and the width of the armature on the viscosity coefficient.

3.3 Resulting Prototype

The system seen in Figure 2 was manufactured using a combination of conventional machining and rapid prototyping (ABS plastic), which was effective for those parts which had to be lightweight. Some structural engineering was carried out to ensure that all the parts had the required rigidity.

The motors were a 35 mm core-less DC motor (model RE-35, 273754, graphite brushes, Maxon Motors AG, Sachseln, Switzerland) for the proximal stage and a small 16 mm motor (model RE-16, 118698, precious metal brushes) with a peak torque approximately 20 times smaller than the larger

motor. The motors were powered by two analog current/voltage amplifiers (model LCAM, Quanser, Markham, ON, Canada). The two shafts were fitted with high-resolution digital encoders (Mercury M1800, MicroE Systems, Bedford, MA, USA) giving 1,638,400 counts per turn. The clutch had 60 neodymium magnets, and the armature parameters were $t_a = 1$ mm, $r = 20$ mm, $w_a = 12$ mm.

It is of paramount importance to have accurate estimates of the system parameters. They were measured as described in Millet et al. (2009) and the results collected in Table 2. The viscosity coefficient was a bit lower than expected from the simulation because of the lack of dimensional accuracy of the rapid prototyping manufacturing process forcing us to increase the gap.

Table 2: Measured system parameters.

Quantity		Value
distal torque constant	k_2	$7.9 \cdot 10^{-3}$ N·m/A
viscous coefficient	b	$8.7 \cdot 10^{-4}$ N·m·s/rad
proximal inertia	J_1	$2.6 \cdot 10^{-5}$ kg·m ²
distal inertia	J_2	$6.4 \cdot 10^{-6}$ kg·m ²
distal dry friction	$C_{2\text{fric}}$	$1.3 \cdot 10^{-4}$ N·m
max torque	$C_{1\text{max}}$	$2.0 \cdot 10^{-1}$ N·m

With proper control, the system matched closely the entire range of human sensorimotor performance, or “perfect transparency” as quantified in Section 2.

4 Control System Design

Referring to Fig. 7, the system has two torque commands and one torque output, which can potentially provides for a number of control options. Here, the primary control objective is output torque regulation and tracking. One effect of this particular control objective is to disconnect the effects of the proximal motor dynamics from the experience of the user. During active exploration, when the handle is back-driven, the power required to move the proximal motor is entirely supplied by the power amplifier, not by the user.

To achieve this objective, the reference torque, τ_d , is compared to the torque produced by the coupler, which is proportional to relative velocity between the proximal and distal motors. The proximal motor is slaved by the compensator $C(z)$ to the torque error, operating in the sampled-data domain with period h . The distal motor is commanded to reduce the error without compensation since it is needed to transiently compensate for the slower response of the proximal stage. More complex schemes to achieve a wider range

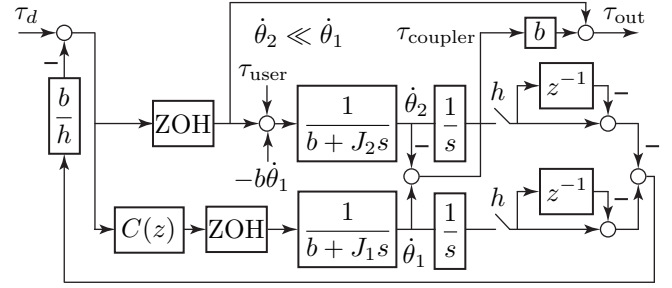


Figure 7: Control Scheme.

of objectives could involve compensation of the distal motor dynamics as well.

The controller has a feedback path to regulate the differential velocity of the two shafts ($\dot{\theta}_1 - \dot{\theta}_2$). If regulation and tracking are sufficiently good, the dynamics of the proximal stage is eliminated from the user’s experience. The transparency is limited by the high-frequency dynamics of the distal stage only. In other words, for a null torque set-point, what is felt is only the distal stage. This scheme relies on an accurately specified viscous coefficient, b , which eddy-current couplers allow.

4.1 Polynomial Pole-Placement

We adopted a discrete-time polynomial pole placement approach because, in contrast with other approaches, this method allows for an exact design (Landau, 1990; Åström and Wittenmark, 1996; Ostertag and Godoy, 2005). The discrete-time controllers are specified by three polynomials $R(z)$, $S(z)$, and $T(z)$, as shown in Fig. 8, where $y(z)$ is the measured plant output, $r(z)$ the reference set-point, and $u(z)$ the control signal.




