

Modeling and Control of an Improved Dissipative Passive Haptic Display

Matthew R. Reed and Wayne J. Book

George W. Woodruff School of Mechanical Engineering
Georgia Institute of Technology
Atlanta, GA 30332

mattreed79@hotmail.com

wayne.book@me.gatech.edu

Abstract—This paper investigates the modeling and control of a new dissipative passive haptic display that uses magnetorheological (MR) brakes as actuators. The device is energetically passive, meaning that it can only remove energy from the system. All motive force must be generated by the user, which guarantees stability of the system and safety of the human operator. A first order system approach is presented as a tool for modeling MR fluid behavior in a low-speed braking device. A simulation was developed that uses the brake model and the equations of motion to predict robot motion based on force input and actuator commands. The accuracy of these models was demonstrated with comparisons to experimental data. Two forms of path following velocity control were successfully implemented and shown to significantly reduce path error in preliminary experiments with human operators.

I. INTRODUCTION

Haptic interfaces comprise a broad range of robotic mechanisms that aid humans in performing certain tasks by interacting directly with the user's sense of touch. These force cues help the user to perform operations faster, more safely, and with greater accuracy. Applications for these devices include teleoperation, vehicle control, medical training devices, computer-assisted surgery, and physical therapy.

Large-scale robotic mechanisms that are capable of generating large forces pose a potential risk to the human operator if they become unstable. Stability of the system and safety of the human operator can be ensured by making the device energetically passive. Passive devices are also useful in delicate operations such as robot-assisted surgery, where it is advantageous to keep the motive force in the hands of the surgeon. Unlike energetically active robots that can add energy to the system via electric motors, hydraulic power, and other mechanisms, passive devices may only dissipate, redirect, or store energy using actuators such as brakes, dampers, and continuously variable transmissions (CVTs).

Several examples of passive haptic displays exist in the research arena. PTER (Passive Trajectory Enhancing Robot) is a parallel five-bar linkage that is actuated by four electromagnetic dry friction clutches located on its central axis [1]. The clutches may brake each of the base joints independently or engage direct or reverse coupling of the base joints. This arrangement makes it possible to lock the tip into any one of four local single degree-of-freedom (SDOF) paths

or generate a net force on the tip from some combination of the four clutch torques. A similar device has also been developed that is actuated by two electrorheological (ER) fluid brakes [2]. An alternative approach attempts to steer the operator using only coupling elements. One such example is Scooter, which is a three degree-of-freedom (DOF) passive haptic display controlled by three motorized castors [3]. All of these devices operate in the horizontal plane and are capable of implementing virtual constraints such as desired trajectories and hard obstacles in the workspace. For a detailed comparison of these and other devices, the reader should refer to [4].

A dissipative passive haptic display is limited in that it may only apply force in a direction that opposes velocity. Therefore it is desirable to have more actuators than degrees of freedom to reduce the effects of this limitation. This is demonstrated by the following two-dimensional example. In Fig. 1(a) and 1(b) the tip velocity, v , points to the right. In Fig. 1(a) two actuators are used which may apply force, F_i , in two directions. The shaded region indicates the range of directions at which force cannot be applied. If, for instance, four actuators are used and the directions of resulting forces are evenly spaced, the region of unachievable force directions can be reduced, as shown in Fig. 1(b).




Figure 1. Illustration of unachievable force directions

Devices such as PTER that rely on mechanical contact to apply force to the tip present several control challenges such as vibration, stiction, and slow response times [1]. This paper presents a new dissipative passive haptic display that is actuated by magnetorheological (MR) fluid brakes. MR fluids are colloidal suspensions of magnetizable particles that undergo rapid and reversible changes in rheological properties with the application of a magnetic field. The use of fluid brakes virtually eliminates the jerky feel of the previous system

while decreasing response time. An empirical model of MR fluid brake behavior is presented and incorporated into a simulation that uses the equations of motion of the robot to predict endpoint motion. Next, two forms of velocity control are presented along with an overall control framework. Finally, some preliminary results from human operator experiments are used to evaluate the performance of the two controllers and two robot configurations.

