

Feeling for Sound: Mapping Sonic Data to Haptic Perceptions

Tom Mudd
Goldsmiths, University of London
New Cross, London
t.mudd@gold.ac.uk

ABSTRACT

This paper presents a system for exploring different dimensions of a sound through the use of haptic feedback. The Novint Falcon force feedback interface is used to scan through soundfiles as a subject moves their hand horizontally from left to right, and to relay information about volume, frequency content, envelopes, or potentially any analysable parameter back to the subject through forces acting on their hand.

General practicalities of mapping sonic elements to physical forces are considered such as the problem of representing detailed data through vague physical sensation, approaches to applying forces to the hand that do not interfering with the smooth operation of the device, and the relative merits of discreet and continuous mappings. Three approaches to generating the force vector are discussed: 1) the use of simulated detents to identify areas of an audio parameter over a certain threshold, 2) applying friction proportional to the level of the audio parameter along the axis of movement, and 3) creating forces perpendicular to the subject's hand movements.

The potential uses of such a device are also discussed, such as 'pre-feeling' as a method for selecting material to play during a live performance, an aid for visually impaired audio engineers, and as a general augmentation of standard audio editing environments.

Keywords

Haptics, force feedback, mapping, human-computer interaction

1. INTRODUCTION

The lack of physical interaction, and particularly physical resistance is a much-discussed aspect of digital music. Despite the increased sophistication and availability of haptic devices, they has been not been a significant interest in incorporating physical forces into musical systems outside of the research community. Significant work has been done within the community however - most notably at the Centre for Computer Research in Music and Acoustics (CCRMA) at Stanford - and various systems have emerged that integrate touch in some

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NIME'13, May 27-30, 2013, KAIST, Daejeon, Korea.
Copyright remains with the author(s).



Figure 1. Novint Falcon haptic device, shown with the ball grip on the left-hand side

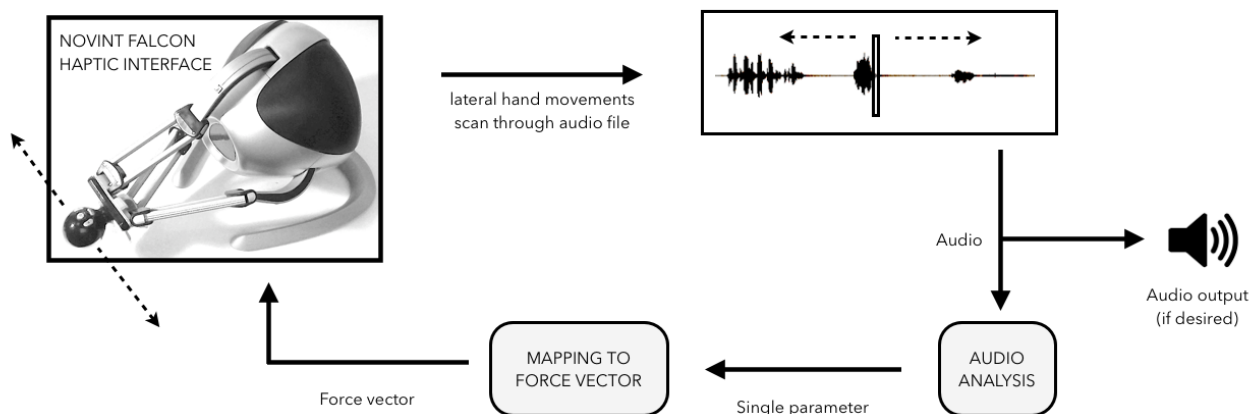
way.

These have tended to approach physicality in one of two different ways: offering resistance in a manner analogous to the resistance found in many acoustic instruments [1] [4], or providing a way of relaying some kind of information back to a subject without requiring visual or audible cues [7].

The project presented in this paper falls into the latter category. It revolves around the use of the Novint Falcon force feedback device (see figure 1) together with the open-source, cross-platform driver 'libnifalcon' [10], to allow a subject to scan through a soundfile for a particular sonic element, and to feel the results physically. The Falcon controller was developed for use in video games and as such represents an affordable - 250 USD at the time of writing - and relatively fast and accurate (forces can be updated every millisecond) haptic device. As with several other haptic devices, it allows the subject to move a handheld grip in three dimensions and to feel a programmable, three dimensional force vector applied to this grip. This allows for both tactile (cutaneous) sensation and kinesthetic feedback through a single interface point.

2. PROJECT OVERVIEW

The ability to find things quickly through the use of touch is a very instinctive process; almost everyone is used to feeling for things and performing actions without the aid of sight or any other sensory feedback other than touch. In this project, sound



is manifested physically as a virtual surface defined by the changing levels of a specified sonic parameter.