Figure 8: Polynomial control structure.

Control synthesis is purely algebraic and brings about a solution systematically, once the discrete model of the plant is known and a closed-loop model transfer function, $H_m(z)$, has been specified. The three control polynomials can be selected to achieve certain performance objectives, including regulation and tracking. The discrete-time plant, here the sam-

pled continuous-time plant, is

$$H(z) = \frac{B(z)}{A(z)}, \quad (6)$$

and the control law is

$$U(z) = \frac{T(z)}{S(z)}r(z) - \frac{R(z)}{S(z)}y(z). \quad (7)$$

The feedback control, $-R(z)/S(z)$, gives a closed-loop characteristic equation,

$$A(z)S(z) + B(z)R(z) = P(z), \quad (8)$$

that takes the form of a Diophantine equation.

The design involves first selecting $P(z)$ to determine the closed-loop regulation behavior through the placement of the closed-loop poles, taking the constraints on the degrees of the polynomials into account. The tracking polynomial, $T(z)$, is then selected to independently determine the tracking behavior.

4.2 Design

If we neglect the velocity of the distal stage, that is, consider it small compared to that of the proximal stage, then the compensator acts on the dynamics of the proximal stage only. The discrete-time transfer function of the plant results from cascading a zero-order hold with an integrator and discrete-time derivation. With a sample period of 0.1 ms, the plant transfer function becomes,

$$H(z) = \frac{0.001671z + 0.001669}{z^2 - 0.9967z}. \quad (9)$$

The desired closed-loop polynomial was selected to correspond to a second order continuous plant with natural frequency and damping, $\omega_0 = 680$ rad/s, $\zeta = 0.9$, respectively. This led to the characteristic polynomial

$$P(z) = z^2 - 1.8800z + 0.8848. \quad (10)$$

The solution of the Diophantine equation (8) gave

$$S(z) = z - 0.98855, \quad R(z) = 1.3240. \quad (11)$$

Perfect model tracking is achieved when

$$T(z) = \frac{P(z)}{B(1)}, \quad (12)$$

and the model was

$$H_m(z) = \frac{0.001625z + 0.001569}{z^2 - 1.898z + 0.9009}. \quad (13)$$

4.3 Control Performance

The control system was implemented on a personal computer running RTAI real-time Linux kernel and the sampling rate was set to 10 kHz. The system was tested by holding the handle with a normal grip and demanding a step of torque. Saturation levels were set to 5 mN·m for the distal motor and to 200 mN·m for the proximal motor, because the distal stage motor normally operates only during the transients, its saturation can be set much higher, viz. by a factor five, without adverse effects. Figure 9a shows the response of the system to a step of amplitude 5 mN·m, where the output torque, τ_{out} (black line), is the sum of the viscous coupler torque, τ_{coupler} , and of the distal stage motor torque, τ_2 .




Figure 9: Step response time domain performance. **a**, Case without saturation. **b**, Case with saturation: the distal motor torque, τ_2 , is limited to 5 mN·m.

As was expected, see Fig. 9a, the proximal motor received a large command, τ_1 , during a short transient to minimize the reaction time and the torque missing in τ_{coupler} was taken up by the torque, τ_2 , produced by the distal stage motor. When saturation was not reached for the distal stage motor, as in the case of Fig. 9a, the mechanical time constant of the system was equal to that of the distal motor, that is about, 1.0 ms. In the cases where the demanded torque caused saturation of the distal stage, the transient response was imperfect, as in Fig. 9b where the step magnitude is twice that of the distal motor saturation level. The steady-state part, however, remains unaffected.

5 Validation

The objective was to verify that the high level degree of transparency of the dual-stage drive had a measurable effect on human performance when using a haptic interface. The purpose of the experiment was to test

the device effectiveness, not human performance. To this end, we selected a task where participants had to perform as signal detectors using a stimulus that had two components. The first component had a fixed magnitude and the other was varied from imperceptible values to easily detectable values. To create an experimental condition that resembled actual use, we asked the same volunteers to detect the presence of weak, high-frequency force perturbations superposed onto a slowly varying background force at random locations in the workspace. At any given time, however, the participants were not aware of the condition under which they were working.

The stimuli were administered through two identical handles, one connected to the proximal motor, and other to the distal motor, see Fig. 10. We tested whether the subjects could detect smaller details in the same signal rendered with the dual-stage interface (Condition A) rather than with the large motor alone (Condition B), similar to a conventional interface. All other parameters remained equal.




