

Haptic Techniques for Media Control

Scott S. Snibbe
541 Mississippi Street
San Francisco, CA 94107 USA
scott@snibbe.com

Karon E. MacLean
Dept. of Computer Science
Univ. of British Columbia
Vancouver, B.C. V6T 1Z4 Canada
maclean@cs.ubc.ca

Rob Shaw
shaw@hapttek.hapttek.com

Jayne Roderick
roderick@interval.com

William L. Verplank
verplank@ccrma.stanford.edu

Mark Scheeff
scheeff@scheeff.com

ABSTRACT

We introduce a set of techniques for haptically manipulating digital media such as video, audio, voicemail and computer graphics, utilizing virtual mediating dynamic models based on intuitive physical metaphors. For example, a video sequence can be modeled by linking its motion to a heavy spinning virtual wheel: the user browses by grasping a physical force-feedback knob and engaging the virtual wheel through a simulated clutch to spin or brake it, while feeling the passage of individual frames. These systems were implemented on a collection of single axis actuated displays (knobs and sliders), equipped with orthogonal force sensing to enhance their expressive potential. We demonstrate how continuous interaction through a haptically actuated device rather than discrete button and key presses can produce simple yet powerful tools that leverage physical intuition.

KEYWORDS: Haptic force feedback, user interface design, interaction techniques, tangible interfaces, media browsing, multimedia control, video editing.

1 INTRODUCTION

We use our haptic (touch) sense of the forces acting on our bodies to perform many everyday tasks like walking, driving, operating machinery, cooking and writing. In doing so, we interact with *physical dynamic systems*, whose components' movement is determined by physical laws. This paper describes several approaches to using virtual, haptically displayed dynamic systems to *mediate* a user's control of various sorts of media. These dynamic systems are constructed as physical task metaphors, rather than as literal representations of the media: e.g. instead of literally rendering the content of individual frames, we use the haptically perceived spinning of a large heavy wheel to indicate a video stream's progress.

Humans are accustomed to manipulating static visual media with physical dynamic systems: pencil and paper, brush and canvas, fingers and clay, chisel and stone. Where these media have migrated to the computer, we engage with a generic mouse or keyboard and have lost distinctive physical sensations. Some research seeks to duplicate the traditional feels with positive results [2, 15]. Others have built haptic dynamic systems for drawing and sculpting that have no direct physical analog [11, 19].

Similarly, traditional physical tools offered now-lost benefits to the manipulation of dynamic sounds and images. Film sound designer Walter Murch observed that the physical properties of editing mechanisms and the media itself enabled a level of control lost in nonlinear digital editing systems [13]: the duration of motion picture film and audiotape is related to physical length or bulk, and physical marks can be scratched and re-found. The spinning mass of a gang synchronization wheel (used to create film audio tracks) allows smooth adjustment of review speed and rapid, accurate location of a critical frame. DJs cling to vinyl for direct access to audio tracks, control over play speed and zero-latency audio response. Naimark performed early experiments to re-introduce the feel of traditional film equipment to digital video [14].

Past haptics research has pursued the manipulative use of force feedback primarily through simulation of medical and musical instruments [4, 10]; and its more exploratory function in the refinement of techniques for haptic rendering (creating a sensation of touching an object that doesn't physically exist [16]). The research described here deploys both knowledge bases in rendering and manipulating a mediating dynamic system, rather than the object of manipulation itself.

Here, we describe an exploration in restoring physicality to nonlinear media, introducing a series of metaphors and techniques for manipulating digital video, digital audio and computer graphics using haptic force feedback. We introduce the general principles of

these metaphors and document the hardware prototypes developed to explore them, then describe the haptic metaphors and behaviors themselves. Finally, we offer informal observations on their use, a note on computational architecture and reference to the continuation of some elements of the work.

2 DESIGN PRINCIPLES

The following design principles are a combination of core interface intuition based on past activities (e.g. as videographer or sound designer) that we brought to our research, and insights that emerged over its course.

Discrete vs. Continuous Control

A manual controller can be a “button” that triggers a discrete action or a “handle” that continuously modulates a parameter. Due to price pressures, arrays of buttons dominate contemporary media control tasks, even for continuous properties such as volume or video scrub rate. A continuous controller like a knob can provide perceptually infinitesimal control for manipulating parameters or media. Its greater size and cost may require it to accommodate more functions, but well-designed haptic feedback can keep it intuitive.

Information Delivery through Touch

Visual interfaces compete for attention with visual media; controls on border toolbars and floating palettes shrink or obscure content. Similarly, auditory interfaces compete with audio content. Maximizing information delivery through a media-independent channel (touch) and thus reducing perceptual noise can enhance both ease-of-use and accuracy of control.

