

Interaction Control for a Brake Actuated Manipulator

Brian Dellon¹
Mechanical Engineering
Carnegie Mellon University

Yoky Matsuoka²
Computer Science & Engineering
University of Washington

ABSTRACT

If a passively actuated physical human-robot interaction (pHRI) device could produce the transparency and controllability of an active device, the safety standard of pHRI could dramatically improve. This is critical since pHRI devices are being developed for rehabilitation or assistance where the user may not be able to provide full strength or speed to counteract accidental movements caused by a motor driven robot. This paper demonstrates effective friction compensation and an interaction controller that increases overall transparency of a passively actuated pHRI device, the Brake Actuated Manipulator (BAM). We analyze the friction compensator's stability with passivity theory, and evaluate the proposed interaction control scheme on the BAM with real time experiments during display of virtual environments composed of non-linear Coulomb or viscous friction. Results show excellent force tracking performance during display, and indicate that passively actuated pHRI devices can be transparent despite notions that brakes are poor actuators.

KEYWORDS: Haptic I/O, Control theory, Manipulators.

INDEX TERMS: H.5.2 [Information Interfaces and Presentation]: User Interfaces; I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search; I.2.9 [Artificial Intelligence]: Robotics.

1 INTRODUCTION

Robotic manipulation by definition implies interaction with the surrounding world, whether for assembly, medical care, or with virtual/distant dynamic environments. A distinction comes when *physical* human contact is made with the device, as the robotic control must consider human safety and uncertain human inputs. This physical human robot interaction (pHRI) is embodied by haptic interfaces, tele-operated manipulators, exoskeletons, and human power augmentation devices. The primary goals of these devices are to provide an experience of contact and dynamic interaction with an environment and to extract/augment/assist human movement and sensing.

One key issue is to ensure safe and stable interaction while providing high fidelity in the face of uncertain human dynamics, non-linear environments, and the effects of sampling time and sensor quantization. Safety is especially critical when an assistive

robot is applied to physical rehabilitation where the operators may already be injured and a large workspace is required. One way to mitigate these safety concerns is through the use of a passively actuated haptic display. Passive actuation through clutches [1], steerable mechanisms [2], or brakes [3] presents the benefit of being inherently stable throughout interaction, therefore conserving safety. This is in contrast to actively actuated systems, which can have large gearings, high effective inertia, or no backdrivability.

For proper dissipative interaction, it has been shown that the device kinematics must be orthogonal to resist any force in Cartesian space [4]. All orthogonal kinematics include at least one linear joint. The drawbacks of linear joints include extra weight and friction that may not be present for revolute joints. Friction in the linear joint remains a sticking point and requires active compensation for good transparency. Currently, we are unaware of any pHRI controllers (known as interaction control) that deal with passive actuation and a device's friction while matching the desired interaction force properly.

Typically pHRI systems are operated with either impedance [5] or admittance [6] control architectures. Causality indicates that a passively actuated, dissipative, device must be controlled as an impedance since the brakes act as programmable friction elements. This constraint is due to the non-invertible nature of frictional force. A given velocity source from an operator can always define a resultant frictional force, but for a given force input, a one-to-one mapping is not necessarily defined for velocity. Consequently, and due to lack of controllability with passive actuation, joint velocities cannot be directly commanded.

One shortcoming of pure impedance control is that the interaction force is only indirectly specified through variation of the controller or virtual environment's parameters. This can pose a problem when the stochastic nature of the environment is weighed against the deterministic nature of system models, causing inaccuracies in interaction force. More sophisticated methods to counteract static and dynamic effects on interaction force have been a topic of much study for motor-driven, active, devices to enable accommodation of simultaneous motion and force control [7-9]. These interaction controllers come in several varieties such as parallel force/position [8] and hybrid force/position [9]. The concept is to close an outer force loop around the inner position or velocity loop with integral action on the force error. This strategy regulates interaction force to a desired steady state value and has been shown to be effective for manipulators interacting with their environment [9].

In general these architectures assume a given trajectory and an unknown dynamic environment. However, for passively actuated pHRI this assumption does not work as well because the human's trajectory is generally unknown. Our solution is to design a new interaction controller to accommodate passively actuated pHRI.

Our philosophy towards haptic rendering in a large workspace requires passive interaction, and more strictly we require all force feedback to be dissipative with respect to the operator. This is in

This work was supported in part by the National Science Foundation under Grant 0238204 and in part by the National Institutes of Health under Grant R21HD47405.

¹e-mail, btd@andrew.cmu.edu

²e-mail, yoky@cs.washington.edu

contrast with other hybrid (active/passive) devices that leverage the controllable brake to increase their z-width [10]. This is done by varying the amount of physical damping in the hybrid interface, thus increasing the region of stability for rendering stiff walls [11]. Though our strategy poses a greater technical challenge in terms of smooth haptic rendering, we can inherently retain safety. This philosophy, however, does not imply that active effects cannot be simulated or produced by other means of affecting the operator such as through visual distortion [12] or perceptual illusion.