II. EXPERIMENTAL TESTBED

A new dissipative passive haptic display has been developed at the Georgia Institute of Technology. It is a two-DOF parallel five-bar linkage actuated by commercial MR fluid brakes (Lord MRB-2107-3). The device was designed in such a way that two configurations could be tested. The three-brake configuration, shown in Fig. 2, features four equal length arms, collinear base joints (A and B), and brakes located at joints A, B, and E. In the four-brake configuration, displayed in Fig. 3, the base joints are separated such that a non-redundant brake can be added at Joint C. In this configuration, the arms AE and BC are shortened slightly in order to improve the shape of the workspace [5]. The joints in both configurations were designed to rotate only through a limited range such that singular positions are prevented.

Position is measured by two rotary optical encoders (Dynamics Research Corporation HS30 series). The brakes are powered by Advanced Motion Controls 12A8 amplifiers and PS4X300W power supply. Force at the handle is measured by a 6-axis Assurance Technologies FT 15/50 force and torque sensor and controller with analog output.




Figure 2. Three-brake configuration




Figure 3. Four-brake configuration

This device has the capability of locking an actuator and constraining the motion of the user to a single kinematic DOF. The actuator effects can be analyzed by observing plots of available SDOF paths in the workspace. Fig. 4 and 5 show these plots for the three-brake and four-brake configurations, respectively. The three-brake configuration offers greater workspace area, while the four-brake configuration has the benefit of less angular spacing between local adjacent SDOF paths. At each point in the workspace the force at the tip generated by each brake lies perpendicular to the local SDOF path that is achieved by locking that brake.

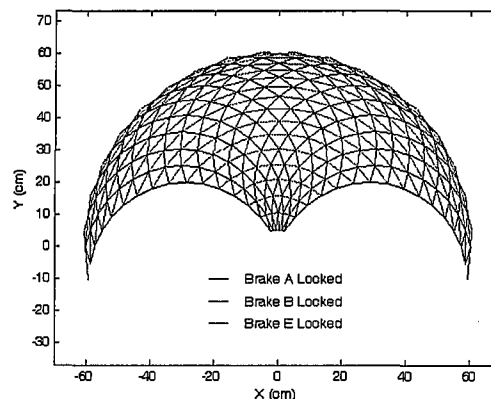


Figure 4. SDOF paths for three-brake configuration




Figure 5. SDOF paths for four-brake configuration

III. MODELING AND SIMULATION

A simulation consisting of two main sections was developed in order to accurately predict the motion of the end effector in response to human force input and actuator commands. The brake model computes the maximum available torque from each brake based upon the applied current and shaft motion. The simulation determines the set of actual braking torques and computes the endpoint motion using the equations of motion for the robot.

A. Actuator Model

Haptic devices such as the one described here typically operate at joint speeds below 60 rpm. Therefore a model is needed that captures the behavior of the MR fluid brake under

these conditions. Each brake was tested by measuring torque in response to shaft speeds controlled by an electric motor with harmonic drive. Initial tests that showed that speed dependence is negligible in the desired operating range and that torque output follows a typical first order system response in reaction to step changes in current. This shows that the dynamics of the brake are based largely on the electrical response since the brake is essentially an RL circuit. It was also observed that the response time increases at very low speeds. This effect can likely be attributed to a certain amount of agitation that must take place in order to initiate the formation and breakdown of particle chains in the fluid.

It is assumed that the amplifiers have a much faster response time than the MR fluid brakes. Therefore each brake and amplifier pair is modeled as a single system. Based upon the initial results, the following first order system model for the MR fluid brake is proposed,

$$\tau(s) = K(I, \theta, \tau_{i-1}) \frac{1}{a(I, \omega)s + 1}, \quad (1)$$

where τ is torque, K is the torque gain, a is the time constant, I is the current, θ is the position, ω is the angular velocity, τ_{i-1} is the torque at the last time step, and s is a differential operator. Equation (1) is not a valid transfer function, but it can be implemented as such in numerical simulation. The terms in appearing in parentheses in the right side of the equation are the inputs to the model and are used to compute the torque gain and time constant at each time step.

