Simulating Cheap Hardware:

A platform for evaluating cost-performance trade-offs in haptic hardware design

Iman Brouwer
Dept. of Mechanical Engineering
University of British Columbia
Vancouver, Canada
ibrouwer@mech.ubc.ca

Karon E. MacLean
Dept. of Computer Science
University of British Columbia
Vancouver, Canada
maclean@cs.ubc.ca

Antony J. Hodgson
Dept. of Mechanical Engineering
University of British Columbia
Vancouver, Canada
ahodgson@mech.ubc.ca

***Abstract*—This paper describes a platform devised to explore the impact on task execution in a virtual environment of the quality, and therefore cost, of the system's haptic hardware. This platform is a complex haptic interface in which hardware quality can be varied in simulation. Software intercepts the position and force signals between the haptic hardware and the virtual environment software, and alters them to supply the effect of increased friction, cogging, backlash, inertia and/or lower force output. All parameters of the introduced effects can be set independently or in combination and on a continuous scale; a primary contribution is the creation of haptically realistic effect models that are stable in combination on complex hardware. This work is part of a larger project in which we will test the effect of the simulated degradations on the ability of surgeons-in-training to learn basic laparoscopic skills.**

I. Motivation & Approach

In recent years, haptic interfaces have moved from research labs into commercial applications ranging from force feedback joysticks to automotive controls and surgical simulators. These items have widely varying price tags – a force feedback joystick costs US$60, a surgical simulator more than 300 times as much. Why? The surgical simulator hardware is more complex than the joystick, and the latter also benefits from a much higher production volume. Here, we focus on the factor of design aim. The joystick is expected only to provide a certain rather crude haptic effect, while the surgical simulator is designed to produce a precise force output at the handle. To achieve this, designers use components known to minimize noise-generating characteristics such as friction, cogging and inertia. A single such motor can cost more than three complete 2-DOF gaming interfaces. There is thus a financial motive to know whether the higher performance obtained actually makes a difference in task execution for this type of application.

A. Goal

The primary question we want to answer is: "How far can we degrade haptic quality before a noticeable difference in user performance occurs?" While many studies suggest that haptic feedback can improve task performance (e.g.[1-3]), the few that have examined task performance as a function of haptic quality suggest that task performance is often not affected by differences in haptic quality even if an obvious degradation is perceptible [4-6].

Our specific context is that of training surgeons in laparoscopic surgical techniques using simulators with force-feedback. These are particularly expensive devices that few teaching institutions can afford, despite their putative benefits. The first step to achieving our larger goal is therefore to create a means by which we can conduct user studies of the impact of hardware performance on surgical task execution.

B. Approach

To study this we have developed an environment in which we can continuously vary hardware "quality" through simulation, by means of a custom software plug-in that intercepts the control loop between the virtual environment and the haptic interface. Through modifying the position signal sent from the hardware to the virtual environment and the force signal sent the other way, a high-fidelity interface can be made to display effects such as increased friction or inertia superimposed on its normal simulation.

A hardware implementation of this setup would entail rebuilding the hardware with different components. A software simulation permits independent and rapid adjustment of each parameter on a continuous scale, avoiding an uncertain and expensive redesign process. However, a hardware implementation would give the highest fidelity possible. For our purposes, high fidelity is relatively unimportant; rather, we require approximate effects which can be scaled to cover the range of variation we might expect with real hardware that spans the range from low to high end. That, combined with the flexibility and implementation time advantages, made us choose the software approach.

C. Hardware and Virtual Environment

Because our software modification technique works by intercepting the force and position signals between the hardware and VR software, it can be easily applied to different kinds of hardware. Since we are studying performance of laparoscopic training simulators, we used Immersion Corp.'s Surgical Workstation (www.immersion.com, Fig. 1; hardware specifications in Table 1). This device has two 5-DOF laparoscopic instruments, each of which move in and out of a

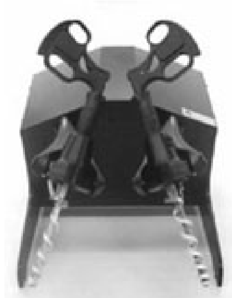

Figure 1. Picture of the Laparoscopic Surgical Workstation

2-DOF pivoting point and rotate around a longitudinal axis. A virtual tool tip opens, closes and rotates relative to the main shaft. All but this last degree of freedom are actuated.