It is generally perceived that brakes produce friction and have slow response times, implying that transparency will not be sufficient for a high fidelity pHRI device. We show here that these drawbacks can be overcome through controller design, and that a passively actuated pHRI device can simulate highly realistic and unimpeded virtual environments. This paper describes two solutions that enhance transparency for passively actuated pHRI such as the Brake Actuated Manipulator (BAM). Because all dissipative devices require at least one linear axis [4], we propose a friction/gravity compensator to increase back drivability for that axis. Further, a new interaction controller is designed and implemented on the BAM to demonstrate its efficacy.

2 BRAKE ACTUATED MANIPULATOR AND ITS FRICTION COMPENSATION

2.1 Brake Actuated Manipulator

The experimental platform used is the large workspace (2m^3) Brake Actuated Manipulator shown in Figure 1. The BAM uses rotary magneto-rheological (MR) brakes (Placid Industry B15), has spherical kinematics, produces 200N maximum resistance at one meter, and is capable of simulating life-sized virtual environments with force feedback. Details of its mechanisms, dynamics, and system identification results are described in [3], whereas this paper narrows the focus to address interaction control and friction compensation. In this section, we address the system friction and its requisite compensation before an effective interaction controller can be implemented.

There are many sources of friction in a power transmission that affect the overall transparency of a haptic device. Mechanical components such as bearings, gearing, cabling, or belting are notorious for introducing Coulomb, viscous, or even chaotic frictional effects into the system, and the controller must be able to compensate for them. The BAM has two main types of joints: revolute and linear. The revolute axes have minimal friction ($<2\text{Nm}$), high backdrivability, and no backlash due to their cable drive transmissions, therefore friction compensation is unnecessary. On the other hand, the linear axis has a belt-driven transmission with viscous and Coulomb friction that affect haptic rendering. To augment these friction effects and gravity, a lightly geared (4:1), small motor (Faulhaber series 3257) was added to the linear axis. The addition of the DC motor changes the previously reported frictional characteristics of the linear joint, thus friction must again be identified.

2.2 Friction Identification and Modeling

As it was demonstrated in [3], passive devices do not have a way to provide automatic velocity control without an active actuator as system identification inputs. Therefore, we used human inputs as a motion source and recorded the interaction with a force sensor (ATI Mini45) on the BAM's grip. The user moved the joint back and forth through its total excursion cyclically at varying speeds. Because we characterized the dynamic model of the BAM [3], it was possible to use a non-constant velocity/force input from a




Figure 1. The Brake Actuated Manipulator (BAM) shown with its Cartesian and spherical coordinate system orientations. The BAM has 2m^3 workspace sufficient for whole limb/body interactions.

user and calculate friction by subtracting any inertial terms. Figure 2A. (friction curve in light gray) shows the measured friction for the BAM. The effective mass for the linear joint plus end effector has previously been reported as $m = 3.56\text{kg}$ [3].

For modeling friction, we initially considered numerous static and dynamic friction phenomena such as stiction, Stribeck effect, frictional lag, pre-sliding displacement, etc. These effects can be traced back to their tribological roots and the microscopic asperities at interfacial surfaces from which they arise [13]. Popular models such as the Dahl and LuGre have been developed to explain these dynamic effects, but their parameters must be carefully identified, they are not robust towards plant variations, and for control purposes parameters must be constrained with respect to one another to retain passivity [14-15]. One must question the necessary complexity of friction compensation for large scale passively actuated pHRI when considering the perceptual threshold for upper extremities, coupled with both large friction/inertial forces and one or two dominating frictional effects. It may be sufficient to lower the system friction just beneath an operator's perceived force threshold for moving their arm in free space. Consequently we posit the use of a simplified model will produce acceptably low values of friction.

For our BAM system (friction shown in Figure 2), the dominating characteristics are Coulomb and viscous effects with limited stiction, and little or no Stribeck effect. Simple static models for Coulomb friction are discontinuous around zero velocity because they depend on the sign of the joint's motion. Therefore, for our first approximation of the linear axis' frictional characteristics, we used a sigmoidal function as a close approximation with continuous zero crossings given by,

$$\hat{F}_f = \hat{F}_c \left(\frac{2}{1+e^{-\alpha \dot{r}}} - 1 \right) + \hat{B} \dot{r}, \quad (1)$$

where \dot{r} is the linear axis' velocity, \hat{F}_c and \hat{B} are the estimated Coulomb and viscous friction levels respectively, and α is a constant that approximates stiction and also effects the response time of the total friction estimate \hat{F}_f . In order to estimate the parameters of (1), non-linear least squares regression was used against the measured data from Figure 2A. The parameter estimates along with standard deviations are listed in Table 1 and

the model fit is shown with confidence bounds in Figure 2A. (black solid and dotted lines).

