ActivePaD Surface Haptic Device

Joe Mullenbach
Northwestern University

Dan Johnson
Design Integrity, Inc.

J. Edward Colgate*
Northwestern University

Michael A. Peshkin†
Northwestern University

Abstract

We present a new surface haptic interface that combines a variable friction device (the Large Area TPaD) with an impedance controlled planar mechanism. This device configuration is novel because it allows control of the frictional force in the static friction regime, control of the direction of force in the kinetic friction regime, as well as a degree of control over the transition between the two regimes. The range of operating modes combined with a large force capability make the device an appropriate platform for exploring surface haptic control algorithms. The design of the system is explained, two major categories of control algorithm are introduced, and the implementation of a virtual dimple is discussed. Experimental data are used to compare a virtual dimple to its physical analog, and to reveal areas for improvement.

Index Terms: H.5.2 [Information Systems]: User Interfaces—Haptic I/O

1 Introduction

The touchscreen interface has become the standard user interface for entire classes of devices. Designers of mobile phones, tablet computers, and instrument panels have all gravitated toward touchscreens over physical buttons and knobs for their natural mappings and programmable input displays. However, the consequence of this shift has been that users of these interfaces no longer receive the affordance cues, e.g., a toggle switch that can physically only move to one other configuration, or the haptic feedback, e.g., the sudden release of a keyboard button that indicates it has been selected, that make physical interfaces natural and easy to learn. The goal of this research is to develop force-based feedback and affordances for users of touchscreen interfaces.

At this point, most touch screen devices that incorporate haptic feedback operate by creating vibro-tactile sensations on the finger. Through the use of an eccentrically loaded motor [1,2] or a piezo-actuated beam [3] these devices have been able to create an added sense of realism to virtual events such as the impact of a bouncing ball or the shaking of a race car.

In order to display a more refined virtual environment, it is also important to be able to convey shape information. The desire for shape information has led to approaches which use an actuated pin array to change the shape of the surface itself [2,4,5]. In real world contact, the perception of shape on the fingertip is a combined effect of skin deflections and changes in force. While physically changing the shape of the surface does allow for both of these cues, Robles-de-la-Torre and Hayward proved that force cues alone are enough to overcome the perception of shape [6]. While this and other work was with kinesthetic displays [7], the idea has since been applied to several surface displays [8,9]. The ActivePaD joins these displays in the use of friction to create lateral force.

*e-mail:colgate@northwestern.edu
†e-mail:peshkin@u.northwestern.edu

IEEE Haptics Symposium 2012
4-7 March, Vancouver, BC, Canada
978-1-4673-0809-0/12/$31.00 ©2012 IEEE

2 Background–Friction

Although in general it is a complex phenomenon, we approximate frictional force with a Coulomb friction model. For a finger in contact with a surface, we represent the interaction with two vector forces. One force, N, acts normal to the surface plane, and the other, F, acts in the plane. There are two cases. In the kinetic case, there is a non-zero relative velocity, and the finger is sliding. The force on the finger is described in Eq. 1 and shown in Figure 1.

$$\overrightarrow{F}_{kinetic} = N * \mu_{kinetic} * \frac{-\overrightarrow{v}}{|v|} \tag{1}$$

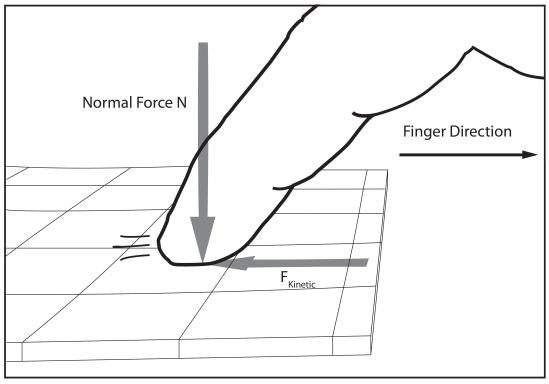

Figure 1: Kinetic Friction

In the static case, the finger is not moving relative to the surface, and is effectively stuck. In this case, shown in Figure 2, the force on the finger is equal to that acting on the surface, provided that it is below the static friction limit (Eq. 2, 3).

$$\overrightarrow{F}_{static} = \overrightarrow{F}_{surface} \tag{2}$$

$$F_{static} \leqq N * \mu_{static} \tag{3}$$

In these equations, μ is the friction coefficient between the finger and the surface, v is the relative velocity vector between the finger and the surface, and the quantity $\frac{\overrightarrow{v}}{|v|}$ is the unit vector direction of motion.

In general, these parameters can be controlled to change the force on the finger, and there may also be a transition between the static and kinetic cases. Normal force may be controlled through actuation of the surface in the vertical direction, or through the application of an attractive force. Yamamoto et al. describe an application where electrostatic forces are used to pull a slider down on the surface, modulating normal force, and thus frictional force [9]. Additionally, several devices have shown the ability to apply the same principle to a bare finger [10].