The model parameters, K and a , were determined using current steps to the brake with speed maintained constant by the motor. Data was collected for a matrix of current and speed values spanning 0, 0.2, 0.4, 0.6, 0.8, and 1.0 ampere and 0.5, 1, 2, 4, 8, 16, 32, and 64 revolutions per minute, respectively. A logarithmic sequence was selected for speed in order to emphasize the region where time constant is known to increase. Data was collected for three consecutive steps for each pair of current and speed values.

The average steady state torque for each set of current and speed values is shown in Fig. 6. The torque gain was computed as the average steady state torque at each current. A piecewise linear and quadratic function of the following form was applied to the data.

$$K = \begin{cases} k_1 I + k_2, & I < .2 \text{ A} \\ k_3 I^2 + k_4 I + k_5, & I \geq .2 \text{ A} \end{cases} \quad (2)$$

The torque gain coefficients, k_i , were determined by a least squares fit of the steady state torque data. The torque coefficient k_2 is equivalent to the off-state (zero current) torque, which is the minimum torque that each brake will output when in motion.

The time constant was determined from the inverse slope of the following equation using the transient region of the step response data.

$$\frac{t}{a} = -\ln \left(\frac{\tau - \tau_1}{\tau_2 - \tau_1} \right) \quad (3)$$

This equation is derived from the standard first order formulation, with τ_1 as the initial torque, τ_2 as the final torque, and t as time. Again the average of three identical current steps was used to compute each value. The resulting data set for rising time constant is shown as a surface plots vs. current and speed in Fig. 7. The falling time constant follows a similar pattern yet is generally smaller because the amplifier can switch to negative voltage to drive the current down faster.




Figure 6. Average torque gain data and fit equation

Examination of Fig. 7 reveals that the time constant remains relatively constant above a certain speed. Below this threshold, the time constant appears to be a linear function of current and a quadratic function of $\log_2(\omega)$. Therefore the following equation is proposed to model the time constant,

$$a = \begin{cases} a_{avg}, & |\omega| \geq \omega_0 \\ a_1 I + a_2 \omega_L + a_3 I \omega_L + a_4 \omega_L^2 + a_5 I \omega_L^2 + a_6, & |\omega| < \omega_0 \end{cases} \quad (4)$$

where ω_0 is the cutoff velocity and $\omega_L = \log_2|\omega|$. Separate functions were generated for the rising and falling time constants using a least squares fit of the relevant data. The surface generated by this formula for the previous data set is shown in Fig. 8. In the simulation, the time constant is extrapolated using (4) if the speed is outside the tested range.

One peculiar aspect of the brake is that it features backlash "by design" [6]. The backlash is simply a constant deadband through which the shaft must travel when the direction of motion changes where the torque output drops to its off-state value. The backlash was implemented in simulation using the following method. When torque crosses zero, indicating a change in the sign of speed, angular backlash boundaries are set and the torque gain is set to the product of off-state torque and the negative sign of speed. When either of these boundaries is crossed, torque gain is set to its normal value.

The validity of the brake model was evaluated by comparing the model output to physical data for a variety of current and speed inputs. Fig. 9 compares the brake model to an example of the current step data used for model development. Here the model closely matches the data. Fig.

10 shows a comparison of the model output to a sinusoidal motion experiment. In this case, the steady state torque magnitude is underestimated by about 1 Nm. However, this is within the range of torque fluctuation that was observed in testing.




Figure 7. Rising time constant data




Figure 8. Rising time constant surface fit

B. Robot Model

Lagrange's equation was used to obtain the equations of motion for the robot. A total of ten generalized coordinates were used in combination with eight kinematic constraints in order to eliminate the nonlinear terms. The equations of motion can then be expressed in the following form,

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} = \mathbf{Q} + \mathbf{A}^T \boldsymbol{\lambda}, \quad (5)$$

where \mathbf{M} is the mass matrix, \mathbf{A} is the constraint matrix, \mathbf{q} is the vector of generalized coordinates, \mathbf{Q} is the vector of generalized forces, and $\boldsymbol{\lambda}$ is a vector of Lagrange multipliers. For a detailed derivation of the dynamic model the reader should refer to [5].