TABLE I
MODEL PARAMETER ESTIMATES

Parameter	Magnitude	Standard Deviation
\hat{F}_c	15.81 N	0.05 N
\hat{B}	7.59 Ns/m	0.04 Ns/m
α	19.72 s/m	0.20 s/m
Variance of the Residuals, $F_{tot} - \hat{F}_{tot}$	$\sigma^2 = 5.47 \text{ N}^2$	

2.3 Friction Compensation

There are two main strategies for model based friction compensation available, feed-forward and feed-back. Because our design philosophy dictates that all haptic feedback must remain dissipative in order to retain safety during interaction, we ruled out the option of feed-forward compensation, which requires the linear joint to track a target trajectory thus specifying velocity directly. This would require an admittance controller which violates causality for our brake actuated device. To perform a stability analysis, the gain K_c was placed on the Coulomb friction term in (1) and the friction compensator becomes,

$$\begin{aligned}\hat{F} &= \hat{F}_c \left(\frac{2}{1+e^{-\alpha \dot{r}}} - 1 \right) \\ \hat{F}_{tot} &= K_c \hat{F} + \hat{B} \dot{r}\end{aligned}\quad (2)$$

In order to ensure (2) remains dissipative, we used passivity theory to analyze the energy exchange between the BAM and controller, then derived a stability criterion bounding K_c . Following [16], using the standard sign conventions for passivity, a system with flow V , effort F , and initial energy $e(0)$ is passive if,

$$\int_0^t F(\tau)V(\tau)d\tau + e(0) \geq 0 \quad (3)$$

for all functions F , V , and $t \geq 0$. The implication from (3) is that our friction compensator (2) must produce overall less force than the joint's inherent friction to remain stable (i.e. $K_c \leq 1$). To satisfy the passivity equation (3) we must compare the true system friction with the amount compensated by (2). However, from Figure 2A., it is clear that approximating the true system friction with a deterministic function, as in (2), will not adequately reflect the lower limits of the system's true friction due to its inherent variability and complex nature. Because of this it is difficult to bound Coulomb friction and K_c may need to be significantly less than one to remain stable $\forall t \geq 0$. To overcome this problem, we used a stochastic approach to define the bounds for K_c . Stability ensured by this bound will be robust to change in friction parameters over time given the fact that the system friction is likely to increase from wear. We first assume the true friction in the linear axis to be stochastically represented by a mean estimate of Coulomb/viscous phenomena with additional zero mean Gaussian process noise w ,

$$F_{tot} = \hat{F}_c \text{sign}(\dot{r}) + \hat{B} \dot{r} + w, \quad w \sim N(0, \sigma) \quad (4)$$

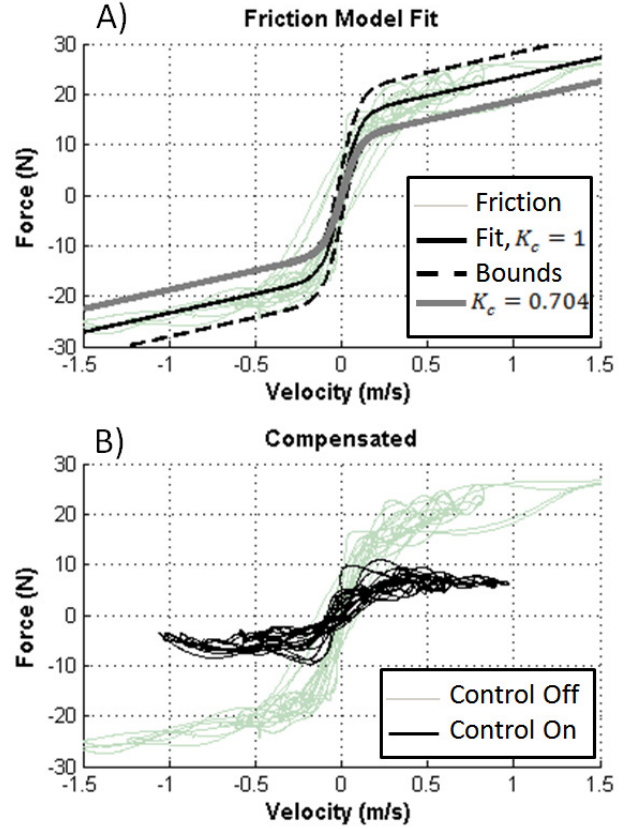


Figure 2. A) The sampled friction data and corresponding model fit, $K_c=1$, with 95% confidence bounds for the linear joint. The curve shown with gain $K_c=0.704$ produces passive compensation. B) After compensation friction is considerably reduced and viscosity is eliminated.

where σ^2 represents the variance of observed friction effects. Now it is possible to carry out the passivity analysis.