The virtual environment (VE) is a simulation of minimally invasive surgery by Reachin Corp. (www.reaching.se) featuring soft tissue interaction. It supplies force feedback in the hardware's yaw, pitch and insertion axes; the models described here are therefore applied to these three degrees of freedom.

D. *Challenges*

The main challenges in implementing these models are caused by the complex dynamics of the haptic hardware. Models that are stable in a computer simulation may not be when displayed haptically, because the virtual models interact with the real hardware dynamics, including the hardware's friction, inertia and coupled kinematics. For an example of 3D interaction in our hardware, the orientation of the tool handle around it's axis alters the dynamics enough to introduce instability in other DOFs. An accurate description of hardware dynamics, when available, can be incorporated into a computer simulation; however, this is usually unobtainable for commercial hardware, and experimental parameter determination is difficult and often inaccurate.

Likewise, our need to simultaneously simulate a variety of hardware degradations vastly complicates our ability to achieve stability. These virtual models interact with one another, in addition to the real hardware. This has influenced both details of the model implementations and imposed limits on their parameterization.

Finally, Reachin's VE's sampling rate is not constant. On our hardware and with our software plugin, it varies between 500 to 2000Hz. Therefore our models must work with all sampling rates within this range.

E. *Paper Outline*

In the next section we discuss the models we use for simulating the degradations. In section III, we discuss how we integrated the different models. Results and model parameters will be presented in section IV, our conclusions in V. Section VI will contain a discussion on possible future improvements in the models and how this work fits into our larger objective of obtaining design parameters for haptic hardware for surgical training.

TABLE 1. LAPOROSCOPIC WORKSTATION SPECIFICATIONS

	Range	Cont. Output	Peak. Output	Sensor Res.
Insertion	170 mm	11.0 N	19.0 N	.008 mm
Pitch	100°	0.47 Nm	0.85 Nm	0.01°
Yaw	100°	0.47 Nm	0.85 Nm	0.01°
Handle Twist	180°	0.04 Nm	0.07 Nm	0.03°
Virt. Tip Twist	Cont.	N/A	N/A	0.7°
Handle Grip	20°	0.15 Nm	0.32 Nm	0.04°

II. HARDWARE MODELS

We have chosen to model several primary effects found in less expensive haptic interface hardware: cogging, inertia, backlash, friction and force saturation. Together with encoder resolution and refresh rate, these are the most prominent quality descriptors for haptic hardware. We did not degrade refresh rate in our experiments because it depends on computing power rather than the haptic hardware, and is rapidly improving, nor did we degrade encoder resolution because the encoders used were not expensive (in fact, our degradations would have benefited from better encoders).

Throughout the paper, the effects are represented as 1-DOF linear (translational) models. We do not describe the rotational variants, also implemented, which are obtainable through a straightforward transformation. Each section begins with a short discussion of previous work, and model element parameterizations are listed in Table 2.

A. *Inertia*

The most straightforward way to simulate inertia is to multiply actual acceleration by the virtual inertia. However, an acceleration estimate obtained by double-differentiating the position signal is too noisy to produce a stable simulation. A common solution is to simulate the virtual inertia's dynamics through integration of a 2nd order system, and virtually couple it to the probe position through a stiff damped spring (e.g.[7]); the damping requires only a velocity estimate. The stiffness of the spring and damper coefficient determine the tightness of the coupling, which ideally is critically damped.

Our implementation: Our system's temporal and position resolutions are such that the velocity signal tends to oscillate between a small number of values; we smoothed it with a 1st order Butterworth filter with a cut-off frequency of 70Hz . The varying sample rate requires frequent real-time adjustment of this filter's coefficients, imposing a ceiling on the coupling's damping. This in turn reduces the stability limit on the spring constant, cutting down the dynamic range of the virtual mass. To increase the stable range of parameters, we low-pass-filtered the resulting interaction force by averaging it over a 25-point window. The resulting model is illustrated in Fig. 2. Parameters are listed in table 2. We expect to be able to increase stiffness of the coupling by applying stability analysis (e.g. [7, 8]) and making the K and B variables dependent on sample rate. The last item can be beneficial since the virtual coupling can be made stiffer at higher update rates and this is exactly when we expect the largest accelerations in the user movements: high accelerations are more likely to occur in free

