CENTER FOR COMPUTER RESEARCH IN MUSIC AND ACOUSTICS
OCTOBER 1995

DEPARTMENT OF MUSIC
REPORT NO. STAN-M-95

THE MOOSE: A HAPTIC USER INTERFACE FOR BLIND PERSONS WITH APPLICATION TO THE DIGITAL SOUND STUDIO

RICHARD BRENT GILLESPIE
SILE O'MODHRAIN

RESEARCH SPONSORED BY
THE CCRMA INDUSTRIAL AFFILIATE PROGRAM,
THE STANFORD OFFICE OF TECHNOLOGY LICENSING
AND THE FULLBRIGHT SCHOLARSHIP PROGRAM

CCRMA
DEPARTMENT OF MUSIC
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305-8180

The Moose: A Haptic User Interface for Blind Persons

Sile O'Modhrain and Brent Gillespie
Center for Computer Research in Music and Acoustics (CCRMA)
Stanford University
sile@ccrma.stanford.edu, brent@ccrma.stanford.edu

October 16, 1995

Abstract

This paper presents our work to date on a haptic interface whose immediate aim is to provide access for blind sound engineers to the graphics-based computer interfaces currently found in digital sound studios. We describe the hardware and supporting software which together reinterprets a Microsoft Windows screen for the haptic senses.

With a haptic interface, screen objects such as buttons, sliders and pull-down menus are presented mechanically to the user's haptic senses (kinesthetic and tactile), where they can be felt, located, identified, and, through the use of the same device for input, activated. We have built a prototype two-axis device which operates much like a mouse, except that it is also able to move under its own power. Using this prototype device we have already implemented an interface for Microsoft Windows and have proven the feasibility and usefulness of the haptic interface approach for non-visual computer access. We expect that haptic interface devices will become standard computer interface tools, supplementing the visual presentation with haptic presentation for all users. A more holistic presentation of information will be made by the computer. This approach is, we believe, particularly valuable in the design of application interfaces for digital audio editing. We have implemented a prototype sound file editing application in which the sound waveform is presented haptically as a contoured groove within the workspace.

1 Introduction

'Haptics' refers to the human tactile (cutaneous) and kinesthetic (muscle movement) senses. A haptic interface is a computer-controlled motorized device to be held in the hand by a user, which displays information to that user's haptic senses. It is an extremely powerful modality for interface design because the same device can be used for both displaying output from the computer and accepting input from the user. Moreover, using haptics significantly reduces the burden on other information channels such as vision and audition, thereby freeing these channels for other tasks. In certain instances it is, we believe, possible to completely substitute haptics for other sensory modalities. In this way, graphical information displayed on a computer's screen can be made accessible to blind persons who at the moment are denied access to standard Graphical User Interfaces (GUIs). For example, by producing forces on the user's hand which are a function of both the user's motions and properties of the icons under the cursor, touchable representations of the screen objects can be created.

In particular, we are interested applying haptics as a sensory substitute for the graphical interface of today's digital audio editing applications. A sound engineer's audio channel is already occupied while using such applications, so our display cannot make good use sound or speech output. Our eventual goal, therefore, is to design a display which only uses haptics. The most natural format to adopt, especially when we wish to grant access to existing third-party applications, is the 'powered mouse', since it so easily adapts to the mouse paradigm. Our mouse, however, also functions as an output device, taking over the role of the screen. Our use of haptics for audio editing effectively brings us back to the analogue audio studio, where sound editing was primarily a hands-on affair–a move which will be welcomed by many a modern sound engineer.

A handful of other research groups are working on powered-mouse type interfaces for GUIs. In Wisconsin, the TRACE Group is developing a computer access tool for the blind based on haptics [Wiker 1991]. At the university of Tokyo, a mouse with both vibrotactile and force feedback has been developed. [Akamatsu, Sato 1994] At Stanford in the 1970s the Optacon was developed by John Linvill, and stands as one of the first commercialized haptic display devices. [Linvill 1973] Other access devices for the blind and deaf-blind based on haptics have also been developed [Loubal 1992], [York 1989] [Frisken-Gibson 1987] [Eberhardt 1993] [Kelly, Salcudean 1994] See [Minsky 1995] for a comprehensive overview of current haptics technology.

Our project is somewhat unique in that it must pay very careful attention to the manner in which information is displayed through the audio channel. Our constant directive is to substitute visual presentation with haptic presentation alone.