If we consider the linear joint as a point with mass m , its change in kinetic energy would be, $\dot{e} = m\dot{r}\ddot{r}$. Performing a force balance with (2) and (4) on the point mass yields,

$$\dot{e} \leq (\hat{F}_c \text{sign}(\dot{r}) - K_c \hat{F} + w) \dot{r} \quad (5)$$

which is the differential form of (3). In order to satisfy the constraint in (5) we require only that $F_c + w \geq K_c \hat{F}$ by noting all frictional terms have the same sign as \dot{r} . The variance, σ^2 , of the process noise can be evaluated by examining the residuals of $F_{tot} - \hat{F}_{tot}$. If F_{tot} is limited by the 95 percent confidence interval then, $|F_c + w| \geq \hat{F}_c - 2\sigma$. It is also true that $|\hat{F}| \leq \hat{F}_c$. Equating these inequalities and solving for K_c arrives at

$$K_c \leq 1 - \frac{2\sigma}{\hat{F}_c} \quad (6)$$

as the criterion for passivity. Evaluation of (6) with the values from Table 1 arrives at $K_c = 0.704$ as the maximum value.

Because the joint's frictional parameters are free to vary $\pm 2\sigma$ implies our compensator will cancel only a proportion of the friction. This proportion is given by, $\frac{\hat{F}_c - 2\sigma}{\hat{F}_c + 2\sigma}$ and reflects how

closely the true parameters are known. If there is a large uncertainty, then passivity can only be guaranteed for smaller values of K_c . For the values in Table 1 this leads to a ratio of 0.54, meaning over fifty percent of system friction is being annulled. The model fit with gain $K_c = 0.704$, used for compensation, is depicted in Figure 2A. (dark gray curve). We have verified experimentally that gains above the ideal value in (6) result in marginal stability, and setting $K_c > 1$ results in runaway.

To evaluate the effectiveness of the friction compensator, friction parameters were again identified using the gain specified by (6) in a similar manner as before. The resulting friction curve is shown in Figure 2B. (black curve) with identified parameters listed in Table 2.

TABLE II
MODEL ESTIMATES WITH COMPENSATION

Parameter	Magnitude	Standard Deviation
\hat{F}_c	5.75 N	0.06 N
\hat{B}	0.0013 Ns/m	0.10 Ns/m
α	13.83 s/m	0.29 s/m

The linear joint's friction compensation makes it significantly easier to back-drive and increases that axis' transparency. For comparison the back-drive force listed in Table 2 is similar to that required to stretch an ordinary (3.5 x 0.125 inch) rubber band to twice its length; This is impressive considering the inertia is ~3.56kg. The reported joint damping during compensation is statistically indifferent from zero and the Coulomb term is reduced to 1/3 of the original value. We have found that the stability criterion is quite conservative and increasing it from 0.7 to 0.9 allows us to reduce Coulomb friction further to 3.37N ($\sigma = 0.07$) while maintaining satisfactory stability.

Before this friction compensator can be treated holistically in the next section, we show how the effects of gravity are also accounted for. The BAM's kinematics are spherical, and this results in the linear joint's center of gravity shifting with respect to r , producing body forces along both the linear and pitch axes. While a spring loaded mechanism counteracts gravity along the pitch axis irrespective of r [3], the linear axis' gravity can be compensated by a combination of motor and MR brake in conjunction with the friction compensation methods above. The total compensation, including gravity, along the linear axis is given by,

$$F_{comp} = \hat{F}_{tot} + mgr \sin \phi \quad (8)$$

For hybrid actuation, a transition point must be established dictating the switch from motor to MR brake. Comparing the sign between F_{comp} and the component of sensed user force along the linear axis, F_r , allows us to distinguish the case where the required compensation is purely resistive. This corresponds to the ideal situation where the motor is not needed and the MR brake produces F_{comp} . When there is no (user) interaction force the sensor will oscillate around zero due to noise; rapid switching between modes and poor static compensation of gravity will result. For this reason we define a bound, δ , on the switching behavior so the controller defaults to braked mode, thus requiring a minimum applied user effort before activating. The motor torque, τ_m , and MR brake torque, τ_b , then become,

$$\tau_m = \begin{cases} 0, & F_{comp}F_r < \delta \\ K_m F_{comp}, & F_{comp}F_r \geq \delta \end{cases} \quad (9.a.)$$

$$\tau_b = \begin{cases} K_b F_{comp}, & F_{comp}F_r < \delta \\ 0, & F_{comp}F_r \geq \delta \end{cases} \quad (9.b.)$$

where K_m and K_b represent gearings from linear to rotational motion. The effort bound, $\delta = 3$, was chosen experimentally to provide both a low threshold for activation and robustness to loss of interaction force from a soft grip.

3 INTERACTION CONTROLLER DESIGN

Using the BAM, we are interested in letting the user "feel" the simulated environment accurately. This means that the impedance of the device felt by the user must match that of the simulated world. To ensure this, the system model should be defined as accurately as possible. This is not trivial when the device has inertia, friction, and knowledge of gearing may be imperfect.

Furthermore, MR brakes do not have a pure linear relationship between input current and output torque. They exhibit hysteresis and have a slow mechanical response time (~15-25ms) relative to a motor. Actuator models have been developed for MR brakes [17], but usage of such does not increase response time nor reproduce faithful interaction in the face of inertia. Therefore, a simple open loop interaction between commanded force and the virtual world will not be sufficient.

