Physical Interface Design for Digital Musical Instruments

Mark T. Marshall

Music Technology Area,
Department of Music Research,
Schulich School of Music,
McGill University,
Montreal, QC, Canada.

April 2008

A thesis submitted to McGill University in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

©2008 Mark T. Marshall

Abstract

This thesis deals with the study of performer-instrument interaction during the performance of novel digital musical instruments (DMIs).

Unlike acoustic instruments, digital musical instruments have no coupling between the sound generation system and the physical interface with which the performer interacts. As a result of this, such instruments also lack the direct physical feedback to the performer which is present in an acoustic instrument. In fact in contrast to acoustic musical instruments, haptic and vibrotactile feedback is generally not present in a DMI contributing to a poor *feel* for the instrument. The main goal of this thesis is to propose ways to improve the overall *feel* of digital musical instruments through the study and design of its physical interface: the instrument body, sensors and feedback actuators.

It includes a detailed study of the existing theory and practice of the design on physical interfaces for digital musical instruments, including a survey of 266 existing DMIs presented since the inception of the NIME conference. From this, a number of differences become apparent between the existing theory and practice, particularly in the areas of sensors and feedback.

The research in this thesis then addresses these differences. It includes a series of experiments on the optimal choice of sensors for a digital musical instrument. This is followed by research into the provision of vibrotactile feedback in a digital musical instrument, including the choice of actuator, modification of actuator frequency response, and the effects of response modification on human vibrotactile frequency discrimination.

Following this, a number of new digital musical instruments are presented, which were created during the course of this work. This includes an instrument designed specifically to follow the results of research in this thesis and also instruments designed as part of larger collaborative projects involving engineers, composers and performers.

From the results obtained in this work, it is shown that a careful design of both the sensor and actuator aspects of the physical interface of a DMI can lead to an instrument which is more engaging and entertaining to play, offering an improved *feel* over that which is present in many digital musical instruments.

Abrégé

Cette thèse porte sur l'étude de l'interaction ayant lieu, en situation de jeu, entre un(e) instrumentiste et un instrument musical numérique (IMN).

A l'inverse des instruments acoustiques traditionnels, il n'existe aucun couplage entre le dispositif de production du son et l'interface sur laquelle agit l'instrumentiste dans le cas des IMN. L'une des implications de cette observation est que ces instruments ne procurent pas la rétroaction tactile normalement présente dans les instruments de musique traditionels. Par conséquent, les IMN sont souvent perçus par leurs interprètes comme manquant d'*âme*, de *personnalité*.

Le but de ce travail de thèse est d'avancer quelques solutions permettant d'insuffler un peu plus *âme* à un instrument musical numérique. Le point focal de la recherche étant l'étude et la conception de l'interface physique (corps de l'instrument, capteurs et dispositifs de rétroaction utilisés) d'un tel instrument.

Ce mémoire présente, en premier lieu, une étude détaillée de la théorie et de la pratique actuelles dans le domaine de la conception d'interfaces physiques pour les IMN. L'inventaire des 266 instruments recensés depuis la création de la conférence NIME constitue l'un des points majeurs de cette partie du travail. En effet, ce tour d'horizon permet de faire ressortir les incohérences entre théorie et pratique. Ces différences sont particulièrement frappantes en ce qui concerne les capteurs et les dispositifs de rétroaction.

Le travail de recherche de cette thèse a donc pour objectif de mieux comprendre comment réduire ces incohérences. Des expériences portant sur le choix optimal des capteurs à utiliser dans un IMN ont donc été menées. Différents dispositifs de rétroaction vibrotactile ont aussi été étudiés en regardant d'abord quels actuateurs utiliser, et en évaluant les effets de la modification de leur réponse en fréquence sur la discrimination fréquentielle de stimuli vibrotactiles chez des sujets humains.

Des exemples d'applications pratiques de ces recherches sont ensuite détaillés. En effet, plusieurs IMN ont été construits lors de cette thèse : des dispositifs conçus dans le cadre des expérienes pré-citées ainsi que d'autres instruments s'inscrivant dans le cadre de projets collectifs regroupant des ingénieurs, des compositeurs et des instrumentistes.

A l'issue de ce travail, il apparaît clairement qu'une attention particulière portée au choix des capteurs et des actuateurs de rétroaction utilisés lors de la conception de l'interface peut améliorer de façon considérable la perception que les interprètes ont d'un instrument de musique numérique. Effectivement, les musicien(ne)s ayant joué des instruments conçus lors de cette thèse ont généralement trouvé l'expérience ludique et agéable, pouvant mieux percevoir la *personnalité* des instruments.

Acknowledgements

There are a number of people and organisations without whom this thesis might never have happened. First and foremost I would like to thank my supervisor Marcelo Wanderley for bringing me to McGill and for invaluable guidance over the course of this work. Thanks also to everyone in the Music Technology Area at McGill, and those in the IDMIL in particular, for their help, encouragement and general support.

I would like to especially thank those people that I worked with in the course of the research projects under which much of the work in this thesis took place: The McGill Digital Orchestra project and the gesture-controlled spatialization project. In no particular order therefor, thanks to Sean Ferguson, Stephen McAdams, Joseph Malloch, Stephen Sinclair, Georgios Marentakis, Xenia Pestova, D. Andrew Stewart, Nils Peters, Heather Hindman, Chloé Dominguez, Jonas Braasch, David Birnbaum, Rodolphe Koelhy, Fernando Rocha and Erika Donald.

I would also like to thank Max Hartshorn for his work on the sensor experiments in this thesis, particularly for his work on the design and running of the final two experiments described in Chapter 4. The Matlab code used to perform the analysis of the modulation signals for those experiments was provided by Vincent Verfaille, who also, along with Bertrand Scherrer, provided invaluable insight on the analysis of low frequency modulations. Thanks also to Dan Levitin for invaluable feedback and discussions regarding those experiments.

The work performed in this thesis was financially supported by a number of organisations. Thanks to the EU Cost Action 287 ConGAS for funding the initial short-term scientific mission which led to this research. Thanks also to the Fonds de recherche sur la société et la culture (FQRSC) of the Quebec government and the Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT) for their support of the work performed as part of the Digital Orchestra project, and to the Canada Council for the Arts, the Natural Science and Engineering Research Council of Canada (NSERC) and CIRMMT for their support of the gesture-controlled spatialization project. Thanks also to both CIRMMT and the McGill University Alma Mater fund for travel support which allowed me to attend a number of international conferences during the course of this work.

Finally, a special thanks to my parents, Mary and Jim, and to Sinead for their support of me throughout this work.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

DMI	Digital Musical Instrument
DOF	Degrees of Freedom
FSR	Force-Sensing Resistor
HCI	Human Computer Interaction
HSD	Honestly Significant Difference
JND	Just Noticeable Difference
LED	Light-Emitting Diode
MIDI	Musical Instrument Digital Interface
NIME	New Interfaces for Musical Expression
OSC	Open Sound Control
USB	Universal Serial Bus

Chapter 1

Introduction

The physical design of traditional musical instruments is a direct result of the ways in which these instruments generate sound. That is, the use of membranes, strings and air columns to create sound informs the physical design of the instrument themselves. To generate the required pitches and timbres the instrument must be built (and played) in a specific way. For digital musical instruments however, these restrictions do not apply.

The use of computer technology allows us to electronically generate any sound and to control parameters of this sound in any way, without any physical restrictions. This allows us freedom in the physical design of new digital musical instruments that the designers of acoustic instruments do not have. From this arises a new question: how best to design the physical interface of a new digital musical instrument?

This thesis deals with exactly this question, examining several issues around the design of physical interfaces for digital musical instruments. Specifically I deal with improving the performer-instrument interaction through the careful design of both sensor and feedback systems and the integration of these systems into complete digital musical instruments.

1.1 Acoustic and Digital Musical Instruments

There are a number of fundamental differences between acoustic and digital musical instruments. Perhaps the most fundamental of these differences arises from the separation of the control system from the sound synthesis synthesis in a digital musical instrument (DMI). In musical instruments the control systems are those portions of the instrument which the performer manipulates to create sound. For acoustic instruments these are integrated with the sound creation systems. The performer creates sound on an acoustic instrument by acting directly on the sound production mechanisms. For a stringed instrument the performer changes the pitch by manipulating the length of the vibrating portion of the string. They produce sound output by adding energy to the vibrating string itself, whether through plucking, striking or bowing. Similar processes are used for other acoustic instruments, including wind and percussion instruments.

For a DMI on the other hand, this situation is different. The performer acts on the sensors which are present in the DMI. These sensors are used to translate physical parameters of the performer's gestures into digital values which are then used to manipulate parameters in the computer-based sound synthesis system. The performer acts on the sensors, the computer reads the sensor values, manipulates synthesis parameters, synthesizes the sound and outputs the sound through a separate speaker system, often located away from the performer.

The production of the sound from loudspeakers that are separate from the instrument results in further differences between a DMI and an acoustic instrument. One resulting difference is that the sound of the instrument does not seem to come from the instrument itself. Cook (2004) cites this as one of the issues which results in a loss of intimacy when playing a DMI compared to playing an acoustic instrument. In terms of musical performance, intimacy can be defined as (Moore, 1988):

> "... the match between the variety of musically desirable sounds produced and the psycho-physiological capabilities of a practiced performer."

While a similar loss of intimacy also exists for some electric instruments (such as the electric guitar or violin) which produce their sound output from separate speakers and amplifiers, the effect is even greater for a DMI. This is due to the fact that a DMI is often entirely silent at the interface, not even generating the quiet sounds that come from an electric guitar or violin. Even for such electric instruments Trueman (1999) notes a sense of detachment felt by performers when playing an electric violin:

> "This sense of detachment can be at once both empowering and distressing... there is a striking loss of intimacy, even with a small amplifier placed nearby."

This loss of intimacy is further exacerbated in digital musical instruments by the lack of direct feedback from the instrument to the performer. Cook (2004) describes a digital musical instrument as a "feed-forward system", where the entire flow of information is from the performer through the instrument, with no feedback from the instrument to the performer. As the sound generator is separated from the controller, the performer receives none on the intrinsic vibrations which are present in an acoustic instrument. The sound production in an acoustic instrument causes vibrations within the instrument body itself. These vibrations provide useful information to the performer about the state of the instrument. In fact, while beginner musicians generally make extensive use of visual feedback when playing their instruments, more experienced performers rely to a much greater extent on tactile and kinaesthetic feedback from the instrument (Keele, 1973).

In many cases, digital musical instruments also lack haptic feedback. Acoustic instruments offer resistances to the performer which must be overcome in playing the instrument. Musical instrument strings are held at tension and require a certain force in order to bend, pluck or bow them. Pianos have keys which require a certain force to actuate them to produce sound. As Gillespie (2001) states:

> "While audition carries meaning regarding the acoustical behaviour of an instrument, haptics carries meaning regarding the mechanical behavior."

In a DMI, often no such resistances are present. Many sensors used in DMIs are designed to be actuated with minimal effort by the user, as they are meant for commercial or industrial control systems which should not require effort to activate. Still other digital musical instruments use sensors which require no physical contact at all, for instance measuring the distance between the performer's hand and the sensor.

In relation to the design of new instruments, the separation of control and sound creation in a digital musical instrument creates some substantial differences when compared to acoustic instruments. In an acoustic instrument the sound creation method chosen influences the physical design of the instrument. Having decided to produce sound in a specific way (for instance through the use of strings) certain constraints are then placed on the designer by this choice. Lower pitches generally require longer strings, polyphony requires multiple strings, increased sound output levels require a resonator of some sort and so on. All of these criteria effect the possible physical forms that the instrument can take. This is not true for digital musical instruments. By separating the control surface from the sound creation, the instrument can take any physical form and be interacted with in any way the designer wishes, while still producing the desired sound. Instruments can be played like a stringed instrument but sound like a woodwind, brass or percussion instrument.

The use of computer sound synthesis systems also allows for freedom in the sound produced by the instrument. The instrument can sound like an existing acoustic instrument, or like a blending of two different acoustic instruments. It can be used to create sounds which are not possible (or just not feasible) with acoustic sound generators. A DMI can even change sound, so that the performer may choose a different sound for different performances or contexts.

These differences between acoustic and digital musical instruments can result in advantages and disadvantages for each type of instrument. Magnusson and Mendieta (2007) performed a survey of musicians of both acoustic and digital musical instruments and compiled a resulting list of frequent positive and negative comments for each type of instrument. While mostly concerned with software-based digital musical instruments, many of the comments are also valid when applied to hardware-based digital musical instruments. Of particular interest for this work were the discussions of the presence/absence of haptic and tactile feedback, latency in digital musical instruments (a delay between performer action and sound production) and the ability to master an acoustic instrument specifically due to its limitations.

This last issue is a particularly interesting one. Digital musical instruments allow for the performer to change not only the sound being produced by the instrument, but also the relationship between the performers gestures and the parameters of the sound synthesis (known as the mapping). This gives digital musical instruments a huge breadth of possibilities for performance (which is cited as one of the advantages of DMIs), but can potentially limit the depth to which a performer can learn the instrument. Acoustic instruments on the other hand have limitations imposed by their sound creation method. This results in less of a breadth of possibilities, but allows the performer to learn the instrument to a much greater depth.

Given these differences and the perceived advantages and disadvantages of digital musical instruments, the focus of the work presented here is on the design of the physical interfaces for DMIs. In particular I focus on providing a closer (more intimate) performer-instrument interaction, such as that presented by many acoustic instruments. The next section discusses the aims of this in greater detail.

1.2 Aims of this Research

When a performer plays an instrument, there is a flow of information both from the performer to the instrument and vice-versa. The performer's gestures communicate information to the instrument and the reaction of the instrument to these gestures (both in terms of sound produced and other physical responses) communicate information to the performer. This communication takes place both in traditional acoustic instruments and in digital musical instruments.

For a digital musical instrument, the flow of information from performer to instrument is accomplished through the use of sensors. Sensors translate aspects of the performer's gestures to electrical signals which can be digitized by the computer and used to control aspects of the instruments sound synthesiser. Communication from the instrument to the performer is accomplished both through the sound produced by the synthesiser and through the use of visual, haptic and tactile feedback.

I propose then that the performer-instrument interaction can be improved through an investigation of these flows of information. This thesis will investigate several aspects of these flows. In particular, the following aims and objectives have been formulated for this thesis:

- investigate existing theory and practice in the design of digital musical instruments
- examine the use of sensors and feedback in existing digital musical instruments
- perform experiments to determine the suitability of sensors for specific tasks in an interface
- develop methods and apparatus for the production of vibrotactile feedback in an interface

- design and develop a number of digital musical instruments to test theories that will be developed on sensors and vibrotactile feedback

- analyse the effectiveness of the design of these instruments, in conjunction with composers and performers

While there is a definite and important relationship between the haptic and auditory feedback channels within a digital musical instrument, investigation of the auditory feedback channel is outside the scope of this thesis. This effect has been taken into account in this work through keeping the gesture-sound mappings constant while varying the other feedback mappings. This allows for isolating the effects of the tactile feedback from those of the auditory feedback from the instrument.

The remainder of this section will address each of these aims in more detail, discussing the reasoning behind each and the methods used to achieve them.

Existing Theory in the Design of Digital Musical Instruments

The design of digital musical instruments is a varied field, which has seen research from a number of disciplines. There has been research into areas such as the use of sensors, the provision of tactile and haptic feedback and the interaction between the performer and the instrument. This research must be taken into account when dealing with the design of new digital musical instruments. Therefore, this work includes a review of existing literature on the design of digital musical instruments.

The Use of Sensors and Feedback in Existing Digital Musical Instruments

The physical interaction between the performer and a digital musical instrument is accomplished through the use of sensors to sense performer gestures and actuators to provide feedback to the performer. As research has taken place into the use of sensors and actuators in DMIs, it is interesting to see how this is reflected in the design of new digital musical instruments. To enable this, I performed a detailed survey on the design of 266 different digital musical instruments presented at the 8 annual conferences on New Interfaces for Musical Expression (NIME) since its inception as a workshop at the ACM CHI conference in 2001.

Determining the Suitability of Sensors for Specific Tasks

Many digital musical instruments are designed without any empirical examination of the suitability of specific sensors for the task required in the instrument. The opportunity therefore exists to examine the suitability of a variety of sensors for specific musical tasks. To achieve this, a series of experiments were performed examining the suitability of sensors for specific musical tasks. These experiments made use of both subjective judgements such as user preference and ease of use ratings and objective measurements such as the accuracy and precision of task performance.

Producing Vibrotactile Feedback

A main aim of this thesis is to examine ways of providing a closer performer-instrument interaction for digital musical instruments. Given that digital musical instruments, unlike traditional instruments, do not generally provide vibrotactile feedback to the performer (due to the removal of the sound source from the instrument), one possible way of improving performer-instrument interaction is the addition of vibrotactile feedback to these instruments. To enable this, I examined a number of devices and methods for producing vibrations, comparing them across a number of different criteria.

Development of digital musical instruments

To properly test the results of the experiments described in this work it was necessary to develop a number of digital musical instruments. These instruments were designed to follow a number of the results of my experiments and provide a test bed for evaluating the research performed for this thesis.

Analysis of Developed Instruments

As part of two larger collaborative projects, a number of instruments developed for this research were used by composers and performers in the production of new pieces of music. This collaboration provided an opportunity to further test the soundness of the design guidelines and technologies developed here within the context of live musical performance.

1.3 Originality and Importance

While many new digital musical instruments are being developed each year, little effort is being made to develop the physical interfaces based on guidelines from quantitative analysis of experimental data. This thesis presents a systematic approach which, coupled with an examination of the physical interface as a whole, provides an important reference for designers of digital musical instruments. The resulting new digital musical instruments have been used in musical performances as part of long-term collaborative projects involving researchers, composers and performers. This has resulted in hundreds of hours of hands-on work with the instruments and offered unique insight into the effectiveness of the approach developed for this work. These performances also provide invaluable feedback on the results of this research within the context for which the instruments are designed, that of live musical performance.

1.4 Layout of this Document

This dissertation is organised into 7 main chapters, as follows:

This chapter offered an introduction to the dissertation topic, as well as a detailed description of the aims of the research, methods used and the importance of the work described here.

Chapter 2 gives a detailed review of available work on the design of the physical interface for digital musical instruments. It presents a number of models for digital musical instruments taken from the existing literature and develops a general model which incorporates aspects of each of them. It also provides a review of the existing research on the components of the physical interface: the instrument body, the sensors and the feedback systems.

Chapter 3 presents a detailed review of the design of the physical interface of existing digital musical instruments. This is accomplished through a survey of 266 instruments presented at the NIME conferences and an examination of the use of sensors and feedback in these instruments. The results of this survey are then discussed in light of the research into these areas presented in Chapter 2.

Chapter 4 discusses a series of experiments to determine the suitability of specific classes of sensors for specific musical tasks. It provides a classification of both sensors and tasks and examines mappings between them in terms of user preference, ease of use, accuracy and precision.

Chapter 5 deals with vibrotactile feedback as it applies to digital musical instruments. It discusses devices to produce vibration and compares them across a number of parameters including frequency response, input signal requirements, availability and cost. A method of measuring device frequency response is provided, along with details of compensating for the frequency response of both the devices themselves and the sensitivity of human skin, to allow for a equal magnitude vibration spectrum to be produced. Finally, an experiment to examine the effects of such compensation on human vibrotactile frequency discrimination is also described.

Chapter 6 describes the Viblotar, an instrument built to allow an evaluation of the results of the research performed in this thesis within the context of musical performance. The sensors and feedback systems of the Viblotar are described in detail. An experiment is performed to examine the effects of the sensors and feedback on performer ratings of the instrument across a number of criteria.

Chapter 7 details a number of digital musical instruments developed as part of large collaborative projects. These instruments offer the further validation of the work on sensors described in earlier chapters of this thesis, as well as providing an additional case study on the use of vibrotactile feedback in DMIs.This chapter also contains a discussion of a number of important issues regarding the design of digital musical instruments which arose as part of these projects.

Chapter 8 provides a summary of this dissertation, presents some general conclusions drawn from the work and goes on to discuss some areas for further research within this topic.

Chapter 2

Instrument Design

In recent years, a community of research has grown around the creation of new digital musical instruments (DMIs), which are instruments consisting of a physical interface (sometimes called a gestural controller) and computer-based sound and feedback synthesis system (Birnbaum, 2007; Wanderley and Depalle, 2004; Miranda and Wanderley, 2006). Such instruments allow for the possibility of controlling a wide range of sounds and for the development of new compositional and performance practices. Also, as their sound synthesis system is separable from the physical interface, these instruments offer a large number of design possibilities that are unavailable to designers of acoustic instruments.

The design of traditional acoustic musical instruments is in many ways dictated by the physics of the method of sound production used in the instrument. That is, certain constraints are placed on the designer of the instrument which influence the physical form of the instrument. For example, to produce a certain frequency range with a vibrating string, we must use a string of specific length and thickness, held at a specific tension[1]. This requirement limits the design possibilities of the instrument.

For DMIs, on the other hand, such limitations do not exist. The sound production is performed by a computer-based sound synthesis system, which can be used to produce any sound. The parameters of this sound (for example its frequency) are available as controls which can be mapped to any input the designer might wish. In a software-based digital musical instrument these parameters might be controlled with the keyboard and mouse, through on-screen sliders and knobs, or with a generic controller such as a graphics tablet or joystick. For a physical DMI, these parameters can be controlled using a sensor (or a combination of sensors) which form part of the physical interface of the instrument.

This de-coupling of the sound creation mechanism from the physical interface in a DMI also affects the feedback from the instrument to the performer. In an acoustic instrument the instrument body is in direct contract with the performer's body and so vibrations caused by the sound production mechanism can travel from the instrument to the performer. These vibrations, together with the physical resistances offered by strings, keys, valves etc. form an important part of the *feel* of the instrument, something which is learned early in training (Chafe, 1993).

In a DMI however, the sound comes from loudspeakers which are located away from the performer and so this vibration transmission in lost. Along with this, most sensors are designed to offer minimum physical resistance to the user, so that they require little effort to manipulate. This can result in a DMI having the *feel* not of a traditional instrument, but more of a computer input device.

---

[1] More correctly, for any given base pitch, we must choose one from a range of string length, thickness and tension combinations

This chapter deals with issues relating to the design of digital musical instruments and pays particular attention to the design of the physical interface of such instruments. It begins with a short historical overview of a number of important digital (and non-digital) musical instruments which have had a major influence on the development of the field of DMI design. This is then followed by the presentation of a model for a digital musical instrument, which details the components of the physical interface, the synthesis systems and the interactions between them. This is followed by a review of existing work on the components of the physical interface of digital musical instruments: the instrument body, the sensors and the feedback systems.

2.1 Precursors of DMI Design

While much of the work presented in this thesis, and the survey in Chapter 3 in particular, concentrates on the digital musical instruments presented at the New Interfaces for Musical Expression (NIME) conferences and workshop, there have been a number of important new musical instruments developed prior to this conference which have had a major influence on the field of DMI design. This section provides a brief overview of the a number of important instruments which have led to the development of the field of DMI design.

One of the most important electronic musical instruments developed is the *Theremin*. Developed in 1920 by Léon Theremin (born Lev Termen), the Theremin is played using hand gestures, but without actual contact between the performer and the instrument. The performer's hands act as ground plates for capacitive sensors in the instrument. One hand controls the pitch of the sound created by the instrument while the other controls the volume.

Developed in 1929, the *Ondes Martenot* can be played using either a keyboard, or by sliding a metal ring worn on the right hand index finger along a strip in front of the keyboard. The position of the ring on the strip corresponds to the pitch of the note produced. However, no sound is produced directly by the performer's gestures on either the keyboard or the strip. Instead notes are activated using the left hand on a series of controls which allow for the selection of the dynamic of the sound. Together the combination of right hand pitch selection and left hand dynamic selection create the sound of the instrument.

As with both the Theremin and the Ondes Martenot, the *Trautonium* offers the performer continuous (rather than discrete) control over the pitch of the sound created by the instrument. In the Trautonium a resistive wire is strung above a metal plate. The performer can create pitches by pressing the wire to the metal plate. The position at which the wire is pressed corresponds to the pitch of the sound created.

While the Theremin, Ondes Martenot and Trautonium concentrated on allowing continuous control of pitch, the *Electronic Sackbut* was designed to allow the performer to affect the timbre of the sound being created. The Sackbut keyboard allowed the performer to move the keys both vertically and laterally. Vertical movements modified the volume and attack shape of the sound. Horizontal movements performed pitch bending. The keyboards was played using the right hand. At the same time, the left hand was used to manipulate the timbre. Each finger of the hand manipulated a separate pressure-sensitive control, controlling aspects such as the main formant of the sound, the basic shape of the waveform and periodicity of the sound.

The Hands, developed by Michel Waisvisz in the 1980s, is a non-contact type instrument, where the performer creates sound using hand gestures in space. The performer wears two handpieces, each of which contain a number of sensors. These sensors measure finger positions, hand orientation and the distance between the hands themselves. One particularly interesting aspect of the development of the Hands is that at one point in time development was frozen. From this point on Waisvisz spent time mastering the instrument, becoming a Hands virtuoso.

Developed in 1987 by Max Matthews and Robert Boie the Radio Baton allows a musician to control a musical performance by moving two batons, each containing a different low-frequency radio transmitter, over a flat receiving surface. The instrument produces 3-dimensional position information for each baton over the receiver. The Radio Baton has been used both as an interface for conducting and also as a percussion instrument.

Laetitia Sonami's Lady's Glove, developed in 1991, uses a number of sensors on a glove to allow the performer to perform music through finger, wrist and arm movements. Gestures such as bending fingers, touching fingers together and moving the arms can be used by the performer to create sound. The Lady's Glove, together with the Hands, can be seen as forerunners of the FM-Gloves and T-Box instruments discussed in Chapter 7.

The Buchla Lightning instruments (developed in 1991 and 1995 by Don Buchla) use two wands held in the performer's hands. These wands emit infrared light which is detected by remote light sensors. The Lightning can then be used to detect the location, acceleration, velocity, and direction of the wands. As with the Radio Baton, the lightning can be used, among other things, as both a conductors interface and a percussion interface.

These instruments (and others like them) provide much of the background for the research presented in this thesis. As such, they also provide examples of the type of instrument being discussed in this work. The next section presents a general model of digital musical instruments, while the remaining sections discuss each part of the model in detail, relating them to existing digital musical instruments.

2.2 The Physical Interface

A number of models for digital musical instruments have been presented in the literature. Bongers (2000) presents a model for a digital musical instrument which is based on a more general human-machine interaction loop (see Figure 2.1). In this model the performer acts on the system (in this case, the digital musical instrument) through motor functions, which are detected by the system using sensors. The system can also act on the performer, through the use of actuators and displays. This model specifically requires memory and cognition on both the part of the performer and the system. If the system lacks this facility, then it becomes a "reactive" rather than "interactive" system[2].

Bongers (2000) also presents an extended version of this model, which models not just the performer-instrument interaction but also the performer-audience and audience-instrument interactions (see figure 2.2). This model is one of the few which considers the interactions involving members of the audience and the system, something which is more common as part of interactive rather than traditional musical performance. As Bongers states:

---

[2] Interestingly, in general computer science a reactive system and an interactive system are considered the same. If a system reacts to user input it is interactive.

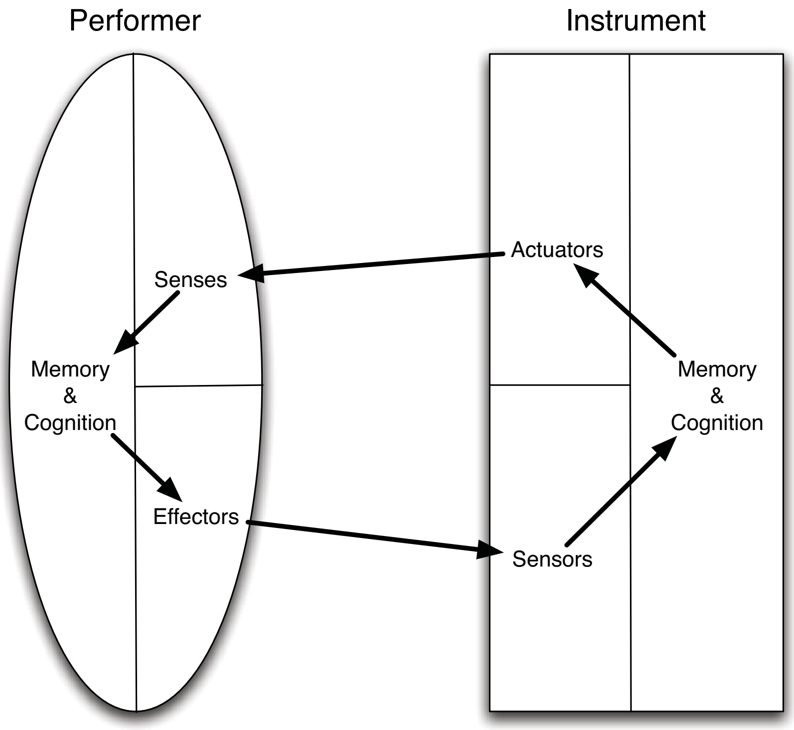

Figure 2.1: A model of the interaction in a digital musical instrument, from that of Bongers (2000)

> “In (musical) performance, there can be two active parties: the performer(s) and the audience. The audience can (and often does) participate by (even subtle and non-verbal) communication directly to the performer(s), which may influence the performance. Apart from this direct interaction between the parties, performer and audience can communicate with each other through the system. The system may facilitate new interaction channels...”

Figure 2.3 shows a model of a digital musical instrument as presented by Wanderley (2001). The instrument consists of a controller and a sound production

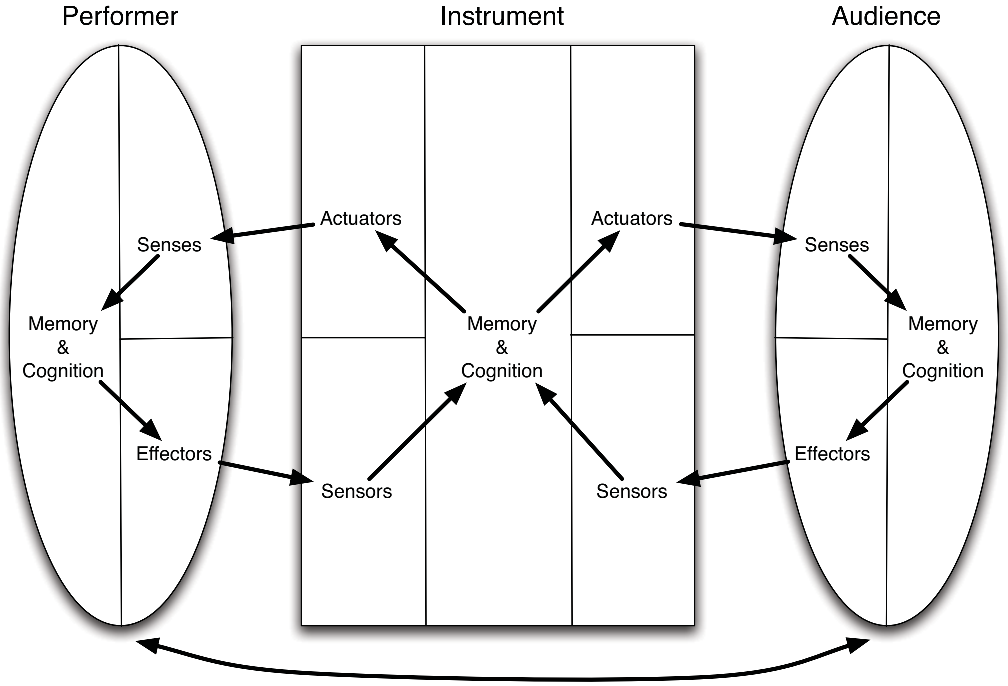

Figure 2.2: A model of the interaction in a digital musical instrument including audience interaction, from that of Bongers (2000)

system, connected by a mapping. Input to the controller is through performer gestures. This model specifies two forms of feedback from the instrument to the performer, *primary feedback* and *secondary feedback*. Primary feedback is the feedback from the controller itself and can include haptic, tactile, visual and even auditory feedback (such as the sound of key's clicking). Secondary feedback is the sound produced by the instrument's synthesis system.

While the models presented by Bongers (2000) and Wanderley (2001) show information flow from the performer to the instrument and back, this is not always the case in digital musical instruments. Yet, in many DMIs the flow of information is unidirectional, always moving from the performer to the instrument and on out

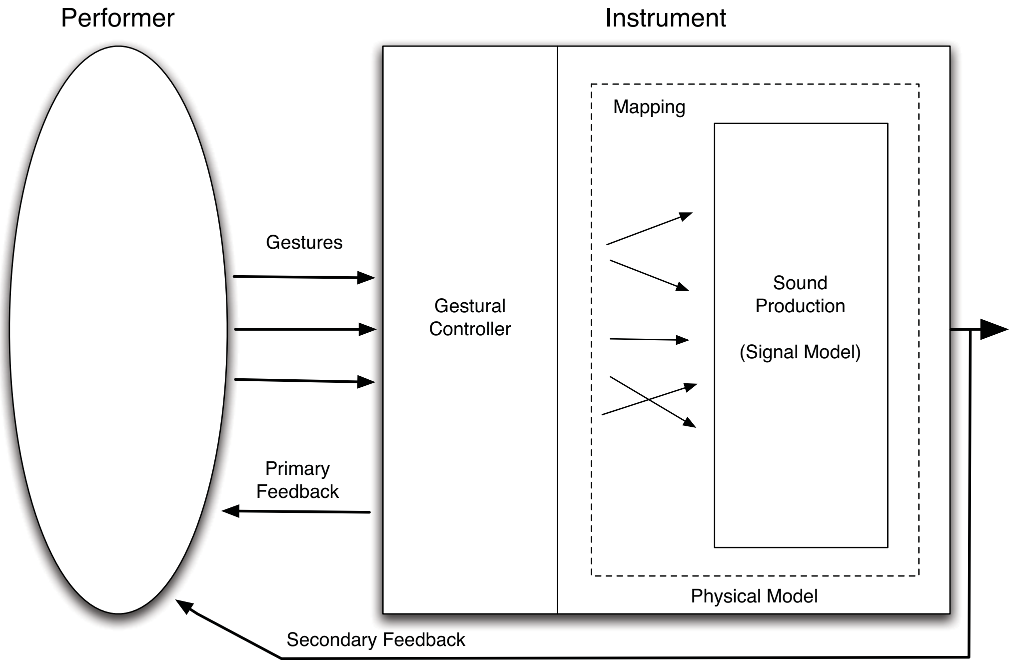

Figure 2.3: A model of the interaction in a digital musical instrument, from that of Wanderley (2001)

through the sound system. Cook (2004) presents a model of such a DMI, which he refers to as a "feed-forward" system. Figure 2.4 gives a representation of such a system. This is not being presented as the ideal system, but rather as an example of the standard configuration for digital musical instruments.

In order to remedy this lack of intimacy, Cook proposes a process of "remutualizing" the design of digital musical instruments, which involves concurrent development of the control, synthesis and feedback aspects of digital musical instruments. He presents a number of instruments developed using this process, some of which include features such as embedded vibrating elements for tactile feedback and embedded speaker systems so that the sound produced by the instrument comes from the instrument itself.

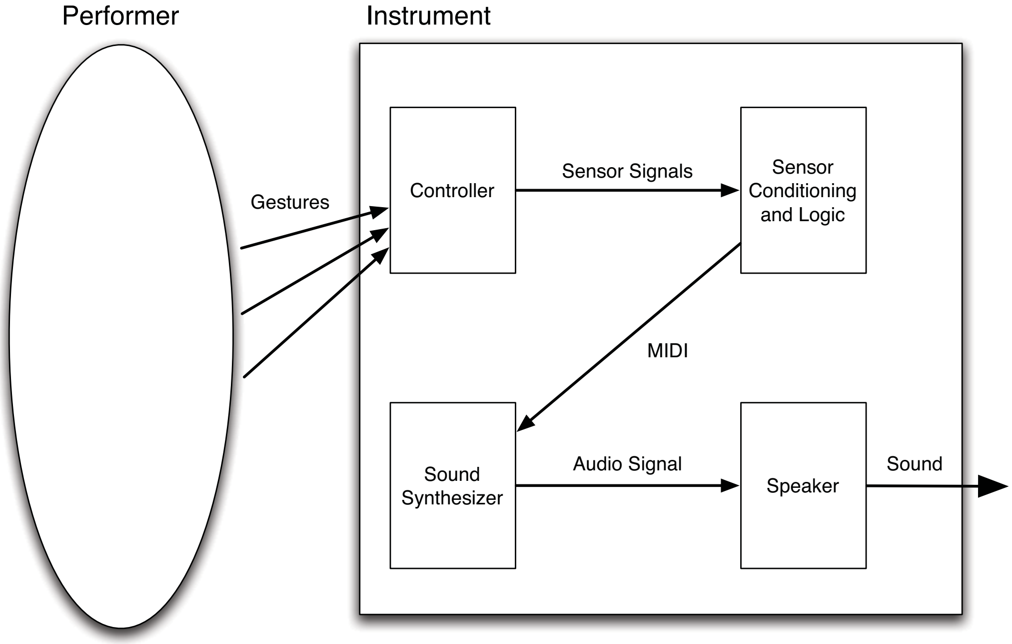

Figure 2.4: A model of the standard design of digital musical instruments, from that of Cook (2004)

A further interesting model was presented by Birnbaum (2007). This model includes a bi-directional mapping between the gestural interface and the *feedback generator*. The *feedback generator* in this case synthesizes both the sound of the instrument and the other forms of musical feedback, including vibrations. The gestural interface in this case then differs from the standard gestural controller. It has both inputs (sensors) and outputs (actuators) allowing it to sense performer gestures and also to produce feedback to the performer. Figure 2.5 shows this model.

The work presented in this thesis is based around a model which incorporates elements of a number of these different models. Figure 2.6 shows this model of a digital musical instrument, which is primarily based on those models presented

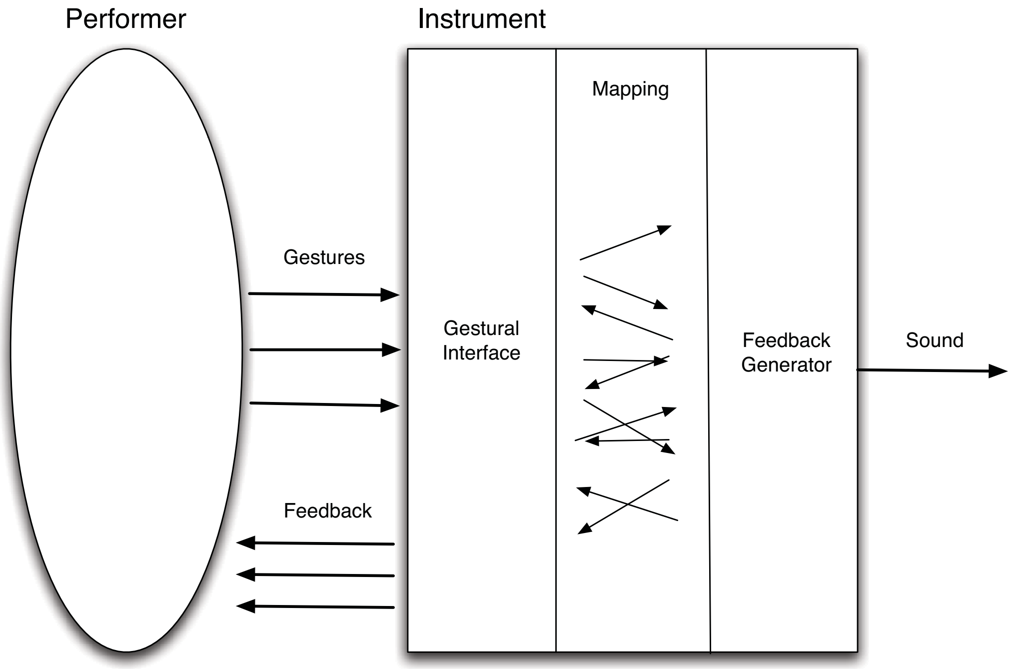

Figure 2.5: A model of a digital musical instrument, including bi-directional mapping and musical feedback generator, from that of Birnbaum (2007)

by Bongers (2000), Wanderley (2001), and Birnbaum (2007). It can be seen as consisting of 3 main components:

The physical interface containing the sensors, actuators and physical body of the instrument.

The software synthesis system which creates both the sonic output of the instrument and any visual, haptic and/or vibrotactile feedback.

The mapping system in which connections are made between parameters of the physical interface and those of the synthesis system.

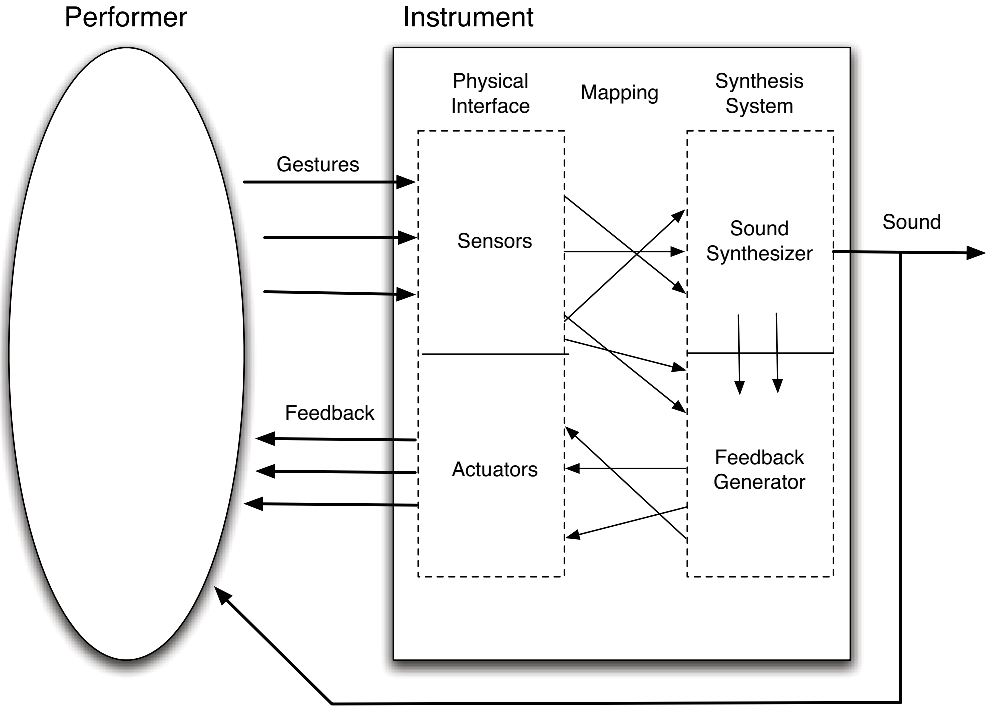

Figure 2.6: A combined model of a digital musical instrument, based on the work of Bongers (2000), Wanderley (2001), and Birnbaum (2007)

The physical interface of a digital musical instrument is the part of the instrument with which the performer is interacting. It consists of the physical body of the instrument, the sensors used to detect performer gestures and any actuators which produce feedback to the performer. It should be noted that the physical interfaces of some digital musical instruments do not contain all of these parts. For instance, instruments which track performer gestures using optical techniques (for instance using cameras or infrared distance sensors) may not have a body or any feedback actuators. Instead they consist entirely of the sensors used for the tracking.

The focus of this research is on DMIs whose physical interfaces contain most (or all) of these parts. The aim is to develop instruments which have the close coupling of the performer and the instrument that is present in most acoustic instruments. However, some work will also be presented dealing with instruments which either have no physical body or no contact between the performer and the instrument body.

2.3 The Instrument Body

It is possible to class digital musical instruments based on their relationship to existing acoustic musical instruments, resulting in the following classes of instrument (Wanderley, 2001; Bongers, 2000):

Instrument-like instruments are instruments designed to reproduce the features of an existing acoustic instrument as closely as possible. The most famous example would be the electronic keyboard, which is designed to be like a piano.

Instrument-inspired instruments have been inspired by an existing acoustic instrument but do not necessarily attempt to faithfully reproduce the features of that instrument.

Augmented instruments are acoustic instruments that have additional sensors added to them.

Alternate instruments are instruments which do not follow the design of a traditional musical instrument in any way. However, in many cases alternate instruments are based on existing objects.

Examples of instruments within each of these classes can be found in Section 3.1. Alternate instruments can be further sub-divided into classes based on certain features, as described by Mulder (2000):

Touch Controllers which are alternate controllers that require the performer to touch a physical control surface.

Expanded-range Controllers that either do not require physical contact between the performer and instrument or require only limited contact. In cases where no physical contact is required, such instruments have only a limited range of effective gestures. This allows the performer to make movements without any musical consequences.

Immersive Controllers which have few or no restrictions on performer movements. The performer is always within the sensing field of the instrument and so all of their movements have musical consequences.

Pirringer (2001) further classified immersive controllers based on the degree of immersion which they provide as either partially immersive or fully immersive. Partially immersive controllers include devices like datagloves that respond to just a single part of the human body. Fully immersive controllers respond to the movements of the whole body.

From these descriptions it is clear that the body design of an instrument-like instrument or an augmented instrument is limited (or at least heavily influenced) by the design of a traditional instrument. Alternate instruments on the other hand can be designed with any physical body imaginable. This allows for the design of the instruments body to be based on factors such as ergonomics, artistic goals, or aspects of musical theory.

2.3.1 Bases for Instrument Body Design

There has been some interest in the ergonomic design of digital musical instruments. In particular, the *aXio MIDI controller* was designed from an industrial design perspective and makes use of human factors techniques to create an ergonomic physical control surface (Cariou, 1994). Mulder (1998) notes that the ergonomics of the design are only evident for one specific gestural vocabulary. This is a factor that does not make it suited as a general purpose controller, but which gives it more in common with traditional instruments, as they generally exhibit a set of gestural constraints on the performer as a result of their design.

Similar ideas were used in the design of *Mr. Feely* as described by Armstrong (2006). This instrument was designed based on the concept of enaction, which revolves around the idea of embodied musical performance. He provides a set of five criteria which are required for embodied musical performance with digital musical instruments, which are as follows:

1. *Embodied activity is situated*, meaning that it arises from the interaction between the performer and their environment.

2. *Embodied activity is timely* and so possess real-time constraints which the performer must meet.