C. Simulation

The robot simulation was implemented using a combination of Simulink models and MATLAB m-files. The challenge in

this simulation lies in computing each of three or four braking torques in a two-DOF device. The torque output of each brake is determined using an adaptation of Karnopp's friction model,

$$\tau_i = \begin{cases} \tau_{i,y}, & |\omega| \geq \delta\omega \\ \tau_{i,y}, & |\omega| < \delta\omega, |\tau_{i,a}| > \tau_{i,y} \\ \tau_{i,a}, & |\omega| < \delta\omega, |\tau_{i,a}| \leq \tau_{i,y} \end{cases} \quad (6)$$

where τ_i is the brake torque, $\tau_{i,y}$ is the yield torque, $\tau_{i,a}$ is the applied torque, ω is the angular velocity, $\delta\omega$ is a very small angular velocity, and i corresponds to the brake (A, B, C, or E). If angular velocity is very close to zero and the magnitude of applied torque is less than yield torque, the model predicts that the brake will stick. Otherwise, the brake is in slip mode and torque is set to the yield torque from the brake model. If the brake is in stuck mode, it is necessary to compute the torque applied to that joint. If two or fewer brakes are stuck, the equations of motion are used to solve for the applied torque by setting the joint acceleration(s) to zero. If three or more brakes are in stuck mode, braking torques are computed using the lumped actuator approach [7].




Figure 9. Comparison of model to current step data




Figure 10. Comparison of model to sinusoidal motion data

The validity of the robot simulation was tested by affixing a constant force torsional spring to the handle and releasing it from rest at several locations. The force measured by the force

sensor was used as the input to the simulation. Fig. 11 compares the model output to the experimental data in the four-brake configuration with brake A locked. Similar levels of accuracy were found when each of the other brakes was locked and in free motion.

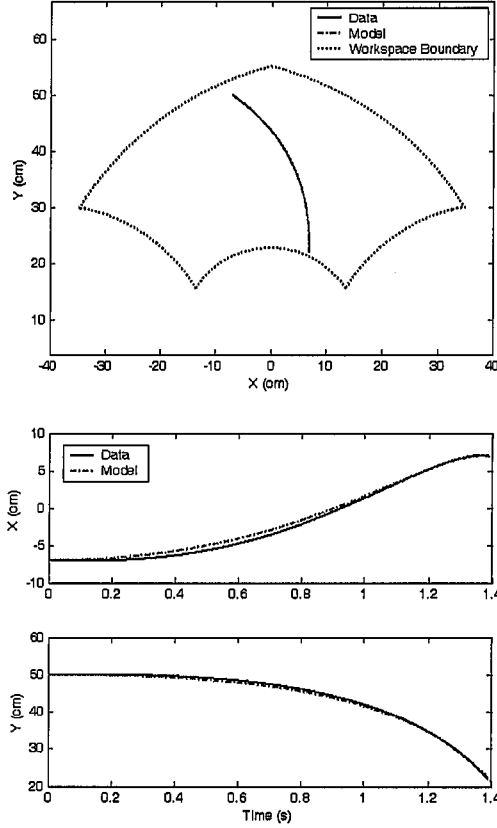


Figure 11. Comparison of robot simulation to experimental data

IV. CONTROL

The controllers investigated in this section focus on the path-following problem. Two forms of velocity control are presented, along with an overall control framework. The main goal of these controllers is to minimize the distance from the end effector to the desired path. Both controllers were tested in simulation and found to significantly reduce path error over the uncontrolled case [5].

A. Velocity Control Formulation

The direction of desired velocity is determined using a velocity field [8]. Outside a specified boundary, the desired velocity always points toward the desired path. Within the boundary, the direction is linearly blended to the direction of the path. The direction of the path and the blend is determined by the direction of force input relative to the path. A sample desired velocity field is shown in Fig. 12, with the desired path set to the line $y = 40$ cm. In this work, the boundary was set at a distance of 5 cm from the path.

Both forms of velocity control actuate only one brake at a time. The controller determines which brake to actuate and the

tip force commanded to that brake. Tip force, F , is related to joint torques, τ , by the following form of the Jacobian, J .