As the subject scans through the soundfile by moving the ball grip (see figure 1) horizontally, forces are fed back to the hand to provide information about the audio at any given point (figure 2). The force is a function of a single audio parameter: volume, noisiness or the volume of a particular frequency.

The strength of the force applied also relates to the vertical position of the subject's hand, with the intensity increasing in strength as their hand is lowered, or decreasing and disappearing as the hand is raised. This allows the subject to easily remove all forces, or work at a level that they are comfortable with (as recommended by Berdahl et al [2]).

The subject can use the grip to select a part of the soundfile and spread the selection across the entire horizontal axis, effectively zooming in on a particular area to allow for more detailed exploration of a smaller segment of sound. The potential applications of a subject being essentially able to feel their way around a sound file are discussed further in section 4.

3. WORKING WITH FORCES

3.1. Approaches to Haptic Waveforms

In this project, haptic feedback is deployed towards a very specific end: the fast and accurate transmission of information about the location of certain perceptual features of sound in a manner that does not significantly disrupt the subject's interaction with the device. Creating mappings that actively help rather than hinder when performing a specific audio task is not necessarily a straightforward procedure. Compared to the body of literature investigating the possible relationships between visual images and haptic sensation there has been fewer quantitative studies convincingly demonstrating the benefits of haptic-enhanced approaches to working with sound ([12], [3] and others at CCRMA are useful exceptions). Berdahl et al conclude that there is no ideal form of haptic feedback [2], and that the most appropriate mappings will vary depending on the nature of the task.

The mapping of audio information to physical information can be implemented in an infinite number of ways. Two specific approaches were implemented for this project, both of which revolve around relaying information about the strength and presence of a certain parameter through forces acting along a single axis, parallel to the movements of the hand. The forces generally act in the opposite direction to the subject's movements, obstructing them in different ways. A third approach that utilises forces normal to the hand movements was

investigated and is also discussed below, but proved problematic and has not been incorporated into the project for the time being.

3.1.1. Discrete Force Mapping

Berdahl et al describe the successful implementation of haptic detents to assist the accurate selection of diatonic pitches along a continuous axis [2]. A detent acts in a similar manner to a small dip in a flat surface in that the hand experiences resistive lateral forces when trying to move out of the dip. In this project, detents are applied to any areas of the audio over a certain threshold, thereby creating a discrete quantity of detents across the waveform.

This requires a certain amount of threshold calibration, as too many detents in a small area of the sound can be very confusing, and too few detents cannot communicate much about the nature of sound. This issue could be alleviated by collecting peaks together across small time intervals and integrating them into a single detent which could then be expanded using the zoom feature described in section 2.

3.1.2. Continuous Force Mapping

The second approach to force mapping involves attempting to map the strength of the audio parameter to the amount of resistance applied on the hand. The greater the parameter, the more friction is experienced, so that the hand can be freely moved through areas with little activity, but experiences resistance in denser areas. With volume as the selected audio parameter for example, silent areas would offer no resistance, but sustained loud areas would be difficult to move through (depending on the vertical position of the hand which varies the overall intensity of the forces). A short peak is experienced as a brief bump. Similar results would apply when scanning for a particular frequency: a single percussive note playing the specific frequency is experienced as a bump, a sustained drone at that frequency will be experienced as drag, and areas that do not include the frequency can be moved through with no resistance. This approach is based on Minsky and Lederman's description of how friction can be simulated effectively through alterations to lateral forces (in the same dimension as the hand movements in this case) [11]. The expectation was that the user would therefore be able to differentiate between different quantities of friction rather than just experiencing a discrete, on/off force.

In practice however, although some nuance could be felt, it was difficult to differentiate the subtleties of the changing

parameter, and the force experienced at a particular point in an audio file varied considerably depending on the speed at which the area was traversed by the subject's hand. This approach worked well when feeling for small areas of activity separated by gaps, and the use of a gate with a variable threshold made navigation more comfortable and the data more intelligible to the hand. Mapping in this manner loses a good deal of the subtlety in the data however and is effectively a variant on the discrete force mapping described above.

Further work could attempt to distinguish different degrees of a parameter through alterations to the nature of the friction experienced. Minsky and Ledermen describe the accurate perception of grated and grid textured surfaces through alterations in lateral force. This could also be applied to the development of multi-parametric approaches with different audio features being connected to different surface types.