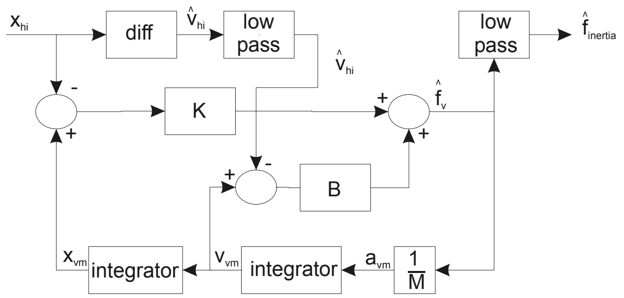

Figure 2. Inertia simulated by a virtual coupling with two low pass filters. Subscript 'vm' indicates that the variable is related to the virtual mass, 'hi' indicates it is related to the haptic interface.

space motion when the VE update rate is high, but low when there is a lot of interaction with the tissue.

B. Backlash

In a system with backlash, motion transfer between two masses occurs within a finite gap, causing a discontinuity and impact upon direction changes. Impact between the two masses can be approximated as occurring through a linear damped spring [9, 10].

Our implementation: We have adopted this model by attaching the virtual coupling to the gap-wall, engaging it when the user interface contacts either edge of the gap. In Fig. 3, the virtual mass (M_{vm}) represents the simulated extra mass of the motor and transmission. The position of the virtual mass is x_{vm}. We assume that there is negligible backlash in our hardware's cable drive and therefore consider the encoder signal an accurate estimate of the probe position x_{hi}, controlled by the user.

To enhance stability, we apply a small amount of viscous damping between the probe and the mass when the probe is within the gap. When the probe is in contact with the mass, the virtual coupling engages the gap wall (1).

$$\begin{aligned} &if\ \ x_{hi} > x_{vm} + 0.5 \cdot d_{gap}: \quad p_{vc} = x_{vm} + 0.5 \cdot d_{gap} \\ &if\ \ x_{hi} < x_{vm} - 0.5 \cdot d_{gap}: \quad p_{vc} = x_{vm} - 0.5 \cdot d_{gap} \end{aligned} \tag{1}$$

p_{vc} denotes the attachment point of the virtual coupling, and is undefined when the probe is not in contact with the mass. The force felt by the user can then be described as:

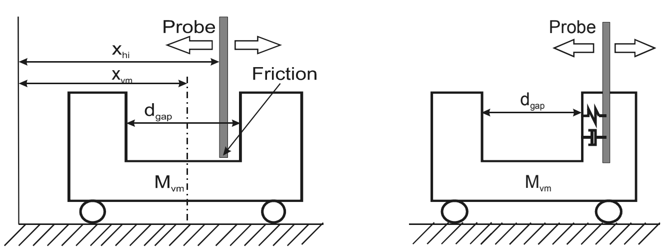

(a): In the gap (b): Outside the gap

Figure 3. Schematic representation of the backlash model in a single translation

$$\begin{aligned} &if\ \ (x_{vm} - 0.5 \cdot d_{gap} < x_{hi} < x_{vm} + 0.5 \cdot d_{gap}): \\ &\qquad F = -B_1 * \dot{x}_{hi} \\ &otherwise:\ \ f = f_{ext} + K(x_{vm} - x_{hi}) + B_2(\dot{x}_{vm} - \dot{x}_{hi}) \end{aligned} \tag{2}$$

in which $B_1 = 0.2 \cdot B_2$ (ratio optimized empirically).

C. Friction

Many friction models are described in the literature; Armstrong-Helouvry *et al.* provides a good overview [11]. Friction is a complex phenomenon and dependent on specifics of material and lubrication.

A "bristle model" is used to accurately simulate microscopic stick-slip contacts in real surfaces [12], but is too computationally expensive for real time processing. Chen *et al.* [13] developed a version for haptic rendering based on a single bristle that produces the dependency between normal and friction force. The authors report mixed results, and we could not implement it because our interaction normal force is unavailable.

Dahl's friction model uses one differential equation [14]. Hayward & Armstrong [15] showed that this model drifts under circumstances that often occur in haptic simulation, and produced a 4-state version dependent only on position. However, the state transition process assumes a constant sampling rate, making it unusable for our system.

Karnopp introduced a friction model that incorporates stick-slip without pre-sliding: i.e. when the friction force is below f_{static}, the relative velocity between surfaces is zero [16]. In two example implementations, the static friction force is made to depend on probe velocity and position[17]. Nahvi & Hollerbach introduced a haptic friction model in which the phase, the haptic interface is allowed only minimal movement due to a spring force. This spring ruptures when the spring force exceeds f_{static}. The transition from slip-stick transition is continuous by choosing the attachment position of the spring such that the static friction force is equal to the slip friction force [18].