In section 2, we introduce the architecture of our prototype interface. We document the design of a Haptic User Interface (HUI) which translates certain elements of the standard Graphical User Interface (GUI) into a form where they can be felt. Section 3 details the specific goals for which this research was undertaken. Section 4 summarizes.

2 Our Prototype Interface

We have designed and built a working prototype which serves to prove the concept and has generated quite a bit of enthusiasm from those who have seen (felt) it. It is basically a powered mouse, giving the user the ability to feel the screen objects under the mouse cursor. This mouse is used to navigate the screen like a regular mouse; but by reflecting forces (produced by its motors) back to the user, it presents a haptic representation of the symbols on the screen. In other words, as the powered mouse is moved, its position is continuously compared against the screen image. If the mouse should alight on an icon, a haptic representation of that icon will be presented by the motorized mouse for the user to feel and explore. The edge of a window, for example, might feel like a detent under the mouse. A button might feel like a patch of sandpaper. Once a desired icon has been found, it may be selected using a typical mouse button. A user can explore a screen, activate 'buttons' from menus, and select other screens which in turn will be mapped and haptically displayed. For text which appears on the screen (file names, button labels, etc.) we have experimented with integrating a speech synthesizer into the interface. We are exploring the presentation of altogether new information through haptics. For example, the next appropriate user action in a given context can be indicated by causing the powered mouse to gravitate to a particular icon.

2.1 Hardware

Figure 1 shows the hardware components of our present planar haptic interface which we affectionately call 'The Moose'. The puck or manipulandum in the center is coupled to two linear voice coil motors through two perpendicularly oriented flexures.

The unique feature of our hardware design is this double flexure. On the present prototype, the double flexure is executed in two pairs of foot-long strips of spring steel. The double flexure conveniently decouples the 2-axis motion of the puck into two single-axis motions at the linear motors. Moments and vertical forces are resisted, yet translations in the horizontal plane are transmitted directly to the motors by the manipulandum. The kinematics of this device are simple and very nearly linear, making forward and inverse kinematic calculations unnecessary. Furthermore, the workspace is flat, square like a mousepad, and free of singularities. The entire workspace is also naturally counterbalanced. This design ensures that very little inertia is added to the motors. Over-limit forces will cause buckling of the flexures, which we consider a safety feature. The only real disadvantage of the double flexure design is added high-frequency structural resonances inherent in the flexures themselves. These resonances will bandlimit the display capabilities. But if chosen high enough by design, they should not interfere with the bulk of haptic object images.

A simple Digital I/O card provides for PC-bus communication to four 12-bit DACs and four quadrature counters. The voltage outputs of two DACs, ranging +/- 5 Volts, feed to two transconductance amplifiers based on the LM12 power op amp and in turn to the motors. A linear position encoder, 150 lines per inch, reads position on each of the motors while the count circuit maintains an up-to-date binary representation of position. Other digital switch inputs such as buttons can be polled from software. Finally, a speech synthesizer linked through the serial port is available.

Future hardware enhancements will include the following: 1) a braille display to take the place of the speech synthesizer for text output, 2) the use of braille cells for "shape" display, and 3) the use of a small voice coil motor for vibration and texture display [Kontarinis 94].

2.2 Software

Various control routines which create haptic effects such as virtual springs, textures, and buttons have been developed and incorporated into our Windows interface. By combining these primitives we have begun to construct a library of "hapticons" each of which correspond to a standard

Windows icon. For example, our haptic checkbox has a frame surrounding the checkbox text and a detent corresponding to the checkbox state indicator. Just as the state indicator changes color when the checkbox is checked, so also our haptic checkbox state indicator changes from a detented spring to a repelling pyramid which is immediately apparent when the profile of the checkbox is examined.

Figure 2 is an outline of the architecture of our software. Our software is divided into three distinct modules:

2.2.1 The Icon Management Class

This module is responsible for "mapping" each new screen as it appears and storing information, such as icon dimensions and icon names, about each icons it finds in a linked list.

2.2.2 The Hapticon Management Class

This module queries the icon manager, using the obtained information to construct a list of corresponding hapticons. It encapsulates the haptic properties of each hapticon in control laws and lookup tables for convenient use by the control module.