The coefficient of friction may be altered nominally by the choice of different materials, textures or lubricants, or it may be

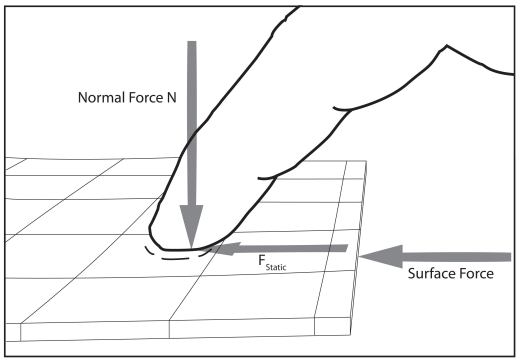

Figure 2: Static Friction

controlled dynamically through the use of a variable friction device. Watanabe and Fukui showed that the coefficient of friction of sandpaper could be reduced by vibrating the surface ultrasonically [11]. Winfield et al. applied this same principle to a sheet of glass to create the TPaD. Varying the amplitude of vibration effectively varies the coefficient of friction [12].

The relative velocity may be changed through the control of the velocity of the surface. Control of the relative velocity allows control over the direction of force. Furthermore, it allows for the transition between static and kinetic friction. Wang et al. created a device that used an overlaying material to pull the finger across the surface [8].

Finally, when the finger is experiencing static friction, the situation is equivalent to using a stylus or thimble and impedance control methods can be applied.

3 ActivePaD Design

With the exception of normal force, the ActivePaD is capable of changing each of these parameters- coefficient of friction, lateral force, and velocity. The touch surface of the ActivePad is a large area TPaD (Tactile Pattern Display), allowing for modulation of coefficient of friction. That surface is then mounted to a planar mechanism which allows motion in all the planar degrees of freedom. The TPaD is 4x4 inches square, while the interface area is 3x3 inches square. The mechanism is fully backdrivable, and is controlled with brushed DC motors fitted with high resolution encoders. The assembled device is shown in Figure 3.

Knowing the kinematic configuration also allows open loop control of surface force through pulse width modulation of motor voltage. Also, the velocity of the surface is derived from encoder readings. Finger position is measured optically. Knowledge of both surface and finger position allows software control of relative velocity. Finally, an LCD is mounted to ground just below the planar mechanism so that it is always stationary, with a 3x3 inch area revealed by looking through the glass of the TPaD. Graphics on the LCD are created through a PC program written in Processing. Control of the motors, as well as reading of the sensors, and calculation of the virtual environment is done on 3 Microchip PIC32 microcontrollers running on a 1 kHz update loop. Separate microcontrollers were used in order to minimize interference of different control interrupts while also providing a modular design that allows reuse of the elements on future projects.

3.1 Kinematics

The ActivePaD utilizes three brushed DC motors with a 3RRR (revolute-revolute-revolute) parallel linkage to directly control the

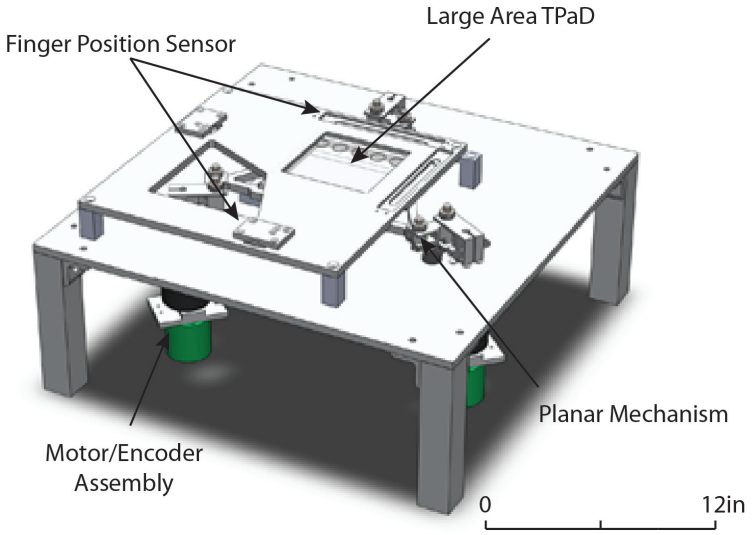

Figure 3: ActivePaD

three planar degrees of freedom of the touched surface. This mechanism was chosen in order to minimize moving mass while also avoiding translational joints. The linkage consists of three legs, each with two passive and one driven revolute joint, and is shown in Figure 4. Each linkage is 1 in. in length and the triangular plate ABC is 7 in. on each side.

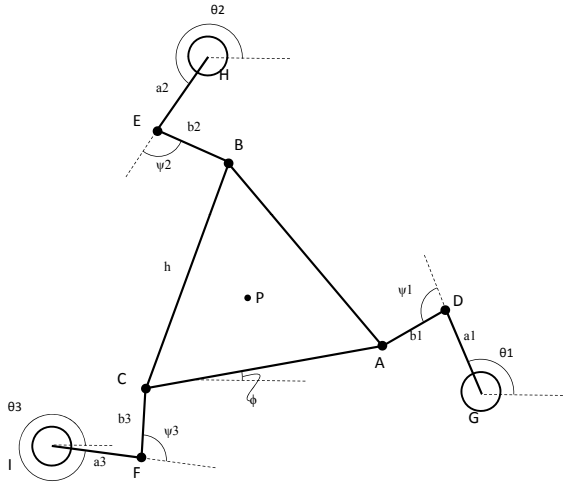

Figure 4: 3RRR Planar Mechanism

With the rigid, triangular plate ABC as the end effector, the inverse kinematics of the parallel system provide a transformation from standard, planar configuration variables $\{X_p, Y_p, \phi\}$ to the configuration variables of the actuators $\{\theta_1, \theta_2, \theta_3\}$. This allows for the determination of motor positions required for a desired end effector pose, as well as motor torques $\{\tau_1, \tau_2, \tau_3\}$ required for end effector forces $\{F_x, F_y, T\}$. In practice, ϕ is always controlled to be zero, while F_y and F_x are used in the virtual environment.

