Haptic Interfaces for Musical Notation and Expression

Jonathan Lane-Smith

Department of Electrical & Computer Engineering
McGill University, Montreal

April 5, 2023

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of

Master of Science

Abstract

Haptics is a growing area of technology that has permeated many fields including robotics, health care, and entertainment. One particularly interesting field that it can contribute to is that of music. This thesis explores two main areas in which haptics can assist or enhance musical performance.

First, two force-feedback devices were used to create the Hapstrument, an expressive bimanual digital musical instrument. The left-hand device controls pitch selection with force feedback helping to guide the user to the desired note. The right-hand device controls excitation by simulating the bowing or plucking of a string and generating the corresponding haptic feedback. A user study was run to evaluate the effectiveness of the Hapstrument and received a wide range of reviews, some high and some low. For the participants whose musical backgrounds aligned with the Hapstrument's capabilities, it was an effective digital musical instrument that uses force feedback to enhance musical expression.

The second part of this thesis investigates how haptics can be used to communicate musical information to blind musicians so that they can "read" music without needing to memorize it. The current methods for learning music with blindness are learning by ear or with Braille music notation, both of which require memorization. By using a haptic vest embedded with vibrotactile motors, musicians feel vibrations which map to specific musical notes, allowing them to “read” and play their instrument in real-time. After developing the mapping technique and running a pilot study on sighted participants, a user study was conducted with ten blind musicians. Over the course of the study, each participant substantially improved their ability to identify vibrations and play the corresponding notes. However, there was still a significant delay, especially when moving between octaves. There are two possible future routes for the haptic vest. First, novice blind musicians could use the vest while initially learning their instrument. Second, the vest could communicate chords instead of individual notes, which musicians could use when playing jazz or contemporary music in an ensemble.

Through both the Hapstrument and the haptic vest, this thesis demonstrates two ways in which haptic interfaces can enhance or assist in the making of music. The results of this research will hopefully inspire future exploration into the integration of haptics with music production and education.

Abrégé

L'haptique est un domaine technologique en pleine expansion qui a pénétré de nombreux secteurs, dont la robotique, les soins de santé et le divertissement. Un domaine particulièrement intéressant auquel elle peut contribuer est celui de la musique. Cette thèse explore deux domaines principaux dans lesquels l'haptique peut aider ou améliorer la performance musicale.

Premièrement, deux dispositifs à retour de force ont été utilisés pour créer l'Hapstrument, un instrument de musique numérique bimanuel expressif. Le dispositif de gauche contrôle la sélection de la hauteur du son, le retour de force aidant à guider l'utilisateur vers la note désirée. Le dispositif de la main droite contrôle l'excitation en simulant l'archet ou le pincement d'une corde et en générant le retour haptique correspondant. Une étude sur des utilisateurs a été menée pour évaluer l'efficacité de l'Hapstrument et a reçu un large éventail d'avis, certains positifs, d'autres négatifs. Pour les participants dont les antécédents musicaux correspondaient aux capacités de l'Hapstrument, il s'agissait d'un instrument de musique numérique efficace qui utilise le retour de force pour améliorer l'expression musicale.

La deuxième partie de cette thèse étudie comment l'haptique peut être utilisée pour communiquer des informations musicales aux musiciens aveugles afin qu'ils puissent "lire" la musique sans avoir besoin de la mémoriser. Les méthodes actuelles d'apprentissage de la musique pour les non-voyants sont l'apprentissage à l'oreille ou la notation musicale en braille, qui nécessitent toutes deux de la mémorisation. En utilisant un gilet haptique équipé de moteurs vibrotactiles, les musiciens ressentent des vibrations qui correspondent à des notes de musique spécifiques, ce qui leur permet de "lire" et de jouer de leur instrument en temps réel. Après avoir mis au point la technique de mappage et mené une étude pilote sur des participants voyants, une étude sur des utilisateurs a été réalisée avec dix musiciens aveugles. Au cours de l'étude, chaque participant a considérablement amélioré sa capacité à identifier les vibrations et à jouer les notes correspondantes. Toutefois, un retard important subsistait, notamment lors du passage d'une octave à l'autre. Deux voies d'avenir sont possibles pour le gilet haptique. Tout d'abord, les musiciens aveugles débutants pourraient utiliser le gilet lors de l'apprentissage initial de leur instrument. Ensuite, le gilet pourrait communiquer des accords au lieu de notes individuelles, ce que les musiciens pourraient utiliser lorsqu'ils jouent du jazz ou de la musique contemporaine dans un ensemble.

Par le biais de l'Hapstrument et de la veste haptique, cette thèse démontre deux façons dont les interfaces haptiques peuvent améliorer ou aider à faire de la musique. Les résultats de cette recherche inspireront, nous l'espérons, l'exploration future de l'intégration de l'haptique dans la production et l'éducation musicale.

Acknowledgements

I am very thankful for the guidance and support of my supervisor, Prof. Jeremy Cooperstock. He allowed me to pursue my individual research interests while giving me valuable direction, which made the past two years very enjoyable. I am especially grateful for the time he spent reading and offering feedback for the various documents I wrote, including the many ethics applications.

I would like to acknowledge my colleagues at the Shared Reality Lab for their suggestions and advice which helped to guide my research. Particularly, Yongjae Yoo supported me right from the start, especially regarding the hardware I used; without him, the results of this thesis may have been much different.

Finally, I am very thankful to my friends and family for their support over the past couple of years. Specifically, I wish to thank my Dad, for letting me borrow his car during my user studies; Rosanna, for letting me borrow her piano; Cameron, for always being there for me unless he is doing fieldwork; and Steph, for her continual love and encouragement.

Contents

1 Introduction **1**

1.1 Thesis Overview 3

2 Background **5**

2.1 Force-Feedback Digital Musical Instruments 5

2.2 Haptics, Music, and Assistive Technology 7

2.2.1 Haptics for Musical Enrichment and Education 8

2.2.2 Complex Communication with Haptics 9

2.2.3 Haptic Communication for the Blind and Deaf Community 11

2.2.4 Haptic Musical Assistive Technology 12

2.2.5 Assistive Technology for Musical Notation 13

3 The Hapstrument **15**

3.1 Design 17

3.1.1 Pitch Selection Interface 18

3.1.2 Excitation Interface 24
3.2 Evaluation 26
3.2.1 Procedure and Recruitment 28
3.2.2 Results 29
3.3 Discussion 31

4 Haptic Interface for Blind Musicians 34

4.1 Dual-Arm Haptic Interface 36
4.1.1 Design 36
4.1.2 Pilot Study 39
4.1.3 Results 41
4.2 Haptic Vest 43
4.2.1 Initial Design 43
4.2.2 Construction 45
4.2.3 Final Design 47
4.3 User Study 49
4.3.1 Methodology 50
4.3.2 Pilot study 54
4.4 Results and Analysis 57
4.4.1 Haptic Vest 58
4.4.2 Learning by Ear 62

4.4.3 Mapping 64

4.5 Discussion 65

4.5.1 Choice of Methodology 65

4.5.2 Results 66

4.5.3 Potential Improvements 67

4.5.4 Comparison with Listening 68

5 Discussion **70**

5.1 Themes 70

5.2 Potential Limitations and Future Research 71

5.2.1 Hapstrument 71

5.2.2 Haptic Interface 73

5.3 Conclusion 75

List of Figures

3.1 Haply 2diy. 17

3.2 Block diagram of the Hapstrument. 18

3.3 Piano-like interface for pitch selection. 19

3.4 Two-octave circular layout used for pitch selection. 20

3.5 Circular layout for the C Major scale. 21

3.6 Pitch selection in Iteration 3, with red annotations added to each of the columns. 22

3.7 Final pitch selection interface with and without columns. 23

3.8 Plucking model. 24

3.9 Pure Data patch. 27

3.10 Distribution of ratings among participants. 32

4.1 Placement of actuators (Left) and full system (Right). 37

4.2 Mapping strategies 1 (Left) and 2 (Right). 38

4.3 Computer keyboard mapping. 39

4.4 Melodies for pilot study. 40

4.5 Test results. 41

4.6 First vest design. 44

4.7 Octave mapping. 44

4.8 Second vest design. 46

4.9 Inside of vest (Left) and side view of vest (Right). 47

4.10 bHaptics vest. 48

4.11 “Backwards N” vest design. 50

4.12 Learning a melody through audio with chunks. 53

4.13 “Forwards N” vest design. 57

4.14 Melodies for blind musicians. 58

4.15 Results of Step 3 and Step 4. 59

4.16 Haptic values of participants. 60

4.17 Comparison of haptic score with audio, age, and years of piano lessons. . . . 63

4.18 Results of individual notes. 64

List of Tables

4.1 Haptic values for each session. 60

4.2 T-test results for haptic values. 61

4.3 Haptic score, audio score, age, and years of piano lessons. 62

List of Acronyms

BLV blind and low vision.

CIRMMT Centre for Interdisciplinary Research in Music, Media, and Technology.

DIY do it yourself.

DMI digital musical instrument.

DoF degrees of freedom.

ERM eccentric rotating mass.

GUI graphical user interface.

HCI human computer interaction.

HMD head-mounted display.

HMP haptic music player.

MIDI Musical Instrument Digital Interface.

NIME New Interfaces for Musical Expression.

OSC Open Sound Control.

SRL Shared Reality Lab.

Contributions

First, this thesis contributes to musical research by demonstrating that low-cost haptic force-feedback devices may be used to effectively model musical excitation strategies such as bowing or plucking, and that these devices can be used to create expressive musical interfaces. Additionally, this research shows that the experiences people have with new digital musical instruments will be heavily influenced by their personal musical backgrounds and expectations.

Second, this thesis demonstrates that blind musicians can use a haptic interface to read musical notes while simultaneously playing their instrument, introducing an alternative to memorization. From just three sessions of 45 minutes, participants were able to learn the vibrational patterns of a haptic vest, reaching speeds of up to 33 notes per minute in a 2-octave range. Participants enjoyed using the vest, with many of them saying that it has significant potential for future development.

Chapter 1

Introduction

Haptics is a growing area of technology that has permeated many fields including robotics, health care, entertainment, and more. Because haptics is separate from the visual and auditory channels, it presents a method of communication that is largely unaffected by the many inputs that humans regularly experience with their eyes and ears. In fields where communication is important and mental processing should not be overloaded, haptics offers a channel in which users can receive additional information while not interfering with audiovisual stimuli.

One particularly interesting field that haptics can contribute to is that of music. The production of music has undergone vast changes with the advent of electric and digital instruments, which provide many new methods of generating music. The traditional relationships between physical shape and sound can be replaced by alternative interfaces with unique ways of selecting pitches and playing notes. However, the natural haptic feedback of traditional acoustic instruments is integral to musical expression, which is the way in which musicians can feel connected with their instruments and the music they are producing. Common features of musical expression such as vibrato or variation in volume often require the musician to vary the amount of force they apply, and subsequently receive, from their instrument. Therefore, haptic feedback should not be overlooked when designing a new musically expressive interface.

Haptic technology can be used to enhance digital musical instruments (DMIs) while not interfering with the instrument's sound or introducing visual feedback. This can allow musicians to feel physically connected with the music they are making, improving their ability to express themselves. However, force-feedback devices can often be expensive, and it can be difficult to generate effective haptic models. This motivates the first objective of this thesis, which is to explore the effect of force feedback on a novel low-cost digital musical instrument and the contribution it makes to overall musical expression.

In addition to musical performance, haptics can be applied to music education. Learning an instrument is a long and tedious process and the majority of instruments take many years to master. Assistive tools, such as metronomes or visual computer programs, can be used to make the learning process easier, but these tools use visual or auditory channels, which may be distracting to a music student. An alternative method is to introduce haptics, which can communicate information to a music student without overloading audiovisual stimuli.