Dynamic Systems for Control

One can employ direct manipulation, a fundamental tenant of contemporary interface design [6, 18], by modeling a task or process with a dynamic physical metaphor and then haptically rendering this metaphor as the process controller. We thus place an abstract tool *between* the application task and the user’s hand: it is the leverage that a paintbrush offers over one’s fingers or a film-editing table over a knife and tape.

Modeless Operation

Modes necessitate a sensible or mental record of a device’s current state; the former are often not provided by electronic and computer interfaces, and the latter increases workload and errors. Modeless interaction is achieved, however, by a consistent and trustworthy physical behavior, like that of a car steering wheel. Haptic feedback facilitates sophisticated tools for which physical intuition can be developed and information subconsciously absorbed.

Application and Interface Communication

Conflicting needs often dictate computational separation of haptic feedback control from application

content. To achieve the low-latency force feedback that provides a sense of task presence and control, we have found it advantageous to locate the dynamic model with the haptic controller. This in turn requires high bandwidth inter-CPU communication, in conflict with the typical design goal of independence between application and interface (Section 6).

3 DEVICES

Here we present the manual controllers used in these interaction experiments, largely single degree-of-freedom (df) devices designed for low cost and ease of integration into embedded contexts. *Engineering prototypes* are actuated, sensed and computer-interfaced with little attention to appearance, used to evaluate a particular technology or as general-purpose development platforms. *Form studies* were used to explore ergonomic design and act out task scenarios. They often have moving parts, but are not actuated or computer-controlled. *Functional prototypes* are working devices with a form factor relevant to their specific application, and required the greatest effort. We have developed most of our haptic behaviors using engineering prototypes, with form studies to develop specific physical application contexts.

Orthogonal Force Sensing

As described elsewhere [12, 16], “shading” force magnitude along one actuated axis can create an illusion of force supplied along a second, orthogonal axis (Figure 1). Conversely, user-applied deflection or pressure can be *measured* along an axis orthogonal to the actuation. Many of our behaviors are designed to supply an actuated response to such a deflection, enhancing the illusion of a second actuated df and adding a valuable, integrated control dimension.

Orthogonal force can be sensed in many ways, depending on the precision required. We have employed at one end a high-performance force-torque sensor (Figure 2a), and at the other a variety of cheap force sensing technologies including force-sensing resistors and optical and Hall-effect measurement of small displacements (Figure 2b).




Figure 1: Orthogonal force illusion, created by projecting the slope of a two-dimensional geometric profile (bottom) onto a single-axis force profile. Arrows indicate direction and magnitude of displayed force. A user perceives the two-dimensional surface as a dip (shown) or a hill (arrows reversed).




Figure 2: Engineering Prototypes. Clockwise from upper left: *Big Wheel* for high quality sensing (F/T sensor and optical encoder are located in series behind motor), large diameter knobs and a variety of orientations (motor/sensors assembly terminates in a post which can be inserted at various angles); *Cheap Force Sensing Wheel*, which measures displacement of a cantilevered mount with Hall-effect and optical sensors; *Brake* with an encoder and particle brake linked via a belt; and the *Slider* for experiments in absolute positioning, with a pressure sensor on the handle to select engagement with the haptic model.

Big Wheel: Multi-axis Force Sensing

For behavior development we used a powerful motor to directly drive a large-diameter wheel while sensing both knob rotational position and the forces exerted by the user on the handle. The knob axis may be aligned at 0°, 45° and 90° relative to the device’s mount, facilitating simulation of different kinds of hand/wheel interaction – e.g. edge versus top surface (Figure 2a). The device has a 90W, 490 mNm Maxon™ DC brush motor and an HP 4000-line optical encoder. The F/T sensor (ATI™ Mini-40 20-1) measures 6 axes of decoupled force (x, y, z at 20-60 N saturation and 1/800 resolution) and torque (roll, pitch, yaw at 1 Nm and 1/32000 resolution). Handles of various shapes and diameters were designed by specifying hand interaction and tangential forces – e.g. 4 N at the rim for a 300-mm diameter platen to simulate a phonograph record.

Different force measurement schemes are possible; e.g. of a specific axis or the maximum from all axes. Observing activity on all axes also helps determine how

a 1-df sensor should be mounted: if the wheel is pressed radially at its rim, the force may be best sensed in the x or y-axis (normal to the knob shaft). Behaviors engaged by touching the top knob surface utilize z-axis force (aligned with the knob shaft).

Brake: Passive Haptic Display

We performed a series of small experiments with a brake (Figure 2c). Since brakes can only remove energy from a system, stability is guaranteed, making them safer and more predictable to inexperienced users. Lower cost and power needs for a given torque make them attractive for consumer products. However, crafting a precise haptic experience can be more challenging: e.g., a position error cannot be corrected with closed-loop control. Instead, one can synchronize application to position to create an illusion of hitting precise targets (*sticky channels*, Section 4.2).