Our approach is to close a force loop between the commanded force from the virtual environment and a force sensor at the BAM's end effector. This error signal contains all of the unmodeled dynamic information including inertial effects, and can be used to modulate the current commands sent to the brake to drive interaction force towards the desired value from the virtual environment. The schematic of our interaction controller is shown in Figure 3 and details are described in the following three subsections.

3.1 Controller Selection

We have chosen to use a PI controller with feed-forward torque commands to act on the interaction force error signal. The control law is stated as,

$$\ddot{u} = J^T \ddot{f}_{ve} + K_p \{J^T(\ddot{f}_{ve} - A\ddot{f}_u) + K_i \int_0^t J^T(\ddot{f}_{ve} - A\ddot{f}_u) d\zeta\} \quad (10)$$

Where K_p , and K_i are diagonal matrices of proportional and integral gain respectively, J^T is the BAM's transposed Jacobian, A is an attitude matrix which transforms sensed force, \ddot{f}_u , into Cartesian space, and \ddot{f}_{ve} is the commanded vector of Cartesian forces from the virtual environment. The gain matrices were chosen empirically to provide responsive accurate behavior, no formal optimality criteria were employed, the gains used are listed in Table 3.

TABLE III
CONTROLLER GAINS

Gain	θ	ϕ	r
K_p	1	1	1
K_i	5	5	10
K_{is}	6	6	11
η	0.25	0.75	0.25

The requirement for a feed-forward term is straightforward: it drives the BAM's joints with the virtual force constraint under open-loop interaction and allows for good force tracking with relatively low feedback gains [9]. Feeding the force constraint forwards also helps to increase the interaction force response time. Without this term the force controller would have a sluggish response relying on integral action to achieve the constraint.

3.2 Considerations of Braked Interaction

Interaction with a brake actuated device such as the BAM can result in unexplored effects that affect the fidelity of the virtual environment. Consider the benchmark of haptic interaction, the virtual wall. Virtual walls are typically rendered as stiff springs, however when rendered with a brake, only the compression phase of the spring can be displayed and the operator will remain 'stuck' inside the constraint. This is largely due to the fact that the controller does not know the user's intended trajectory in advance.

In addition, for a very stiff wall the final penetration depth determines the maximum constraint force and the user's force may saturate below this value upon subsequent pushes while they remain 'stuck.' The consequence is that \vec{u} will quickly become very large due to integral windup along the constrained axes.

In order to counteract the integral windup we use a back information propagation scheme from [18] augmented with dynamic saturation in between the MR brake and the controller output. The upper and lower limits of this saturation are dynamically set by \vec{f}_{ve} . If either limit is reached an additional integral term with a faster time constant than K_i stabilizes the feedback loop. To accomplish this we modify (10) resulting in,

$$\vec{e}_s = \begin{cases} 0, & |\vec{u}| \in \{\eta |J^T \vec{f}_{ve}|, (I + \eta) |J^T \vec{f}_{ve}|\} \\ J^T \vec{f}_{ve} - \vec{u}, & \text{otherwise} \end{cases} \quad (11.a.)$$

$$\vec{u} = J^T \vec{f}_{ve} + K_p (J^T (\vec{f}_{ve} - A \vec{f}_u) + K_i \int_0^t J^T (\vec{f}_{ve} - A \vec{f}_u) d\zeta) + K_{is} \int_0^t \vec{e}_s d\zeta \quad (11.b.)$$

Where \vec{e}_s is the error vector from saturation, I is the identity matrix, K_{is} is the integral saturation gain matrix with $K_{is} \geq K_p K_i$, and η is a matrix of constants which dictates the percentage of commanded force within which the controller can modulate brake torques. In this way the integral action can never drive above the saturation limit, nor command a brake torque with opposing sign to \vec{f}_{ve} . The chosen values for η and K_{is} are listed in Table 3.

In order to combat the 'stuck' wall effect, our solution is to switch off the control and allow a free mode when $\vec{u} \cdot J^T A \vec{f}_u \geq 0$. This works because all force feedback is defined dissipatively countering the user's impending motion. These consequences are a focus of ongoing research with passive haptic devices [16] and our free mode is similar to the passivity observers employed by other authors [19]. When free mode is activated the accumulating integral effects are reset to zero in preparation for the next constraint.