Haptics can also be used in assistive technology. People with visual or auditory disabilities may be unable to use one of the communication channels that most humans use on a regular basis. This may limit their ability to receive pertinent information, requiring them to find an alternative method of doing so. To communicate through haptics, the information would have to be given a specific mapping, and it would take time for the brain to learn how to decode this information. However, with enough time and practice, the brain can learn how to read and understand haptic messages quickly and efficiently.

Building on this theme, the second research focus of this thesis combines music education and assistive technology by exploring whether a haptic interface can be used to communicate musical notation to blind musicians so that they can "read" music and play it in real-time. Because blind musicians often cannot visually read sheet music, the most common methods for learning a musical composition are to learn by ear or with Braille music notation, both of which require memorization. The haptic interface provides an alternative channel that the brain can use for reading music without the need for memorization.

1.1 Thesis Overview

This thesis is divided into five chapters: the introduction, a review of the relevant background literature, two chapters that explore the two main research objectives of this thesis, and a discussion.

Chapter 2 gives a comprehensive summary of the background literature relevant to both research objectives of this thesis. First, it discusses previous work that has been done with force-feedback DMIs, consisting mainly of submissions to a conference on New Interfaces for Musical Expression (NIME). Second, this chapter discusses the various ways that haptics has been used for musical interaction and communication, and the work that has been done to assist the disabled community in learning and playing music.

Chapter 3 presents the Hapstrument, a bimanual haptic interface for musical expression. Using two low-cost 2-DoF force-feedback devices, a DMI was created with the ability to select pitches and play notes by simulating the act of bowing and plucking. For the left hand, haptic feedback was used to help guide the user to the correct note, while for the right hand, it was used to generate a realistic sensation of excitation.

Chapter 4 explores how a haptic interface may be used to communicate musical information to blind musicians so that they can "read" music without the need for memorization. A number of interfaces and mapping strategies were explored and this culminated in a user study with blind musicians, to validate the effectiveness of the interface.

Chapter 5 discusses the research completed in this thesis by giving a summary of the overarching themes and stating the potential limitations. Furthermore, this chapter discusses directions for possible future research building upon the results of this thesis. Finally, this chapter concludes this thesis by revisiting the original objectives and discussing whether they were successfully met.

Chapter 2

Background

2.1 Force-Feedback Digital Musical Instruments

The first objective of this thesis, to explore the effect of force feedback on a novel low-cost DMI and its contribution to musical expression, was influenced by former DMIs and their use of force feedback. These interfaces, mainly developed by the NIME community, use haptic effects as a means of communication, expression, or to imitate the characteristics of acoustic instruments. Reviewing this literature inspired the research team with new ideas regarding force feedback in DMIs, and helped them to determine realistic constraints for the project.

Among the relevant literature, a force-feedback musical system was developed which used a brake-augmented ball pen stylus on a sticky touch-sensitive surface [1]. Haptic effects were used to simulate various textures, which could be mapped to musical parameters.

An enhanced handle for a Phantom Omni combined vibrotactile feedback with force feedback to simulate bowing [2]. By augmenting the handle with a vibration actuator, the high-frequency portion of acoustic bowing was included in the haptic response.

A force-feedback device was built from a surplus disk drive and controlled with a simple microcontroller [3]. Having only a single degree of freedom (1-DoF), this device used various haptic illusions, such as a slope or clutch illusion, to simulate various terrains and map these to audio parameters.

Sheffield et al. used two rotational force-feedback knobs to convey physical models for various musical applications [4]. These knobs could be “linked” together so that the user could feel on one knob the effect of the movement of the other knob. One of the inspiring applications provided in this paper is using the knobs to simulate plucking harp strings. This paper inspired the research team to use force feedback to imitate the plucking sensation with the Hapstrument, the novel DMI of this thesis.

Howard et al. created a music synthesis system that allows virtual instruments to be controlled with a force-feedback joystick and force-feedback mouse [5]. The authors used force feedback to simulate various excitation methods, such as plucking and bowing a string. This paper inspired the research team to look beyond just plucking and to also include the sensation of bowing in the Hapstrument.

Steiner used a haptic joystick and haptic mouse to create a new DMI entitled StickMusic [6]. The joystick was used to select the timbre and the mouse was used to select the pitch and amplitude. Haptic feedback assisted with note and timbre selection. This paper helped to guide the research team in the way they separated functions between the two force-feedback devices of the Hapstrument.

Berdahl and Kontogeorgakopoulos created the FireFader [7], a 1-DoF force feedback device that has been used in several musical projects in the NIME community [8–10]. The FireFader's prevalence and low cost show that inexpensive force-feedback devices can be used to create successful and effective DMIs.

2.2 Haptics, Music, and Assistive Technology

The second objective of this research, to explore whether a haptic interface can be used by blind musicians to "read" music, builds upon a number of areas of research. First, haptic technology has been applied in many musical applications to enrich musical experiences and assist in education. Second, haptics has been used to communicate language and other complex information to humans through various mapping strategies. Third, haptic interfaces have been developed for communication with the blind and deaf community. Fourth, haptics has been applied to musical applications for people with disabilities. Finally, assistive technology has been developed regarding musical notation. Reviewing this literature provided relevant information which aided the development of the haptic interface described in this thesis.

2.2.1 Haptics for Musical Enrichment and Education

There are many examples, both in research and in commercial products, of haptics being used to enrich musical experiences for both audiences and musicians. Turchet, a designer of musical haptic systems [11], created a wearable device that can allow a listener to physically feel music while listening, improving engagement and enjoyment [12]. A haptic glove was developed in order to amplify the mood of film music, such as calmness or excitement, through the use of vibrations with various intensities and frequencies [13].

Regarding musicians, virtual musical instruments have been augmented with haptics so that they feel more realistic and are easier to learn [14]. Mid-air haptic feedback, produced with ultrasonic vibrations, was used to create a touchable piano in virtual reality [15]. As discussed in Section 2.1, there are many examples of force feedback being used to enhance musical instruments, and by extension, the experience of the performer.

Another application of haptics in music is for teaching musical instruments. Research has been done in this area, shown by haptic bracelets that can be used for learning multi-limb rhythm skills [16], an electronic glove with vibration motors on each finger to aid in learning simple piano pieces [17], and an interactive haptic feedback system to improve a violinist's form when using a violin bow [18]. These haptic interfaces can help a musician learn various skills or techniques, although none of them address the challenge of reading sheet music.

2.2.2 Complex Communication with Haptics

As the haptic interface in this thesis was required to communicate information to users as quickly and accurately as possible, related research was explored to inspire actuator placement and mapping strategies.

One of the earliest tactile languages was Vibratese, developed in 1957 [19]. Five vibrotactile actuators were placed on the chest, in an "X" formation, and each actuator had three possible intensities and durations, resulting in 45 different combinations. Vibratese could communicate the full English alphabet and the numbers 0-9, and with enough training, rates of up to 67 words per minute could be reached [19]. Among users, the majority of the mistakes came from misidentifying the intensity levels of the vibrations.

Since then, there has been much research into optimizing vibrotactile information transfer. Rabinowitz et al. explored the effect of intensity, frequency, and contact area on the information transfer of vibratory stimuli [20]. Luzhnica and Veas found ways to optimize vibrotactile reading such as by reducing simultaneous vibrations to improve accuracy [21].

Brewster and Brown coined the term tactons, referring to structured messages that can communicate information through haptics [22]. Tacton construction can use features such as location, frequency, intensity, or duration to create unique "symbols" with specific meanings. The introduction of tactons led to many more wearable interfaces for haptic communication.

Jones et al. designed a wireless tactile display of 16 vibrating motors to be mounted on the lower back [23]. Eight tactile patterns were created using temporal sweeps or blinks. The patterns were used to communicate directional information or alerts and were designed to mimic their instructions, e.g., a sweep to the left to indicate moving to the left. A user study was conducted and showed the patterns to be very successful at communicating the information quickly and accurately.

De Vargas et al. used a novel technique for communicating speech to users with haptic armbands [24]. The locations of the bands were mapped to specific regions of the mouth, and by varying the locations and characteristics of the vibrations, the bands could communicate distinct English phonemes.

Haptic interfaces have even been developed for human-animal interaction. Golan et al. created a vibrotactile vest to be worn by dogs so that handlers can communicate with them without the need for visual or audible cues [25]. The haptic commands consist of vibrations that differ in waveform or location. This vest can be used by police dogs, deaf dogs, or the dogs of owners with speech impairments.

With much relevance to this thesis, haptic interfaces have been created specifically for musical communication. Turchet and Barthet designed haptic wearables to be worn on the chest, feet, or arms, in order to facilitate communication between performers, conductors, and sound engineers [26]. This research was also applied to ensembles of visually-impaired performers [27].

A haptic belt was developed to convey musical phrases to a musician in real-time using ten tactons [28]. These tactons had unique vibrational patterns and users would memorize which phrase corresponds with each tacton. Building upon this research, an entire full-body suit was developed to allow for more effective and meaningful information delivery [29]. This research demonstrates that haptic wearables can effectively communicate real-time musical information.

2.2.3 Haptic Communication for the Blind and Deaf Community

There are several examples of haptics being used to communicate information to blind or deaf people. McDaniel et al. developed a haptic belt that conveys nonverbal cues to blind individuals during social interactions, so that they may know the directions and distances of other people in the room [30].

Nicolau et al. designed a haptic Braille-reading device called UbiBraille [31]. This device consists of six vibrotactile actuators that are placed on a user's fingers to communicate Braille characters. The results of a user study showed that participants' recognition rates were significantly influenced by their level of expertise with Braille. The most proficient users were able to read at a speed of one character per second.

Duvernoy et al. created a haptic communicator device called HaptiComm for deafblind communication [32]. Users would feel taps or swipes on their hands that would communicate letters of the Australian deafblind tactile fingerspelling alphabet. Unlike the other examples in this section, the actuators in the HaptiComm are not vibrotactile, but simply "poke" the hand at specific locations.

Vuijk et al. developed a mobile game called PatRec for teaching haptic communication to people with deafblindness [33]. Using a chair interface with a 3×3 array of vibrotactile motors, users had to guess the correct pattern based on what they felt, and they received scores based on their answers. This process was further gamified by introducing leaderboards and badges, leading to more entertainment and engagement with the system.

Eagleman and Novich developed a haptic vest which communicated speech to deaf users through 32 vibrotactile actuators situated around the torso [34]. With enough training, users were able to sense the vibrations and immediately translate them into words, substituting the sense of touch for the sense of hearing. This research shows great promise for the field of haptic interfaces, demonstrating that sensory substitution can make tactile communication fast and efficient.

2.2.4 Haptic Musical Assistive Technology

Haptic assistive technology for the blind and deaf community has also spread to the field of music. Haptic music players (HMPs) are interfaces that use vibrotactile actuators so that deaf people may experience music, communicated through vibrations of varying intensities and frequencies. HMPs have been developed that use chairs [35], sleeves [36], gloves [37], or other types of interfaces for haptic communication. Remache-Vinueza et al. give a summary of existing HMPs and their various attributes [38].

Vibrotactile interfaces have been used to improve musical education for the deaf. Petry et al. developed a vibrotactile music-sensory-substitution device in the form of a bracelet that can be worn by deaf musicians [39]. This device communicates rhythms through vibrations, simplifying the process of learning music and giving musicians greater confidence in their playing. Similarly, Tranchant et al. created a haptic platform that vibrates in rhythm to assist deaf dancers in synchronizing their movements to the beat [40].

2.2.5 Assistive Technology for Musical Notation

There has also been significant research to assist the blind and low vision (BLV) community regarding musical notation. The first research in this area was from Louis Braille himself, who developed Braille music notation. This system communicates notes, key and time signatures, and other musical symbols [41].

In recent years, there have been a number of developments that help BLV musicians to learn Braille music notation. Payne et al. created SoundCells, a browser-based tool for quickly notating print and Braille music scores with text [42]. Borges and Tomé created Musibraille, a software designed to help children learn musical notation [43].