3.1.3. Forces Normal to the Direction of Movement

Although this approach has not been fully implemented it may be useful to detail the preliminary investigations. As subject moves their hand horizontally, forces act vertically, pushing the hand upwards in proportion to the amplitude of the audio parameter. This felt highly disruptive to the smooth operation of the device however, and made it difficult to focus on a particular area of the sound, as the hand would be constantly pushed around unexpectedly. Such a disruptive engagement with the interface does not seem appropriate for a project like this where the primary concern is to feed back information, but it may be beneficial in other contexts, such as new instrument development (see [14] for example on the benefits of resistance in instrument design).

3.2. Simplifying Audio Data for Haptic Representation

In their research on mapping visual data into the haptic realm, Fritz et al [6] highlight the importance of simplifying the data. This must likewise be engaged with in projects wishing to represent audio data as physical forces, as the full content of the audio would be very difficult to translate into discernible physical sensations. It is therefore necessary to extract meaningful elements from the audio data and present them in a way that the hand can usefully interpret.

The decision to present only a single audio parameter at a time provides an initial simplification of the audio data. Even with just a single dimension however, it is difficult to present the same level of detail in the haptic domain (as noted in the section 3.1.2 on continuous force mapping for example). As mentioned above, one approach to this that proved useful was to remove values below a certain threshold. In scanning for a particular frequency for example, the frequency would be present at a low level in many noisier sounds making it difficult to distinguish a clear perceivable occurrence of that frequency from these less prominent occurrences. The use of a threshold gate simplifies the data into a set of discrete areas of activity which are more easily perceivable by the hand than continuous fluctuations.

The approaches described in this section all represent attempts to present audio information through the abstraction of an artificial surface. This is in contrast to more literal attempts to represent sonic data through touch, such as [3], which attempts to recreate the sensation of feeling a sound as one would experience it when in contact with the sound generating object.

4. POTENTIAL APPLICATIONS

4.1. Pre-Listen / Pre-feel

The ability to scan through a sound file without hearing it or seeing it could be of use in certain models of live performance. When sampling a live instrument for example, there may often be segments of sound amongst stretches of silence. For a performer who is interested in taking this sound and utilising it during the performance in some way, it can be important to know where the sections of activity lie in the soundfile, and to be able to quickly move from one to another. Three methods are generally employed: onscreen displays representing the waveform, playing the sound (with or without alteration) and working forwards or backwards from there based on the audio feedback, and using a separate headphone channel that allows the performer to audition a sound privately prior to playing it through applying it in the performance. The potential drawbacks associated with each of these - reliance on a visual display, possible unwanted sounds, and having to focus on sounds other than those in the performance - are removed by receiving information about the sound through silent haptic sensations. The information is passed back to the performer through the same device they are using to control the sound in the first place.

4.2. A Tool for the Visually Impaired

Being able to feel the key features of a sound file with the hand and make selections by feeling for events and silences may provide improved access to audio editing tools for people with visual impairments.

Although there have been a variety of studies on relating sonic elements to physical sensations [3] [5] [12], there has been significantly more work done in translating visual data into physical information [6] [15] and how this may benefit people with visual impairments. Visual data has the advantage that it is inherently spatial, whereas audio - despite being physical in nature - is time-based. This project is not attempting to present sound in its entirety as a haptic experience, but merely to present some information about that sound in a physical way that allows it to be navigable by touch alone. It therefore shares some of the concerns of the visual mapping approaches: time is stretched out along a horizontal axis as is seen in a waveform display, creating a spatial representation of some aspect of the sound (volume being the most obvious, as is used in traditional waveform displays, but any audible parameter that can be graphed over time is potentially presentable). Further development in this direction could therefore draw on literature that deals with the issue of converting visual data to haptic sensation.

The system was recently presented at a participatory design workshop for visually impaired musicians and audio production specialists as part of the Design Patterns for Inclusive Collaboration project at Queen Mary, University of London. This ongoing project will be an important source of feedback and suggestions as to how effective particular forms of haptic feedback are for navigating sound files, and for suggestions as to how improvements might be made.

4.3. Quick Information for Audio Editing

The application of haptic feedback to more conventional audio editing applications or digital audio workstation environments could provide certain benefits. Hecht et al [8] highlight the faster response times firstly from haptic impulses in relation to visual or audio cues, and secondly from the combination of different modes of feedback compared to any individual mode - the most effective being the combination of visual, audible and haptic sensations. Whether this is directly applicable to audio

editing environments remains to be seen. Chu's *TouchSound* project [5] provides some movement in this direction through the use of haptic knobs and sliders in the context of a digital audio workstation.