Slider: Absolute Positioning

We constructed a device with limited range of motion (Figure 2d), affording absolute rather than relative positioning. This constraint is an opportunity to exploit muscle memory, as we do when operating a radio dial

without looking: specific destinations are stationary relative to the device’s base. Haptic landmarks such as bumps or textures further “anchor” locations.

In a high-end audio mixing board slider, a small motor drives the slider open-loop via a toothed belt. To eliminate cogging and obtain position readout, we replaced the original actuator with a Micromo™ 1524 motor/encoder, geared 6.3:1 and with post-quadrature position resolution of ~1.25 cnts/mm; and added a handle with a force-sensitive resistor to sense squeeze pressure. Despite belt compliance, this system (with an 85-mm range of motion) worked well for the experiments of Section 4.2.

Tagged Handles: Discrete & Continuous Control

The versatility of force feedback means that a device can change behavior while retaining the same appearance, compromising predictability. Further, a generic handle might not be appropriate for a given task. We developed a concept where behavior is



Figure 3: Tagged Handles. Detachable physical knobs with distinctive shape can be connected to a single haptic device (left). A functional prototype with top-mounted textural tags aids in switching functions, by pressing different tags.

determined by the handle attached to it (Figure 3, left), by combining force feedback and tagged objects [7]. Handles are electronically tagged; a reader near the motor shaft ascertains which handle is current. Thus, the discrete selection of application behavior is coupled with the continuous control of our haptic displays.

Although described elsewhere [9] we mention this concept here in the context of our media control experiments. The functional prototype used in the applications of Section 4.3 has five textured buttons mounted over FSR pads (Figure 3, right). One PIC microcontroller sends force measured from each button to the base via a wireless transmitter; another in the base communicates serially with the host computer. A Maxon™ 20W brush motor drives the wheel.

Rock-n-Scroll

The Rock-n-Scroll, designed for game-like interactions where a control is held continually, is a finger-sized actuated wheel mounted on a second passive, sprung axis that swings parallel to the thumb joint (Figure 4). Thus the thumb both pushes down on the wheel (rocks) with deflection sensed; and rotates the wheel at its edge (scrolls) with wheel rotation sensed. Micromo™ 1524 motor/encoders geared 6.3:1 were used for both axes,




Figure 4: Rock-n-Scroll. A form study (left) of Lego and rubber bands helped to tune ergonomics and act out scenarios. A functional prototype (right) has a left/right actuated scroll axis and a non-actuated in/out rock axis for orthogonal force sensing.

although only the encoder was used for the passive axis. Design specifics proved subtle, involving optimizing of dimensional constraints between mechanism and grip. The main parameters of mechanism iteration were spring stiffness, outer wheel diameter and friction on the wheel's edge. The prototype shown proved sensitive to hand size, and its reliance on thumb motion raises ergonomic concern not addressed here.

4 CONTROLLING MEDIA VIA HAPTIC INTERFACES

We used our devices in three categories of experiments for controlling digital media, aiming to construct modeless dynamic systems with the immediacy of real-world physical controls. Techniques for *haptic navigation and control* aid in navigating any digital media stream. *Haptic annotation* refers to physical marking of content, by manual or automatic processes. Finally, methods for *functional integration* point towards physical forms that can incorporate a variety of metaphors and techniques into a final real-world application. In practice, we have implemented the metaphors across several different devices, but we show each with only one or two for clarity.

4.1 Haptic Navigation and Control

Haptic Clutch

We can feel the road through a steering wheel and control a piano's hammer action with a key. In the same way, we can perceive and manipulate a complex virtual model through a single-axis wheel, and thus increase a behavior's power and expressiveness [4]. Selective engagement with the virtual model may require an additional user input channel, such as force or position.

We built several applications on the principle of a *haptic clutch*. Here we simulate the clutched engagement of a concentric pair of wheels (Figure 5): the outer wheel's motion corresponds to that of the physical wheel turned by the user. Pressing down on the physical wheel engages it with the virtual inner wheel;

pressing is measured with a force sensor orthogonal to the outer wheel's rotation. Both virtual wheels are modeled as inertial elements with bumps on their facing surfaces, which correspond to features in the media. The wheels couple when the bumps mesh, with the manipulated wheel driving the inner one.