Singular configurations of the ActivePaD occur when any of the parallel chains becomes fully straightened, i.e., when $\psi = 0$. It is critical that the system remains well away from singular configurations across its entire workspace. The ActivePaD was designed to be capable of translating 0.5 inches in any direction from the center

position without encountering a singular configuration.

3.2 Actuator Sizing

There were two design requirements that determined actuator sizing. First, in order to control the direction of applied frictional forces, the ActivePaD needs to be able to create surface velocity greater than the velocity of the finger. A simple targeting task was presented to a group of users and their fingertip motion profiles were recorded. These profiles were then analyzed using Lagrangian methods to determine the motor torque required to meet them. The low mass of the mechanism resulted in relatively low required motor torque.

Second, the plate must be able to resist friction forces of the finger. The surface force required to resist the frictional forces of the fingertip was measured, and the motor torque necessary to maintain this force was estimated. This torque was much higher, and was used to determine minimum torque requirements for the actuators. Maxon 148877s brushed DC motors each with a maximum torque of 184 mNm were chosen. Figure 5 shows the response of a motor to a typical requested motion profile.

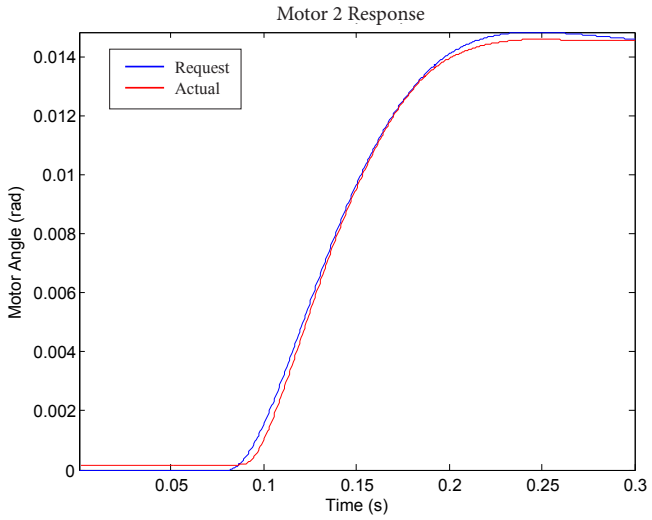

Figure 5: Single Motor Response

High resolution encoders were used to create the high-gain motion controller needed on the ActivePaD. The Gurley R137 encoders used on ActivePaD provide 720,000 quadrature counts per revolution, which is monitored by dedicated counters (LS7166) capable of receiving quadrature at up to 25MHz. This allows for the motors to move quickly without the system missing any counts.

3.3 TPaD Design

Marchuk et. al discussed the abilities of large area, higher resonant mode TPaDs in the reduction of surface friction coefficient [13]. It was seen that where nodal lines intersected on the surface, the TPaD had negligible ability to control the coefficient of friction between the surface and a fingertip. Although the friction reduction effect was also reduced at nodal lines themselves, the effect was still present and controllable.

To exploit this result, the TPaD used on the ActivePaD is designed to excite a resonant mode without intersections, described as a 0xn, or banded, mode. The banded mode is encouraged by placing the piezo actuators in a line along one edge of the device, which also serves to create a large area of transparent glass, through which graphical feedback is given. Figure 6 shows the TPaD with salt gathered on the nodal lines on its surface. This horizontal nodal line pattern with no intersecting nodal lines is what is referred to as the banded mode.

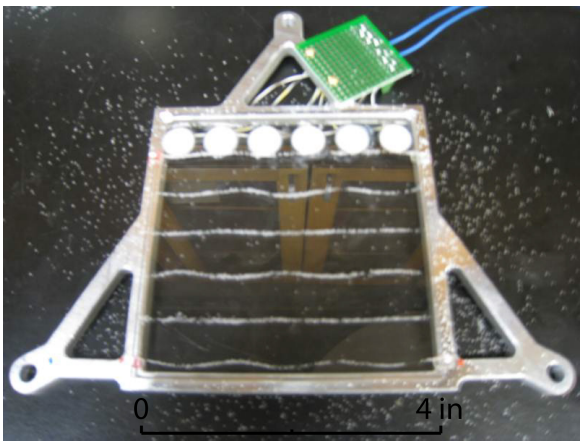

Figure 6: Large Area TPaD Showing Nodal Lines

3.4 Finger Position Sensing

The ActivePaD uses an optical finger position sensor consisting of two laser line generators and two linear photodiode arrays set up in a frame similar to an infrared touchscreen. The laser line is shined past the probing finger and onto the linear sensors. The center of the shadow cast from the finger is then calculated, weighted by light intensity. Cylindrical lenses are used to fan the laser, and Fresnel lenses are utilized to bend the fanned out laser into parallel light, which serves to increase the available sensing area without moving the lasers inconveniently far from the workspace, as shown in Figure 7.

