The Effects of Real and Computer Generated Friction on Human Performance in a Targeting Task

Christopher Richard and Mark Cutkosky
Stanford University Center for Design Research
Bldg. 560 Panama Street
Stanford, CA 94305-2232
crichard@cdr.stanford.edu

ABSTRACT

We present the results of haptic interaction with real and virtual Coulomb friction. Tests were conducted to explore the effects of various types of friction on human performance in a Fitts-style targeting task. The results indicate that haptically rendered friction affects subject performance in a manner quite similar to that of real physical friction. Furthermore, moderate low-stiction friction tends to improve subject performance both in terms of speed and accuracy. High-stiction friction, however, degrades subjects performance, especially in terms of speed.

INTRODUCTION

Friction is present to some degree in all mechanisms; nevertheless, it is often absent or greatly simplified in virtual simulations of mechanisms. The presence of friction undoubtedly introduces complications into mechanical systems. A mechanism with high friction will require more energy to operate and will therefore be less efficient. Stiction, friction in which the static value is higher than the kinetic value, is often the bane of controls engineers. In servo-systems using integral control, friction can cause the system to “hunt” for its goal position. There are however some benefits associated with friction. Because friction dissipates energy, it tends to have a stabilizing effect similar to that of viscous damping.

Does the presence or absence of friction in virtual simulations affect the user of a haptic interface? More specifically, how does friction affect human performance in positioning or targeting tasks? To answer these questions, we have implemented a haptic rendering of a block sliding along a surface with friction interacting between the two. The design of our apparatus is such that we can easily switch between real and simulated friction. The implementation allows us to change the mass of the sliding block as well as the friction model parameters, such as the level of the static friction, the level of the dynamic friction and a viscous damping coefficient. The friction model used for this work is the based on the model developed by Karnopp (1985) and discussed in Richard et al. (1999).

With our haptic rendering of friction implemented, we conducted a series of human subject tests to examine the effects of friction on subject performance in a targeting task. Similar to the method used by Fitts (1954), subjects moved between targets of various sizes and spaced at various distances. We examined changes in subjects' task completion times, error rates, and average error magnitudes as they moved a cursor through a virtual environment with various types of frictional resistance. In this paper, we provide the details of our haptic friction rendering and present results of subject performance as they acquired targets in environments with real physical friction, and various types of haptically rendered friction.

Friction Models applied to haptics

There are several models of friction presented in the literature. Armstrong-Helouvry et al. (1996) provide a thorough review. For haptic rendering, the three models reported most often in the literature are variations of:

- the bristle model (Hessig and Friedland, 1991),
- the Dahl model (Dahl, 1976)
- and the Karnopp Model (Karnopp, 1985)

Chen et al. (1997) render friction and adhesion in a manner similar to the bristle model. Using a single bristle in their haptic rendering, the authors indicated that they were less than satisfied with the fidelity of the implementation.

Hayward and Armstrong (2000) introduced a friction model for haptics based on the Dahl model, but with several improvements. Most notable among them is that the Hayward-Armstrong variant does not exhibit the spurious position drift that can occur in the original Dahl model. In its most general form, the Hayward-Armstrong model displays four regimes of friction dubbed sticking, creeping, slipping and sliding. In its simplest incarnation the model is quite similar to the single bristle model of Chen et al.(1997).

Salcudean and Vlaar (1994) and Berkelman (1999) simulate haptic friction in a manner based upon the Karnopp model. The implementation is well suited for emulating the stick-slip nature of friction. In both implementations however kinetic friction was ignored and only viscous damping was rending once the system was sliding.

EXPERIMENTAL APPARATUS

The apparatus used for this work is the same as the one used by Richard et al.(1999) for friction identification. It is shown in Figure 1. It was modified by the addition of a mouse shell riding on a low friction linear bearing. A wrist rest around the base of the mouse prevented subjects from dragging their fingers along the apparatus structure in order to generate more friction as they performed the task. Because of the wrist rest, subject motion was generated more from the elbow and shoulder than it is on a standard computer mouse.