3.3 Incorporating Friction Compensation

The friction compensator designed in section II was incorporated in our controller (Figure 3). The force controller acts as high level input to the friction compensator. By augmenting (8) with the row of \vec{u} , corresponding to the linear joint's brake command, u_r , we obtain the final control signal for the linear axis,

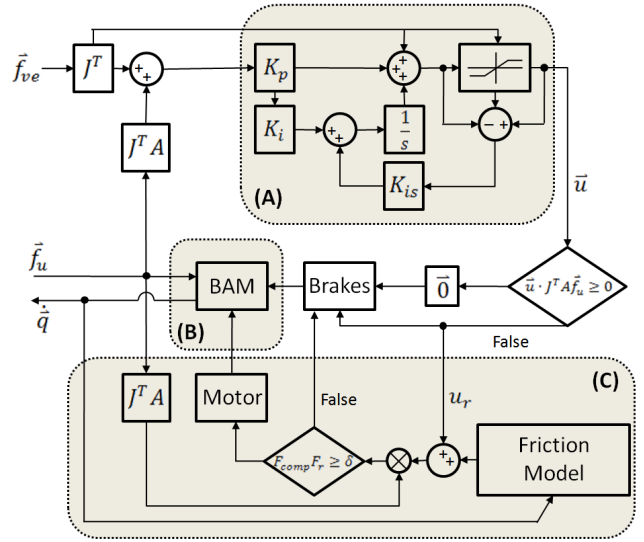


Figure 3. Schematic of interaction controller. Note, input from virtual environment is chosen to satisfy $\vec{f}_{ve} \cdot \vec{f}_u \leq 0$. Section (A) denotes the PI force controller with feed forward force signal and back information propagation. Block (B) is the plant including system friction, and (C) is the friction compensation sub-system.

$$F_{comp} = \hat{F}_{tot} + mgr \sin \phi + u_r \quad (12)$$

The force control along the linear axis, u_r , always acts in the same direction as the linear joint's friction. When the commanded force along this axis is inside the friction envelope, compensation decreases and natural frictional characteristics are used to render the force-feedback. Otherwise, the motor is used to compensate. This proceeds until the commanded value is large enough for the brake to take over rendering as in (9).

4 INTERACTION CONTROL PERFORMANCE

4.1 Validation Methods

To evaluate the performance of our interaction controller, a trajectory following task was defined within a virtual environment running on the BAM. The virtual environment displayed two conditions: (1) viscous and (2) non-linear Coulomb friction effects. Viscosity is an effect for simulating motion through water, other dense liquids, and dynamic interaction with clay. Here we use Cartesian viscosity for test purposes with damping set to 44Ns/m (0.25 lb*s/in). The virtual environment's output was set to,

$$\vec{f}_{ve} = -B_{ve} \dot{\vec{q}} \quad (13)$$

Where \vec{q} is the generalized device coordinates and B_{ve} is a diagonal viscosity matrix.

Coulomb friction is intrinsic to the behavior of the BAM as the MR brakes act like programmable friction elements. Viscosity is actually rendered as Coulombic friction with a linear coefficient. Non-linear Coulomb friction forces therefore comprise a general behavior for any number of virtual environments. We set the virtual environment's output to,

$$J^T \vec{f}_{ve} = -C_{ve}(t) \text{sign}(\dot{\vec{q}}) \quad (14)$$




Figure 4. Plots comparing the linearity of virtual viscosity for the revolute and linear axes with and without interaction control. Cartesian damping was set to 44Ns/m with $r=0.76$ m for trials with the pitch axis.

where C_{ve} is a diagonal matrix with time varying sinusoidal terms.

The controller was implemented in software running at a rate of 1kHz in windows with a quad core Xeon processor @ 2.00Ghz. Since the yaw and pitch axes share similar transmission construction we present only the pitch axis performance as one is representative of the other. To exclude any potential coupling effects the revolute axis is tested independently of the linear joint.

A desired sinusoidal trajectory was defined, represented on screen by an oscillating dot, while another dot represented the user's current position. The user manipulated the BAM to match the control dot's position. To provide comparisons, interaction force/torque with and without a controller were measured. In the case of 'control off' neither interaction control nor friction compensation were enabled; virtual environment forces are simply commanded open loop using a look up table to drive the MR brakes.

4.2 Validation Results

Figure 4 shows both velocity and force response for independent measurements of each type of joint under the viscous condition. Ideally the viscosity curve follows a straight line with slope equal to the diagonal element of B_{ve} when transformed into Cartesian space and projected onto a single axis. Figure 4 shows how the open loop curves deviate significantly from the commanded values. By transforming the joint space information into Cartesian space a line can be fit indicating the average viscous behavior - an estimate of B_{ve} . From linear regression the pitch axis' average viscosities are 33.94 ± 0.085 Ns/m ($R^2 = 0.92$) and 45.69 ± 0.56




Figure 5. Plots comparing force tracking in a non-linear Coulomb friction environment. Dotted line represents the desired interaction force, note how its sign flips periodically, this is related to the operator's velocity reversals. With interaction control (dark line) the operator's force profile closely matches the commanded value. Linear axis (bottom plot), horizontal lines indicate the friction envelope.

Ns/m ($R^2 = 0.98$) for control off/on cases respectively. Average viscosities for the linear axis are 49.33 ± 0.24 Ns/m ($R^2 = 0.83$) and 44.77 Ns/m ± 0.081 ($R^2 = 0.98$) for the control off/on cases respectively. For the controller off case, the linear axis performance shows large amounts of friction present. The interaction controller drives the system toward the commanded value of 44Ns/m, and R^2 shows that most of the measured variance is accounted for by the linear relationship of viscosity. Even with the controller, there is still some deviation from the commanded force at low velocity. This occurs during the deceleration phase of motion when the MR brakes are not effective for inertial compensation. The pitch axis is representative of yaw direction, except open loop error along the pitch axis tends to be larger due to the effects of gravity.