Figure 10: Experimental setup. In Condition A, the handle was connected to the output of the dual-stage drive. In Condition B, an identical handle was connected directly to the shaft of the proximal motor and the coupler was disconnected. The device was mounted on a heavy steel slab to prevent possible vibrations arising from the rotation of the proximal motor to find a path to the hand.

5.1 Methods

5.1.1 Subjects

Six members of the laboratory, five male and one female, volunteered their time. They were aged from 23 to 56 with a median age of 28.

5.1.2 Stimulus

The slowly varying signal was produced by simulating dry friction and the fast signal was obtained by

modulating the amplitude of the force according to

$$\begin{cases} \tau_C \left[1 - h_0 \sin \left(2\pi \frac{\theta}{\lambda} \right) \right], & \text{when } \theta_0 < \theta < \theta_0 + \Theta, \\ \tau_C, & \text{otherwise,} \end{cases} \quad (14)$$

see Figure 11. The friction torque was calculated according to the algorithm described in (Hayward and Armstrong, 2000) and its level, τ_C , was fixed at 5 mN·m (0.14 N at the handle periphery). From trial to trial, the onset of the texture, θ_0 , changed randomly in the range 0–1.57 radian (0– $\pi/2$), and its extent, Θ , was fixed at 0.035 radian (1.22 mm) and its period, λ , at 0.005 radian (0.17 mm).




Figure 11: Stimulus signal. A small burst of oscillation of size Θ , amplitude h_0 , and spatial period λ , was superposed onto a larger friction force, τ_C .

When turning the handle, it felt as if its motion was impeded by a brake and it was like trying to detect a slight imperfection. For Conditions A and B, preliminary tests enabled us to determine that the thresholds of detection for h_0 were below 0.1 and 1.0, respectively (0.014 N and 0.14 N at the handle periphery). The period, λ , and amplitudes, h_0 , simulated very fine and very small textures, which was a test for the device's capabilities. A better device should allow users to perform better.

Rendering a given texture with a given device, however, imposes limits that are not to be exceeded (Campion and Hayward, 2005). The lowest spatial period, λ , that can be reliably synthesized is constrained by $(\alpha/\lambda)vT < 1$ where v is the velocity and $\alpha > 2$, and T is the sampling period. For a velocity of at most 2.9 rad/s (0.1 m/s at the handle periphery), with $T = 10^{-4}$ s and $\alpha = 10$, the smallest spatial period should be 2.9 mrad (0.1 mm). The other constraints in were easily met by our prototype. The width, Θ , of the texture should also be greater than 3λ .

5.1.3 Procedure

We adopted a two-alternative forced choice testing procedure where h_0 was randomly selected in the sets $\{0.0, 0.01, 0.03, 0.053, 0.08, 0.1\}$ mN·m and $\{0.0, 0.2, 0.4, 0.6, 0.8, 1.0\}$ mN·m, for Condition A and B,

respectively. Each stimulus was presented 10 times, therefore the participants completed 60 trials for each condition, which could be done in a few minutes. In both conditions, the setup was concealed inside an enclosure having the same visual aspect from both sides. Only the two handles were visible. To switch from Condition A to Condition B, all there was to do was to turn the set-up around and to disconnect the coupler from the proximal motor.

The participants were asked if they felt anything at all, in which case they had to respond ‘no’, or a small texture, in which case they would answer ‘yes’, and indicate where they had felt it.

5.1.4 Results

We fit the experimental data with cumulative Gaussian distributions of the form $f(x) = [1 + \text{erf}((h_0 - \mu)/\sqrt{2}\sigma)]$ where μ and σ^2 were the detection mean and variance, respectively, using a non-linear least-square procedure. The threshold points were extracted from the amplitude of the stimulus texture that corresponded to a 50% probability texture detection, for each subject, in the two configurations. The average detection threshold across subjects was 0.036 when using the dual-stage drive and 0.27 when using the standard configuration.

5.1.5 Discussion

All participants, by-and-large, produced typical psychometric curves shown in Fig. 12. Participant SD, who was very cautious when using the improved device, made decisions without hesitation when using the standard drive, but the opposite was true for participant SE. Such differences are not surprising considering that haptic interfaces, no matter how well they are constructed, are instruments and that users can choose to focus on various aspects of their operation, even in a simple detection task. They can also refer to their past experience to various degrees. Whether naive or experienced, their relationship with the instrument and their skill in using it can vary and the effects of changing an instrument design can be different from person to person.