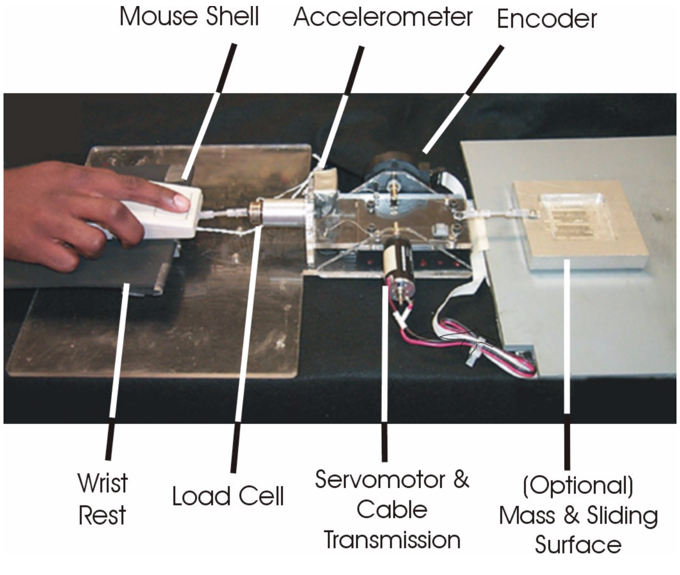

Figure 1. Experimental Apparatus

DESCRIPTION OF THE HAPTIC RENDERING

Karnopp's Model

Our implementation of virtual friction is based on the Karnopp friction model (Karnopp 1985, Richard et al. 1999). We begin by generating the equation of motion for a block sliding on a surface with a single degree of freedom

$$F_a - F_f = m\ddot{x} \tag{1}$$

where

F_a is the force being applied to the block;
F_f is the friction between the block and the surface;
m is the mass of the block, and;
$\ddot{x}$ is the acceleration of the block.

According to the Karnopp friction model, the block must be in one of two possible states: STUCK or SLIDING. In the STUCK state, the friction force, F_f, between the block and the surface on which it is sliding is a function of the applied force, F_a. If the magnitude of the force applied to the block is less than that required to overcome the static friction, F_s,the block remains stationary. The frictional force in this case is then equal to the applied force.

$$F_f = -\mathrm{sgn}(F_a)F_a \qquad (F_a < F_s) \tag{2}$$

If, however, the applied force exceeds the static friction force, the block begins to move; it has entered the SLIDING state. In the SLIDING state, the friction force between the block and the surface depends on the block's velocity, $\dot{x}$. The friction force in the SLIDING state is

$$F_f = -(\mathrm{sgn}(\dot{x})F_d + b\dot{x}) \qquad (|\dot{x}| > \Delta v) \tag{3}$$

where

F_d is the dynamic value of the Friction;
b is the viscous damping coefficient, and;
Δv is a velocity threshold.

If the magnitude of the block's velocity falls below the threshold, Δv, the block re-enters the STUCK state and its velocity is set to zero.

Simulating mass

In our virtual simulation, we wish to simulate the inertial and frictional forces of the block as it slides. Since the mass of the block we wish to simulate is known, we can choose to simulate the mass as an impedance (where acceleration is the input variable and force is the output variable) or as an admittance (where force is the input variable and acceleration is the output). To simulate the inertia of the block using impedance causality, we would need an estimate of the user's input acceleration. This is unfortunate as accurate acceleration estimates can be difficult to attain. Digital encoders provide reasonable position sensing but introduce errors when the position signal is differentiated to obtain velocity or acceleration estimates. Accelerometers are also prone to errors due to drift and noise.To simulate the inertia of the block using admittance causality, we would need to know the force being applied to the block. While force sensors are often used successfully in control situations, simulation via admittance causality is best suited for haptic interfaces that are not backdrivable.

What then is an effective way of simulating both inertia and frictional forces with a haptic interface? By using a virtual coupling (Colgate and Stanley, 1995, Adams and Hannaford 1999) we are able to effectively simulate both the friction and inertia of the block without estimating either the user's input acceleration or the user's input force.