Figure 5 shows the force history for trials with simulated non-linear Coulomb friction for both types of joints. We can see how the desired interaction force follows the sinusoidal command while the user executes a separate sinusoidal motion at a different frequency. Without interaction control performance is only acceptable, resulting in RMS errors of 3.93 Nm and 8.15N for the pitch and linear axes respectively. The error is taken as the difference between the commanded interaction force and the recorded force for each case (on/off). Activating the controller RMS error drops to 2.07 Nm for pitch, and 3.94 N for the linear axis. In the case of the linear axis with no force control the response is dominated by the friction in the system. With both

friction compensation and interaction control the force is tracked even inside the friction envelope, represented as horizontal lines in the bottom half of Figure 5. The horizontal lines correspond to the estimated Coulomb friction in the joint before compensation following Table 1.

We conducted the Integral of Time multiplied by the Absolute value of Error (ITAE) performance index, as another metric to examine effectiveness [20]. Error is calculated in the same fashion as for the RMS calculations above. The ITAE index emphasizes errors occurring later in the response, this helps reduce bias between the time data collection begins and when the operator starts tracking the dot. Because the ITAE accumulates error it represents overall tracking efficiency as well as the steady state error of the interaction controller. This metric helps to examine the effectiveness of integral action in our controller. The ITAE scores for all trials shown in Figures 4 & 5 are listed in table 4.

TABLE IV
ITAE PERFORMANCE INDICES

Controller	Viscous Field		Non-Linear Coulomb	
	On	Off	On	Off
Pitch Axis	456	2.16e3	0.62e3	1.90e3
Linear Axis	436	2.63e3	1.18e3	3.94e3

The ITAE values are encouraging with respect to the performance of the interaction controller. For the viscous field condition cumulative tracking error decreases by an order of magnitude, with comparable ITAE scores along both types of joints. In the non-linear Coulomb case ITAE drops by as much as seventy five percent. It's interesting to note that the change in ITAE is approximately double for the non-linear Coulomb condition. This is a reflection of the controller's response to smoothly varying (ramp type) inputs of the viscous environment versus the step like nature of a non-linear Coulomb input. A step input necessarily contains high frequency content leading to overshoot, longer settling times, and a higher ITAE score.

5 DISCUSSION

We have shown effective force matching for benchmark virtual environments with a passive haptic device using a novel interaction controller. This control architecture provides the necessary tools for dealing with the behavior of passively actuated dissipative devices in a stable manner without knowledge of the user's trajectory; previously unexplored by other interaction controllers [7-9].

Two consequences of control with dissipative devices were indicated namely, human force saturation and 'sticky' walls. We provide solutions via dynamic saturation and observation of user force. Because non-linear Coulomb effects are accurately tracked a vast number of virtual environments composed of non-linear dissipative elements can be faithfully experienced. Furthermore, because a force loop is closed between the virtual environment and the user's hand, device dynamics are fed back to produce a component of the dissipative force.

A method for stable hybrid friction and gravity compensation under the conditions of additive process noise in the system's frictional characteristics is shown. The stability criterion is conservative, resulting largely from unknown operator dynamics which affect the behavior of the system's friction. As demonstrated in Figure 2 and Tables I and II, a major benefit of our design philosophy is that we maintain safety during compensation while decreasing the back driving force. Although




Fig 6. Limit cycles at the limits of the BAM's workspace, $r = 1.25m$, large interaction forces are commanded. Here a Cartesian viscosity of 135Ns/m is attempted and damped oscillations result.

friction is not completely cancelled Weber's law tells us that for heavy objects rendered in large scale environments (chairs, doors, pots, books), as long as the interaction force is sufficiently high this small amount of friction will be imperceptible [21].

It's also worth noting that choosing the switching condition in (9) conditioned on force instead of velocity, creates a scenario which favors limit cycles over runaway instability. This is beneficial when considering safety, not only from a practical point of view, but because it helps with static cancellation of gravity. A slight stutter was sensed when the joint is manipulated around the effort limit δ in (9). This is being addressed through more accurate transmission models, and it disappears during regular haptic rendering even with the current model.

The effective inertia for the linear joint remains isotropic over its excursion. This is not so with the revolute axes, because their effective inertia varies with the arm's center of mass. This anisotropy of inertia shows its effects in the force response at the limits of the BAM's workspace. At this limit sensor quantization and integral competition with friction also play a role. For a given Cartesian velocity the angular velocity is related to the distance from revolution, thus motion at high excursions produce low joint velocities. This can inject high frequency noise into the control signal. Conversely for a given Cartesian force the resistive torque must be higher further away, hence the controller will also need to modulate over larger amplitudes.

