Surgical Navigation System and Method Using Audio Feedback

Kristen Wegner

Computer Aided Surgery Incorporated
New York, U.S.A.

Abstract

We discuss an experimental audio feedback system and method for positional guidance in real-time surgical instrument placement tasks. This system is intended for future usability testing in order to ascertain the efficacy of the use of the aural modality for assisting surgical placement tasks in the operating room.

The method is based on translating spatial parameters of a surgical instrument or device, such as its position or velocity with respect to some coordinate system, into a set of audio feedback parameters along the coordinates of a generalised audio space. Error signals that correspond to deviations of the actual instrument trajectory from an optimal trajectory are transformed into a set of audio signals that indicate to the user whether correction is necessary.

An experimental hardware platform was assembled using commercially available hardware. A system for 3-D modelling, surgical procedure planning, real-time instrument tracking and audio generation was developed. Prototype software algorithms for generating audio feedback as a function of instrument navigation were designed and implemented. The system is sufficient for future usability testing.

This technology is still in an early stage of development, with formal usability and performance testing yet to be done. However, informal usability experiments in the course of the basic engineering process indicate the use of audio is a promising alternative to, or redundancy measure in support of visual display technology for intra-operative navigation.

Introduction

The overwhelming thrust of research in the area of augmented reality surgical planning and execution has been concerned with the development of visual tools and interfaces that will allow surgeons to comprehend volumetric, functional, and navigation-assisting data. Planning and intra-operative navigation systems ostensibly designed to reduce the risks and unknowns in an operating room have proven to have shortcomings when applied in real-time situations. These shortcomings, such as inadequate frame and refresh rates, poor resolution or detail [3], and various aetiologies of so called simulator sickness [16] typically render them useless except for simulated or experimental proof-of-concept surgeries. In order for the anticipated future of augmented reality assisted surgery [11] to truly become a reality either Moore's Law must be broken (i.e. massive improvements in computing architecture must be made) or alternative approaches for solving these problems must be discovered.

We believe the application of tactical audio - defined as an audio feedback expedient for achieving a goal - in augmented reality systems may provide a way to overcome the shortcomings of the state of the art. Audio feedback confers a number of advantages for human-machine interface in the operating room. The aural modality has been relatively unexplored for usage in the operating room for other than for simple indicators or gages such as heart rate monitors, etc [7]. The feasibility of using audio feedback as a basis for more complex applications becomes more apparent when one considers that audio feedback can extend the available information transmission bandwidth significantly, in part because humans are capable of processing audio information in parallel, for instance polyphony in music [2, 9, 12]. The omnidirectional nature of auditory spatial perception permits the 3-D localisation of sounds emitted from any point in space, notwithstanding occlusions¹, in sharp contrast with the limited viewing frustum of the human visual apparatus. There are many benefits to using audio from an engineering standpoint in addition to the fact that the auditory sensory modality is comparatively rich in parallel bandwidth. For example, the computational requirements for generating a 3-D audio signal in real-time are substantially smaller than for real-time 3-D graphics [12]. Thus the disorders provoked by rendering and scanning latency, such as eyestrain [6] may be transcended.

The fact that audio feedback technology avoids so many of the shortcomings of visual systems has already made it an attractive area of exploration for a wide range of non-medical applications [1, 4, 8, 10, 14]. It is perceived that intelligently applied audio feedback can considerably enhance the utility of modelling and planning technology to users who cannot tolerate the encumbrance of graphical display hardware, and whose visual faculties have pre-existing obligations, such as addressing the task at hand. Our goal is that using a tactical audio system, the operator never needs to take his/her eyes off the patient.

Hardware Implementation

This system and method are intended for assisting instrument placement for procedures in which there is significant knowledge of the context and extents of the target and access for a procedure in computer memory prior to the execution of the procedure. In general this concerns procedures in which anatomical geometry has

been rendered before the operation and is registered to the patient. The lesion is localised in the data, and there is some optimal trajectory and/or manipulation of the instrument with respect to the lesion (for example, tumour stereotaxis) or some placement of some anatomical object such as a bone fragment (for example, craniofacial reconstruction). We further specify that this procedure plan is of such a complexity or precision that it cannot be fully remembered or known by the surgeon during the operation without intermittent recourse to the pre-recorded trajectory (i.e. some mnemonic is needed). The positional and orientational states of objects are dependent upon prior placements and therefore cannot be planned except in terms of some set of constraints, such as proportions (i.e. a number of consecutive measurements are needed). In the former case, we define as a procedure requiring an insertion trajectory, and the latter as a measurement task.