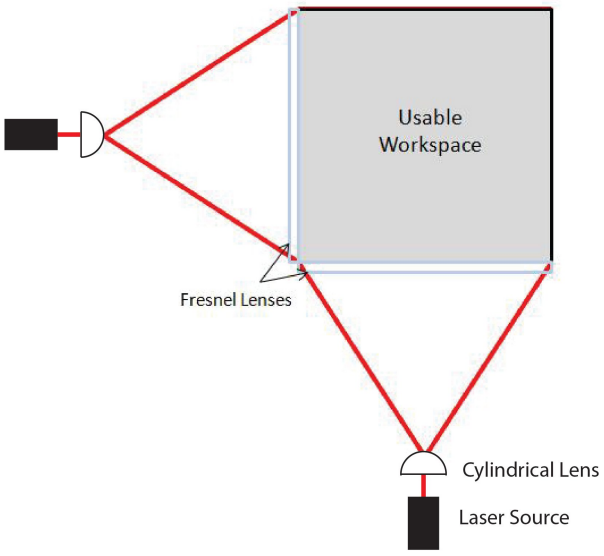

Figure 7: Finger Position Sensing Configuration

The update rate of the finger position sensors is 100Hz, and is dependent upon the saturation time of the photodiodes during one read cycle. Faster update rates could be reached with higher intensity light sources. This would decrease the time needed to saturate the photodiode arrays, thus increasing the read frequency.

In this system, alignment is critical, so a frame was machined from a single piece of aluminum to mount the lasers and photodiode assembly. The sensing area of the finger position sensor, and thus the usable workspace of the ActivePaD is 3x3 inches in size. The finger position sensing takes place 1.6mm above the surface of the TPaD, which allows for a good approximation of the finger pad location, however finger tilting may result in misreading the location of the center of the finger pad by up to 2mm.

4 Control Strategies

The control strategies for the ActivePaD can be categorized into two general groups: those that rely on kinetic friction to force the finger, and those that rely on static friction. We discuss the advantages and disadvantages of these control strategies, as well as strategies for transitioning between them. Although the device also allows control of a rotational degree of freedom, only the two translational degrees of freedom are used in the current control strategies.

4.1 Force Control via Kinetic Friction

As was previously discussed, to take advantage of kinetic friction, the magnitude of the relative velocity between the finger and the plate must remain greater than zero. In this mode, the magnitude of the frictional force can range from roughly 0.1 to 0.6 times the normal force, although these numbers are known to vary from person to person, and also according to skin conditions such as hydration or the presence or absence of lubricants such as skin oil [13,14]. Kinetic force control is further broken down into two strategies: passive and active.

4.1.1 Passive Kinetic

In passive kinetic force control mode, the finger is allowed to move freely, but the surface is fixed to ground. In this mode, we are able to modulate the magnitude of the frictional force, but the direction is determined by the velocity direction of the finger. Figure 8 shows a plot representing all possible force values. Assuming a positive x finger travel direction, the only possible forces are within a range on the negative x axis. Taken alone, this is a simple control strategy, but is still able to produce engaging effects. Using this type of control only, Winfield et al. were able to create textures such as roughness, fish scales or smooth bumps [12]. Levesque showed that passive kinetic control significantly increased user performance in targeting tasks, and had a positive effect on the sense of realism users experienced as well as increasing engagement and enjoyment [15]. While passive kinetic forcing was not the primary objective of the ActivePaD's construction, its use is appropriate in certain situations, and combining it with other control strategies may lead to novel haptic effects.

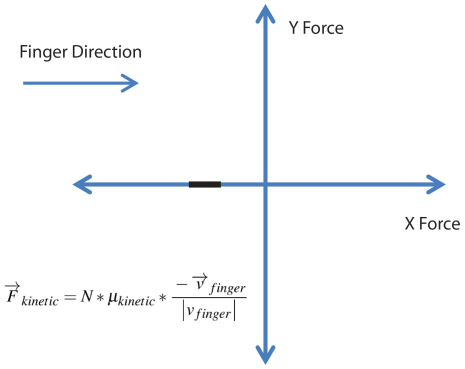

Figure 8: Line Representing Space of Forces that can be Produced with Passive Kinetic Force

4.1.2 Active Kinetic

Active kinetic force control is defined as modulation of both the magnitude and direction of the frictional force. The addition of control over direction of force allows the effect not just to resist the motion of the finger, but to guide it as well. A limitation of passive control of course is that if the finger is not moving, there is no force. In practice, this means that persistence of the virtual environment is not maintained. The ActivePaD was designed so that it can travel faster than the finger so that it is possible to completely reverse the direction of kinetic friction. This creates the potential for forcing in every direction for a range of magnitudes as is shown in Figure 9. A major limitation of using active control on the ActivePaD is in the range of motion. Because the workspace extends only 0.5 inches from the centered position in any direction, the surface can maintain a constant velocity for only a short time. This effectively limits not only the size of the effects, but their duration as well.