These effects contribute towards damped oscillation when a step in the commanded resistance is large. The Cartesian damping limit on our system is empirically found to be 122Ns/m at maximum excursion of 1.25m on the linear axis. Figure 6 shows damped oscillation in the force response for a virtual environment displaying high levels of damping ($>135Ns/m$ @ 1.25m). At low velocity force is still tracked reliably, the first zero crossing injects high frequency noise and oscillation results. Reducing the proportional gains found in Table III to values less than unity helps to control overshoot and reduce this effect considerably. To overcome sensor quantization a dead band around the encoder resolution may be implemented similar to [22] to improve rendering. Gain scheduling can be used to account for the change in effective inertia. Because the change in this quantity is known a priori the proportional gain matrix may be varied according to the parallel axis theorem. Finally, incorporating a model based observer for feedback will reduce controller action and require smaller gains. In particular at high excursions where the BAM's

dynamics play a more dominant role we will create smooth interaction for any generic dissipative environment.

In conclusion, our interaction control for passively actuated devices, such as the BAM, provides the user with accurate force information from the simulated environment despite un-modeled non-linearities in the MR brakes and joint friction. The techniques presented here can easily be extended to work with other passively actuated pHRI devices. Human intuition would dictate that brake actuated devices have poor fidelity and cannot produce smooth interaction, mostly due to their inherent frictional nature. However, here we have shown an exciting prospect that brakes can have high transparency and acceptable force tracking in non-linear environments despite their drawbacks.

REFERENCES

- [1] D. K. Swanson and W. J. Book, "Path-following control for dissipative passive haptic displays," in *Proceedings of 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2003, pp. 101-8.
- [2] J. E. Colgate, *et al.*, "Cobots: robots for collaboration with human operators," in *IMECE*, 1996, pp. 433-439.
- [3] B. Dellon and Y. Matsuoka, "Modeling and System Identification of a Life-Size Brake-Actuated Manipulator," *IEEE Transactions on Robotics*, vol. 25, pp. 481-491, 2009.
- [4] Y. Matsuoka and W. T. Townsend, "Design of lifesize haptic environments," in *Experimental Robotics VII*, R. a. Singh, Ed., ed, 2001.
- [5] N. Hogan, "Impedance control: An approach to manipulation, Parts I-III," *ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 107, pp. 1-24, 1985.
- [6] D. E. Whitney, "Force Feedback Control of Manipulator Fine Motions," *ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 99, pp. 91-97, 1977.
- [7] S. Chiaverini, *et al.*, "A Survey of Robot Interaction Control Schemes with Experimental Comparison," *IEEE Transactions on Mechatronics*, vol. 4, pp. 273-285, 1999.
- [8] S. Chiaverini and L. Sciacovico, "The Parallel Approach to Force/Position Control of Robotic Manipulators," *IEEE Transactions on Robotics and Automation*, vol. 9, pp. 361-373, 1993.
- [9] M. H. Raibert and J. J. Craig, "Hybrid position/force control of manipulators," *ASME Journal of Dynamic Systems, Measurement, and Control*, vol. 103, pp. 123-133, 1981.
- [10] F. Conti, *et al.*, "A hybrid actuation approach for haptic devices," in *Proceedings of Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2007, pp. 367-372.
- [11] J. E. Colgate, Schenkel, G., "Passivity of a class of sampled-data systems: application to haptic interfaces," *Journal of Robotic Systems*, vol. 14, pp. 37-47, 1997.
- [12] B. Dellon and Y. Matsuoka, "Feedback Distortion to Augment Controllability of Human Limb Motion," in *Virtual Rehabilitation*, 2008.
- [13] B. Armstrong-Helouvry, *et al.*, "A survey of models, analysis tools, and compensation methods for the control of machines with friction," *Automatica*, vol. 30, pp. 1083-1183, 1994.
- [14] H. Olsson, *et al.*, "Friction Models and Friction Compensation," *European Journal of Control*, vol. 29, pp. 176-195, 1998.
- [15] M. Mahvash and A. Okamura, "Friction Compensation for a Force-feedback Telerobotic System," in *Robotics and Automation*, Orlando, FL., 2006, pp. 3268-3273.
- [16] B. Brogliato, *et al.*, *Dissipative Systems Analysis and Control*, 2 ed.: Springer-Verlag, 2007.
- [17] M. R. Reed and W. J. Book, "Modeling and control of an improved dissipative passive haptic display," New Orleans, LA, USA, 2004, pp. 311-18.
- [18] A. Visioli, *Practical PID Control*: Springer London, 2006.[19] B. Hannaford and J. Ryu, "Time Domain Passivity Control of Haptic Interfaces," in *International Conference on Robotics & Automation*, Seoul, Korea, 2001.
- [20] W. S. Levine, *The Control Handbook*: CRC Press, Inc, 1996.
- [21] H. E. Ross and D. J. Murray, *E.H. Weber: the sense of touch*. New York: Academic Press, 1978.
- [22] R. Isakov, *et al.*, "Influence of sensor quantization on the control performance of robotics actuators," in *International Conference on Intelligent Robots and Systems*, San Diego, CA., 2007.