Our implementation: Since the DOF of our haptic interface associated with tool insertion already has noticeable real friction, we tried to imitate its feel. We modified Karnopp's model to incorporate a proportional position-based controller between the probe and object that reaches maximum static friction (stuck state) at a pre-sliding displacement of 100 μm on the insertion (0.2 degrees in rotation). This model is similar to Nahvi's, with two differences: our friction force is independent of normal force (which value we don't know), and the slip-stick transition is effected by attaching the spring at the mass's last position before it entered the stuck state (Nahvi's method led to instability for our system). Parameters are shown in Table 2.

D. Cogging torque

DC brushed permanent magnet motors are the most common actuators used for haptic interfaces. Ideally, their output torque would be independent of the position of the rotor. In low quality motors, cogging may cause torque

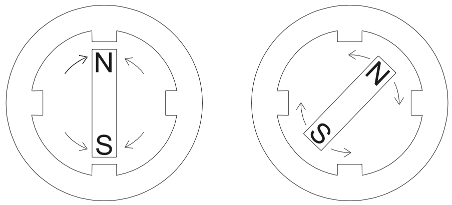

Figure 4. Cogging: A stable (left) and unstable (right) detent position of a permanent magnet (brushless) motor. A fluctuating torque can be felt due to the magnetic attraction between the permanent magnet rotor and stator teeth.

fluctuations as the motor rotates. Caused by the preferential alignment of rotor and stator, it can be felt as a series of opposing and aiding torques as the motor is turned when unpowered (Fig. 4).

Our implementation: We produced a torque-angle shape match to experimentally obtained cogging data [19-22] which resulted in a sinusoidal relationship between torque and motor angle.

E. Torque saturation

Electromotors are usually described by both continuous and peak maximum torque outputs; the peak torque can only be exerted for a limited time because of heat generated. Thus, while a motor has two design torque limits, the lower limit will be expressed in hardware as overheating and eventual damage to the motor rather than a haptically perceptible performance reduction.

Our Implementation: We applied a single cut-off limit for motor torque: i.e., when in effect, the motor force is clipped to the imposed saturation level.

III. MODEL INTEGRATION

We integrated our models in two stages. First, we combined the various degradations into a single DOF model so as to maximize simulation fidelity and stability. Next, we extended this 1-DOF model to the 3-DOF movement of the instrument tool-tip.

1-DOF Integration: To the greatest extent possible, we based our integration on the actual physical location of the respective degradations in a typical haptic hardware system We first simplified the reality of Fig. 5 by lumping the mass of the motor and transmission. Backlash is then defined as the play between the user's probe and this lumped mass, and

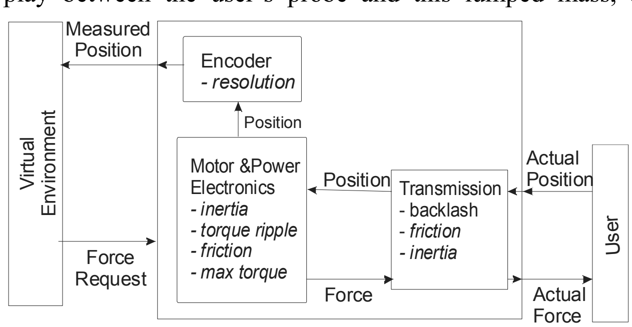

Figure 5. Flow diagram of forces and positions in a virtual reality system with haptics, and the fators that limit haptic fidelity.

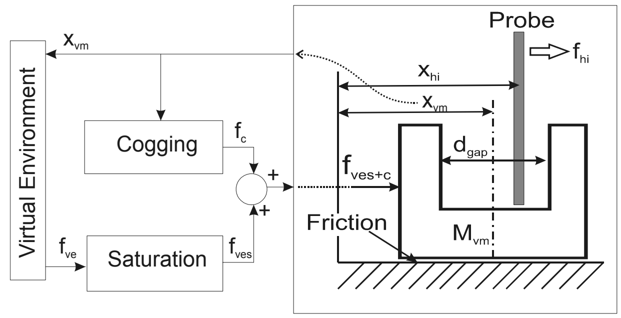

Figure 6. The basic physical representation of our model integration.

friction as the movement-opposing force between this mass and the 'ground'. Forces from the virtual environment are transferred through this backlash mechanism. As a result, the user feels forces from the virtual environment and from the degradation models only while the probe is in contact with the mass (Fig. 6).