3. *Embodied activity is multi-modal*, involving the use of a number of distinct sensorimotor modalities at the same time.

4. *Embodied activity is engaging*, meaning that the sense of embodiment requires the performer to be present and that it consumes a large portion of the performers attention.

5. *The sense of embodiment is an emerging phenomenon*, which means that the sense of embodiment is not present in the beginning but grows as the performer's competence increases.

Within the area of embodied musical performance, the idea of *instrumentality* is a primary concern. In the specific case of the design of a new digital musical instrument (such as the aforementioned *Mr. Feely*) this means that the instrument should have the *feel* of a traditional instrument, but also that its "material embodiment should be indicative of a specific purpose" (Armstrong, 2006). In this case, the physical design of the instrument body informs the performer as to how the instrument should be played.

A further example of ergonomic design of the body of a DMI can be seen in the BentoBox (Hatanaka, 2003). The BentoBox was designed with the aim of creating an instrument that could be played in small spaces and using headphones. The idea was to create an instrument which could be played when commuting on public transport. This aim resulted in a list of requirements including small size, the ability to hold and play the instrument using just the hands and control it using only smaller finger, wrist and hand movements rather than larger arm movements. Using techniques from product design, a process of development including requirements analysis, rapid prototyping and user testing resulted in an initial prototype instrument which met the necessary requirements for a portable musical instrument.

The physical operation of moving mechanical systems has also formed the basis of the design of some alternate instruments. Sinyor (2006) designed a number of instruments whose bodies were based around such systems. The Gyrotyre is an instrument based on a rotating wheel (Sinyor and Wanderley, 2005). It consists of a small diameter wheel which is attached to a handle. The handle is held by the performer and the wheel spun. This causes forces which make moving the instrument in certain directions much easier, while making movement in other directions much harder. The SpringWave is constructed around a long spring (made from a toy known as a Slinky) which is suspended horizontally, fixed at both ends (Sinyor, 2006). Vibrations and deformation of the spring are sensed using a combination of sensing methods and used as control parameters for sound synthesis systems.

One interesting aspect of these instruments is that the choice of interacting with moving mechanical systems strongly informs the design of the rest of the instrument. For the Gyrotyre for example, the choice of a rotating wheel generating centrifugal forces as the main form of interaction results in a reduction in the number of possible ways of constructing the instrument. In some ways this is similar to the constraints in body design for acoustic instrument caused by their sound generation mechanisms, but here it is the interaction, not the sound generation, which informs the design.

One area where the physical body design of DMIs has seen a lot of research is that of tangible musical instruments. Such instruments use tangible objects which are physically manipulated by the performer to create sound. Examples include the reacTable* (Jorda et al., 2005), the TAI-CHI project (Crevoisier and Polotti, 2005) and the PebbleBox and CrumbleBag (OâĂŹModhrain and Essl, 2004; Essl and OâĂŹModhrain, 2006). All of these instruments create sound through the interaction of the performer with tangible objects, which requires much thought on the form of the objects and how this affects the performer-object interaction.

The reacTable* allows multiple performers to interact with colourful physical objects and to create sound by placing, moving, rotating and relating these objects on a luminous round table surface. The choice of the table and objects which make up this instrument was led by the intention of producing an instrument which was intuitive, easy to master (for adults or children) and suitable for both novices and experts.

The TAI-CHI project is based around the production of interfaces based on everyday objects, which allow for natural interaction without the need for extra hand-held devices. Interaction with the objects is sensed using acoustic tracking, which means that the sound of the users interaction with the object is used as the control for the system. An example of such an interface is a table which senses where the user is tapping it based on how long it takes the sound of the tapping to travel through the table. This allows for easily augmenting almost any object, with the result that according to Crevoisier and Polotti (2005):

> “a new expressive dimension for musical instrument is introduced by the possibility to communicate a message not only with the sound, but also with the symbolic nature of the object that is chosen for the interface.”

The PebbleBox and CrumbleBag instruments developed by Essl and OâĂŹ-Modhrain (2006) were designed to utilize:

> “familiar sensorimotor experiences for the creation of engaging and playable new musical instruments.”

These interfaces are based on real-world interactions which are similar to the virtual interaction being created by the sound synthesis system. Therefore, to control synthesis of friction sounds or coarse grain collision sounds, the authors propose the use of interfaces which make use of friction or collision interactions. In some ways the results of this are similar to those of the co-design of synthesis and interface proposed by Cook (2004), where the design process for both the synthesis and the interface continually feed back into each other to produce an interface which is more closely connected to the sound synthesis.

Within the category of alternative instruments, there are two more specific sub-categories of instrument which are also of interest, namely *collaborative instruments* and *bodiless, or open-air instruments*. Collaborative instruments present interesting challenges in the design of their bodies, as they must be suitable for playing by two or more performers simultaneously. While some tangible instruments, such as the aforementioned reacTable*, inherently allow this there have been some other DMIs designed with this specific goal in mind.

The Tooka is an example of such an instrument (Fels and Vogt, 2002). The Tooka is a two-performer wind instrument, consisting of a long flexible tube with a mouthpiece and keys at each end. Two performers play simultaneously, a process which results in interesting demands on the performers and the instrument itself. The body of the instrument had to be designed in such a way as to allow two performers to comfortably manipulate it and to stand up to stresses from possible conflicting movements by the performers. Such instruments also present interesting challenges for the performers, both with regard to cooperation during performance and communication and collaboration during practice and rehearsal.

The OROBORO provides an interesting variation on a collaborative instrument in that it not only requires the performers to work together to perform, but also makes use of "interpersonal haptic feedback" to transmit feedback from the primary hand of each performer to the secondary hand of the other (Carlile and Hartmann, 2005). Again, the aim of creating a collaborative instrument influences the body design of the interface, as it must allow the two players to not only comfortably perform together, but also allow for the easy transmission of feedback from one performer to the other.

Bodiless instruments (sometimes referred to as open-air instruments) are those instruments for which the performer is not necessarily acting on a physical instrument body. In many cases, the performer interacts with the instrument through gestures made in the air, which are tracked and used to manipulate synthesis parameters. Many such instruments use video cameras and movement tracking software to track performer gestures (for examples, see Hornof and Sato (2004) or Mase and Yonezawa (2001)), while others use ultrasonic or infrared sensing built into some form of central transmitter (e.g. Rich (1991), Livingstone and Miranda (2005) or Suzuki et al. (2008)).

One of the most common forms for such an instrument is that of a glove. Glove-based instruments allow the performer to play in the air, using hand, arm and finger motions. In some cases the body of the instrument is actually the body of the performer as they press on their own body to actuate the sensors. Such instruments present interesting design challenges as they often require detection of small finger movements and movements for which the performer has no feedback other than the sense of their own muscles. Examples of such interfaces include the Lady's Glove (Sonami, 2008), Scanglove (Kessous and Arfib, 2003), Genophone (Mandelis and Husbands, 2004), GRASSP (Pritchard and Fels, 2006) and GloveTalk-II (Fels and Hinton, 1995). A glove-based instrument developed as part of the research for this thesis, the FM Gloves, is described in detail in Chapter 7.

2.4 Sensors

Bongers (2000) states that sensors are the “sense organs of the machine”. What this means is that in terms of human-machine interaction, sensors allow the machine to detect the actions of the human, just as our sense organs (eyes, ears) let us detect the responses of the machine. In a digital musical instrument it is the sensors which allow the instrument to detect the gestures of the performer and use those gestures to create sounds. As the interaction and the sound generation are not based on physical systems like those of traditional instruments, but rather on mappings of gestural parameters to sonic parameters, DMIs rely on sensors. While sensors form such an important part of a DMI there has (with some notable exceptions) not been much research into determining the best sensors for use in controlling specific parameters in a DMI.

2.4.1 Classification of Sensors

In order to study and compare sensors, it can be useful to first classify them. There have been a number of attempts to do so, in both the field of digital musical instruments and also in the broader field of sensor technologies. An in-depth classification has been provided by White (1987), in which sensors are classified using 6 parameters:

1. the quantity being measured (the measurand)
2. the technological specifications of the sensor (such as the range, resolution, accuracy etc.)
3. the means of detection (whether biological, mechanical, physical etc.)

4. the conversion phenomena (e.g. photoelectric, chemical transformation, electromagnetic etc.)
5. the material of the sensor
6. the fields of application

Such a classification scheme provides a detailed analysis of a sensor and allows for comparison between sensors based on a large number of criteria. There have also been a number of extensions and adaptations of this classification, including Fraden (2004) and Pallas-Areny and Webster (2001).

Within the field of DMI design, Bongers (2000) categorizes sensors based on the human output modalities which they detect, with a specific focus on those modalities generally used in instrumental performance. This allows sensors to be classified into the following categories:

- muscle-action sensors
- blowing sensors
- voice sensors
- other sensors

While both blowing and vocalization are technically performed using muscle action (within the throat), Bongers treats them separately. Such separation could be justified based on the differences in how they interact with the instrument, including differences in feedback from physical muscle action (such as hitting, pressing, pulling etc.) and the less physical acts of blowing or speaking/singing.

The final class of sensors (other sensors) are those which detect changes in the state of the body. These include factors which are within the direct control of a human being (such as bio-electricity from muscle movements) and those which are not (such as blood pressure, temperature etc.).

Vertegaal et al. (1996) provides another classification of sensors for digital musical instruments, which is based on the type, range and resolution of the sensor and also of the feedback provided by the sensor. The parameters used for classification in this case are:

1. physical property sensed
2. resolution of sensing
3. direction of sensing
4. type and amount of feedback provided

The physical property being sensed includes properties such as position and force. The resolution is represented on a continuum from 1 to infinity. The direction of sensing determines whether the sensors is, for example, linear or rotary for position or isometric or isotonic for force. Finally, the type and amount of feedback provided by the sensor includes tactile, kinaesthetic and visual feedback, each of which is once again represented by a continuum from 1 to infinity.

Some issues arise with this particular classification. Resolution will differ from one particular sensor to another (even within the same class) depending on how they are manufactured. Furthermore, the resolution is also dependant on external factors, such as the how the sensor is used and the electronic systems to which it is connected. Consider, for example, the Force Sensing Resistor (FSR): the FSR can be used to sense force pressing on the sensor (giving it a high resolution) or as a simple touch switch (giving it an on/off or binary resolution).

Similarly, some sensors may be used to sense a number of physical properties. This makes the classification based on the property sensed dependant on the implementation of the sensor. Section 3.2.1 provides more information on this, along with specific examples.

Finally, while measures of the types and amounts of feedback given by each sensor are provided by Vertegaal et al. (1996), no indication is given of how this is calculated. If sensors are to be classified based on the amount of each feedback provided, a metric is required to allow calculation of this parameter.

2.4.2 Comparing and Evaluating Sensors

Once a classification of sensors has been decided upon, it becomes possible to compare different classes of sensors in order to determine the most suitable one for a particular application. To this end, Vertegaal et al. (1996) produced a mapping from classes of sensors to classes of musical function. For that work, sensors were classified as just described, while musical functions were classified based on a simple, three class system consisting of *absolute dynamical*, *relative dynamical* and *static* functions. Absolute dynamical functions are those which change often and where the aim is to select an absolute value from those available. An example of this would be selection of a pitch to play on an instrument. Relative dynamical functions also change often over time, but do so relative to some baseline, rather than by selection of an absolute value. An example of this would be the modulation of a given pitch to produce a vibrato. Finally, static functions change rarely. Examples of this would include tuning selection or key selection. From these classifications, they produced a graphical representation of the suitability of specific classes of sensors for specific musical functions. However, no experimental evidence was provided to validate the resulting mappings.

Following on from this, Wanderley et al. (2000) attempted to experimentally evaluate the use of sensors for the control of a single specific musical function. The function chosen for evaluation was that of pitch modulation which can be classified (based on Vertegaal et al. (1996)) as a relative dynamical function. They examined control of vibrato using a linear position sensor, a force sensing resistor and the tilt of a stylus on a Wacom tablet. The participants played two notes by moving the stylus from one point marked on the tablets surface to another, modulating the second note to produce vibrato. For the FSR and linear position sensor, modulation was performed with the secondary hand (i.e. not the hand manipulating the stylus) while the tilt was performed using the same stylus (and therefore same hand) as the note selection.

They found that the FSR received the highest preference rating, which is consistent with the mapping described by Vertegaal et al. (1996). However, they found that the linear position sensor out-performed the tilt movement. Based on the classification used, the tilt movement is a rotary position sensor, which (again according to Vertegaal et al. (1996)) should out-perform the linear position sensor for a modulation task. This result could indicate a problem with the original mapping, but it is also possible that the difference between this result and the theory they were testing is due to the use of two hands for the FSR and linear position sensor and only one hand for the tilt.

It is possible that the physical separation of the two tasks (note selection and modulation) to different hands better followed the perceptual structure of the task the participants were being asked to perform (Jacob et al., 1994). This means that as the task can be seen as being composed of two separate sub-tasks, of note selection and pitch modulation, an input device which mirrors this structure would be the most usable. In this particular case, this would give an advantage to both the FSR and linear position sensor methods when compared to the stylus tilt method.

The separation of the task into two parts also fits with research in high degree of freedom human-computer interaction tasks. Masliah (2001) found that users prefer to separate translation and adjustment tasks and perform better at combined tasks when they approach each part separately. For instance, studies have shown that performance in a 6 degree-of-freedom docking task, consisting of a translation and rotation in 3D space, is improved when the user approaches the task as two separate sub-tasks (Masliah and Milgram, 2000). In terms of musical instruments, a similar example would be the performance of a note with vibrato. This task can be separated into two sub-tasks, consisting of first selecting the note and then modulating it to add the vibrato. Interestingly, for many traditional acoustic musical instruments these sub-tasks are performed using the same "input device" and so are not necessarily separated in what may be the optimal perceptual structure. This offers some interesting possibilities for the design of digital musical instruments, where such separation can be easily created through the design of the instrument.

2.5 Feedback

When playing an acoustic instrument the performer receives feedback from the instrument through a number of channels (see Figure 2.6). This feedback includes visual, haptic, sonic and tactile feedback. As previously mentioned, when performing on a digital musical instrument some of these channels of feedback can be missing.

Within the field of DMI design there has been much interest in the potential to "unchain" the performer from the physical constraints of instruments, through the use of non-contact sensing technologies (Rovan and Hayward, 2000). Instruments such as the Theremin, the Buchla Lightning (Rich, 1991) and the Twin Towers (Tarabella et al., 1997) allow performers to control aspects of sound synthesis using "open air" gestures. Each of these instruments uses different forms of sensing, but allows the performer to play sounds without touching the instrument itself. For the Theremin, the performer controls pitch and amplitude by varying the distance between their hands and two antennae. For the Buchla Lightning, the performer holds a baton in each hand and creates sound by moving the batons within the field of view of the instrument itself. For the Twin Towers, playing involves the performer moving their hands within a certain volume of air above a number of infrared rangefinders. For all of these instruments, the use of non-contact sensing techniques results in the loss of the tactile and haptic feedback from the instrument itself, causing the performer to have to rely on the other channels of feedback.

Performers are then forced to rely on visual and sonic feedback from the instrument as well as proprioceptive cues from their own body. While this might seem adequate, there are a number of possible issues with these feedback channels.

Studies of human performance have shown that while beginners generally rely on visual feedback, those who have mastered their instrument make use of haptic and tactile feedback (Keele, 1973). In a performance setting visual feedback can be inadequate or impractical. For instance, there can be more important visual cues such as interaction with other performers or with the audience, or reading a score. Also, the physical feedback channels are more tightly coupled than visual and auditory channels (Rovan and Hayward, 2000).

Even with DMIs which have a physical body for the performer to interact with there can be limitations to the physical feedback provided by the instrument. First and foremost, the sound from a digital musical instrument is generated by a computer system and comes from a speaker system which is generally located away from the performer. Performers interacting with an acoustic instrument receive vibrations from the instrument which are directly caused by the sound generating mechanism and so are directly linked to the state of the instrument itself. The lack of these vibrations in a digital musical instrument results in a reduction in the amount of information available to the performer through the tactile feedback channel (Chafe, 1993; Armstrong, 2006). This lack of feedback from the instrument can result in a disconnect between the performer and the instrument, a situation which is made worse by the lack of any sense of the sound coming from the instrument itself (Cook, 2004).

As mentioned at the beginning of this chapter, DMIs often also lack much of the haptic feedback which is present in acoustic instruments. Keys and strings require the use of force to manipulate them, membranes “push” back when struck[3].

---

[3] Interestingly, Berdahl et al. (2008) describes the Haptic Drum, an instrument which consists of a woofer loudspeaker with a sunglass lens attached to its cone. This lens is struck with a drumstick. The system senses the strike and sends feedback to the performer using the woofer.

The performance of an acoustic instrument requires a certain amount of physical effort, which is often much greater than that required in the performance of a DMI. Yet such a lack of effort, while perhaps useful in the design of systems for general human computer interaction, is not ideal for the design of a digital musical instrument. As Ryan (1992) states, "Effort is so closely bound to expression in playing traditional instruments", that digital musical instruments which can be played with minimal effort may not allow for a useful level of expression. In fact, he states that it may be more useful to design an instrument which requires an enormous effort to play than one which requires almost none.

2.5.1 Vibrotactile Feedback

One of the most straightforward methods of providing vibrotactile feedback to the performer is to embed the sound generation in the instrument. This has the dual advantage of providing vibrotactile feedback to the performer and also causing the instrument's sound to come from the instrument itself (Cook, 2004). The BoSSA (Bowed Sensor Speaker Array) described by Trueman and Cook (2000) is an example of such a system. For BoSSA, the instrument body consists of 12 loudspeakers mounted in a (roughly) spherical enclosure, to which the various sensors are attached. The instrument is played seated, with the speaker enclosure between the performer's legs, in a way which is similar to the cello. This arrangement allows the performer to feel the vibrations created by the instrument, as well as causing the sound to radiate from the instrument itself. It also has the added advantage of allowing the instrument to radiate sound in a directional manner, which is more

---

This feedback can simulate the vibration of a drum membrane. It can also be used to enable techniques which are difficult to play on an acoustic drum, such as one-handed drum rolls.

consistent with that of an acoustic instrument.

Armstrong (2006) also acknowledges the importance of having the sound radiate from the instrument itself, both for the performer and for the audience. He states that the

> "perceptual localisation of the origin of the sound is an important indicator of the instruments phenomenal presence, both for the performer, fellow performers, and the audience."

However, he points out that in some cases this can be difficult to accomplish, such as when large amplifiers and speakers are required in order to allow the instrument to be used without further amplification. In such cases, he recommends that an external amplifier and loudspeaker be used, but placed as close to the performer as possible. By placing the speaker on the floor near the performer it is possible to feel the vibrational energy of the instrument through the legs and torso.

The Viblotar, one of the instruments developed as part of the research for this thesis makes use of embedded speakers and amplifiers in the instrument body to provide both vibrotactile feedback to the performer and to locate the sound production within the instrument. A detailed description of the instrument and its design criteria is provided in Chapter 6.

Another possible approach for creating vibrotactile feedback, useful when it is not possible to mount speakers to the instrument body, is the use of vibrotactile actuators. A number of different devices are available which can be used to create vibrations within a digital musical instrument. These devices vary in size, cost, availability and the type and freedom of control which they offer. Chapter 5 includes a survey and comparison of a number of such devices.

In one example of the generation and use of vibrotactile feedback, Chafe (1993) examined the use of a vibrotactile actuator to allow for closer control of physically modelled sounds. He created a controller with vibration feedback to allow performers to sense the modes of vibration in the lips of a brass instrument model. He found that performers were more easily able to control the model when using this controller. It enabled them to remain within the range of parameters which gave a stable system.

For “open air” instruments, Rovan and Hayward (2000) describe the development of vibrotactile actuators and a typology of tactile sound events which can be used to add vibrotactile feedback. They used a variety of tactile signals to allow performers to determine their position in a virtual space. Signals used included different spectral envelopes to generate a continuous texture with ridges caused by noise bursts to indicate zone crossings in space. Vibrations were passed to the performer using both a vibrating ring placed on the finger and vibrational actuators under their feet. A software system was developed which allowed control of a number of parameters of Tactile Simulation Events (TSEs), including the frequency, waveform, envelope, duration, amplitude, number of repetitions and delay between repetitions.

Using these TSEs, the authors performed an experiment to determine which features of vibration can be perceived by the performer and in which ways. They found that performers could sense 8 to 10 discrete frequency steps between 70 and 800 Hz, but that the use of “larger-scale audio gestures”, such as rapidly rising or falling pitch curves were more perceptible and more memorable than using discrete pitches. They also found that spectral content performed well as a vibrotactile cue.

By increasing spectral content, they were able to generate a range of textures, from smooth textures using pure sine tones, to rough textures using a noisy spectrum. Finally, they found that short tone burst events, consisting of fast attack and decay envelopes, were very useful for noting boundary crossings in the instrument performance space.

Human skin senses vibration through four different types of receptor. These receptors sense different types of vibration based on the area, frequency and amplitude of the stimulus. The FA/SA system, developed by Birnbaum (2007), models each of these separate channels of mechanoreception. It includes functions which extract perceptually meaningful sonic features from audio signals and map them to perceptually meaningful features of the vibration signal. It has been used in a number of instruments to provide vibrotactile feedback. One example of this is the BreakFlute, which is a flute-like tactile display device that uses breakbeat music samples (Birnbaum and Wanderley, 2007). In such a case direct vibrotactile feedback from the audio signal might not be meaningful and so a new vibration signal can be generated from aspects of the signal itself.

2.5.2 Haptic Feedback

Haptic feedback involves the creation of forces which resist the movements of the performer. These forces can be used to create the feeling of interaction with virtual objects and surfaces and/or to simulate the effort required in the physical interactions present in a musical instrument. Much research into haptics is in the areas of telepresence and teleoperation, but there has also been some significant work on the use of haptic feedback for musical interfaces.

One of the earliest applications of haptic feedback to the design of digital musical instruments is that described by Cadoz et al. (1984). The authors describe a haptic feedback device which they designed called the *retroactive touch transducer key*. This device was inspired by a piano keyboard key, but offers a much larger displacement and contains a motorized actuator which can be used to provide forces to the key to allow it to resist movement. Together with their *Cordis* system, this device allows them to simulate certain musical interactions through physical modelling of both the sound synthesis and the physical interaction.

Another keyboard-like haptic system was developed by Gillespie (1992). This system models the performer and instrument as dynamical systems which interact through a port (in this case the haptic device). It has been used to simulate the action of a number of different keyboard-based instruments, including that of the grand piano. Nichols (2000) developed a violin-like haptic controller, which can sense the violinist's bow stroke and also simulate the friction and vibration of the string on the bow.

Chu (2002) examined the use of haptic feedback to provide information about positioning within audio tracks to a user manipulating an audio editing system. This system generated signals for a haptic knob , which included features such as detents, pops, textures and springs. Some of these features, most notably the detents and textures, are similar to features used in the system developed by Rovan and Hayward (2000), but in this case are connected to a force rather than tactile feedback system.

OâĂŹModhrain (2000) evaluated the effects different haptic signals on the accuracy of performance of Theremin melodies using a haptic device called the Moose. These signals included simulations of springs (both positive and negative), constant forces and a viscous condition. The addition of any of these feedback signals proved to offer improvements in performance accuracy over performance without any force feedback. Interestingly, she noted that the addition of any form of force feedback, even that created by attaching an elasticated band between the performer's hand and the antenna of the Theremin produced an increase in the ease of performing with the instrument.

Finally, the DIMPLE system, created by (Sinclair, 2007) is a software environment allowing the run-time creation of a physically dynamic, haptically-enabled virtual scene. This system can generate both haptic and vibrotactile signals, where the vibrotactile signals are generated based on the output of the sound synthesis system being used. It allows users to interact with the virtual scene through a number of different haptic interfaces (Sinclair and Wanderley, 2007). One application discussed for the system is the creation of friction models which can be used to simulate aspects of bowing a stringed instrument.

2.6 Conclusion

Section 2.2 described the various parts of a digital musical instrument: the physical interface, the mapping and the synthesis systems. From these parts, the one which we are most interested in here is the physical interface. The physical interface is the portion of the instrument with which the performer physically interacts. In order to improve the performer-instrument interaction, one possible course of action is to consider the design of the physical interface. Specifically, we can examine the sub-parts of the physical interface and determine possible areas of research for each of them.

For the instrument body, the areas of industrial design and ergonomics can provide guidelines in the design process (Cariou, 1994). In addition to this, ideas such as those of *enaction* and the development of the idea of *instrumentality* can be used to develop DMI bodies which are more easily played, or which feel similar to those of traditional musical instruments (Armstrong, 2006). In particular the freedom of design of the body of digital musical instruments can allow for the use of techniques from the study of ergonomics to create instruments which reduce the risk of performance-related injury that is present in many traditional instruments. In fact, much research is taking place in the application of ergonomics to more traditional instruments, although the design of these instruments is much more restricted by their methods of sound generation than is true of digital musical instruments (Marmaras and Zarboutis, 1997; Storm, 2006).

The use of sensors in digital musical instruments allows for a number of possible avenues of exploration. It is possible to create new sensors, perhaps using common, low-cost materials such as rubber, paper and conductive pigments (Jensenius et al., 2006; Koehly et al., 2007; McElligott et al., 2002). Another area, already mentioned in Section 2.2 is the evaluation of sensors for specific musical functions. While some work has taken place in this area, there is still a need for detailed empirical research into factors such as user preference, quantitative measurement of sensor performance and the effects of learning and previous musical experience on sensor usability.

For the provision of feedback in digital musical instruments,we can examine and evaluate the use of a variety of different actuators for the provision of vibrotactile and/or force feedback. We can develop new actuators which can be used to provide more controllable or higher levels of feedback than is available with existing actuators (Yao, 2004). It is also possible to examine the creation of optimal signals for vibrotactile feedback, by taking into account both the response of human skin and also of vibrational actuators to different frequencies of vibration.

The next chapter examines the application of the research discussed in this chapter in existing digital musical instruments. It includes a detailed survey of 266 digital musical instruments presented in the 8 yearly international conferences on New Interfaces for Musical Expression, including the design of the instrument bodies, sensors and feedback systems for these instruments.

Chapter 3

A Survey of Existing DMIs

This chapter provides an in-depth survey of existing digital musical instruments, accomplished through a detailed literature review of all of the papers and posters from each of the 8 years of the conference on *New Interfaces for Musical Expression*. In total, this survey encompassed 577 papers and posters, containing descriptions of 266 different instruments. Some papers described multiple instruments and some instruments were described in multiple papers. Those instruments described in multiple different papers (generally in different years) usually involved descriptions of new applications or design improvements over the original. For reference, a breakdown of the number of papers presented at each NIME conference is given in Table 3.1.

This survey focuses on the specific area of physical interface design for these

	2001	2002	2003	2004	2005	2006	2007	2008	Total
Papers	14	48	49	54	77	131	105	99	577
Instruments	26	41	32	29	42	36	31	29	266

Table 3.1: Number of papers and instruments presented at each NIME conference

instruments. No account was taken of mapping or synthesis systems. Also, instruments which were completely software-based (using only the keyboard and mouse in a standard human-computer interaction paradigm) were not included. This chapter is divided into 3 main sections, corresponding to the 3 components of the physical interface described in Section 2.2.

3.1 Instrument Body Design

As discussed in Section 2.2, the design of the body of a digital musical instrument is generally dependant on the class of the instrument itself. The classes of instrument which are examined in this survey are based on their relationship to existing acoustic instruments, as instrument-like controllers, instrument-inspired controllers, extended instruments and alternate controllers.

As noted by Miranda and Wanderley (2006), the classification system has some issues, as it is not exhaustive and classes may overlap. Also, the alternate controllers class can be seen to be very broad, as it includes any instrument which does not fit into the other classes. A more thorough classification might be possible using these same classes, but presenting instruments on a continuum between the discrete points represented by the classes themselves (Miranda and Wanderley, 2006; Manning, 2004). However, a discrete classification system allows for a straightforward comparison of instruments and is more consistent with the existing literature in this area. Therefore, for the purposes of this research the previously described discrete classification system will be used.

Table 3.2 shows the number of instruments for each class across the 8 years of the NIME conference. On examination, it is clear that (as expected) the alternate

	2001	2002	2003	2004	2005	2006	2007	2008	Total
Instrument-like	1	2	2	2	2	4	4	1	18
Instrument-inspired	2	4	1	1	-	3	2	1	14
Extended instrument	2	4	5	3	5	7	6	5	37
Alternate controllers	21	31	24	23	35	22	19	22	197
Total	26	41	32	29	42	36	31	29	266

Table 3.2: Classes of instruments presented at the NIME conferences, by year

controllers class contains the majority of instruments presented. In total, they make up more than 74% of the instruments found by this survey. The remaining instruments are spread over the classes of extended instruments (14%), instrument-like controllers (6.8%) and instrument-inspired controllers (5.2%). The remainder of this section will discuss each of these instrument classes and provide examples of instruments from each class.

3.1.1 Extended Instruments

An interesting example of an extended instrument is the Mutha Rubboard. This is an instrument designed around a rubboard (or washboard or frottoir) which is often used in Zydeco music (Wilkerson et al., 2002). It was specifically designed with experience washboard players in mind and the main aim was to "maintain their natural relationship with the instrument". The Mutha Rubboard uses a traditional washboard and keys, to which a number of sensors have been added. This design allows the instrument to be played using existing techniques, but in a number of different ways. For instance, it is possible to play the instrument as an acoustic washboard, as an electric washboard (using the built in piezoelectric pickups), or as an extended washboard, with control of effects and other sounds through the capacitive sensing of the washboard keys.

Wind instruments have often been used for the creation of extended instruments within the NIME community. This may be a result of many wind instrument performers having "spare bandwidth" as described by Cook (2001). This allows such performers to manipulate controls other than those which are inherent within their instrument. Examples from the papers presented at NIME conferences include 4 extended saxophones (Burtner, 2002; Schiesser and Traube, 2006; Favilla et al., 2008)[1], a trumpet (Kartadinata, 2003), flutes (Palacio-Quintin, 2003; da Silva et al., 2005), trombones (Farwell, 2006; Lemouton et al., 2006) and a tuba (Cáceres et al., 2005).

A number of augmented stringed instruments have also been presented, including a guitar (Bouillot et al., 2008), 2 violins (Bevilacqua et al., 2006; Overholt, 2005) and a cello (Freed et al., 2006).

3.1.2 Instrument-like Controllers

Instrument-like controllers, those which attempt to model the gestural interface of an acoustic instrument as closely as possible, are the least common controllers found in this survey. When developing new digital musical instruments based on acoustic instruments, it seems that it is more common to attempt to extend or improve the capabilities of the acoustic instruments control surface than to copy it completely. However, there have been some notable exceptions to this.

The FrankenPipe is one such instrument. The FrankenPipe is based on a set of bagpipes, to which a number of sensors have been added (Kirk and Leider, 2007).

---

[1] Although the Gluisop and Gluialto presented by Favilla et al. (2008) could almost be considered two versions of the same instrument, as they differ so little in sensing

Unlike an extended controller however, the FrankenPipe is designed not to make any acoustic sound, but purely as a digital instrument with the form of an acoustic instrument. Such an instrument can still sense the traditional acoustic instrument performance gestures, but uses them to control a digital synthesis system. As noted by Miranda and Wanderley (2006), this type of instrument provides a control surface which is as close as is possible to the acoustic instrument.

An unusual example of an instrument-like controller is the Croaker (Serafin et al., 2006). Unlike other instrument-like controllers which are based on more well-known acoustic instruments, the Croaker emulates one of Luigi Russolo's Intonarumori (noise intonters), the Gracidatore. The Intonarumori were a series of 27 instruments built around 1913 by the Italian Futurist composer and painter Luigi Russolo that worked as acoustic noise generators to create a variety of everyday noise sounds, from rumbles to screeches (Serafin, 2005). The original Gracidatore (or Croaker) was a mechanical instrument which used a toothed wheel mounted on a crank to excite a metal string. An external lever allows controlling the tension of the string thus offering some pitch control.

The digital Croaker allows the same form of interaction, through a crank and a lever. It allows (through its synthesis system) for the simulation of the types of sounds created by the original Croaker instrument. The digital Croaker also allows the possibility to the performer of controlling a variety of other sounds, whether based on those of other Intonarumori or completely different sounds.

3.1.3 Instrument-inspired Controllers

One of the first examples of an instrument-inspired controller from the NIME conferences is the Accordiatron, which is based on the traditional squeeze-box or concertina Gurevich and von Muehlen (2001). The Accordiatron allows for a number of similar performance gestures to those performed when playing a concertina, in that it can sense squeezing, button presses and twisting of the ends. For an acoustic concertina, some gestures which are part of the performance technique are not essential to the performance of the instrument. These gestures are not directly involved in producing sound. In particular the twisting of the hands are not an essential part of concertina performance, but occurs nonetheless. For the Accordiatron these gestures are involved in the control and creation of sound, offering the instrument extra degrees-of-freedom which are not found in the acoustic instrument.

Also of interest is the EpipE, a controller based on the Uilleann pipes, which are traditional Irish pipes similar to the bagpipes (Cannon et al., 2003). The EpipE uses a variety of sensing technologies to provide a control interface which is similar to that of the Uilleann pipes, but with some controls removed, or made more easy to use. For instance, as the instrument is not producing sound using the pumping of air (as the acoustic instrument does) the bellows is removed, reducing the effort necessary to control the instrument. Compare this with the FrankenPipe bagpipe-like controller (see 3.1.2) which senses all of the acoustic instrument performance gestures.

3.1.4 Alternate Controllers

As can be seen from Table 3.2, the most common type of controller presented at NIME is the alternate controller. This is likely due in part both to the broad nature of the category itself and to the wealth of design possibilities offered by digital musical instruments. As such, the alternate controllers presented at the NIME conferences have covered a large range of different designs.

The Ski, by Huott (2002), presents an example of an alternate controller with a physical body. In this case, the body is a wooden structure resembling a large ski. It is played upright in either a sitting or standing position, using a number of position sensitive touch pads as controls. This upright playing position, coupled with the wooden construction of the instrument, can give the Ski a visual impact which is a somewhat like that of a traditional instrument when being played.

Examples of more unusual alternate controllers include the Gyrotyre and the T-Stick. The Gyrotyre is designed around a rotating bicycle wheel attached to a handle (Sinyor and Wanderley, 2005). It makes use of the physical behaviour of a simple dynamic system (in this case a spinning wheel) to allow the performer to play sounds with a number of different mappings. Most interestingly, the mechanics of the motion of the wheel result in certain inherent proprioceptive and force feedback to the performer based on how the wheel is spun and moved.

The T-Stick is made from a long thin PVC tube, to which a variety of sensors have been added (Malloch and Wanderley, 2007). The T-Stick can sense fingering information on multiple capacitive-sensing strips on the body, pressing force, torsion, impacts, as well as acceleration and orientation. The shape of the instrument allows it to be held in the hands in front of the body, shaken, spun, or (with the use of a spike similar to that on the base of a cello) to be played in an upright standing or seated position.

The NIME conferences have also seen a number of non-contact alternate controllers, including a number of systems based on tracking performer movements using a video camera. An example of one such instrument is the Iamascope+, which uses performer movements in front of a camera to generate both visuals and sounds (**?**). Other examples include the vision-based mouth interface described by Lyons et al. (2003) and EyeMusic, which tracks eye movement to create sound (Hornof and Sato, 2004).

Several glove-based controllers have also been presented at NIME. These have included systems based around custom-made gloves, such as the Genophone (Mandelis and Husbands, 2004), VIFE_alpha (Rodríguez and Rodríguez, 2005) and MusicGlove (Hayafuchi and Suzuki, 2008), or a combination of a custom glove and a commercially-available glove, such as the Scanglove presented by Kessous and Arfib (2003) or the GRASSP system (Pritchard and Fels, 2006).

For instance, Genophone uses a custom-made glove with bend sensors on each finger to allow the performer to perform with sounds which have been generated using an Artificial Life paradigm. The VIFE_glove used in the VIFE_alpha system consists of force sensing resistors (FSRs) mounted on the tip of each finger, allowing the performer to manipulate virtual “sonorous objects” in a real-time 3D rendering. When these objects collide they generate specific sound events.

Both the Scanglove and GRASSP systems make use of a pair of gloves. The Scanglove consists of a 5DT$^{\mathrm{TM}}$Dataglove worn on the non-preferred hand and a custom glove worn on the preferred hand. The custom glove consists of FSRs and bend sensors. In performance, the 5DT glove is used to recognise symbolic hand signs which are mapped to pitch values. The custom glove is used to trigger notes at the pitch set by the 5DT glove. It is also used to manipulate several continuous parameters of the scanned synthesis system used by the instrument.

GRASSP uses a Cyberglove$^{\mathrm{TM}}$on the right hand and a custom glove on the left. These gloves are used to control a speech and singing synthesis system. The custom glove has a series of nine touch sensitive switches, two on each finger and one on the thumb. By touching one of the other triggers with the thumb a plosive sound is generated. Other vocal sounds are generated using postures of the right hand.

Several datasuit- or exolskeleton-based instruments have also been presented at NIME. Afasia, by Jorda (2001), uses potentiometers mounted at the joints of an exoskeleton suit to track the movement of the performer's joints. It also makes use of touch sensitive contacts on the performer's torso, which are activated by pressing with a gloved finger. The Meta-Instrument 3 also uses an exoskeleton to track performer gestures (de Laubier and Goudard, 2006). Once again, potentiometers are mounted at the joints to measure rotation. The Meta-Instrument 3 also has a series of pressure-sensitive buttons and sliders mounted on pads for the hands.

Unlike the exoskeletons used by the Meta-Instrument and Afasia, the BodySuit is a datasuit-based instrument (Goto and Suzuki, 2004). It consists of a black bodysuit worn by the performer, with a total of 12 bend sensors mounted at the joints. The performer can use the BodySuit to control sound and video with large-scale body movements.

Each of these datasuit- or exolskeleton-based instruments are examples of immersive controllers, as defined by Mulder (2000). Applying the sub-classification sued by Pirringer (2001), the Meta-Instrument 3 can be considered to be partially immersive, in that it only tracks the movements of the arms and the performer can "escape" from it by letting go of the arms of the instrument. Afasia and the BodySuit, on the other hand, can be considered fully immersive, as they track the movements of the whole body and so performer movements always have a musical effect. It is also interesting to note that due to their exoskeleton-like construction, Afasia and the Meta-Instrument 3 both mechanically restrict the range of movement of the performer, which is not true of the BodySuit.

Finally, there have been a number of collaborative instruments presented at NIME, including two-performer instruments such as the Tooka (Fels and Vogt, 2002) or OROBORO (Carlile and Hartmann, 2005)[2] and instruments designed for a larger number of performers, such as the Jam-O-Drum (Blaine and Forlines, 2002), the Beatbugs (Weinberg et al., 2002) and the reacTable* (Jorda et al., 2005).

The Jam-O-Drum makes use of 6 drumpads connected to a MIDI drum module. Each player is in control of a single drumpad. A software system generates a MIDI percussion score with spaces in the score for the players to perform. Visuals are also generated to indicate which player (or group of players) is currently invited to play. This allows the system to steer the interaction into situations where a single player is performing solo, a subgroup of players is playing together, or all of the players are performing at once.

A Beatbug is a bug shaped musical controller, which is held in the hand and used to control percussive rhythmic motifs. Players perform by striking the Beatbug to trigger sounds and manipulating the bugs antennae to control aspects of the sound. Multiple Beatbugs are connected in a network, designed to allow children

---

[2] Both of these instruments were discussed in Section 2.3

Sensor	Occurences	Property Sensed
FSR	68	Force
Accelerometer	56	Acceleration
Video Camera	54	
Button/Switch	51	Position (On/Off)
Rotary Potentiometer	31	Rotary Position
Microphone	29	Sound Pressure
Linear Potentiometer	28	Linear Position
Infrared Distance Sensor	27	Linear Position
Linear Position Sensor	23	Linear Position
Bend Sensor	21	Rotary Position (Bending)

Table 3.3: Most popular sensors from NIME instruments

participate in the process of making music.

The reacTable* allows a number of performers to play together on an interface based around tangible objects placed on a transparent table. A camera under the table tracks the nature, position and orientation of these objects. By manipulating these objects, performers change the state of a sound synthesis system, creating different sounds. A projector, also mounted under the table, is used to present visual feedback on the state of the system to the performer through animations.

3.2 Sensor Use

Sensors provide the means of capturing the performer's gestures for a digital musical instrument.They form an integral part of the interaction between the performer and the instrument. This section will discuss the use of sensors in existing digital musical instruments. It includes a count of the number of instruments making use of each sensor and examples of exceptional sensors or uses of sensors.

Table 3.3 shows the most popular sensors in digital musical instruments presented at the NIME conferences, along with the number of instruments in which one or more of each particular sensor was found. Note that this is not a count of the number of sensors used (as an instrument may include multiple copies of a particular sensor), but instead offers a measure of the relative popularity of particular sensors. The total sum of sensor types used was 595 sensors, implying that on average each instrument used 2.25 sensors [3].

Interestingly, FSRs are the most popular sensor, used in 26% of all instruments, followed by accelerometers, which are found 21% of the instruments surveyed. These sensors, while easily available, are not generally associated with traditional computer music interfaces, such as on MIDI fader boxes and keyboards. Such interfaces more often make use of rotary or linear potentiometers and buttons/keys. The popularity of these sensors may be due to their ability to offer continuous real-time control (Cáceres et al., 2005), as well as the ability of many accelerometers to measure multiple parameters (such as acceleration, rotation, energy) with between 1 and 3 degrees of freedom (see section 3.2.1 for more details).

Buttons, which are the third most common sensor, are most often used in the surveyed digital musical instruments to allow mode changes, rather than the control or generation of musical parameters (see for example Jorda (2001), Wilkerson et al. (2002) or Singer (2003)). One exception to this is the Tooka (Fels and Vogt, 2002), where buttons are used to select notes in the same way as keys on a wind instrument. Similarly, the Accordiatron (Gurevich and von Muehlen, 2001) uses buttons either as note triggers or to trigger clusters of notes. Finally, the Tenori-on (Nishibori and Iwai, 2006) uses buttons both as note triggers to generate a note

---

[3] It should be noted that as many of the instruments based around video cameras (40 out of 54) used only 1 sensor, the average number of sensor types per instrument for non camera-based instruments is probably slightly higher than this.

when pressed and as triggers for sounds in a loop (in a way which is similar to programmable loops in a drum machine).

3.2.1 Sensing Multiple Physical Properties

As already mentioned, some sensors can be used in multiple ways, allowing them to sense different physical properties. In the previous example of the FSR, it can either be an isometric force sensor, or (by using it as a touch switch) a linear position sensor. It is also possible that the same sensor can be used to sense both of these properties at once.

Another example that can be seen in Table 3.3, is the accelerometer. An accelerometer is an acceleration sensor. It is acted upon by the force of gravity. While this force does not actually change, movements of the sensor cause an *apparent* change which the sensor can measure. This allows the accelerometer to be used to measure acceleration.

An accelerometer can however also be used to measure rotation. A simple calculation performed on the acceleration value allows calculation of the orientation of the sensor relative to the Earth's center of gravity[4]. This ability to measure multiple properties complicates the process of classifying the sensor [5].

It is also possible to extract further properties from the signal from some sensors. Again taking the accelerometer as an example, by integrating the acceleration data we can extract velocity data. By integrating this velocity data we can extract position data. Similarly, velocity and acceleration data can be extracted from

---

[4] This depends on the range of the sensor. This calculation is only correct if the acceleration measured is within the range of ± 1g

[5] Further examples of sensors which can be used to sense multiple properties can be found on the SensorWiki at http://www.sensorwiki.org

position sensors through a process of differentiation.

In many cases it is possible to sense each of these multiple physical properties from a particular sensor at the same time. A single acceleration signal can be used to calculation acceleration, velocity and rotary position. For example the T-Stick uses a pair of 3-axis accelerometers mounted in the ends of the instrument to extract both acceleration and orientation data (Malloch and Wanderley, 2007).

The overall result of this issue in classifying sensors may be that any classification of physical property sensing is dependant on the particular implementation.

3.2.2 Combining Sensors

As can be seen from Table 3.3, there are issues with classifying a device such as a video camera based on the parameter it senses. This is due to the fact that a video camera is in reality a matrix of simpler sensors. It is composed of an matrix of visible light sensors, from which an image signal is produced. However this raises the question of whether the camera is a visible light sensor, or some other form of sensor.

Similar issues exist with some commercial sensors which are (in effect) a combination of 2 or more sensors to allow for sensing of a number of different unrelated parameters. While the accelerometer discusses in the previous section can extract different physical properties from a single measurement, some sensors perform separate measurements for each parameter, often reporting them in separate signals.

An example of such a sensor is a magnetic position tracker, which reports position and orientation data of moving objects, such as those used by Marshall et al. (2002), Couturier and Arfib (2003) and Gadd and Fels (2002). Such sensors can also be made by combining existing sensors (such as pressure and position sensors) by locating one on top of the other. For instance, the Viblotar (described in detail in Chapter 6 and presented at NIME in Marshall and Wanderley (2006)) makes use of a linear position sensor mounted on top of an FSR to create a sensing strip which senses both the position and pressure of a performer's touch.

These sensor combinations can prove useful in determining multiple parameters of a gesture, but offer difficulties in classifying the parameter sensing, direction and resolution of the sensors themselves. The issue becomes one of deciding between classifying each of the parts of the sensor separately, or classifying the whole sensor based on the integrated nature of the gesture which it is sensing.

3.2.3 Custom Sensors

An interesting aspect of the use of sensors in digital musical instruments is the development of completely custom sensors. Unlike the previously described sensors made by joining 2 or more existing sensors, these are entirely new sensors designed with a specific purpose in mind. A number of examples of such sensors have been used in instruments presented at the NIME conferences.

One of the first examples presented at NIME of such sensors are the *Prexels* presented by McElligott et al. (2002). The authors created a sensor using a conductive polymer which allows them to sense force applied to the surface of the sensor[6]. Arrays of such sensors, placed either on a chair or a floor tile, were used to allow the control of effects on the sound of an acoustic instrument through shifting the performers center of mass.

---

[6] While this was the first example of such sensors presented at NIME, a similar sensor design using conductive rubber was used a decade earlier, in the second version of the Continuum, a continuous musical keyboard designed by Hakken et al. (1992)

For the Hyperbow, Young (2002) developed a custom linear position sensor system which makes use of a resistive strip run along the length of the violin body, through which are sent two square wave signals of different frequencies. By measuring the amplitudes of the these square wave signals (as measured at the bow), the position of the bow on the violin can be found.