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = J^T \begin{bmatrix} \tau_A \\ \tau_E \\ \tau_B \\ \tau_C \end{bmatrix} \quad (7)$$

J^T can be expressed in the following form [5],

$$J^T = \left[(J_1^{-1})^T \quad (J_2^{-1})^T \right] = \begin{bmatrix} J_{Ax} & J_{Ex} & J_{Bx} & J_{Cx} \\ J_{Ay} & J_{Ey} & J_{By} & J_{Cy} \end{bmatrix}, \quad (8)$$

where J_1 and J_2 are given by the following.

$$J_1 = \begin{bmatrix} -L_{AE} \sin \theta_A - L_{ED} \sin(\theta_A + \theta_E) & -L_{ED} \sin(\theta_A + \theta_E) \\ L_{AE} \cos \theta_A + L_{ED} \cos(\theta_A + \theta_E) & L_{ED} \cos(\theta_A + \theta_E) \end{bmatrix} \quad (9)$$

$$J_2 = \begin{bmatrix} -L_{BC} \sin \theta_B - L_{CD} \sin(\theta_B + \theta_C) & -L_{CD} \sin(\theta_B + \theta_C) \\ L_{BC} \cos \theta_B + L_{CD} \cos(\theta_B + \theta_C) & L_{CD} \cos(\theta_B + \theta_C) \end{bmatrix} \quad (10)$$




Figure 12. Sample velocity field

Since the robot has more actuators than degrees of freedom, (7) cannot be used to solve for a unique set of joint torques from the applied tip force. However, it may be used to analyze the contribution of each brake torque to the tip force using the following equation.

$$|F_i| = |\tau_i| \sqrt{J_{ix}^2 + J_{iy}^2} \quad (11)$$

This relationship was used to normalize commanded force by determining the minimum force magnitude at maximum current that can be exerted by each brake at all points in the workspace. It is also used to relate the force commanded by the controller to a corresponding desired joint torque. The current commanded to the brake is then determined by the inverse of the torque gain equation (2).

B. SDOF Velocity Control

The SDOF velocity controller operates by simply selecting the local SDOF path that most closely matches the desired

velocity vector. The schematic shown in Fig. 13 highlights the key parameters of SDOF control. The desired path in this case is a line at angle α in the workspace. The direction, β , of desired velocity, v_{des} , is determined by the velocity field method. In this figure, four local SDOF paths, p_i , are available. The local SDOF path for each brake lies perpendicular to the force vector contributed by that brake and can be computed by the following.

$$\angle p_i = \tan^{-1}\left(\frac{J_{iy}}{J_{ix}}\right) \pm 90^\circ \quad (12)$$




Figure 13. SDOF velocity control schematic

The controller actuates the brake with the local SDOF path that points in the direction closest to the desired velocity vector. However, the local SDOF path direction must lie in the 90° span between vectors pointing in the direction of the tangent to the desired path and perpendicular toward the desired path. In the case shown in Fig. 13, local SDOF paths p_2 and p_3 are valid, and p_2 would be selected. If a valid local SDOF path exists, the maximum normalized force is commanded to that brake. If there is no local SDOF path in the valid range, one of two control actions is taken. If the input force points away from the desired path, the system is locked by commanding full current to all brakes. Otherwise, no brakes are actuated.

C. Proportional Velocity Control

Proportional Velocity Control was adapted from Swanson's velocity ratio controller [8]. The main difference between the two controllers is that proportional velocity control operates in tip space while velocity ratio control operates in joint space. A schematic of the control technique and the necessary parameters are shown in Fig. 14. The velocity is translated from global coordinates into the space of desired velocity. The goal of the controller is to drive $v_{\beta y}$, the y-component of velocity, v , in the desired velocity space, to zero by applying an appropriate force to the tip. This method assumes that the coupling effects on the link accelerations are negligible.