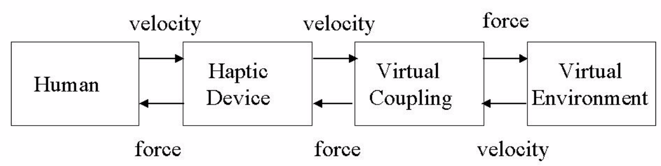

Figure 2. The virtual coupling (block diagram adapted from Adams and Hannaford (2000))

A block diagram is helpful in explaining the functionality of the virtual coupling (Figure 2). It connects the physical haptic device (which in our case has an impedance causality) to the virtual environment (presented with admittance causality.) For friction rendering we used a virtual coupling that consisted of a virtual spring (K_{VC} = 25000 N/m)and dashpot (B_{VC} = 10Ns/m).

The algorithm for the friction rendering is as explained below. The haptic rendering was implemented on a 200Mhz Pentium II running at a servo rate of 1000Hz. The initial values for the position and velocity of the haptic device and of the block are set to zero. The initial force in the virtual coupling is also set to zero. The block is initially in the STUCK state. Once each time step the following actions are taken:

- Obtain position and velocity of the haptic interface.
- Calculate the friction force and state of the block based on previous state and velocity of the block.
- Calculate the net force being applied to the block based on force in the virtual coupling and the frictional force.
- Calculate the block's acceleration based on the applied force and the block's mass,

$$\ddot{x} = \frac{F_a - F_f}{m} \tag{4}$$

- Integrate to find the block's velocity and position

$$\begin{aligned} \dot{x} &= \int \ddot{x}\, dt + \dot{x}_0 \\ x &= \int \dot{x}\, dt + x_0 \end{aligned} \tag{5}$$

- Calculate the force in the virtual coupling

$$F_{vc} = K_{vc}(p - x) + B_{vc}(v - \dot{x}) \tag{6}$$

where

p is the position of the haptic device, and;

v is the velocity of the haptic device.

- Apply the force in the virtual coupling to the human operator.

EXPERIMENT DESCRIPTION

In a targeting task, the user of a haptic interface is required to acquire a target by moving a cursor to the location of the target and pressing a button or pulling a trigger. How does the presence or absence of friction, either in the mechanism itself or in the virtual environment, affect a user's performance in a targeting task? The answer to this question may depend on many factors. For example, one might expect a moderate amount of friction to be helpful. It may dampen a subject's involuntary hand movements and help him or her acquire a target more quickly. Conversely, the presence of significant friction may actually slow a subject; more work and larger muscle forces are required to move the system when large amounts of friction are present. Finally, if the value of static friction is significantly higher than the dynamic value of friction (a high stiction type of friction), we may expect a subject to have difficulty homing in on the target. High stiction may cause the subject to "hunt" around the target area much like a servo-system hunts around its goal position in the presence of stiction.

A targeting task is useful for comparing the effects of real versus virtual friction. Asking subjects to discriminate between real and virtual friction in a subjective manner (e.g. exploring two environments, one with real friction and one with virtual friction and asking subjects which is which) will nearly always favor the real friction over the virtual friction. It is too stringent a criterion on which to base the quality of a haptic simulation. Humans are quite adept at exciting incipient instabilities in haptic simulations and at detecting subtle phenomenon caused by sensor resolution and actuator limitations.

Rather than asking subjects to compare a real friction and virtual friction directly, we present them with a targeting task. The idea being that the task at hand will mask any irrelevant shortcomings of the simulation and provide a comparison of the real phenomena and the virtual phenomenon under the context of the task. If the subjects' performance with simulated friction is similar to their performance with real friction, we can infer at a minimum, that the virtual friction similarly affects the user's performance for the task of interest. In fact, subjects did comment that the real and simulated friction were very close.

To determine how the presence of Coulomb friction affects a user's performance in a targeting task, and to compare the effects of real friction and simulated friction in an objective manner, we conducted a series of Fitts-type targeting tasks (Fitts 1954). In a Fitts test, subjects are required move back and forth between two targets in rapid succession. Fitts defines the index difficulty, I_d, for targeting task as:

$$I_d = -\log_2 \frac{W_s}{2A} \qquad W_s = W - D \tag{7}$$

where

A is the distance between the targets;
W is the width of the targets, and;
D is the diameter of the cursor.