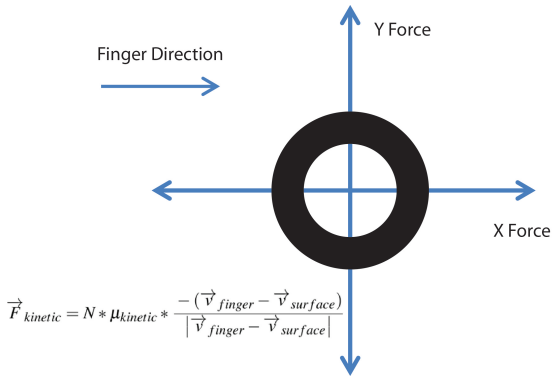

Figure 9: Ring Representing Space of Forces that can be Produced with Active Kinetic Force

4.2 Force Control via Static Friction

Static force control is predicated on the assumption that the finger and the plate are not moving relative to one another. In other words, the finger is stuck to the surface, which is the usual case in force feedback haptics. Because of this, the force applied to the plate is equal to the force on the finger. The range of forces is from zero to roughly 0.8 times normal force in every direction as is shown in Figure 10 [16]. While the force on the finger is equal to the force applied to the surface, the force applied to the surface is not directly the force of the motors, but is complicated slightly by inertia and joint friction.

$$F_{Surface} = F_{Motors} - m * a - F_{JointFriction} \quad (4)$$

These additional forces could be modeled and accounted for in the force command, although this has not yet been done. As such, the force on the finger is assumed to be the force applied at the motors. This sort of absolute control over force is in some ways an advantage over kinetic forcing, in that as long as it is below the static friction limit, it does not vary from finger to finger, and it does not vary according to surface conditions such as moisture or lubrication for the same finger.

Impedance control has been implemented, and it has been explored by using both the finger position and plate position as the feedback signal. Since we are assuming that the finger is stuck to the surface, the two approaches should produce the same result, but of course the sensors (laser shadow for the finger, encoders for the plate) have different resolutions and the compliance of the fingertip ensures some disagreement. Thus, the results are, in practice, different.

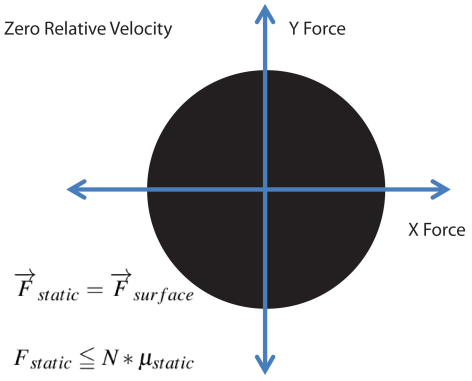

Figure 10: Disc Representing Space of Forces that can be Produced with Static Force

4.2.1 Impedance Control- Finger Position

In this mode, shown in Figure 11, the force on the surface is controlled based on finger position. Using finger position provides a natural mapping for non-moving effects such as a virtual hill or valley, because the user expects the force level to be purely a function of finger location. However, if the normal force is not high enough, and the plate force exceeds the maximum friction force, the plate will slide, and will continue to slide until it runs out of range. This unexpected transition to kinetic friction is felt as a change in force, and the persistence of the effect is lost.

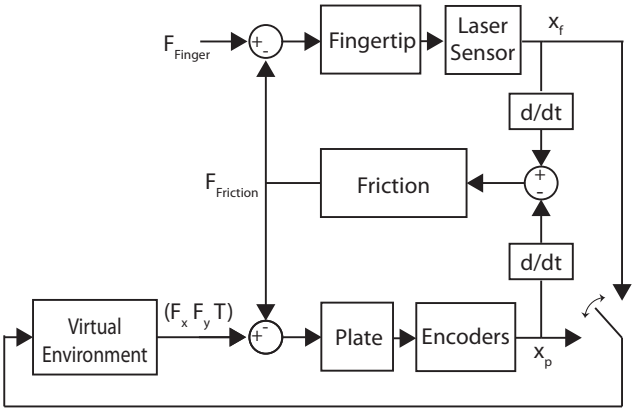

Figure 11: Finger Position Impedance Control

4.2.2 Impedance Control- Plate Position

In this mode, the force on the surface is controlled based on plate position. Thus the finger has to move a physical object in order to move the virtual object. This mode is well suited for virtual effects that depend on relative position. An advantage of this mode is that less care needs to be taken in making sure that the finger is actually stuck to the surface. Consider the same case of insufficient normal force to keep the finger in static friction. In this control mode, nothing drastic happens. The finger slides off, and the plate returns to its natural resting position, and does not run out of range. There is assurance that the user is actually feeling the appropriate force, because the plate will only move if the finger is pushing it.

4.3 Combining Strategies

Each control strategy has strengths and weaknesses, and the control algorithm should be designed to use multiple strategies to their best advantage. For example, if the user is using a finger scanning motion to explore the surface, passive kinetic control may be all that is necessary to represent an effect. If the user is using a more deliberate exploratory motion, static force control may be better suited to display a greater range of forces and force directions. One can then imagine a single interaction where a user starts by scanning across the surface to find the gross location of the feature, and then homes in on it to explore it precisely. An effective control algorithm may be to display the feature with passive kinetic control while the finger velocity is above a certain threshold, but then transition to static force control for low speed display.

4.3.1 Transition Strategies

In order to create such an algorithm, strategies must be developed to create the transition from kinetic to static friction and vice versa.