A surgical target path is expressed in terms of two or more spatial coordinates, the values of which are indicative of the desired position and velocity of the instrument used in the surgical procedure along the surgical target path. This further comprises a way for translating values of the two or more spatial coordinates obtained from the measurement into corresponding values of two or more coordinates of an audio space. In particular, each of the two or more coordinates of the audio space may correspond to an audio theme recognisable by the surgeon. This, for example, this can be a consonant harmonic structure, such as a major triad, each tone of which corresponds to values along a specific spatial coordinate. Spatial coordinates are broadly considered as positional ($x-y-z$) or angular coordinate, velocity, acceleration, torques, etc.

In addition, this system further comprises means for automatically selecting different coordinates of the audio space based upon the surgical target path and the sensed surgical execution, so that a surgeon can know how close the hand held instrument is moving to a desired path simply by listening to the audio feedback. Notably, unlike visual systems that require full attention from the human for certain periods of time, an operating surgeon can correct his movements virtually without a distraction. Naturally, this feedback can be further supplemented by a corresponding visual feedback means for advising the surgeon.

Development Platform

We assembled an experimental audio feedback system using commercially available hardware and a software system we wrote for 3-D modelling, surgical procedure planning, real-time instrument tracking and audio generation. Prototype software algorithms for generating audio feedback as a function of instrument position relative to a pre-planned trajectory were designed and implemented.

The hardware system is based upon a Pentium PC integrated with a high-performance audio engine, the Lake DSP1 Huron Digital Audio Convolution Workstation [5]. This system is fitted with a Polhemus2 Insidettrak electromagnetic position tracking device. The PC implements a control system which acquires data from the Insidettrak, compares this data with the surgical plan, and then according to the specific audio feedback algorithm implemented sends commands to the DSP programs running upon the Huron. This system is depicted in Fig. 1.

Figure 1: Experimental audio feedback hardware system consisting of: (A) sensor subsystem, (B) computational subsystem, (C) sound synthesis and filtration subsystem, (D) amplification and public address subsystem, (E) graphics subsystem.

1 <http://www.lake.com/>

2 <http://www.polhemus.com/>

Software implemented upon the hardware system previously described consists of two main processes: PC executables and DSP executables. The PC executables acquire data from the tracking device(s), compare it with the preoperative plan, and use the resulting error function via some specific audio feedback algorithm to control the DSP program. We depict this function in Fig. 2.

Figure 2: Software functional block diagram (lettering A,B,C and D corresponds to the hardware blocks as depicted in Fig. 1).

The basis of our synthesis approach is a wave-table lookup oscillator program running on the Motorola DSP56002 chipset. Sound synthesis is implemented using banks of these oscillators. These oscillator banks serve as the building blocks for creating synthesizers using additive synthesis as well as other synthesis techniques. These synthesizers are controlled by the PC host using the specific audio feedback algorithm to control the various parameters. In addition to sound synthesis we employ a subsystem for filtering or spatializing the signals output from the oscillator banks. This subsystem was developed using Lake DSP's proprietary sound field simulation and auralization software.

We discuss a number of different approaches we are in the process of implementing on the aforementioned system towards the sonification of some function of instrument position in surgical placement tasks. These include a vernier-derived technique, 3-D spatialization, audio gages or callipers, use of wave-terrain and granular synthesis techniques for audio rendering of 3-D volumes, and acoustical physical modelling for simulating force feedback.

Audio User Interface Methodologies

A number of methodologies have emerged for employing audio feedback in computer and embedded applications. We define these as audio gages, mimetic audio, symbolic audio, audio for positional guidance, and musical instrument interfaces.

The most basic methodology is to use a simple monophonic signal, modulated in some way, as an alternative to readouts and gages for industrial equipment. This has proven useful in critical situations such as an aeroplane cockpit when the operator is incapable of comprehending in a single glance the states of all the readouts, perhaps some gages need not even be considered except when they reach some critical value. Audio feedback provides a very effective way to overcome this visual overload resulting from cluttered or complex display systems. Some notable implementations have included warning systems for civil aircraft and audio feedback systems for medical equipment [7, 8].