The concept is that a targeting task becomes more difficult as the amplitude of the required motion, A, becomes larger, and as the required accuracy increases (W_s becomes smaller)

For our friction based Fitts tests, subjects were instructed to move a cursor vertically on the screen by manipulating the computer mouse attached to the experimental apparatus. Their task was to move the cursor back and forth between two colored regions on the screen (See Figure 3). They were instructed to click the mouse button once the cursor was completely within the bounds of the green region. If subject clicked the mouse button with the cursor completely within the bounds of the green region, a valid click was registered. The green target would turn red, and the red target would become green. If the subject clicked the mouse button while the cursor was not within the bounds of the green target, an error click was recorded. After making a valid click, the subject would then proceed to the new green target and click the mouse button at the appropriate time. Eight valid clicks were required to complete one segment of the Fitts test. Subjects were asked to complete the task as quickly as possible.

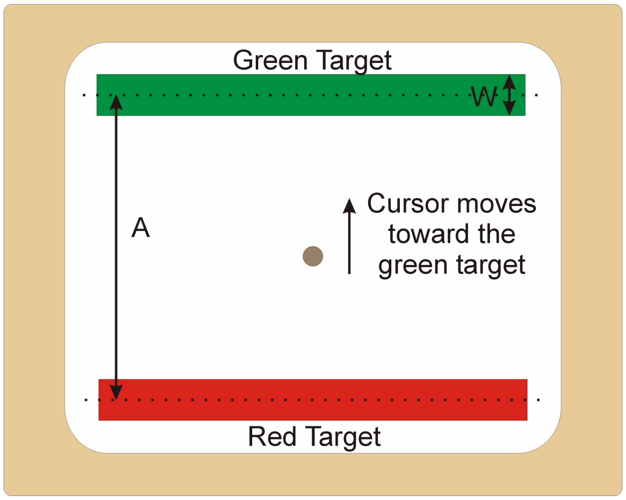

Figure 3. Screen seen by subjects during the Fitts test.

Twenty right-handed subjects (10 male and 10 female) participated in this experiment. Each subject completed 45 blocks of Fitts tests. The 45 blocks included nine different indices of difficulty (see Table 1) and five variations of frictional resistance to the subjects motion (see Table 2). Before beginning the experiment subjects practiced the baseline test, and the test with real friction added to the system (an aluminum block sliding on a rubber pad). After practicing, subjects began the experiment by completing nine blocks of Fitts tests with the baseline resistance. Subjects completed nine blocks of the remaining four resistances (real friction, simulation, high stiction, and viscous damping) in random order. In each case the order of the blocks was also presented in random order. After each block, we recorded the subject's completion time, number of errors, and average error magnitude (in millimeters).

Table 1: Indices of difficulty for Fitts test

	A *pixels(mm)*	W_s *pixels(mm)*	I_d
1	228(342)	12(18)	5.25
2	228(342)	40(60)	3.51
3	228(342)	68(102)	2.75
4	300(450)	12(18)	5.64
5	300(450)	40(60)	3.91
6	300(450)	68(102)	3.14
7	380(570)	12(18)	5.98
8	380(570)	40(60)	4.25
9	380(570)	68(102)	3.48

Table 2: Frictional resistances used for the Fitts test

	Resistance	Description
1	Baseline	No frictional resistance is added. Subjects interact only with the friction (approx. 0.8N) and mass (approx. 1.4kg) inherent in the haptic interface
2	Real Friction	A 0.5kg block of aluminum sliding on a rubber pad. (Friction Force is approx 3.5N)
3	Simulated Friction	A simulated level of friction that approximately matches case 2.
4	High Stiction	Similar to case 3 except that the value of static friction is set to a higher value (approx. 7.0N)

Table 2: Frictional resistances used for the Fitts test

	Resistance	Description
5	Viscous Damping	Subjects interact with virtual viscous damping of with a damping coefficient of approx. 35Ns/m and mass of 0.5 kg.

RESULTS

We wish to examine the subjects' results on the Fitts test by comparing their performance on a given test with added friction with their performance on the same test under the baseline condition. For each index of difficulty and for all added resistances, we subtract the subjects' baseline performance on the same index. The results are a subject's time difference, difference in number of errors and difference in average error magnitude.

As is often the case when human subjects are involved, the person to person variability was high. For example, in some cases most of the subjects showed improvement in performance, but a few show a decrease in performance. Average scores and variations were therefore not particularly instructive. Nonetheless, some clear qualitative trends can be seen in the data. These trends are best illustrated with scatter plots in Figures 4-7.