Finally, Lussier-Dalpė et al. used a head-mounted display (HMD) to assist low-vision musicians in reading sheet music [44]. With the HMD, users employed various approaches to read music, such as magnification, highlighting, or contrast adjustment. An exploratory study was conducted with this technology and the results showed that while the HMD can be used to overcome a number of problems, reading sheet music still remains a complex and difficult task.

Chapter 3

The Hapstrument

The first objective of this thesis is to explore the effect of force feedback on a novel low-cost digital musical instrument and to what degree haptic feedback can contribute to musical expression. The Hapstrument is a DMI that was created for this objective. This work began as a collaboration between Derrek Chow, Sahand Ajami, and the author of this thesis, as part of a haptics course entitled CanHap 501 [45]. During this course, the Hapstrument was envisioned, designed, and created. Three iterations occurred through which ideas were tested and refined to improve the Hapstrument. A small informal evaluation occurred near the end of the course. After the course had finished, further improvements were made and an official formal evaluation was undertaken to properly analyze the effectiveness of the instrument. A short video showcasing a sample performance with the Hapstrument can be found at `https://www.youtube.com/watch?v=WpKuaUBec8M`.

During CanHap 501, the author was tasked with designing and creating the pitch selection interface, while Derrek and Sahand worked on the excitation interface, which was split into plucking and bowing. The initial bowing technique had several significant flaws so the author redesigned it to be the model described in this chapter. He also conducted an informal evaluation and completed a sample performance with the Hapstrument. Once the course was finished, the author planned a formal evaluation procedure, conducted a user study with 11 participants, and analyzed the data. He wrote and submitted a paper about the Hapstrument to NIME 2023, a conference on new interfaces for musical expression, and the associated content forms the core of this chapter. Sahand Ajami, Juliette Regimbal, and Jeremy Cooperstock proofread and gave valuable feedback for the NIME submission.

3.1 Design

The Hapstrument is a bimanual DMI consisting of two Haply 2diy devices, a low-cost pantograph-inspired force-feedback device produced by Haply Robotics in Montreal [46], shown in Figure 3.1.

Figure 3.1: Haply 2diy.

We decided to set one Haply 2diy to control the pitch of the instrument, and the other 2diy to control volume and timbre through physical models that use force feedback to simulate plucking and bowing. We chose pitch to be controlled by the left-hand 2diy and excitation to be controlled by the right-hand 2diy because this division of tasks is the standard configuration for most stringed instruments, such as violin and guitar.

As illustrated in Figure 3.2, the end effector positions of the two Haply 2diy devices are passed from the *Haply hAPI* library to Processing through serial communication, and force feedback is calculated and delivered. Processing generates visuals for both devices, which are presented through a graphical user interface (GUI). Sound control commands from Processing are passed to *Pure Data* as Open Sound Control (OSC) messages, and external speakers play the synthesized output. Through this setup, users will hear the music they are making while simultaneously feeling force feedback from the instrument and seeing the positions of the end effectors in the GUI.

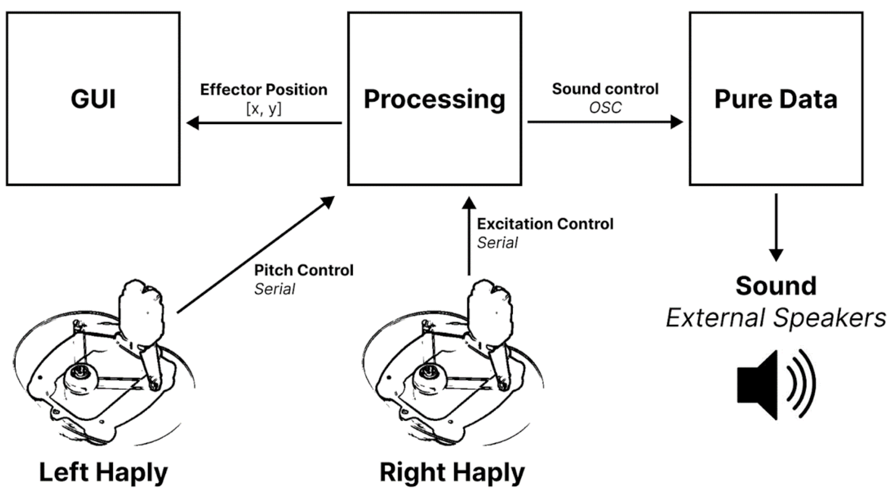

Figure 3.2: Block diagram of the Hapstrument.

3.1.1 Pitch Selection Interface

The purpose of the pitch selection interface is to provide musicians with an easy way of selecting a pitch/frequency for the synthesized sound. Through an iterative design process, a number of pitch selection interfaces were created and tested.

The first design employed a piano layout, as illustrated in Figure 3.3. Unlike the quantization of a conventional piano to discrete notes, the pitch could also be modified continuously by moving the end effector horizontally below the representation of the keyboard. This design was inspired by the Haken Continuum [47], which supports continuous pitch selection using a keyboard-like interface. While moving the end effector above the piano, the pitch remained constant, thus permitting the musician to transition between non-adjacent notes. The ability to have both discrete and continuous notes provided the musician with multiple ways of expressing themselves. Finally, with this design, the Haply 2diy generated a haptic "bump" when moving over a line, giving musicians force feedback to help them move from note to note.

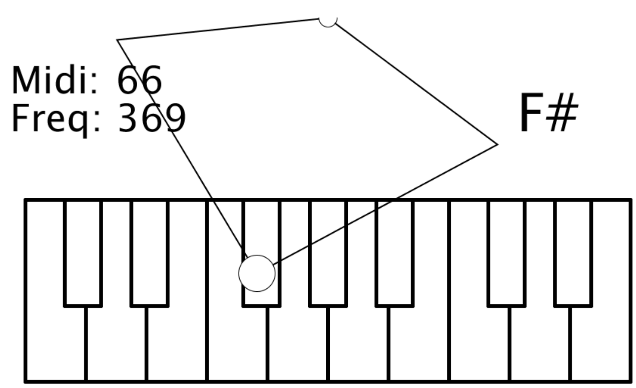

Figure 3.3: Piano-like interface for pitch selection.

We tested this design and discovered some significant limitations. Most notably, it was very difficult to quickly move from a note to any note not directly beside it. A piano interface ultimately has its design because, from an ergonomic perspective, it fits the affordances of human hands very well. For a single end effector, a keyboard representation is not the ideal interface. As a result, for the second iteration, we replaced the piano interface with a 2-octave circular layout, as shown in Figure 3.4.

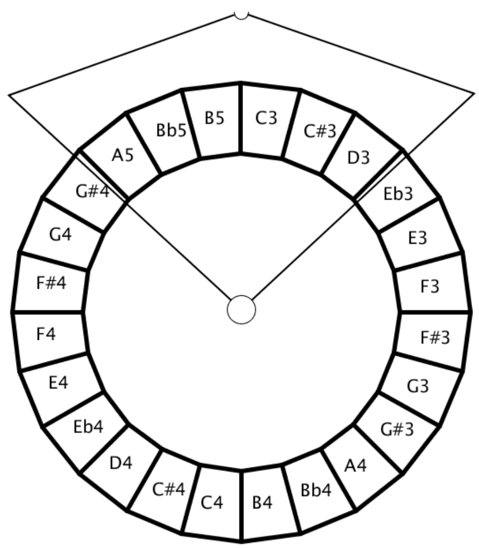

Figure 3.4: Two-octave circular layout used for pitch selection.

This layout made it much easier to move between different notes, with the pitch remaining constant inside the circle and continuous pitch selection occurring outside the circle. Furthermore, a feature was implemented which allowed a specific key and scale to be selected so that only the relevant notes would be available to play, as shown in Figure 3.5.

Similar to the piano keyboard, all of the lines in this interface provided force feedback as “bumps”, so that users could tell when they moved from one note to another. However, the force feedback of the bumps seemed to decrease higher up on the surface of the Haply

2diy. We believed that this was due to an increase in the internal friction of the 2diy, so the strength of the force feedback was programmed to be greater when the end effector is near the top of the interface.

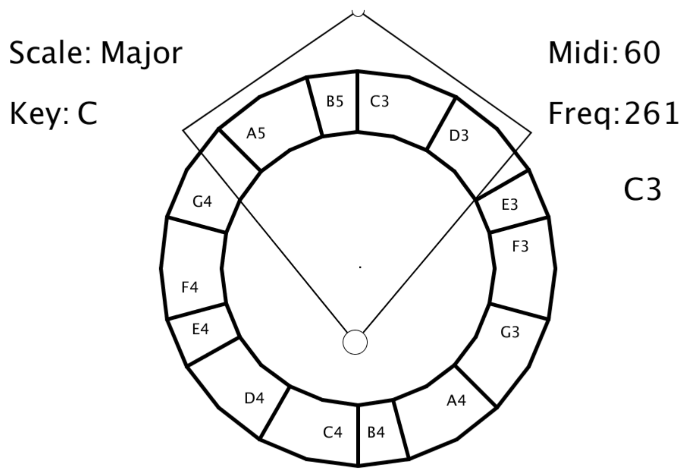

Figure 3.5: Circular layout for the C Major scale.

For the third iteration, a number of changes were made to the pitch selection interface, as shown in Figure 3.6. Three columns were added so that users may select keys and scales while using the interface. Moving the end effector over column 1 selects the desired key while doing so over column 2 selects the scale.

Column 3 shows the saved key-scale combinations. A user can move the end effector over one of these segments to select it, and then any change to the key or scale will be saved there. This allows a user to easily switch between three key-scale combinations, without needing to manually select the key and scale each time. As with the other lines of the interface, there is a "bump" sensation when the end effector moves across any of these lines, which helps the

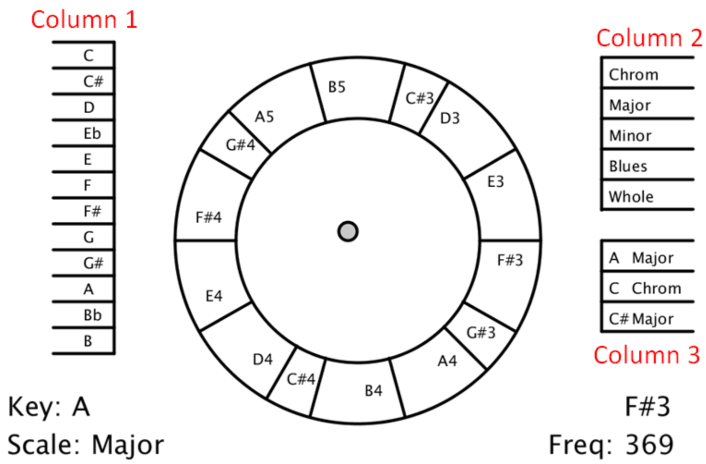

Figure 3.6: Pitch selection in Iteration 3, with red annotations added to each of the columns.

user to know when a different key/scale has been selected.

The second change for this interface was to make the circular rings actual circles instead of 24-sided polygons. This was done to fix a bug in which if the end effector crosses through the “circle” at an intersection of the lines, the program might not detect it and force feedback might not be delivered properly. By implementing a real circle, this bug was eliminated and the interface felt much more smooth and polished.

Third, the ability to implement vibrato was added to the pitch selection interface, giving musicians additional options for expressing themselves. When the end effector is inside one of the discrete ring segments, the user can shake it back and forth, and the amplitude and frequency of the movement will map to the properties of the vibrato.

The fourth and final change to the third iteration is the introduction of what we call the "elastic force", for when the end effector is in a discrete segment. We noticed that while moving around the ring, the end effector would often accidentally move outside into the continuous zone. To make it easier to prevent this from happening, an elastic force was added in the direction of the center of the segment, where the magnitude of the force is dependent on the distance from the segment's center.