Kinetic to Static In order to make the transition from sliding to being stuck, the relative velocity between the finger and surface must be zero, and the lateral finger force must be below the maximum static frictional force. An initial strategy has been to control the plate velocity to match the velocity of the finger. A second strategy has been to slow the finger down to match the speed of the plate. This is done by increasing the kinetic frictional force on the finger by reducing the TPaD amplitude. For either strategy to be successful, however, there should be no sudden changes in force felt by the finger through the transition.

Static to Kinetic For the finger to transition from a stuck state to a sliding state, the finger must overcome the maximum static frictional force. This can be done one of two ways- by increasing the lateral finger force, or by decreasing the maximum static frictional force. With the TPaD effect, the coefficient of friction and thus the static frictional force can be reduced. Again, in general the transition is a noticeable event, and the correct algorithm must be developed as to minimize the spike in force.

5 Example Effect

While the device is capable of any number of active effects such as pulling the finger along some desired trajectory, the initial focus is on simulations of physical surface features such as bumps or depressions. Doing so allows geometry to serve as a guide for the effect design, and also creates a standard against which we can judge the quality of the virtual effect. The first case to be tested was that of a circular depression or dimple in the surface. The dimple is a potential well such that as the surface is explored, the finger tends to fall into it and get drawn to the center. The physical feature was formed from a piece of sheet metal, and is shown in Figure 12.

The virtual feature is an algorithm with two regions based on finger position. The strategy is to allow the finger to slide freely when it is outside of a certain radius, Rad, of the target, but once it has crossed the radius, to transition to static friction, and pull it into the center. In the surrounding region we use passive kinetic control, simply turning the TPaD on high so that the friction force is as small as possible. Once the finger has crossed into the effect region, the TPaD is transitioned to “off”. Within the effect region, we use static force control based on the force profile of a physical dimple.

```
If (xFinger - xCenter) < Rad
   tpadON = FALSE;
   Impedance control based on finger position
Else
   tpadON = TRUE;
   Plate is stationary
```

The potential profile of the dimple was modeled as a Gaussian [Eq. 5] to allow for smooth transitions.

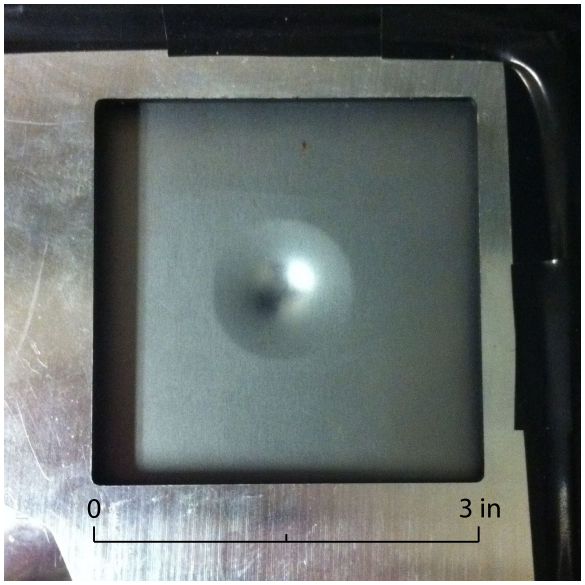

Figure 12: Physical Surface Dimple

$$g(x) = -A * e^{\frac{-(x)^2}{w^2}} \tag{5}$$

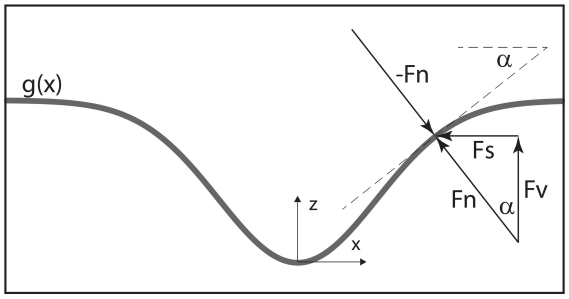

Figure 13: Gaussian Profile

Assuming a completely static arrangement as is shown in Figure 13, the lateral force F_s is related to the vertical force F_v by the relationship $tan(\alpha)$.

$$\overrightarrow{F}_s = \overrightarrow{F}_v * tan(\alpha) \tag{6}$$

Since $tan(\alpha)$ also describes the slope of the curve, the relationship becomes:

$$\overrightarrow{F}_s = \overrightarrow{F}_v * \frac{dg(x)}{dx} \tag{7}$$

Finally, a constant downward vertical force is assumed, and the forcing function becomes Eq. 8. Note that the size of the effect can be scaled by changing the coefficient C. The final force profile is shown in Figure 14.

$$\overrightarrow{F}_s = C * \frac{dg(r)}{dr} \tag{8}$$

For illustration purposes, velocity data were taken of the author completing a targeting task. A target, the finger position, and starting point were displayed on a computer screen away from the ActivePaD, and a metronome was playing for time referencing. Starting from the left of the target, the task was to trace the finger on

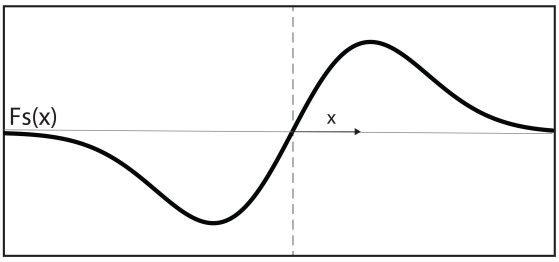

Figure 14: Virtual Effect Force Profile

the screen from the starting position to the center of the target in a constant amount of time (0.6 seconds). Four cases were compared, and 15 trials were taken of each: the physical dimple, flat sheet metal, the virtual dimple, and the glass surface with the TPaD on. Furthermore, the size and intensity of the virtual dimple were altered to demonstrate its effect. Figure 15 shows the shape of the mean and standard deviation of the speed profiles by plotting finger speed against finger position in the x direction.