5. SUMMARY

This paper introduces a project investigating haptic mapping strategies to assist with the navigation and selection of areas of a sound file. The ability to receive feedback about any measurable sonic parameter without the need for visual or audible feedback suggests possibilities for 'pre-feeling' when working with sound in live performance, aiding visually impaired people in using audio tools, and enhanced audio editing through multi-modal feedback. Continuous and discrete mappings were used to translate the audio parameters into physical forces, and although at this stage no rigorous testing has been carried out, more discrete mappings appear to provide a more comprehensible - albeit reduced - representation of the audio information.

6. FUTURE DEVELOPMENTS

Development and implementation of a more advanced method for conveying continuously changing audio parameters through the haptic interface would be greatly beneficial to a project of this nature, as currently it is difficult to ascertain much more than the presence or absence of a parameter above a certain threshold. This does not necessarily have to be approached in a linear fashion whereby force increases with the level of the parameter; more nuanced mappings may be possible through more abstract approaches. As mentioned above, this could be based around the work done on approaches to representing visual data through haptic sensation. The rapidly developing field of haptic interfaces for touchscreen devices may also suggest useful ways of maximising the amount of information that can be gleaned through haptic devices [9] [13].

7. REFERENCES

1. Berdahl, E., Verplank, B., Smith III, J. O. Niemeyer, G. A Physically-Intuitive Haptic Drumstick. In *Proceedings of the 2008 International Computer Music Conference*, Copenhagen, Denmark, 2007.
2. Berdahl, E., Niemeyer, G., Smith III, J. O. Using Haptics to Assist Performers in Making Gestures to a Musical Instrument, In *Proceedings of the Ninth International Conference on New Interfaces for Musical Expression*, Pittsburgh, PA, June 4-6, 2009.
3. Birnbaum, D. M. and Wanderley, M. M. A systematic approach to musical vibrotactile feedback. In *Proceedings of the 2007 International Computer Music Conference*, Copenhagen, Denmark, 2007.
4. Cadoz, C., Lisowski, L. and Florens, J. A Modular Feedback Keyboard Design, *Computer Music Journal*, Vol. 14, No. 2, 47-51, MIT press, 1990.
5. Chu, L. L. User Performance and Haptic Design Issues for a Force- Feedback Sound Editing Interface, in in *Proc. of the Conf. on Human Factors in Computing Systems, 2002*, Boston, MA, USA, 2002.
6. Fritz, F., Way, F., Barner, F. Haptic Representation of Scientific Data for Visually Impaired or Blind Persons. In *ICCHP'12 Proceedings of the 13th international conference on Computers Helping People with Special Needs*, Volume Part II, 517-520, 2012.
7. Hayes, L., and Michalakos, C. Imposing a Networked Vibrotactile Communication System for

Improvisational Suggestion. *Organised Sound*, 17, pp 36-44, 2012.

8. Hecht, D., Halevy, G. and Reiner, M., Multimodal Virtual Environments: Response Times, Attention, and Presence, in *Presence: Teleoperators and Virtual Environments*, Vol. 15, No. 5, Pages 515-523, 2006.
9. Lim, M. J., Shim, H., Lee, J., Kyung, K. Perceived intensity of click sensation for small touchscreen devices, *ITS '12 Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces*, Cambridge/Boston, MA, USA, Nov, 2012.
10. Machulis, K. *libnifalcon: Open source driver for the Novint Falcon*, <http://libnifalcon.nonpolynomial.com>, accessed 12th December, 2012.
11. Minsky, M. and Lederman, S.J. Simulated haptic textures: Roughness. In *Proceedings of the ASME International Mechanical Engineering Congress: Dynamic Systems and Control Division, Vol. 2* (Haptic Interfaces for Virtual Environments and Teleoperator Systems), DSC-Vol. 58, 421-42, 1996.
12. O'Modhrain, S. *Playing by Feel: Incorporating Haptic Feedback into Computer-Based Musical Instruments*, Ph.D. Thesis, Stanford University, 2000.
13. Pasquero, J., Luk, J., Little, S., MacLean, K. Perceptual Analysis of Haptic Icons: an Investigation into the Validity of Cluster Sorted MDS, *HAPTICS '06 Proceedings of the Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE Computer Society Washington, DC, USA, 2006.
14. Ryan, J. Effort and expression. *Proceedings of the 1992 International Computer Music Conference*. San Francisco CA, USA: International Computer Music Association, 1992.
15. Shi, Y. and Pai, D. K. Haptic Display of Visual Images, In *Proceedings of the IEEE Virtual Reality Ann. International Symposium (VRAIS 97)*, IEEE CS Press, 1997.

8. APPENDIX

The software itself can be downloaded from the following location: tommudd.co.uk/haptic. Note that it is only compatible with the Novint Falcon device together with the 'libnifalcon' driver [10]