A custom mechanical tilt sensor was created for Bangarama, a system to allow the creation of music using headbanging (Bardos et al., 2005). This simple sensor used a small free-swinging, which was worn on a cap, to measure if the head was level or tilted. Transitions from level to tilted and vice-versa were used to trigger musical events.

Freed (2008) presented a number of new sensors developed from piezoresistive fabric, along with two controllers made with custom sensors based around fibre and malleable materials. The controllers described were a Kalimba with custom made force sensors and the Tablo, a fabric-based multitouch controller.

Finally, Koehly et al. (2006) described the construction of several custom sensors using paper and rubber which had been impregnated with conductive ink or pigments. They provide details of the performance of such sensors in relation to specific physical changes and measurements of the reliability and repeatability of some of the sensors under measurement conditions.

3.3 Feedback

Feedback in a digital musical instrument can be considered to be either passive or active. Passive feedback is a direct result of the physical characteristics of the system, such as the noise of a switch being pressed, or the position of the slider on

	2001	2002	2003	2004	2005	2006	2007	2008	Total
Vibrotactile	1	1	2	1	3	4	1	2	15
Haptic	-	3	1	2	1	-	1	-	8
Visual	1	2	2	3	6	4	1	2	21
Additional Sonic	-	1	-	2	3	2	-	2	11
Temperature	-	-	-	1	-	-	-	-	
Total	2	7	5	9	13	10	3	6	55

Table 3.4: Types of active feedback provided by instruments by year. Several instruments provided a more than one type of feedback and so the totals do not indicate how many instruments provided feedback, but rather how many times each type of feedback was produced. The total number of instruments providing active feedback would be less than the total of 55 shown in this table.

a linear potentiometer. Active feedback is a direct response of the system to the performer's actions, such as the sound generated by the instrument, or a graphical display indicating the current note being played (Bongers, 2000; Miranda and Wanderley, 2006).

While all digital musical instruments inherently provide some passive feedback, some designers choose to implement active feedback in the instrument to communicate extra information to the performer. This active feedback can take the form of graphical display, vibrotactile feedback systems, or haptic feedback systems. In this section I review the use of active feedback in those instruments presented at NIME. I begin with an overview of the frequency of use of various forms of feedback in those digital musical instruments, followed by a discussion of specific implementations of each form.

Table 3.4 shows the results of this part of the survey. As can be seen from the results, visual feedback (usually in the form of graphical displays or projections) is the most common form of active feedback provided by these instruments. While

Method	Occurences
Loudspeakers	10
Vibrating Motors	3
Patented Vibrotactile Actuator	2

Table 3.5: Methods of providing vibrotactile feedback

the majority of instruments which provided active feedback produced only one form of feedback, some did produce two, usually in the form of both visual and sonic feedback (for example see Mäki-Patola et al. (2005) and Lock and Schiemer (2006)).

3.3.1 Vibrotactile Feedback

As discussed in Chapter 2 there are a number of different methods of providing vibrotactile feedback to the performer of a digital musical instruments. Table 3.5 shows the methods used in the instruments surveyed and the number of occurrences of each method.

The simplest and most common method of providing vibrotactile feedback is the embedding of loudspeakers within the instrument itself. An example of the use of loudspeakers in this fashion can be seen in the SqueezeVox Lisa, presented by Cook (2005). A speaker embedded in this accordion-based instrument allows the dual purposes of projecting the sound output from the instrument and providing vibrotactile feedback to the performer. These have previously been noted by the designer as important aspects in reducing the disconnect between the performer and instrument (Cook, 2004).

Similarly, the Viblotar instrument (presented at NIME by Marshall and Wanderley (2006) and described in detail in Chapter 6) use embedded speakers to produce the sound output of the instrument at the instrument body itself, resulting in vibrotactile feedback for the performer.

An interesting example of the use of loudspeakers and loudspeaker voice coils can be see in the *Cutaneous Grooves* system described by Gunther et al. (2002). This is not a digital musical instrument as such but rather a system for “tactile listening”, allowing people to experience a composition through the feeling of vibrations on their skin. A series of small voice-coil-based actuators placed on the shoulders, elbows, wrists, thighs and back of the knees of the listener, together with a larger speaker based pack worn at the base of the back provide vibration signals to the “listener”.

Several other controllers/instruments have made use of vibrating motors to produce vibrations to the performer. These have included the SoundStone and PeteCube. The SoundStone is a 3D wireless music controller with a form similar to that of a large stone, held in the performers hand and interacted with through shaking and pressing gestures (Bowen, 2005). The SoundStone uses a built-in vibrating motor to present information from the synthesis system to the performer, such as pulses to indicate the controller having reached a specific limit, or vibrations indicating a strike on a drum in a virtual percussion patch. The PeteCube is a system designed for multi-modal feedback, including vibration, sound and images (Bennett, 2006). It is a plastic cube with visible light sensors (light-dependant resistors in this case) on each face. Contained inside the cube are two vibrating motors of each with a different mass. Due to the difference in mass, these motors can each vibrate at different amplitudes, allowing a range of vibration signals to be produced.

Method	Occurences
Voice-coil Motors	4
Patented Force-feedback actuator	2
Fluid Brake	1
Servomotor	1

Table 3.6: Methods of providing haptic feedback

A final method used for the creation of vibrotactile feedback is the use of commercial controllers, such as the vibrotactile mouse used by the Cymatic system (Howard et al., 2003). This mouse makes use of a patented vibration actuator, developed by Logitech, to provide vibrations at a range of frequencies and amplitudes to the user. In the Cymatic system it is used to present information on the state of the physical modelling software instrument to the performer in real time. The same mouse is also used in the StickMusic system described by Steiner (2004).

3.3.2 Haptic Feedback

As with vibrotactile feedback, there are several different possible methods of providing haptic feedback to the performer. Table 3.6 shows the various actuators used in NIME instruments to produce this feedback.

The most common actuator used to produce haptic feedback in these instruments was to make use of a voice-coil motor[7]. The MIKEY keyboard, a force-feedback keyboard that can be used to simulate a variety of keyboard instrument actions, is an example of this (Oboe and De Poli, 2002). A further example is the

---

[7] It should be noted that a voice-coil motor is not the same as a loudspeaker voicecoil. A voice-coil motor is most commonly used to move the read-write heads in a computer harddrive and is used to provide a strong force, rather than the low force vibrations produced with a loudspeaker voicecoil.

Plank, which is a force-feedback actuated controller that can be used to provide a variety of haptic illusions when controlling a scanned synthesis system (Verplank et al., 2002).

As with the use of a vibrotactile mouse to produce vibrotactile feedback, the Cymatic system also makes use of a commercial controller to provide haptic feedback to the performer (Howard et al., 2003). In this case, a Microsoft SideWinder force-feedback joystick is used. The StickMusic system, on the other hand, which also provides haptic feedback through a commercial force-feedback joystick makes use of a Saitek Force joystick (Steiner, 2004). Both of these devices make use of patented force-feedback actuators to produce haptic forces.

Finally, two more unusual methods of providing haptic feedback can be seen in the Damper system, which uses a fluid brake for haptic feedback (Bennett et al., 2007) and the vBow, which couples the feedback from a servomotor to the input from a rotary encoder to provide a force-feedback violin bow interface (Nichols, 2002).

3.3.3 Visual Feedback

Many digital musical instruments offer some level of passive visual feedback through the sensors that make up the interface, such as the position of the slider on a linear potentiometer, or the position of the performer's finger on a linear position sensor. Some DMIs on the other hand offer additional active visual feedback to the performer. These range from graphical displays for instruments with a large software component, to embedded LEDs, to virtual reality (VR) headset systems. Visual feedback systems have also been used to provide feedback to the audience, or as

Method	Occurences
Graphical display	13
LED display	5
VR/3D display	2
Lasers	1

Table 3.7: Methods of providing visual feedback

part of the output of the instrument itself. Table 3.7 shows a list of the most common methods of providing visual feedback and the number of instruments which implemented them.

As Table 3.7 shows, the most common method is to use a graphical display, either through a monitor, touchscreen or projector. For instance the system described by Tanaka (2004) allows the performer to create music using a PDA, which provides visual feedback through its built-in screen. Projected visual feedback has been used in a range of instruments and systems, including a touch-screen based instrument by Bottoni et al. (2007), the Orbophone (Lock and Schiemer, 2006) (described in Section 3.3.4) and several developed by Levin (2005).

Levin (2005) also describes an interface using 3D/VR glasses to provide immersive visual feedback in Hidden Worlds. A VR display was also used by Mäki-Patola et al. (2005) in their virtual reality instruments, but using a CAVE (CAVE Automatic Virtual Environment), which is a projection-based virtual reality display.

LED (Light Emitting Diode) displays have been used in several instruments. The Beatbugs make use of several differently-coloured LEDS to present timing information about notes in the current phrase, as well as the status of parts of the BeatBug itself (Weinberg et al., 2002; Weinberg and Driscoll, 2005). A single colour-changing LED is used to provide status information in the SoundStone

(Bowen, 2005). The blocks which make up the Block Jam tangible instrument each contain a matrix of 16 LEDS which are used to indicate the function of the block to the performer(s) as well as providing feedback on the location and movements of the virtual "cue ball" (Newton-Dunn et al., 2003). As a last example of LEDs for visual feedback, the Tenori-on instrument has a matrix of buttons, each containing an LED. The LEDs light up when the buttons are pressed, indicating the status of that button within the grid (Nishibori and Iwai, 2006).

Finally, one of the more unusual methods found of providing visual feedback was that used by the Termenova (Hasan et al., 2002). The Termenova uses an array of red lasers, which are broken by the performers hands in order to create sound. To allow the performer to see the lasers and thus provide visual feedback on the pitches or effects to be played, a thin layer of theatrical mist is used. This mist shows the red colour of the lasers both to the performer and to the audience.

3.3.4 Additional Sonic Feedback

Additional sonic feedback involves a sound production system (either loudspeakers or headphones) which form a part of the instrument itself, where the aim is to produce only sound and not vibrotactile feedback to the performer. While in some cases vibrotactile feedback may be a side effect, it was not the main goal of the embedded sound production.

In several cases, speakers were added to instruments in order to create a completely integrated portable instrument. A prime example of this approach is the Tenori-on, which works as a totally integrated instrument, not requiring any additional computer hardware to make sound (Nishibori and Iwai, 2006). Similar aims have also led to the use of headphones on some instruments, such as those presented by Tanaka and Gemeinboeck (2006) and Schacher (2008).

In other cases, speakers have been embedded in instruments in order for the sound output to be localized to the instrument itself but with the synthesis happening elsewhere. Instruments which make use of this method include the Orbophone and the A20. The Orbophone is a collaborative instrument which senses movement in space around itself and projects video and audio from built-in video projector and speakers (Lock and Schiemer, 2006). The A20 a polyhedron-shaped tangible instrument with multi-channel audio output, through speakers mounted on each of it's faces (Bau et al., 2008).

3.3.5 Temperature Feedback

One instrument presented at NIME made use of a tactile feedback modality other than the vibrotactile feedback already described in Section 3.3.1. The Thermoscore system, developed by Miyashita and Nishimoto (2004), used Peltier devices on the keys of a piano to provide thermal feedback to the performer, based on information stored in a special score. The aim of this feedback was to convey "the feels existence, emotion, and 'body warmth' of the composer to the performer".

3.3.6 Additional Passive Feedback

While most instruments which wish to provide haptic feedback do so through an active force feedback system, it is also possible to provide additional passive haptic feedback to the performer. In the papers surveyed here, there were 3 examples of instruments providing this form of feedback. All 3 instruments made use of springs in order to do this. These were the MATRIX, an array of touch sensitive rods which provide haptic feedback by means of a spring connected to the bottom of each rod, which pushes against the performers hands (Overholt, 2001), the Tymbalimba (Smyth and Smith, 2002) an instrument with a mechanical interface which simulates the buckling action of the ribs of the cicada, and the G-Spring, a controller based around a large spring (normally used to open a garage door), which is bent and twisted by the performer using their hands (Lebel and Malloch, 2006).

3.4 Discussion

In Section 3.2, I examined the use of sensors in the 266 new digital musical instruments presented at the NIME conferences. While a wide variety of sensors were used by these different instruments, the ten most common sensors (shown in Table 3.3) represent over 65% of the sensor types used in these instruments.

Looking at the two most common sensors, the FSR (used in 26% of instruments) and the accelerometer (used in 21%), we find that these sensors are used in different ways and to perform different tasks across these instruments. For instance, the FSR is used to modify the current sound volume in the Bento Box (Hatanaka, 2003), an example of a parameter modulation task as described by Wanderley et al. (2000) or relative dynamical function as described by Vertegaal et al. (1996). Yet other instruments provide examples of this sensor used to control other classes of task, for example selecting the center frequency of a bandpass filter in the SCUBA (Cáceres et al., 2005), an example of a parameter selection task (Wanderley et al., 2000) or absolute dynamical function (Vertegaal et al., 1996).

Similarly for the accelerometer we see examples of its use to control both parameter selections, such as the amount of blending between two video clips in the Electronic Sitar (Kapur et al., 2004) and parameter modulations, such as modulating the stored velocity values in the Gyrotyre's MIDI score player mapping (Sinyor and Wanderley, 2005).

Also of note is the use of these sensors to sense different physical parameters. Again taking the accelerometer as an example, in both the Gyrotyre and Electronic Sitar controllers it is used as a tilt sensor (classified as a rotary position sensor by Vertegaal et al. (1996)). Yet, in the TGarden (Ryan and Salter, 2003) or the Hyperbow (Young, 2002) it is used as an acceleration sensor. Looking at the FSR, we can find it used as a continuous force sensor in the SCUBA, the T-Stick (Malloch and Wanderley, 2007) and the Metasaxophone (Burtner, 2002) and as a velocity sensitive trigger in the Beatbugs (Weinberg et al., 2002).

Overall, this shows that there is no single specific way of using many sensors. Not only can these sensors be used to sense different gestures, but they can be used to control different parameters. There is currently no standard method for deciding on the connection between sensors, gestures and musical tasks. The question then arises as to how we choose the best sensor for control of a specific task in a digital musical instrument. If we take specific classes of musical tasks and classes of sensors based on the physical parameter sensed, can we determine which class of sensor is most suitable for which class of task? This question forms the basis of the next chapter.

Examining the issue of the use of feedback in digital musical instruments, we see from Table 3.4 that only 15 instruments (representing less than 6% of instruments presented) offer any form of active vibrotactile feedback. Yet as previously mentioned, several authors have stated the importance of this feedback to the performer in establishing the feel of the instrument. Also, it should be noted that while several non-contact instruments were presented at NIME (instruments which lack tactile feedback even more than most DMIs), none of these instruments provide any active vibrotactile feedback. Yet, studies by OâĂŹModhrain (2000) and Rovan and Hayward (2000) have shown that such feedback can be extremely valuable to the performers of such instruments. Chapters 6 and 7 describe a number of digital musical instruments developed in the course of this research which include vibrotactile feedback. This includes both instrument-like vibrations in a DMI with a physical body and vibrations as state information in contact-less alternate instruments.

Finally, from Table 3.5 we can see that a number of different types of vibrotactile actuators have been used in those instruments which do provide vibrotactile feedback. This raises the question of which of these devices is most suited to providing this feedback and how we can evaluate them. Can these devices produce vibrations across the whole range of frequencies which human skin can sense? Are the amplitudes of these vibrations above the threshold of perception? How do these vibrations compare in frequency and amplitude to those of an acoustic instrument? Chapter 5 will deal with the provision of vibrotactile feedback and will attempt to answer these questions.

3.5 Conclusion

This chapter provided a detailed review of the design of the physical interface of digital musical instruments presented at the international conferences on New

Interfaces for Musical Expression since the initial workshop in 2001. This included a survey of 266 instruments presented in 577 papers and posters and examined the classes of the controllers for these instruments, their use of sensors and the provision of active feedback to the performer by these instruments. Overall this survey has shown that there is no consensus on the use of sensors for specific classes of tasks, and that many instruments are lacking in any active feedback which could be of much use in improving the feel of the instrument for the performer.

The remainder of this thesis deals with each of these issues. To begin this, the next chapter deals in detail with the use of sensors in digital musical instruments and describes a series of experiments to determine the optimum choice of sensor for a number of common tasks.

Chapter 4

Sensors

Sensors allow digital musical instruments to react to the performer's gestures. They convert physical energy into electrical form, which can then be measured and digitized by the computer. Sensors exist which can be used to measure any known physical parameter and can often do so with a range far beyond our human senses.

When designing a new digital musical instrument we can choose to make use of almost any performer gesture to control our instrument. Once we have decided on a gesture to measure, we are then often faced with a choice of sensors which can be used to measure aspects of that gesture. This chapter focusses on the place of sensors in digital musical instruments and more specifically on how to choose the optimal sensor (or sensor/gesture combination) to control a particular type of parameter in an instrument.

I begin with a discussion of the classification of sensors (based on that discussed in 2.4.1) and musical function and research into relationship between these classes of sensors and musical functions. This is followed by the description of a series of experiments which examine the suitability of sensors for specific musical tasks and a discussion of the application of the results to the design of digital musical instruments.

4.1 Sensors and Musical Function

As this chapter is concerned with the connection between sensors and musical function in digital musical instruments, it is useful to begin with a discussion of the classification of sensors and musical function.

4.1.1 Sensor Classification

As previously discussed in 2.4.1, there are numerous possible ways in which we can go about classifying sensors. One of the most common, which is often used in engineering literature is to classify sensors based on the physical property which they measure (e.g. visible light sensors, magnetic field sensors). Another possiblity is to categorize them based on the way in which the human interacting with them is influencing the world. That is, sensors are classified based on which of our physical communication channels, or *output modalities* they can be used to sense (Bongers, 2000). One problem with this particular method is that sensors can end up in numerous classes, as they can be used to sense several output modalities.

The first experiment described in this chapter makes use of the following classification, based on that presented by Vertegaal et al. (1996). In this case, sensors are classified based on the type of physical property sensed and the direction in which it is sensed. Further sub-division is achieved depending on the resolution of the sensing. For out purposes, this classification then provides 8 main classes of sensor:

1. Position sensors
2. Rotary position sensors
3. Velocity sensors
4. Rotary velocity sensors
5. Isometric force sensors
6. Isotonic force sensors
7. Isometric rotary force sensors
8. Isotonic rotary force sensors

The distinction between isometric and isotonic force in this clasiffication is based on whether or not movement is required. That is, isometric force sensors require no movement, whereas isotonic force sensors do require movement. This leads to one possible issue with this classification, in that if a sensor requires a movement then that sensor is more likely a position sensor, rather than a force sensor. In fact, devices which the authors classify as isotonic force sensors may more correctly be though of as position sensors *which implement (usually spring-based) force feedback*.

To remove this issue and to concentrate entirely on the sensing (rather than inherent feedback) qualities of the sensors, this chapter will use a simplified version of this classification, with the following classes:

1. Linear Position sensors

2. Rotary position sensors
3. Velocity sensors
4. Force sensors

4.1.2 Musical Function

Vertegaal et al. (1996) also defines three categories of musical function, based on the amount and type of change the parameter goes through. These categories are:

absolute dynamic functions change regularly over time and have values which are directly selected by the performer

relative dynamic functions change regularly over time but their values are modulated from a baseline by the performer

static functions change infrequently and often involve only limited options

Examples of *absolute dynamic* functions include pitch selection and amplitude selection. *Relative dynamic* functions include pitch bend and vibrato. Key selection and tuning are examples of *static* functions.

This classification of musical function can be applied to possible musical tasks to allow us to evaluate sensor and task combinations. Wanderley and Orio (2002) provide a list of possible musical tasks for evaluating sensors in digital musical instruments. These tasks include playing *isolated tones* at different pitches and loudnesses, *basic musical gestures* such as trills, glissandi and vibrato, and *musical phrases* such as scales, arpeggios and simple melodies. The result of applying the above classification to some of these tasks is shown in Table 4.1.

Task	Class
Note selection	Absolute Dynamic
Vibrato	Relative Dynamic
Scale Playing	Absolute Dynamic
Melody Playing	Absolute Dynamic

Table 4.1: Examples of classified musical tasks

It is also possible to join multiple single tasks together, creating a complex task (Wanderley and Orio, 2002). Examples of such complex tasks include playing a melody with vibrato on certain notes, or playing an arpeggio with glissandi.

4.2 Experiment 1: User Evaluation of Sensors for Specific Musical Tasks

As already stated, previous work has attempted to show a mapping between sensors and classes of musical task (Vertegaal et al., 1996). In that work, the authors classified sensors by the form of input that they sensed (linear position, rotary position, isometric force etc.), the resolution of this input sensing and the types of intrinsic feedback provided by the sensor. They also classified musical tasks by the range and form of input they required (static, absolute dynamic and relative dynamic). From these classifications, they proposed a mapping of the suitability of specific classes of sensors for specific classes of task.

This section will discuss an experiment performed to evaluate the suitability of a range of sensors for specific musical tasks. The experiment described here makes use of a modified version of the categorisations provided by Vertegaal et al. (1996) and attempts to evaluate empirically whether the mapping from sensor type to musical task proposed by Vertegaal et al. (1996) holds.

To allow for this evaluation, I look at user preference ratings for sensors when performing basic musical tasks. The hypothesis is that for some tasks, certain classes of sensor will be easier to use than others and so will receive higher preference ratings from users.

4.2.1 Participants

A total of 11 participants took part in this experiment. The participants were all graduate students in Music Technology and their areas of specialisation ranged from acoustics and physical modelling to interaction design to music information retrieval. Eight of the participants had extensive musical instrument training, while the remainder either did not play an instrument, or had previously played for a period of less than two years and had since stopped. Five participants had experience of playing electronic instruments, whether software or hardware in form.

4.2.2 Design and Materials

The experiment examined the use of specific sensors for specific musical tasks. In total 5 sensors were examined for 3 tasks. The sensors used (and their classification) are shown in Table 4.2. The tasks, based on those suggested by Wanderley and Orio (2002) and classified based on Vertegaal et al. (1996), are shown in Table 4.3. The first 2 tasks (melody playing and vibrato) are simple tasks, while the last task (melody with vibrato) is a complex task.

Participants used each sensor for each task. Tasks were performed in order, but sensor use was randomized within each task. Each sensor was presented attached

Sensor	Sensor Class
Linear potentiometer (fader)	Linear position
Rotary potentiometer	Rotary position
Linear position sensor (ribbon controller)	Linear position
Force sensing resistor	Force
Bend sensor	Rotary position

Table 4.2: List of sensor devices used in the experiments and their associated classes

Task	Class
Melody Playing	Absolute Dynamic
Vibrato	Relative Dynamic
Melody with Vibrato	Complex

Table 4.3: Classified musical tasks

to the table in front of the participant. The participant manipulated the sensor with their primary hand, using their secondary hand to press the spacebar key of a computer keyboard, which caused the system to output sound.

The signal from each sensor was read using an Ethersense analog to digital converter. This sampled the sensor input at a rate of 500Hz and with a 16-bit resolution. This converter was connected using an ethernet cable to an 17-inch Apple PowerBook. For the melody task, the output of the sensor was mapped to a one octave frequency range, subdivided in semitones. For the vibrato task, a portion of the sensors range was mapped continuously over a range of +/- 1 semitone. Finally for the complex task, the sensor range was again mapped over one octave subdivided in semitones, allowing the participants to both play notes and modulate the frequency by +/- 1 semitone. Synthesis was performed in Max/MSP, using a simple waveshaping synthesis system based on Chebychev equations[1].

Participant ratings of each sensor for each task was gathered in Max/MSP using a single on-screen slider. This slider allowed the participants to rate each sensor for each task in a range of 0 - 127, which represented a range of ease of use from *Very Difficult* to *Very Easy*. A video recording was also made containing the interaction with the system and the audio from the system and user themselves, to allow for later analysis.

4.2.3 Procedure

Subjects arrived at the lab and were given an Information/Consent form to read over and sign. Subjects were shown the experimental interface and told that they would be attempting to perform 3 different tasks on this interface using 5 different sensors and asked to rate the ease of use of each sensor for each task. Each sensor was explained to them in turn, to ensure they understood how the sensors worked. They were also informed that we were testing the sensors, not the participants themselves, and that any difficulties performing the tasks would be due to the sensors.

The tasks took place in order. Within each task, the participant was presented with the sensors in a randomized order and asked to attempt to perform the task with each sensor. Each attempt was considered complete when the participant decided that they had performed the task sufficiently well or that they would be unable to perform the task with that sensor. Participants were given a 5 minute break between each task and shorter breaks between each sensor.

---

[1] The synthesis patch used was the waveshaping demonstration patch *cheby.pat* supplied as an example with Max/MSP 4.5

Finally, participants were debriefed verbally after each task was complete and asked to comment on any particular strengths and weaknesses of the sensors for that task.

4.2.4 Data Analysis

Results were analyzed using the Statistical Package for the Social Sciences software (SPSS). Analysis was performed using a 3 × 5 (tasks × sensors) factorial ANOVA with pairwise comparisons performed using Tukey's Honestly Significant Difference (HSD) to determine specific significant differences. Before the ANOVA analysis all outliers were removed and the data was checked for normality using the Shapiro-Wilk test.

4.2.5 Results

A number of significant effects were found. Firstly, there was a significant effect of sensor on the ease of use ratings [$F(4,40) = 26.74$, $p < 0.001$]. This means that certain sensors were rated higher or lower independant of the task being performed. In particular the bend sensor was rated significantly lower than all other sensors, independant of the task [Tukey's HSD, $p < .05$ for all comparisons] and the Linear Position Sensors was rated signficantly higher than the other sensors [Tukey's HSD, $p < .05$ compared to all sensors].

There was also a significant effect of the sensor × task interaction [$F(8,80) = 9.65$, $p < 0.001$]. This implies that certain sensors receive different ratings depending on the task. For the melody task, the linear position sensor rated significantly better than all other sensors [Tukey's HSD, $p < .01$ compared to bend and FSR, $p < .05$ compared to linear and rotary potentiometers]. This was also true for the complex task [Tukey's HSD, $p < .01$ compared to bend and FSR, $p < .05$ compared to linear and rotary potentiometers]. For the vibrato task, the FSR and the linear position sensor rated significantly higher than the other sensors, with no significant differences in rating between them [Tukey's HSD, $p < .01$ for both FSR and linear position sensor compared to all other sensors]. Figure 4.1 shows the mean ratings for each sensor across the tasks.

Examining the vibrato task in more detail, some participants rated the linear position sensor as much easier to use than others. After examining the video footage I found that participants could be divided into 2 groups, depending on the method they used to perform the vibrato using the linear position sensor. Those who performed it using two fingers to hit two different notes (as with alternating keys on a keyboard to produce a trill) rated the linear position sensor as easier to use than the FSR [6 participants, $M_{FSR} = 68.2$, $M_{LPS} = 79.8$], whereas those who slid or rolled their finger on the sensor (similar to performing vibrato on a stringed instrument) rated the linear position sensor lower than the FSR [5 participants, $M_{FSR} = 69.8$, $M_{LPS} = 62.0$]. However, while this effect was noticeable, it was not statistically significant [$F(1,9) = .47$, $p = .51$].

Overall, these results would seem to indicate that certain sensors are more useable for the control of certain musical functions. However, the difference in rating between sensors of the same class (e.g. the linear position sensor and linear potentiometer) indicate that the classification system being used does not completely explain the preferences of the participants. This difference in preference could be due to a number of factors, including the facts that the linear potentiometer requires more effort to move, exhibits some inertia and stays in position rather than

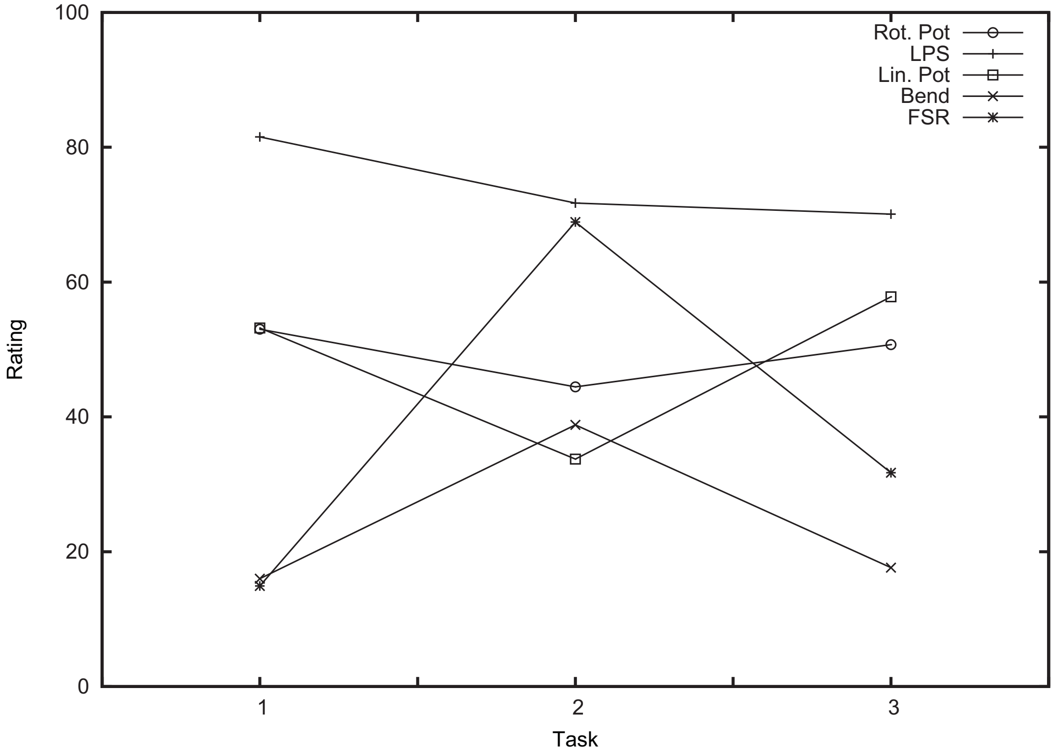

Figure 4.1: Mean ratings for each sensor across tasks.

returning to zero when released.

None the less, we can draw some guidelines from these results which may be used to inform the design of new digital musical instruments. In particular it seems that force sensors receive high usability ratings for modulation of parameters (relative dynamic tasks) and that position sensors (whether linear or rotary) are easiest to use for value selection (absolute dynamic tasks). However, the fact that participants used different techniques for the parameter modulation task with the linear position sensor and that the preferred sensor for that task depended on which method was used would indicate that not only the choice of sensor but the combination of sensor and gesture effects users' overall preference for a given combination of sensor and task. The next section presents an experiment which examines this issue for the performance of a parameter modulation task.

Also of interest is that the choice of sensor for the combined task appears to be dominated by the melody portion of the task, rather than the vibrato. While the FSR and Bend sensors both show higher ratings for the combined task compared to the melody task, the ratings for the combined task are still closer to those of the melody task than the vibrato task.

4.3 Experiment 2: Sensors, Gestures and Musical Experience

This experiment aimed to examine a number of issues which arose from the previous experiment. Firstly, user preference ratings for the parameter modulation task with the linear position sensor depended on the gesture which the participant used to control this sensor. This experiment examined two different methods of interacting with this sensor and compared them with the FSR for the same task. It also examined the effect of previous musical experience on user preference ratings for different sensor and sensor/gesture combinations.

This experiment tested parameter modulation using different sensors and gestures. To examine the effect of previous musical experience on the results, all participants were pianists and violinists. Participants were asked to perform a pitch modulation using 3 different methods, the *pressing method* which involved increasing and decreasing pressure on an FSR, the *sliding method* which involved sliding there finger back and forth along a section of a linear position sensor and the *rolling method* which involved rolling their finger back and forth on a section of the linear position sensor. These final two methods (the sliding and rolling methods) were both used by participants in the first experiment when performing with the linear position sensor. Note that the most popular method from the first experiment (tapping with two fingers) was not used in this experiment as it does not modulate the pitch but rather alternates between two seperate pitches.

The hypothesis for this experiment was that violinists would perform better than pianists overall since they presumably have much more experience with the performance of pitch modulation (through vibrato). I also hypothesized that violinists would both prefer and perform better with the rolling method since it is most similar to the way they produce vibrato on a violin. For this experiment a single objective measure of performance was used, which was a measure of the stability of the speed of modulation. This allows for comparisons between user's subjective rating and objective performance which may also prove useful.

4.3.1 Participants

27 right-handed musicians with at least 8 years of musical experience on their instrument were recruited from McGill University and paid CAD$10. Handedness was assessed using the Edinburgh Handedness Inventory (Oldfield, 1971). Nine subjects were dropped from the final analysis because an equipment error rendered their data unusable. Of the remaining 18 subjects there were 9 pianists and 9 violinists.

4.3.2 Design and Materials

The experiment followed a mixed design. The between subjects factor was instrument played (2 levels, piano or violin) and the within subjects factor was the method of modulation (3 levels, sliding, rolling, pressure).

For the *pressure method*, subjects used an FSR. Increased finger pressure on the sensor raised the pitch while a decrease in finger pressure lowered the pitch back down. Modulations were produced by applying pressure to the sensor in a pulsating motion. This means that the pitch could only be modulated between the base pitch and a higher pitch using this method. This is similar to vibrato on fretted stringed instruments, such as the guitar. The other methods could both raise and lower the pitch allowing for a modulation more like that of fretless string instruments, wind instruments or the human voice. The FSR was mounted flat on a small block of wood without any padding.

For the sliding and rolling method, subjects used a linear position sensor. A light finger pressure is required to activate the sensor. Moving the position of ones finger to the right causes pitch to increase, while movement to the left causes pitch to decrease. For the sliding method subjects produced modulations by sliding their finger back and forth across a limited portion of the sensor (about 3cm wide). Too large a spacing would have made it difficult for the subject to maintain the modulations at the speeds we were looking for, while too small a spacing would have made the modulations harder to control. The spacing we chose was determined on the basis of pilot testing in our laboratory.

Scaling on the sensor was changed for the rolling method from 3cm to < 1cm. Subjects here could produce modulations by simply pivoting their finger back and forth at a demarcated point on the sensor (similar, though not identical to, the way vibrato is produced on a violin). For all 3 methods, subjects alternated using the index finger of either their right or left hand, depending on the conditions currently being tested. The sensor was placed at elbow level to the subject who's forearm rested on the table at approximately a 90 angle to the upper arm.

The experiment was run on a 17-inch Macintosh Powerbook G4. The sensor output was converted to a computer-usable format using an AVR-HID analogue-to-digital converter, at a rate of 100Hz and with a 10-bit resolution (Marshall, 2005). The visual programming environment Max/MSP was used to map the signal onto a musical output. The musical tone that subjects were able to modulate was created with the sound synthesis software Tassman 4[2]. The sound used was preset tone #44: Simple Sine Lead, a straightforward sine-wave tone, modified to have no built in pitch variation, set to a Middle G (196Hz). Subjects produced the tone by holding down on the spacebar of the laptop with one hand (using the opposite hand to modulate the tone). They were able to hear the tone through the laptop speakers which were set to a comfortable volume. At the beginning of the experiment subjects were informed that they would be producing pitch modulations using 3 different methods. They were told for each modulation they produced to try and maintain as constant a *rate* as possible.

To calculate the frequency of the modulations over time (as it changes over the course of the recording), an additive analysis was performed from the short-time Fourier transform of the measured control signal. Only one peak was selected in the 1-7Hz range, using 192 sample Blackman-Harris 92dB windows, with a hop

---

[2] Produced by Applied Acoustics Systems. See http://www.applied-acoustics.com/tassman.htm for more information.

size of 16 samples. The “instantaneous” frequency f_c was then obtained at 7Hz and adequately represented the periodicity of the control signal, as the only remaining components of the additive analysis were at multiples of f_c (i.e. the signal was harmonic with fundamental frequency f_c) and had much lower amplitudes (i.e. contributing mainly to the control curve shape, which was between sinusoidal and triangular, but not changing the modulation centre or depth). This algorithm was adapted from the 2nd order sinusoidal model (Marchand and Raspaud, 2004). From this was calculated a participant’s mean speed for a specific modulation, and their standard deviation from that mean.

The presentation order of the different methods for both trials was randomized with a latin-square design. This allows a different presentation order for each of the participants.

While all subjects tested were right-handed, right-handed violinists traditionally produce vibrato with their left hand. Pianists, however, finger and play notes with both hands and presumably have little to no experience with vibrato, so we should expect them to both perform better with and prefer using their right hand. All subjects were tested with both hands to see if there was really any difference.

Subjects were instructed to create modulations at both slow and fast speeds. A recorded sample of a slow and a fast modulated pitch (produced using the linear position sensor with the sliding method) was played to subjects a few times at the beginning of the experiment to demonstrate what we meant by a slower or faster modulation. The aim was not to have subjects match the speeds precisely, but to obtain generally similar slow and fast modulations throughout the experiment. The order of left-right slow-fast modulation tasks was also randomized throughout the experiment. See Table 4.4 for a sample experimental design.

Trial 1			Trial 2		
Method	Hand	Speed	Method	Hand	Speed
Sliding	Left	Slow	Rolling	Left	Fast
		Fast			Slow
	Right	Slow		Right	Fast
		Fast			Slow
Rolling	Right	Fast	Pressure	Left	Fast
		Slow			Slow
	Left	Slow		Right	Slow
		Fast			Fast
Pressure	Left	Slow	Sliding	Right	Slow
		Fast			Fast
	Right	Slow		Left	Fast
		Fast			Slow

Table 4.4: Sample order of task presentation.

4.3.3 Procedure

Subjects arrived at the lab and were given an Information/Consent form to read and sign. Subjects were shown the experimental interface and told that they would be producing modulations on this interface using 3 different methods. They were played the slow and fast modulation samples and told to try and maintain a uniform rate for each modulation they produced. In the first trial subjects were introduced to the sliding, rolling and pressure methods. They were verbally and visually instructed on how to perform each method and then given up to a minute to practice before we began recording their output.

After going through all three methods once, subjects were then given the Queens Musical Background Questionnaire (Cuddy et al., 2005) and the Edinburgh Handedness Inventory (Oldfield, 1971) to verify their musical background and handedness. No subjects were discarded. The second trial was exactly the same as the first except the order of the tasks was counterbalanced. After each method subjects were given a questionnaire which asked them to rate how difficult the method was, whether they preferred it for slow or fast modulations, whether they preferred it with their right or left hand[3], and how they preferred it in comparison to the other methods. They were also given an opportunity to add any comments they might have with regard to the method in question. After the second trial was completed subjects were given monetary compensation and debriefed as to the nature of the experiment.

4.3.4 Data Analysis

Results were analysed using the Statistical Package for the Social Sciences software (SPSS). The subjective questionnaire data and objective measurements were analysed separately at first. Correlation analyses were then performed between the two sets of data. Both sets were analysed using repeated measures ANOVAs with t-tests for specific relevant comparisons. Post-hoc tests were performed using Tukey's Honestly Significant Difference (HSD). Before the ANOVA analysis all outliers were removed and the data was checked for normality using the Shapiro-Wilk test.

4.3.5 Results

Questionnaire Data

Ease of Use and Preference ratings were in most cases very similar, suggesting that participants did not really distinguish between the two measures. The overall

---

[3] For hand preference, participants were asked to separately rate each hand on a Likert scale. These values were later translated to a single scale ranging from 1 for left to 9 for right.

differences in preference ratings between the violinist and pianist groups were marginally significant [$F(1,21) = 3.39$, $p = .08$], while there was no significant difference between instrument groups for the ease of use rating [$F(1,21) = 0.49$, $p = .49$]. The violinist group rated the pressure method significantly lower than the pianists [$t(21) = 3.06$, $p < .01$], the sliding method marginally lower [$t(21) = 1.95$, $p = .07$], and the rolling method insignificantly higher [$t(21) = -.73$, $p = .48$]. This does not fully confirm the hypothesis, which was that violinists would prefer the rolling method to pianists, nevertheless the trend is in the right direction.

There was a significant effect of method used across all subjects [Preference: $F(2,42) = 7.34$, $p < 0.01$, Ease of Use: $F(2,42) = 5.38$, $p < .01$]. Post-hoc tests indicated that users rated the pressing method significantly higher than the other methods in terms of both ease of use [Tukey's HSD, rolling: $p < .05$, sliding: $p < .05$] and preference [Tukey's HSD, rolling: $p < .05$, sliding: $p < .01$]. No significant differences were found between the sliding and rolling methods for either ease of use or preference ratings.

Looking at the instrument groups individually, the pianist group preferred the pressure method, followed by the sliding method and then the rolling method (see Figure 4.2a) [Preference: $F(2,24) = 5.95$, $p < .01$, Ease of Use: $F(2,24) = 4.30$, $p < .05$]. Post-hoc tests show that the pressing method was rated significantly higher than the sliding and rolling methods in terms of both ease of use [Tukey's HSD, sliding: $p < .05$, rolling: $p < .05$] and preference ratings [Tukey's HSD, sliding: $p < .05$, rolling: $p < .01$]. Once again, no significant differences were found between the sliding and rolling methods for either rating.

The violinist group preferred the pressure method, followed by the rolling method and then the sliding method (see Figure 4.2b). In this case the differences

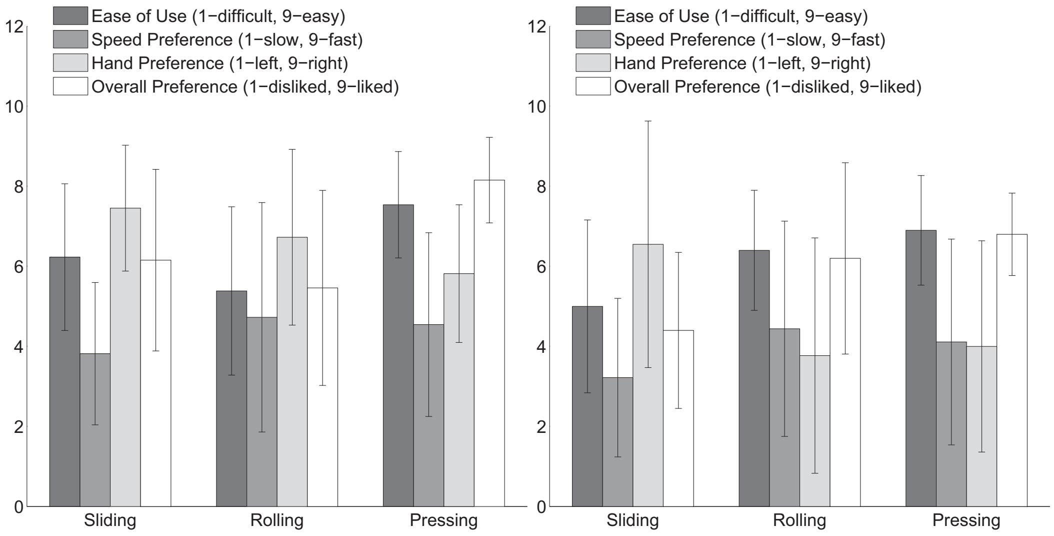

(a) Pianists (b) Violinists

Figure 4.2: Mean questionnaire responses.

were marginally significant [Preference: $F(2,18) = 3.79$, $p < .05$, Ease of Use: $F(2,18) = 3.50$, $p = .05$]. Post-hoc tests indicated only that the sliding method was rated significantly lower than the other methods in terms of both preference [Tukey's HSD, pressing: $p < .05$, rolling: $p < .05$] and ease of use [Tukey's HSD, pressing: $p < .05$, sliding: $p < .05$].

Pianists show significantly greater ratings for their right hand than violinists [$F(1,18) = 10.29$, $p < .01$]. Speed preference hovers generally around the middle with a slight preference in both groups on all the methods for slower speed modulations [average rating: 4.17]. There were no significant differences in speed preference between violinists and pianists [$F(1,18) = .38$, $p = .55$] or between the different methods [$F(2,38) = 1.31$, $p = .28$]. However both groups rated slower speed modulations slightly higher with the sliding method. This makes sense as the finger moves over a somewhat larger distance with this method.

Recorded Data

The standard deviation dependent variable measures the extent to which modulation samples deviate from a constant rate. Repeated measures ANOVAs were performed with instrument as the between groups factor. No significant difference between groups was found for the fast modulations [Left Hand: $F(1,12) = .71$, $p = .42$, Right Hand: $F(1,13) = .01$, $p = .93$], slow modulations [Left Hand: $F(1,15) = .51$, $p = .49$, Right Hand: $F(1,14) = .34$, $p = .57$], or on an overall average of each group's modulations [Piano: $F(2,20) = .41$, $p = .63$, Violin: $F(2,16) = 1.18$, $p = .33$].

There was also no significant overall difference in precision between left hand and right hand modulations in both the pianist [Fast: $F(1,12) = 0.06$, $p = 0.82$, Slow: $F(1,13) = 0.31$, $p = 0.59$] and violinist [Fast: $F(1,13) = 2.28$, $p = 0.16$, Slow: $F(1,16) = 1.71$, $p = 0.21$] groups.

Fast modulations were significantly less precise than slow modulations for both pianists [Left Hand: $F(1,13) = 15.70$, $p < .01$, Right Hand: $F(1,12) = 4.99$, $p < 0.05$] and violinists [Left Hand: $F(1,13) = 13.80$, $p < 0.01$, Right Hand: $F(1,15) = 24.86$, $p < .001$]. There were also significant overall differences in mean modulation frequency for each method [$F(2,32) = 34.623$, $p < .001$]. The rolling method was performed the fastest at an average frequency of 3.99Hz, while the sliding and pressing methods were played significantly slower at a frequency of 2.65Hz and 2.82Hz respectively [Tukey's HSD, $p < 0.01$ for both sliding and pressing]. An analysis using Pearson's correlation coefficient reveals significant correlations between the mean frequency and frequency deviation on each method [Sliding: $r(16) = .77$, $p < .001$, Rolling: $r(16) = 0.73$, $p < .001$, Pressure: $r(15) = 0.84$, $p <$

.001]. This is in keeping with the negative relationship found between modulation speed and precision.

There was a significant difference in precision based on the method used [F(2,20) = 55.84, $p < .001$]. The rolling method was significantly less precise than the sliding method [Tukey's HSD, $p < .01$] and the pressing method [Tukey's HSD, $p < .01$], with no significant difference between the precision of the pressing and sliding methods. There was no effect of instrument on precision [F(1,10) = .06, p = .81], nor of an interaction between instrument and method [F(2,20) = .70, p = .51].