In one implementation of this virtual dynamic system, video frame rate is coupled to motion of the virtual




Figure 5: Haptic Clutch. The user selectively engages the physical outer wheel with the virtual inner wheel by pressing down, and imparts momentum to the inner wheel by spinning and releasing. With a light pressure, the user can feel bumps slip by.

inner wheel. If the user pushes down firmly, the two wheels engage as a single rigid body that the user can rotate in either direction to shuttle between frames. If downward force is relaxed, the released inner wheel continues to spin with imparted momentum while the real outer wheel continues to move with the user's hand. The video's speed is thus continuously variable from slow advance to extreme fast-forward. The user can push down and shove the inner wheel to increase the frame rate, or push steadily without shoving to brake the inner wheel, with a satisfying slip as bumps fly by at decreasing speeds. The medium's physicality is thus restored – the user must exert force and dissipate the flywheel's momentum in order to stop the video.

This behavior can be produced in multiple ways; one implementation is described by the dynamic system of Eq. 1, shown in Cartesian coordinates. The kinematic state of the virtual outer wheel is mathematically equated to that of the real wheel, which in turn is directly controlled by the user (i).

$$\begin{aligned}
 \ddot{x}_o &= \ddot{x}_m & (i) \\
 F_{clutch} &= f_{\perp} h \sin(2\pi n(x_o - x_i)) & (ii) \\
 M_i \ddot{x}_i + B_i \dot{x}_i &= F_{clutch} & (iii) \\
 F_{act} &= F_{clutch} + B_o \dot{x}_o & (iv)
 \end{aligned} \quad [1]$$

Throughout this paper, x_m is the measured state of the physical wheel and is numerically differentiated for velocity. \ddot{x}_o and \ddot{x}_i are the simulated positions of the outer and inner rings respectively, expressed in normalized units (1.0 = one revolution). F_{clutch} is the force transmitted between the two virtual rings (ii). f_{\perp} is derived from the measured orthogonal applied force, and controls the degree of engagement with the virtual inner ring. h is a constant defining bump "height" and n is frames per wheel revolution (typically 10-30). (iii) models the inner ring's state: M_i and B_i are its virtual mass and damping. (iv) derives F_{act} , the actuator force

where B_o is virtual damping applied to the outer ring to increase its stability. At each time-step, we solve this system using Euler's method to determine the new state of the virtual inner wheel. Parameter values are sensitive to hardware as well as the desired behavior.

The reader may observe that due to the mathematical symmetry of (ii), this system initially appears unable to transfer energy to the virtual inner wheel. If we assumed a stationary controlling x_m and a nonzero initial \dot{x}_i while braking, F_{clutch} would integrate to zero. However, in reality the user's control is not rigid or uniform. The compliant interaction between hand and mechanical system induces small bump-linked movements of x_m and f_{\perp} , which impart asymmetry to F_{clutch} 's periodicity and a corresponding net energy transfer.

Browsing digital audio and voicemail with a wheel and an absolute slider were similar to the video experience and also brought new challenges. Audio required pitch correcting for comprehensibility at arbitrary play rates [1]. We also applied the clutch in two dimensions using a Phantom: for perusing an image, we replaced the video's frame bumps with height map based on edges, brightness or color. Similarly, a three-dimensional grid can be selectively engaged by pressing a button or force sensor mounted on the surface of a stylus. We discovered that clutching is less compatible with direct-editing tasks like painting and sculpting, because the clutch forces interfere with the interaction forces generated by virtual material laydown and removal [19]. However, the structure imposed by the clutch might be well suited to CAD applications.

Haptic Fisheye

Here we manipulate an intermediate virtual model by continuously varying haptic resolution based on the user's orthogonal pressure. As with graphical fisheye views [17], this supplies immediate access to fine and coarse details of the manipulated model as a non-modal and continuously varying process (Figure 6). In video browsing, a strong applied force either decreases or




Figure 6: Haptic Fisheye. In this version, the user presses the knob to increase the resolution of browsed media, rather than its speed.

increases the rate of frames passing under the user’s fingers, displayed as a fine texture; the choice of polarity depends on the user’s default viewing mode. As pressure changes, the rate slows, until the frames are felt as individual ticks. Thus the user can rapidly browse an entire video and still find individual frames or scenes by altering the applied force.

The resolution-proportional-to-pressure version of the haptic fisheye can be implemented as Equation 2:

$$\begin{aligned} \dot{x}_s &= \dot{x}_m K(1 - f_{\perp}) & (i) \\ x_s &= \int \dot{x}_s dt & (ii) \\ F_{act} &= h \sin(2\pi n x_s) & (iii) \end{aligned} \quad [2]$$

\vec{x}_s describes the media stream’s motion: we constrain \dot{x}_s to \dot{x}_m , scaled in a non-Newtonian manner by the momentary orthogonal force f_{\perp} , then numerically integrate. F_{act} , the computed actuator force, displays a sinusoidal texture related to stream rate while h , bump height and n , the number of detents (or frames) per revolution, are constant.