The direction of each force vector is computed using the Jacobian. However, due to passivity constraints, the force contribution from each brake must act in a direction that opposes velocity. This constraint is indicated by the dotted line shown in Fig. 14. Therefore the following method is used to calculate force angle.

$$\angle F_i = \begin{cases} \tan^{-1}\left(\frac{J_{iy}}{J_{ix}}\right), & F_i \cdot v < 0 \\ \tan^{-1}\left(\frac{J_{iy}}{J_{ix}}\right) + 180^\circ, & F_i \cdot v \geq 0 \end{cases} \quad (13)$$




Figure 14. Proportional velocity control schematic

Once the force direction of each brake is computed, a matrix, \mathbf{H} , of influence coefficients is formed. Each column of the influence matrix corresponds to a different brake. The first row corresponds to $v_{\beta x}$ and the second row correlates to $v_{\beta y}$. Each cell is assigned a value of 1 if the force from that brake will act to increase the magnitude of that velocity component, a value of -1 if the force will act to decrease the magnitude or change the direction of the velocity component, and a value of 0 if the force will have no effect. Although not applicable to this robot, Fig. 14 shows five forces in order to demonstrate all possible cases. In this example, the influence matrix is given by

$$\mathbf{H} = \begin{bmatrix} 1 & 0 & -1 & -1 & -1 \\ -1 & -1 & -1 & 0 & 1 \end{bmatrix}. \quad (14)$$

It is most desirable to select a brake with the influence coefficients $[1 \ -1]^T$. In this case, the brake acts as a coupling actuator that transmits energy from the undesired direction to the desired direction. If this option is unavailable, priority shifts to brakes with influence coefficients of $[0 \ -1]^T$, and then to $[-1 \ -1]^T$. If any other set of influence coefficients are computed, then no brakes act to decrease the magnitude of $v_{\beta y}$, and no brakes are actuated. If multiple brakes share the best pair of influence coefficients, the brake with corresponding force pointing in the direction closest to the desired velocity is selected. In concurrence with the controller goal, that brake is then actuated according to the following control law,

$$F_i = K_y |v_{\beta y}|, \quad (15)$$

where K_y is the proportional velocity controller gain. This controller can also lock the system under either of two conditions. The first results when no brake is selected to actuate and the input force points away from the desired path. The second occurs when $v_{\beta x}$ is less than zero.

V. PRELIMINARY TESTING

In a haptic device, it is necessary to investigate human factors such as workload and smoothness in addition to position error. Therefore, a set of preliminary tests with human operators was conducted in order to evaluate the two forms of velocity control and the two robot configurations.

A. Experiment Design

The system was configured as shown in Fig. 15, with the monitor displaying the desired path and the start and end points. Each operator performed three trials for each combination of two separate line segments, two configurations, and five controllers. The operators consisted of three male graduate students with an interest in haptics research.

The five controllers are outlined in Table I. The SDOF controller was tested as previously described and also with a gap surrounding the desired path within which no control action is taken. This capability was added to reduce the vibration encountered by the rapid switching of brakes when the tip closely follows the desired path. The proportional velocity controller was tested with two gains. The high gain ensures maximum normalized force is commanded by each actuated brake. The low gain was chosen such that there was no detectable vibration.



Figure 15. Experimental setup for human operator tests

B. Results

Three performance metrics were used to quantify the results from the human operator experiments. The path average position error measures controller accuracy and allows comparison among different paths. Also examined was average tip force magnitude, which Swanson found correlated strongly with user workload [4]. Additionally, Swanson observed a strong inverse correlation between the frequency content of acceleration and perceived smoothness. A weighted sum of the discrete Fourier transform (DFT) of acceleration was used as this performance measure.

The goal of these controllers is to minimize all three performance metrics. The performance metrics are

summarized in Fig. 16, 17, and 18. All trials for each combination of controller and configuration are combined and presented as box and whisker plots. The middle line in each box represents the median of the data. The upper and lower boundaries of the box represent the upper and lower quartiles. The whiskers span the entire range of data excluding outliers. Outliers are defined as points that are a distance of more than 1.5 times the span of the box away from the box boundary.

Fig. 16 compares the average path error for each combination of controller and configuration. Both controllers reduced path error significantly, regardless of the gain used. However, it is worthwhile to note that the SDOF velocity controller with uncontrolled gap resulted in far greater consistency in accuracy than the other systems.

Fig. 17 and 18 summarize the average force data and DFT sum of acceleration data, respectively. As expected, adding an uncontrolled gap to the SDOF controller increased smoothness and decreased workload by sacrificing a negligible amount of accuracy. Likewise, by lowering the gain of the proportional controller, a significant reduction in workload and increase in smoothness was realized with little change in accuracy. In fact, the low gain proportional controller exhibited similar levels of workload and smoothness to the uncontrolled experiment. Overall the proportional controller exhibited lower workload and higher levels of smoothness than the SDOF controller.