We conducted a small informal evaluation near the end of the course to receive feedback and suggestions for improvements. We observed that the presence of the outer columns discouraged the user from moving the end effector outside of the ring into the continuous zone, due to fears that the end effector would accidentally go into these columns.

To fix this issue, a "button" was added to the top right of the interface. By moving the end effector onto this button, it toggles whether the columns are visible, providing the user with the ability to select keys/scales but also to hide these columns when they are not needed. The final pitch selection interface is shown in Figure 3.7.

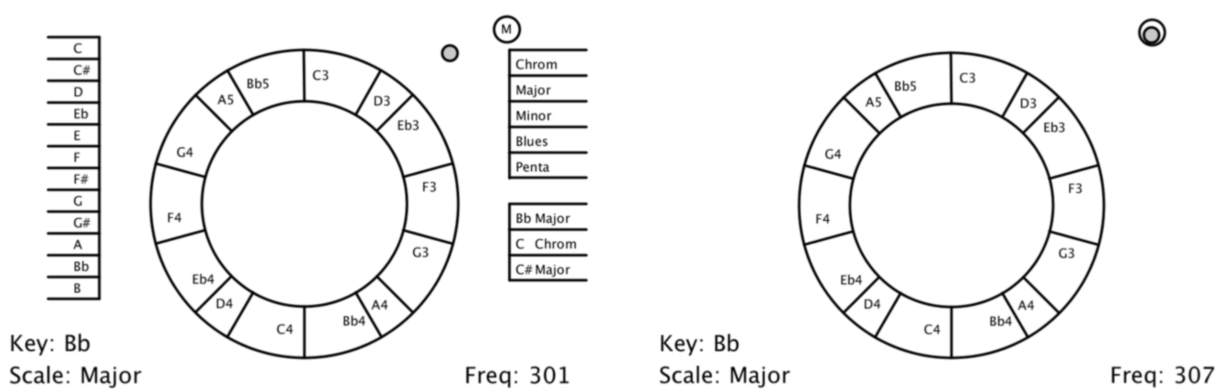

Figure 3.7: Final pitch selection interface with and without columns.

3.1.2 Excitation Interface

While the left-hand Haply 2diy focuses on pitch selection, the right-hand 2diy focuses on note excitation. For this purpose, two approaches were selected: the simulation of a pluck, such as when playing guitar, and the sensation of moving a bow across a string instrument such as a violin or cello.

For the plucking simulation, a vertical line was drawn to represent a string, and the haptic response was modelled as extending a spring. However, past a threshold distance away from the center of the string, the force will be removed suddenly, causing the user to feel the plucking effect. This was made more realistic by adding a subtle vibration at the moment after the pluck, to simulate the resonance of an acoustic instrument. The GUI of the plucking model, shown in Figure 3.8, contains a representation of a string which bends and vibrates to imitate the characteristics of a real acoustic string as it is plucked.

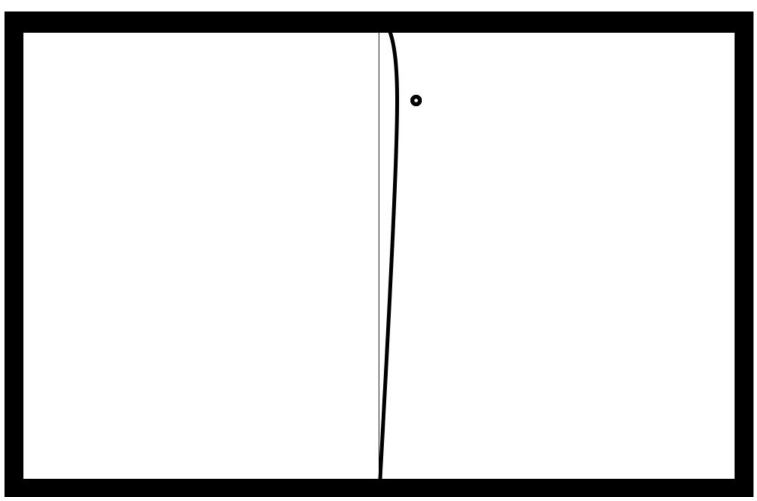

Figure 3.8: Plucking model.

For the bowing approach, we note that users playing a real violin or cello experience a resistive force dependent on how hard the bow is being pressed against the string. Unfortunately, the Haply 2diy does not have a sensor to detect how hard the handle is being pressed against its surface. However, among stringed instruments, there is often a correlation between the speed of the bow and its force against the string.

Therefore, we chose to calculate the resistive force as proportional to the speed of the end effector, as modelled by the following equation, where f is force feedback, $\dot{x}$ is speed, and c is an experimentally determined constant.

$$f = -c\dot{x} \tag{3.1}$$

However, this in itself was not sufficient to replicate the desired sensation, as bowing produces a vibrational force that gives acoustic haptic feedback a “rough” sensation. To address this, we introduced small, random variations in the force feedback to the Haply 2diy. Specifically, every millisecond, both the x and y components of the bowing force were slightly changed by a random amount within a predefined range. This procedure did not change the average force but simply introduced some rough texture to the Hapstrument.

Finally, to prevent the force from being volatile and changing too quickly, we set the force to be a weighted sum of the current and previous values. By adjusting the weights, we found a mix that produces a believable and enjoyable bowing model.

Based on encouraging initial feedback from the members of the research team, we incorporated both models into the Hapstrument, with the top half of the Haply 2diy workspace being used for plucking and the bottom half for bowing. As the end effector moves between the zones, there is a smooth interpolation between these models for haptic feedback, GUI display, and sound output.

At first, a linear interpolation was used; however, it was difficult to only pluck or bow because unless the end effector was at the very top or bottom, it would be a combination of the two. To address this issue, a sigmoid interpolation was chosen, which yielded better results.

The volume of the synthesized sound in Pure Data was set to be proportional to the velocity of the end effector while bowing or at the moment of plucking, similar to the resistive force mentioned earlier. We set the default timbre of the instrument to be a synthesized sine wave. If the volume goes above a certain threshold, a sawtooth wave is gradually introduced to add spectral content at higher frequencies. This provides more opportunities for expressivity while still focusing on the haptic interface. The Pure Data patch is shown in Figure 3.9.

3.2 Evaluation

Two evaluations took place for the Hapstrument: a small informal evaluation during CanHap 501 and a more comprehensive formal user study afterwards. For the initial evaluation, two participants were given time to play with the Hapstrument and try to musically express

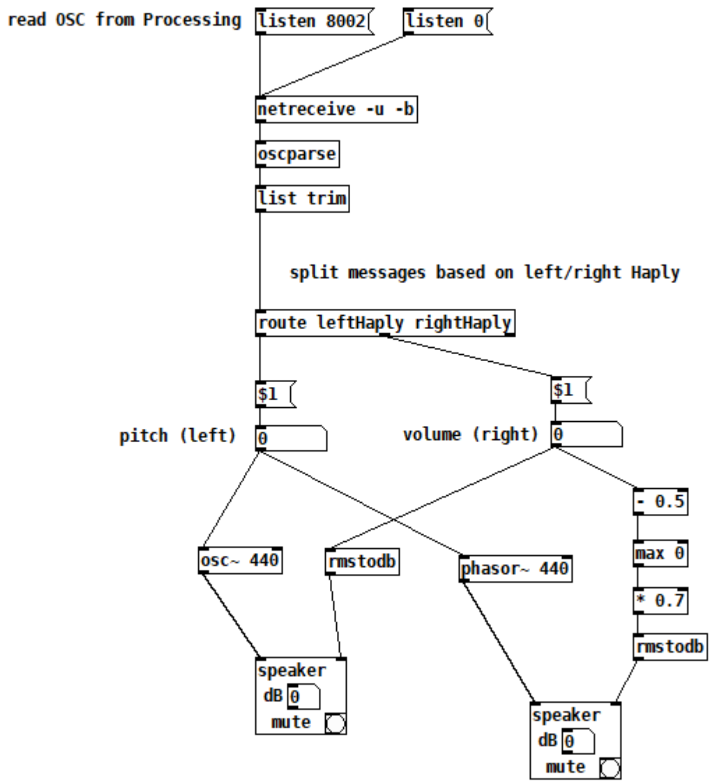

Figure 3.9: Pure Data patch.

themselves. The participants completed questionnaires before and after the experiment to give information about their musical background and their opinions of the instrument.

Overall, the participants enjoyed the plucking and bowing aspects of the instrument but had difficulty using the pitch selection interface. They also disliked some parts of the Haply 2diy hardware, such as its metal surface which was described as “grating” when moving the end effector on it. In response to this initial evaluation, a small change was made to the pitch selection interface as described in Section 3.1.1, and the metal surfaces of both devices were covered in a layer of tape to reduce the friction with the end effector and make the overall experience smoother.

3.2.1 Procedure and Recruitment

The formal user study was undertaken among members of the Centre for Interdisciplinary Research in Music, Media and Technology (CIRMMT) community at McGill University. Eleven people were recruited for evaluating the Hapstrument (5F, 6M, ages 22 to 36, average 28.4). These participants had a range of musical backgrounds. Some had no formal training, some grew up playing regularly, and some had a degree in music.

Before the test, participants were given an introductory survey, asking them about their musical skill, their exposure to DMIs, and their experience with haptic force-feedback devices. During the half-hour test, participants were introduced to the Hapstrument and given a brief explanation of its functionality. They were then given time to explore the instrument and try it out. Following this, they were asked to complete several tasks, such as playing a scale, plucking, bowing, and transitioning between the two. Next, participants were given 15–20 minutes to play with the Hapstrument and prepare a performance for the end of the session, in which they would play a composition with *musical expression*. This was defined to them as the ability to invoke an emotional response or bring the music to life.

Participants played their composition and completed a post-test questionnaire, consisting of open-ended questions about what they liked and disliked, as well as eight Likert scale questions regarding the following aspects of their experience:

- Pitch selection interface
- Feeling of plucking
- Feeling of bowing
- Plucking/bowing interface
- Overall experience with the Hapstrument
- Ability to express yourself
- Impact of haptic feedback on expressing yourself
- Impact of mixing plucking and bowing on expressing yourself

Throughout the study, the verbal observations of participants were written down. An analysis was conducted with this information along with the results of the post-test questionnaire.

3.2.2 Results

The pitch selection interface received mediocre reviews overall, having an average value of 3.3 when rated from 1 to 5. The idea of using a scale/key system to select distinct tones was described by a participant as “Eurocentric”. Some participants liked this system and said that it made it easy to play a piece in the chosen key. However, other participants described it as very limiting for them.

When a scale is selected, some of the segments may be larger than others due to the distances between notes in a scale, as seen in Figure 3.6. Some participants liked this, saying that it is good for visualizing music scales. However, the smaller segments were often harder to select, and many participants had difficulty with this.

The haptic feedback of this interface, which consisted of “bumps” when moving over lines, had generally good reviews. Most participants liked it, saying that it improved spatial awareness. However, other participants reported that the additional force required made navigating the interface more cumbersome. Other positive aspects of the pitch selection interface included the ability to add vibrato and the presence of the continuous zone outside of the ring.

The excitation interface had slighter higher reviews, receiving an average value of 3.7. People generally enjoyed plucking but had a range of preferences for the strength of the haptic resistance, which some described as unrealistic. Some participants enjoyed the vibrational force that is delivered after plucking to imitate resonance but others found it distracting and “mushy”. A few participants suggested that it would be nice to be able to “lift up” the Haply 2diy handle so that you could move it across the string without plucking, although the 2diy lacks that ability.

Participants generally enjoyed bowing more than plucking. The haptic resistance made it feel lifelike, which helped users to musically express themselves. However, some of them disliked the roughness of the bowing and preferred for the force to simply be resistive. The ability to switch between plucking and bowing was one of the novelties of the Hapstrument, and 8 of the 11 participants enjoyed this feature. However, one user disliked it because “it can't be done on a real instrument", while others found it too difficult to use while focusing on the pitch selection interface.