Frictionless Shuttle

The *frictionless shuttle* is one of the few models where the physical wheel moves when not touched by the user. Under user control, it provides evenly spaced haptic detents that correspond to video frames. One full revolution might correspond to 30 frames of video. If the user lets go of the wheel while it is moving, it will continue on its own at the same rate. If this were implemented with a non-actuated wheel, the rate would diminish due to friction; but here we can maintain a strict correspondence between the rotation of our wheel and the advancing frames. Thus, the user can initiate any rate from a single frame advance to rapid fast-forward. While slightly dangerous, this behavior exemplifies the type of magical behavior possible with haptic feedback – in this case removing friction from a mechanical system.

While the wheel is touched, haptic feedback is a sinusoid of fixed spatial frequency:

$$F_{act} = h \sin(2\pi n x_m) \quad [3]$$

Untouched, velocity is maintained at the let-go rate with PID control on error between target velocity \dot{x}_t and measured velocity \dot{x}_m at release:

$$\begin{aligned} e &= \dot{x}_t - \dot{x}_{m|release} & (i) \\ F_{act} &= K_p e + K_i \int e dt + K_d \frac{de}{dt} & (ii) \end{aligned} \quad [4]$$

We tried several methods of pinpointing wheel release. The most successful compared orthogonal user force over a sample window; the signal is smoother in non-contact. Handle capacitance can sometimes be monitored [5], but this might constrain both continuous rotation and handle material. If only a position signal is available, we found we could watch for a smooth slowdown, then search back to the last occurrence of hand jitter noise and use that \dot{x}_m .

4.2 Haptic Annotation

We used haptic annotation to mark, highlight or delimit significant segments of video and audio material, a function useful at varying complexity and abstraction to professionals and casual browsers alike. Most simply, we literally represent the media’s form – e.g. frames of video or temporal audio intervals. The next level can be automatically extracted – scene breaks, activity, color, brightness, location or time. The highest level requires human intervention to indicate qualities such as actor, mood, genre, etc.

Any iconic representation requires a mapping from parameter to sensation [3]; “hapticons” require a relatively abstract correspondence and are challenging to make both perceptible and memorable. Allowing users to design their own maps, as done here, is one laborious and individualized approach. Developing a more universal language is an ongoing project.

Foreshadowing

In viewing, a visual mark typically appears at the moment of annotation and inevitably is overshoot, particularly with unfamiliar footage or annotations. We *foreshadow* marks haptically by gradually increasing the amplitude of a pre-annotation before reaching the mark from either direction. The wheel is used as a conventional spring-centered video shuttle knob, where deflection sets frame rate.

A texture whose intensity gradually increased in magnitude or frequency proved most effective (Figure 7). E.g., an annotation texture is overlaid as a vibration on the spring force, with frequency increasing as the mark approaches. A user may make new marks while browsing by firmly pressing down on the wheel,




Figure 7: Haptic Foreshadowing. As the user approaches marks in the stream, a texture is overlaid on the wheel’s spring restoring force; the texture gradually rises and falls around the point of interest, alerting the viewer to the upcoming event. Users can also add marks by pressing down.




Figure 8: Alphabet Browser. Hierarchically arranged audio tracks are quickly indexed without a screen using haptic and audio feedback. At high rates of rotation (above) spoken letters accompany detents; at slower rates (below), individual artists and then titles are spoken.

engaging a force sensor. We tried an active haptic nudge as an alternate marking method, but this interfered with the dynamic act of browsing – nudging is most effective when the user’s hand is static and receptive. Modifying viscosity and friction similarly interfered with navigation.

Alphabet Browser

Browsing through large media collections like MP3 recordings is tedious. Screen interfaces are the norm, but portable device screens are small, difficult to read and can unduly divert attention. The *alphabet browser* uses a haptic knob with an auditory display to browse audio collections eyes-free (Figure 8), and might be most useful in driving or portable contexts where visual attention is least available. Turning the knob activates a spoken and felt alphabetic index. When the knob is turned rapidly, one hears the first letter from each entry - ‘A’, ‘B’, ‘C’; full titles emerge at slower rates. Haptic detents enhance the audio feedback, aid navigation and indicate volume of material under each heading: alphabet letters get strong clicks, individual titles gentle clicks that fuse with rapid rotation.

We implemented an alphabet browser for an MP3 audio player on Bigwheel (artist names are traversed alphabetically and selected by pressing); and for a voice mail collection using the slider. We observed users adjust scroll rate continuously to control the amount of artist or caller name revealed as the search narrowed, suggesting an optimal but probably nonlinear relation between scroll speed and list traversal rate which will depend on typical entry lengths. Audio feedback alone provided some utility, but haptic annotation seemed to improve user’s speed, accuracy and confidence of navigation as well as their aesthetic appreciation. Other application possibilities include a haptic dial integrated with a cell phone for the fast retrieval of numbers, a car audio control for radio channels and audio tracks, and email on a wireless PDA.