Plotting the difference in error (either number of errors or error magnitude) versus the difference in time we obtain a graphical comparison of performance with added friction versus baseline performance. Figure 4 shows how one can interpret each quadrant of the error/time plane. Near the origin, the presence of added friction makes little difference in subject performance. In quadrant 1, added friction negatively impacts subject performance. Data in this quadrant indicates that the presence of friction caused subjects to perform more slowly and with less accuracy. Quadrant 3, on the other hand, indicates that the presence of added friction results in performance enhancement. Data in this quadrant indicate that friction helped the subject perform the task more quickly and with fewer errors.

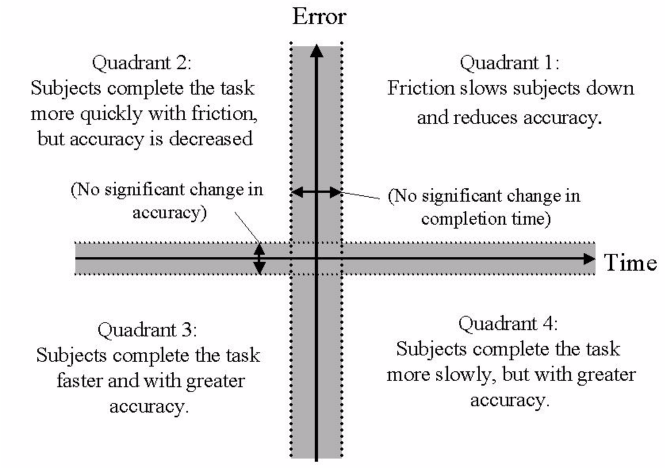

Figure 4. Possible effects of friction on subject performance.

Figures 5, 6, and 7 compare subjects' performance for the three most difficult indices of difficulty 1, 4, and 7. For index 1 (Figure 5) we see a clustering of data points in quadrant 3. This indicates, that for this index, friction generally helps subjects perform the Fitts tests. Completion times were 17% faster for real friction and 23% faster for simulated friction. Subjects committed an average of 1.4 fewer errors for real friction and 2.8 fewer errors for simulated friction. . There is however a small cluster of data points in quadrant 1. Most of these data represent high stiction cases.

For index 4 (Figure 6) the results are similar to those for index 1. We see more clustering in quadrant 1, but again this clustering is mostly due to the high stiction case. Subjects took 64% longer to complete index 4 and committed 1.35 more errors under the high stiction case.

For index 7 (Figure 7) we see most of the data clustered around the origin indicating that the presence or absence of friction has made little difference from a performance standpoint. We do see, however, some high stiction outliers in quadrant 1 as was the case in indices 1 and 4.

It is not surprising that the high stiction case degrades subjects' performance. As mentioned previously, high stiction can cause a servo-system to "hunt" about its goal position. A similar behavior was observed in several subjects as they attempted to location the cursor within the bounds the of the target. Figures 8 and 9 show the position trajectory for a subject completing index 4 under baseline conditions and under the high stiction case. In the high stiction case we see the subject repeatedly missing the target by overshooting while trying to adjust to the goal position. Because the static value of the

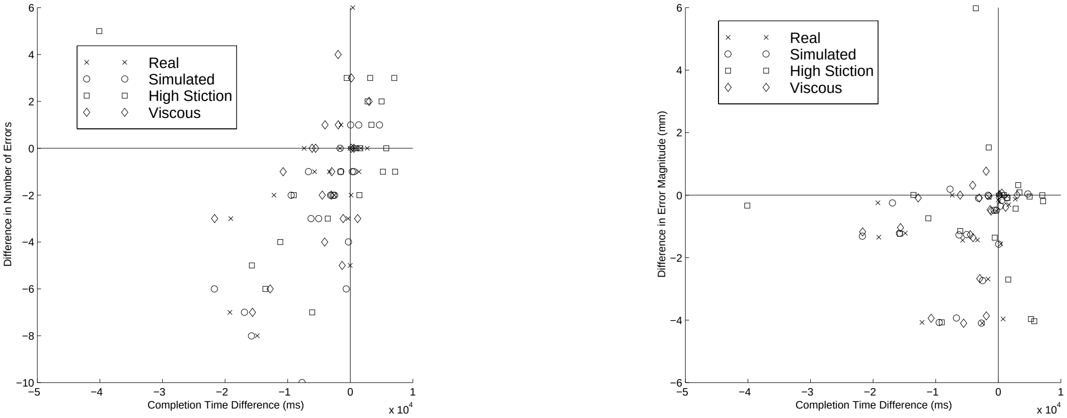

Figure 5. Subjects' performance relative to their baseline case for index of difficulty #1 (a) Number of Errors versus Time (b) Average Error Magnitude versus Time.