To maximize perceptual fidelity of the different models, we further modified this physical model by removing the virtual coupling from all models except inertia and backlash: this coupling is an artifact necessary to simulate inertia but also low-pass filters the other degradations as well as the forces coming from the virtual environment. Therefore all force signals, except for the inertial force, are exerted directly on the probe. A switch signal produced by the backlash sub-model allows all forces to pass unmodified when the probe is in contact with the wall, and blocks all forces when the probe is in the backlash-gap. Finally, we made friction force depend on the position and velocity of the user probe rather than the simulated mass (Fig. 7). Richard took the latter approach with a relatively stiff 1-DOF haptic interface [23], but it led to a muddy-feeling friction in our system.

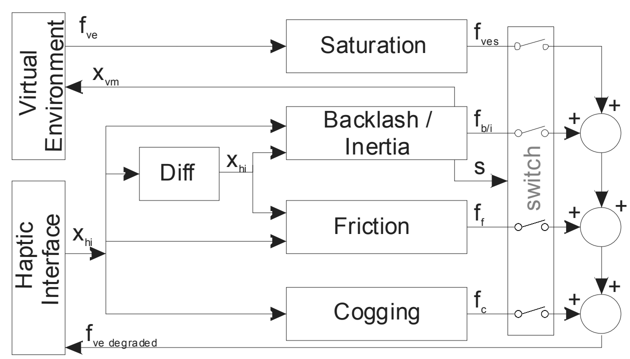

Figure 7. Flow diagram for final integrated model, illustrating backlash switch mechanism

3-DOF Extension: Not surprisingly, our backlash-inertia sub-model was the hardest to stabilize at higher force levels; it is both velocity-dependent and discontinuous, and sensitive to kinematic coupling. To overcome this, we had to significantly lower the stiffness of the virtual coupling until a time constant T=150ms was reached.

IV. RESULTS AND DISCUSSION

The models above were implemented on a dual-processor Xeon PC running at 2.0GHz with 2 GB of memory. Table 2 lists key model parameters used in the integrated version of the models; the values were chosen through a combination of realistic levels expected to be seen in inexpensive hardware components, and constraints imposed by simulation stability. Some of the more interesting features of the individual simulations are discussed below.

TABLE 2 MAIN PARAMETERS OF THE MODELS ON SIMULTANUOUSLY

Effect	Parameter	Translation	Rotation
Inertia	Mass	≤ 0.1 kg	≤ 1mkgmm
	K	≤ 150 N/m	≤ 1 Nm/rad
	B	≤ 0.05	≤ 0.04
Backlash	Gap Width	≥ 1 mm	≥ 2º
Cogging	Amplitude	0.6 N	0.04 Nm
Friction	Pre-sliding	≥5 millirad	≥100μm
	Stick Velocity	≥ 6º/s	≥5mm/s
	$\frac{\max(f_{stick})}{f_{slip}}$	1.05	1.05

Friction: shows a measured probe trajectory segment with only the friction degradation turned on. The friction model transitions from the slip to the stuck state just before t = 23.9s. The friction torque drops significantly, and then resumes (glitch just before 23.9s) because the user is still moving slowly in the same direction, elongating the virtual coupling spring. Once the probe changes direction (23.93s), the friction force changes sign as well and grows until the model re-enters the slip state at roughly t=24.24s. The glitch at t=23.9s is not realistic, but we found it is not noticeable.

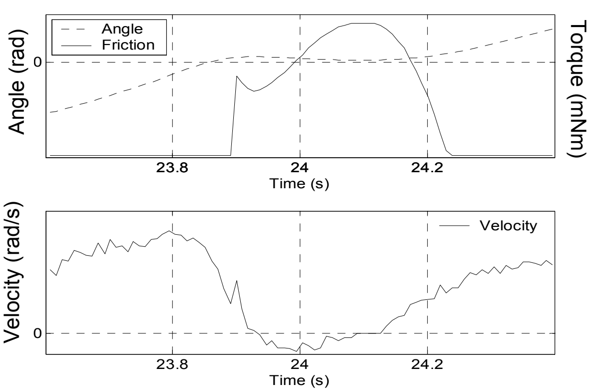

Figure 8. Friction model: actual probe position (left axis) and output friction force (right axis). VE forces are turned off.