Sticky Channels

Conventional manual interfaces such as channel-change

knobs have detents at channel boundaries, but with current television remotes, users generally must remember a numerical association with a channel. With active force feedback we can customize the feel of individual detents to reflect frequency of use, genre and other characteristics. Like wagon trails, *sticky channels* are ruts that get worn into the haptic landscape.

$$i = \text{floor}(dx_m) \quad (i) \quad [5]$$

$$F_{act} = h_i \sin(2\pi dx_m) \quad (ii)$$

In Equation 5, the current channel number i is computed from hand position assuming a regular spacing as an index into an array of channel detent strengths h_i . We applied this construct to switching television channels, digital audio tracks and voicemail recordings, on our wheels, the tagged handle, slider and brake. In our scenarios the annotation might be made *a priori* based on popularity, genre based on station ID or predicted user preference; or a user could set favorites explicitly. The haptic cues seemed to facilitate navigation and generated positive response from heavy TV viewers chosen outside our group. The slider provided a redundant cue of absolute position, speeding navigation, but its display set was limited.

Video Carousel

We extended *sticky channels* to a three-dimensional graphical ring of TV channels for the Brake and Rock-n-Scroll (Figure 9). With the brake, a channel initially fills the entire video screen with dynamic content. With




Figure 9: Video Carousel with Sticky Channels. Haptic detents mark channels on a graphically displayed ring (top, zoomed view). With a motor, favorite channels have a stronger attractive force; with the brake, more friction. However, when the brake doesn’t stop exactly on a frame boundary (bottom), the frames must be slewed into alignment.

handle rotation the visual channel slides to one side with a click, and the adjacent channel slides in. We found that users appreciated “channel snapping”, i.e. bringing the nearest channel to the screen center on a pause. With active feedback, we corrected the resulting mismatch between visual and manual position by servoing the handle into place. With the brake, we had to slew frames into their centered position as handle velocity decreased, anticipating a full stop.

At higher velocities, the visual display gradually zooms to a view of a ring of available channels. The current channel is live, while stills updated at multi-second intervals represent the others. With Rock-n-Scroll, we used the scroll axis to change channels and the rock axis for zoom. This decoupling of functions, separating time and velocity dependence, proved the most popular.

Absolute Media Browsing

We used the slider to tie absolute position to the current position in a media stream. For digital movies, the slider slowly advances with the playing movie. However, the user can pull the slider forward or back, feeling a spring connecting the handle to their current play position and gently returning to that position when let go. For voicemail, this method is implemented hierarchically – at the top level, the slider browses individual messages, playing the brief introduction and haptically providing information about the time and importance of the call. Squeezing the slider pops the user down into an absolute traversal of an individual message – the slider is drawn to the far left and then advances linearly with the message. As in our video example, the user can pull the slider forward or back to review, fast forward or listen slowly. Squeezing again pops the application back up into multi-track browsing.

This implementation proved overly modal for haptically inexperienced users, who were uncertain whether they were in list or individual message mode. While state was evident haptically with a light touch, some grasped the handle too tightly to notice this; haptically sophisticated users found it much more intuitive. A similar scheme for digital music files garnered little enthusiasm: casual listeners wanted to listen straight through rather than jump within a track.

Super-Sampling

The differing resolutions of the haptic device and browsed material can complicate implementations. A typical stream contains thousands to millions of elements, whereas the haptic display’s position resolution is at best thousands of counts per revolution. Media elements must be filtered or super-sampled over the haptic servo interval to produce a suitable output [20]. To retain control over individual elements, we coupled the physical device’s position to the exact media position with a virtual spring [16]. With encoder-

count-sized jumps, the spring pulls the probe along the media. The probe’s position is computed using real numbers, and is thus not aliased at the encoder resolution and has the well-behaved derivatives crucial for multirate display of digital audio.

4.3 Functional Integration

With two final projects, we aimed to prototype a complete device concept bringing together diverse functionality in a seamless and modeless manner.

Tagged Handles

We designed a suite of five behaviors for the tagged handle wheel to prototype a broadband universal remote; our goal was to provide a consistent tactile interaction across disparate media. Buttons on the wheel’s base select digital media target (e.g. audio library, TV or video-on-demand), while the functions applicable to those targets (e.g. sticky channel behavior or frictionless shuttle) are assigned to the textured pads on the wheel’s face. When the media target changes (from TV channels to audio tracks), metaphor and haptic feedback do not.

In general, users found this method of applying browse tools to different media genres intuitive. However, the device itself was unsuccessful because it required physically or visually searching for a tag on the rotating knob face. Implementing the same classes of behavior on a side-mounted wheel [9] seems promising.