Mimetic Audio Interfaces

A mimetic form of audio feedback is used quite frequently in computer games and environment simulations. For example, when the user's avatar or cursor interacts with an object within the simulated environment, a sampled audio clip is triggered. These sounds are intended to represent or mimic sounds that would result from real physical interactions, such as a car impacting a brick wall. In this methodology sound is used more as a redundancy measure to give credence to the graphical simulation, than as a parallel informational channel.

Symbolic or “Earconic” Audio Interfaces

A common methodology applies audio feedback in the symbolic manner of icons in graphical user interfaces. With this approach programmatically generated music and sampled audio clips function as auditory icons, so-called “earcons”, which may be manipulated using a mouse or other 2-D controller within an audio desktop space [4, 26]. A more primitive example of this methodology is the use of a sound scheme in a graphical desktop environment such as Microsoft ® Windows 95 where audio clips or synthesised sounds are triggered as the result of actions such as clicking upon an icon or a system error [4, 17, 18].

Musical Instrument Interfaces

The musical instrument interface methodology is based, in part, on the understanding that the interface design methodology employed by musical instruments may serve as a model for systems which aim to provide positional guidance using of audio feedback. Musicians who play variable pitch instruments such as the violin, the trombone, or the Theremin control the acoustical aspects of their performance by varying their hand position relative to the instrument body. Sensitivity to position measurable to fractions of a millimetre is necessary for certain notes to sound correctly upon an instrument such as the violin.

The musical instrument interface presents one or more axes of control. The violin, for example, presents at least three axes of manipulation. Firstly, the musician’s fingering upon the string controls gross pitch. Secondly, another axis perpendicular to the string, parallel to the horizontal plane of the instrument body controls a smaller range of pitch called bend. Lastly, the other axis perpendicular to the string and parallel to the plane of the instrument varies amplitude and the spectral composition of the sound (via bowing). The latter two axes are small in comparison to the string axis, but are important in precisely shaping the resulting sound. There are further other axes of control but these three serve to illustrate the feasibility of this design methodology. It is worth noting that a trained violinist can accurately, precisely and repeatably place his/her hand in the same position in order to create a particular sound. This facility is also well illustrated by the classic electronic musical instrument known as the Theremin, invented in 1928 by Leon Theremin. The Theremin is unique in the fact that the musician controls pitch and amplitude by moving his/her hands in the air, relative to two antennae [19].

Figure 3: Virtuoso thereminist, Clara Rockmore, performing upon the Theremin.

Our “tactical audio” surgical navigation methodology is based on a conceptual inversion of this paradigm: instead of using position to control sound; sound is used to provide feedback as to position in manual placement tasks.

Precursors to Tactical Audio

Design methodologies using tactical audio - which we define as an audio feedback expedient for a achieving a goal, such as for providing spatial guidance for placement tasks - are of a more speculative nature despite the observation that they are similar in concept to the gage and musical instrument methodologies. These systems employ simple algorithms that translate values of two or more spatial coordinates into corresponding values of two or more coordinates of an audio space. Some experimental applications have included a 3-D auditory “visualisation” system used for simulating spacecraft maintenance tasks in zero gravity (see also [10]). This system uses frequency beat interference between two sinusoids as a means for providing feedback for properly positioning a circuit board in a training simulation [14]. Other related applications have included a 3-D visualisation system using spatialised audio [27].

Mapping Geometry to Sound

There are a number of obstacles to developing a feasible methodology for tactical audio user interfaces. For example, numerous psychoacoustic phenomena - especially the nonlinearity of human hearing[16], and the problem of determining a frame of reference for the user. The problem of determining reference frames can be stated as the decision as to the appropriate mapping of virtual space to the user's coordinate system, and to the real world.

The use of reference frame mapping in interface design is actually quite common. Consider how computer users have no trouble overcoming the positional translations which occur between a mouse which must be manipulated upon the horizontal plane of a table top, and the corresponding cursor which is projected upon a plane perpendicular to the table top. The interface has a surprising intuitiveness despite the fact that the axes may be reversed and offset, and the magnitude of movement scaled. The interface is simple and consistent; the expected outcome of shifting the mouse does not diverge too significantly from the actual movement of the cursor. See Fig. 4.