The Haply 2diy itself had a number of issues throughout the evaluation process. The magnetic handle would sometimes fall off, and some participants disliked its ergonomics and said that their wrists felt sore. Furthermore, the 2diy itself generated a minor rattling sound when the end effector moved across the surface, which some participants found to be distracting and irritating.

Overall, the results of the post-test questionnaires exhibited significant variation across participants, with average ratings almost uniformly spread between 2 and 5 out of a maximum of 5, as shown in Figure 3.10. Some gave the Hapstrument high ratings across the board, indicative of promise for future development, while others consistently rated it poorly, suggesting that from their perspective, the instrument suffered insurmountable issues.

3.3 Discussion

The key to understanding the range of evaluations lies in the musical backgrounds of the participants. Although we did not observe a correlation between the musical education level of our participants and their post-test questionnaire results, the type of music they played was largely predictive of their enjoyment of the Hapstrument: those who focused on atonal composition and experimental music generally disliked the Hapstrument and gave an average

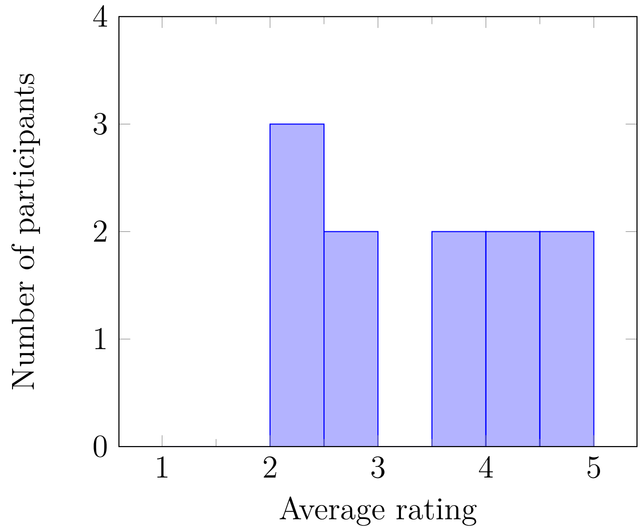

Figure 3.10: Distribution of ratings among participants.

rating of 2.5, while those who played classical music appreciated its pitch selection interface, which matched their musical expectations, and they gave an average rating of 3.8.

Participants with prior exposure to haptic devices favourably viewed the haptic sensations from bowing and plucking, saying that it was able to guide them and assist them in expressing themselves. On the other hand, those who had never used such devices often found the force feedback distracting, with some even recommending that all haptic feedback be removed.

Our main research goal in this study was to explore the impact of force feedback on a novel low-cost digital musical instrument and to find the extent to which haptic feedback influences musical expression. Based on the evaluation we conducted of our DMI, some people found it effective and realistic and were able to meaningfully express themselves, with the force feedback having a large contribution towards this goal. However, the negative reviews from different participants would indicate otherwise.

Ultimately, the evaluation of the Hapstrument demonstrates that the preferences people have towards new digital musical instruments will heavily depend on their backgrounds and the expectations that they bring into their experiences with the instrument, an observation also seen by Young and Crowley [48]. For the participants whose expectations matched what the Hapstrument could offer, it was an expressive and effective low-cost DMI. One participant even described it as the "most successful software instrument I've seen coming out of CIRMMT for the last 4 years". However, for the participants who had different expectations, this DMI was not the right choice for them.

Chapter 4

Haptic Interface for Blind Musicians

The second objective of this thesis is to investigate whether a haptic interface can be used so that blind musicians can "read" music through the sense of touch, eliminating the need for memorization. Several interfaces and mappings were explored in order to find a solution that is easy to learn and use. First, a dual-arm haptic interface was created and a small pilot study was run to compare the results of three mapping strategies. Second, a custom haptic vest was built and more mappings were explored. Finally, a commercial vest was chosen and the haptic communication strategies were refined.

With this final version, a user study was conducted, first with sighted participants to address any significant methodological issues, and then with blind musicians. A statistical analysis was undertaken to interpret the quantitative data from the study and determine how easily musicians could use the vest to play individual notes as well as simple melodies.

This analysis was used with the qualitative data from questionnaires to discuss the results of the haptic vest and consider future directions.

This was a solo project and all of this research was done independently. However, the author received advice from his supervisor, various colleagues at the Shared Reality Lab (SRL), and several blind musicians and teachers with whom he connected.

4.1 Dual-Arm Haptic Interface

The first design for this project was a dual-arm haptic interface, consisting of 12 vibrotactile actuators placed on a user's arms and shoulders. Users could map the locations of the vibrations to specific musical notes, which they could then play. For this interface, three mapping strategies were designed, and a small informal user study was conducted to determine the effectiveness of the haptic interface with each mapping. While the overall project focused on blind musicians, this subproject only recruited sighted participants, so that the blind population could be recruited for a later study. This subproject had three main research questions:

- RQ1: Are arms an effective location for communicating musical information?
- RQ2: How do the mapping strategies compare in terms of effectiveness and overall satisfaction?
- RQ3: What challenges exist in using haptic interfaces to read sheet music?

These questions were motivated by the objective to determine the best location and mapping of a haptic interface for this application and to identify the challenges that users may experience while using it.

4.1.1 Design

The interface used 12 motors altogether, as there are 12 notes in a scale. Based on the actuator placement for the work by West et al. [29], the motors were placed on the user's

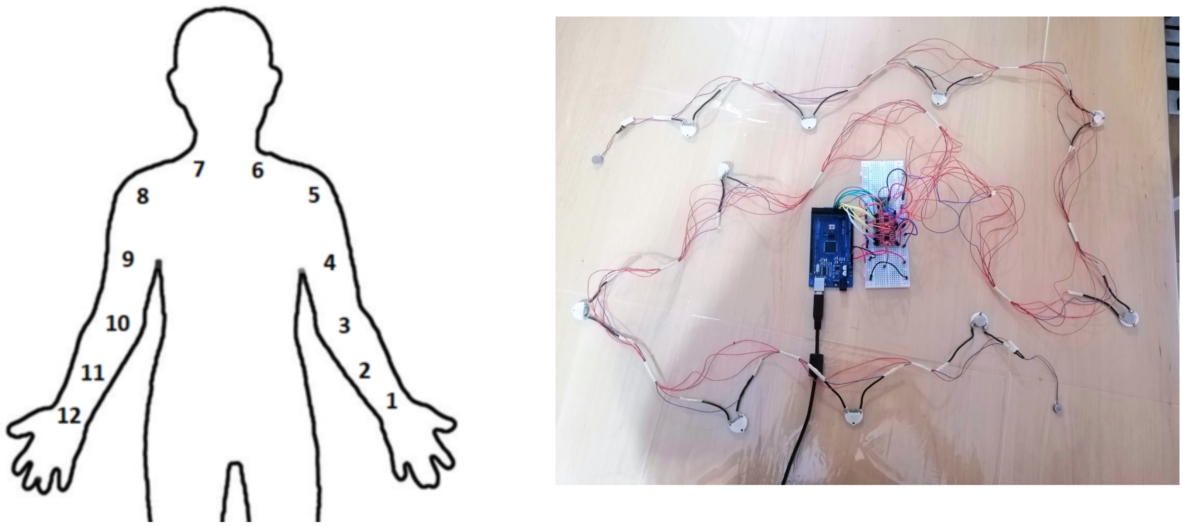

Figure 4.1: Placement of actuators (Left) and full system (Right).

arms, from the wrist to the base of the neck and back down on the other side. The system was comprised of an Arduino Mega 2560 along with three dual motor drivers, each controlling four small eccentric rotating mass (ERM) vibration motors. The system can be seen in Figure 4.1.

To mount the system onto a user, the Arduino and breadboard were first placed on the user's lap. Next, low-adhesive tape was used to gently attach the actuators to the user's skin or shirt. This allowed the motors to remain fairly snug against the body but did not make it painful when they were removed.

Three mapping strategies were designed for the interface. Two of them, shown in Figure 4.2, used exclusively the haptic interface for communication, while a third strategy also made use of audible voice commands.

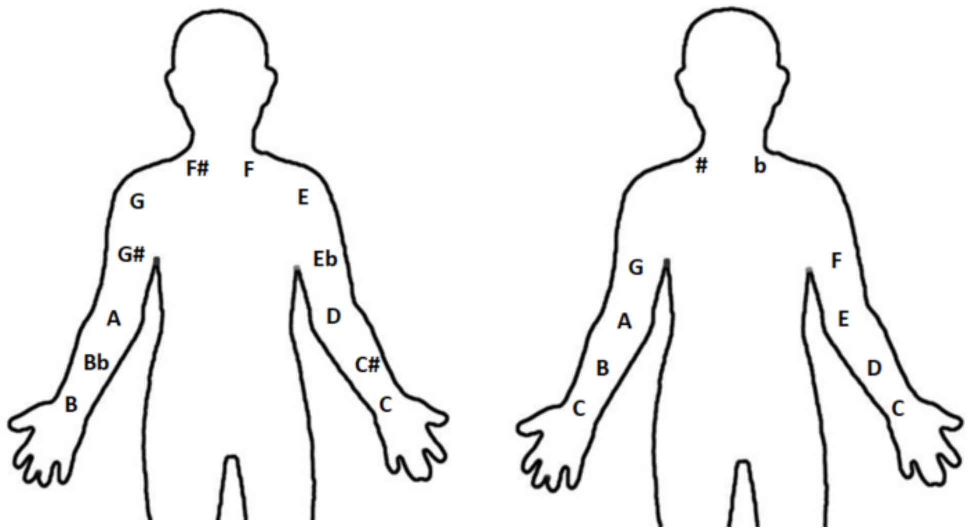

Figure 4.2: Mapping strategies 1 (Left) and 2 (Right).

For the first strategy, each actuator was mapped to a specific note, from C on the left wrist up to B on the right wrist, covering a full octave. The intensity of the actuators did not vary, although the duration that an actuator vibrated corresponded to how long the user should play the note.

The second mapping strategy took advantage of the fact that most melodies only use notes in a particular key. The 8 actuators closest to the hands (4 on each arm), were mapped to the 8 notes of the C Major scale. To allow for accidentals, the actuators at the base of the neck were used to indicate “flat” or “sharp”. For example, the note C# would result in the “C” actuator vibrating at the same time as the “sharp” actuator.

In the final mapping strategy, the name of the note was verbally spoken and the two actuators at the base of the neck vibrated to indicate the duration of the note. This strategy

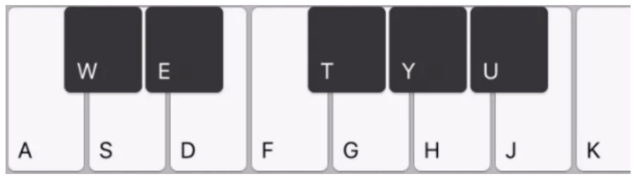

Figure 4.3: Computer keyboard mapping.

took advantage of the prior musical knowledge of participants but did not benefit from the spatial relationships used by the haptic actuators.

In order for the computer to send commands to the Arduino during the pilot study, code was written so the computer keyboard could act as a piano keyboard, as shown in Figure 4.3. By entering the desired note into the Arduino's serial monitor, the corresponding vibrational signal would be sent to the hardware.

4.1.2 Pilot Study

The pilot study had three participants due to COVID restrictions, which is not enough for a statistical analysis but enough to receive valid feedback about the interface. The main methodology of the pilot study was as follows: For each of the three mapping strategies, the user was given ten minutes to practice using it. The researcher played notes on the computer, causing vibrations to occur on the interface, and the user attempted to play the notes on the piano, receiving verbal feedback on whether they were correct. Next, a test occurred, in which a simple melody was played on the computer and the user had to attempt to reproduce it on the piano. The next note was only played once the previous note had

Figure 4.4: Melodies for pilot study.

been played correctly. Once the melody was finished, the test moved on to the next mapping strategy, until all three mapping strategies were completed. Because each test required a simple melody, three distinct monophonic melodies of equal difficulty were written, each having ten quarter notes and three half notes with two of the notes being accidentals. These melodies are shown in Figure 4.4.