2.2.3 The Control Module

The control module is responsible for executing the control loop. Its action is embedded in the Windows message loop. It polls the current moose position, moose button status, and current window identifier. If the current window has changed, it initiates the mapping of the new window by the icon manager and requests the hapticon manager to update its hapticon list. The control module constantly passes the current moose position to the hapticon manager and receives a force appropriate for whatever icon lies at that position. The control module then outputs that force to the moose. Moose button clicks and moves are passed through to the Windows mouse button control routine.

3 Specific Goals

We are currently experimenting and developing a palette of haptic effects which will be used to explore and allow comparisons among various haptic substitutes for graphic objects- i.e. detents for buttons, solid blocks for inaccessible objects, compliant and non-compliant borders for windows, and so on. We hope that this research will result in a characterization of graphic interface objects and that a corresponding library of haptic effects for representing these objects will be created. In time we hope that a common practice for haptic interface design will arise.

The project has attracted the interest of Neil Scott and his team at Stanford's Center for the Study of Language and Information, who are eager to incorporate our work into their Total Access Port (TAP) system. TAP is aimed at developing a generic adaptive interface port through which interface device signals can be intercepted and hence made available to whatever access device a disabled user finds most appropriate. The system's broad goal is to provide the individual with one personalized interface which they can bring with them to whatever computer they need to use. The haptic interface is a very realistic option for conveying the contents of a GUI to a blind person. Scott's team will substitute our icon management module with their own vision recognition routines which obtain the icon information directly from the video signal rather than from the Windows environment. The advantage of this system is that it will be platform-independent and will thus allow blind people to use the same access device for any number of computers running any type of operating system.

Having proven the feasibility of substituting haptics for graphics in a prototype of a general computer user interface for the blind, we now turn our attention to designing a haptic interface for sound processing applications.

Our project's goal for some time has been to provide the ability to cursor through an audio stream as one does with a tape head on magnetic tape. We are now in fact able to both haptically explore and hear the sound signal simultaneously. The contour of the recorded sound on the tape is felt as a virtual groove embedded in the workspace while the same signal under the hand is heard. We have most recently added a third axis to our display platform for the purpose of playing back the sound signals being edited. By using the DACs already available in our haptic display hardware, we sidestepped the need to develop the requisite real-time audio tools for the PC. Existing PC sound cards are simply not capable of real-time audio processing. Also, we require more than just sound triggering, which is all that MIDI driven synthesizers could have provided.

We now have a system which duplicates much of the analogue tape machine interface. It allows for the display of a soundfile both haptically and graphically and allows one to cursor through the sound, examining it's contour in any one of four "zoom" levels. The user may 'throw' the virtual tape, which will continue to move past their hand as if it had inertia. See [Gillespie 93] for a discussion of haptic display with interactive dynamics. As it scrolls past, the stored audio signal will be played back. The user may drop marks and feel for previously marked passages, and perform cut cut and splice operations in an intuitive manner. But beyond this, we have programmability –the opportunity to explore transformations between the haptic and audio presentations of the media under scrutiny. In many respects, we can do better than the analogue tape machine. By implementing a looping-buffer process, we can cursor through the signal while allowing for dynamic relationships between the tape velocity and sounded pitch. When cursoring stops, for example, the output loops on the segment of audio lying under the cursor at its real pitch. Moreover, the speed at which the cursor passes over the signal in this case only determines the streaming speed and not the pitch of the audio output.

Other types of pre-processing of the audio signal would, we believe, be worth exploring. An audio event detector [Chafe 1986] could be used to place haptic "landmarks" on the events of special interest in the signal. For example, speech, silences, tone onsets, and decay anomalies can all be given special haptic characteristics. Other confusing content could be screened out.

4 Summary

Haptic technology is particularly well suited to solving the problem posed by this project –namely to make digital sound editing applications accessible to blind sound engineers. Unlike existing access technologies for the blind, most of which rely heavily on speech output, it makes no demands on the auditory channel. The ears are left free for the task of editing sound.

A new motorized mouse takes over the function of the conventional mouse while conjoining the output role the screen. With our powered mouse, a blind user can navigate and interact through application's window. Sighted users may also realize advantages in speed and dexterity.

Using Windows Internals, we have transcribed the visual information of the screen and made it available to the haptic senses. The Windows environment enables inquisitive software such as ours to access all information on the screen. Our software simply gathers that information and displays it haptically. The real advantage of our haptic interface over a speech screen reader is that information about the icon and window topology is presented directly and immediately rather than through time-consuming descriptive language.