Further conclusions can be made by comparing the results between the two robot configurations. As Fig. 16, 17, and 18 indicate, the four-brake configuration did not offer any measurable performance gains over the three-brake configuration. Also, the path error for the uncontrolled four-brake configuration varied over a much wider range than the uncontrolled three-brake configuration. This suggests that the dynamics of the three-brake configuration are such that it is more intuitive for the operator to maneuver accurately than the four-brake configuration.

VI. CONCLUSION

This work introduced a modeling procedure that characterizes the behavior of MR fluids in a low-speed braking device. The comparison of the brake model to experimental data indicates that the first order model accurately predicts braking torque for steps in current and is satisfactory for other types of inputs. The actuator model was incorporated into the equations of motion of an arbitrary parallel five-bar linkage. Comparisons of the simulation output to experimental data showed that the simulation is an effective tool for predicting robot motion and evaluating control techniques.

A preliminary round of testing with human operators was used to evaluate each control technique and configuration. All of the controllers offered a similar improvement in accuracy over the uncontrolled case. The main differences were observed in workload and smoothness, where the low gain proportional velocity controller was superior. Also, the results suggest that the four-brake configuration offers no performance advantages over the three-brake configuration. It should be emphasized that these results are preliminary, and may vary under a different set of testing conditions.

TABLE I. CONTROL METHODS

Number	Type	Gap / Gain
1	None	--
2	SDOF without gap	0 cm
3	SDOF with gap	.50 cm
4	Proportional (high gain)	10,000 Ns/m
5	Proportional (low gain)	5 Ns/m




Figure 16. Average Path Error by Configuration and Controller




Figure 17. Average Tip Force by Configuration and Controller




Figure 18. DFT Sum of Acceleration by Configuration and Controller

This work also introduced a new dissipative passive haptic display that features MR fluid brakes as actuators. This strategy eliminates many of the hassles associated with friction brakes including slow response times, stiction, and vibration. The two controllers presented here are the first to be implemented on this device and the results are very promising. Work will likely continue in the development of control techniques, particularly in the area of obstacle avoidance. A more comprehensive investigation of the performance of this robot will likely follow. Plans are also underway to use this device as a passive master in the teleoperation of an active slave.

ACKNOWLEDGMENT

This work was partially supported by the HUSCO/Ramirez Chair in Fluid Power and Motion Control.

REFERENCES

- [1] Swanson, D. and W. Book, "Torque feedback control of dry friction clutches for a dissipative passive haptic display," *Proceedings of the IEEE International Conference on Control Applications*, pp. 736-741, Anchorage, AK, September 2000.
- [2] Sakaguchi, Masamichi, Junji Furusho, and Naoyuki Takesue, "Passive Force Display using ER brakes and its control experiments," *Proceedings of the IEE Virtual Reality Conference*, (Yokohama, Japan), March 2001.
- [3] Moore, Carl A., Michael A. Peshkin, and J. Edward Colgate, "Design of a 3R cobot using continuously variable transmissions," *Proceedings of the IEEE International Conference on Robotics and Automation*, (Detroit, Michigan), May 1999.
- [4] Swanson, D., *Implementation of Arbitrary Path Constraints Using Dissipative Passive Haptic Displays*, PhD thesis, April 2003.
- [5] Reed, M., *Development of an Improved Dissipative Passive Haptic Display*, Master's thesis, Georgia Institute of Technology, November 2003.
- [6] Lord Corporation website, <http://www.mrfluid.com>, Accessed 26 October 2003.
- [7] Swanson, D., E. Romagna, W. Book, and A. Barraco, "Influence of actuator dynamics on passive haptic interface performance," *Proceedings of the International Conference on Advanced Intelligent Mechatronics*, pp. 440-445, Atlanta, GA, September 1999.
- [8] Swanson, D. and W. Book, "Path-Following Control for Dissipative Passive Haptic Displays," *International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, Los Angeles, CA, 2003.