Figure 4: Simple transformation of the mouse to monitor interface.

Mapping geometry into sound places a larger challenge upon the cognitive faculties of the user than the simple transformations of the mouse-to-cursor interface. This involves what could be described as transmodal mappings instead of homomodal mappings - translating one or more axes within the same modal space. In order for the system to present a coherent interface to the user, the interface designer must determine which dimensions of the initial modality are to be mapped to the resulting modality. For systems requiring such a high degree of usability as systems intended for use in the operating room, the chosen mapping scheme must be intuitive to the surgeon. Perceptual issues are important if the transformation is desired to be as lossless as possible due to the differing, even incomparable perceptual resolutions, ranges, nonlinearities, etc. [20, 21]. We believe it is possible to overcome these obstacles. Some of the mapping schemes we have prototyped are described below.

Experimental Approaches

As we discussed earlier, the intuitiveness, in short the usability of the user interface embodied by the audio feedback algorithm will make or break the system. In the course of the engineering process we implemented a variety of audio feedback algorithms. We will use these prototypes at a later date to perform formal usability tests in order to determine which factors and which, if any, specific approach holds the most promise for audio user interface design for surgical navigation systems. We discuss some of these design approaches.

Beat Interference Method

Beat interference is perhaps the simplest approach for indicating to a user the variation of some component of instrument position with respect to some component of a desired position indicated by the pre-planned trajectory. This approach is a member of a class of approaches defined by an interface in which one or more coordinates of some function of instrument position are mapped to one or more coordinates of a generalised musical space. Using an adaptation of the Vernier technique [22,23], two reference parameter are used, sinusoids which we designate as **A** and **B**. Sinusoid **A** is fixed at some arbitrary frequency, f_A and functions as a reference or gnomon. The frequency of sinusoid **B**, f_B , which we shall call the "cursor", varies proportionally with the some error function which represents for example the difference of $f_B(\mathbf{x})$ from $f_A(\mathbf{x})$. The user corrects for error by trying to close the frequency gap between **A** and **B**.

In the context of an actual interface, for example providing feedback for error within a Cartesian space, this could take the form of three beat interference sets, one for each coordinate. The set of three reference pitches, the reference pitch set $\{f_A(\mathbf{x}), f_A(\mathbf{y}), f_A(\mathbf{z})\}$ could be chosen in order to form a consonant triad. This would imply that: $f_A(\mathbf{x}) \neq f_A(\mathbf{y}) \neq f_A(\mathbf{z})$. In navigating, the goal would be to bring the cursor pitch

set $\{fB(x), fB(y), fB(z)\}$ into some pre-defined consonant state with the reference pitch set, that is: $fB(x)=fA(x)$, $fB(y)=fA(y)$, $fB(z)=fA(z)$. Fig. 5 depicts this relationship in terms of a Cartesian coordinate system.

Figure 5: Spatial position mapped to pitch-space.

This approach produces a reasonably intuitive interface, but there are certain drawbacks. For example, it is obvious that there are solutions for which it would appear that the cursor triad were approximately in harmony with the reference set, yet in terms of the actual coordinates, the positioning would be in error, for example, if $fB(x)=fA(y)$ or $fB(x)=fA(z)$. It becomes obvious that there are many other similarly deceptive combinations. Some scheme for excluding these ambiguities must be devised if this approach is to achieve an acceptable level of usability.

3-D Audio Spatialization Method

Using 3-D spatialization is similar to the beat interference method in that one or more coordinates of some function of instrument position are mapped to one or more coordinates of a generalised audio space. In this case instrument position mapped to 3-D audio space about the user's head. This approach cannot stand alone as an interface but can be employed as a redundancy measure and extension to any other algorithms. The particular strengths of the Huron hardware with respect to audio digital signal processing [24], in accordance with nature of 3-D localisation in human hearing, permits such an approach using our system.