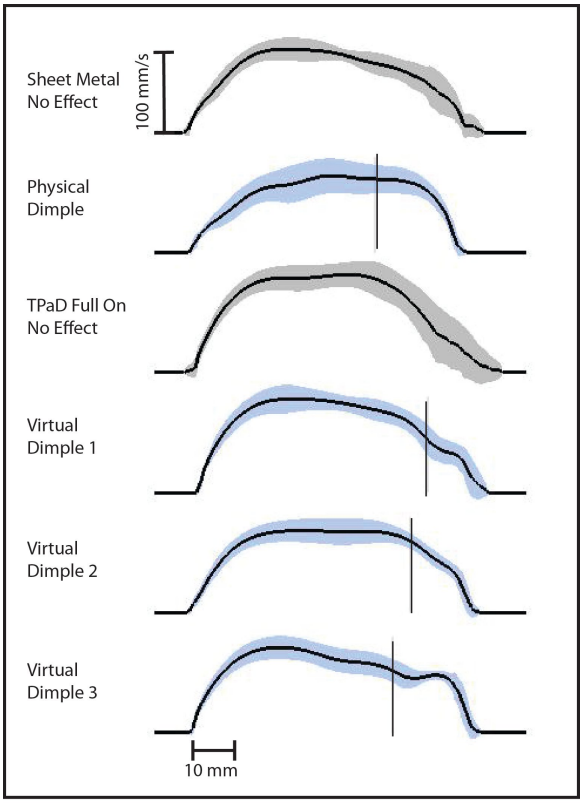

Figure 15: Mean and Standard Deviation Speed Profile of Finger Approaching Dimple Boundaries and Continuing Through Effect.

As the finger traveled from left to right, the vertical lines mark the point where the effect begins. On the flat sheet metal, the finger accelerates in the beginning, and then tends to gradually slow down as it nears the target. The physical dimple is similar, except that as it comes into the effective region, the dimple keeps the finger from decelerating for a distance, and then brings it to a sudden halt.

From a qualitative perspective, this effect feels very smooth, as if the finger just naturally falls into the center. There are no rough

edges or discontinuous areas, and that is reflected in the velocity profile. Comparing the virtual effect to this, there is a feeling of actively being grabbed and pulled to the center.

The virtual dimples increase in width and force magnitude as they get higher in number. This can be noted on the plots by the starting location moving to the left, and the speed taking a more dramatic departure from its natural descent as the graphs go down.

Each dimple plot (physical as well as virtual) shows a "knee" extending to the right where velocity is kept higher than in the no-effect case. The shape and location of this knee is instructive. For the physical dimple, the knee starts at the very top of the curve, so that the finger does not start decelerating until it is in the dimple area. For virtual dimple 1, the knee starts after the finger has already slowed significantly. For virtual dimple 3, the knee actually accelerates the finger before bringing it to a rest. In general, these discontinuous events were very noticeable, and in the case of virtual dimple 3, somewhat unpleasant. Of all the virtual effects, dimple 2 qualitatively felt the best. We notice from Figure 15 that it also had the least abrupt change in velocity.

While this virtual dimple was solely a function of finger position, the apparent correlation between the "natural feel" of the virtual dimples with the shape of the speed profiles suggest that it should be a function of finger speed as well. Our hypothesis is that the physical dimple felt good not only because it was a physical object, but also because it did not cause any abrupt changes in finger speed. A future algorithm design will try to exploit this. For example, instead of a preset force profile, perhaps it would be more compelling to force the finger according to some reference velocity profile. Alternatively, we can think of the finger as an active element. It is possible that the finger is reflexively responding to the lateral forces that it is experiencing. This is an area for further investigation.

6 Commentary and Future Work

Since this is a new device, and we are only beginning to explore the creation of virtual effects with it, this is a short commentary on the challenges presented, and opportunities for improvement.

What makes an effect compelling? One thing that makes for a compelling effect is tight control of force, since real-world effects are characterized by impedances, which are dynamic mappings of motions into forces. A major limitation of the current setup is that we are unable to measure forces directly on the fingertip. In the future, we intend to explore techniques of making fingertip force measurements.

A constant normal force was assumed for the rendering of the virtual effects. While this assumption was necessary because there are no force sensors on the ActivePaD, we know that in practice, normal force varies significantly. Furthermore, normal force varies in a predictable way for a physical dimple, and we are unable to recreate that force. This may be another reason for the discrepancy between the physical and virtual dimples.

As was previously discussed, the stretch of the skin is a powerful cue for interpreting the shape of an object, but to this point, attention has not been paid to it in algorithm design. Consider the finger sliding into the physical dimple. As it is sliding, the skin is stretched so that the pad is trailing the motion of the finger. This continues as the finger slides to the bottom of the dimple. This suggests that the net force vector on the finger pad never actually reverses its direction. Now consider the virtual dimple. When the finger encounters the effect zone, the plate moves so as to pull the finger toward the center of the dimple. In this case, the force vector on the finger pad is reversing direction. This effect may be the reason that transition events from kinetic to static control modes are currently quite noticeable. In any case, a strictly static model does not seem to capture all of the important aspects.