While there are no significant correlations between method preference and method precision, some slight trends are apparent when examining the data. In the piano group, the rolling method is both preferred less and performed less precisely than the other two methods (see Figure 4.3a). This suggests a relationship between preference and performance precision. However when we look at the data from the violin group this relationship disappears (see Figure 4.3b). Even though the rolling method is performed less precisely than the sliding method, it is preferred more. It is likely here that another factor besides performance precision has influenced method preference.

4.3.6 Discussion

Overall, results show that participants prefer the pressing method and that this is true of both violinists and pianists. While violinists did exhibit slightly higher preference ratings for the rolling method when compared to the pianists, the effect is too small to allow for a rejection of the null hypothesis. Moreover there is no evidence that violinists produce modulations any more precisely than pianists.

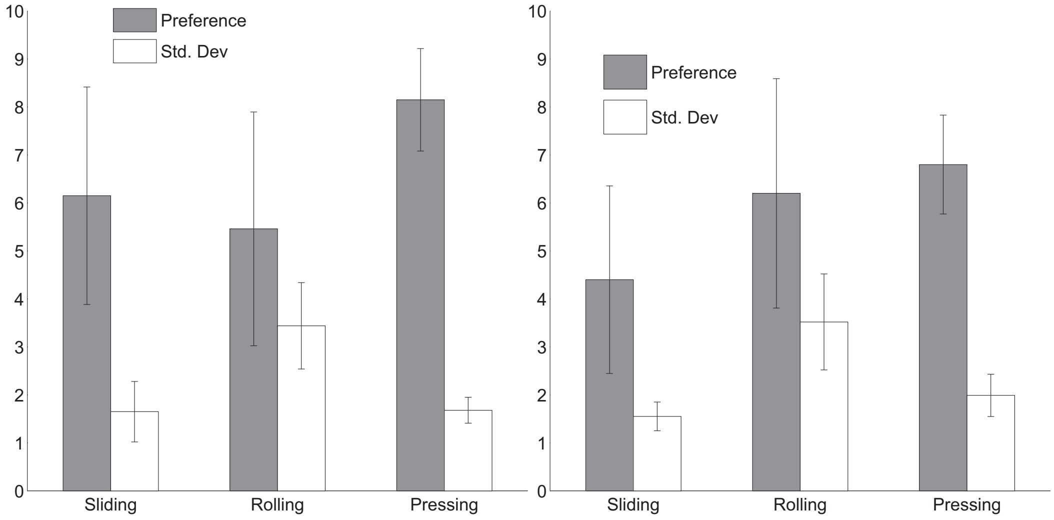

(a) Pianists (b) Violinists

Figure 4.3: Subjective preference compared to deviation from the participants' mean achieved frequency (scaled for comparison).

Therefore it is not possible to directly conclude that preference for (or performance on) novel mappings is influenced by previous learned musical experience.

Conversely, it is not possible to conclude the opposite, that previous learned musical experience has no effect on preference or acquisition of new musical mappings. It may be that the mapping used was in reality not all that similar to the traditional method of violin vibrato production. Violinists typically hold their instrument up to their chin, with their arm curved upwards and their hand perpendicular to their body. In contrast, subjects here were asked to perform the task with their hand flat on an elbow high table. A number of the violinist subjects even asked if they could hold the sensors in the manner they would normally hold a violin, but were not allowed to do so.

The fact that pianists show a slight correlation between precision and preference (see Figure 4.3a) suggests they relied to a certain extent on a subjective assessment of their own performance in making their ratings. This correlation was absent in violinists, which means that something else may have been informing their preference ratings. Perhaps the violinists' previous vibrato experience has, in some indirect way, produced an effect on their method preference. Although it is difficult to discern from this data the nature of the effect.

As can be seen in Figure 4.2b, violinists gave a lower preference rating to the sliding method than the rolling method, despite the fact that their sliding method modulations were more precise. When examining the data more closely, both the pianists and the violinists perform significantly more precisely on the sliding method with their right hand. It is the only method where this is the case, and it only occurs for fast modulations. We can assume that right-handed pianists (or any right-handed control) would expect to perform better with their right hand. However this may not be the case with right-handed violinists. These musicians have extensive experience creating vibrato with their left hands and therefore might expect to perform better with their left hand. During the experiment, a number of violinists expressed hesitation when asked to produce modulations using their right hand. Some were doubtful they could perform accurately. Perhaps this resulted in them giving lower ratings to this method.

The factor most clearly responsible for mediating performance (and possibly in turn preference) was the speed of the modulation. Fast modulations were found to be significantly less precise than slow ones. Moreover there were significant differences in mean modulation speed between the different methods, and these differences were significantly correlated with differences in accuracy. Modulation speed therefore represents a serious potential confounding variable. Differences in preference and performance between the different methods may primarily be the result of how fast the modulations were played on each method.

In this experiment the three methods were played at different speeds despite the fact that subjects were told to try and maintain consistent speeds throughout all experimental trials. While some random variance is to be expected, the fact that subjects averaged faster speeds for some methods and slower speeds for others suggests they felt more comfortable performing at these speeds. Just as different instruments naturally elicit their own specific set of movements and gestures, these three methods naturally elicit modulations at specific frequencies. From a holistic perspective modulation frequency on a given method can even be seen as an integral component of the mapping itself, a function of the interface's design.

Nevertheless modulation speed is something that should be controlled for, it just needs to be done carefully. Pilot testing would be necessary to determine an experimental speed that works well for each method being evaluated. There is no sense in testing each method with a modulation frequency of 7Hz if two of the methods only work well at about 5Hz. In addition the results of any such experiment would need to be interpreted cautiously. Even if a method tests very well at 3Hz it may not be suitable if the intent is to mimic violin vibrato, as this type of vibrato is traditionally produced at a speed of 5-7Hz (Papich and Rainbow, 1974).

From this we can see the possibility of examining subjects performance of modulations at a variety of different fixed speeds. This would allow us to determine whether speed is a factor in the production of modulations with each of these different gesture mappings. This question forms the basis of the third experiment, described in the next section, which is centred on gaining objective measures of the performance of pitch modulation using each gesture mapping, at a variety of fixed speeds.

4.4 Experiment 3: Objective Measurement of Performance

The third experiment examined the ability of participants to perform modulations at different speeds using each method of production. In particular the aim was to improve upon the methodology of the previous studies by including multiple measures of objective performance. In addition to measuring precision as in the previous study, measures of accuracy and modulation depth were also performed to provide a more thorough objective assessment.

This experiment also offered an opportunity to validate some of the results from the previous experiment, while controlling more carefully for the effects of speed. As was noted in the previous experiment, the speed of the modulations influenced precision far more than any other variable. Simply playing subjects a sample speed and asking them to attempt to mimic the speed was not enough to control for this. In this experiment subjects were asked to attempt to match the speed of a sample modulation which was being played to them at the same time. The main interest was in discovering whether or not the higher precision of the sliding and pressure methods relative to the rolling method held with more stringent control on speed. Also of interest was whether the effects were robust enough to hold up under a range of speeds.

4.4.1 Participants

There were 10 participants in this study, each of whom was compensated CAD$10 for taking part. All participants were musicians with at least 8 years of experience on their instrument. There were no restrictions as to the instruments played by the participants. Once again all participants were right-handed.

4.4.2 Design and Materials

This experiment follows a within-subjects design. The factors examined are the method of modulation (3 levels, sliding, rolling, pressure) and speed of modulation (6 levels, 1 Hz, 2Hz, 3Hz, 4Hz, 5Hz, 6Hz).

The sensor setup used by the participants to perform each method was the same as those described in the previous study. The experiment was run using a 15-inch MacBook Pro computer. As a result of the equipment problems for some subjects during the first study, the sensor output was converted to a computer-usable format using an Electrotap Teabox sensor interface (Allison and Place, 2005), which offers a higher resolution and faster sampling rate than the interface used in the first study. Data was sampled at 4000Hz, with a 12-bit resolution. This was connected to the computer using an S/PDIF optical cable.

Max/MSP was once again used to process the incoming sensor data. It was also used to produce the sound output, using a simple sine wave generator. However, the patch created in Max/MSP was also able to produce sample signals at each of the speeds which were to be tested. These signals were also output as a sine tone, but at a fundamental frequency which was a perfect fifth higher than that of the sound output from the participants sensors. Again, subjects produced the tone by holding down on the spacebar of the laptop with one hand while using the opposite hand to modulate the tone.

At the beginning of the experiment subjects were informed that they would be producing pitch modulations using 3 different methods of production at 6 different speeds. They were told for that for each modulation they were to produce there would be a sample playing, the speed of which they would attempt to match.

The sensor input was sampled in Max/MSP at a rate of 8000 Hz and recorded to an audio file for later processing. Processing was performed in Matlab, using the same algorithm as was used in the first experiment. However, due to the higher sampling rate of the recorded signals, some parameters were changed and additional filtering was added. Firstly, the data was low-pass filtered using a 4th-order Butterworth low-pass filter with an f_c of 100 Hz to remove any high-frequency noise before processing. It was then resampled to a sampling rate of 2000 Hz. For the analysis we used a window length of 4196 samples and a hop size of 100 samples, resulting in the output being determined at a rate of 20 Hz. The algorithm then determined and recorded the mean modulation frequency produced, the standard deviation from this mean (as a measure of precision), the deviation of the signal from the target frequency (as a measure of accuracy) and the RMS amplitude of the signal (which gives a measure of the depth of modulation).

It is important here to note the difference between the precision and the accuracy of the modulation. The accuracy of the modulation is a measure of how close to the target speed the performance was. This is measured as the deviation from the target frequency. The precision, on the other hand, is a measure of how much variation occurred over the course of the production. This is determined as the deviation of the signal from the mean speed the participant produced (i.e. the standard deviation). It should also be noted that as we are calculating deviations the result is actually the inverse of the precision and accuracy. This means that a more accurate or precise performance will result in a lower deviation (from the target and mean frequencies respectively).

As with the previous study, the order of the presentation of the combination of methods and speeds was randomized throughout the experiment.

4.4.3 Procedure

Subjects arrived at the lab and were given an Information/Consent form to read over and sign. Subjects were shown the experimental interface and told that they would be producing modulations on this interface using 3 different methods. They were also told that in each case the system would produce a modulated signal and that they would have to try to match as close as possibly the speed of that modulation. A sample was played for them at this time so that they would know what to expect.

Subjects were introduced to the 3 methods of modulation and were then verbally and visually instructed on how to perform each method. They were then allowed up to a minute to practice before we began recording their output. After completing two of the three trials, the subjects were given the Queens Musical Background Questionnaire Cuddy et al. (2005). Upon finishing the experiment, subjects were given monetary compensation and debriefed as to the nature of the experiment.

4.4.4 Data Analysis

Results were analysed using SPSS. All data was analysed using a 3 × 6 (Method × Speed) factorial ANOVA, with Post-Hoc tests performed using Tukey's HSD. Before the ANOVA analysis all outliers were removed and the data was checked for normality using the Shapiro-Wilk test.

4.4.5 Results

Accuracy

The largest significant factor found for the accuracy of performance was the speed [$F(5,35) = 39.83$, $p < .001$]. This factor accounted for most of the variance in accuracy ratings [$\eta^2 = .66$]. Post-hoc tests showed that each speed is significantly more accurate than any higher speed [Tukey's HSD, $p < .05$ in all cases].

Significant effects were also found for the method used [$F(2,14) = 38.07$, $p < .001$] and the method × speed interaction [$F(10,70) = 3.95$, $p < .001$]. Specifically, the pressing and rolling methods were both significantly more accurate than the sliding method [Tukey's HSD, $p < .01$ for both methods], but were not significantly different from each other. Post-hoc tests on the interaction showed significant differences between low speed (≤ 3 Hz) modulations using any method and high speed (> 3 Hz) modulations using any method [Tukey's HSD, $p < .05$].

Figure 4.4 shows the mean achieved frequency versus the target frequency for each method. Note the increased distance from the target line for each method as the target frequency increases.

Examining the effect of method at each speed separately showed that at low speeds there was no significant difference between the accuracy of each method.

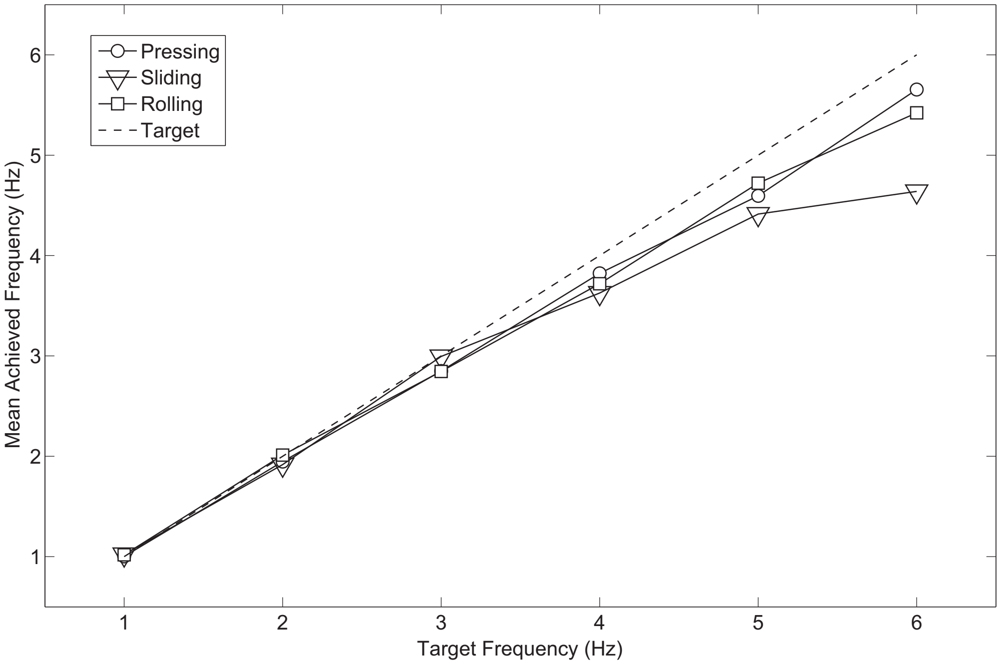

Figure 4.4: Mean achieved frequency versus target frequency for each method

At high speeds however (> 3 Hz), both the pressing and rolling methods prove to be significantly more accurate than the sliding method [Tukey's HSD, $p < .05$ for each comparison]. This means that using these methods the participants' achieved speeds were closer to the target speed than when using the sliding method. There were no significant differences found between the pressing and rolling methods at any speed.

Precision

The largest significant factor found to effect precision is again that of speed [F(5,35) = 55.69, $p < 0.001$, $\eta^2 = 0.75$]. Again, post-hoc tests showed that performance at any speed was significantly more precise than performance at higher speeds [Tukey's HSD, $p < .05$ for all comparisons]. We also found significant effects of method of modulation [F(2,14) = 9.24, $p < 0.01$]. Interestingly, in this case the sliding method proved to be significantly more precise than either the pressing or rolling methods [Tukey's HSD, pressing: $p < .01$, rolling: $p < .01$]. Once again the scores for the pressing and rolling methods were not significantly different. There was also a significant effect of the method × speed interaction [F(10,70) = 3.12, $p < 0.01$]. As with the accuracy measurement, the method × speed interaction indicates that low speed modulation using any method is more precise than high speed modulation using any method [Tukey's HSD, $p < .05$ for all comparisons].

There was no significant difference between any of the methods when controlling low speed (≤ 3 Hz) modulations. Above 3 Hz there were significant differences between the precision of the sliding method and that of the other two methods. In this case however, the sliding method is significantly more precise than the other methods [Tukey's HSD, $p < .05$].

Figure 4.5 shows the deviation from the mean achieved frequency for each method at each of the target frequencies. A lower deviation indicates a higher level of precision. It can clearly be seen that at higher frequencies the sliding method is more precise than the other methods [$p < 0.01$].

Modulation Depth

Looking at the modulation depth showed a significant effect of speed [F(5,35] = 10.39, $p < .001$] and of the method × speed interaction [F(10,70) = 2.09, $p < .05$]. There was also a marginally significant effect of method [F(2,14) = 4.20, $p = .08$].

Post-hoc tests showed that higher speed modulations generally had a lower depth than low speed modulations. Modulations at frequencies of ≤ 3 Hz had a significantly higher depth than those of > 3 Hz [Tukey's HSD, $p < .05$ for all

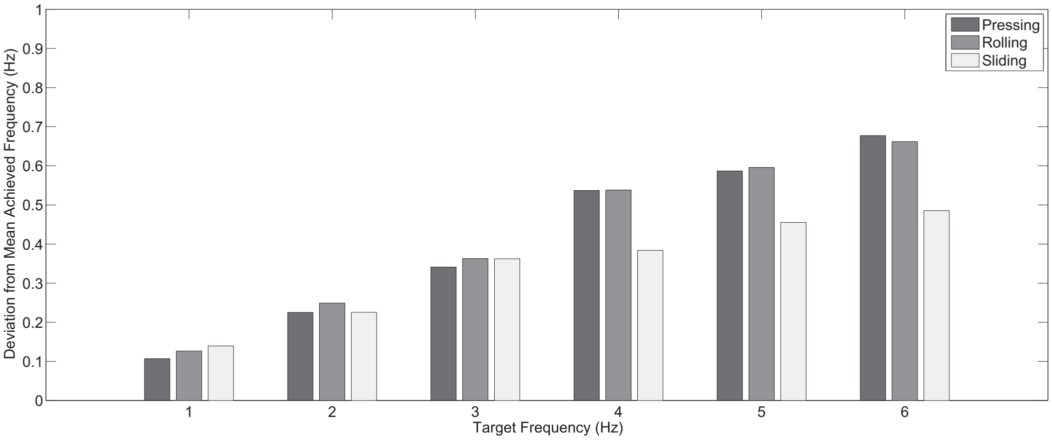

Figure 4.5: Deviation from the mean achieved frequency for each method at each of the target frequencies. A lower score indicates a higher level of precision.

comparisons].

Examining the modulation depth for each method depending on speed, there was no significant effect of speed on depth for the pressing method [F(5,35) = 1.47, p = .25] or for the rolling method [F(5,35) = .96, p = .45]. While there is a decrease in depth as speed rises, it is not a significant decrease. The sliding method on the other hand showed a significant difference in modulation depth due to speed [F(5,35) = 16.50, p < .001]. Post-hoc tests showed significant differences in depth of the modulation between speeds below 4 Hz and speeds from 4 Hz upwards [Tukey's HSD, p < .05 for all comparisons]. Figure 4.6 shows the modulation depth as a function of frequency for each of the mappings.

4.4.6 Discussion

Firstly, as noted in the previous experiment, the most significant effect on the precision of performance is that of the speed. Interestingly, this also holds for both accuracy and depth (although the power of this effect is much less for depth). This

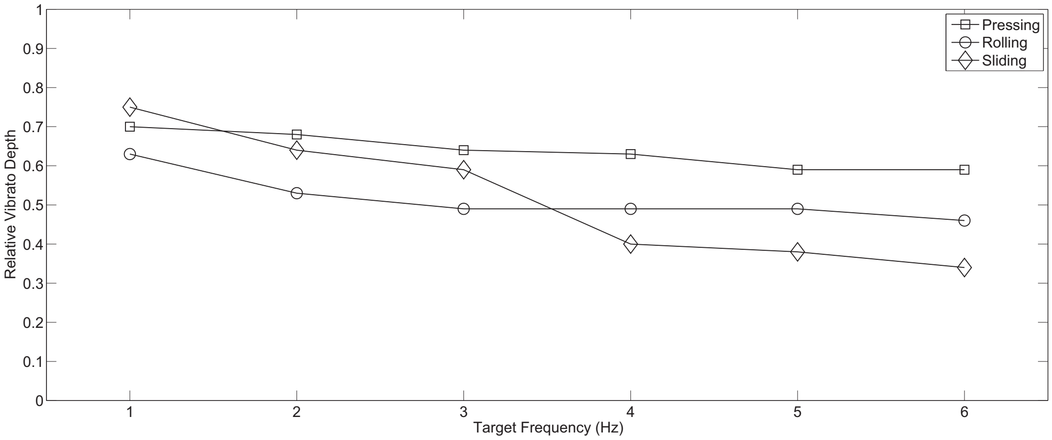

Figure 4.6: Modulation depth for each mapping as a function of frequency, normalized over the range +/- 1 semitone.

would indicate that when deciding on a control for a modulation task we must be aware that there is a decrease in performance at higher speeds. Also interesting to note is that there appears to be a cut-off point between 3 and 4 Hz, as accuracy, precision and depth all vary significantly above and below this point.

There also appears to be a link between the method used and the precision of the modulation being performed, as noted in the first experiment. In this experiment the sliding method was significantly more precise, although only for modulations at speeds above 3 Hz. At the same time, we can see a decrease in both modulation depth and accuracy for the sliding method at these speeds.

Another possible reason for this could be the availability of intrinsic visual feedback for the sliding method when compared to the other methods. When sliding their finger back and forward over a distance it is possible for the participant to see where their finger is and to use this feedback to ensure that they are consistent in how they play. This visual feedback is only available with the linear position sensor and only when moving over a distance along the sensor (as is the case with the sliding method but not with the rolling method).

The decrease in modulation depth at higher speeds for the sliding method may also be the result of the mechanics of the hand/arm movement. In order to increase the rate of the modulation while still maintaining some control over it, the performers must decrease the magnitude of the sliding movement. It is also possible that as the instructions for the experiment emphasized the speed of modulation as the primary interest, the participants may have purposely reduced the width to allow them to concentrate on the speed. Further experimentation where the width of the modulations is fixed could allow us to see if the accuracy of the sliding method would decrease even further and whether the precision of performance would also suffer.

Taken together, these results could indicate that the sliding methods is more suited to slower modulations, rather than the higher speeds. For instance, as already mentioned, most violin vibrato is in the 5 Hz - 7 Hz range, a range for which the sliding method would not be suited due to its reduced accuracy and depth at this range.

Overall there were no significant differences between the pressing and rolling methods for any of the examined factors (precision, accuracy or depth). If we were to choose between those two methods for a modulation of pitch in an interface then performer preference ratings (such as those in the first experiment or the previous work already discussed) would seem to be a good indicator of which is the most suited.

Finally, it should be noted that the reduced precision of the pressing and rolling methods at higher frequencies could be compensated for by the performer and might disappear with practice. Again, this provides a possible area for further research.

4.5 General Discussion

These experiments raised a number of interesting points which should be taken into account when deciding on the sensors to use in a digital musical instrument. Firstly, it would seem that the use of human-computer interaction techniques can help us derive guidelines for the use of sensors in digital musical instruments. The experiments described in this chapter showed significant relationships between a number of factors (including ease of use, precision and accuracy) and specific sensor and task combinations. Similar test procedure could prove extremely useful for the evaluation of specific sensors, or of prototype interfaces for digital musical instruments.

One of the most interesting results from these experiment seems to be that it is not just the sensor but the combination of sensor and gesture which is the important factor when deciding on the control method for a specific parameter. All of the experiment described here showed differences between the same sensor used with a different gesture, including significant differences in ease of use and preference ratings as well as accuracy and precision of performance. Designers of digital musical instruments should take care to examine different gestures for the sensors being tested. The fact that in the first experiment participants decided on different gestures to interact with the linear position sensor also shows the importance of user evaluation of prototypes. Users may not interact with an interface in the way the designer expects and so what seems like an optimal interface from the designers perspective may not be from the performer's perspective.

The results discussed here also indicate the importance of considering all of the parameters of a task when designing an interface. As seen in the third experiment, the speed of the modulation participants were asked to perform had the largest effect on the ability of participants to perform the task. While there was no significant difference between any of the methods tested at frequencies up to 3 Hz, above this point the accuracy and precision varied significantly for each method.

Finally, the experiments presented in this chapter show the importance of considering both user preference and ease of use ratings together with objective performance measurements in evaluating sensors for a digital musical instrument. Neither the preference ratings nor the objective measurements showed the complete picture in and of themselves. It takes a combined analysis of all of these factors to be able to determine the most suitable sensor for control of a specific task.

4.6 Conclusions

This chapter concentrated on a series of experiments to examine the suitability of sensors for specific tasks in a digital musical instrument. These experiments showed that a combination of user preference ratings and objective performance measurements determined from user testing can offer insights into the optimal choice of sensor for a specific task in a DMI. However, the experiments also show the importance of considering all aspects of the use of the sensors, including parameters of the task such as the speed and the gestures used to interact with them, in deciding on the best sensor for a specific task.

The next chapter deals with vibrotactile feedback in a digital musical instrument and more specifically on devices and methods of creating vibrotactile feedback and methods of measuring and compensating for the response of these devices over the range of frequencies for tactile sensation.

Chapter 5

Vibrotactile Feedback

Most traditional musical instruments inherently convey an element of tactile feedback to the performer in addition to auditory and visual feedback. Reed instruments produce vibrations which are felt in the performer's lips, string instrument vibrations are felt through the fingers on the strings, or through contact between the performer's body and the resonating body of the instrument. This tactile feedback leads to a tight performer-instrument relationship which is not often found in digital musical instruments.

Studies have shown that while beginners make extensive use of the visual feedback provided by musical instruments, in expert performance it is the tactile and kinaesthetic which is the most important (Keele, 1973). The majority of digital musical instruments provide only auditory and visual feedback to the performer, which results in a less complete sense of the instrument's response to the player's gestures than is available with traditional instruments (Chafe, 1993). It has also been stated that only the physical feedback from an instrument is fast enough to allow a performer to successfully control articulation (Puckette and Settel, 1993).

	2001	2002	2003	2004	2005	2006	2007	2008	Total
Loudspeaker	-	-	1	-	3	3	1	2	10
Vibrating Motor	1	1	-	-	-	1	-	-	3
Custom Actuator	-	-	1	1	-	-	-	-	2
Total	1	1	2	1	3	4	1	2	15

Table 5.1: Instruments presented at NIME providing vibrotactile feedback and the acuators used to do so.

Yet, if we examine the results of the survey discussed in Chapter 3 we can see that of the 266 new digital musical instruments examined only 55 provide any form of additional active feedback to the performer. From Table 5.1 we can see that in total, only 15 of the 266 instruments (less than 6%) surveyed provided active vibrotactile feedback.

When we further examine those instruments that do provide vibrotactile feedback, we see that there are 3 types of actuator used to produce this feedback: loudspeakers, vibrating motors and custom patented actuators. Each of these actuators has distinct advantages and disadvantages and yet little discussion has taken place on the optimal choice of actuator.

Finally, there are physiological and mechanical issues with the production of vibrotactile feedback. The response of human skin to vibrations differs depending on the frequency of vibration. Similarily, most vibration actuators produce different levels of vibration at different frequencies. These responses can be compensated for, yet of those instruments surveyed none do so.

This chapter addresses these issues. It deals with the creation of vibrotactile feedback, discussing the types of device available to generate vibration. Details are given of a method of measuring and evaluating these devices and a number of available devices are measured and compared. I then describe details of compensating for the frequency response of both the actuator and the human skin. Finally an experiment is performed to evaluate the effects of such compensation on the ability to detect changes in the frequency of vibrations.

5.1 Producing Vibrotactile Feedback

In order to create vibrotactile feedback we require mechanical actuators which can produce vibrations whose frequency and amplitude are within the range of human tactile sensation. Verillo (1992) provides a detailed description of the fundamental characteristics of such sensation. From this we can extract the requirements for a system to create vibrotactile feedback.

5.1.1 Feedback System Requirements

Frequency Response

Tactile (or vibrotactile) feedback results from contact between the body of the performer and the vibrating body of the musical instrument. Mechanoreceptors in the skin are sensitive to these vibrations. The fingers are capable of sensing vibrations in the region of 40 Hz to 1000 Hz and are most sensitive at 250 Hz (Verillo, 1992). These frequencies are within the audible range and are also frequencies which are among those produced by acoustic instruments. Figure 5.1 shows equal-sensation curves for the skin of the right hand to vibrations within this frequency range, from Verillo (1992). We can also see that the skin is most sensitive to vibrations at 250 Hz. Note that this response will vary with the size of the contact area.

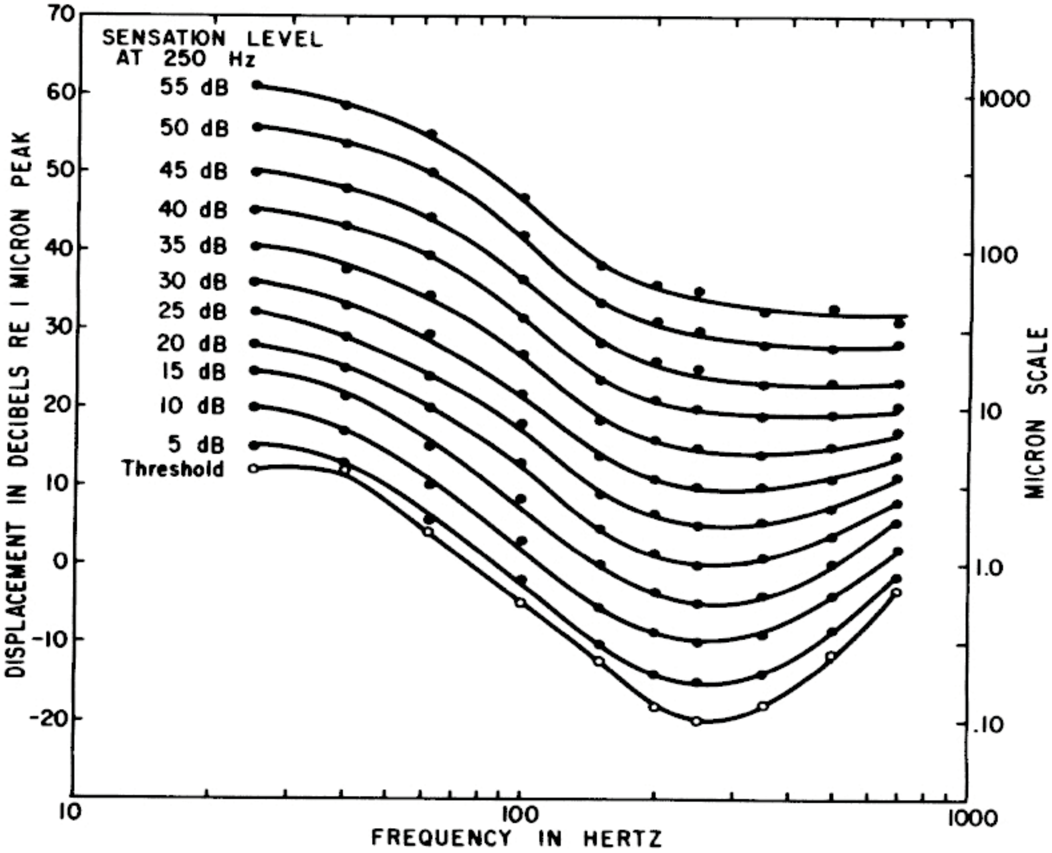

Figure 5.1: Equal-sensation curves for the skin on the hand, from Verillo et. al, Perception & Psychophysics, vol. 6, p. 371, 1969. Reproduced with permission from the Psychonomic Society. Copyright 1969, Psychonomic Society.

While we can sense vibrations within this frequency range, actual frequency discrimination is quite poor. Just Noticeable Difference (JND) measurements on tactile sensation of frequencies have been calculated at 30% (Goff, 1967), with improved discrimination at frequencies below 100 Hz when compared to those above 100 Hz (Mowbray and Gebhard, 1957). A more recent study by Pongrac (2008) found a JND of 18% across the frequency range. Note however that each of these values is far larger than the JND of frequency perception of the auditory system, which can discriminate frequencies differing by only 0.3% (Verillo, 1992).

From this information, we can derive the frequency response requirements for a vibrotactile feedback system. If we wish to make use of the frequency of vibration as a control parameter, then such a system should provide the ability to reproduce frequencies within the range of 40 – 1000 Hz. To provide the maximum level of frequency discrimination it should also allow for a change in frequency of at most 18%.

If frequency is not to be a control parameter, but kept fixed then our system could instead be designed to produce only a single frequency. Ideally this would be 250 Hz, the frequency at which the skin is most sensitive. Indeed, commercial actuators known as tactors use this strategy. More information on tactors can be found in Section 5.1.2.

Amplitude Response

As can be seen from Figure 5.1, the sensitivity of the skin to vibration varies greatly with frequency. In order to produce vibrations which are perceived as being of the same magnitude across the frequency range requires a variation in vibration amplitude of almost 60 dB. It can also be seen from this figure that lower frequency vibrations must have a larger amplitude than higher frequency vibrations in order to produce the same perceived amplitude (Verillo, 1992).

Again, the amplitude response requirements for a vibrotactile feedback system will vary depending on whether frequency is used as a control parameter or not. For a variable frequency system, the actuator should allow for at least a range of 60 dB in vibration amplitude, in order to allow us to create perceptually equal magnitude vibrations across the frequency range. For a fixed frequency system there is no minimum amplitude range. However as amplitude is one of the main controllable parameters in such a system a large range of amplitude control would be preferred.

Spectral Content and Texture

Tactile perception also allows for a sensation of texture. Rovan and Hayward (2000) reported that by varying the spectral content of the vibration signal presented to participants, they were able to produce a range of perceived textures. These varied in a continuum from a smooth texture with a sine wave to rough textures when using a noise signal.

Control of Vibration Parameters

In order to maximize control of the vibrotactile feedback system, it should provide for independent control of as many control parameters as possible. For a fixed frequency system, vibration amplitude is the sole control parameter. However, for a variable frequency system both frequency and amplitude are available as control parameters. In such a case, the system should ideally provide for independent control of frequency and amplitude.

Also of interest regarding control of vibration parameters is the control signal for the actuator. These can vary from Alternating Current (AC) voltages, to Direct Current (DC) voltages, to Pulse-Width-Modulated (PWM) signals. AC signals offer the greatest level of control, including amplitude, frequency and waveform. This is in direct contrast to DC signals, for which only the voltage level is variable and PWM signals for which only the duty cycle is variable. It should however be noted that some textural effects are also possible with DC and PWM signals, by quickly varying the signal.

The control signal and control parameters of an actuator also influence how easily it can be produced in a digital musical instrument. Complex AC signals can be derived from the sound synthesis output, but often require an amplifier in order to drive the actuator to significant amplitudes. DC and PWM signals can be often created directly from the microcontroller which is performing analog-to-digital conversion for the sensors in the DMI, without needing an additional connection to the computer system.

The ideal choice of control system and control parameters again depends on the system design. For a variable frequency system, a control signal which allows independent control of amplitude, frequency and waveform is ideal. This would seem to suggest the use of an AC control signal. For systems where the frequency is fixed, or where we wish to use only amplitude or frequency (but not both), we do not require independent controls and so DC or PWM signals may be more suitable.

Sound Production

Some vibrotactile actuators will also produce audible sound, as the frequencies of vibration which they are producing are within the audible range. If we wish to produce acoustic instrument-like vibrations this can be useful, as we can produce both sound and vibration within the instrument from the same actuators. For systems where the vibration produced is not to be related to the sound synthesis, but rather to the state of the system, this sound can be unwanted. It is possible that the sound produced by the actuators could even distract the performer from the actual instrument sound. In such cases, actuators which do not produce sound would be more suitable.

5.1.2 Devices for Vibrotactile Feedback

A number of different types of devices are available to produce vibro-tactile feedback in digital musical instruments. This section describes a number of these devices including their methods of operation, advantages and disadvantages and examples of their use in existing systems.

Tactors

The term *tactor* can be used to describe any device which creates tactile sensation. Therefore any of the devices described in this section can be classified as a tactor. Indeed, examples exist within the literature of tactors which use vibrating motors (Lindeman et al., 2005; Noma et al., 2006), piezoelectric elements (Gemperle et al., 2001) and voice coils (Gallace et al., 2005). However, there also exists a class of commercially sold devices called tactors, which are specifically designed to have a resonant frequency (and thus greater levels of vibration) at 250 Hz, the frequency at which human skin is most sensitive to vibration. In this work, the term tactor will be used to refer to such a device. This kind of tactor was used in experiments in adding tactile feedback to the LaserBass (Bongers, 1997), but was found to offer too low an amplitude of vibration and to have a delay when driven, which made it unsuitable for a vibrotactile element of a digital musical instrument. Tactors are commonly used to convey information to users in simulations and interfaces for the blind (Yanagida et al., 2004) or tactile information systems (Gemperle et al., 2001).

Piezoelectric elements

Piezoelectric elements are ceramic elements created from layers of piezoelectric crystals, which vibrate when an electrical current is applied to them. Normally used as sound producing devices in low-cost buzzers or to produce high frequency sound in high efficiency speaker systems, they can also be used as vibrotactile transducers. They do not seem to have found use as vibrotactile sound producers in digital musical instruments, but have been used in other tactile displays, such as the Optacon system for tactile representation of text (Linvill and Bliss, 1966). Other tactile interfaces have tried to use piezoelectric elements, but found the sound generated to be too loud for their requirements (Gemperle et al., 2001).

Voice coils

Voice coils are coils of wire which create a magnetic field when current is driven through them. Usually, the voice coil is held in the magnetic field of a permanent magnet. By applying an AC voltage to the voice coil, the coil can be caused to move within the field of the permanent magnet. The most common use of a voice coil is to drive the diaphragm of a loudspeaker in order to reproduce audio signals. Voice coils have been used to provide tactile feedback in a number of digital musical instruments, including those described by Chafe (1993), Rovan and Hayward (2000) and Birnbaum (2004). As they can be driven using the same audio signal (or a new signal derived from this signal) as that creating the audio output of the instrument, they are easily used as vibrotactile devices in digital musical instruments.

Motors

DC motors can be controlled by modulating the DC voltage applied to them. By either varying the voltage itself, or using a Pulse Width Modulated (PWM) voltage signal the speed of rotation of the motor can be controlled. When a weight is attached to the rotating shaft of the motor in such a way that the weight is not evenly dispersed around the centre of the shafts rotation this results in vibrations, the frequency and magnitude of which vary proportionally to the speed of the motor's rotation. Such unbalanced motors are commonly used as vibrational alarms in pagers, mobile phones and many vibrotactile game controllers and have also been used in developing vibrotactile feedback systems for a variety of computer systems including those described by Lindeman et al. (2005) and Noma et al. (2006).

Solenoids

In engineering terms, a solenoid is a transducer which converts electrical energy into linear motion. Solenoids consist of an electromagnetically inductive coil, wound around a movable steel or iron slug (termed the armature). When a current is passed through the coil it will create an electromagnetic field, which will either repel or attract the armature. By applying an AC signal to the coil the armature can be caused to repeatedly move in and out, resulting in a vibration. The Tactile Ring (Bongers, 1997), which contains a miniature solenoid has been successfully used to add tactile feedback to a number of instruments, including the LaserBass and the SonoGlove.

5.2 Evaluating and Comparing Actuators

There are several characteristics which must be examined in order to fully evaluate and compare actuator devices. These include mechanical characteristics and control characteristics. Mechanical characteristics, such as the frequency response of an actuator, must be measured for each specific actuator device. Control characteristics are more consistant across devices within a specific class and can be determined from the technical specifications of the device.

In order to fully determine the mechanical characteristics of a specific actuator, we must perform several measurements:

- Frequency response from 40 Hz to 1 kHz.
- Frequency resolution.
- Amplitude response over the entire range of amplitude input

Similarily for the control characterstics, we must examine the following:

- Number of control variables

- Dependency of control variables
- Type of control signal

This section deals with the comparison and evaluation of 5 actuators, one from each class described in Section 5.1.2. It begins with a discussion of the control parameters of these devices. This is followed by a description of the measurement of the mechanical characteristics of the devices. Finally, it includes a discussion of the choice of the optimal actuator based on the requirements of the specific vibrotactile feedback system being designed.

5.2.1 Control Characteristics

The two main control characteristics of interest for each actuator are the type of control signal used to drive the actuator and the independence of control of amplitude, frequency and spectral content. Table 5.2 shows a summary of these actuator characteristics.

Actuator	Control Signal	Amplitude Control	Frequency Control	Spectral Control
Loudspeaker	AC signal	Independant	Independant	Independant
Motor	DC or PWM signal	Dependant	Dependant	None
Piezo Disc	AC signal	Independant	Independant	Independant
Tactor	AC signal	Independant	Independant	Limited
Solenoid	DC	None	Independant	None

Table 5.2: Control characteristics of actuators

Three of the actuators tested can be driven using an AC signal: the loudspeaker voice-coil, the piezo disc and the tactor. All three offer independent amplitude and frequency control. Apart from the tactor all also offer independent spectral control. While the tactor can be driven using any AC signal, the output of the tactor is not a direct copy of the input signal. For instance, when driven using a square wave signal the tactor was found to output a noisy sine wave-like signal (Brown et al., 2005).

The motor and solenoid offer much less control than the AC signal driven actuators. The amplitude and frequency of the motor directly depend on one another. Higher input levels create both higher frequency and higher amplitude vibrations. It also does not offer any control of spectral content. The solenoid on the other hand offers no control of amplitude at all. It does offer control of frequency, but no control of spectral content.

5.2.2 Mechanical Characteristics

Methods and Procedure

In order to accurately measure the amplitude and frequency response of each of the devices an apparatus was built making use of a small, lightweight accelerometer mounted to a thin wooden board. During testing each device is also attached to the board, on the opposite side directly above the accelerometer. Vibrations from the device under test pass into the board and are measured by the accelerometer.

The accelerometer used was a PCB Piezotronics ICP accelerometer, model 352C22. This is a piezoelectric ceramic shear accelerometer, in a small (3.6mm × 11.4mm × 6.4mm) and lightweight (0.5 gm) aluminium enclosure. It provides a sensitivity of 10 mV/g and is accurate to within 5% over a frequency range of 1Hz to 10kHz.

The output signal from the accelerometer was connected to a PCB Piezotronics ICP Signal Conditioner, model 480E09. This acts as a power supply and signal amplifier, providing a gain of 10 for the accelerometer signal, with an extremely low noise level (<= -125 dB between 10 Hz and 10 kHz). Analog to digital conversion of the amplified voltage was performed using a National Instruments PCI-6036E with a 16-bit resolution (resulting in an overall Signal-to-Noise Ratio (SNR) of 96 dB) and a sampling rate of 10 kHz. Finally, control and data logging was performed using National Instruments LabView 7.1 software. Analysis was performed with Matlab, using the same functions as in Chapter 4.

Each device under test was placed in turn on the upper side of the wooden board. Devices were attached using an adhesive wax, which was also used to attach the accelerometer. A control signal was then sent to the actuator, depending on the specific test being performed and the actuator under test. The control signals used were as follows:

Frequency Response Measurement: For AC signal controlled actuators (the tactor, piezo disc and voice-coil) the control signal was a fixed amplitude sine wave. Test frequencies were created at 1/4 octave intervals between 40 Hz and 1030 Hz. For the PWM controlled actuator (the motor), a PWM voltage was applied. The pulse width was varied over 20 different values, spread logarithmically between calibration values. The calibration values were determined as those pulse widths which results in 40Hz and 1000Hz vibrations. For the DC controlled actuator (the solenoid) the control signal was a DC voltage which switched regularly between 0V and 5V. The frequencies of switching were varied between 40 Hz and 1000 Hz at 1/4 octave intervals.

Amplitude Response Measurement: The amplitude response was calculated as the range between the minimum and maximum vibration amplitudes produced by the actuator. The minimum amplitude was calculated as the lowest amplitude of vibration produced by the actuator which was above the threshold of human detection. For AC signal controlled actuators, the amplitude response was measured at a single frequency. Maximum control signal amplitudes were determined based on the RMS power ratings of the actuators. For the motor, the maximum vibration amplitude was measured at 100% pulse width (i.e with a constant DC voltage). As the solenoid is only capable of producing a single amplitude of vibration it was not measured.

Frequency Resolution Measurement: This was measured based on the smallest detectable change in frequency for each actuator. As resolution is more important at low frequencies, it was measured at the lowest frequency the actuator is capable of producing. For the AC signal controlled actuators, the frequency of the control signal was increased in steps of 0.1 Hz until a change in vibration frequency was detected. The difference between the detected frequency and the starting frequency was then measured. For the PWM controlled actuator, the pulse width was increased in steps of 0.1% until a change was detected. Similarly for the DC controlled actuator, the switching frequency was also increased in steps of 0.1 Hz until a change was detected.

The remaining characteristics of the actuators were determined based on observation of the devices under test by the test coordinator. The particular test methods were chosen to be consistent with those which exist in the literature and which have been successfully used to determine the frequency response of particular actuators (Teh and Featherstone, 2007).

Results

Figure 5.2 shows the measured vibration frequency responses for 4 of the 5 actuator types tested. No measurement is given for the solenoid as it was unable to produce any vibrations above 45 Hz. While high-speed solenoids capable of frequencies up to 1000 Hz do exist, they are relatively uncommon and I was unable to acquire one for comparison purposes. As such, the discussion of frequency response results in this section will not include any discussion of the solenoid. It will however be included in the other results and discussion later in this section.

Of the remaining actuators, all are capable of producing vibrations within the desired frequency range for human vibrotactile sensation. Interestingly, while the tactor has been specifically designed to produce a peak in its response at 250 Hz, several other actuators show peaks at or near this frequency. In particular, the 55mm diameter loudspeaker used shows a peak at 260 Hz and another at 520 Hz. The piezo electric disc also shows a peak in its response, although at just over 300 Hz.

The vibrating motor and the loudspeaker voice-coil offer the broadest range of vibration frequencies. The vibrating motor also shows the least variation in amplitude across the frequency range. Both the piezoelectric disc and the tactor perform poorly in the lower and upper frequency ranges.