The ability to localise sound sources in 3-D space around the listener's head is a function of intensity and phase differences between the signals from the left and right ears. The impact of the shape of each individual's head and external ear, or pinnae, on the reflected sound waves received by the inner ear is crucial for sound localisation. The pinna has a significant influence on shaping the spectral envelope of incident sound. This spectral shaping is dependent upon the 3-D origin of the sound source with respect to the listener's head and pinnae. The auditory cortex determines 3-D spatial position from the unique signature the pinnae place upon the acoustic pressure wave. The interference characteristics of head and pinnae shape on the transference of sound to the ear canals is a function that can be modelled and employed by an audio convolution algorithm to simulate the placement of sounds in 3-D space. In practice, speaker arrays and sensitive miniature microphones inserted into the ear canal make it possible to derive a set of HRTFs. Four other parameters, in addition to the parameters of inter-aural time delay, head shadow, pinna response, and shoulder echoes, comprise the Head-Related Transfer Function (HRTF). These include head motion, vision, intensity, and response caused by the local acoustical environment [28].

In the context of the beat interference method previously described, it would be desirable to add some measure of redundancy in order to remove or reduce the ambiguities, in short to improve the intuitiveness of the interface. 3-D spatialization could be used, for example, to filter the sound sources in order to simulate spatial movement. The pitches of the reference triad might be placed at some memorable position within the user's 3-D audio world, for example the origin of the 3-D audio space, which corresponds to the centre of the user's head. Cursor pitches would then move out and around the user in the case of error, or in the case of a correct placement of the instrument would come to rest at the origin. The direction required to zero the audio cursor would correspond to the direction required to correct the placement of the instrument.

Callipers Method

This approach concerns the requirements of taking measurements intraoperatively, such as for craniofacial reconstruction. On-the-fly positioning or measurement tasks are simplified. Instead of using hardware Vernier callipers, rulers or other measurement devices, all manner of measurements may be taken and recorded using an audio feedback system, a stylus or speech recognition system and a footswitch. With measurements, taken between two points in 3-D space, the surgeon samples an initial point upon the anatomy by placing the stylus in

the desired location, for example at nasion in Fig. 6, and then activating a footswitch or speaking some command such as “origin”.

Figure 6: Measurement between two points upon the anatomy.

This 3-D location forms the origin of a spherical gradient of sound events propagated in 3-D space. As the surgeon moves the stylus through this sound field, sound events are triggered at the passing of each concentric measurement increment. For example, each millimetre increment triggers an audible click, and each centimetre increment triggers a speech synthesiser to speak the radius in centimetres from the origin. See Fig. 7.

Figure 7: Spherical gradient of sound events.

When the desired radius has been located, its position may be recorded by use of some input device, either speech or switch, or the sound field may be turned off. For complex measurement tasks, such as required for minimising multiple skull fractures or placing and wiring multiple bone fragments, measurements may be automatically accumulated, and labelled. In this way the surgeon might simply say the word “nasion” and the feedback system would automatically propagate a sound field around that location.

Wave Terrain Synthesis

Wave-terrain synthesis is another approach that we are exploring for audio rendering of 3-D volumes, more specifically, 3-D surfaces. This technique proceeds from the principle of wave-table lookup synthesis. It is possible to extend the basic principle of wave-table scanning as implemented in a sine wave oscillator to the scanning of 3-D anatomical object surfaces for the purpose of generating an audio representation of such objects.

A typical wave-table can be plotted as a two dimensional function, **wave_table(x)**, using **x** as the index. A 2-index wave table, or wave terrain, can be plotted as a function **wave_terrain(x,y)** on a 3-D surface, for instance the surface of an anatomical object model. The **z**-point of this function corresponds to the waveform value for a given pair (**x,y**). [25]

In implementing wave terrain synthesis in an audio feedback system for surgery, the rigid body angle of the instrument is measured with respect to the surface of the anatomical object. This angle defines a normal to a surface region to be sampled from the anatomical object as a wave terrain (see Fig. 8). This surface region is scanned using a periodic scanning function. This scanning process generates a stream of waveform amplitude values that are streamed to the DSP.

Figure 8: The surgical instrument describes an angle to the anatomical surface. The wave-terrain (outlined) is sampled relative to this angle.

The signal generated depends on both the wave terrain and the scanning trajectory. The trajectory may take any number of forms, such as a straight linear trajectory, or an elliptical function. When the trajectory is a periodic function, the resulting waveform exhibits a static spectrum. This spectrum will be relatively homogeneous as the instrument passes across homogeneous surfaces, but will vary significantly upon passing across a surface in which the surface function changes abruptly. This change in spectrum can assist a user attempting to comprehend or “visualise” the surface terrain of an object.