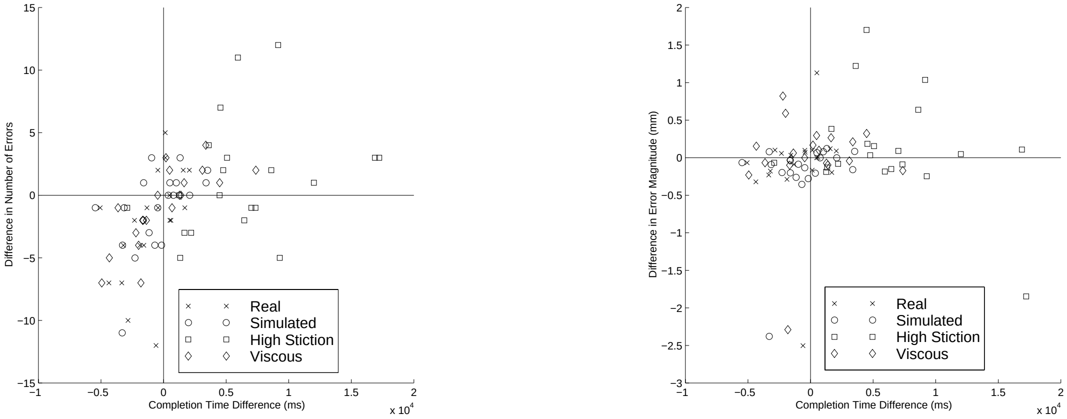

Figure 6. Subjects' performance relative to their baseline case for index of difficulty #4 (a) Number of Errors versus Time (b) Average Error Magnitude versus Time.

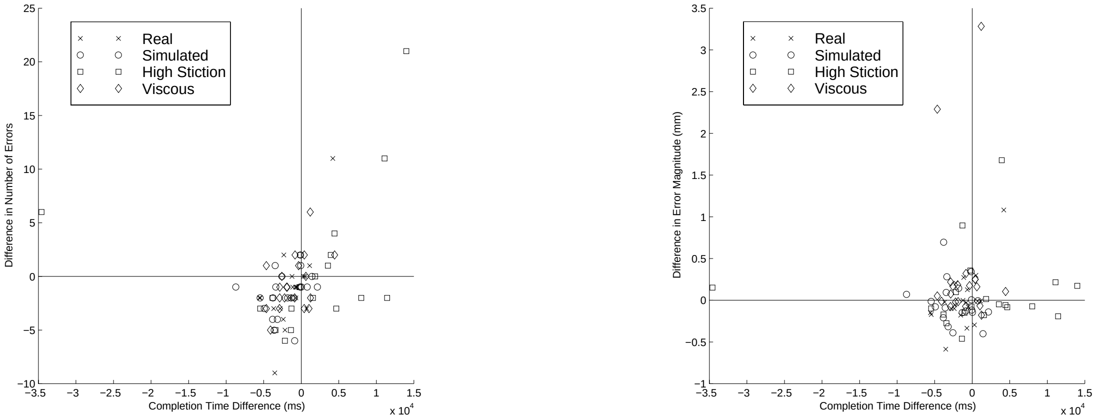

Figure 7. Subjects' performance relative to their baseline case for index of difficulty #7 (a) Number of Errors versus Time (b) Average Error Magnitude versus Time.

friction is significantly higher than the dynamic value, subjects found that the amount of friction necessary to “break-away” was more than they wanted to accelerate the mass once it was free. Because of this overshooting and undershooting, we see that it takes the subject longer to complete each cycle of motion in the presence of stiction (compare the subject’s seven target entries in Figure 8 with the three target entries in Figure 9 for the same 800 ms of time).