Participants completed a pre-test survey before the experiment, asking them about their experience with piano and general music theory, and a post-test survey afterwards, with questions about their overall experience and preferences towards the different mappings. During the experiment, the mappings and melodies were presented in counter-balanced order, to reduce unintentional bias.

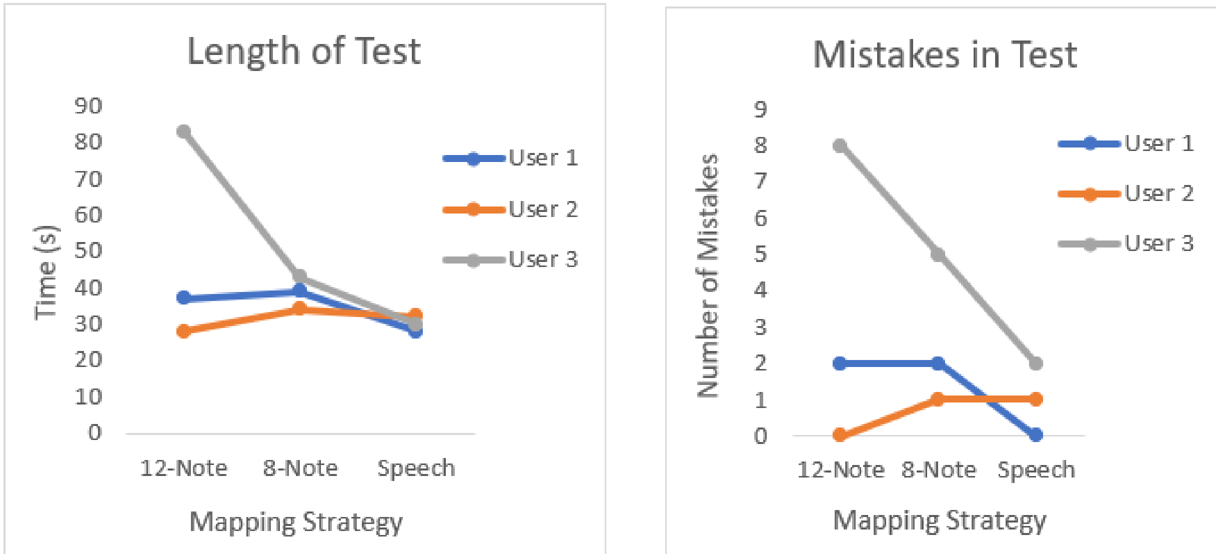

Figure 4.5: Test results.

4.1.3 Results

Once the tests were completed, the audio recordings were used to extract information regarding the time it took to complete each melody and the number of mistakes made. This information was used to create the graphs in Figure 4.5.

From the post-test surveys and comments made during the test, there were some common themes and challenges among the participants, which, along with the quantitative results, indicated some clear areas for change in future designs.

The speech strategy tended to be the most successful design, but it added additional stimuli to the auditory channel, which is an undesirable outcome for the interface. A mapping strategy with only haptics is preferable for the next design.

For the 12-note and 8-note mapping strategies, it was often difficult to detect the exact locations of the vibrations; participant 3 in particular had a lot of trouble with this and was often guessing, which suggests that the actuators should be placed further apart or on areas of the skin with higher sensitivity. Participants tended to play with their left hand when feeling vibrations on their left arm, and their right hand when otherwise. This was not an issue for this study because they were playing piano, but it suggests that for different instruments, such as trumpet or violin where the arms do not function symmetrically, it may introduce a new challenge. For the above reasons, it would seem as though the arms are not the optimal location for this application.

Participants found it quite challenging when multiple motors vibrated simultaneously, which occurred for accidentals with the 8-note strategy, and they were often only able to detect one of the vibration locations. This suggests that future designs should refrain from communicating two pieces of information via vibration at the same time.

One solution to this problem is to simply encode the information together, similar to how the 12-note mapping strategy conveyed all 12 notes individually instead of communicating the “base” note as well as whether it is sharp/flat. Another solution is to send the data sequentially so that a user is first able to distinguish the base note, followed by whether it is an accidental.

This interface was confined to a single octave, which is quite limiting when playing any non-basic melody. The next design should allow for multiple octaves. Only monophonic melodies were used during this study, but based on the difficulty users had with just a single note at a time, polyphonic melodies seem unfeasible.

Finally, in this study, users were given only ten minutes with each design of the interface, which is far from enough time to become comfortable using it. Future studies should give users significantly more time to practice and gain familiarity with the system.

4.2 Haptic Vest

Based on the results of the dual-arm interface, it was chosen to instead use a haptic vest to communicate musical notation, inspired by similar interfaces in related literature [19, 23, 29] as well as affirmative feedback from Andy Slater, a musician and composer with blindness.

4.2.1 Initial Design

The first vest design, shown in Figure 4.6, was based on the 8-note strategy from earlier. The main notes were arranged on the user's back, similar to the location of Geldard's motors with Vibratese [19], while the sharp and flat actuators were placed on the front of the vest, near the waist. The horizontal placement of the latter motors was chosen so that the right actuator corresponds to a "sharp", which raises the note up (to the right on a piano), while the left actuator conveys a "flat", which moves the note to the left. Placing these motors on the front of the vest kept them further from the main note actuators on the back, which made it easier for users to distinguish their locations.

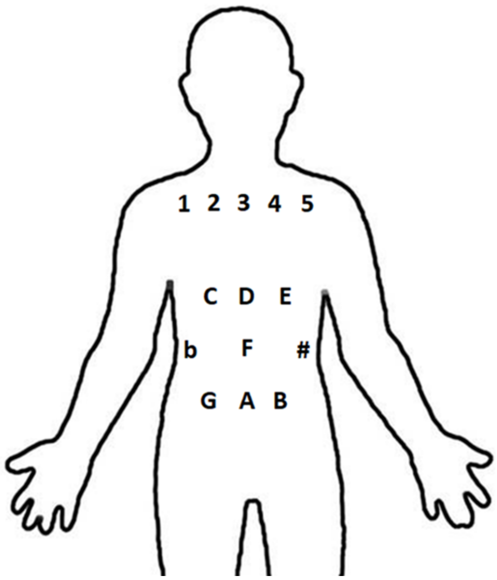

Figure 4.6: First vest design.

Five actuators were placed in a row on the back between the shoulders, to indicate octave. By adjusting the intensities of these motors, the perceived point of vibration could be anywhere along this row to inform the user of where the note is within a 2-octave range. The user did not need to discern the exact location of the vibration, as the main purpose of these actuators was just to communicate the octave of the note. This technique can be visualized in Figure 4.7.

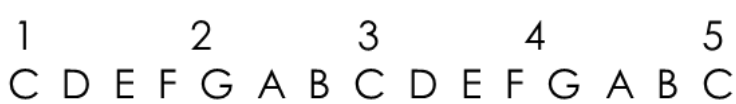

Figure 4.7: Octave mapping.

This interface would communicate notes by first sliding the point of vibration along the top row of actuators from the location of the previous note to that of the next note. The direction of movement would help the user to know which octave it will be in. Next, the main note would vibrate on the person's back. Finally, if the note is sharp or flat, the associated motor would vibrate. Through this design, multiple octaves were possible and each piece of information was delivered sequentially, addressing the issues of the dual-arm haptic interface.

Before constructing the vest, a slight change was made to the design, as shown in Figure 4.8. Instead of each of the seven notes having its own location, sweeps were used, which can be effective for communicating information through haptics [23]. Nine actuators were used instead of seven, arranged in a 3×3 grid. C was a downward sweep, such that the actuators labelled 2, 5, and 8 would vibrate in succession. D was a diagonal sweep downwards and to the left, E was a leftwards sweep, and so on, moving in a circle. This mapping gave each note a specific direction so that a user would not have to discern where exactly a vibration occurred, but rather where the vibrations moved in relation to each other.

4.2.2 Construction

This haptic vest was built by adding four additional actuators and another motor driver to the hardware of the dual-arm interface and attaching it to a simple vest. After working through many electrical wiring issues, the prototype was created, shown in Figure 4.9. The Arduino and breadboard were taped to the outside of the vest, with the wires running

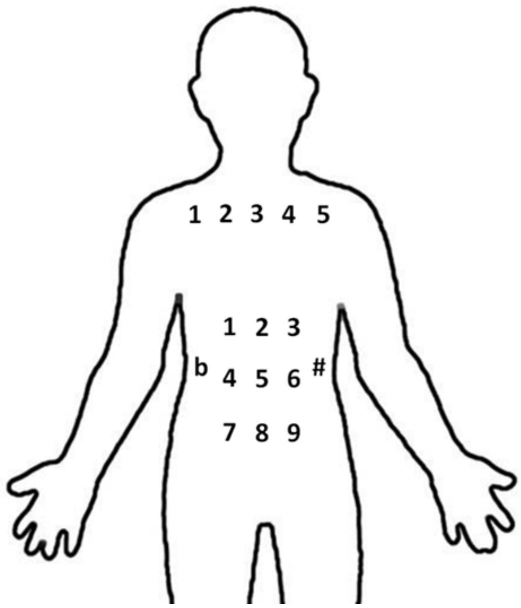

Figure 4.8: Second vest design.

underneath and up inside the vest to their respective motors, which were taped to the inner mesh. A USB cable was used to connect the Arduino to the computer, which sent vibratory commands to the interface.

To test the prototype, code was written to select a random note in the 2-octave range and play the corresponding vibrations on the vest. With some training, the author of this thesis was able to identify the haptic sensations and translate this information into musical notation, albeit with a fairly significant delay.

Unfortunately, there were some substantial hardware issues with this system. The motors were quite small and the vibrations would sometimes feel quite weak, even when the motors were vibrating at their maximum intensity. Furthermore, the vest was too loose in some

Figure 4.9: Inside of vest (Left) and side view of vest (Right).

areas and several actuators were not snug against the body, making them difficult to detect. This vest was not one-size-fits-all, so if it were used for tests on participants of different heights and sizes, it would need to be significantly augmented with various straps to tighten it.

4.2.3 Final Design

Given these challenges, it was chosen to take a different route by abandoning this prototype and instead using the bHaptics TactSuit X40, a commercial haptic vest with 40 ERM actuators, arranged as a 4×5 array on both the front and back, as shown in Figure 4.10.

With this new hardware, the previous mapping design was applied. Because bHaptic's API provides a very easy way of controlling each actuator, several alternative designs were

Figure 4.10: bHaptics vest.

also tested. This culminated in three major changes. First, the decision was made to return to the original design of having specific actuators map to notes, instead of using directional sweeps. If there were at least five seconds between sweeps, they were easy to identify, but if they occurred in quick succession, it would feel overwhelming. This is because three motors vibrated for each sweep and the vibrations all occurred in the same spatial region.

Second, the octave "range" technique had some success but initial users just wanted to know whether the note was in the top or bottom octave. Communicating the exact position of that note within the octave only made it more confusing. As a result, it was chosen to use the top row of actuators to either perform a sweep to the left, if moving down an octave, or a sweep to the right, if going up. If there was no octave change, then no vibrations occurred.

This functionality was moved to the front of the vest to give more room for the main note actuators on the back, which would vibrate immediately after the octave indication.

Third, the author considered using the bHaptics Tactosy armbands to communicate accidentals, with the left arm being used for “flat” and the right arm for “sharp”. Some initial testing was done with the full system but it took a significant amount of time to mentally think through all three parameters: octave, note, and accidental. Additionally, the armbands were very bulky and interfered with arm placement during piano playing. Consequently, the accidental indicators were removed altogether so that the system only communicated two octaves of the main diatonic notes in the C Major scale using the mapping shown in Figure 4.11.