We have explored the new opportunities afforded by haptic display within a prototype audio editing environment. We have used the old analog tape and reel metaphor to design an interface to an audio editing application. Similar principles could be used to effectively display the information in graphs and diagrammatic figures. We foresee that in such applications, haptic interface technology will not only help blind people but will also supplement graphics for sighted computer users.

References

[Akamatsu, Sato 1994] Akamatsu, M., Sato, S. and MacKenzie, I.S. *Multimodal mouse: A mouse-type device with tactile and force display* Presence, Vol 3 No. 1, Winter 1994 pp 73-80.

[Chafe 1986] Chafe, C. and Jaffe, D. *Source separation and note identification in polyphonic music.* IEEE Conference on Acousics, Speech, and Signal Proccesing, Vol 2, 1986, pp. 1289-1292.

[Eberhardt 1993] Eberhardt, S.P., et al. *OMAR: a haptic display for speech perception by deaf and deaf-blind individuals.* Proceedings of IEEE Virtual Reality Annual International Symposium. Held: Seattle, WA, USA, 18-22 Sept. 1993. p. 195-201

[Edwards 1989] Edwards, A.D.N. *Soundtrack: an auditory interface for blind users.* Human Computer Interactions vol.4, no.1, 1989 p. 45-66.

[Frisken-Gibson 1987] Frisken-Gibson, S.F., et al. *A 64-solenoid, four-level fingertip search display for the blind.* IEEE Trasactions on biomedical engineering. (Dec. 1987) vol.BME-34, no.12, p. 963-5.

[Gillespie 93] B. Gillespie, M. Cutkosky. *Interactive dynamics with haptic display* In Proceedings of the 1993 ASME Winter Annual Meeting, New Orleans, pp. 65-72.

[Kelly, Salcudean 1994] Kelly, A.J., and Salcudean, S.E. *On the development of a force-feedback mouse and its integration into a graphical user interface.* in Proceedings ASME Winter Annual Meeting, DSC-Vol. 55-1, 1994.

[Kontarinis 1994] Kontarinis, D. A. and Howe, R. D. *Tactile display of vibratory information in teleoperation and virtual environments* Presence, in press.

[Linvill 1973] Linvill, J. *Research and development of tactile facsimile reading aid for the blind: the optacon* [Washington] U.S. Dept. of Health, Education and Welfare, Office of Education, Bureau of Education for the Handicapped, 1973.

[Loubal 1992] Loubal, P.S. *Fingertip maptracing devices for the blind.* Proceedings of Conference on Technology and Persons with Disabilities. Held: Los Angeles, CA, USA, 18-21 March 1992. p. 315-18

[Minsky 1995] Minsky, M. *Computational Haptics: Texture* PhD Thesis, MIT Media Labs, 1995.

[Wiker 1991] Wiker, S.F., et al. *Development of tactile mice for blind access to computers: importance of stimulation locus, object size, and vibrotactile display resolution.* Proceedings of the human factors society 35th annual meeting. vol.1 Held: San Francisco, CA, USA, 2-6 Sept. 1991. p. 708-12

[York 1989] York, B.W., et al. *Tools to support blind programmers.* 17th Annual ACM Computer Science Conference. Held: Louisville, KY, USA, 21-23 Feb. 1989. p. 5-11