Within actuator groups, there is also some level of variation in the frequency response depending on the size of the actuator. The loudspeaker, vibrating motor

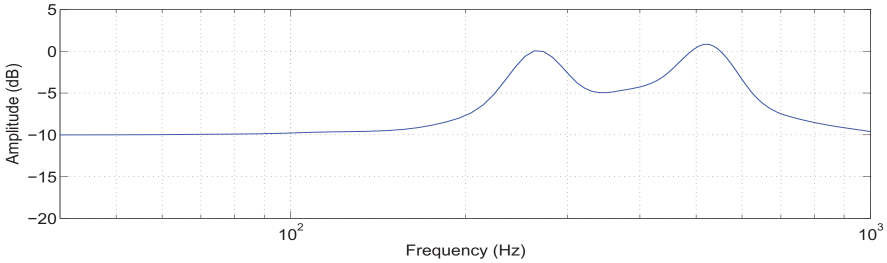

(a) Loudspeaker voice-coil

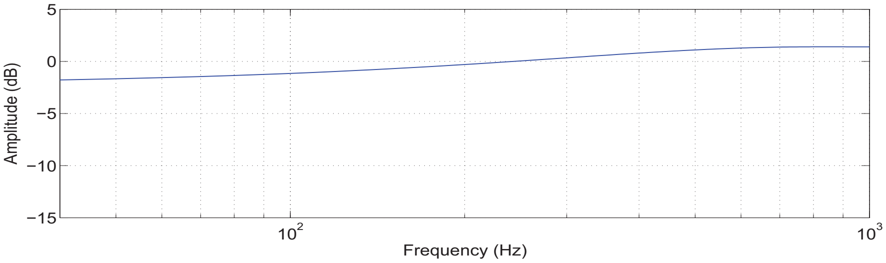

(b) Vibrating motor

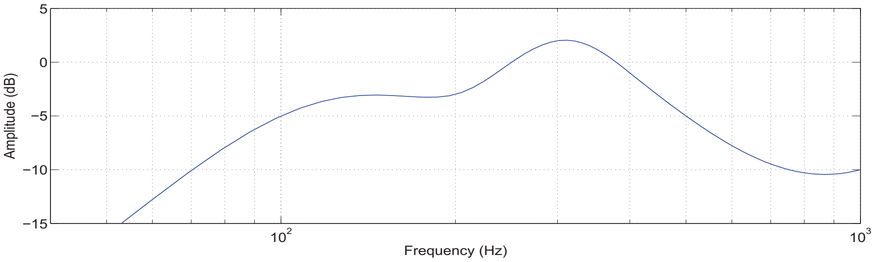

(c) piezoelectric disc

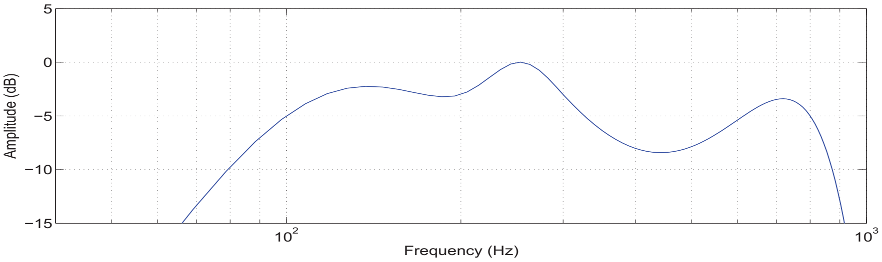

(d) Tactor

Figure 5.2: Measured vibration frequency response for the actuators under test

and piezo disc all show changes in frequency response depending on the size of actuator used. Larger actuators provide improved low frequency response, sometimes at the expense of high frequency response. Figures 5.3 and 5.4 show the frequency response measurements for two different sizes of vibrating motor and piezo disc respectively. Note how for the piezoelectric discs the response appear similar in shape, but to have shifted along the frequency axis.

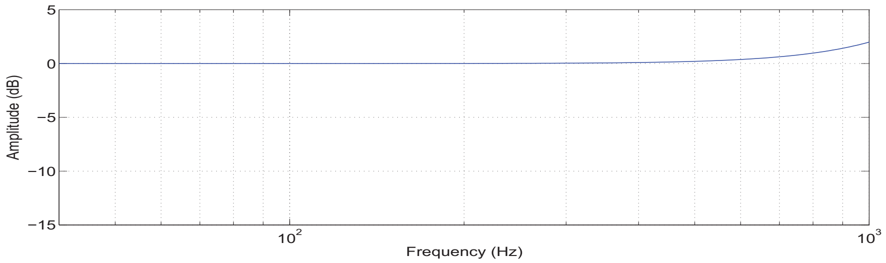

(a) Large vibrating motor

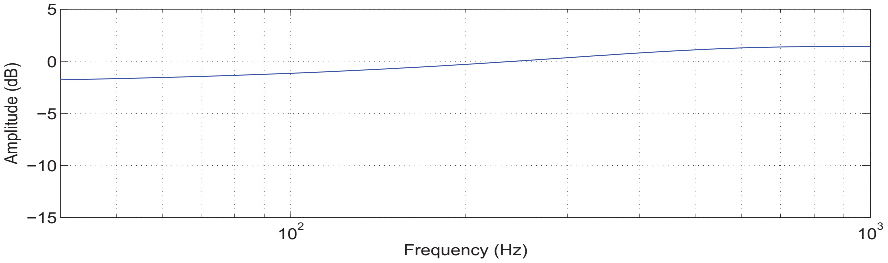

(b) Small vibrating motor

Figure 5.3: Effect of motor size on measured frequency response

The amplitude response of the actuators measures how wide a range of amplitude of vibration they are capable of producing. Table 5.3 provides a list of the amplitude ranges provided by each device. Measurements are given in dB relative to the minimum vibration amplitude produce by each actuator.

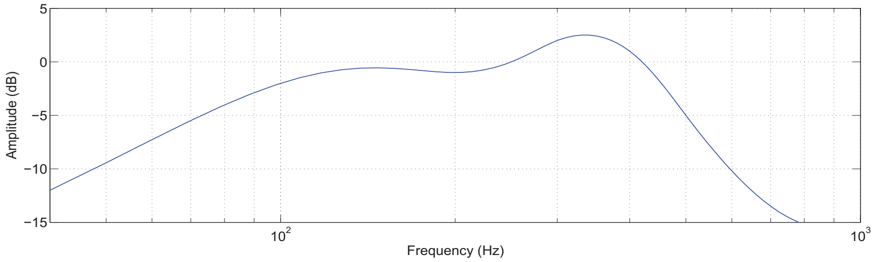

(a) Large piezoelectric disc

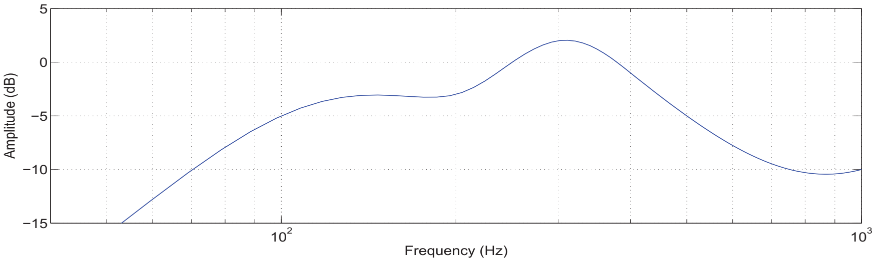

(b) Small piezoelectric disc

Figure 5.4: Effect of piezoelectric disc size on measured frequency response

Actuator	Amplitude Range
Loudspeaker	25.6dB
Motor	3.2 dB
Piezo Disc	22.1 dB
Tactor	48.7 dB
Solenoid	0 dB

Table 5.3: Measured amplitude response of each actuator at 250 Hz.

As can be seen from Table 5.3, the motor and solenoid each offer a very small range of amplitude control. The solenoid in fact offers only a single amplitude value. All three AC signal driven devices (the tactor, piezo disc and loudspeaker voice-coil) offer larger ranges of vibration. In particular, the tactor offers a very large range of amplitude control, although the full range is only available at a frequency of 250 Hz. While the voice-coil and the piezo disc both have variations in amplitude range at different frequencies, the variations are on the order of around 3 dB. For the tactor, these variations are much larger, even up to 25 dB at some frequencies. In the case of both the piezo disc the limiting factor on the amplitude response is the power rating of the device. The higher power rating the disc has the higher amplitude a signal can be fed to it and the larger the vibration amplitude.

For the loudspeaker voice-coil the maximum amplitude produced is limited by the power rating of the device, together with the maximum excursion of the voice-coil. The further the voice-coil is free to move, the larger a vibration amplitude it can produce. Larger loudspeakers often offer both a higher power rating and a larger maximum excursion, allowing for a much larger range of vibration amplitudes to be produced. Note however that the vibration amplitude when the actuator is in contact with human skin will also depend on a number of factors, including at-rest static skin pressure, damping, or skin impedance (Birnbaum, 2007).

Table 5.4 shows the measured frequency resolution data for each of the actuators measured. Measurements are given in percentage frequency change. This was measured from a base frequency of 40 Hz for all actuators except the solenoid. The measurement for the solenoid was made with a base frequency of 10 Hz.

Actuator	Frequency Resolution
Loudspeaker	$< 1\%$
Motor	4.3%
Piezo Disc	1.2%
Tactor	$< 1\%$
Solenoid	5.6%

Table 5.4: Measured frequency resolution of each actuator

As can be seen from these results, all actuators are capable of producing changes in vibration frequency which are smaller than the JND for human frequency detection.

5.2.3 Discussion

The results of this experiment indicate that a number of the devices under test are capable of producing useful vibrotactile feedback. The choice of actuator for a specific vibrotactile feedback system then can be made from these actuators based on the specific requirements of the system itself.

For example, to simulate acoustic instrument vibrations we require independent control over the frequency, amplitude and spectral content of the vibration. These requirements would limit the available options to the piezo disc and the loudspeaker voice-coil. On the other hand, a system which requires only a single variable with which to communicate information could use any of the devices tested, although a faster solenoid might be needed than the one examined here.

For a system with control over multiple parameters of the vibration, the options are slightly narrower. The voice-coil, piezo disc and tactor all offer independent control of the amplitude and frequency of vibration. This is in direct contrast to the motor, for which the frequency and amplitude are directly linked, and the solenoid, with a fixed amplitude and control only over frequency.

One interesting result of this experiment is the variations in response of each actuator across the frequency range. While the vibrating motor produces vibrations whose amplitudes are within 3 dB of each other, all the other actuators had large variations in vibration amplitude at different frequencies. If we wish to use both frequency and amplitude of vibration as separate parameters then we must compensate for this tendency. We must ensure that a change in vibration frequency does not result in a change in vibration amplitude.

The next section of this chapter deals with modifying actuator frequency responses. It includes a discussion of compensating for the frequency response of a specific device and also of compensating for the response of human skin to vibrations.

5.3 Modifying Actuator Frequency Responses

As already mentioned, none of the actuators tested produced a flat frequency response. However, for those actuators which offer control of both the frequency and amplitude of vibration we can modify the frequency response at specific frequencies in order to produce a flat (or near flat) response.

If we examine a hypothetical vibrotactile feedback system, there are 3 points at which we can think about the amplitude of the vibrations:

1. The amplitude of the vibration signal in the software which generates the excitation signal.

2. The *physical* amplitude of vibration produced at the actuator (i.e. the displacement caused by the actuator).

3. The *perceived* amplitude of vibration, caused by the response of the human skin to the vibration.

By analysing the response of the actuator and the skin we can introduce an equalisation element into the software which allows us to compensate for the response of the actuator and/or the skin. This can allow us to create vibrations which have the same *physical* amplitude, or the same *perceived* amplitude. In addition to this, it could allow us to further modify the frequency response of the system in order to achieve specific effects, such as modelling the vibration response of an existing instrument.

5.3.1 A Response Modification System

One way in which we can modify is to use a bank of bandpass filters. By setting the center frequency, amplitude and damping factor of each filter one can modify the response in whatever way one wishes. Together, these bandpass filters form a parametric equaliser. Such equalisers are commonly used in audio systems to modify the overall sound, allowing for the production of a flat frequency response even when used in a venue with large resonances.

To allow for further examination of the effects of response modification on vibrotactile feedback, I developed a vibrotactile feedback system in Max/MSP which contains a parametric equalizer (implemented using the *peqbank~* object from Jehan et al. (1999)). This system allows control over the frequency, amplitude and waveform of the vibration signal. Available signals are sine, square, triangle and sawtooth waves or an arbitrary waveform using a signal input.

The chosen vibration signal is passed through the *peqbank∼* object which contains a bank of bandpass filters and can be used to modify the response of the system as desired. Should we wish to modify the response in stages, such as to modify first for the actuator response and then for skin response, then multiple *peqbank∼* objects can be used. Visualization of the filter response is provided using a *filtergraph∼* object.

5.3.2 Using Response Compensation

In order to compensate for the frequency response of an actuator we first require a measurement of the actuator's response. This can be obtained as described in Section 5.2. The settings for the parametric equalizer can then be obtained as the inverse of this response. Figure 5.5 shows the measured frequency response of a 55 mm loudspeaker and the compensation curve calculated from this response.

The response of human skin to vibrations is also not flat across the frequency range. Verillo (1992) provides what can be considered "equal-sensation" curves, similar to the equal loudness curves for hearing. From these we can extract an average of the skin's response to vibration at different frequencies. As with the actuator response, we can determine the necessary compensation for the skin as the inverse of this response. Figure 5.6 shows the compensation curve determined from the skin's response.

By using these response compensation curves we can fully decouple the amplitude and frequency parameters for a vibrotactile feedback system. This allows

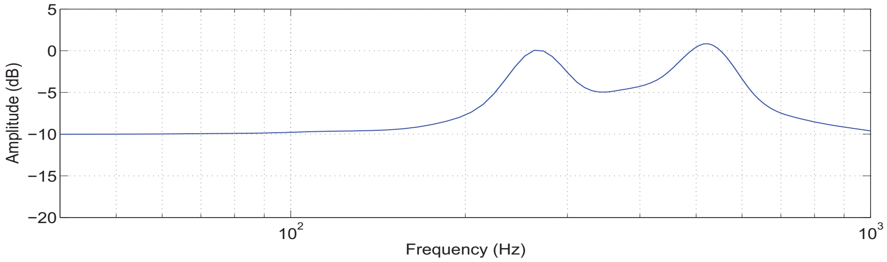

(a) Measured Frequency Response of 55mm Loudspeaker

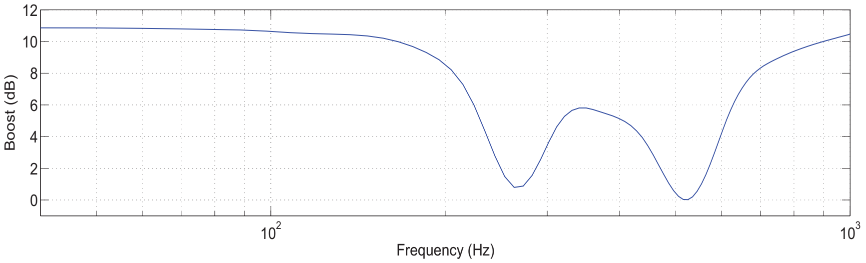

(b) Frequency Response of Compensation Filter

Figure 5.5: Measured response of a loudspeaker and calculated frequency compensation curve

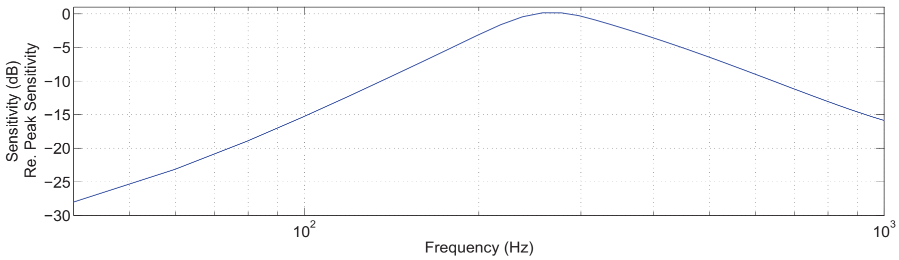

(a) Vibration Frequency Response of Skin

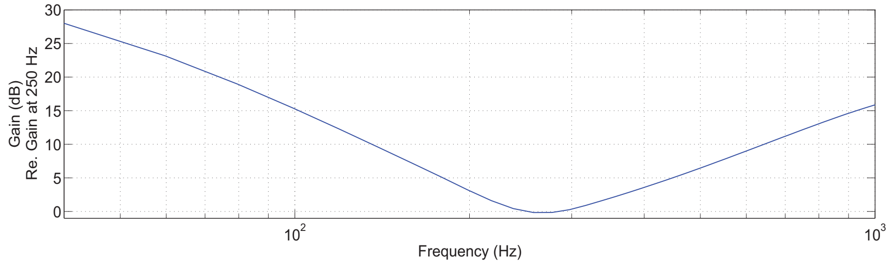

(b) Frequency Response of Compensation Filter

Figure 5.6: Vibration response of human skin and calculated frequency compensation curve

us to independently control these parameters. Such a system can also be used to tailor the response of the feedback system. This could then allow us to (for instance) model the vibration response of an acoustic instrument in a digital musical instrument. The response system can also be used to generate vibrotactile feedback directly from the sound synthesis system output of a DMI, which can then be modified to take into account both the actuator used and the sensitivity of the performer to vibrations. These approaches have been examined using the Viblotar instrument, which is discussed in more detail in Chapter 6.

5.4 Experiment: Actuator Response and Frequency Perception

The previous sections have detailed methods of both measuring and then compensating for the vibration response of various actuators, as well as human tactile sensation. While this is interesting from an engineering perspective, the question arises as to the effects of such compensation on users' perception. This section describes an experiment to examine this issue.

This experiment examines the effect of such compensation on the perception of the frequency of vibrations. As already mentioned, human discrimination of vibration frequency is poor, with measured values of JND between 15% (Pongrac, 2008) and 30% (Goff, 1967). The aim of the experiment described here is not to measure the JND, but rather to measure the ability to detect frequency changes in this range and the effects of compensation of this ability.

5.4.1 Participants

A total of 10 participants took part in this experiment. The participants were all graduate students in Music Technology. Nine participants had 5 or more years of musical training.

5.4.2 Design and Materials

This experiment examined the ability of the participants to perceive changes in the frequency of vibrations presented at the tips of their fingers. A 55mm diameter loudspeaker was used to produce the vibrations. This actuator is the same as that used in the measurements described in Section 5.2. It's vibration frequency response is shown in Figure 5.2a. This actuator was chosen as it allows independent control of the frequency and amplitude of vibration.

There were 3 test conditions and the participants performed all 3 conditions. Each condition consisted of 40 vibration stimuli. The stimuli were vibrations with a duration of 2 seconds. The task for the participants was to determine whether or not there was a change in the frequency of vibration after 1 second. Of the 40 stimuli, 20 contained a frequency change equal to or greater than the recorded JND for vibrotactile perception as reported by Pongrac (2008) and 20 contained no change. The changes all began at 250 Hz or 500 Hz and were of ± 0%, 18%, 20%, 25% or 30%. All signals were created in Max/MSP and produced by a MacBook Pro laptop computer.

The 3 test conditions examined were as follows:

Uncompensated vibration signals: all of the signals were created with the same signal amplitude in Max/MSP. No compensation was made for the actuator or skin response.

Actuator compensated signals: the amplitude of the vibration signal at each frequency was modified to compensate for the measured response of the actuator. This condition best reflects the setup used in existing JND experiments such as that of Pongrac (2008). This results in an equal physical amplitude of vibration at each frequency.

Fully compensated signals: the amplitude of the vibration signal at each frequency was modified based on both the actuator response and the "equal sensation" curves for human skin response described by Verillo (1992). This results in vibrations of equal sensation (i.e. equal perceived amplitude) being presented to the participants.

The presentation order of the test conditions and of the stimuli within each condition was randomised. Participants were asked for each presented stimulus to decide if there had been a change in frequency. They were also asked to rate their confidence in each answer using a 5-point Likert scale.

While audio levels produced by the actuators were generally too low to be heard, participants wore ear plugs to remove the possibility of being influenced in their decisions by the sound of the actuator.

All participants were right handed. Each participant was given the choice of using either hand to feel the vibrations of the actuator. Nine participants chose to use their left hand, with one choosing the right hand. This choice allowed them to activate the sound and notate their results using their right hand, while feeling the vibrations with their left hand.

5.4.3 Procedure

Subjects arrived at the lab and were given an Information/Consent form to read and sign. Subjects were shown the experimental interface and told that they would be presented with a number of vibration signals and asked to determine whether a change of frequency had occurred within each signal. They were also informed that they would be asked to rate their level of confidence for each answer. They were then presented with a number of example signals, featuring large, obvious changes in frequency. These examples were used in order to ensure that the participants were aware of how a change in frequency feels.

Participants took part in each test condition, in a randomised order, with a short break between each condition. After the experiment was complete, participants were also given an opportunity to add any comments they might have with regard to the experiment. They were then given monetary compensation and debriefed as to the nature of the experiment.

5.4.4 Data Analysis

Results were analysed using the Statistical Package for the Social Sciences software (SPSS). Data was analysed using repeated measures ANOVAs with post-hoc tests performed using Tukey's Honestly Significant Difference (HSD). Before the ANOVA analysis all outliers were removed and the data was checked for normality using the Shapiro-Wilk test.

5.4.5 Results

Frequency Change Detection

Figure 5.7 shows the frequency change detection results for each condition. A significant difference was found in the effect of compensation on the ability of the participants to determine whether or not a frequency change had occurred [F(2,18) = 15.2, $p < .001$]. In particular, both the actuator and full compensations resulted in significant improvements in the ability of the participants to detect frequency changes, when compared to the uncompensated signals [Tukey's HSD, $p < .001$ for both].

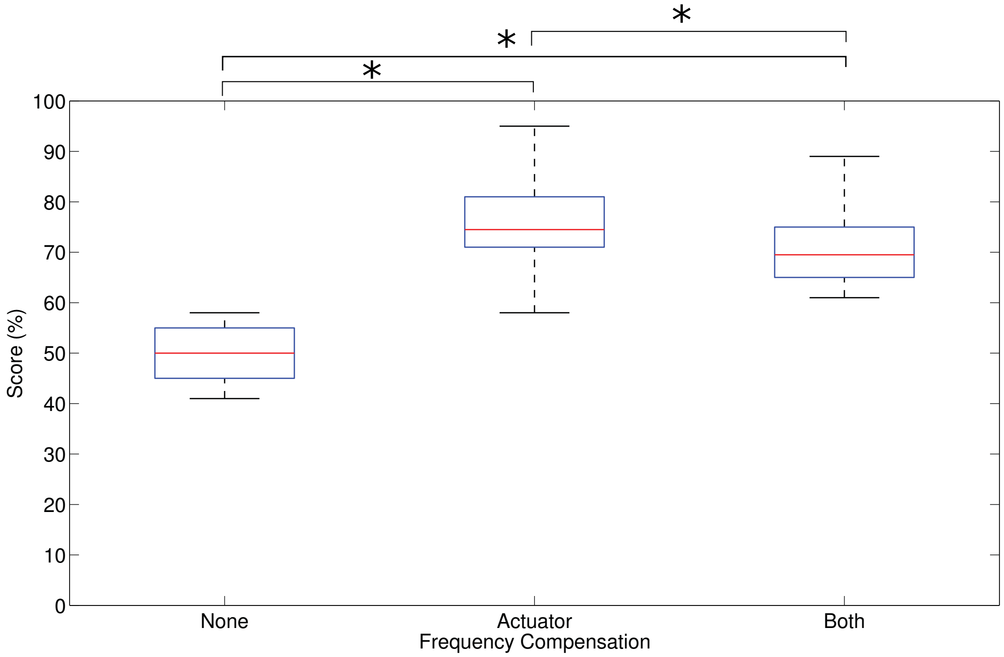

Figure 5.7: Frequency change detection results for each compensation condition. A * indicates a significant difference. Red lines indicate median values, while blue lines indicate lower and upper quartile values. Whiskers extend to 1.5 times the interquartile range.

There was also a smaller but significant difference between results achieved with the actuator compensation compared to the full compensation (Tukey's HSD, $p < .05$]. Interestingly, the actuator compensation resulted in a significantly higher percentage of correct judgements than the full compensation [$M_{actuator} = 75\%, M_{full} = 68\%$].

Confidence Ratings

Figure 5.8 shows the confidence ratings for each condition. Once again there was a significant effect of the compensation method on participant confidence ratings [F(2,18) = 13,24, $p < .001$]. Both actuator and full compensation conditions received higher confidence ratings than the uncompensated condition [Tukey's HSD, $p < .001$ for both comparisons].

There was no significant difference between the confidence ratings for the actuator and full compensation methods. The full compensation condition received slightly higher confidence ratings than the actuator compensation condition [$M_{actuator} = 3.94, M_{full} = 4.02$].

5.4.6 Discussion

The results of this experiment clearly indicate that the compensation methods used have a perceptible effect on the vibrations being produced. In particular, compensating for the actuators vibration response produces a significant improvement in both frequency discrimination and confidence ratings. Interestingly, this method of compensation is generally used in experiments to measure the JND of vibrotactile frequency perception, such as Pongrac (2008). However, based on the

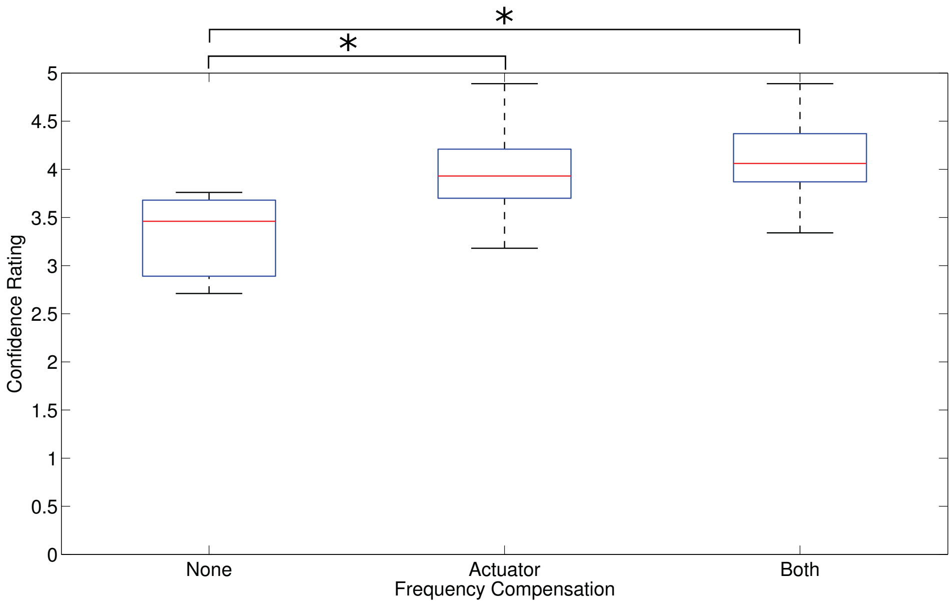

Figure 5.8: Confidence ratings for each compensation condition. A * indicates a significant difference. Red lines indicate median values, while blue lines indicate lower and upper quartile values. Whiskers extend to 1.5 times the interquartile range.

results of the survey of performed in Chapter 3, it seems that this form of compensation is not used in those digital musical instruments which provide active vibrotactile feedback.

From the frequency change detection results, we can also see that the full compensation (compensating for both the actuator and the skin) results in significantly worse results. When presented with vibrations which are at the same *perceptual* amplitude participants were less successful at identifying frequency changes than when presented with vibrations at the same *physical* amplitude.

One possible reason for this would be that when presented with the actuator compensated signal participants are perceiving an amplitude change and treating it as a frequency change. With this form of compensation participants are presented with vibrations whose amplitudes are physically the same (i.e. the displacement of the actuator is the same distance). However, the skin's sensitivity to such vibrations differs depending on the frequency and so the result could be a perceived amplitude change. This amplitude change may then be mistaken for a frequency change, generating a false positive which could skew the results.

To examine this possibility we performed a post experiment trial with each participant to test the confusion of amplitude and frequency changes. The same apparatus, methods and participants were used as in the original experiment. In this case however, the stimulus consisted of 40 vibration signals with either a frequency change or an amplitude change, but not both. Participants were asked to determine whether a frequency change had occurred or not.

For the ten participants, there was a mean rate of error of $M_{error} = 35.2\%$. While this shows that participants could differentiate frequency and amplitude changes to a reasonable extent, this error rate is still high enough that it could affect the results in the original experiment. This means that on average participants mistook an amplitude change for a frequency change 35% of the time. In the actuator compensated condition the possiblitiy arises that for certain frequency changes the particpants would not perceive the frequency change. However, they could perceive an amplitude change and then mistake this change for a frequency change. This could result in a correct answer, but based on an incorrect judgement of the perceived change. Again, this might explain why participants performed better with actuator compensation than with full compensation.

Interestingly, existing attempts to determine the JND of vibrotactile frequency detection in the literature do not seem to take into account the response of the skin to vibration frequency. This leaves the possibility that similar confusion between perceived amplitude and frequency changes could also have affected the results of these studies. This is an area that could benefit from further research which is outside the scope of this work.

Finally, it is also interesting to note that participants were more confident of their answer when compensation was present. Both levels of compensation produced a large increase in confidence level when compared to the uncompensated signal. Also, while there is no significant difference between the confidence ratings for actuator and full compensation, we can see that participants were most confident of judgements made using full compensation, even though they were slightly less accurate with this method.

5.5 Conclusion

This chapter discussed technical aspects of the use of vibrotactile feedback in digital musical instruments. While acoustic instruments provide vibrotactile feedback to performers directly due to the sound production mechanisms which they use, most digital musical instruments do not. In order to provide such feedback, an active vibrotactile feedback production system must be used.

The design of such a system is dependent on the way in which humans sense vibrations. The optimal choice of actuator to create the feedback must then be based on the goals of the system itself, in conjunction with psychophysical aspects of vibration perception. Issues such as frequency discrimination of vibrations, vibration detection thresholds and the ability to differentiate between changes in amplitude and frequency all affect the design of vibrotactile feedback systems. In this chapter I described methods of evaluating actuators and compared a number of common actuators in light of these psychophysical aspects.

The results of this comparison indicated that a number of actuators are capable of being used for vibrotactile feedback. Some, such as the voice-coil, piezo disc and tactor, are capable of producing the full range of vibration frequencies to which humans are sensitive. They also offer control over the amplitude and to some extent the spectral content of the vibrations. These actuators then could be used to create acoustic instrument-like vibrations in a DMI.

Other actuators have reduced frequency ranges, no spectral control and often have no amplitude control or have a coupling of the amplitude and frequency parameters. These actuators (such as vibrating motors) may be more useful to provide an extra channel of information to the performer, whether directly related to a system parameter, or communicating system state changes to the performer.

When examining the frequency response of actuators in more detail we find that while many can cover the full range of vibrations which humans can sense, they do not do so with an equal amplitude at each frequency. This causes changes in amplitude of vibration with frequency changes. If we wish to use both parameters separately in our vibrotactile feedback system then we require a method of compensating for this effect.

The frequency compensation system described in Section 5.3 performs this purpose. This system can be used not only to compensate for the actuators, but also for the frequency response of the skin itself. An experiment was conducted to verify the perceptual effects of this system. Results indicate that such compensation has a perceptible effect, providing a significant improvement in vibration frequency discrimination.

What remains then is to examine the use of such vibrotactile feedback in digital musical instruments themselves. In the remaining chapters of this thesis I discuss a number of instruments developed over the course of this research, some of which make use of vibrotactile feedback systems.

The next chapter deals with the Viblotar. The Viblotar is a new digital musical instrument which was developed specifically to allow examination of the research discussed in this thesis within the context of instrumental performance. It contains a vibrotactile feedback system which is used to create vibrations based on the sound it is producing, in a way which is similar to that of acoustic instruments.

Chapter 6

The Viblotar

The Viblotar is an instrument designed to allow evaluation of the results presented in earlier parts of this thesis within the context of performance with a digital musical instrument. In particular, it was designed to follow the results of the experiments described in Chapter 4 in the choice of its sensors. As the Viblotar contains an internal sound and vibrotactile feedback production system, it also offers an opportunity to evaluate some of the componenets of the vibrotactile feedback system discussed in Chapter 5.

Having found strong indications of the suitability of specific sensors for both note selection and note modulation tasks in the first experiment described in Chapter 4, I decided to develop an instrument that was played almost entirely using these tasks and to implement it using the sensors which the experimental data indicated were best suited to these tasks. In addition to this, the decision was made to implement vibrotactile feedback in the Viblotar using built-in speakers driven by the sound synthesis system. As already discussed, this has the dual benefit of providing both vibrotactile feedback and localizing the sound production to the instrument itself.

6.1 Design

Based on the overall aims of the project, the following design requirements were arrived at:

- The instrument should be playable using only note selection and note modulation gestures.
- The gestures should be sensed using the sensors which had shown to be most suitable for these particular tasks.
- The instrument should provide vibrotactile feedback to the performer and this feedback should be directly related to the sound produced by the instrument.
- It should be possible to play existing musical pieces using the instrument, which should offer a good range of continuous pitch control.
- The overall design of the instrument should aim towards an instrument in which all aspects of the system are integrated (i.e. no external systems are needed to perform with the instrument).

With these overall goals in mind I decided to develop an instrument inspired by a traditional monochord. A number of cultures have developed monochord musical instruments, for use in education, research or as performance instruments (Hopkin, 1996). One such instrument is the Dan Bau, a Vietnamese monochord, which is used to play melodies using the harmonics of the string. The instrument is played using the right hand both to touch the string at one of seven nodes and to pluck the string to excite the harmonics of that node. The left hand is used to modulate the note produced by increasing or decreasing the tension of the string, which is connected to a length of bamboo at that end. Figure 6.1 shows a traditional Dan Bau.

Figure 6.1: A traditional Dan Bau. Image copyright by DanTranh.com, used with permission.

The design of the Viblotar was then based on the Dan Bau, but offering some additional methods of playing not possible on the acoustic instrument[1]. This would make the Viblotar an instrument-inspired controller, as it would not create the sounds of the acoustic instrument, but its methods of interaction are directly inspired by the Dan Bau and other monochord-style instruments.

6.1.1 The Physical Interface

The body of the instrument is constructed from 19mm thick pine boards. The overall shape of the instrument is a long rectangular box and is designed to allow

---

[1] See http://www.mcld.co.uk/oddmu/danbau/ and http://www.thanhcammusic.com/ for more information, pictures and videos of the Dan Bau

performance of the instrument when placed on a desk or a keyboard stand. The body also acts as an enclosure for the speakers, allowing a better frequency response than that of the speakers in free air. The overall dimensions of the body were arrived at by balancing the size of the sensors, the speakers and the necessary playing surface with dimensions calculated to maximize the frequency response of the speakers. The resulting enclosure is 910mm wide, 250mm deep and 270mm high.

These dimensions were determined by modeling the speakers in an enclosure using MathCad models with the Thiele-Small parameters of the speakers[2], which were determined by small signal analysis (Thiele, 1971; Small, 1973). This resulted in the enclosure shown in Figure 6.2.

Figure 6.2: A view of the front and top of the Viblotar. In performance, the two front-mounted loudspeakers point towards the audience. The long linear position sensor can be seen on the left, with the two FSRs on the right.

The sensors are placed on the top of the instrument, with the speakers on

---

[2] See http://www.quarter-wave.com/ for details of the models used and their derivation

the front, facing away from the performer and towards the audience. To allow for interaction with the instrument a long linear position sensor and a long force sensitive resistor (FSR) (overall length 480mm) were placed on top of each other towards the right side of the instrument. Two smaller square FSRs (38mm on each side) were placed to the left of these sensors, one above the other. The layout of the sensors on the top of the instrument can be seen in Figure 6.2.

Electronics

The electronics for the Viblotar consist of three major sections. Each of these sections is concerned with a specific aspect of the instrument interaction. All electronics are mounted internally, on top of the board which is the bottom of the body. The sections of the electronics are:

Sensors and Signal-Conditioning Circuits This section of the electronics is concerned with taking the interaction from the user and converting it to a voltage in the range of 0 to 5V. The long FSR and linear position sensors only require a 5V supply in order to operate in this way, whereas the smaller square FSRs require a signal conditioning circuit to allow them to produce a voltage rather than the resistance which is inherent in the sensor. A simple voltage-divider was created for these FSRs using two resistors.

Analog-to-Digital Conversion The analog-to-digital conversion system used for the Viblotar is based on an Arduino Mini with an Arduino Mini USB adapter[3]. It samples the sensors at a rate of 120 Hz and with a 10-bit resolution. The sampled values are then sent to the computer over a USB

---

[3] See http://www.arduino.cc for more information.

connection. It also provided the +5V signal required for the sensor and signal conditioning circuits.

Internal Sound Production For the current version of the Viblotar, the synthesis is performed on a computer. However, the actual sound production is performed in the instrument using the two 80mm diameter loudspeakers mounted on the front of the instrument body. Two small audio amplifier circuits were created which are capable of providing 1W of power into these 8 Ω speakers from a 5V supply. The power for these amplifiers was taken from 6 1.5V batteries and run through a regulator to provide a stable 5V power supply. This system allows for the production of a maximum of 93dB (SPL) of sound output at a distance of 1m from the instrument.

6.1.2 Mapping and Synthesis

The synthesis engine for this system consists of a physical model running in the Max/MSP environment. The model comes from the PeRColate package by Dan Trueman and R. Luke DuBois[4] which is a port to Max/MSP of instruments from the Synthesis ToolKit (STK) by Perry Cook and Gary Scavone[5]. The physical model currently being used is a hybrid model called the blotar (thus the name Viblotar for a ***Vi**brating **Blotar***). This model is a hybrid of an electric guitar model and a flute model. This allows for a large range of sounds to be produced by the system, while the use of physical modeling gives a more traditional instrument-like sound to the Viblotar.

As one of the main aims of this project was to evaluate the results of the

---

[4] Available from http://music.columbia.edu/PeRColate/

[5] See http://ccrma.stanford.edu/software/stk/ for more information

previous experiments on mapping of sensors to musical tasks, the majority of the gesture to sound mapping was implicit in the system design. As such it was already determined that position along the linear position sensor would map to the frequency of the output sound. Pressure on the smaller square FSRs would map to pitch bend up and down from the output sound frequency. In the initial version of the Viblotar, pressure on the long FSR (mounted under the linear position sensor) was mapped directly to amplitude. For the current version of the Viblotar, amplitude is based on the force of striking gestures on the long FSR.

The use of physical modeling synthesis allows these mappings to be made in an intuitive way. For instance, instead of mapping the output from the long FSR to a variable in an equation for synthesis, the physical model allows us to map it to the pluck strength. The physical model takes this pluck strength and applies it to the necessary equations to generate the variables for the synthesis. This approach makes it easier to interface the controller with the synthesis and allows for a much better understanding of the way which the instrument will react to user actions (Hunt et al., 2000).

The current mapping of the controls are as follows:

- Position on the linear position sensor controls the frequency of the note played. This is a continuous linear mapping increasing from left to right. The total frequency range is from 100Hz to 1000Hz (just over 3 octaves).

- The force of striking gestures on the long FSR controls both the pluck strength and the overall output sound volume.

- Slower touching gestures on the long FSR damp the sound which is currently being played.

- Pressure on the smaller FSRs is mapped to a ±10% change in the frequency of the played note. Again this is a linear mapping from the sensor output to the frequency change.

The Viblotar is played by manipulating the sensors on top of the instrument. The performer creates notes using striking- or plucking-type gestures on the linear position sensor. These gestures trigger sounds from the synthesis system. The position of the pluck gesture on the linear position sensor controls the pitch of the note. The amplitude of the note is dependent on the force of the pluck, as detected by the long FSR which is placed under the linear position sensor. The force of the gesture also effects some aspects of the timbre of the sound, through changing the *pluck strength* parameter of the physical model.

Once a note has been triggered, the performer can manipulate it in two ways. Firstly, the pitch of the note may be modified using the two square FSRs on the left side of the instrument. One FSR is used to raise the pitch, the other to lower it. This allows for effects such as pitch bend and vibrato. Secondly, the performer may damp the currently playing note by pressing the linear position sensor with their finger or hand. The force of this pressing gesture, again detected by the long FSR under the linear position sensor, is used to control the amount of damping applied. The detection of these two different playing gestures on the same sensors (damping and plucking), is based on the attack time of the gesture itself.

6.2 Producing Instrument-like Vibrations

As mentioned in the previous chapter, one possible use of a vibrotactile feedback system in a digital musical instrument is to produce vibrations that are based on the sound the instrument is producing. In an acoustic instrument the sound production mechanism also produces the vibrations that the performer feels. If we wish to provide vibrations in a DMI that are produced in a similar way to those of an acoustic instrument, these vibrations must then be directly linked to the sound production. Such a link can be achieved by deriving the vibrotactile feedback signal from the sound synthesis output of a DMI.

In order to physically produce these vibrations then, an actuator is needed which meets the following requirements:

1. Capable of producing the full frequency range of human tactile sensation.

2. Offer independent control of frequency, amplitude and waveform.

3. Offer a large range of amplitude control (to allow for instrument dynamics).

4. Driven by an audio signal, or a signal easily derived from an audio signal.

Examining the results presented in Section 5.2, we can see that voice-coil, the tactor and the piezoelectric element each meet these requirements to different extents. Of these, the voicecoil offers the greatest range of frequency and amplitude control. Also of interest is that if we use a voicecoil in the form of a loudspeaker, then the system can also be used as the main sound production method of the instrument. This not only adds sound-related vibrotactile feedback to the instrument but also co-locates the sound production into the instrument itself (Cook, 2004; Armstrong, 2006).

6.2.1 Vibrotactile Feedback from the Sound Synthesis System

As discussed in Section 5.3, it is possible to create vibrotactile feedback using a loudspeaker and software such as Max/MSP which can create and modify audio signals. While the system described in that section was designed to compensate for actuator and human skin vibration response, it can also be used to produce instrument-like vibrations in a digital musical instrument such as the Viblotar.

Figure 6.3 gives an overview of the components of the Viblotar. The output of the sound synthesis system is used to drive both the *external* sound production and the vibrotactile feedback (and *internal* sound production) components. The internal sound production mechanism consists of the amplifier and loudspeakers embedded in the instrument body. The external sound production would be any amplifiers or external loudspeakers, which could be used to provide amplified sound for performance in a larger space. In many cases the internal and external sound production would be driven using the same signal, so that the external sound is an amplified version of the internal sound. However, the use of separate internal and external sound production mechanisms allows for some interesting effects which will be discussed in more detail in Section 6.2.2.

The vibrotactile feedback generator for the Viblotar is the same as that described in Section 5.3. In this case, the output from the sound synthesizer is fed to the input of the response modification system. This signal then passes through the parametric equalizer sections before being output through a digital to analog converter (DAC). The output of this DAC is a line level audio signal which is fed to the hardware of the Viblotar's vibrotactile feedback component. There it is

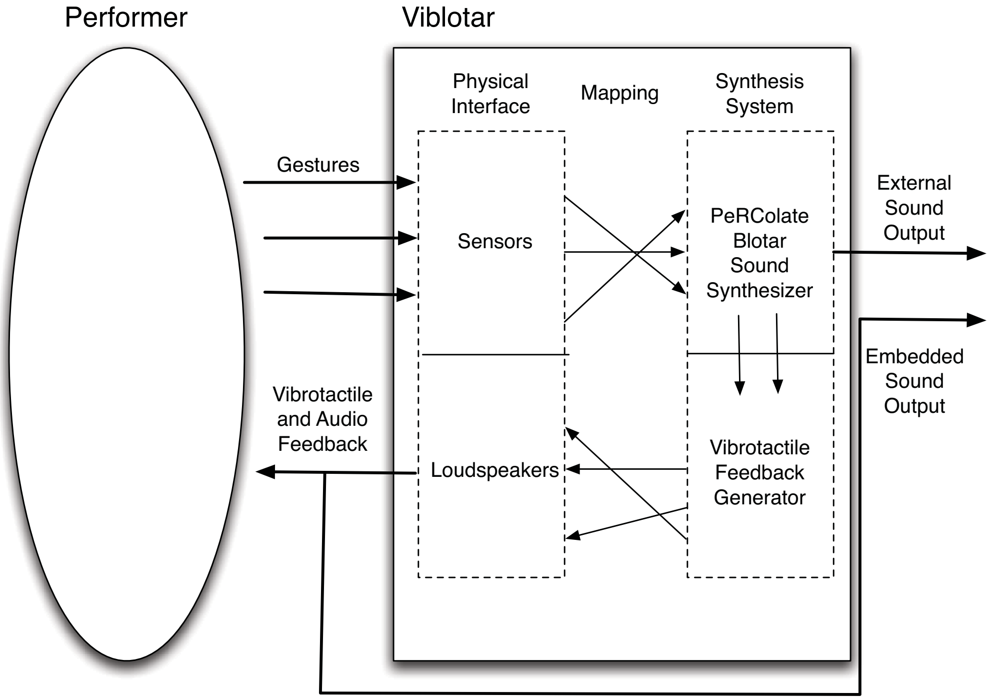

Figure 6.3: The overall structure of the Viblotar. This design is based on the DMI model shown in Figure 2.6.

amplified and output through the embedded loudspeakers.

When the sound synthesis signal is fed directly through the vibrotactile feedback generator, without any modification of the signal, then the vibrotactile feedback provided by the Viblotar is directly related to the sound of the instrument. The sound produced by the embedded loudspeakers is the sound of the instrument itself and this sound causes vibrations in the instrument. However, it is also possible to modify the signal used to drive the vibrotactile feedback component. In this case, the vibrotactile feedback would still be related to the sound produced by the instrument, without being directly caused by it. By using the unmodified signal to drive the external sound production and a modified signal to drive the vibrotactile feedback and internal sound production we can create a number of interesting feedback effects.

6.2.2 Modifying the Vibration Response

The availability of both internal and external sound production mechanisms in the Viblotar allows 3 main modes of operation:

Internal sound production only: in this mode of operation, all of the instrument's sound is generated within the instrument itself, by the built in loudspeakers. This is closest to how an acoustic instrument such as the acoustic guitar works.

Internal and external sound production: this mode offers two sound sources. The first is the instrument itself, through the embedded loudspeakers. The second source is an external (and possibly amplified) loudspeaker. This mode of operation is based on instruments such as the electric guitar or electric violin.

Modified internal sound production: when using both internal and external sound production it is possible to modify the signal used for internal sound production, creating a difference between the sound created internally by the instrument itself and that produced by the external system.

When using different signals for each sound generating mechanism, we can perform a number of interesting effects, including:

- Compensation for the frequency response of the loudspeakers and/or human skin.

- Simulation of the frequency response of a different instrument.
- Production of only those frequencies for which the skin is sensitive.

Each of these effects can be performed for the internal sound production and vibrotactile feedback portion of the instrument, while still producing the unmodified sound from the sound synthesis system through the external sound production mechanism.