Despite these inter-subject variations, it can then be concluded that the performance detection of the subjects was improved by almost an order of magnitude. It was the case for participants SA, SB, SC, and SF) and to a slightly lesser extend for SD and SE. Lastly, the thresholds corresponded to force variations of 5.2 mN and 38.3 mN for Conditions A and B respectively, the former figure being in line with the human performance threshold (0.5 g experienced dynamically).




Figure 12: Results and fits for Conditions A and B.

6 Conclusion

Using a dual-stage approach, we designed an electromagnetic drive for use in haptic interfaces and other applications where the maximum torque delivered is decoupled from its effective inertia. In terms of haptic interfacing, this means that transparency is improved since the ‘tool’ interposed between the source of torque and its point of delivery causes parasitic forces originating from inertia and friction to be commensurate with perceptual thresholds. In other words, the torque specified is for all practical purposes independent from both the movements of the hand and from those of the device.

In sum, the interface achieves two orders of magnitude of improvement in transparency over existing designs measured in terms of the magnitude of parasitic forces owing to friction and inertia. The inertial load seen by a finger is only of the order of one or two grams. Yet, the design does not limit the torque magnitudes that can be produced, and can operate in its present configuration over a bandwidth that exceeds one kilohertz. For the task that we have selected, the improved transparency allowed users to detect details that were ten times smaller in magnitude than when using a conventional design.

Acknowledgments

The authors wish to thank Andrew H. Gosline for advice on magnetic design of viscous couplers and Bruno Sauvet for help with prototype manufacturing. This work was supported by the ANR (Agence Nationale de la Recherche, France), project PACMAN “perception haptique des échelles micro et nanoscopiques”. Additional funding was from the European Research Council, Advanced Grant PATCH, agreement No. 247300 (to V.H.).

References

- Anwar, S. (2004). A parametric model of an eddy current electric machine for automotive braking applications. *IEEE Transactions on Control Systems Technology*, 12(3):422–427.
- Åström, K. and Wittenmark, B. (1996). *Computer-controlled systems: theory and design*. Prentice Hall, New York.
- Balakrishnan, A., Joines, W., and Wilson, T. (1997). Air-gap reluctance and inductance calculations for magnetic circuits using a schwarz-christoffel transformation. *IEEE Transactions on Power Electronics*, 12(4):654–663.
- Baud-Bovy, G. and Gatti, E. (2010). Hand-held object force direction identification thresholds at rest and during movement. In Kappers, A., van Erp, J., Tiest, W. B., and van der Helm, F., editors, *Haptics: Generating and Perceiving Tangible Sensations*, volume 6192 of *Lecture Notes in Computer Science*, pages 231–236. Springer Berlin / Heidelberg.
- Berkelman, P. J. and Hollis, R. L. (2000). Lorentz magnetic levitation for haptic interaction: Device design, performance, and integration with physical simulations. *International Journal of Robotics Research*, 19(7):644–667.
- Bicchi, A., Peshkin, M., and Colgate, J. E. (2008). Safety for physical human-robot interaction. In *Springer Handbook of Robotics*, chapter 57, pages 1335–1348. Springer-Verlag.
- Boulet, B. (1990). Modeling and control of a robotic joint with inparallel redundant actuators. Master’s thesis, Department of Electrical Engineering, McGill University, Montréal, Qc, Canada.
- Campion, G., Gosline, A. H., and Hayward, V. (2008). Passive viscous haptic textures. In *Proceedings of the 16th Symposium on Haptic Interfaces For Virtual Environment And Teleoperator Systems*, pages 379–380.
- Campion, G. and Hayward, V. (2005). Fundamental limits in the rendering of virtual haptic textures. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2005*, pages 263–270.
- Chew, M., Hong, G.-S., and Zhou, W. (2004). Series damper actuator: a novel force/torque control actuator. In *Proceedings of the 4th IEEE/RAS International Conference on Humanoid Robots*, pages 473–478.
- Conti, F., Khatib, O., and Baur, C. (2007). A hybrid actuation approach for haptic devices. In *Proceedings of the Second Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2007*, pages 367–372.
- Elleau, P., Chavanne, J., and Faatz, B. (2000). Design considerations for a 1Å SASE undulator. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 455(3):503–523.
- English, C. and Russell, D. (1999). Implementation of variable joint stiffness through antagonist actuation using rolamite springs. *Mechanisms and Machine Theory*, 34(1):27–40.
- Gosline, A. H. and Hayward, V. (2008). Eddy current brakes for haptic interfaces: Design, identification, and control. *IEEE/ASME Transactions on Mechatronics*, 13(6):669–677.
- Hayward, V. and Armstrong, B. S. R. (2000). A new computational model of friction applied to haptic rendering. In Corke, P. and Trevelyan, J., editors, *Experimental Robotics VI*, volume 250 of *Lecture Notes in Control and Information Sciences*, pages 403–412.
- Hayward, V. and Astley, O. R. (1996). Performance measures for haptic interfaces. In Giralt, G. and Hirzinger, G., editors, *Robotics Research: The 7th International Symposium*, pages 195–207, Heidelberg. Springer Verlag.
- Heald, M. A. (1988). Magnetic braking: Improved theory. *American Journal of Physics*, 56(6):521–522.