On the back, the notes were arranged in a “Backwards N” formation and F used two actuators simultaneously vibrating at a slightly lower intensity. This design was chosen so that the C-D-E and G-A-B sequences moved in the same direction. Because the vest’s array of actuators has longer columns than rows, a vertical alignment was used. F, which is in the middle of the octave, was placed in the center of the vest. As with the dual-arm interface, the durations of the vibrations corresponded with the lengths of the notes.

4.3 User Study

From the initial study done with the dual-arm haptic interface, it was noted that a user would need much more than 10 minutes to gain familiarity with the system. Therefore, the

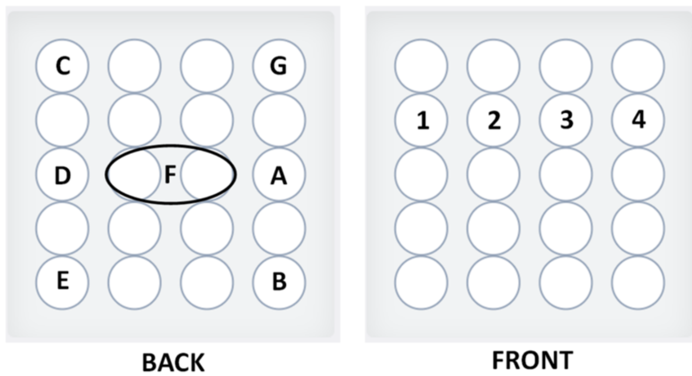

Figure 4.11: "Backwards N" vest design.

study was designed to have three separate sessions of 45 minutes each. With this structure, participants can learn the interface without experiencing mental fatigue from doing it all in a single day.

While the vest is not specified for any particular instrument, the user studies were run with an electric piano keyboard, as it is a common instrument and can connect to a computer through MIDI, allowing numerical data to be collected throughout the test. Using the bHaptics API, a Python script was written to control the vest's functionality during the test while also logging the details of the vibrations and MIDI data to a text file, to be parsed and analyzed later.

4.3.1 Methodology

In this study, participants were sent an introductory questionnaire before the test, to ask them about their musical background and abilities. During the experiments, the following procedure was used.

- Step 1 - Initial test (only first session)
- Step 2 - Introduction to vest
- Step 3 - Testing phase 1
- Step 4 - Testing phase 2
- Step 5 - Learn melody with haptics (or audio)
- Step 6 - Learn melody with audio (or haptics)

First, participants were told the names of notes, such as C or G, and would attempt to play that note on the piano as quickly as possible. This exercise only occurred during the first session, to quantify a participant's familiarity with the notes of the piano.

Second, participants were introduced to the vest and given an explanation of how it works. They could then play notes on the piano as they desired, and the corresponding vibrations would automatically be applied to the vest.

Third, once they understood how the interface works and felt comfortable with it, they would go through a "testing" phase. They would feel vibrations through the vest and would attempt to play the corresponding note on the piano. If they were correct, an audible "ding" would occur; otherwise, they would hear a low-pitched "beep". Once they played the correct note, the program would apply the next one to the vest. The notes would be randomly selected within the 2-octave range, leading to frequent jumps.

Fourth, the testing phase would continue, but each note would be randomly chosen up to a major fourth away from the previous note, causing the intervals to be smaller and making the music feel more "melodic". Furthermore, the duration of the vibrations would vary, corresponding to either a quarter or half note.

The previous steps would take 25 minutes overall. For the next 20 minutes, the participant would attempt to learn two different melodies — one with haptics and one without. For the haptic melody, participants would feel the melody played on the vest and would attempt to replicate it on the piano, with the program waiting for the correct note before applying the next note to the vest. After doing this several times, if the participant could play the melody comfortably, the code would be changed so that the program does not wait for them to play the correct note but instead plays through the melody at a set speed. If they do this successfully, the speed would be gradually increased as they become more familiar with the melody. However, this could be difficult, especially if a user starts to fall behind, so playback controls were implemented with the bottom C to G of the piano. Pressing these keys had the following effects:

- C - Go to the beginning of the melody
- D - Go back one bar
- E - Go back one note
- F - Play
- G - Pause

Initially, the plan was to implement these controls as foot pedals, but in order to keep the system simple, they were integrated with the piano. With these playback controls, a user could manipulate the melody to help them learn it easily and efficiently.

The final step of the experiment was to learn a similar melody just through audio, to compare the results of the haptic melody with the participant's ability to learn by ear. The inclusion of both methods provided a comparison that can provide insight into the participant's skill with the vest.

To learn a melody with audio, it was split into four small chunks. Participants learned the melody in sections, by listening to them and attempting to replicate them on the piano. Once played correctly, participants would move on to the next section. This process is shown in Figure 4.12. Finally, after each of the three sessions, participants completed an end-of-session questionnaire, asking them about their experience and opinions of the vest.

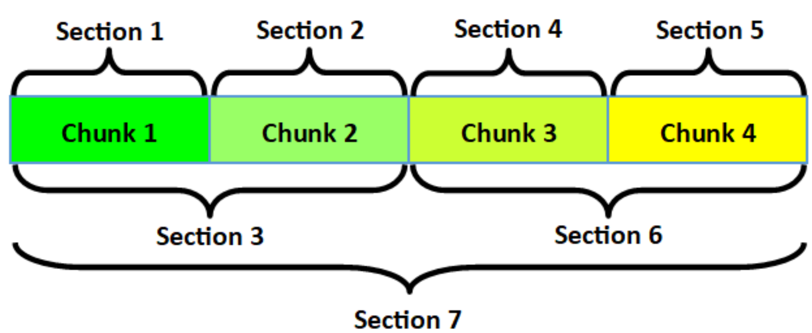

Figure 4.12: Learning a melody through audio with chunks.

In the above procedure, several variations occurred to reduce bias. For half of the participants, the melody-learning phase of the first and third sessions started off with the haptic melody and finished with the audio melody, while the second session was reversed. For the other half of the participants, this process was inverted, such that the first and third sessions began with the audio melody.

The melodies, each having the same number of notes and being of equal difficulty, were evenly distributed using a Latin square design so that they were learned through both haptics and audio during all three sessions among the participants. Using this method, each person learned all six melodies, but in a different order.

4.3.2 Pilot study

A pilot study was conducted with sighted participants before testing the interface on blind musicians. Through McGill University's Centre for Intelligent Machines, six participants (2F, 4M, ages 23 to 37, average 27.3) were recruited for the study. As per the inclusion criteria, each of them was able to identify and play piano notes by name.

In this pilot study, only two sessions occurred per participant instead of three, as the purpose of this study was not so much to collect data but to find any significant usability issues or methodological flaws in the experimental design. Once this study was finished, a simple data analysis was conducted to determine the amount of time it took for participants to feel vibrations and play the correct notes. Along with the participants' verbal comments and online questionnaires, this information suggested common themes and problems with the interface.

Overall, participants enjoyed the study and considered it very promising for blind musicians. However, they had difficulty determining the locations of some of the vibrations on their backs. The D and E vibrations were often confused with each other, as well as the A and B. As a result, a slight change was made so that for the bottom left and right actuators, the motor horizontally beside them also vibrates, which makes the sensations easier to distinguish from the middle left and right actuators.

Some participants disliked the “Backwards N” mapping design. One person suggested that the B and G could be switched, so the mapping resembles a W instead. With this design, the C and B are closer together, which makes sense considering their close proximity in music. Another person suggested doing a “Forwards N” mapping, such that the top row and bottom row are swapped. With this design, the C-D-E and G-A-B actuators go upwards as pitch increases, which is a more natural mapping. The argument behind the “Backwards N” design was that we read sheet music from top left to bottom right so the mapping should follow this pattern, but given the better reasoning of the “Forwards N” design, the latter was chosen for the experiments in the main study with blind musicians.

Participants in the pilot study found the octave sweep to be distracting from the main notes on their backs, even though the vibrations were presented sequentially instead of simultaneously. During Step 4 of the final session, the average time it took for participants to play a note after feeling a vibration with and without an octave change was 3.8 and 2.4 seconds, respectively. Evidently, the presence of the octave sweeps caused a significant increase in time for participants to determine which note to play. Furthermore, when feeling the sweeps, some participants could only tell that they occurred and not in which direction they were moving, which sometimes led to confusion.

A few participants recommended that a different mechanism be used for indicating octaves. Following a suggestion, a new system was designed that can reduce the number of octave sweeps needed for communicating a melody. This design was inspired by Braille music notation, in which symbols are used to indicate specific octaves and only appear when the interval between the current and next note either is greater than a fifth or is a fourth or fifth but moving to a different octave, where each octave begins with a C. If no symbol occurs, musicians will play the note that is closest to the previous note. For example, if a melody is going from a C down to an A, no symbol will occur even though it is in a different octave, because the interval is less than a fourth. The musician will just read an A and will choose the lower A because it is closer than the higher A.

Based on this design, the haptic vest's new octave system consists of downward sweeps, either on the left or the right side of the front of the vest, which act as symbols for the two possible octaves that the interface can communicate. If the interval of the next note satisfies the constraints described earlier, then a sweep will occur to indicate whether the octave of the note is the top or bottom. This approach decreases the total number of required sweeps in a melody because if the notes are close together, there will not be a sweep each time the music moves into a different octave. Furthermore, this system fixes the problem of users

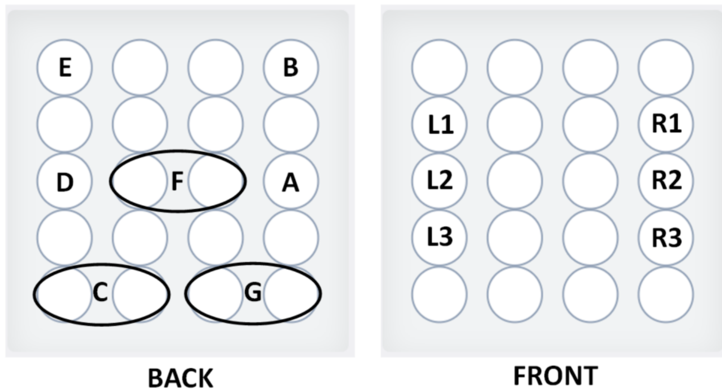

Figure 4.13: "Forwards N" vest design.

having difficulty determining which direction a sweep is moving, because a different set of actuators is used for each octave and the directions are always downwards. Although this design is certainly more complex than the previous system, it was chosen for the remainder of the studies given its potential benefits. The full vest mapping is shown in Figure 4.13

Because of the new system, the melodies had to be changed so that all of them have the same number of octave indications. Six 8-bar melodies were composed, as shown in Figure 4.14, each consisting of fourteen quarter notes and eight half notes with three octave indications.

4.4 Results and Analysis

With the new changes, a study was conducted with blind musicians. In total, 10 participants were recruited (5F, 5M, ages 22 to 72, average 50.2) in both Montreal and Toronto, through the Montreal Association for the Blind, Institut Nazareth et Louis-Braille, and the Canadian

Figure 4.14: Melodies for blind musicians.

Council for the Blind: Toronto Visionaries. Each participant in the main study had three sessions with the vest, providing quantitative data through their playing and qualitative data through questionnaires and verbal feedback.

4.4.1 Haptic Vest

Once the studies were complete, code was written to analyze the log files and calculate the average amount of time it took for participants to play the correct note after feeling the vibrations, as well as the average number of mistakes. In Step 3, notes had significant jumps so participants had more difficulty than in Step 4, where notes were closer together to sound more melodic. The participants' results for these steps can be seen in Figure 4.15.