Tables 3, 4, and 5 show the number of subjects that performed better and worse than their base case. Performance improvements or degradation is based on task completion time, number of errors committed, and the average error magnitude. In terms of time (Table 3), subjects were considered to have performed better if their completion time was more than 0.5sec faster than their baseline time. They were marked as performing worse if their time was more than 0.5 sec. slower than their baseline time. We see from the table that subjects performed much worse in the high stiction condition.

In terms of the number of errors committed (Table 4), subjects were considered to have improved if they committed at least 1 fewer error than they did on the baseline case. They were marked worse if they committed as least one more error than they did on the baseline case.

For error magnitude, subjects must have had an average error magnitude at least 1 mm smaller than their base case to register an improvement; if their average magnitude as greater than 1 mm larger than their base case, their performance was declared worse.

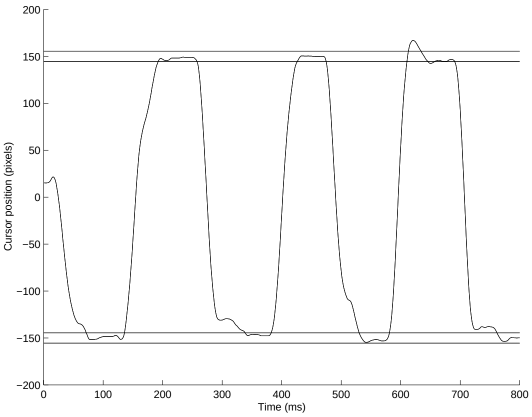

Figure 8. Subject’s trajectory on a the baseline case. The subject was able to acuire the target 7 times in 800ms.

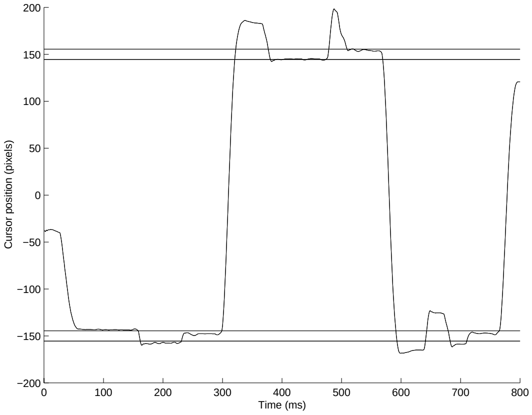

Figure 9. Subject’s trajectory on a High Stiction case--note the difficulty positioning the cursor with the target bounds.

Table 3: Number of subjects (out of 20) that performed better and worse based on task completion time.Tests with highest index of difficulty are highlighted.

Index	*Real*		*Simulated*		*High Stiction*		*Viscous*	
	Better	*Worse*	*Better*	*Worse*	*Better*	*Worse*	*Better*	*Worse*
1	10	4	14	3	9	11	15	3
2	6	10	5	9	3	9	8	7
3	8	4	5	5	5	9	7	4
4	10	5	10	7	1	19	9	8
5	3	8	6	4	3	9	6	5
6	5	6	3	4	3	7	2	5
7	15	3	13	3	8	10	12	5
8	8	6	12	3	8	9	11	3
9	9	4	7	5	8	5	10	4

Table 4: Number of subjects that performed better and worse based on the number of errors committed.

Index	*Real*		*Simulated*		*High Stiction*		*Viscous*	
	Better	*Worse*	*Better*	*Worse*	*Better*	*Worse*	*Better*	*Worse*
1	11	2	16	3	10	7	11	5
2	1	8	3	9	5	7	3	7
3	4	5	5	5	3	5	4	5
4	14	5	10	7	8	10	10	7
5	7	2	6	6	5	5	6	6
6	5	2	3	7	3	4	4	5
7	14	3	15	3	12	7	12	6
8	6	7	9	6	6	11	9	4
9	4	5	5	5	5	5	4	9

Table 5: Number of subjects (out of 20) that performed better and worse based on average error magnitude. Tests with highest index of difficulty are highlighted.