Preview Button

We found the *preview button* to provide one of the most intuitive ways to combine discrete and continuous control. We installed pressure sensors on the surface of normal pushbuttons, allowing the user to preview the button’s action before committing (Figure 10). For example, when using a row of preview buttons to select radio stations, a light touch gently fades the sound up



Figure 10: Preview Buttons. This engineering prototype shows buttons retrofit with force sensors, so that a light touch can provide a preview of a button’s behavior – for example fading up an audio track, or gradually enlarging a picture-in-picture.

while the previously selected channel continues – if it passes muster, a firm push engages the track. The preview button can also be used for TV picture-in-picture so that a video inset grows based on the pressure, providing a preview of an alternate channel. The preview button can be outfit with haptic feedback, so that the preview is textural – haptic annotation representing the genre of media or a coarse representation of an audio signal.

5 OBSERVATIONS

These prototypes have been used by ourselves and ~50 fellow researchers over the project’s 12-month life; frequent, informal sessions with regulars, novices, enthusiasts and skeptics formed a crucial aspect of our iterative design approach. In-depth studies in most cases would have been premature, but we nevertheless obtained critical insights that integrate the smaller lessons interspersed in the previous descriptions.

- These dynamic system metaphors promise to give users functional integration together with simplification. Integration worked well, for example, with tagged handles, and people adapted easily to a multiple-behavior model.
- It is abundantly clear that we need to know more about haptic language: the perceptibility, salience and associability of complex haptic signals.
- Hand-crafting and quality of haptic experience were essential to our techniques’ acceptance, with many users simply enjoying the feel of the tools themselves.
- Textures generally worked better than forces for emphasis and annotation. Varying compliance, viscosity or inertia was less salient than, e.g., noise frequency. With the clutch and fisheye, which use textural marks, users were able to rapidly locate individual frames.
- Careful physical and visual affordance design is critical for these close-coupled applications. We had to explain where and how to interact with prototypes implemented on the general-purpose platforms, and errors were common. Physically customized versions often eradicated these problems.
- Compliantly mounted haptic displays reduce the impact of changes in texture and feature size. Some of Rock-n-Scroll’s applications suffered from the rock axis absorbing subtle haptic signals and reducing its controllability.

The type and amount of haptic feedback to include in a complete system remains an open question. Balancing its limitations, we did find that passive force feedback eliminated fear and surprise from some novice users. Certain metaphors worked better with the brake because its features are so solid. Stickiness seemed to register subconsciously for some, who found themselves stopping on “favorites” without knowing why.

6 COMPUTATIONAL ARCHITECTURE

Our applications ran on two-CPU systems of 1998 vintage, with one processor dedicated to haptic feedback while the second managed the media. The haptic server was a PC-based system running QNX, a real-time UNIX. The media server was a PC with a digital disk recorder for video experiments, an SGI O2 for three-dimensional video, or an Apple running MAX for digital audio. Inter-CPU communication employed a custom RS-232 serial protocol that proved simple, reliable and just fast enough.

The haptics software used a custom architecture [8]. Requiring an environment where a non-programmer could rapidly prototype custom dynamic systems, we created a system that could hide details of scheduling, communication, thread and device management but register a callback function (one line to half a page) to implement the low-level haptic model at as low a level as desired. This approach is at odds with current trends in commercial haptics software architectures, where developers use a high-level toolkit of primitives such as springs and boundary-representation but cannot modify implementation. As harsh as the economic constraints are, we believe that for haptics to become successful in the mass market, such hand-crafting is necessary.

7 CONCLUSIONS AND FUTURE WORK

With Interval Research’s impending demise, this work was curtailed as it approached a full ripening. In some cases, such as the Tagged Handles and Rock-n-Scroll, we iterated partway to a solution and felt close to an optimal form. In others, such as the Slider, we have only conceptual sketches showing a device in the side of a cell phone or remote control – physical prototyping is essential to see if these form-factors would really be pleasing and practical. Some concepts continue in altered contexts (e.g. driving controls and media browsing) in current projects at UBC.

Our own examination of the cost and feasibility of embedded implementations of our techniques combined with evidence of other products shipping in 1999 demonstrate that embedded haptic feedback will be in our future; power requirements are the greatest challenge, particularly for portable displays. We are investigating novel power schemes specifically for haptic displays.

We also believe that it is possible to completely encapsulate techniques such as the fisheye or clutch into a general-purpose haptic device such as a mouse. In this case only high-level information need pass to the application and these techniques could become part of a commodity product requiring no special communication to the host.

In designing haptic media controllers, we want to maximize both the rate and vocabulary of information

transfer. What types of forces and magnitudes can be combined without interference, capture or blocking? What do (or could) sensations mean to users? The development of a haptic language and of a flexible, multimodal realtime control platform is a foci of interdisciplinary work at UBC.