As was already discussed in Chapter 5, neither the actuators used to produce vibrotactile feedback nor the human skin offer a flat response to vibrations across the frequency range. By having separate control over the frequency content of the signal sent to the vibrotactile feedback system we can compensate for these responses. For instance, if the instrument is to generate low frequency sounds it is possible that the loudspeakers used may have a reduced response at these frequencies. By modifying the signal sent to the loudspeakers we could increase the output amplitude for these low frequencies.

Modification of the vibrotactile feedback signal can also be used to modify the vibration response in such a way as to make it more like the response of a different instrument. It is possible to increase or reduce the response at certain frequencies or within certain frequency bands. This could, for instance, be used to produce low frequency vibrations for an instrument with a poor low frequency response. It could also be used, together with measurements of the vibration response of an existing musical instruments, to simulate the resonances of the body of other instruments in the Viblotar.

Finally, by modifying the feedback signal, we could restrict the sound produced by the internal sound production mechanism (and thus the vibrations created) to only those frequencies to which the human skin is sensitive. This results in the internal sound production being used mostly for vibration production, while the actual sound production occurs outside of the instrument itself. In fact, it would even be possible to restict the internal sound production to frequencies which are too low to be audible, thus using it solely for vibration generation.

It is also possible (and perhaps even advisable) to combine a number of these effects together. For instance, when attempting to simulate the resonances of another instrument it may well be necessary to apply compensation for the actuator so that the target response is produced by the system.

6.3 Measuring Instrument Vibrations

For some of the effects mentioned in the previous section, and indeed to enable a mechanical evaluation of the vibrotactile feedback system used in the Viblotar, it is necessary to be able to measure the vibrations of a given instrument, whether acoustic or digital. This section describes a method of measuring instrument vibrations and provides examples and comparisons of the vibration of an acoustic guitar and the Viblotar. The measurement method described in this section is based on that used by Askenfelt and Jansson (1992), who measured the vibration response of a number of stringed instruments at different points on the instrument body.

The aim of the measurements made here are to compare the vibrations of an acoustic instrument (an acoustic steel stringed guitar) with a new digital musical instruments (the Viblotar). In particular, we are interested in showing certain common traits between these two different instruments. Questions of particular interest are:

1. Do these instruments produce vibrations above the threshold of human detection?
2. Are there similarities in the spectral content of these vibrations?
3. Are the spectra of the vibrations related to the note being played?

6.3.1 Methods and Procedure

All vibration measurements were made with the instrument in normal playing position. As in the measurements described in Section 5.2, a PCB Piezotronics ICP accelerometer, model 352C22 was used for all vibration measurements. The output signal from the accelerometer was connected to a PCB Piezotronics ICP Signal Conditioner, model 480E09. Analog to digital conversion of the amplified voltage was performed using a National Instruments PCI-6036E with a 16-bit resolution and a sampling rate of 100 kHz. Finally, control and datalogging was performed using National Instruments LabView 7.1 software. Analysis of the recorded signals was again performed with Matlab.

For each instrument, the accelerometer was attached at the measurement position using adhesive wax. Each instrument was held in the playing position. All measurements were performed using a single pitch, corresponding to the open low E string of the guitar. This gives a frequency of 82 Hz. Multiple measurements were made for each instrument. These measurements were averaged during the analysis stage to reduce the effect of any artefacts from single notes.

For the guitar, the procedure was as follows: the accelerometer was attached to the instrument on the top plate, near the bridge. The instrument was held in the playing position, with the neck resting in the left hand, but no fingers pressed to the fingerboard. The low E string was plucked using a pick at the specified dynamic level and allowed to resonate until no detectable vibrations were present. This was repeated 10 times.

For the Viblotar, the procedure was similar. The instrument was held in the playing position, with the body of the instrument resting on the performer's legs, as shown in Figure 6.4. The left hand was allowed to rest on the left side of the instrument, near the FSRs. The right hand was also allowed to rest on the instrument, directly below the linear position sensor. For the purpose of this experiment, the Viblotar mapping was modified so that a touch at any point on the sensor produced the desired note. The linear position sensor is touched using one of the fingers of the right hand. The note is allowed to resonate until no detectable vibrations are present. To ensure no accidental damping or modulation of the note occurs, these functions of the mapping system were also disabled for the duration of the test. As with the guitar, this procedure was repeated 10 times.

6.3.2 Results

Figure 6.5 shows the average vibration spectrum measured for the acoustic steel string guitar. Notice the peaks fundamental and each of its harmonics. The spectrum shows especially large peaks at the 2^{nd} and 4^{th} harmonics. Note also how the vibrations in the lower frequencies are above the threshold of human vibrotactile detection.

The average vibration spectrum for the Viblotar is shown in Figure 6.6. As with the guitar, it shows peaks at the harmonics of the note played. Unlike the guitar,

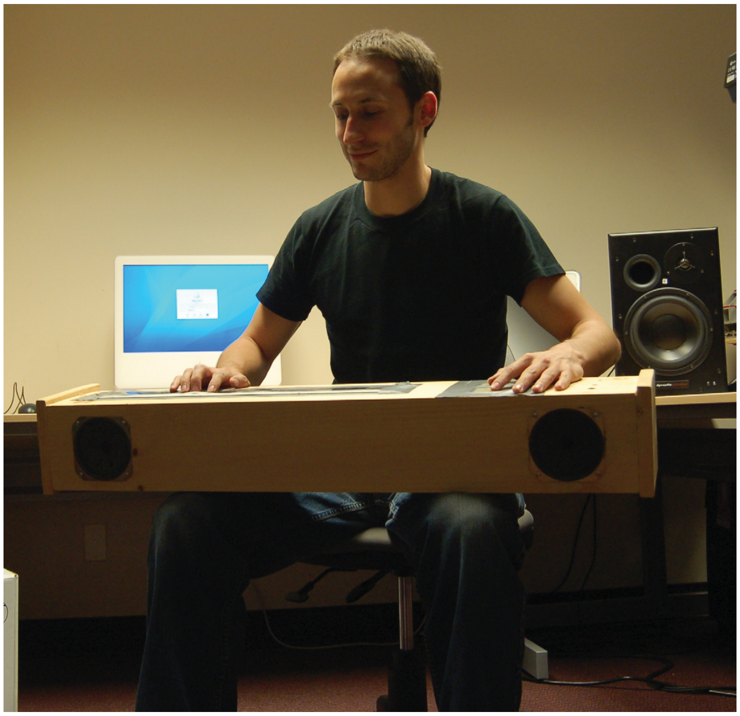

Figure 6.4: The Viblotar in the playing position.

there are also peaks in the spectrum at non-harmonic frequencies. These peaks are due to the flute portion of the hybrid guitar/flute model used in the blotar synthesis. Similar to the guitar, the lower frequencies are above the threshold of detection. Unlike the guitar, a number of higher frequencies are also well above the threshold of detection.

Examining both spectra, it can be seen that both instruments produce vibrations above the threshold of detection. There are also a number of similarities in the spectra, each producing detectable vibrations at a number of frequencies which are harmonics of the note being played.

Having examined the vibrations produced by these instruments, we can see that both produce vibrations which would be felt by the performer. Also, the

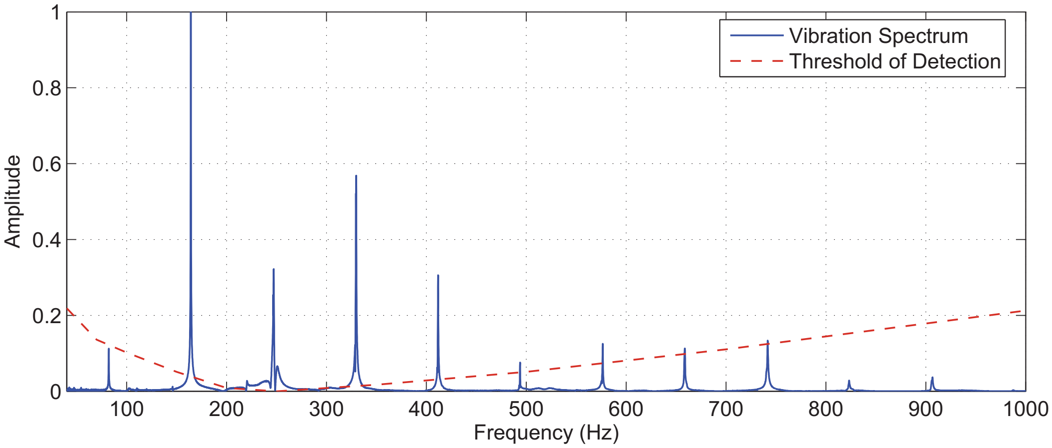

Figure 6.5: Average vibration spectrum of an acoustic steel string guitar playing open low E (82 Hz), as measured near the bridge.

vibrations produced by the Viblotar are similar to those produced by an acoustic instrument. This then raises the question of whether these vibrations affect the *feel* of the Viblotar for the performer. The experiment described in the next section attempts to deal with this question.

6.4 Experiment: Performer Evaluation

Chapters 4 and 5 dealt with the use of sensors and feedback in digital musical instruments. There remains however the task of evaluating this research within a complete digital musical instrument. This section describes an experiment which attempted to evaluate the effects of this work on the *feel* of the instrument. While the concept of the *feel* of an instrument is one which is often mentioned by performers it is difficult to objectively evaluate. Therefore, for this experiment a measure of the *feel* of the instrument is determined based on a number of different characteristics, which participants are asked to rate:

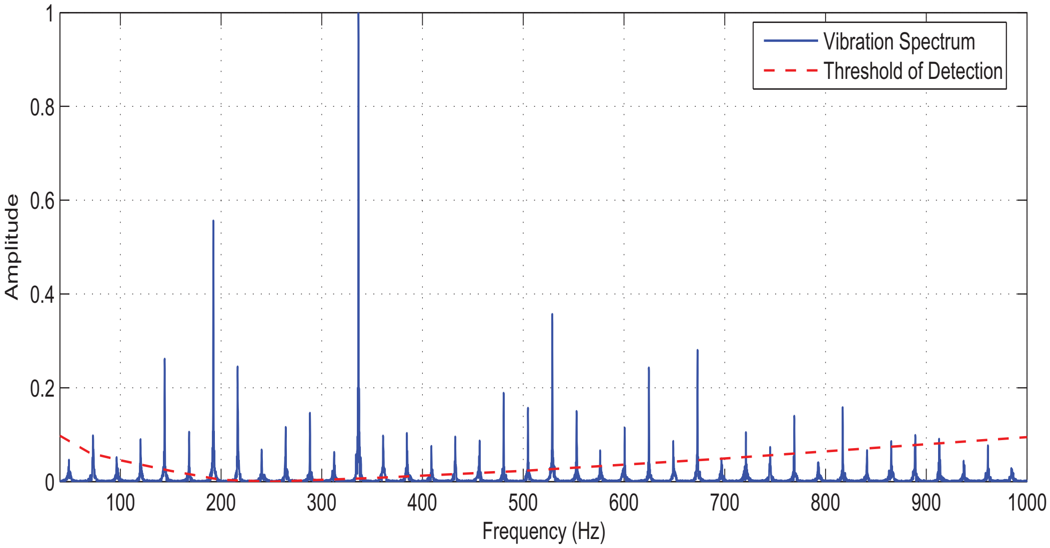

Figure 6.6: Average vibration spectrum of the Viblotar playing a frequency of 82 Hz, as measured on the top

Ease of use: how easy the instrument is to perform with.

Controllability: how much the performer was in control of the instrument.

Engagement: how much of the performer's attention was put into playing the instrument.

Entertainment: how entertaining the instrument is.

Potential for further performance: how much potential the instrument offers for further performance.

6.4.1 Participants

The participants were 5 graduate students from McGill University. All participants were experienced musical performers, having completed at least an undergraduate degree in music performance. Two of the participants had previous experience playing digital musical instruments, while the others did not. None of the participants were familiar with the Viblotar.

6.4.2 Design and Materials

The aim of this experiment was to examine how the choice of sensors and feedback affected the *feel* of the instrument. To evaluate this we asked performers to play the Viblotar in two different configurations:

1. With external sound production and no vibrotactile feedback.
2. With internal sound and vibrotactile feedback production.

In the external sound production configuration, the synthesized sound is output using a pair of loudspeakers which are placed in front of the performer at a distance of 1 meter. This removes all vibrotactile feedback from the instrument and dissociates the sound from the instrument itself. The result of this is a configuration like existing digital musical instruments.

With the internal sound production, the sound is produced using the two loudspeakers which are in the body of the instrument itself. This results in vibrotactile feedback to the performer and in the sound coming from the instrument in a way most like an acoustic instrument. For both configurations the sound volume was maintained at the same level (90dB peak, A-weighted), measured using a Radio Shack 33-2055 digital SPL meter.

These configurations allow for an examination of the effects of vibrotactile feedback and embedded sound production on performer ratings of the instrument.

In addition to examining the effects of vibrotactile feedback, this experiment also offers the opportunity to evaluate the choice of sensors used in the Viblotar. The sensors in this instrument were chosen based on the results of the experiment described in Section 4.2. The aim of that experiment was to determine the sensor which were easiest to use for the tasks of note selection and modulation. High ease of use ratings for the Viblotar would offer validation of the results of that experiment.

Overall, the hypotheses for this experiment are:

1. The Viblotar should be easy to use. Performer ease of use ratings of the Viblotar should be high in both configurations.
2. Vibrotactile feedback should improve the *feel* of the instrument. Some performer ratings should be higher for the internal sound production configuration.

6.4.3 Procedure

Subjects arrived at the lab and were given an Information/Consent form to read over and sign. Subjects were then introduced to the Viblotar and its playing interface. The sensors used on the Viblotar were explained, along with the parameters that they control. They were then given a demonstration of playing the instrument.

Subjects were informed that they would be playing the instrument in two different configurations. They were not told what the difference between each configuration was. They were told that for each configuration they would be allowed to play the Viblotar for 20 minutes and then asked to rate the instrument on several criteria. They were shown the list of criteria and each item was explained to them. The order of presentation of the configurations was randomized. All ratings were performed on a 5-point Likert scale.

Participants then spent 20 minutes performing with the instrument in the first configuration. Once the time was up, they rated that configuration on each of the criteria being examined. This process was then repeated for the second configuration.

Finally, participants were debriefed verbally after the experiment and asked for any comments they had on the instrument or either configuration. The differences between each configuration was also explained at this point.

6.4.4 Data Analysis

Results were analyzed in Matlab. As the data was found not to follow a normal distribution the analysis was performed using the Wilcoxon signed rank test.

6.4.5 Results

Differences between Configurations

There was a marginally significant improvement in engagement for the configuration with vibrotactile feedback [p = .07] (Figure 6.7). This was the only significant difference found in this experiment. However, there were also two non-significant differences found between configurations.

Firstly, there was a slight improvement in entertainment ratings for the vibrotactile feedback configuration [$M_{without} = 3.0$, $M_{with} = 3.4$] (see Figure 6.8). In contrast to this, there was a slight deterioration in ratings of the controllability of

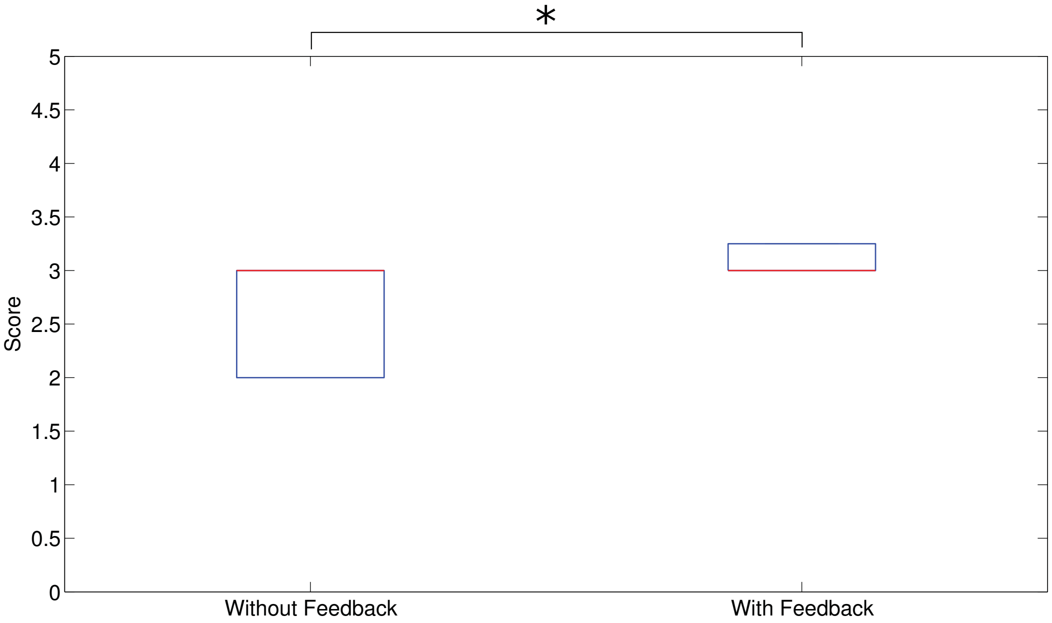

Figure 6.7: Participant ratings of engagement with the Viblotar, with and without vibrotactile feedback. A * indicates a significant difference. Red lines indicate median values, while blue lines indicate lower and upper quartile values. Whiskers extend to 1.5 times the interquartile range.

the instrument for the vibrotactile feedback configuration [$M_{without} = 3.8$, $M_{with} = 3.4$] (see Figure 6.9).

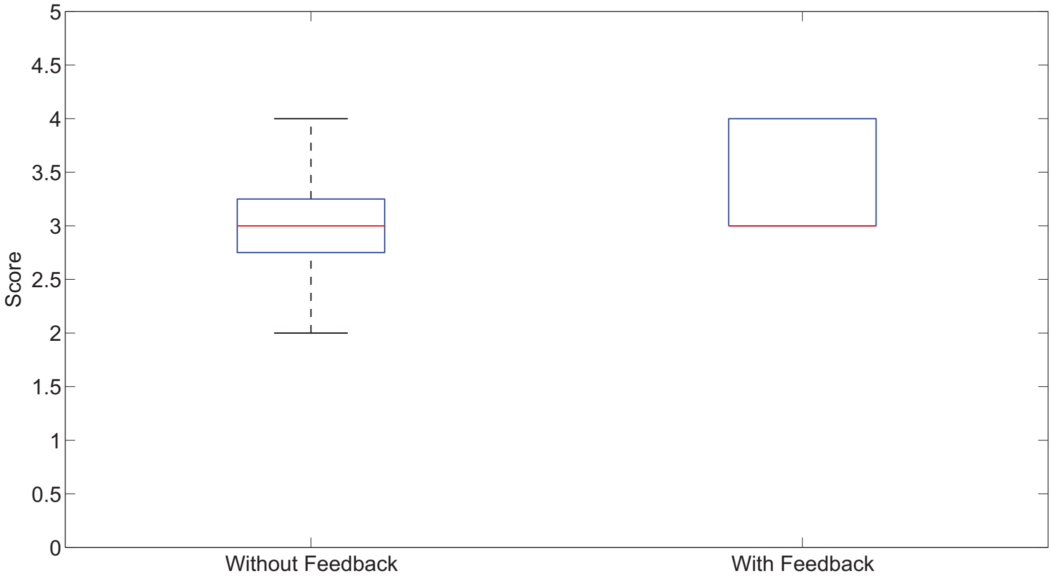

Figure 6.8: Participant entertainment ratings of the Viblotar, with and without vibrotactile feedback. A * indicates a significant difference. Red lines indicate median values, while blue lines indicate lower and upper quartile values. Whiskers extend to 1.5 times the interquartile range.

There were no significant differences in user ratings of the configurations for ease of use or potential for future performance.

Overall Ease-of-Use

Overall, ease of use ratings were high [$M_{ease} = 4.4$]. The ratings were identical for both configurations. In fact, each participants gave the same rating to both the configurations.

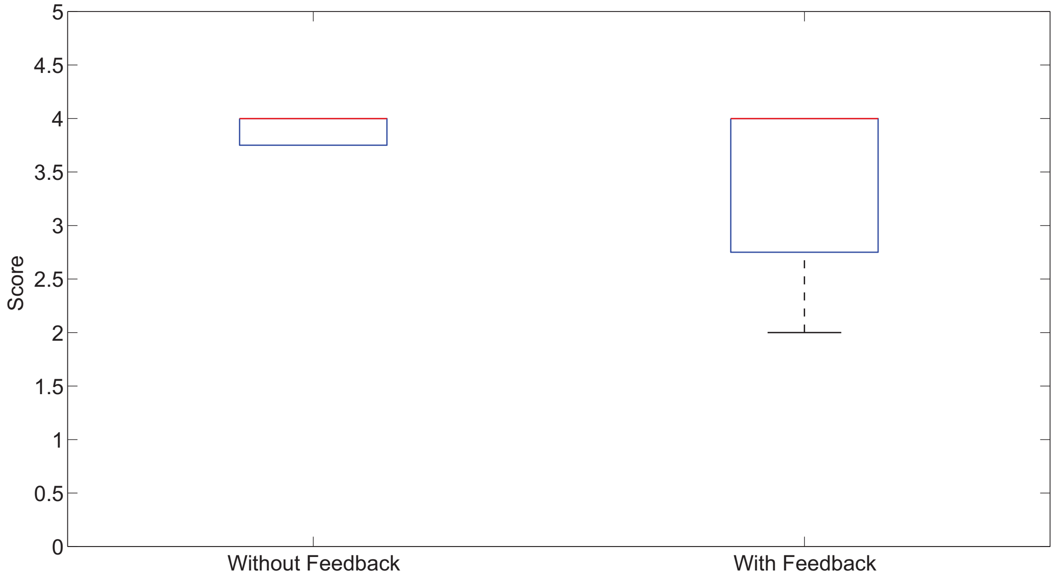

Figure 6.9: Participant ratings of the controllability of the Viblotar, with and without vibrotactile feedback. A * indicates a significant difference. Red lines indicate median values, while blue lines indicate lower and upper quartile values. Whiskers extend to 1.5 times the interquartile range.

6.4.6 Discussion

This experiment provided an opportunity to validate the results of the work presented earlier in this thesis within the context of the performance of a digital musical instrument. While the experiment is only preliminary, the results do show some promise.

Firstly, the ease of use ratings for both configurations were high. A mean ease of use of 4.4 out of 5 was received by each configuration. This indicates that the sensors chosen provide an easy to use interface, as was suggested by the results of the experiment in Section 4.2. The fact that each participant gave the same ease of use rating for both configurations would also seem to confirm that this result is due to the combination of sensors, gestures and tasks, as it was unaffected by the presence or absence of vibrotactile feedback.

Looking at the effects of vibrotactile feedback, we find a number of criteria which change when this feedback is present. Firstly, there was a marginally significant improvement in engagement when feedback was present [t(4) = 2.45, p = .07]. Participants found themselves more engaged with the instrument when vibrotactile feedback was present. They were more involved in the performance of the instrument, spending more of their attention on the instrument.

Interestingly, participant rating of controllability dropped with the addition of vibrotactile feedback [$M_{without} = 3.8$, $M_{with} = 3.4$]. Participants felt less in control of the instrument when the feedback was present. One participant commented on noticing changes in the sound for the internal sound production configuration that had not been noticed for the other configuration. This could indicate that the vibrotactile feedback channel was providing extra information to the performers that was not present in the other configuration, so that they noticed changes which they would otherwise have missed. Such extra information could be extremely useful for developing expert performance technique. It is also possible to consider that a reduction in controllability might result in an increase in the challenge involved in performing the instrument. This could have an effect on the overall performance potential of the instrument in the longer term.

Finally, there was a small increase in entertainment ratings for the configuration with internal sound and feedback generation [$M_{without} = 3.0$, $M_{with} = 3.4$]. Together with the significant increase in engagement this would seem to indicate that the playability, or indeed the *feel* of the instrument is improved when vibrotactile feedback is present.

6.5 General Discussion

The Viblotar provides some interesting insight into the work which was described in Chapters 4 and 5. The first experiment described in Chapter 4 provided some indication of the ease of use of certain sensors for the tasks of pitch selection and pitch modulation in a digital musical instrument. The remaining experiments in that chapter further examined the pitch modulation task. The Viblotar was designed following the results of these experiments.

The Viblotar is controlled using gestures for pitch selection and pitch modulation. The pitch selection gesture also incorporates note triggering and amplitude control within it. By selecting controls that correspond with those musical functions examined in the experiments of Chapter 4, it is possible to further examine those results within the context of musical instrument performance.

The expectation arising from the results of Chapter 4 is that the Viblotar would be an easy to use instrument. This expectation is upheld by the results presented in this chapter. Performer ease of use ratings for the Viblotar had a mean of 4.4 out of 5. This would seem to validate the results of Chapter 4. By careful evaluation of the specific sensors and tasks for a digital musical instrument, we can produce an instrument which is easy to use, offering a “low entry fee” (Wessel and Wright, 2002).

However, this then leaves the issue of expert performance. As well as the “low entry fee” already mentioned, Wessel and Wright (2002) state that a digital musical instrument should also have “no ceiling on virtuosity”. Instruments which are too easy to use may seem more like toys and less like instruments. Hunt (2000) found that users enjoy performing with instruments which offer more of a challenge. For the Viblotar, the addition of vibrotactile feedback resulted in reduced controllability ratings. This might indicate that the instrument becomes more challenging with the feedback present, as it provides more information about the state of the instrument to the performer.

The addition of internal sound generation to the Viblotar produced a number of effects. It localized the sound to the instrument itself and it added vibrotactile feedback to the instrument. Looking at the results of the experiment in Section 6.4, we can see that this resulted in a marginally significant increase in engagement, along with a small (although not significant) increase in entertainment. This would seem to indicate that there is an improvement in the feel of the instrument for the performer when vibrotactile feedback is present.

Interestingly, the additional vibrotactile feedback also resulted in a slight (and again not significant) decrease in performer controllability ratings. In post-experiment debriefing, one of the participants explained that they thought the sound synthesis had changed between configurations. On further examination it was discovered that the participant had noticed changes in the sound under the vibrotactile feedback configuration which had not been noticed under the other configuration. More information was being presented to the performer by the extra feedback channel. It seems that this extra information was causing the performer to feel less in control of the instrument than in the other configuration.

However, a number of issues still remain to be addressed. A longer term evaluation, perhaps with more participants, could lead to much insight into the playability of the Viblotar. Keele (1973) states that vibrotactile feedback is used more by expert performers than beginners. As the participants in the experiment in this chapter were all novice Viblotar players, it is possible that they were not making use of the vibrotactile feedback in the same way as an expert performer would. A longer term experiment examining the changes in user ratings over a longer period of time would allow the participants to increase their skill with the instrument. Such an experiment might also lend insight into the effects of the vibrotactile feedback on the *feel* of the instrument, through changes in participant ratings over time.

Also, further work with performers and composers resulting in musical performances could also offer useful information on the design of new digital musical instruments. In fact, this issue is the focus of the next chapter, which discusses a number of instruments and interfaces developed as part of large collaborative projects. These instruments were developed in collaboration with composers and performers and have been used in a number of performances of new musical works.

Overall, the Viblotar has provided a useful testbed for examining the results of the work presented in earlier chapters of this thesis. In particular, it seems to provide validation for the experimental results presented in this work when applied to the construction of a new digital musical instrument.

6.6 Conclusions

This chapter described the Viblotar, a new digital musical instrument designed to allow for musical performance-based evaluation of the results of the work described in earlier chapters of this thesis. The Viblotar was designed following the results of the work presented in Chapters 4 and 5. Experimental evaluation of the Viblotar offers validation of the methods of evaluating sensors for specific musical tasks which were presented in Chapter 4, as well as the use of vibrotactile feedback systems described in Chapter 5.

The next chapter presents a number of instruments and interfaces which were developed both as part of the work for this thesis and also within the framework of large, collaborative, interdisciplinary research projects. These instruments provide further insight into the design of the physical interfaces for digital musical instruments. In particular, the next chapter deals with issues such as instrument aesthetics, reliability and use in musical performance.

Chapter 7

Collaborative Development of Digital Musical Instruments

This chapter details a number of digital musical instruments which were developed during the course of the research as part of larger collaborative projects. I begin with a discussion of the context of the development of these instruments. This is followed by detailed descriptions of the instruments themselves. Finally, I present a discussion of a number of important issues discovered during this work, including issues such as reliability of instruments, instrument aesthetics and the evolution of instruments through a collaborative design process.

7.1 Context

Within this chapter I will discuss the development of 4 new digital musical instruments, which were developed as part of two larger projects. These projects involved collaborations between researchers, composers and performers on the development of systems for use in live musical performance. This section will provide details of these projects, including their aims, the participants and the scope of the work I performed as part of the projects.

7.1.1 The McGill Digital Orchestra

The McGill Digital Orchestra was a research/creation project supported by the *Appui à la recherche-création* program of the *Fonds de recherche sur la société et la culture (FQRSC)* of the Quebec government and the *Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT)*. It involved collaboration among researchers, composers and performers within the Schulich School of Music at McGill University. The goal of the project as stated on the project website was[1]:

> "to develop new creative resources that allow composers and performers to expand and renew their artistic practice through the interaction of live performance and digital technologies, and to utilize these tools in the composition and performance of a number of new works"

The project had a duration of 3 years and culminated in a performance of new musical works during the 2008 MusiMarch festival in Montreal. In total 21 people were involved in this project, including 6 faculty members and 15 graduate students.

This project resulted in the development of 5 new digital musical instruments:

- Two T-Sticks, a tenor and an alto. The basic design of the T-Stick was described in Chapter 3.

---

[1] http://www.music.mcgill.ca/digitalorchestra

- The Rulers, an instrument developed that evokes the gesture of plucking or striking a ruler that is fixed at one end.

- The T-Box, which makes use of ultrasound sensing to detect the performer's gestures in the air above the instrument. This will be described in Section 7.3.

- The FM Gloves, a pair of custom designed datagloves, which will be described in Section 7.2.

These instruments were developed using an iterative and collaborative design process, with members of the research project collaborating on the physical interface design, mapping, software development, synthesis and performance techniques for each instrument. The instruments were then used in the creation and performance two new musical works, *The Long And The Short Of It* by Heather Hindman and *sounds between our minds* by D. Andrew Stewart.

In particular, my work on this project concentrated on the development of 2 instruments, the T-Box and the FM Gloves. This included physical interface design and hardware and software development. These instruments were used in the performance of Heather Hindman's *The Long And The Short Of It*, a piece for cello and digital musical instruments. The performers involved in this piece were Erika Donald (Cello), Xenia Pestova (FM Gloves) and Fernando Rocha (T-Box). Full credits for this piece and D. Andrew Stewart's *sounds between our minds* can be found in Appendix B.

7.1.2 Gesture Controlled Sound Spatialization

The second collaborative project was titled *Compositional Applications of Auditory Scene Synthesis in Concert Spaces via Gestural Control* and was jointly funded by the *Canada Council for the Arts* and the *Natural Science and Engineering Research Council (NSERC)* of Canada. The aim of the project was to develop novel compositional and technological methods for the advanced use of the multidimensional nature of auditory space in music composition.

The project involved 4 faculty researchers (3 from McGill and 1 from Rensselaer Polytechnic Institute), a post-doctoral researcher and 5 students. This project also had a duration of 3 years and culminated in the performance of a new work for small ensemble, live electronics and gesture controlled spatialization at the 2008 MusicMarch festival.

There were 3 main areas of research within this project:

- Development of a new multichannel sound reproduction system.
- Design of new gesture control techniques for the optimal control of the multidimensional spatial attribute.
- Examination of the cognitive aspects of auditory scene analysis and developing new auditory scene synthesis systems.

My work for the project was within the scope of the second area: the design of new gesture control techniques. This involved the creation of a number of new interfaces for the gesture control of spatialization, as well as the evaluation of existing interfaces and metaphors. Two of these new interfaces were used in the final concert: a system for the virtual direct manipulation of sound sources using hand gestures (described in Section 7.4) and a system for non-conscious control of spatialization using cello performer gestures (described in Section 7.5).

This project resulted in two new musical pieces, both written by Sean Ferguson. The first, *Miroirs*, was performed at the 4th International Conference on Enactive Interfaces, in Grenoble, France, in November 2007. This piece was for solo cello and gesture controlled spatialization and was played by Chloé Dominguez. A physically-modeled virtual dancer, created by Chi-Min Hsieh and Annie Luciani, was also part of the performance. This piece made use of the cello performer gesture system described in Section 7.5.

The second piece was performed at the 2008 MusiMarch festival in Montreal. This piece, titled *Ex Asperis*, was also composed by Sean Ferguson and was for solo cello, gesture-controlled spatialization, live electronics and small ensemble. Both the cello performer gesture system (again performed by Chloé Dominguez) and the sound source manipulation system (performed by Fernando Rocha) were used in this piece. The McGill Contemporary Music Ensemble, directed by Denys Bouliane, provided the ensemble portion of the performance. Full credits for this piece and the project itself can be found in Appendix C.

7.2 The FM Gloves

The FM Gloves (or Fortier-Marshall Gloves) were originally designed by Pierre-Yves Fortier, for use in performance as a DJ-style interface for the control of loops and samples. For the Digital Orchestra project it was decided to work on a modified version of these gloves, which could be used to control a variety of synthesis systems. There were a number of specific aims in the development of the new FM-Gloves:

Optimize the range and reliability of the sensors used. A major issue with the use of sensor gloves and bodysuits is the reliability of the sensors used. Sensors such as bend sensors and FSRs can be used to provide interesting information about performer gestures when attached to joints or fingertips. However, these sensors were generally not designed for such use and can be placed under severe stress by doing so (Interlink Electronics, Inc.). The aim in this case was then to modify the gloves in such a way as to improve the reliability of the sensors, through careful mounting and connection of the sensors.

Maximize the number of available continuous control variables. The original gloves contained an array of buttons worn on the performers left wrist. These buttons were intended to be used to change the state of the synthesis system, such as by changing the current sample or effect being played. For the Digital Orchestra project, we were more interested in the use of continuous rather than discrete control of a variety of variables in the sound synthesis system. As such, the decision was to remove the buttons and add a number of continuous sensors in their place.

Allow for larger-scale control gestures than initially used. Originally, the gloves were designed to allow for interaction on the performer's body, or other surfaces such as a tabletop. The sensing on the hands consisted of FSRs on the fingertips and some bend sensors on the fingers. Additionally, there was an array of buttons on the left arm and a linear position sensor on the torso, which was manipulated with the hands. This resulted in a small control space and small gestures which could make it difficult for an audience to perceive what the performer was doing. Therefore, we decided to add sensors for large scale control gestures such as arm movements.

Provide active vibrotactile feedback. As discussed in Chapter 2, non-contact immersive instruments such as the FM Gloves have a reduced number of feedback channels to the performer. For many of the control gestures used, there is little or no haptic and tactile feedback. As research has already shown that the addition of vibrotactile feedback can be used to provide an extra channel of information to the performer about the instrument, we chose to implement some active vibrotactile feedback in the FM Gloves. The choice of this feedback was based on input from the performers after testing initial prototype configurations of the FM Gloves.

7.2.1 Physical Interface

The FM Gloves consist of two black skin-tight gloves which are augmented with a number of sensors. These sensors are connected to a small beltpack, which digitizes the signals and transmits them wirelessly to a Kroonde Gamma system [2]. The use of a wireless sensor system allows the performer to have greater freedom of movement than is available with wired systems. Figure 7.1 shows the FM Gloves.

A variety of different sensors are used on the FM Gloves, with different arrangements of sensors on each hand. For the left hand, there are two types of sensors used: FSRs and a 2-axis accelerometer. Four small FSRs are mounted on the tips of the glove's fingers, to measure fingertip pressing force. The 2-axis

---

[2] See http://www.la-kitchen.fr/kitchenlab/kroonde-en.html for more information.

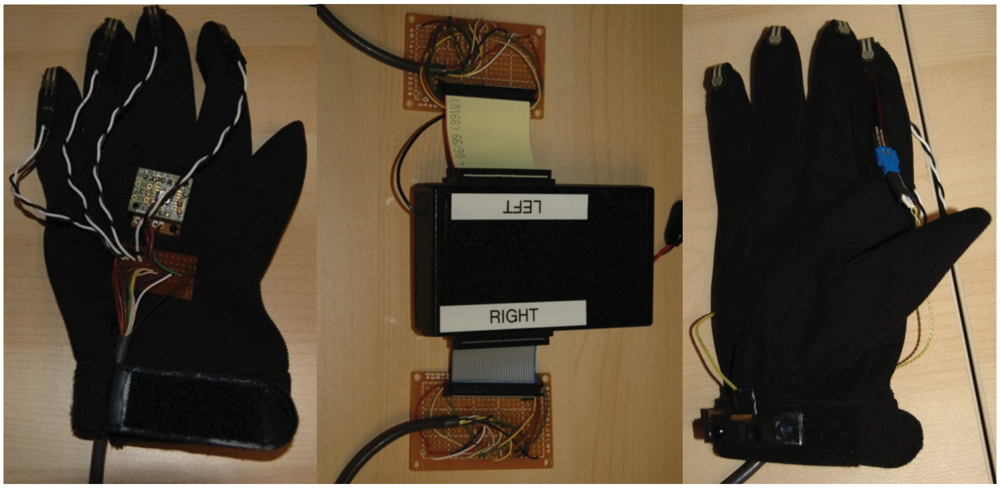

(a) Left Glove (b) Interface (c) Right Glove

Figure 7.1: The FM Gloves.

accelerometer is placed on the back of the hand and is used to measure inclination in two dimensions.

For the right hand four small FSRs are again placed on the fingertips. In addition, there is a bend sensor attached to the underside of the index finger. This is used to measure the amount of bend in this finger. Finally, an infrared distance sensor is mounted on the inside of the wrist. This is used to give a measure of the distance between the right hand and the performer's body.

Vibrotactile feedback is provided using a small vibrating motor also attached to inside of the right wrist. The motor was chosen for a number of reasons. While with the Viblotar the aim was to produce vibrotactile feedback directly from the sound synthesis, the aim here is to provide additional information about the state of the instrument, without additional sound production. Only a single channel of information is required, which can be provided using the motors speed of vibration[3]. Also, the motor is small and light, but can provide a large amplitude vibration. It can be driven with a simple DC voltage signal, available from the sensors on the gloves. Finally, it does not require any cable connection to the computer, or any additional amplifiers or power sources and so does not weigh down or restrict the movements of the performer.

7.2.2 Synthesis

The synthesis system used for the FM Gloves by the Digital Orchestra is based around the software Logic Pro from Apple, Inc. In particular, it is created using a combination of software instruments, synthesizers, effects and plugins that are available with Logic Pro 7. The basic sound synthesis is performed using the Sculpture instrument, which generates a variety of physically-modeled plucked string sounds. A variety of effects are then added to the sound produced by Sculpture, including delay and reverb. Additionally, a number of audio samples are also mixed into the output sound.

The synthesis system for the FM Gloves is much more complex than that used in the Viblotar (which was described in Chapter 6). Much of this complexity is due to the need to meet specific sonic goals of the composer for the instrument. The synthesis system and the mapping (described in the next section) were developed by composer Heather Hindman, specifically for her composition for the project.

As with the Viblotar, physical modeling synthesis forms the main part of the synthesis system, due to the ease of mapping of gestural parameters to the meaningful parameters offered by such a synthesizer.

---

[3] As noted in Chapter 5 the speed of the motor will affect both the amplitude and frequency of vibration at the same time.

7.2.3 Mapping

The mapping system for the FM Gloves is also more complex than that of the Viblotar. It uses a multi-layered mapping system based on that described by Malloch et al. (2008). Various parts of this system were developed by myself in conjunction with Stephen Sinclair, Joseph Malloch and Heather Hindman.

Parameters from the glove controllers are read into the Max/MSP software as Open Sound Control messages. Within Max/MSP, various operations are carried out on the parameters, such as linearization, scaling and smoothing. For the FSRs, in addition to the pressure values, the software recognizes striking gestures similar to those on the Viblotar. When the performer quickly strikes one of their fingers against a surface, a trigger is generated with a velocity value based on the force of the strike. These triggers, together with the other parameters extracted, provide a list of semantically meaningful gesture parameters.

These gesture parameters are then mapped onto musically meaningful parameters of the synthesis system. This includes a variety of one-to-one, one-to-many and many-to-one mappings. These mappings are accomplished using the Digital Orchestra Toolbox and the Mapper, developed by Joseph Malloch and Stephen Sinclair (Malloch et al., 2008).

The resulting mapping allows the performer to generate plucked-sound pitches with striking gestures using the fingers of the right hand. The pitch of the note is dependent on the finger struck. The range of pitches available is determined by the amount of bending of the index finger of this hand. The FSRs on the left hand act as continuous controls for the different effects applied to the sound, including reverb, delay and damping.

The two axes of tilt from the accelerometer on the left hand are used to change the “material” parameter in the physical model, which changes the timbre of the sound through a range including sounds like nylon strings, metal strings, wooden bars and glass bars. These axes are also used to control a number of filters which are processing the output sound.

Finally, overall volume level is controlled using the infrared distance sensor on the inside of the right wrist. By varying the distance of this sensor from their body, the performer can control the volume. The volume varies from low levels near the performer’s body to high levels at the limit of extension of the performer’s arm. The vibrotactile feedback from the motor is also directly mapped to the distance of the right hand from the performer’s body, using the infrared sensor.

7.2.4 Interaction

Interaction with the FM Gloves involves a number of different types of gesture. In keeping with one of the design goals, there are both small scale and large scale gestures. Small scale gestures include pressing the fingers against the thumb and body. Large scale gestures include the movement of the right hand towards and away from the body to control the volume and the rotation of the left hand in the space in front of the performer to control timbral aspects of the sound.

To create sound, the performer must make a plucking or striking gesture with one of the fingers of their right hand. The pitch of the note generated depends on both the finger used, and the amount of bend of the index finger of the right hand. The timbre of the sound is controlled using the tilt of the left hand and the pressure on the left hand fingers. These gestures can be used to set the timbre

Figure 7.2: Performer Xenia Pestova practices with the FM Gloves

before the sound is created and to modify the sound once it is playing.

Volume is controlled using the distance of the right hand from the body, with a lower volume when close and a higher volume when far away. This gesture was chosen to be representative of the size of the sound, where a small, closed-in gesture represents a small (and therefor quiet) sound and a larger, more open gesture represents a larger sound. The aim was for this gesture to be intuitive for both the performer and the audience.

The choice of gestures and mappings for interacting with the FM Gloves was arrived at using a collaborative process. Numerous testing sessions took place, involving myself, the composer and performers. The result of this is that the current methods of interaction with the FM Gloves evolved based on the input of all those involved in the project. Figure 7.2 shows performer Xenia Pestova practicing with the FM Gloves.

7.3 The T-Box

The T-Box (formerly known as the Tralf) is a third generation instrument, based on initial versions first by Geof Holbrook and later by Geof Holbrook and Eileen TenCate. It makes use of ultrasound sensing to detect the performer's gestures in the air above the instrument.

The T-Box hardware has gone through a number of stages over the course of its development. The basic design uses ultrasonic transmitter and receivers and senses the distance between them using the amplitude of the signal at the receiver, with the incoming signal passed through an envelope follower to extract the amplitude.

In the first two versions of the instrument, both the generation of the 40 kHz square wave for the transmitters and the envelope following of the received signal was performed using analog circuitry. This resulted in some issues regarding the stability of the square wave and also the response of the envelope follower. The envelope followers attack and decay coefficients were fixed, while the oscillator exhibited some drift and required regular retuning.

The design goals then for the current version of the T-Box, which was built for the Digital Orchestra project, were as follows:

Improve stability of ultrasound sensing. As already mentioned, the analog circuitry used to generate the ultrasonic signal in the original Tralf instrument was unstable. It required regular tuning of the square-wave oscillator. Therefore, one design goal for the project was to improve the stability of the ultrasonic signal by migrating from analog circuitry to digital microcontroller-based circuitry.

Allow for control of envelope following. The measurement of the distance of the hands above the instrument was performed by envelope-following the output of the ultrasound receivers. Again in the original Tralf this was performed using analog circuitry. The attack and decay times for the envelope follower were set by the values of a number of resistors and capacitors in this circuitry, meaning that they could not be controlled during performance. Therefore, we decided to also migrate the envelope following to software on the microcontroller and to allow for control of the attack and decay time by the performer.

Improve the aesthetics of the instrument. One issue which has not been discussed much in the literature in relation to the design of DMIs is the aesthetics of the instrument body. Many digital musical instruments look like prototypes, with visible wires and electronic components. This can make the instrument look fragile and reduce the appeal of the instrument for people to perform. To counteract this, an improved aesthetic design of the T-Box was made a major design goal.

7.3.1 Physical Interface

The T-Box's physical form is that of a rectangular wooden box, mounted on top of a microphone stand. Four ultrasound receivers are mounted on top of this box. The control for the envelope follower is one the back, facing the performer. Figure 7.3 shows the T-Box, including the hand pieces.

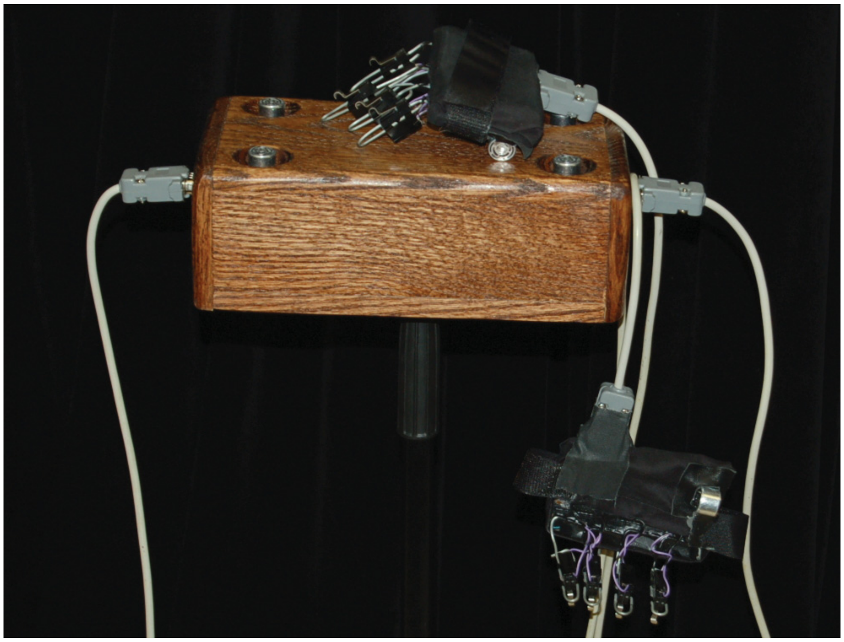

Figure 7.3: The T-Box with hand pieces, set up on a microphone stand for performance.