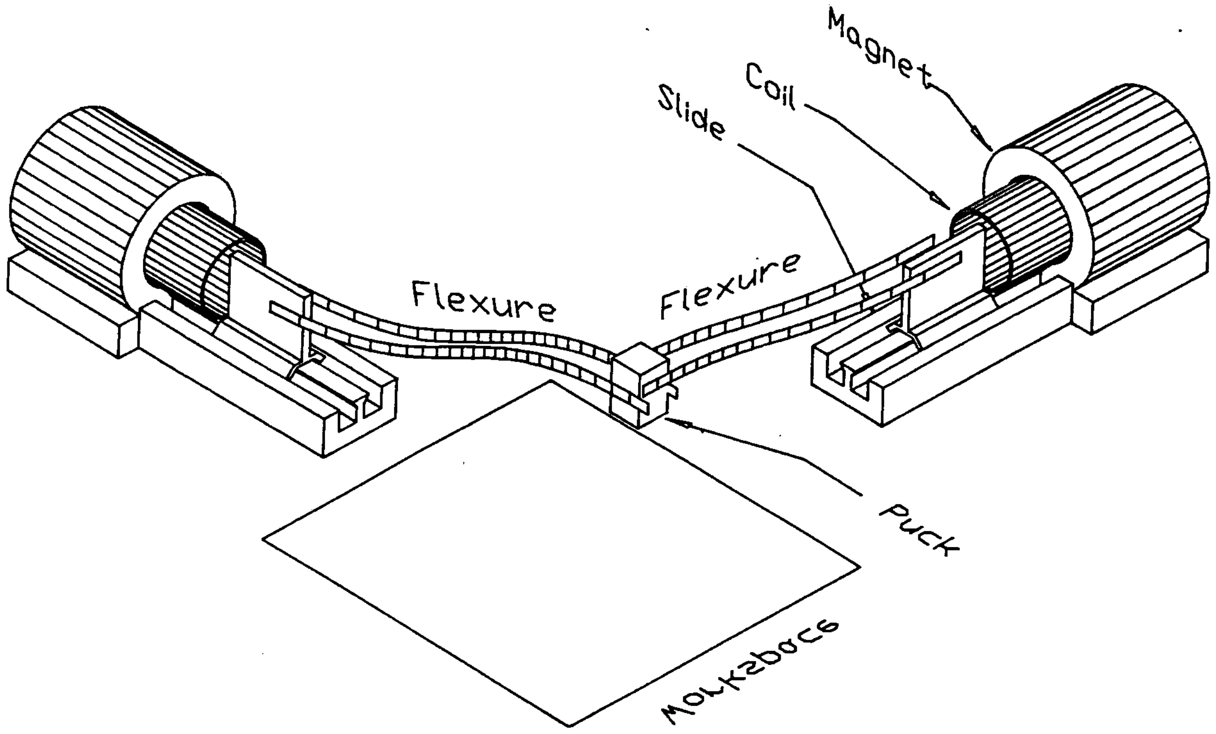

Figure 1. The Moose

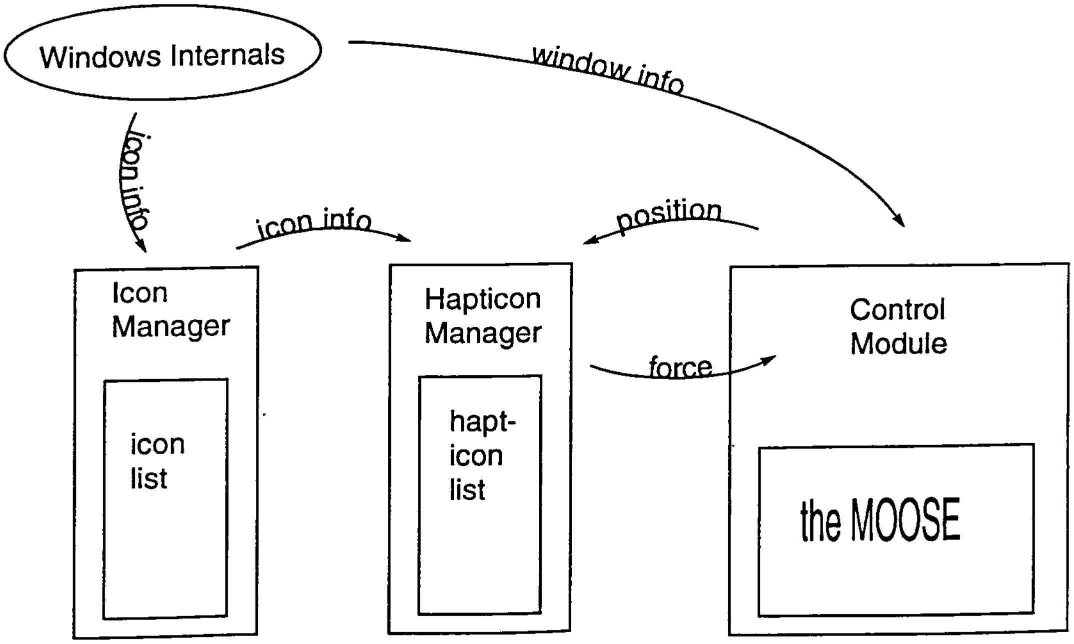

Figure 2 Software Architecture