What are the important aspects of control? Most of the things that make for a good controller design also make the haptic effects more believable. Overshoot, delayed response, and excess vibration especially have been found to be problematic for a natural feel . Conversely, adding too much derivative gain is also apparent to the user, and gives a sticky or viscous feeling.

Tight spatial and temporal integration between the haptic and visual displays is also important to maintain the object illusion. Further research is devoted to exploring under which conditions the brain integrates these separate cues into a mental model of a single object.

Do we strive for realism or performance? While the goal up to this point has been to recreate familiar physical objects in a virtual environment, this is not a necessary condition to increase task performance. For example, using passive kinetic control of a TPaD, Levesque measured user performance in a targeting task with and without haptic feedback [15]. He found that user performance increased with the haptic feedback despite the fact that the target effect was not designed to mimic a physical object or to feel natural in any way. Although it was not measured experimentally in this paper, target accuracy seems to increase with increased effect magnitude. This intuitively makes sense, but increased force does not necessarily correspond to a more natural feeling effect. Metrics must be further developed for measuring the overall quality of any given surface haptic effect. Force-motion characteristics, task performance, and qualitative impressions should all be considered in future work.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant No. IIS-0964075 as well as a grant from the Center for the Commercialization of Innovative Transportation Technology (CCITT), a USDOT Tier II University Transportation Center operated within The Transportation Center at Northwestern University

References

[1] Linjama, Jukka, and Topi Kaaresoja, "Novel, minimalist haptic gesture interaction for mobile devices," *Proceedings of the Third Nordic Conference on Human-Computer Interaction*, pp. 457–458, 2004.

[2] S. Kim, K. Kim, B. Soh, G. Yang, and S. Kim, "Vibrotactile rendering for simulating virtual environment in a mobile game," *Consumer Electronics, IEEE Transactions on*, vol. 52, no. 4, pp. 1340–1347, 2006.

[3] I. Poupyrev, S. Maruyama, and J. Rekimoto, "Ambient touch: designing tactile interfaces for handheld devices," in *Proceedings of the 15th annual ACM symposium on User interface software and technology*. ACM, 2002, pp. 51–60.

[4] H. Iwata, H. Yano, F. Nakaizumi, and R. Kawamura, "Project feelex: adding haptic surface to graphics," in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 2001, pp. 469–476.

[5] Paris S. Wellman, William J. Peine, Gregg E. Favalora, and Robert D. Howe, "Mechanical design and control of a high-bandwidth shape memory alloy tactile display," in *International Symposium on Experimental Robotics*, 1997

[6] G. Robles-De-La-Torre and V. Hayward, "Force can overcome object geometry in the perception of shape through active touch," *Nature*, vol. 412, no. 6845, pp. 445–448, 2001.

[7] M. Minsky, "Computational haptics: the sandpaper system for synthesizing texture for a force-feedback display," *Massachusetts Institute of Technology, Cambridge, MA*, 1996.

[8] D. Wang, K. Tuer, Mauro Rossi, and Joseph Shu, "Haptic overlay device for flat panel touch displays," in *Proceedings of the 12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, 2004.

[9] A. Yamamoto, S. Nagasawa, H. Yamamoto, and T. Higuchi, "Electrostatic tactile display with thin film slider and its application to tac-

tile telepresentation systems," *IEEE Transactions on Visualization and Computer Graphics*, pp. 168–177, 2006.

[10] O. Bau, I. Poupyrev, A. Israr, and C. Harrison, "Teslatouch: electrovibration for touch surfaces," in *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. ACM, 2010, pp. 283–292.

[11] T. Watanabe and S. Fukui, "A method for controlling tactile sensation of surface roughness using ultrasonic vibration," in *Robotics and Automation, 1995. Proceedings., 1995 IEEE International Conference on*, vol. 1. IEEE, 1995, pp. 1134–1139.

[12] L. Winfield, J. Glassmire, J. Colgate, and M. Peshkin, "T-pad: Tactile display through variable friction reduction," in *Second Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. IEEE Computer Society, 2007, pp. 421–426.

[13] N. Marchuk, J. Colgate, and M. Peshkin, "Friction measurements on a large area tpad," in *Haptics Symposium, 2010 IEEE*. IEEE, 2010, pp. 317–320.

[14] D. Highley, M. Coomey, M. DenBeste, and L. Wolfram, "Frictional properties of skin," in *The Journal of Investigative Dermatology*. No. 69, 1977, pp. 303–305.

[15] V. Levesque, L. Oram, K. MacLean, A. Cockburn, N. Marchuk, D.Johnson, J. Colgate and M. Peshkin, "Enhancing physicality in touch interaction with programmable friction," *Proc. ACM Conference on Human Factors in Computing Systems* (CHI '11), Vancouver, Canada, May 2011, pp. 2481-2490.

[16] X. Dai, J. Colgate, and M. Peshkin, "LateralPaD: a surface haptic device that produces lateral forces on a bare finger," in *Submitted to Haptics Symposium 2012*.