Two hand-pieces are connected to the sides of the box by flexible cabling. Each hand-piece is worn by the performer and contains an ultrasound transmitter and 4 switches, one for each finger. The ultrasound transmitter faces downwards, perpendicular to the fingers. Figure 7.4 shows a close-up of one of the hand-piece.

The electronic systems for the T-Box are inside the box itself. These include a microcontroller to generate the 40 kHz square wave for the ultrasound and a second (more powerful) microcontroller which converts the signal from the ultrasound receivers to a digital value. An Arduino Mini is used as the second microcontroller, performing analog to digital conversion of the voltage from the receivers and envelope following on this signal. It also detects presses on the switches on each hand piece. All of this data is then sent to the computer as serial data over a USB

Figure 7.4: One of the T-Box hand pieces, showing the ultrasound transmitter and the finger switches.

connection, using an Arduino Mini USB adapter. Data is transmitted over USB at a rate of 115200 bps, which allows an update rate of $\sim$ 200Hz with the 10-bit resolution used by the microcontroller.

7.3.2 Synthesis

As with the FM Gloves, the synthesis system for the T-Box was based around Apple, Inc.'s Logic Pro. The physical modeling synthesizer Sculpture provided the basic sound synthesis. The sounds produced by this were then run through a large number of effects, including a modulating delay and a number of filters.

The system allowed for the creation of 12 semitone pitches per octave over a range of 5 octaves. Timbral shaping was accomplished using the "material" control in sculpture in combination with delays and filters. A vibrato was also applied to the pitches being generated within sculpture. Finally, a global volume control was provided. Again, the synthesis system was created by composer Heather Hindman.

7.3.3 Mapping

The mapping system consisted of the same multi-layered structure used for the FM Gloves. Once again, the Digital Orchestra Toolbox and Mapper were used to allow for easier mapping of physical parameters to synthesis parameters.

The incoming serial data is read into Max/MSP. This data is then processed in a number of ways. Firstly, the signals from the ultrasound receivers (or more correctly the envelopes of these signals) are linearized and converted to values which are directly proportional to the distance between the transmitter and receiver[4]. This provides a meaningful physical value which can be mapped to synthesis parameters.

The binary data from the switches on the hand pieces is also processed. This data is used to calculate fingering patterns, similar to those used on some wind instruments. Each switch is treated as a single binary digit and a number representing the current fingering pattern is generated for each hand from the combination of these bits. This allows us to generate a possible 16 fingerings with each hand, although some of these are difficult for the human hand to make and are not used. In the end, a total of 12 patterns per hand are used. This allows the right hand to produce all of the semitone pitches in an octave, while the left hand switches are used to change register and also to turn on and off certain effects and filters.

The parameters of the physical model and also of the effects are controlled by

---

[4] Note that these signals are not linear without processing and that the output is affected by both the distance between the transmitter and receiver and also by the directivity of the transmitters. This means that a higher output will be received when they are closer to each other, or when they are pointed most directly at each other.

the ultrasound receivers. One receiver controls the "material" control to modify the timbre of the physical model. Another modifies the vibrato rate and also (in combination with certain fingering patterns on the left hand) the amount of delay added to the signal. A third controls the cutoff frequency of a filter, within a range defined by the currently selected register. The fourth receiver controls the vibrato depth.

Finally, a global volume control is provided based on the sum of the signals from all four receivers. In practice it is possible to activate one or two receivers with each hand, allowing all four to be manipulated with two hands. The final mapping was designed collaboratively by myself, Heather Hindman, Stephen Sinclair and Joseph Malloch.

7.3.4 Interaction

When playing the T-Box, the performer stands facing the instrument, which is placed at waist height by adjusting the stand. On the face of the box nearest the performer is a control knob which can be used to change the attack and decay of the envelope follower, along for a range of effects from staccato to sustained.

The performer moves their hands, wearing the hand pieces, in the air above the instrument. Parameters change from low values when far from the instrument, to higher values when closer. This includes the volume, which is highest when closest to the box. Interestingly this is the opposite of how the volume control works on the FM Gloves, but was chosen as most natural by the group. Figure 7.5 shows the T-Box being played.

The original intention was for the performer to manipulate a single parameter

Figure 7.5: Stephen Sinclair demonstrating the T-Box at Wired magazine's NextFest 2007 in Los Angeles, USA.

at a time with each hand. This would be accomplished by moving the hand in the vertical axis above a single receiver. However, during the design process it was found that interesting effects could be achieved by tilting the hands in such a way as to activate multiple controls at once. We also found that performer's used gestures such as rapidly shaking their hand from side to side over a receiver, so as to generate an intermittent signal. Finally, other unexpected performance gestures discovered during the development process included circling over all of the receivers with one hand and rapidly passing horizontally over a single receiver so as to generate a pulse. These gestures allow for modulation of parameters and for the creation of effects such as re-triggering, which are important for expressive performance.

7.4 Manipulation of Spatial Sound Sources Using Hand Gestures

For the gesture-controlled spatialization project, we identified 3 main roles which are useful for the performance of gesture-controlled spatialization (Marshall et al., 2007). These are:

Spatial Performer - performs with sound objects in space by moving sound sources in real-time using gesture

Instrumental Performers - indirectly manipulate parameters of their own sound sources through their performance gestures on their own acoustic instrument

Spatial Conductor - directly controls large-scale (room and environment) parameters of the spatialization system using gesture

For the final concert as part of this project, we implemented two of these: the instrumental performer and the spatial performer. This section will describe the spatial performer system, with the instrumental performer system described in the next section.

The goal for the spatial performer system was to allow the performer to directly control the positions of sound sources in space. This includes controlling single sound sources and groups of sound sources. One of the most obvious ways to do this was to use hand gestures, resulting in *direct manipulation* (or perhaps *virtual direct manipulation*) system.

To do this, we designed a pair of datagloves using sensors to recognize hand postures and a 6-DOF (Degree-of-Freedom) magnetic tracker system to track the position and orientation of the hands. The goals for this system were:

Allow movement of single sources. The system should allow the performer to grab and move single sound sources within the virtual room.

Allow movement of a group of sources. The system should also allow the performer to select and move all of the sound sources as a group.

Allow manipulation of a group of sources. In addition to moving a group of sources, the performer should be able to perform operations such as rotation and scaling on the group as a whole.

Allow both continuous (current) control and ballistic control. The performer should be able to move and place sources directly (a form of current control) and also to *throw* sources, which will then possess a certain direction and velocity. These sources will slowly lose velocity, due to inertia, until the come to a halt.

7.4.1 Physical Interface

The spatial performer system consists of a pair of sensor-equipped gloves. Each glove has two FSRs, a small round one on the thumb and a larger square one on the palm of the hand. Attached to the back of each hand is a sensor from a Polhemus Liberty motion tracker system. This sensor uses a magnetic field to sense the position and orientation of the hand within 3D space.

AS with the FM Gloves, a cable connects each glove to a beltpack, worn at the base of the performer's back. Unlike the FM Gloves, this beltpack contains an Arduino-based circuit which performs the analog to digital conversion for the sensors and transmits the data over a USB connection to the computer. Separate cables run from the Polhemus liberty sensors to the main Polhemus Liberty unit, which is connected by a USB connection to the same computer.

Two feedback systems are provided for the performer: auditory and visual feedback. Auditory feedback is provided through the loudspeaker array within the room itself. This has been tested with 4-, 8- and 24-channel loudspeaker systems and was used with the 24-channel system in the final concert. Visual feedback is provided using a computer monitor, located in front of and below the performer. The provides a top-down view of the virtual room and virtual sound sources.

7.4.2 Spatialization System

The spatialization system used by this system is ViMic, developed by Braasch et al. (2008). The ViMic-System, which has undergone improvement as part of this project, combines a *tonmeister's* know-how of various sound recording techniques with the knowledge of room acoustics, sound propagation and spatial perception.

The development of the ViMic system for this project was performed by Nils Peters.

ViMic is a computer-generated virtual environment, where gains and delays between a virtual sound source and virtual microphones are calculated according to their distances and the axis orientations of the microphone directivity patterns. Besides the direct sound wave, a virtual microphone signal also contains early reflections and an adequate reverberation tail. This depends both on the sound absorbing and reflecting properties of the virtual surfaces and on the geometry of the virtual enclosed space. Sound sources and microphones can be spatially placed and moved in 3D as desired.

The system architecture was designed to comply with the expectations of audio engineers and to create sound imagery similar to those associated with standard sound recording practice. There are a number of adjustable parameters which can affect this auditory virtual environment, including:

- Position of the sound sources [X,Y,Z]
- Radiation patterns of the sound sources
- Ratio of direct and reverberant sound waves
- Absorption properties of the walls
- Room size [X,Y,Z]
- Position of the microphones [X,Y,Z]
- Directivity patterns of the microphones [continuously adjustable from omni via cardioid to figure 8]

For the implementation of the spatial performer system, we make use only of the ability to change the positions of the sound source, although other parameters are used in the instrumental performer system.

7.4.3 Mapping and Interaction

As already mentioned, the system uses a *direct manipulation* approach to mapping. The performer may select and move a single sound source using a *pinch* gesture, by pressing the index finger against the thumb. This attaches a single sound source to the performer's hand, which will then follow the movement of the hand for as long as the pinch is maintained. Once the performer stops pressing the index finger against the thumb, the source is dropped at the current position.

To manipulate the group of sound sources, the performer makes a *grasping* gesture with both hands. This involves pressing the fingers against the FSR mounted on the palms of the hands. Once this is done, the performer is then manipulating the sources as a group, until they release this gesture. When the performer has grasped the group of sources, they may then move the group by moving their hands. The center-point of the group of sources then follows the movement of the center-point of the hands.

As well as moving the group of sources, the performer may also scale the size of the group by adjusting the distance between their hands. Rotation of the group is also possible, by changing the orientation of the hands relative to each other around the midpoint. Note that control of these parameters (position, scale and orientation) is integrated, allowing the performer to control multiple parameters with one gesture, for instance scaling and rotating at the same time.

Finally, the performer may *throw away* the group of sound sources, by releasing the grasping gesture while in the process of moving or scaling the group. The sources then continue to move in the direction in which they were already moving, with a velocity determined by that of the performer's hands. This can be used to cause all of the sources to leave the virtual room. This effect was found to be compositionally interesting, particularly to indicate the end of a section of the piece. Once the performer is done with the sounds, he throws away the sounds before moving on to the next set.

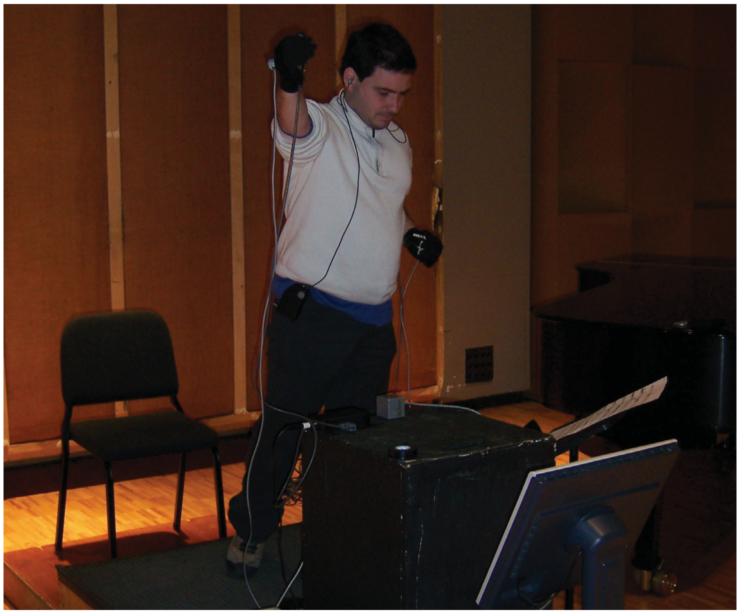

Figure 7.6: Fernando Rocha performing with the system for virtual direct manipulation of spatial sound sources during a rehearsal for the performance of Sean Ferguson's *Ex Asperis*.

Figure 7.6 shows Fernando Rocha performing this system. Note that he is standing on a raised platform on stage, to allow the audience to see the gesture as he makes them. This has proven to be useful, as seeing the gestures allows the audience to better perceive the spatial effects (Marentakis et al., 2008).

Various aspects of the mapping and control system were developed in collaboration with Joseph Malloch, Nils Peters and Marlon Schumacher.

7.5 Non-Conscious Gestural Control of Spatialization

Along with the spatial performer system just described, we also developed a system to allow an instrumental performer to control aspects of the sound spatialization system. This system, developed for cellist Chloé Dominguez by Joseph Malloch and I, allows for a performer to manipulate spatialization parameters *without their conscious control.*

If conscious control is desired, gestures must be chosen such that they can be performed without disturbing the instrumental performance, and it is assumed that the performer has spare attention for this task. For non-conscious control, the mapping relationships between performer movement and spatialization effect becomes an indirect compositional process rather than instrument augmentation or performer interpretation. Rather than asking the performer to deliberately manipulate spatial parameters, the composer or designer must plan instrumental movement with thought to the spatial effect as well as acoustic sound production. Research has shown that such non-concious and ancillary gestures are very repeatable and so offer a useful method of controlling specific effects.

In order to allow for non-conscious control, we needed to extract parameters of the performer's normal playing gestures. The aim was to extract as much data as possible using sensor systems that were as unobtrusive as possible. Overall, the system needed to interfere with normal playing as little as possible. The design goals for this system were therefore:

Maximize information from cello performance gestures. The system needed to extract as much information as possible from the cellist's performance gestures. This meant that the system should track both hands and possibly some body movements.

Minimize interference with playing technique. We required the system to still allow the performer to play using normal techniques, with a minimum of interference from the sensors. This required the sensors to be small and light and all cables to be kept out of the performer's way.

System design should prioritize ergonomics. Any sensors used should be designed so that they do not cause fatigue. They should be easy to attach and remove. Sensor weights should be low. Cables should not restrict movement.

7.5.1 Physical Interface

The instrumental performer system was implemented using an XSens Xbus kit[5]. This is an inertial tracking system, which can provide 3 axes of acceleration, gyroscope and magnetic field data from a number of modules. We used two modules, one on each forearm of the performer. Modules were placed in pouches attached to elasticated arm bands.

---

[5] See http://www.xsens.com/ for more information

Cables run from the modules to a beltpack, worn on the performer's back. Elasticated arm bands on the performer's upper arms were used to hold the cables against the arms, keeping them out of the way when playing. The beltpack transmits the sensor data to the computer, using either a wireless or wired connection. Figure 7.7 shows cellist Chloé Dominguez wearing the modules and arm bands during a rehearsal.

Figure 7.7: Cellist Chloé Dominguez wearing the Xsens Xbus-based modules and arm bands during a rehearsal of Sean Ferguson's *Ex Asperis*.

Initially, we had planned to use the wireless (Bluetooth) connection, as this would reduce the number of cables used and increase the performer's freedom of movement. Unfortunately, issues with reliability of the signal in concert spaces resulted in our using the wired connection. A cable from the beltpack therefore connects the performer to the computer.

Initial prototypes of the system also used a force-sensitive floor under the performer's seat, to allow us to sense movements of the performer's center of mass. This was not however used in the final concert system.

7.5.2 Gesture Tracking

An examination of the angular and velocity information from the sensors during a number of motion capture sessions allowed us to determine a number of different features which could be easily extracted. These were:

- Relative position of playing on the fingerboard, extracted from left arm rotation data.
- Current string being bowed, measured from bowing-arm rotation data.
- Overall bowing energy.
- Overall energy of left hand
- Recognition of some specific performance gestures (e.g. large pizzicato)

These features provide a large amount of information about the performer's gestures, while requiring only 2 modules to be worn.

7.5.3 Mapping and Interaction

For the two concert pieces, *Miroirs* and *Ex Asperis*, we primarily made use of the energy measurements for each hand. An overall performance energy was determined as a combination of the two hand's energy values, with some weighting applied.

These energy levels were mapped to a number of different parameters at different parts of the pieces. This included a mapping of overall energy level to the amount of delay applied to the cello sound before spatialization (and therefor the mix between acoustic and recorded cello sound). Also included was a mapping of the energy of pizzicato playing to both the amount of sampled pizzicato sounds being played and the spread of these sounds into the space using the spatialization system.

These mappings moved the emphasis for control of the sounds from the performer to the composer. The performer need only play the piece from the score and need not think about the control gestures. The composer, on the other hand, must think about the gestures when composing. They must write the piece to cause certain gestures so that certain control effects can be realized. Figure 7.8 shows the system in use in performance.

7.6 Discussion

The work performed in the course of designing the digital musical instruments discussed in this chapter has led to a number of interesting issues which bear further discussion. This section will discuss these issues are their relevance to digital musical instrument design in general and the work performed in this thesis in particular.

Figure 7.8: Cellist Chloé Dominguez performing with the non-conscious control system during a rehearsal of Sean Ferguson's *Ex Asperis*.

Collaborative Design

The collaborative design process used in the development of the instrument discussed here involved a number of people with different areas of expertise. This included engineers, composers and performers. The process used allowed each person to provide input on the design of the instruments at each stage of development. This is a freedom which digital musical instruments allow which is not present for traditional musical instruments.

From the outset, the composers and performers who will eventually be working with the instruments can influence their overall design. The fact that a digital musical instrument can produce any sound through its synthesis system allows the composer to design the sound of the instrument for the specific piece or pieces being written. The instrument sound can be designed around the composer's ideas for the composition, rather than vice-versa.

This design process also offers control to the performer. There are no preconceived notions of how the instrument should be played. During practice sessions performers can experiment with the instrument, finding interesting gestures or sounds. Those elements of the instrument which they find interesting to perform with can be kept or expanded, those which are uninteresting can be removed.

The fact that all members of the project are discovering the instruments at the same time also leads to an element of communication of ideas between groups. Interesting aspects of the instrument discovered by the performer's can be passed on to the composer. Such aspects may influence or be incorporated into the composition itself.

Instrument Reliability

As we worked with the instruments described here, the issue of reliability came to the fore. Existing acoustic instruments have developed to the point where they are extremely reliable. Faults with instruments are rare. Even those instruments which have parts which can break regularly (such as strings on a guitar or violin) have been designed to allow the performer to easily repair them. Changing a string on an instrument for example is a relatively simple proposition on most acoustic stringed instruments.

With digital musical instruments, there are often many complex electronic components which can be fragile. Sensors which have not been designed for the purpose for which they are being used can break quite easily. For instance, the FSRs used on the fingertips of the FM Gloves are designed for mounting on a solid non-flexible surface (Interlink Electronics, Inc.). Yet we use them on the fingertips, where there is no solid backing and the leads of the sensor can become bent. In such a case, the reliability of the instrument can suffer.

Repairing a broken sensor can require some electronics knowledge and skills. It can also require complex equipment. For a DMI to be used regularly for practice and performance, such breakdowns and repairs are a problem. The instruments must then be designed to be as reliable as possible. In those cases where a sensor is likely to break (such as the already mentioned FSRs in the FM Gloves), the instrument should be designed so that they can be easily replaced, without requiring electronics training and equipment. For the FM Gloves, this meant a mounting system where the FSRs could be unplugged when broken and a new one plugged in.

Instrument Aesthetics

Another interesting issue which arose from this research revolved around the aesthetics of digital musical instruments. Many DMIs are developed and performed by engineers. They can look like early prototypes, with exposed components and wiring. We found during the course of this work that such instruments can be intimidating for musicians. They look fragile and so people are almost afraid to touch them.

Such was the case for the early versions of the T-Box. With the circuitry and wiring exposed, people were worried about breaking it and so were very cautious when interacting with the instrument. Once the current wooden body was created, this worry went away. People were no longer afraid to touch the instrument, they became much more free in their interaction with it.

The improved aesthetics of the instrument also had another effect. When demonstrated to the public, the T-Box was recognized as an instrument. People were drawn to it and were interested in playing it.

Sensor Choice for Digital Musical Instruments

The two instruments developed for the Digital Orchestra project, the T-Box and FM Gloves, also provided some validation of the research on sensors and musical functions discussed in Chapter 4. Unlike the Viblotar, these instruments were not designed to follow the results of the experiments described in that chapter. However, over time the design process used resulted in an evolution of the instruments to the point where the sensor and parameter choices matched those predicted by the experiments.

For example, the distance between one ultrasound transmitter and receiver in the T-Box (an example of linear position sensing) was mapped to the center frequency of a filter (an absolute dynamic function). This matches the results of the experiment in Section 4.2, where for the absolute dynamic task (task 1) the linear position sensor performed significantly better than all other sensors. This also holds true for the other ultrasound receivers, all of which were mapped to absolute dynamic functions.

This is also the case for the linear position type sensor in the FM Gloves. In this case, the infrared distance sensor is used as a linear position sensor, measuring the distance between the performer's hand and body. The output from this sensor is mapped to the volume of the sound output, another absolute dynamic function.

The FM Gloves also contain a number of FSRs, which are pressure sensors. As shown by the experiments in Chapter 4, such sensors are highly suited for parameter modulation tasks (or relative dynamic functions). While the right hand

FSRs are used to recognize striking gestures, those of the left hand are used to modulate effect parameters from a starting value.

Vibrotactile Feedback for System State Information

The FM Gloves also provide a case study in the use of vibrotactile feedback for information regarding the state of an instrument. The infrared distance sensor used on the left hand has a non-linear response with distance. Once the object being sensed passes within a certain distance of the sensor the distance output of the sensor starts increasing rather than decreasing. Figure 7.9 shows the response of the sensor with distance.

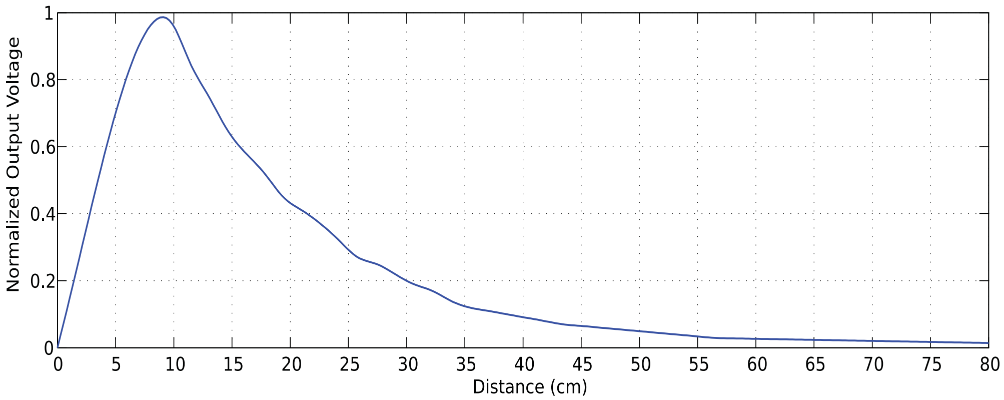

Figure 7.9: Output of infrared distance sensor relative to the distance to the object being tracked.

For the FM Gloves the output of this sensor is mapped to the volume of the sound, with smaller distances between the hand and the body giving quieter sounds. However, if the hand passes within the cutoff distance for the sensor then the distance reading will increase instead of decreasing (and will do so quickly) resulting in a much louder sound than desired. The only knowledge the performer has of this cutoff point will come from their proprioceptive sense of the position of their own body.

By using a simple vibrating motor, which is directly controlled by the output of the sensor, we can provide an additional vibrotactile feedback signal which can be used to warn them of the approach of this point. This allows them to determine where the point is using a combination of their own proprioceptive sense and the signal from the motor, which is easier than using the proprioceptive alone.

Gesture-tracking vs Physical Digital Musical Instruments

The systems developed for the gesture-controlled spatialization project provide examples of gesture tracking systems rather than digital musical instruments. Such systems track parameters of human movements themselves rather than having physical forms which the performer acts upon.

While these systems are still designed using sensors, the mappings between sensors and parameters are less direct than those examined in Chapter 4. As such, the results of the experiments in that chapter cannot be readily applied to these systems. The focus in their design is on the movement which we wish to track, rather than a combination of the movement and the parameter being controlled.

However, it is possible to apply some of the work on vibrotactile feedback from Chapter 5 to these systems. For instance, vibrotactile feedback similar to that used in the FM Gloves could be used for the system for virtual direct manipulation of sound sources. In this system, the performer can grab and manipulate virtual objects. However, as there is no physical object being manipulated, there is no feedback from the object to the performer except the visual.

This can make it difficult, for instance, to know when you are about to mistakenly drop a virtual object. With a physical object, there would be a physical feedback to indicate losing your grip on the object. For a virtual object, no such feedback exists. The first indication of losing the object is when it no longer moves on the visual feedback display. By using vibrotactile feedback, we can provide an indication of the loss of grip on the object. The feedback could be directly controlled by the sensors on the gloves, such that a reduction in pressure on the sensor would cause an increase in vibration.

7.7 Conclusions

This chapter detailed a number of instruments which were developed as part of larger collaborative projects. The design of these instruments made use of the research described in previous chapters. These instruments also benefited from a collaborative design process, which involved composers and performers as well as engineers. The collaborative design process raised a number of important issues regarding the design of DMIs.

These issues included the importance of reliability in a digital musical instrument. Instruments which are reliable enough to demonstrate, or even to practice with, may not be reliable enough to perform. Not only this, but an instrument which cannot be used reliably by a *non-technical person* can create serious difficulties for performers and composers.

The second issue raised was that of aesthetics. Many new DMIs look like prototypes, with visible electronic components and wires. Such instruments can look both fragile and overly technical to musicians. This can discourage them from interacting with the instrument. By improving the aesthetic of the instruments design, one can create a DMI which looks “instrument-like” and encourages interaction.

The instruments discussed in this chapter provided validation of some of the work discussed in earlier chapters. In particular, it is interesting to note that the parameters controlled by specific sensors in both the T-Box and FM Gloves evolved through the collaborative design process to match those predicted by the experiments discussed in Chapter 4.

Finally, the gesture tracking systems developed for the gesture-controlled sound spatialization project show that some of the work presented in this thesis can also be applied to system which do not necessarily have a full physical interface. While such gesture tracking systems may not benefit from the work on sensor choice in this thesis, there are still opportunities for the use of vibrotactile feedback as a communication channel in these systems.

The next chapter will provide an overall discussion of the work presented in this thesis. This includes an analysis of the importance of this work, as well as a discussion of areas of future work which emerge from it.

Chapter 8

Discussion, Conclusions and Future Work

This chapter provides a discussion of the work presented in this thesis. It includes a detailed discussion on the results of the work presented here, comments on interesting issues which have arisen from this research and presents a number of possibilities for future research which arise from this work.

8.1 Discussion and Conclusions

The work described in this thesis dealt with the improvement of the performer-instrument interaction in digital musical instruments. In particular it concentrated on the physical interface of digital musical instruments and the components of the physical interface: the sensors, actuators and the body. This section will discuss the major results of this thesis, as they relate to the components of the physical interface.

8.1.1 Sensors

A number of experiments were described in Chapter 4 on the optimal choice of sensors for specific tasks in a digital musical instrument. Initially, these experiments aimed to determine whether we could empirically validate the mappings of sensors to tasks presented by Vertegaal et al. (1996). If not, the aim was to attempt to empirically determine new mappings between classes of sensors and classes of musical task.

Three experiments were performed, each examining different aspects of the choice of sensor for specific musical tasks. They dealt with issues such as the effects of previous musical experience on sensor preference, the use of objective measurements of performance along with subjective user ratings and the effects of gesture parameters on both objective measures and subjective ratings.

Overall, these experiments give a number of guidelines for the choice of sensors for specific tasks in a digital musical instrument. While no overall mappings of sensor to task could be determined, the results indicate that it is possible to use a methodology such as the one used here to determine optimal sensor/task mappings. However, such work would likely have to be performed on an instrument by instrument basis.

As well as this, these experiments indicate that we must consider the sensor and gesture used to activate it together, rather than just the sensor itself. There were a number of differences between the two methods of activating the linear position sensor. We cannot therefore just choose a sensor, we must think about how that sensor will be used in the instrument. Evaluation should then be performed with the sensor on the instrument (or a mockup of the instrument). This allows the participants to use the sensor as they would do on the actual instrument.

The parameters of the musical task are also an important part of the evaluation. As seen from these experiments, the largest effect on the measurement results was caused by the speed of the modulation. This factor significantly effected all of the objective measures. As shown in the second experiment, participants were also more comfortable with each method at specific speeds. We must therefor carefully consider such variables and examine them as part of any evaluation.

Also these experiments showed the importance of using both subjective user ratings together with objective measures of performance. While in some cases there were no significant differences in preference ratings between two methods, there were significant differences in objective measures. Conversly, for those cases where there were no significant differences in objective measures, there were often differences in the subjective ratings.

It should be noted that these experiments examined initial ease of use and preference for the sensors and methods from participants who had not had much practice with the sensors. Following the results of these experiments in the choice of sensors would be expected to result in an instrument which is easy to use from the beginning. In fact, there seems to be evidence from the experiment with the Viblotar in Chapter 6 that an instrument designed along these lines receives high ease of use ratings from performers. This is useful for an instrument where the aim is to have it easy to use for beginners. For instance, instruments such as the reacTable* (Jorda et al., 2005) are designed with this aim in mind and so might benefit from this approach.

However, for instruments which are to be more like acoustic instruments, requiring some amount of time spent practicing in order to reach a level of expertise, it would also be important to look at the development in objective and subjective measures over time. In particular, it may be valuable to look at the effects of training and learning over time.

Finally, further validation for these results can be seen in the instruments described in Chapter 7. The FM Gloves and the T-Box were developed through an iterative, collaborative design process involving engineers, composers and performers. Unlike the Viblotar, their interfaces were not designed to follow the results of the experiments in Chapter 4. Yet, examining the final configuration of these instruments we can see that their sensor/task mappings have evolved to match those predicted as resulting in the best ease of use and preference ratings by the experiments that were performed in Chapter 4.

8.1.2 Actuators

Chapter 3 discussed a survey of new digital musical instrument presented at the NIME workshop and conferences. One of the most interesting results of this was the lack of instruments providing physical feedback to the performer. While a number of works discussed in Chapter 2 dealt with the importance of such feedback to the performer of a musical instrument, fewer than 9% of new digital musical instrument presented at NIME included any.

Towards remedying this, Chapter 5 discussed methods of providing vibrotactile feedback to the performer of a digital musical instrument. In that chapter, I described a variety of actuators which could be used to produce such feedback. In order to choose the most suitable actuator for a specific implementation of vibrotactile feedback, it is necessary to have a method of comparing and evaluating such actuators. Chapter 5 also described a method of doing just that.

The method described there involved a combination of physical response measurements (frequency response, amplitude response) together with control characteristics (such as number of independent control variable, or control signal type). Using the methods described in that chapter it is possible to evaluate a variety of actuators and determine the optimal one for a specific implementation. For instance, to implement vibrotactile feedback which is based on the sound produced by the instrument (as is the case in an acoustic instrument), we would require independent control of amplitude, frequency and waveform, together with a frequency response range of at least 40 to 1000 Hz and a large range of amplitude variation available. Examining the results of the measurements in that chapter, we can see that a loudspeaker voicecoil would be the optimal choice in this case. Similar measurements and comparisons can be made for other implementations.

Along with measuring and comparing actuators, Chapter 5 also described a method of modifying the frequency response of an actuator. Such response modification can be used to achieve a number of effects, such as forcing an actuator to produce a linear frequency response, creating vibrations only within the range of human sensation, or even modifying the response to be more like another actuator or instrument. To examine the effectiveness of such a system, I also performed an experimental evaluation of the effect of response modification on vibrotactile frequency discrimination.

Looking at the results of this experiment, we see that there is a significant improvement in frequency discrimination when any form of compensation is used. Interestingly, we also see a significant improvement in discrimination when only using actuator compensation is used rather than combined actuator and skin compensation. It seems that when only using actuator compensation, participants are noticing a change in the perceived amplitude of the vibration and misconstruing this as a change in frequency. This is resulting in a correct judgment of a frequency change, but based on the wrong evidence. Examining this further, I found that on average participants will mistake an amplitude change for a frequency change over 35% of the time.

This result is particularly interesting as studies which have examined the just noticeable difference (JND) for vibrotactile frequency sensation used only actuator compensation. The result of this experiment would seem to indicate that a more accurate measurement might result from using the combined actuator and skin compensation.

While Chapter 5 examined the technical issues and perceptual effects of vibrotactile feedback production, there still remained the issue of its effects on an actual digital musical instrument. The Viblotar, which was described in Chapter 6, provided a useful testbed for examining this issue. The Viblotar includes vibrotactile feedback which is generated by playing the synthesized sound of the instrument through loudspeakers which are built into the instrument body. This results in a situation which is closest to that of an acoustic musical instrument, where the sound and vibration are generated by the same mechanism. Experimental evaluation of the Vilbotar showed that the presence of vibrotactile feedback had an effect on a performer ratings of engagement, entertainment and control of the instrument.

While these results are interesting, they do not offer as much indication of the improvements we might expect from adding vibrotactile feedback. However, vibrotactile feedback is more often used by expert performers than novices. The performers who took part in this experiment, while all experienced musicians, were all new to the Viblotar. It is possible then that a long term experiment, examining the changes in these ratings over time, might find more significant changes than did this experiment.

Finally, while vibrotactile feedback can be used to provide instrument-like vibrations, it can also be used as an extra channel of information from the instrument to the performer. An example of this is provided by the FM Gloves, which were described in Chapter 7. These gloves use a vibrating motor to provide a n extra channel of information to the performer about the state of the instrument.

8.1.3 Instrument Body

While the sensors and actuator portions of the physical interface have received the most examination in this work, there have been some interesting findings regarding the design of the instrument body. Two areas of particular interest are that of reliability and aesthetics.

The issue of reliability effects all of the portions of the physical interface of a DMI. The sensors and actuators must work reliably in order for the DMI to be considered an instrument. The collaborative projects described in Chapter 7 have particularly highlighted this. Composers and performers require their instruments to be reliable. They must work correctly the majority of the time in order to allow the musicians to spend the necessary time developing the required performance skills for the instrument.

This issue means that care must be taken then designing the instrument body. The placement and attachment of sensors can place stresses on them which can result in breakages or failures. For the FM Gloves, we found that the FSRs on the fingertips in particular were subject to breakages due to their mounting on the fingertips. In this case, work on finding the sturdiest method of mounting and connecting the sensors resulted in a major reduction in such issues. Further work resulted in a mounting method which allowed them to be removed and replaced with no need for any additional equipment or skills. FSRs were merely unplugged and new ones plugged in. Such work improves the overall reliability of the instrument and should a sensor break, makes the replacement process as simple as replacing a broken string on a stringed instrument such as a guitar or violin. The result is more playing time for the musicians, leading to a better overall experience.

The issue of instrument aesthetics is also an important one, yet one that is often overlooked in the design of digital musical instruments. Often, new digital musical instruments are created as experiments, whether for new sensing technologies, mapping systems or synthesis techniques. Such instruments do not receive the same focus on aesthetics which acoustic instruments have.

From working with composers and performers and demonstrating to members of the public, we have found that the aesthetics of an instrument have a major effect on how people approach the instrument. DMIs which have lots of visible wires and components can look fragile and so people will be restrained in how they interact with them. There is a fear of damaging the instrument. Such was the case with some of the early versions of the T-Box.

Yet, once the T-Box circuitry was placed in its current wooden body, peoples reaction to the instrument changed. It became an instrument, rather than a delicate piece of electronics. This encouraged people to interact with the instrument and to interact more freely with it. There was no longer the same fear of damaging the instrument.

By taking care to deal with the aesthetics of the instrument, making it more “instrument-like” in its appearance, we can produce an instrument which encourages interaction and which attracts people to it.

8.2 Relevance

The research discussed in this work provides insight into a number of aspects of the design of new digital musical instruments. In particular it highlights the interdisciplinary nature of the design of such instruments, through the application of methods, techniques and information from fields such as human-computer interaction, experimental psychology, perception, engineering and music performance. As such, the work may have relevance not just within the field of music technology but also within these broader fields.

Specifically for digital musical instrument design, the work here has highlighted the discrepancy between the theory and practice within the field. While much has been written regarding the importance of the correct choice of sensors, or the importance of physical feedback to the performer, Chapter 3 shows that in practice this is often ignored.

Much of the work in Chapter 4 dealt with the application of human-computer interaction and experimental psychology techniques to the selection of optimal sensor/task pairings. As well as being applicable to the design of new DMIs, this work may also benefit the design of other real-time interactive systems, such as computer game interfaces or robot control systems. In particular the findings regarding the importance of a combination of objective and subjective metrics could prove valuable to those areas where often only one of these is used (if any).

Chapter 5 offers interesting insight into the creation of vibrotactile feedback. Vibrotactile feedback systems have been used in areas such as teleoperation, mobile and wearable computing and aviation. The work in that chapter provides methods and criteria for evaluating and comparing actuators for vibrotactile feedback, to allow for the selection of the optimal actuator for a specific task. The results on vibrotactile frequency discrimination could also prove useful in these and other areas. Experiments on the just-noticeable difference (JND) of vibrotactile frequency discrimination generally compensate only for the actuator frequency response, rather than for the combination of the actuator and human vibrotactile perception. Yet the experiment in Chapter 5 shows a significant difference in frequency discrimination between the actuator compensation and combined actuator and skin compensation conditions.

Finally, the work in Chapters 6 and 7 offers useful information for the design of new digital musical instruments from interaction with the musicians who are the end users of such instruments. In particular, the importance of both the reliability of such instruments and also their aesthetic appeal may be of interest to instrument designers who may normally focus on the use of interesting technology rather than the improvement of the overall instrument.

8.3 Future Work

There are a number of areas of possible future work which arise from this research. In this research I made use of evaluation methods from human-computer interaction. In particular I concentrated on task-based evaluation. Task-based evaluation examines the suitability of a system (in this case a DMI) for performing specific tasks. For example, for the sensor experiments, I examined the suitability of a number of sensors for note selection and pitch modulation tasks. However, a trend has arisen in recent years in the field of human-computer interaction towards experience-based evaluation, using techniques such as discourse analysis (Stowell et al., 2008), biosensors (Picard and Daily, 2005) and non-verbal reporting (Isbister et al., 2007).

Given this trend, the possibility exists to perform additional analysis of sensor/task combinations (or indeed complete digital musical instruments) using these techniques. Musical performance is inherently experiential. Performance is not generally task-based, unlike much general human-computer interaction. As such analysis of the overall experience of performance with a digital musical instrument may lead to new insights into the design of such instruments. As the work in Chapter 4 showed that both subjective judgments and objective performance measures are necessary to arrive at the optimal choice of sensor, the addition of experiential evaluation to the task-based evaluation methods may result in better overall choices or design for the instrument itself.

The work discussed in Chapter 5 showed a significant effect on human vibrotactile frequency discrimination when a frequency response compensation was introduced. This, together with the difference found between actuator-only compensation and combined actuator and skin response compensation, offers much possible future research into the effects of such compensation.

The evaluation of the effects of vibrotactile feedback on the Viblotar in Chapter 6 was a short-term experiment involving novice performers. The results did not show many significant differences between the ratings with and without vibrotactile feedback. This may be due to the short term nature of the experiment. Longer term studies of the effects of vibrotactile feedback, which is known to be used by expert performers, might produce more significant results and a better understanding of the effects of such feedback. This is an area offering much potential for further study.

Finally, there is much potential in the area of collaborative design of digital musical instruments. The process of collaborative design, involving both designers and end users, is an area which has recently begun to receive much interest, including the formation of an international journal on "CoCreation in Design and the Arts"[1]. Yet, little research has been presented in this journal or elsewhere on the use of collaborative design processes in the design of new digital musical instruments.

---

[1] http://www.tandf.co.uk/journals/titles/15710882.asp

Appendix A

Experiment Materials

A.1 Sensor Experiments: Chapter 4

A.1.1 Experiment 1

Information for Prospective Participants

Dear Prospective Participant:

We are interested in studying the usefulness of a number of sensor devices for performing specific musical functions. We are hoping to derive a mapping from particular categories of sensor to these musical functions. While some mappings have been documented, little empirical evidence exists to confirm these mappings.

In the present study you will be asked to perform three (3) simple musical tasks using a variety of sensors. You will receive training in the use of each sensor and will be allowed ample practice time with the sensor before being asked to perform the tasks. Each task can be performed as many times as you wish until you have successfully completed the task or have decided that you cannot complete it.

It should be noted that due to the nature of the experiment, successful completion of a given task might prove extremely difficult or impossible with certain of the sensors. It is important to realize that unsuccessful completion of a task is most likely due to the inadequacy of the sensor for the particular task and is not a failure on behalf of the participant. It is the suitability or otherwise of the sensor which is being evaluated.

On completion of each task you will be asked to rate the ease of use of the sensor for the given task. You will also be asked for any comments you wish to make regarding the sensor and the task. The performance of the tasks will also be recorded using audio recording equipment. This audio recording will be analyzed to allow us to further examine the suitability of the sensor to the task.

Participation in this experiment will require less than one hour. You may discontinue participation at any time. If you require a break during the test session, please inform the test administrator.

If you would like to learn the results of the study please contact Mark Marshall at [redacted] or by email at [redacted]. Thank you for your interest in this study.

Sincerely

____

Mark T. Marshall, M.Sc.

Participant Consent Form

My participation in the study “Evaluation of Sensor Usability for Specific Musical Tasks” is voluntary. I have read and understand the accompanying information regarding the study. I understand that I may discontinue the study at any time and that my name will not be disclosed at any time during the analysis or dissemination of the findings.

Participant’s Name: _____________________________________

__________________________________ ________________

Signature Date

Instructions

Task 1: Note Selection

For this task, you will be using each sensor in turn to play a melody. The range of the sensor has been divided evenly into 12 semitones. Whenever you press the button with your left hand a note will be played, the pitch of which corresponds to the current position of the sensor in your right hand.

Task 2: Pitch Modulation

For this task, you will be using each sensor in turn to perform a pitch modulation on a specific note in a 4-note riff. The range of the sensor has been mapped onto a range of +/- 1 semitone. Whenever you press the button with your left hand a note will be played from the 4-note riff. The riff will progress automatically to the next note with each button press. You should attempt to perform a pitch modulation of the last note of the riff using the sensor.

Task 3: Note Selection and Pitch Modulation

This task combines the previous 2 tasks. Again the range of each sensor is divided evenly into 12 semitones. Whenever you press the button with your left hand a note will be played, the pitch of which corresponds to the current position of the sensor in your right hand. You should attempt to play the melody using this sensor, modulating the pitch of the final note of the melody.

A.1.2 Experiments 2 and 3

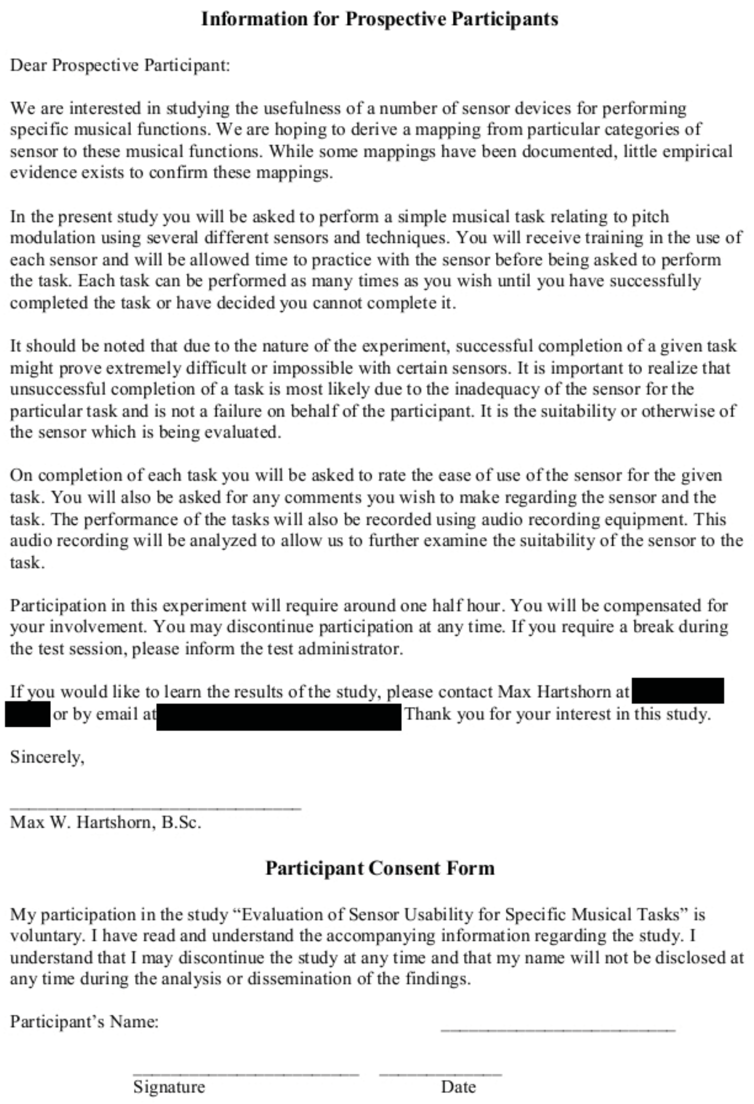

Information for Prospective Participants

Dear Prospective Participant:

We are interested in studying the usefulness of a number of sensor devices for performing specific musical functions. We are hoping to derive a mapping from particular categories of sensor to these musical functions. While some mappings have been documented, little empirical evidence exists to confirm these mappings.

In the present study you will be asked to perform a simple musical task relating to pitch modulation using several different sensors and techniques. You will receive training in the use of each sensor and will be allowed time to practice with the sensor before being asked to perform the task. Each task can be performed as many times as you wish until you have successfully completed the task or have decided you cannot complete it.

It should be noted that due to the nature of the experiment, successful completion of a given task might prove extremely difficult or impossible with certain sensors. It is important to realize that unsuccessful completion of a task is most likely due to the inadequacy of the sensor for the particular task and is not a failure on behalf of the participant. It is the suitability or otherwise of the sensor which is being evaluated.