Haptic-Audio Integration

Current approaches for haptic and acoustic simulation focus upon either haptic or audio simulation as entirely separate modelling and simulation processes. The resulting simulation systems proceed from haptic or acoustic-biased approaches to the engineering process, instead of proceeding from knowledge of the geometry and physical nature of objects. These simulation systems are unconvincing in their realism.

In an extension to the sonification of 3-D volumes, specifically the wave terrain approach, both haptic and acoustical feedback would be derived from a single general geometric and physical model. Deriving haptic and acoustic feedback from a general physical model will result in more realistic simulations, because these two sensory modalities, while perceived through different physiological apparatus, proceed, in the natural world, from the physical properties of materials. Such a direct correspondence between physical interactions with objects, and their resulting haptic and acoustic effects, will result in a greater degree of immersion for the user in simulated environments.

Material descriptions, which formerly only described geometric and optical, or possibly haptic parameters will be revised to include parameters such as 3-D physical surface texture (as opposed to 2D texture or “bump” maps), material physical properties. Graphical, haptic and acoustic subsystems will read these parameters and render them to the appropriate display hardware (e.g. heads up display, phantom, and headphones).

For surgical simulation incorporating haptic feedback this approach provides a greatly increased sense of realism. For example, when a scalpel scrapes against bone, the resulting haptic forces and acoustic waveforms will be simulated correctly, directly in proportion to the user interaction forces, as opposed to using discrete sampled sound clips or simple haptic models. Such an approach also opens up the exciting possibility for enhancing, amplifying, or otherwise processing subtle haptic and acoustic bandwidths out of the range of perception in the natural world. For example, in simulating the sound and reaction forces of a biopsy needle passing through subtly different tissue types, the resulting feedback would be scaled to fall within perceptible ranges. Rotating the needle, extending and retracting the biopsy core, and applying differential pressure would all result in noticeable changes in the simulated reaction forces on the biopsy needle, and sounds generated by the procedure. This feedback could be used as an aid to help surgeons better understand the interplay of forces between instrument and tissue, resulting in more precise cutting and placement tasks.

During operation a haptic subsystem would analyse general object physical properties and user interaction forces, and based upon these parameters return haptic force feedback data to a haptic display device (e.g. a Sensable Technologies, Inc.³ Phantom). This would include various functions for enhancing or amplifying the simulated reaction forces.

The acoustic subsystem analyses object properties and user input forces, and returns waveform data to an audio-rendering device (e.g. a Lake DSP Huron). The 3-D physical texture map (identical to that used for the haptic simulation) is sampled over time, and is implemented as a wave terrain. As the cursor scans or otherwise

³ <http://www.sensable.com/>

interacts with this wave terrain over time, a unique waveform is generated which bears the signature of the interaction forces. See Fig. 9.

Figure 9: A cursor tracks at a constant rate and force across the texture plane, generating the resulting synthetic haptic and audio waveforms.

Real-world Application

Informal usability testing [12,13] in the course of the engineering process indicate that multi-channel spatialised audio position feedback can be of assistance in basic spatial placement tasks. Much basic research needs to be done, but the potential for this new technology is evident.

We are currently in the early stages of developing a prototype tactical audio system as an extension to the Sofamor Danek Group's⁴ StealthStation system for image-guided stereotactical neurosurgery. In the near future we plan on pursuing this project in conjunction with the Saint Louis University (SLU) Medical School, Neurosurgical Division⁵, and R. Bucholtz, M.D. Through usability testing with surgical residents at SLU, and at a later date, the performance of audio-guided neurosurgical test surgeries, we plan to develop an optimally-intuitive tactical audio system design methodology, as well as a novel and powerful surgical tool.