Backlash: shows a measured trajectory segment of the user-controlled probe and the virtual mass with backlash turned on. A t=29.5s, the probe is pushing agains one wall of the gap, dragging the virtual mass closely behind it. When the probe stops, the virtual mass continues until the other wall of the gap hits the probe. When the probe starts moving in the other direction, this repeats itself. The backlash gap-width in this example is 1mm, and the simulated mass increment 0.2kg.

While our backlash model is structurally similar to that used in non-haptic simulation [9, 10], its parameters are unrealistically low: K=600 N/m. As a rough comparison, a 1 cm^2 contact area of a 1 cm^3 steel block has a K (EA/L) value of 200×10^7 N/m. This is reflected in the backlash model's feel: there is a clearly perceptible play in the gap, but the impact is not as crisp as one would expect. One remedy (untried) might be a force impulse on impact with the mass, as described for crisp simulation of virtual walls [17] However, the small gap creates a serious risk of wall-to-wall oscillation.

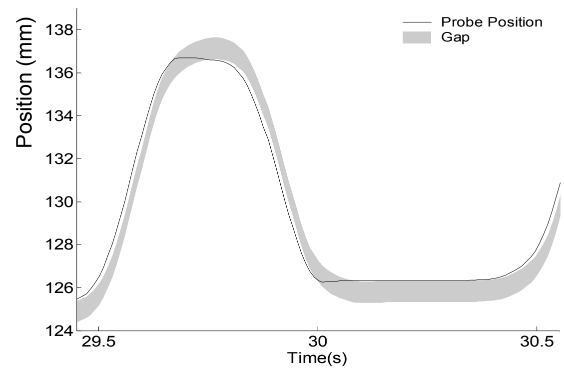

Figure 9: Backlash: probe position relative to the backlash-gap in the virtual mass.

Computational Load: The CPU effort required to simulate the various model aspects for six degrees of freedom on the computer described previously are listed in Fig. 10. Values were obtained by recording the time required to run each degradation independently and without the VE for 10,000 cycles, then computing mean update time.

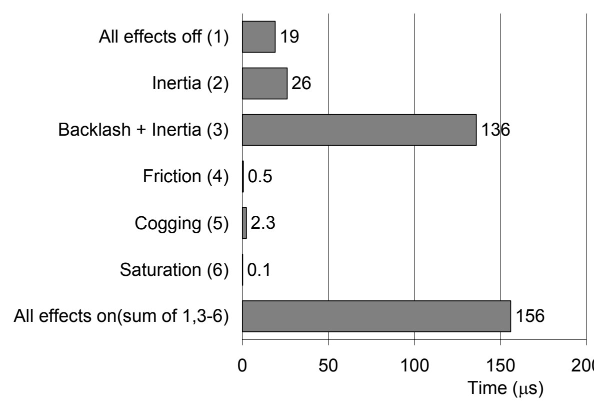

Figure 10. Chart with computing times of the various degradations.

V. CONCLUSIONS

We have modeled and implemented inertia, backlash, cogging, friction, and force saturation and shown that it is possible to implement degradation factors of haptic hardware under the following circumstances: multiple models working simultaneously, some with non-linearities, on hardware with complex and unknown dynamic properties and at varying sampling rates. While the models can be further improved upon to either extend the range of model parameters (e.g. mass), or the fidelity of the effects, the current model parameters fit the range we need to test for and we feel that the fidelity is high enough for our purposes.

VI. Future Work

Several aspects of the reported models are potentially improvable. Stiffening the virtual coupling will increase the fidelity of the inertia / backlash model; to do this stably requires application of more sophisticated filtering methods and a better understanding of the actual hardware dynamics. A better velocity estimate for the friction model will make the transition from slip to stick state more reliable and possible at a lower velocity.

This work is part of a larger study exploring how task execution is influenced by changes in haptic performance in the context of laparoscopic surgical training. A first experiment will determine how far we can degrade haptic quality before we notice a difference in task execution metrics, a result we expect to be task dependent. We are especially interested in learning which of these effects have the greatest influence.

References

[1] R. J. Adams, D. Klowden, and B. Hannaford, "Virtual Training for a Manual Assembly Task," *haptics-e*, vol. 2, 2001.