Index	*Real*		*Simulated*		*High Stiction*		*Viscous*	
	Better	*Worse*	*Better*	*Worse*	*Better*	*Worse*	*Better*	*Worse*
1	17	1	17	2	14	5	13	5
2	4	8	6	7	6	7	6	7
3	5	5	5	5	3	7	4	6
4	9	10	14	5	9	11	9	10
5	10	3	4	10	9	5	8	8
6	5	3	4	8	4	5	4	7
7	15	5	13	7	12	8	7	13
8	10	5	9	7	10	9	7	7
9	6	6	6	5	6	6	6	7

DISCUSSION AND CONCLUSIONS

In this paper, we present results obtained with simulated friction in a haptic interface. The method allows a backdrivable haptic device to simulate both inertial and frictional forces without measuring acceleration or force.

Using our haptic rendering of friction as well as real physical friction, we have compared the effects of friction on human performance in a targeting tasks. Our results indicated that a moderate amount of low-stiction friction can improve human performance in a Fitts type targeting task in terms of both speed and accuracy. This result was observed whether the subjects were experiencing real or virtual friction. Anecdotally, subjects indicated a preference for the frictional cases over the baseline case. Many indicated that the frictional cases “felt better.” Subjects also noted that simulated and real low-stiction cases felt very simular.

Not surprisingly, the results also indicate that high stiction friction negatively impacts subject performance, particularly with regard to task completion time, as subjects had difficulty acquiring narrow targets in a high stiction environment.

ACKNOWLEDGMENTS

We wish to thank Matt Hage and Michael Koenig for their assistance with the design and implementation of the Fitts test as well as their assistance with conducting the human subject tests.

REFERENCES

Fitts, P. M., 1954, “The Information Capacity of the Human Motor System in Controlling the Amplitude”, Journal of Experimental Psychology, V. 47, n. 6, 1954. pp 381 - 3916

Richard, C., Cutkosky M. R., and MacLean K. E., 1999, “Friction Identification for Haptic Display,” Proc. of the ASME Dyn. Sys. and Ctrl., Vol. 67, pp. 327-334.

Colgate, J.E., Stanley, M.C., 1995, “Issues in the haptic display of tool use,” Proc. IEEE/RSJ Int. Conf. Intel. Robot Sys. Pittsburgh, PA, pp. 140-145.

Hayward, V., and Armstrong, B., 2000, “A New Computational Model of Friction Applied to Haptic Rendering,” *Experimental Robotics VI*, Peter Corke and James Trevelyan (Eds), pp. 403-412, Lecture Notes in Control and Information Sciences, Vol. 250, Springer-Verlag

Chen, J., DiMattia, C., Taylor II, R. M., Falvo, M., Thiansathon, P., and Superfine, R., 1997, “Sticking to the Point: A Friction and Adhesion Model for simulated surfaces.” Proc. of the ASME Dyn. Syst. and Ctrl., DSC-Vol. 61, pp. 167-171.

Dahl, P. R., 1976, “Solid Friction Damping of Mechanical Vibrations,” AIAA Journal, v. 14, n. 12, pp. 1675-1682.

Karnopp D., 1985, “Computer simulation of stick-slip friction in mechanical dynamic systems,” ASME Journal of Dynamic Systems, Measurement and Control, Vol. 107, pp. 100-103.

Hasser C. J., Goldenberg, A. S., Martin, K. M., and Rosenberg, L. B., 1998, “User Performance in a GUI Pointing Task with a Low-Cos Force-Feedback Computer Mouse,” Proc. of the ASME Dyn. Syst. and Ctrl., DSC-Vol. 64, pp. 151-156.

Salcudean, S. E. and Vlaar, T. D., 1997, “On the emulation of stiff walls and static friction with a magnetically levitated input/ output device,” Trans. ASME J. Dyn. Syst., Meas., Ctrl., vol. 119, no. 1, pp. 127-132.

Adams, R. J., and Hannaford, B.,1999, “Stable Haptic Interaction with Virtual Environments,” IEEE Trans. Robotics and Automation, vol. 15, no. 3, pp.465-474.

Haessig, D. A., Friedland, B., 1991, “On the modeling and simulation of friction.,” ASME J. Dyn. Syst. Meas. and Ctrl., Vol. 113, pp. 354-362.

Berkelman, P., 1999, “Tool-Based Haptic Interaction with Dynamic Physical Simulations using Lorentz Magnetic Levitation,” Ph.D. Thesis, Robotics Institute, Carnegie Mellon University.