References

1. Brewster S. Providing a Structured Method for Integrating Non-Speech Audio into Human-Computer Interfaces. Ph.D. Thesis. University of York, U.K., 1994
2. Buser P, Imbert M. Audition. (Kay R trans) Bradford Books, Inc., New York, 1992
3. Earnshaw R., Vince J., Jones H. Virtual Reality Applications. Academic Press, Inc., San Diego, 1995
4. Edwards A. Soundtrack: An Auditory Interface for Blind Users. Hum Comp Int 1989; 4(1)
5. Lake DSP Pty. Ltd. Huron Digital Audio Convolution Workstation Chassis. Lake Datasheet No. D60110-4., Sydney, 1996
6. Mon-Williams M. et al Binocular Vision in a Virtual World: Visual Deficits Following the Wearing of a Head-mounted Display. Ophth & Phys Opt 1993; 13(4):387-391
7. Patterson R. Alarm Sounds for Medical Equipment in Intensive Care Areas and Operating Theatres. Report No. AC598, Institute for Sound and Vibration Research, University of Southampton, 1981
8. Patterson R. Guidelines for Auditory Warning Systems on Civil Aircraft. Report No. 82017. Civil Aviation Authority, London, 1982
9. Plomp R. Aspects of Tone Sensation: A Psychophysical Study. Academic Press, Inc., Boston, 1976
10. Potter S., Wegner K., Karron D., et al. Teleoperation Strategies for Exploring and Utilising the Resources of Space. In Space Manufacturing 11, The Challenge of Space: Past and Future. Space Studies Institute, pp 364-369, Princeton, NJ, 1997

⁴ <http://www.sofamordanek.com/>

⁵ <http://medschool.slu.edu/departments/surgery/div/neuro/>

11. Satava R. Virtual Reality and Telepresence for Military Medicine. *Comp Bio & Med* 1995; 25(2):229-236.
12. Wegner K, Karron D. Surgical Navigation Using Audio Feedback. In Morgan KS et al (eds): *Medicine Meets Virtual Reality: Global Healthcare Grid*. IOS Press, Ohmsha, Washington, D.C., pp 450-458, 1997
13. Wegner K, Karron D. Tactical Audio for Surgical Placement Tasks. Submitted for publication to *Jour Comp Aid Surg*, 1998.
14. Wenzel E., Fisher S, Stone P., Foster S. A System for Three- Dimensional Acoustic "Visualisation" in a Virtual Environment Workstation. *Proceedings of Visualization '90*. IEEE Press., New York, NY, 1990
15. World L. The Reality of Cybersickness. *IEEE Comp Graph & App* 1995; 15(5):95.
16. Zwicker E, Fastl H *Psychoacoustics: Facts and Models*. Springer-Verlag, New York, NY, 1990
17. Edworthy J. et al. Improving Auditory Warning Design: Relationship Between Warning Sound Parameters and Perceived Urgency. *Hum Fact*, 1991, pp 205-231.
18. Haas E. A Pilot Study on the Perceived Urgency of Multitone and Frequency-modulated Warning Signals. *Proceedings of the Human Factors Society 36th Annual Meeting*, 1992, pp 248-252.
19. Glinsky A. The Theremin in the Emergence of Electronic Music. Ph.D. Thesis, New York Univ., New York, NY 1992.
20. Burdea G. *Force Feedback and Touch Feedback for Virtual Reality*. John Wiley & Sons, Inc. New York, NY 1996.
21. Strybel T., Manligas C., Perrott D. Minimum Audible Movement Angle as a Function of the Azimuth and Elevation of the Source. *Hum Fact*. 1987, 34(3):267-275.
22. Bonitz M. *Nuclear Instruments and Methods*. 1963, 22:238-252.
23. Michel H. Le 'vernier' et son inventeur Pierre Vernier d'Ornans. In *Mémoires de la Société d'émulation du Doubs* 8, 1913, pp 310-373.
24. Lake DSP Pty. Ltd. Huron DSP Board, Lake Datasheet D60108-4, Lake DSP, Sydney, 1996.
25. Roads C. *The Computer Musci Tutotial*. MIT Press, Cambridge, MA, 1996.
26. Vanderheiden G. Nonvisual Alternative Display Techniques for Output from Graphics-based Computers. *Journal of Visual Impairment and Blindness*, 1989.
27. Mulder J., Dooijes E. Spatial Audio in Graphical Applications. In Gobel M., Muller H., Urban B. Eds, *Visualization in Scientific Computing*. Springer-Verlag, New York, 1995, pp 215-229.
28. Begault D. *3-D Sound for Virtual Reality and Multimedia*. Academic Press, Boston, MA, 1994.