[2] L. Moody, C. Baber, and T. N. Arvanitis, "Objective surgical performance evaluation based on haptic feedback," in *MMVR Conference*, 2002.

[3] C. R. Wagner, N. Stylopoulos, and R. D. Howe, "The role of force feedback in surgery: analysis of blunt dissection," 10th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2002., 2002.

[4] C. Richard and M. R. Cutkosky, "Friction modeling and display in haptic applications involving user performance," Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on, 2002.

[5] K. E. MacLean, "Emulation of Haptic Feedback for Manual Interfaces",Ph.D. Thesis, Mechanical Engineering, MIT,Cambridge,1996.

[6] M. O'Malley and M. Goldfarb, "Comparison of Human Haptic Performance in Real and Simulated Environments.," Proceedings of the IEEE 10th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2002.

[7] J. E. Colgate, M. C. Stanley, and J. M. Brown, "Issues in the haptic display of tool use," Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems. 'Human Robot Interaction and Cooperative Robots', Pittsburgh, PA, USA, 1995.

[8] J. E. Colgate and G. Schenkel, "Passivity of a Class of Sampled-Data Systems: Appliation to Haptic Interfaces," American Control Conference, 1994, 1994.

[9] J. C. Gerdes and V. Kumar, "An impact model of mechanical backlash for control system analysis," Proceedings of the American Control Conference, 1995.

[10] T. Jukic and N. Peric, "Model based backlash compensation," American Control Conference, Arlington, VA, USA, 2001.

[11] B. Armstronghelouvry, P. Dupont, and C. C. Dewit, "A Survey of Models, Analysis Tools and Compensation Methods for the Control of Machines with Friction," *Automatica*, vol. 30, pp. 1083-1138, 1994.

[12] D. A. Haessig and B. Friedland, "On the Modeling and Simulation of Friction," *Journal of Dynamic Systems Measurement and Control-Transactions of the Asme*, vol. 113, pp. 354-362, 1991.

[13] J. Chen, C. DiMattia, R. M. I. Taylor, M. Falvo, P. Thiansathaporn, and R. Superfine, "Sticking to the point: A friction and adhesion model for simulated surfaces," Proceedings of the Sixth Annual Symposium on Haptic Interfaces for Virtual Environment and Teleoperator System, Dallas, TX, USA, 1997.

[14] P. R. Dahl, "Solid Friction damping of mechanical vibrations," *AIAA*, vol. 14 pp. 1675-1682, 1976.

[15] V. Hayward and B. Armstrong, "A new computational model of friction applied to haptic rendering," in *Experimental Robotics*, vol. VI, *Lecture notes in control and information sciences*, P. Corke and J. Trevelyan, Eds., 2000, pp. 403-412.

[16] D. Karnopp, "Computer simulation of stick-slip friction in mechanical dynamic systems," *Transactions of the ASME*, vol. 107 pp. 100-103, 1985.

[17] S. E. Salcudean and T. D. Vlaar, "On the emulation of stiff walls and static friction with a magnetically levitated input/output device," *Journal of Dynamic Systems Measurement and Control-Transactions of the Asme*, vol. 119, pp. 127-132, 1997.

[18] A. Nahvi, J. M. Hollerbach, R. Freier, and D. D. Nelson, "Display of friction in virtual environments based on human finger pad characteristics," American Society of Mechanical Engineers, Dynamic Systems and Control Division, Anaheim, CA, USA, 1998.

[19] R. P. Deodhar, D. A. Staton, T. M. Jahns, and T. J. E. Miller, "Prediction of cogging torque using the flux-MMF diagram technique," *Industry Applications, IEEE Transactions on*, vol. 32, pp. 569-576, 1996.

[20] M. Benarous and J. F. Eastham, "The effect of the distribution of the magnetisation in brushless DC machines on cogging torques," Electrical Machines and Drives, 1999. Ninth International Conference on (Conf. Publ. No. 468), 1999.

[21] T. Ishikawa and G. R. Slemon, "A method of reducing ripple torque in permanent magnet motors without skewing," *Magnetics, IEEE Transactions on*, vol. 29, pp. 2028-2031, 1993.

[22] C. Studer, A. Keyhani, T. Sebastian, and S. K. Murthy, "Study of cogging torque in permanent magnet machines," IEEE Industry Applications Conference, 1997. Thirty-Second IAS Annual Meeting, IAS '97.,, 1997.