Finally, some of these techniques may ultimately be redesigned in a passive mechanical form, eliminating the need for powered devices and opening the door to portable devices.

ACKNOWLEDGMENTS

This paper describes much of the final 12 months of work by the haptics team at Interval Research in 1998-99. Tina Blaine, Cy de Groat, John & George Kembel, Camille Norment, Tim Perkis, Terry Winograd and Tricia Wright helped with ideas and early prototypes. Kim Johnson and Oliver Bayley made early contributions to the metaphors for marking and browsing media. Jesse Dorogusker, John Ananny, Brad Niven and Lee Felsenstein helped build most devices.

REFERENCES

- [1] M. Covell, M. Withgott, and M. Slaney, "Mach1: Nonuniform Time-Scale Modification of Speech," Proc. of the IEEE Int'l Conf. on Acoustics, Speech, and Signal Processing, 1998.
- [2] T. A. Galyean and J. F. Hughes, "Sculpting: An Interactive Volumetric Modeling Technique," Proc. of SIGGRAPH, vol. 25(4), pp. 267-274, 1991.
- [3] W. Gaver, "Auditory Interfaces," in Handbook of Human-Computer Interaction, M. G. Helander, T. K. Landauer, and P. Prabhu, Eds., 2nd ed. The Netherlands: Elsevier Science, 1997.
- [4] B. Gillespie, "The Virtual Piano Action: Design and Implementation," in Proc. of Int'l Computer Music Conf., Aarhus, Denmark, pp. 167-170, 1994.
- [5] K. Hinckley and M. Sinclair, "Touch-Sensing Input Devices," in Proc. of ACM CHI'99 Conference on Human Factors in Computing Systems, 1999.
- [6] E. L. Hutchins, J. D. Hollan, and D. A. Norman, "Direct manipulation interfaces," in User Centered System Design: New Perspectives on Human-Computer Interaction, D. A. Norman and S. W. Draper, Ed. Hillsdale, NJ: Lawrence Erlbaum Associates, 1986.
- [7] H. Ishii and B. Ullmer, "Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms," in Proc. of CHI '97, Atlanta, GA, 1997.
- [8] K. E. MacLean and S. S. Snibbe, "An Architecture for Haptic Control of Media," in Proc. of the 8th Ann. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME / IMECE, Nashville, TN, DSC-5B-3, 1999.
- [9] K. E. MacLean, S. S. Snibbe, and G. Levin, "Tagged Handles: Merging Discrete and Continuous Control," in Proc. of CHI '2000, The Hague, Netherlands, 2000.
- [10] A. Madhani, G. Niemeyer, and J. K. Salisbury, "The Black Falcon: A Teleoperated Surgical Instrument for Minimally Invasive Surgery," in Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS), Victoria, BC, Canada, 1998.
- [11] T. Miller and R. C. Zeleznik, "The Design of 3D Haptic Widgets," ACM Symposium on Interactive 3D Graphics, pp. 97-102, 1999.
- [12] H. B. Morganbesser and M. A. Srinivasan, "Force shading for shape perception in haptic virtual environments," in Proc. of the 5th Ann. Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems, ASME/IMECE, Atlanta, GA, DSC:58, 1996.
- [13] W. Murch and F. F. Coppola, In the Blink of an Eye : A Perspective on Film Editing. Los Angeles: Silman-James Press, 1995.
- [14] M. Naimark, "Movimap and a 3D Panorama," SPIE, vol. 3012, pp. 297-305, 1997.
- [15] W. Plesniak and R. Pappu, "Coincident Display Using Haptics and Holographic Video," in Proc. of CHI'98, 1998.
- [16] D. D. Ruspini, K. Kolarov, and O. Khatib, "The Haptic Display of Complex Graphical Environments," in Proc. of SIGGRAPH '97, Los Angeles, CA, pp. 345-352, 1997.
- [17] M. Sarkar, S. S. Snibbe, O. Tversky, and S. Reiss, "Stretching the rubber sheet: A metaphor for viewing large layouts on small screens," in Proc. of UIST, pp. 81-92, 1993.
- [18] B. Schneiderman, "Direct manipulation: A step beyond programming languages," IEEE Computer, vol. 16, 8, pp. 57-63, 1983.
- [19] S. S. Snibbe, S. Anderson, and W. L. Verplank, "Springs and Constraints for 3D Drawing," in Proc. of the 3rd Phantom Users' Group, Dedham, MA, MIT AI Lab TR #1643, 1998.
- [20] J. Watkinson, The Engineer's Guide to Motion Compensation: Petersfield: Snell & Wilcox, 1994.