- Hollis, R. L., Salcudean, S. E., and Allan, A. P. (1991). A six-degree-of-freedom magnetically levitated variable compliance fine-motion wrist: Design, modeling, and control. *IEEE Transactions on Robotics and Automation*, 7(3):320–332.
- Khatib, O. (1990). Reduced effective inertia in macro-/mini-manipulator systems. In Miura, H. and Arimoto, S., editors, *Robotics Research 5*, pages 279–284. MIT Press.
- Landau, I. D. (1990). *System Identification and Control Design*. Prentice Hall, Englewood Cliffs, N.J.
- Lauria, M., Legault, M.-A., Lavoie, M.-A., and Michaud, F. (2008). Differential elastic actuator for robotic interaction tasks. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 3606–3611.
- Millet, G., Haliyo, S., Regnier, S., and Hayward, V. (2009). The ultimate haptic device: First step. In *Proceedings of the Third Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009*, pages 273–278.
- Morrell, J. B. and Salisbury, J. K. (1997). In pursuit of dynamic range: Using parallel coupled actuators to overcome hardware limitations. In *Experimental Robotics IV: the 4th International Symposium, 1995*, volume 223 of *Lecture Notes in Control and Information Sciences*, pages 263–273. Springer-Verlag.
- Ostertag, E. and Godoy, E. (2005). RST-controller design for sinewave references by means of an auxiliary diophantine equation. In *Proceedings of the 44th IEEE Conference on Decision and Control and the European Control Conference (CDC-ECC '05)*, pages 6905–6910.
- Pratt, G. A. and Williamson, M. M. (1995). Series elastic actuators. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 399–406.
- Quimby, D. C. and Pindroh, A. L. (1987). High field strength wedged-pole hybrid undulator. *Review of Scientific Instruments*, 58(3):339–345.
- Salcudean, S. E. and An, C. (1989). On the control of redundant coarse-fine manipulators. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1834–1840.
- Sharon, A. and Hardt, D. E. (1984). Enhancement of robot accuracy using endpoint feedback and a macro-micro manipulator system. In *Proceedings American Control Conference*, pages 1836–1842.
- Sharon, A., Hogan, N., and Hardt, D. E. (1988). High bandwidth force regulation and inertia reduction using a macro/micro manipulator system. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 126–132.
- Shrivastava, R. and Kumar, S. (2009). An alternative approach for calculation of braking force of an eddy-current brake. *IEEE Transactions on Magnetics*, 45(1):150–154.
- Stocco, L. and Salcudean, S. E. (1996). A coarse-fine approach to force-reflecting hand controller design. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 404–410.
- Stramigioli, S. and Duindam, V. (2001). Variable spatial springs for robot control applications. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1906–1911.
- Sugano, S., Tsuto, S., and Kato, I. (1992). Force control of the robot finger joint equipped with mechanical compliance adjuster. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2005–2013.
- Wall, S. A. and Harwin, W. (2001). A high bandwidth interface for haptic human computer interaction. *Mechatronics*, 11(4):371–387.
- Wouterse, J. H. (1991). Critical torque and speed of eddy current brake with widely separated soft iron poles. *Electric Power Applications, IEE Proceedings B*, 138(4):153–158.
- Yao, H.-Y. and Hayward, V. (2010). Design and analysis of a recoil-type vibrotactile transducer. *Journal of The Acoustical Society of America*, 128(2):619–627.
- Zinn, M., Khatib, O., Roth, B., and Salisbury, J. K. (2008). Large workspace haptic devices – a new actuation approach. In *Symposium on Haptic interfaces for virtual environment and teleoperator systems*, pages 185–192.
- Zinn, M., Roth, B., Khatib, O., and Salisbury, J. K. (2004). A new actuation approach for human friendly robot design. *International Journal of Robotics Research*, 23(4):379–398.