On completion of each task you will be asked to rate the ease of use of the sensor for the given task. You will also be asked for any comments you wish to make regarding the sensor and the task. The performance of the tasks will also be recorded using audio recording equipment. This audio recording will be analyzed to allow us to further examine the suitability of the sensor to the task.

Participation in this experiment will require around one half hour. You will be compensated for your involvement. You may discontinue participation at any time. If you require a break during the test session, please inform the test administrator.

If you would like to learn the results of the study, please contact Max Hartshorn at [redacted] or by email at [redacted] Thank you for your interest in this study.

Sincerely,

____
Max W. Hartshorn, B.Sc.

Participant Consent Form

My participation in the study "Evaluation of Sensor Usability for Specific Musical Tasks" is voluntary. I have read and understand the accompanying information regarding the study. I understand that I may discontinue the study at any time and that my name will not be disclosed at any time during the analysis or dissemination of the findings.

Participant's Name: ________________________

________________________ _____________
Signature Date

Instructions

Task 1: Pitch modulation using a sliding technique on a linear position sensor.

The sensor you will be using for this task is a linear position sensor, which measures the relative position of your finger(s) on the sensor. When position is sensed to have moved to the left pitch is decreased. When position is sensed to have moved to the right pitch is increased. For the purposes of this experiment we would like you to create a vibrato or 'trill' effect using the sensor. This effect should be created by placing 2 fingers from your right hand on the sensor and rapidly sliding them back and forth.

On your left is a button that will allow you to progress through a short 4-note melody. Each time you press the button, one note in the melody will be emitted. In this task, please play through the 4-note melody with your left hand and, on the fourth note, while holding down the button, create a 'trill' effect using the sliding technique on the linear position sensor. You are allowed up to five minutes to practice with the sensor before attempting to complete the task.

Task 2: Pitch modulation using a rocking technique on a linear position sensor

The sensor you will be using for this task is a linear position sensor, which measures the relative position of your finger(s) on the sensor. When position is sensed to have moved to the left pitch is decreased. When position is sensed to have moved to the right pitch is increased. For the purposes of this experiment we would like you to create a vibrato or 'trill' effect using the sensor. This effect should be created by placing your right middle and index fingers on the sensor and rapidly rocking them back and forth so you alternate between which finger is touching the sensor.

On your left is a button that will allow you to progress through a short 4-note melody. Each time you press the button, one note in the melody will be emitted. In this task, please play through the 4-note melody with your left hand and, on the fourth note, while holding down the button, create a 'trill' effect using rocking technique on the linear position sensor. You are allowed up to five minutes to practice with the sensor before attempting to complete the task.

Task 3: Pitch modulation using a pressure sensor.

The sensor you will be using for this task is a pressure sensor. It receives its input from the relative pressure exerted upon it and uses the change in pressure to modulate a pitch. Pushing down on the sensor will cause pitch to increase while lessening pressure on the sensor will cause pitch to decrease. For the purposes of this experiment we would like you to create a vibrato or 'trill' effect using the sensor. The effect should be created by placing 2 fingers from your right hand on the sensor and rapidly increasing and decreasing the pressure you exert on it.

On your left is a button that will allow you to progress through a short 4-note melody. Each time you press the button, one note in the melody will be emitted. In this task, please play through the 4-note melody with your left hand and, on the fourth note, while holding down the button, create a 'trill' effect using the pressure sensor. You are allowed up to five minutes to practice with the sensor before attempting to complete the task.

Post-Experiment Questionnaire

The following questions have do with the experiment you just participated in, and your personal musical background. Please answer them in as much detail as you would like. If you need more space you can write on the back.

1) Age: _______ Gender: M/F_______

2). Is English your first language? Yes / No

3) Was there a particular task you liked most, what was it you liked about it?

4) Was there a task you liked least, what did you dislike about it?

5) Have you ever taken any music lessons? ______________

(NOTE: ANY kind of lessons count, including high school band class)

****If YES, complete #6-11; if NO, proceed to #12**

6) At what age did you start music lessons?______________________________

7) Private or group/classroom lessons?______________________________

8) What instrument(s)?______________________ # of years of training: _______

9) Years of training in total:____________________

10) Royal Conservatory Grade Level:_________

11) If not Royal Conservatory, what method of training?______________________

12) If you have brothers and/or sisters, have they had music lessons?__________________

13) Are your parents involved in music in any way (sing, play an instrument, avid listeners, etc.)? ____________________
____

14) Do you consider yourself musical?________

15) Are you currently involved in musical activities (i.e., are you currently playing an instrument)?__________________________

16) Do you listen to music? If so, how often? (e.g., everyday for about 3 hours)___________________________________
____

17) What type of music do you usually listen to? (e.g., classical, rock)_______________________

18) What is your favorite type of music?____

19) Did you often listen to music as a child? Did your parents often play/listen to music in your home?_______________
____

20) If so, what type of music?____

21) To the best of your knowledge, are you tone deaf? Yes/ No/ Somewhat/ Don't know

22) To the best of your knowledge, are you an absolute pitch possessor? Yes/No/Don't know

23) Lastly, are there any final comments or suggestions you would like to add about the experiment you just participated in?
____
____
____
____

A.2 Feedback Experiment: Chapter 5

McGill
Schulich School of Music
École de musique Schulich

Schulich School of Music
of McGill University
555 Sherbrooke Street West
Montreal, Quebec H3A 1E3

École de Musique Schulich
de l'Université McGill
555, rue Sherbrooke Ouest
Montréal (Québec) H3A 1E3

☎ (514) 398-4535
📠 (514) 398-8061
www.mcgill.ca/music

Information for Prospective Participants

Dear Participant,

This experiment aims to examine the effects of frequency compensation on vibrotactile frequency discrimination. You will be asked to touch a vibrating speaker using the fingers of one hand and to determine for each of a series of vibration signals whether or not a change in vibration frequency has occurred. There will be 40 signals of 1 second in duration in each series. You will be asked to perform this experiment with 3 different series of signals. Your participation will require 45 minutes to 1 hour. You will be compensated $10 for your time

The aim of this test is not to evaluate your individual performance, but to determine whether or not frequency compensation has an effect on human vibrotactile frequency discrimination. Only your ratings of each series will be analyzed. Your ratings will be anonymized and grouped with that of the other participants. Your name will not be disclosed at any time.

You may discontinue participation in this study at any point during the process. If you need a break, kindly inform the test administrator.

If you would like to learn about the results of the study, please contact Mark Marshall at [redacted]. We thank you for your interest in this study.

Sincerely,

____

Mark Marshall, PhD Candidate, Music Technology Area, Department of Music Research.

Supervisor: Prof. Marcelo M. Wanderley,
Music Technology Area, Department of Music Research. Tel: [redacted]

Participant Consent Form

My participation in the study "The effect of frequency compensation on vibrotactile frequency discrimination" is voluntary. I understand that I may discontinue participation at any point during the experiment and that my name will not be disclosed at any time during the analysis or the dissemination of findings.

Participant's name _______________________________________

Signature ____ Date ____________

A.3 Viblotar Experiment: Chapter 6

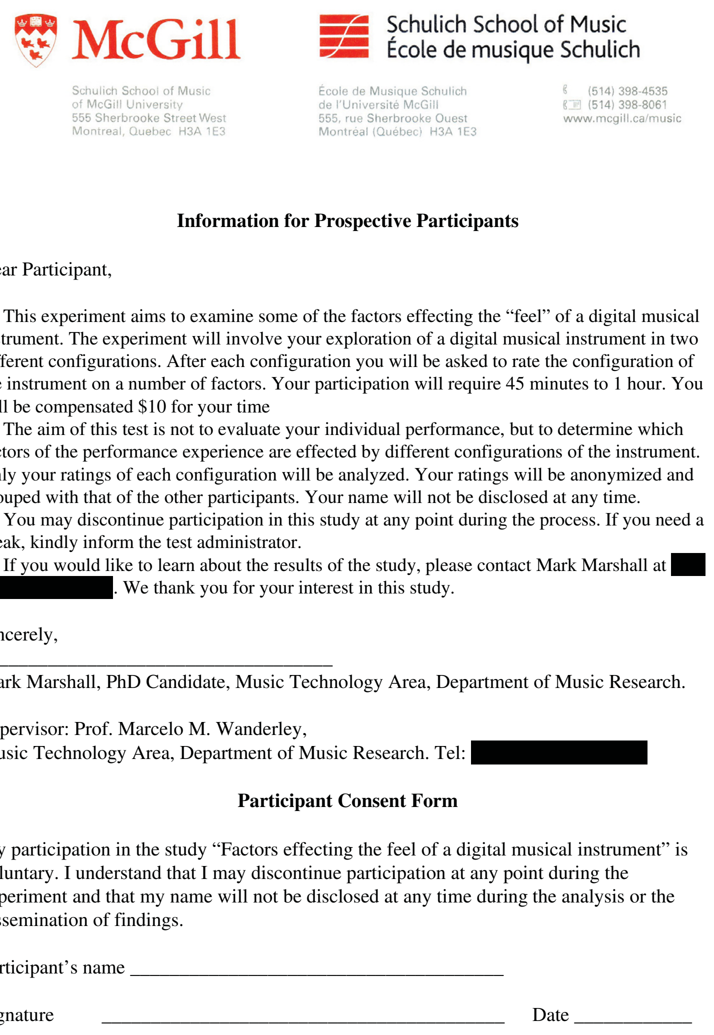

McGill

Schulich School of Music
École de musique Schulich

Schulich School of Music
of McGill University
555 Sherbrooke Street West
Montreal, Quebec H3A 1E3

École de Musique Schulich
de l'Université McGill
555, rue Sherbrooke Ouest
Montréal (Québec) H3A 1E3

(514) 398-4535
(514) 398-8061
www.mcgill.ca/music

Information for Prospective Participants

Dear Participant,

This experiment aims to examine some of the factors effecting the "feel" of a digital musical instrument. The experiment will involve your exploration of a digital musical instrument in two different configurations. After each configuration you will be asked to rate the configuration of the instrument on a number of factors. Your participation will require 45 minutes to 1 hour. You will be compensated $10 for your time

The aim of this test is not to evaluate your individual performance, but to determine which factors of the performance experience are effected by different configurations of the instrument. Only your ratings of each configuration will be analyzed. Your ratings will be anonymized and grouped with that of the other participants. Your name will not be disclosed at any time.

You may discontinue participation in this study at any point during the process. If you need a break, kindly inform the test administrator.

If you would like to learn about the results of the study, please contact Mark Marshall at [redacted]. We thank you for your interest in this study.

Sincerely,

____

Mark Marshall, PhD Candidate, Music Technology Area, Department of Music Research.

Supervisor: Prof. Marcelo M. Wanderley,
Music Technology Area, Department of Music Research. Tel: [redacted]

Participant Consent Form

My participation in the study "Factors effecting the feel of a digital musical instrument" is voluntary. I understand that I may discontinue participation at any point during the experiment and that my name will not be disclosed at any time during the analysis or the dissemination of findings.

Participant's name ____

Signature ____ Date ____________

Please rate the instrument in it's current configuration for each of the following criteria:

1. Ease of use: how easy the instrument is to perform with.

Very Difficult				Very Easy

2. Controllability: how much you feel in control of the instrument.

Very Uncontrollable				Very Controllable

3. Engagement: how much of your attention was put into playing the instrument.

Very Disengaged				Very Engaged

4. Entertainment: how entertaining the instrument is.

Very Low				Very High

5. Potential for further performance.

Very Low				Very High

Appendix B

The McGill Digital Orchestra: Credits

B.1 Project Description

The McGill Digital Orchestra is a research/creation project supported by the Appui à la recherche-création program of the Fonds québécois de la recherche sur la société et la culture (FQRSC). The objective of the Digital Orchestra research-creation program is to develop new creative resources that allow composers and performers to expand and renew their artistic practice through the interaction of live performance and digital technologies. Team members include Professors Denys Bouliane and Sean Ferguson of the Composition Area, Professors Marcelo Wanderley, Gary Scavone and Philippe Depalle of the Music Technology Area, and Professor AndrÃľ Roy of the Performance Department. All participants are members of the Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT). The grant has a duration of three years and will culminate with a performance of new works during the 2008 MusiMarch Festival.

B.2 Participants

Researchers

- Denys Bouliane (Composition)
- Philippe Depalle (Music Technology)
- Sean Ferguson (Composition)
- André Roy (Performance)
- Gary Scavone (Music Technology)
- Marcelo M. Wanderley (Music Technology)

Research Assistants

- Bertrand Scherrer (PhD Student, Music Technology)
- Chloé Dominguez (DMus student, Performance)
- David Birnbaum (MA Student, Music Technology)
- Erika Donald (DMus student, Performance)
- D. Andrew Stewart (DMus student, Composition)
- Fernando Rocha (DMus student, Performance)
- Heather Hindman (MA student, Composition)
- Joseph Malloch (PhD student, Music Technology)

- Kent Walker (PhD student, Sound Recording)
- Mark T. Marshall (PhD student, Music Technology)
- Rodolphe Koehly (PhD student, Music Technology)
- Simon de Leon (MA student, Music Technology)
- Stephen Sinclair (PhD student, Music Technology)
- Xenia Pestova (DMus student, Performance)

B.3 Funding Organisations

- Fonds québécois de la recherche sur la société et la culture (FQRSC)
- Centre for Interdisciplinary Research in Music Media and Technology (CIRMMT)

B.4 Compositions and Performances

Two new works were composed as part of this work, *The Long and the Short of It* by Heather Hindman and *sounds between our minds* by D. Andrew Stewart. Both pieces were premiered in Pollack Hall, McGill University as part of the 2008 MusiMarch festival.

The Long and the Short of It

Composed by: Heather Hindman

Performed by: Erika Donald (Cello), Xenia Pestova (FM Gloves), Fernando Rocha (T-Box)

Mapping & Synthesis: Heather Hindman

Digital Musical Instruments: Mark T. Marshall, Pierre-Yves Fortier, Geof Holbrook, Eileen TenCate

Software Development: Joseph Malloch, Mark T. Marshall, Stephen Sinclair

Technical Support: Digital Composition Studios (DCS), Input Devices and Music Interaction Laboratory (IDMIL), Schulich School of Music, McGill University

Project Coordination: Sean Ferguson, Marcelo M. Wanderley

sounds between our minds

Composed by: D. Andrew Stewart

Performed by: Erika Donald (Soprano T-Stick), Xenia Pestova (Rulers), Fernando Rocha (Tenor T-Stick)

Mapping & Synthesis: D. Andrew Stewart

Digital Musical Instruments: Joseph Malloch, David Birnbaum

Software Development: Joseph Malloch, Stephen Sinclair, D. Andrew Stewart

Technical Support: Digital Composition Studios (DCS), Input Devices and Music Interaction Laboratory (IDMIL), Schulich School of Music, McGill University

Project Coordination: Sean Ferguson, Marcelo M. Wanderley

Appendix C

Compositional Applications of Auditory Scene Synthesis in Concert Spaces via Gestural Control: Credits

C.1 Project Description

This project aims to develop novel compositional and technological methods for the advanced use of the multidimensional nature of auditory space in music composition. The development of a novel auditory virtual environment (AVE) is part of the project. The design of the AVE is based on a 24-channel loudspeaker system and virtual microphone control. The system will be used to deliver the spatial aspects of compositions accurately over a large listening area. Within this project, the way to best control the parameters of the AVE via gestures of the musicians will also be investigated. This way, the musicians will be free to interpret the spatial aspects of the score. New and innovative Acoustic Scene Synthesis methods will be developed to allow use to fuse and segregate acoustic objects into different auditory streams. At the end of the project, the newly developed methods will be demonstrated in a composition for a small music ensemble by Sean Ferguson. The composition will be performed in Montreal at the 2008 MusiMarch Festival.

C.2 Participants

Researchers

- Stephen McAdams (Music Technology, McGill)
- Marcelo M. Wanderley (Music Technology, McGill)
- Sean Ferguson (Composition, McGill)
- Jonas Braasch (School of Architecture, Rensselaer Polytechnic Institute)

Research Assistants

- Georgios Marentakis (Post Doc., Music Technology)
- Nils Peters (Ph.D. Candidate, Music Technology)
- Mark T. Marshall (Ph.D. Candidate, Music Technology)
- Joseph Malloch (Ph.D. Candidate, Music Technology)
- Marlon Schumacher (Ph.D. Candidate, Music Technology)
- Tristan Matthews (Undergraduate, Music Technology)

C.3 Funding Organisations

- Natural Science and Engineering Research Council of Canada (NSERC)
- Canada Council for the Arts (CCA)

C.4 Compositions and Performances

There were two new works written and performed as part of this project. These were *Miroirs* and *Ex Asperis* both written by Sean Ferguson. *Miroirs* was performed at the 4th International Conference on Enactive Interfaces, in Grenoble, France, in November 2007. *Ex Asperis* was premiered in Pollack Hall, McGill University at the 2008 MusiMarch festival.

Miroirs

Composed by: Sean Ferguson

Performed by: Chloé Dominguez (Cello)

Mapping, Synthesis and Spatialization: Sean Ferguson, Nils Peters

Gesture Control System: Mark T. Marshall, Joseph Malloch

Rendered Virtual Dancer: Chi-min Hsieh, Annie Luciani

Software Development: Sean Ferguson, Nils Peters, Joseph Malloch, Mark T. Marshall

Technical Support: Digital Composition Studios (DCS), Input Devices and Music Interaction Laboratory (IDMIL), Schulich School of Music, McGill University

Ex Asperis

Composed by: Sean Ferguson

Performed by: Chloé Dominguez (Cello), Fernando Rocha (Gesture Control), The McGill Contemporary Music Ensemble

Conducted by: Denys Bouliane

Mapping, Synthesis and Spatialization: Sean Ferguson, Nils Peters, Marlon Schumacher, Joseph Malloch

Gesture Control System: Mark T. Marshall, Joseph Malloch, Marlon Schumacher, Nils Peters, Georgios Marentakis

Software Development: Sean Ferguson, Nils Peters, Marlon Schumacher, Joseph Malloch, Mark T. Marshall

Technical Support: Digital Composition Studios (DCS), Input Devices and Music Interaction Laboratory (IDMIL), Schulich School of Music, McGill University

Bibliography

Allison, J. T. and Place, T. A. (2005). Teabox: a sensor data interface system. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 56–59.

Armstrong, N. (2006). *An Enactive Approach to Digital Musical Instrument Design*. PhD thesis, Princeton University.

Askenfelt, A. and Jansson, E. (1992). On vibration sensation and finger touch in stringed instrument playing. *Music Perception*, 9(3):311–350.

Bardos, L., Korinek, S., Lee, E., and Borchers, J. (2005). Bangarama: Creating music with headbanging. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 180–183.

Bau, O., Tanaka, A., and Mackay, W. E. (2008). The a20: Musical metaphors for interface design. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*, pages 91–96.

Bennett, P. (2006). Petecube: a multimodal feedback interface. In *Proceedings* *of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 81–84.

Bennett, P., Ward, N., OâĂŹModhrain, S., and Rebelo, P. (2007). Damper: a platform for effortful interface development. In *Proceedings of the 2007 Conference on New Interfaces for Musical Expression (NIME07)*, pages 273–276.

Berdahl, E., Steiner, H.-C., and Oldham, C. (2008). Practical hardware and algorithms for creating haptic musical instruments. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*, pages 61–66.

Bevilacqua, F., Rasamimanana, N., Fléty, E., Lemouton, S., and Baschet, F. (2006). The augmented violin project: research, composition and performance report. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 402–406.

Birnbaum, D. (2004). The touch flute: Exploring roles of vibrotactile feedback in musical performance. Technical report, McGill University - Music Technology.

Birnbaum, D. (2007). Musical vibrotactile feedback. Master's thesis, McGill University.

Birnbaum, D. and Wanderley, M. M. (2007). A systematic approach to musical vibrotactile feedback. In *Proceedings of the 2007 International Computer Music Conference (ICMC07)*, volume II, pages 397–404.

Blaine, T. and Forlines, C. (2002). Jam-o-world: evolution of the jam-o-drum multi-player musical controller into the jam-o-whirl gaming interface. In *Proceedings of the 2002 conference on New interfaces for musical expression (NIME02)*.

Bongers, B. (1997). Tactile display in electronic musical instruments. In *IEE Colloquium on Developments in Tactile Displays*.

Bongers, B. (2000). Physical interfaces in the electronic arts: Interaction theory and interfacing techniques for real-time performance. In Wanderley, M. M. and Battier, M., editors, *Trends in Gestural Control of Music*, pages 124–164. IRCAM - Centre Georges Pompidou, Paris, France.

Bottoni, P., Caporali, R., Capuano, D., Faralli, S., Labella, A., and Pierro, M. (2007). Use of a dual-core dsp in a low-cost, touch-screen based musical instrument. In *Proceedings of the 2007 Conference on New Interfaces for Musical Expression (NIME07)*, pages 394–395.

Bouillot, N., Wozniewski, M., Settel, Z., and Cooperstock, J. R. (2008). A mobile wireless augmented guitar. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*, pages 189–192.

Bowen, A. (2005). Soundstone: A 3-d wireless music controller. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 268–269.

Braasch, J., Peters, N., and Valente, D. L. (2008). A loudspeaker-based projection technique for spatial music applications using virtual microphone control. *Computer Music Journal*, 32(3):55–71.

Brown, L. M., Brewster, S. A., and Purchase, H. C. (2005). A first investigation into the effectiveness of tactons. In *Proc. of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems (WHC'05)*, pages 167–176.

Burtner, M. (2002). Noisegate 67 for metasaxophone. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 71–76.

Cáceres, J. P., Mysore, G. J., and Treviño, J. (2005). Scuba: the self-contained unified bass augmenter. In *Proceedings of the 2005 conference on New interfaces for musical expression (NIME05)*, pages 38–41.

Cadoz, C., Luciani, A., and Florens, J.-L. (1984). Responsive input devices and sound synthesis by stimulation of instrumental mechanisms: The cordis system. *Computer Music Journal*, 8(3):60–73.

Cannon, C., Hughes, S., and OâĂŹModhrain, S. (2003). Epipe: exploration of the uilleann pipes as a potential controller for computer-based music. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME03)*, pages 3–8.

Cariou, B. (1994). The axio midi controller. In *Proceedings of the 1994 International Computer Music Conference (ICMC94)*, pages 163–166.

Carlile, J. and Hartmann, B. (2005). Oroboro: a collaborative controller with interpersonal haptic feedback. In *Proceedings of the 2005 International Conference on New Interfaces for Musical Expression (NIME05)*, pages 250–251.

Chafe, C. (1993). Tactile audio feedback. In *Proceedings of the 1993 International Computer Music Conference (ICMC93)*, pages 76–79.

Chu, L. (2002). Using haptics for digital audio navigation. In *Proceedings of the 2002 International Computer Music Conference (ICMC02)*, pages 118–121.

Cook, P. R. (2001). Principles for designing computer music controllers. In *ACM CHI01 Workshop on New Interfaces for Musical Expression*.

Cook, P. R. (2004). Remutualizing the musical instrument: Co-design of synthesis algorithms and controllers. *Journal of New Music Research*, 33(3):315–320.

Cook, P. R. (2005). Real-time performance controllers for synthesized singing. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 236–237.

Couturier, J.-M. and Arfib, D. (2003). Pointing fingers: Using multiple direct interactions with visual objects to perform music. In *Proceedings of the 2003 conference on New interfaces for musical expression (NIME03)*, pages 184–187.

Crevoisier, A. and Polotti, P. (2005). Tangible acoustic interfaces and their applications for the design of new musical instruments. In *Proceedings of the 2005 International Conference on New Interfaces for Musical Expression (NIME05)*, pages 97–100.

Cuddy, L. L., Balkwill, L.-L., Peretz, I., and Holden, R. R. (2005). Musical difficulties are rare: A study of "tone deafness" among university students. *Annals of the New York Academy of Sciences*, 1060:311–324.

da Silva, A. R., Wanderley, M. M., and Scavone, G. (2005). On the use of flute air jet as a musical control variable. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 105–108.

de Laubier, S. and Goudard, V. (2006). Meta-instrument 3: a look over 17 years of practice. In *Proceedings of the 2006 conference on New Interfaces for Musical Expression (NIME06)*, pages 288–291.

Essl, G. and OâĂŹModhrain, S. (2006). An enactive approach to the design of new tangible musical instruments. *Organised Sound*, 11(3):285–296.

Farwell, N. (2006). Adapting the trombone: a suite of electro-acoustic interventions for the piece rouse. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 358–363.

Favilla, S., Cannon, J., Hicks, T., Chant, D., and Favilla, P. (2008). Gluisax: Bent leather band's augmented saxophone project. In *Proceedings of the 2008 Conference on New Interface for Musical Expression (NIME08)*, pages 366–369.

Fels, S. and Hinton, G. (1995). Glove-talkii: An adaptive gesture-to-formant interface. In *Proceedings of 1995 Proceedings of the SIGCHI conference on Human factors in Computing Systems (SIGCHI95)*, pages 456–463.

Fels, S. and Vogt, F. (2002). Tooka: explorations of two person instruments. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 116–121.

Fraden, J. (2004). *Handbook of Modern Sensors: Physics, Design and Applications*, chapter 1. Springer-Verlag New York, Inc., third edition.

Freed, A. (2008). Application of new fiber and malleable materials for agile development of augmented instruments and controllers. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*, pages 107–112.

Freed, A., Wessel, D., Zbyszynski, M., and Uitti, F. M. (2006). Augmenting the cello. In *Proceedings of the 2006 conference on New Interfaces for Musical Expression (NIME06)*, pages 409–413.

Gadd, A. and Fels, S. (2002). Metamuse: Metaphors for expressive instruments. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*.

Gallace, A., Tan, H. Z., and Spence, C. (2005). Tactile change detection. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 12–16.

Gemperle, F., Ota, N., and Siewiorek, D. (2001). Design of a wearable tactile display. In *Proceedings of the Fifth International Symposium on Wearable Computers (ISWC 2001)*, pages 5–12.

Gillespie, B. (1992). The touchback keyboard. In *Proceedings of the 1992 International Computer Music Conference (ICMC92)*, pages 447–448.

Gillespie, B. (2001). Haptics. In Cook, P. R., editor, *Music, Cognition, and Computerized Sound: An Introduction to Psychoacoustics*, pages 229–246. The MIT Press, Cambridge, MA, USA, first mit press paperback edition.

Goff, G. D. (1967). Differential discrimination of frequency of cutaneous mechanical vibration. *Journal of Experimental Psychology*, 74(2):294–299.

Goto, S. and Suzuki, T. (2004). The case study of application of advanced gesture interface and mapping interface: Virtual musical instrument "le superpolm" and gesture controller "bodysuit". In *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, pages 207–208.

Gunther, E., Davenport, G., and OâĂŹModhrain, S. (2002). Cutaneous grooves: composing for the sense of touch. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*.

Gurevich, M. and von Muehlen, S. (2001). The accordiatron: a midi controller for interactive music. In *Proceedings of the 2001 Conference on New Interfaces for Musical Expression (NIME01)*, pages 1–3.

Hakken, L., Abdullah, R., and Smart, M. (1992). The continuum: A continuous music keyboard. In *Proceedings of the 1992 International Computer Music Conference (ICMC92)*, pages 81–84.

Hasan, L., Yu, N., and Paradiso, J. A. (2002). The termenova: A hybrid free-gesture interface. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*.

Hatanaka, M. (2003). Ergonomic design of a portable musical instrument. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME03)*, pages 77–82.

Hayafuchi, K. and Suzuki, K. (2008). Musicglove: A wearable musical controller for massive media library. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*, pages 241–244.

Hopkin, B. (1996). *Musical Instrument Design: Practical Information for Instrument Making*. See Sharp Press.

Hornof, A. and Sato, L. (2004). Eyemusic: Making music with the eyes. In *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, pages 185–188.

Howard, D., Rimell, S., and Hunt, A. (2003). Force feedback gesture controlled physical modelling synthesis. In *Proceedings of the 2003 Conference on New Musical Instruments for Musical Expression (NIME03)*, pages 95–98.

Hunt, A. (2000). *Radical User Interfaces for real-time musical control*. PhD thesis, University of York.

Hunt, A., Wanderley, M. M., and Kirk, R. (2000). Towards a model for instrumental mapping in expert musical interaction. In *Proceedings of the 2000 International Computer Music Conference (ICMC00)*, pages 209–212.

Huott, R. (2002). An interface for precise musical control. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 1–5.

Interlink Electronics, Inc. *FSR Integration Guide & Evaluation Parts Catalog*. Camarillo, CA, USA.

Isbister, K., Höök, K., Laaksolahti, J., and Sharp, M. (2007). The sensual evaluation instrument: Developing a trans-cultural self-report measure of affect. *International Journal of Human-Computer Studies*, 65(4):315–328.

Jacob, R. J. K., Sibert, L. E., McFarlane, D. C., and Preston Mullen Jr., M. (1994). Integrality and separability of input devices. *ACM Transactions on Computer-Human Interaction*, 1(1):3–26.

Jehan, T., Freed, A., and Dudas, R. (1999). Musical applications of new filter extensions to max/msp. In *Proceedings of the 1999 International Computer Music Conference (ICMC99)*, pages 504–507. International Computer Music Association.

Jensenius, A. R., Koehly, R., and Wanderley, M. M. (2006). Building low-cost music controllers. In Kronland-Martinet, R., Voinier, T., and Ystad, S., editors, *CMMR 2005 - Proc. of Computer Music Modeling and Retrieval 2005 Conference, LNCS 3902*, pages 123–129. Springer-Verlag, Berlin Heidelberg.

Jorda, S. (2001). New musical interfaces and new music-making paradigms. In *Proceedings of the 2001 ACM CHI Workshop on New Interfaces for Musical Expression (NIME01)*.

Jorda, S., Kaltenbrunner, M., Geiger, G., and Bencina, R. (2005). The reactable*. In *Proceedings of the 2005 International Computer Music Conference (ICMC2005)*, pages 379–382.

Kapur, A., Lazier, A. J., Davidson, P., Wilson, R. S., and Cook, P. R. (2004). The electronic sitar controller. In *Proceedings of the 2004 Conference for New Interfaces for Musical Expression (NIME04)*, pages 7–12.

Kartadinata, S. (2003). The gluiph: a nucleus for integrated instruments. In *Proceedings of the 2003 Conference on New interfaces for Musical Expression (NIME03)*, pages 180–183.

Keele, S. W. (1973). *Attention and Human Performance.* Goodyear Publishing Company.

Kessous, L. and Arfib, D. (2003). Bimanuality in alternate musical instruments. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME03)*, pages 140–145.

Kirk, T. and Leider, C. (2007). The frankenpipe: A novel bagpipe controller. In *Proceedings of the 2007 Conference on New Interfaces for Musical Expression (NIME07)*, pages 301–304.

Koehly, R., Curtil, D., van de Ven, T., and Wanderley, M. M. (2007). Carbon black loaded paper: An intelligent substrate for electronic sensors design. In *Proceedings of the IARIGAI 2007 Conference*.

Koehly, R., Curtil, D., and Wanderley, M. M. (2006). Paper fsrs and latex/fabric traction sensors: Methods for the development of home-made touch sensors. In *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression (NIME06)*, pages 230–233.

Lebel, D. and Malloch, J. (2006). The g-spring controller. In *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression (NIME06)*, pages 85–88.

Lemouton, S., Stroppa, M., and Sluchin, B. (2006). Using the augmented trombone in "i will not kiss your f.ing flag". In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 304–307.

Levin, G. (2005). A personal chronology of audiovisual systems research. In *Proceedings of the 2005 International Conference on New Interfaces for Musical Expression (NIME05)*, pages 2–3.

Lindeman, R. W., Yanagida, Y., Noma, H., Hosaka, K.-i., and Kuwabara, K. (2005). Design of a wireless tactor system for haptic feedback in virtual reality. In *Proceedings of IEEE Virtual Reality 2005*, pages 285–286. IEEE.

Linvill, J. G. and Bliss, J. C. (1966). A direct translation reading aid for the blind. *Proceedings of the IEEE*, 54(1):40–51.

Livingstone, D. and Miranda, E. R. (2005). Orb3: adaptive interface design for real time sound synthesis & diffusion within socially mediated spaces. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 65–69.

Lock, D. and Schiemer, G. (2006). Orbophone: a new interface for radiating sound and image. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIMe06)*, pages 89–92.

Lyons, M. J., Haehnel, M., and Tetsutani, N. (2003). Designing, playing, and performing with a vision-based mouth interface. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME03)*, pages 116–121.

Magnusson, T. and Mendieta, E. H. (2007). The acoustic, the digital and the body: A survey on musical instruments. In *Proceedings of the 2007 Conference on New Interfaces for Musical Expression*, pages 94–99.

Mäki-Patola, T., Laitinen, J., Kanerva, A., and Takala, T. (2005). Experiments with virtual reality instruments. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 11–16.

Malloch, J., Sinclair, S., and Wanderley, M. M. (2008). A network-based framework for collaborative development and performance of digital musical instruments. In *CMMR 2007, - Proc. of Computer Music Modeling and Retrieval 2007 Conference*, volume 4969 of *LNCS*, pages 401–425, Berlin Heidelberg. Sprigner-Verlag.

Malloch, J. and Wanderley, M. M. (2007). The t-stick: from musical interface to musical instrument. In *Proceedings of the 2007 Conference on New Interfaces for Musical Expression (NIME07)*, pages 66–70.

Mandelis, J. and Husbands, P. (2004). Don't just play it, grow it!: breeding sound synthesis and performance mappings. In *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, pages 47–50.

Manning, P. (2004). *Electronic and Computer Music.* Oxford University Press, revised and expanded edition.

Marchand, S. and Raspaud, M. (2004). Enhanced time-stretching using order-2 sinusoidal modeling. In *Proceedings of the 7th International Conference on Digital Audio Effects (DAFx'04)*, pages 76–82.

Marentakis, G., Malloch, J., Peters, N., Marshall, M., Wanderley, M. M., and McAdams, S. (2008). Influence of performance gestures on the identification of spatial sound trajectories in a concert hall. In *Proceedings of the 14th International Conference on Auditory Display*, Paris, France.

Marmaras, N. and Zarboutis, N. (1997). Ergonomic redesign of the electric guitar. *Applied Ergonomics*, 28(1):59–67.

Marshall, M. T. (2005). The avr-hid: a 6-channel 10-bit (or 8-channel 8-bit) usb analog-to-digital converter.

Marshall, M. T., Malloch, J., and Wanderley, M. M. (2007). Gesture controlled sound spatialization for live musical performance. Presented at the 2007 International Gesture Workshop, Lisbon, Portugal.

Marshall, M. T., Rath, M., and Moynihan, B. (2002). The virtual bodhran - the vodhran. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 169–170.

Marshall, M. T. and Wanderley, M. M. (2006). Vibrotactile feedback in digital musical instruments. In *Proceedings of the 2006 International Conference on New Interfaces for Musical Expression (NIME06)*, pages 226–229.

Mase, K. and Yonezawa, T. (2001). Body, clothes, water and toys: media towards natural music expressions with digital sounds. In *Proceedings of the CHI'01 Workshop on New Interfaces for Musical Expression (NIME01)*.

Masliah, M. R. (2001). *Measuring the Allocation of Control in 6 Degree of Freedom Human-Computer Interaction Tasks*. PhD thesis, University of Toronto.

Masliah, M. R. and Milgram, P. (2000). Measuring the allocation of control in a 6 degree-of-freedom docking experiment. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 25–32, New York, NY, USA. ACM Press.

McElligott, L., Dixon, E., and Dillon, M. (2002). 'peglegs in music': processing the effort generated by levels of expressive gesturing in music. In *Proceedings* *of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 1–5.

Miranda, E. R. and Wanderley, M. M. (2006). *New Digital Musical Instruments: Control and Interaction beyond the Keyboard*. A-R Editions.

Miyashita, H. and Nishimoto, K. (2004). Thermoscore: a new-type musical score with temperature sensation. In *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, pages 104–107.

Moore, F. R. (1988). The dysfunctions of midi. *Computer Music Journal*, 12(1):19–28.

Mowbray, G. H. and Gebhard, J. W. (1957). Sensitivity of the skin to changes in the rate of intermittent mechanical stimuli. *Science*, 125(3261):1297–1298.

Mulder, A. (1998). *Design of Virtual Three-dimensional Instruments for Sound Control*. PhD thesis, Simon Fraser University.

Mulder, A. (2000). Towards a choice of gestural constraints for instrumental performers. In Wanderley, M. M. and Battier, M., editors, *Trends in Gestural Control of Music*, pages 315–335. IRCAM - Centre Pompidou, Paris, France.

Newton-Dunn, H., Nakano, H., and Gibson, J. (2003). Block jam: A tangible interface for interactive music. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME03)*, pages 170–177.

Nichols, C. (2000). The vbow: A haptic musical controller human-computer interface. In *Proceedings of the 2000 International Computer Music Conference (ICMC00)*, pages 274–276.

Nichols, C. (2002). The vbow: Development of a virtual violin bow controller haptic human-computer interface. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*.

Nishibori, Y. and Iwai, T. (2006). Tenori-on. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 172–175.

Noma, H., Hashida, Y., Susami, K., Hosaka, K.-i., Yanagida, Y., Lindeman, R. W., and Kume, Y. (2006). A study of mounting methods for tactors using an elastic polymer. In *Proceedings of the Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 343–345. IEEE Computer Society.

Oboe, R. and De Poli, G. (2002). Multi-instrument virtual keyboard - the mikey project. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 23–28.

Oldfield, R. C. (1971). The assessment and analysis of handedness: The edinburgh inventory. *Neuropsychologia*, 9:97–114.

Overholt, D. (2001). The matrix: A novel controller for musical expression. In *Proceedings of the 2001 Conference on New Interfaces for Musical Expression (NIME01)*.

Overholt, D. (2005). The overtone violin. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 34–37.

OâĂŹModhrain, S. (2000). *Playing by Feel: Incorporating Haptic Feedback into Computer-Based Musical Instruments*. PhD thesis, Stanford University.

OâĂŹModhrain, S. and Essl, G. (2004). Pebblebox and crumblebag: Tactile interfaces for granular synthesis. In *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, pages 74–79.

Palacio-Quintin, C. (2003). The hyper-flute. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME03)*, pages 206–207.

Pallas-Areny, R. and Webster, J. G. (2001). *Sensors and Signal Conditioning*, chapter 1. John Wiley & Sons, Inc., second edition.

Papich, G. and Rainbow, E. (1974). A pilot study of performance practices of twentieth-century musicians. *Journal of Research in Music Education*, 22(1):24–34.

Picard, R. W. and Daily, S. B. (2005). Evaluating affective interactions: alternatives to asking what users feel. Presented at Evaluating Affective Interfaces: Innovative Approaches. CHI 2005.

Pirringer, J. (2001). Elektronische musik und interaktivität: Prinzipien, konzepte, anwendungen. Master's thesis, Technical University of Vienna.

Pongrac, H. (2008). Vibrotactile perception: examining the coding of vibrations and the just noticeable difference under various conditions. *IEEE Multimedia Systems*, 13(4):297–307.

Pritchard, B. and Fels, S. (2006). Grassp: gesturally-realized audio, speech and song performance. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 272–276.

Puckette, M. and Settel, Z. (1993). Nonobvious roles for electronics in performance enhancement. In *Proceedings of the 1993 International Computer Music Conference (ICMC93)*, pages 134–137, San Francisco. International Computer Music Association.

Rich, R. (1991). Buchla lightning midi controller: a powerful new midi controller is nothing to shake a stick at. *Electronic Musician*, 7(10):102–108.

Rodríguez, D. and Rodríguez, I. (2005). Vife alpha v.01 real-time visual sound installation performed by glove gesture. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIME05)*, pages 252–253.

Rovan, J. and Hayward, V. (2000). Typology of tactile sounds and their synthesis in gesture-driven computer music performance. In Wanderley, M. M. and Battier, M., editors, *Trends in Gestural Control of Music*, pages 355–368. Ircam - Centre Pompidou.

Ryan, J. (1992). Effort and expression. In *Proceedings of the 1992 International Computer Music Conference (ICMC92)*, pages 414–416.

Ryan, J. and Salter, C. (2003). Tgarden: wearable instruments and augmented physicality. In *Proceedings of the 2003 conference on New interfaces for musical expression (NIME03)*, pages 87–90.

Schacher, J. C. (2008). davos soundscape, a location based interactive composition. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*, pages 168–171.

Schiesser, S. and Traube, C. (2006). On making and playing an electronically-augmented saxophone. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression*, pages 308–313.

Serafin, S. (2005). Acoustics of the intonarumori. Available Online. Acessed July 18th 2008.

Serafin, S., de Götzen, A., Böttcher, N., and Gelineck, S. (2006). Synthesis and control of everyday sounds reconstructing russolo's intonarumori. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 240–245.

Sinclair, S. (2007). Force-feedback hand controllers for musical interaction. Master's thesis, McGill University.

Sinclair, S. and Wanderley, M. M. (2007). Extending dimple: a rigid body haptic simulator for interactive control of sound. In *Proceedings of the 4th International Conference on Enactive Interfaces (ENACTIVE'07)*, pages 263–266.

Singer, E. (2003). Sonic banana: a novel bend-sensor-based midi controller. In *Proceedings of the 2003 Conference on New Interfaces for Musical Expression (NIME03)*, pages 220–221.

Sinyor, E. (2006). Digital musical instruments: A design approach based on moving mechanical systems. Master's thesis, McGill University, Montreal, QC, Canada.

Sinyor, E. and Wanderley, M. M. (2005). Gyrotyre: a dynamic hand-held computer-music controller based on a spinning wheel. In *Proceedings of the*

2005 Conference on New Interfaces for Musical Expression (NIME05), pages 42–45.

Small, R. H. (1973). Vented-box loudspeaker systems–part 1: Small-signal analysis. *Journal of the Audio Engineering Society*, 21(5):363–372.

Smyth, T. and Smith, J. O. (2002). Creating sustained tones with the cicada's rapid sequential buckling mechanism. In *Proceedings of the 2002 Conference on New Instruments for Musical Expression (NIME02)*, pages 112–115.

Sonami, L. (2008). Lady's glove. http://www.sonami.net/lady5.html.

Steiner, H.-C. (2004). Stickmusic: Using haptic feedback with a phase vocoder. In *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, pages 203–204.

Storm, S. A. (2006). Assessing the instrumentalist interface: Modifications, ergonomics and maintenance of play. *Physical Medicine and Rehabilitation Clinics of North America*, 17(4):893–903.

Stowell, D., Plumbley, M. D., and Bryan-Kinns, N. (2008). Discourse analysis evaluation method for expressive musical interfaces. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*, pages 81–86.

Suzuki, K., Kyoya, M., Kamatani, T., and Uchiyama, T. (2008). beacon: Embodied sound media environment for socio-musical interaction. In *Proceedings of the 2008 Conference on New Interfaces for Musical Expression (NIME08)*.

Tanaka, A. (2004). Mobile music making. In *Proceedings of the 2004 Conference on New Interfaces for Musical Expression (NIME04)*, pages 154–156.

Tanaka, A. and Gemeinboeck, P. (2006). A framework for spatial interaction in locative media. In *Proceedings of the 2006 Conference on New Interfaces for Musical Expression (NIME06)*, pages 26–30.

Tarabella, L., Bertini, G., and Sabbatini, T. (1997). The twin towers: a remote sensing device for controlling live-interactive computer music. In *Proc. of the 2nd International Workshop on Mechatronical Computer Systems for Perception and Action*.

Teh, Y. H. and Featherstone, R. (2007). Frequency response analysis of shape memory alloy actuators. In *International Conference on Smart Materials and Nanotechnology in Engineering (SMN'07)*.

Thiele, N. (1971). Loudspeakers in vented boxes: Part 1. *Journal of the Audio Engineering Society*, 19(5):382–392.

Trueman, D. (1999). *Reinventing the Violin*. PhD thesis, Princeton University.

Trueman, D. and Cook, P. R. (2000). Bossa: The deconstructed violin reconstructed. *Journal of New Music Research*, 29(2):121–130.

Verillo, R. T. (1992). Vibration sensing in humans. *Music Perception*, 9(3):281–302.

Verplank, B., Gurevich, M., and Mathews, M. (2002). The plank: Designing a simple haptic controller. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*.

Vertegaal, R., Ungvary, T., and Kieslinger, M. (1996). Towards a musician's cockpit: Transducers, feedback and musical function. In *Proceedings of the*

1996 International Computer Music Conference (ICMC96), pages 308–310, San Francisco, CA. ICMA.

Wanderley, M. M. (2001). *Performer-Instrument Interaction: Applications to Gestural Control of Sound Synthesis*. PhD thesis, University Paris 6.

Wanderley, M. M. and Depalle, P. (2004). Gestural control of sound synthesis. *Proceedings of the IEEE*, 92(4):632–644.

Wanderley, M. M. and Orio, N. (2002). Evaluation of input devices for musical expression: Borrowing tools from hci. *Computer Music Journal*, 26(3):62–67.

Wanderley, M. M., Viollet, J.-P., Isart, F., and Rodet, X. (2000). On the choice of transducer technologies for specific musical functions. In *Proceedings of the 2000 International Computer Music Conference (ICMC2000)*, pages 244–247.

Weinberg, G., Aimi, R., and Jennings, K. (2002). The beatbug network: a rhythmic system for interdependent group collaboration. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 106–111.

Weinberg, G. and Driscoll, S. (2005). iltur âĂŞ connecting novices and experts through collaborative improvisation. In *Proceedings of the 2005 Conference on New Interfaces for Musical Expression (NIMe05)*, pages 17–22.

Wessel, D. and Wright, M. (2002). Problems and prospects for intimate musical control of computers. *Computer Music Journal*, 26(3):11–22.

White, R. M. (1987). A sensors classification scheme. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 34(2):124–126.

Wilkerson, C., Ng, C., and Serafin, S. (2002). The mutha rubboard controller. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*, pages 82–85.

Yanagida, Y., Kakita, M., Lindeman, R. W., Kume, Y., and Tetsutani, N. (2004). Vibrotactile letter reading using a low-resolution tactor array. In *Proceedings of the 12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 400–406.

Yao, H.-Y. (2004). Touch magnifying instrument applied to minimally invasive surgery. Master's thesis, McGill University.

Young, D. (2002). The hyperbow controller: Real-time dynamics measurement of violin performance. In *Proceedings of the 2002 Conference on New Interfaces for Musical Expression (NIME02)*.