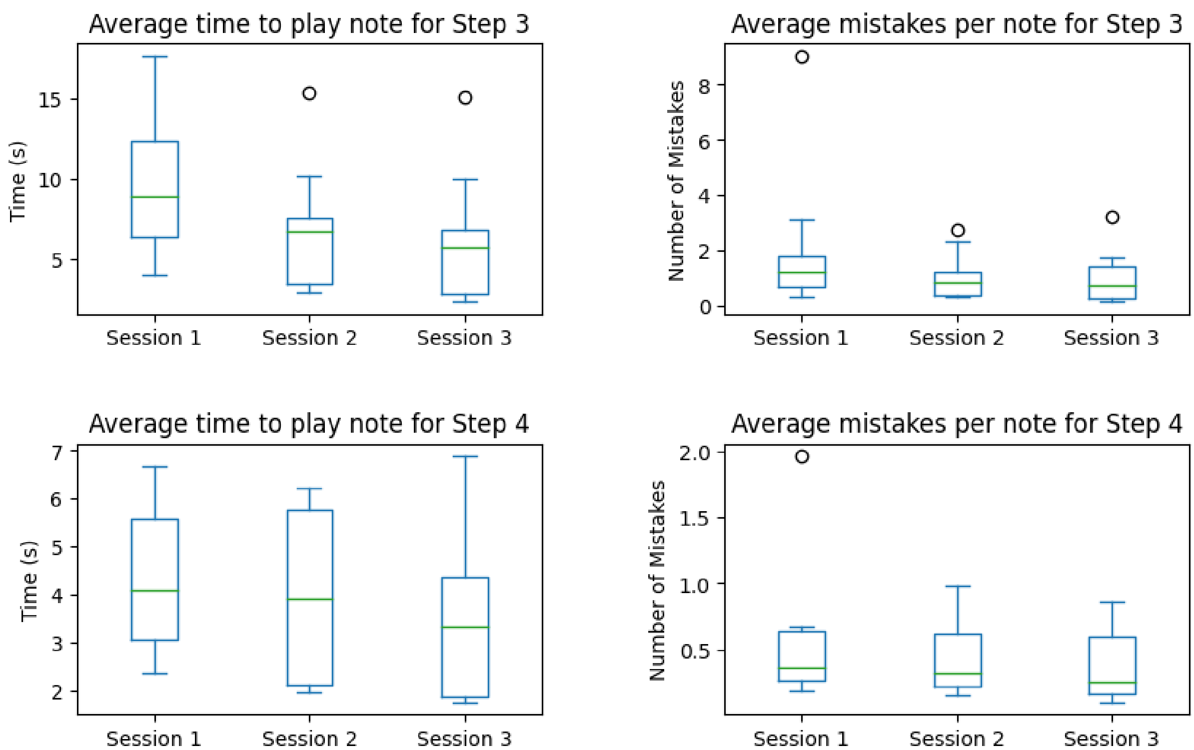

Figure 4.15: Results of Step 3 and Step 4.

During each experiment, participants learned a melody with the haptic vest and were assigned a value that takes into account the speed at which the melody played on the vest, the number of mistakes they made, and the amount of time it took for them to learn it. A weighted sum was taken of this value with the average time results of Steps 3 and 4 to give each participant a "haptic value" for each session, as shown in Table 4.1. A high value implies poor performance with the vest, while a low value signifies good results. These values can be visualized in Figure 4.16.

Table 4.1: Haptic values for each session.

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10
Session 1	19.5	8.7	34.1	23.8	15.6	30.4	64.1	32.0	15.5	7.9
Session 2	9.0	6.9	25.4	22.2	12.3	25.3	56.6	24.9	6.9	6.3
Session 3	7.1	5.7	22.9	26.2	9.8	12.6	50.6	17.8	5.8	5.3

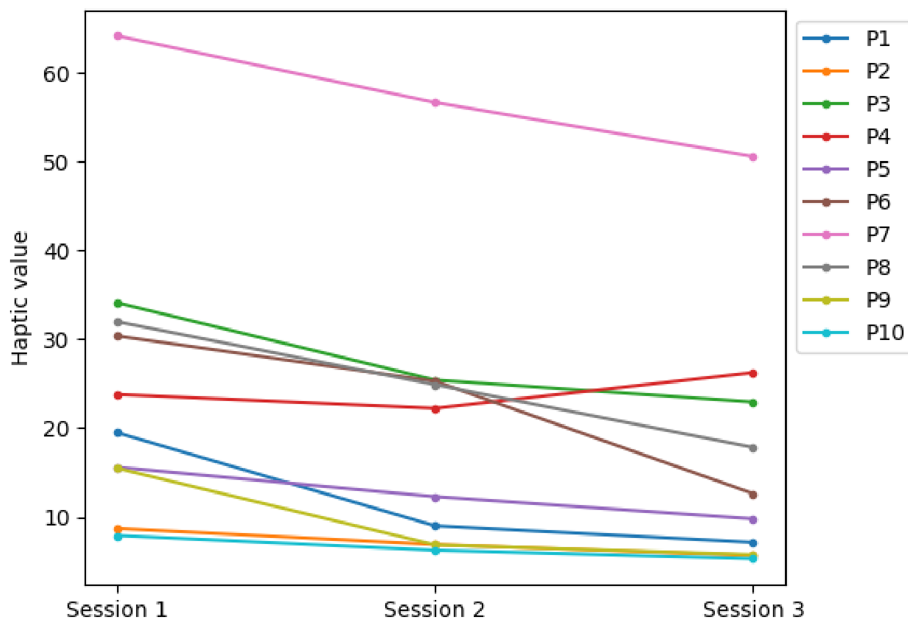

Figure 4.16: Haptic values of participants.

To determine whether participants improved while using the vest, the sessions were compared using paired two-sample t-tests, with the results shown in Table 4.2.

For both comparisons of the initial session with the second and third sessions, $p < 0.002$, showing that participants clearly improved after the first day. Improvement occurred after

Table 4.2: T-test results for haptic values.

	First Session	**Second Session**	**Third Session**
Mean	25.2	19.6	16.4
Variance	273.4	237.8	200.5
P(T<=t) two-tail with 1st session		0.0005	0.0017
P(T<=t) two-tail with 2nd session	0.0005		0.0508
P(T<=t) two-tail with 3rd session	0.0017	0.0508	

the second session as well, although with less confidence, as $p = 0.0508 > 0.05$. Participant 4 did worse in the third than in the second session, showing that not everyone improved each time. However, this result seems to be an outlier compared with the rest of the data.

To give each participant an overall haptic score, which is used in the next analysis, a weighted sum was taken of their haptic values from the three sessions, as shown in Equation 4.1.

$$haptic\_score = 0.2 * value1 + 0.35 * value2 + 0.45 * value3 \tag{4.1}$$

In this equation, the haptic value of the first session is given a low weighting because participants used this time to initially learn the system. In the second and third sessions, participants were more familiar with the vest and had much better scores, so it was chosen to weigh these haptic values higher.

4.4.2 Learning by Ear

Participants also received a score to indicate how well they learned the melody by ear, without the vest. This score was based on how long it took for them to learn the audio melody and how many mistakes were made in the final performance. The haptic and audio scores of each participant, along with their ages and years of piano lessons, are shown in Table 4.3. Visualizations of these data are shown in Figure 4.17. Participant 7 was not able to learn the melody by ear and did not provide her age, so her results are not included.

Table 4.3: Haptic score, audio score, age, and years of piano lessons.

	P1	**P2**	**P3**	**P4**	**P5**	**P6**	**P8**	**P9**	**P10**
Haptic Score	10.3	6.7	26.0	24.4	11.8	20.6	23.1	8.1	6.2
Audio Score	13.2	6.8	72.4	52.5	5.9	10.6	50.4	16.7	27.3
Age	51	49	59	55	60	72	56	28	22
Years of Lessons	10	15	2	6	14	6	10	3	2

Pearson's correlation test was run to determine whether there is a relationship between the audio scores and haptic scores of participants. This test provided a correlation coefficient of $r = 0.76$, suggesting a moderate correlation. This is what was expected, as people with high audio scores are very familiar with notes on the piano and may be better able to translate external information into musical notes. This test was also run to compare the haptic scores with the ages of participants, providing a value of $r = 0.67$, again suggesting a moderate correlation. The ability to learn tends to decrease with age, so this result is understandable.

Surprisingly, a comparison of participants' haptic scores with the number of years they

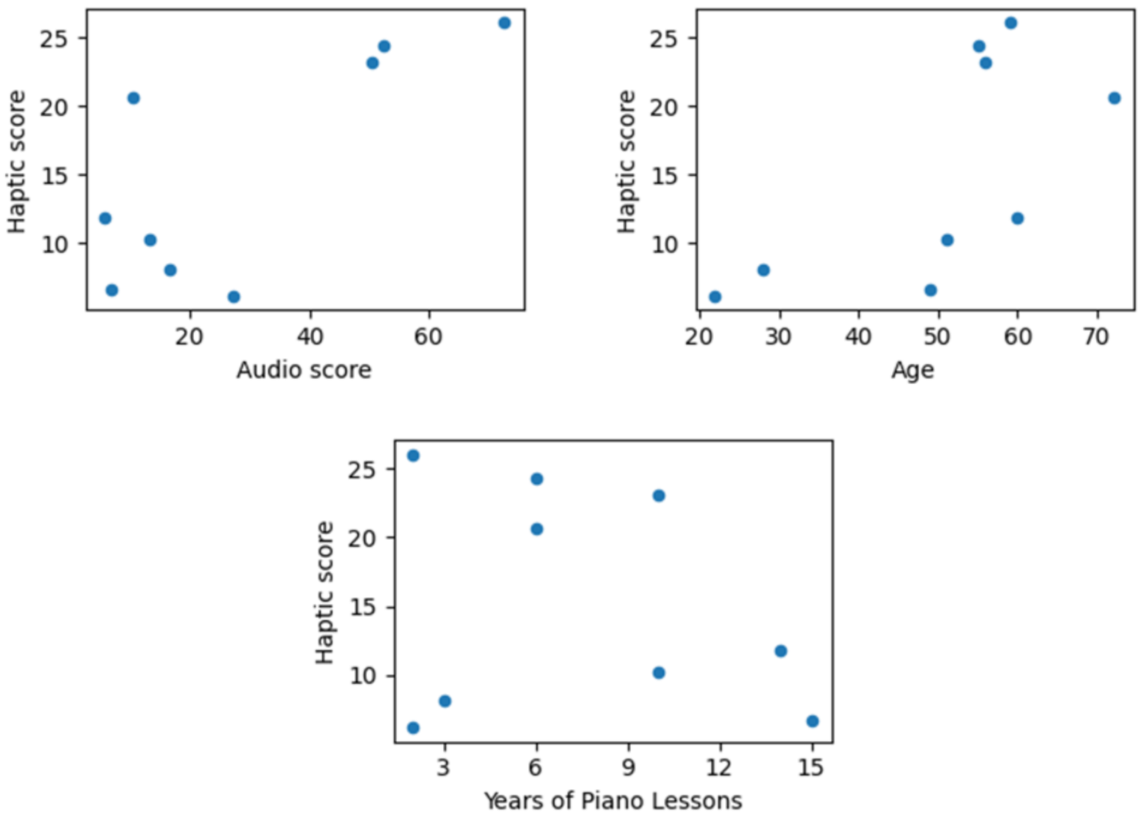

Figure 4.17: Comparison of haptic score with audio, age, and years of piano lessons.

had piano lessons gave a value of $r = -0.24$, a much weaker correlation than expected. However, Participants 9 and 10 may be outliers in this regard — without their data, $r = -0.85$, indicating a strong negative correlation. These two participants were also significantly younger than the other participants, suggesting that while years of piano lessons are certainly a major factor, age may be an even bigger factor. Nevertheless, $n = 9$ for these analyses, which may not be high enough to be confident about these conclusions.

4.4.3 Mapping

Finally, to determine the success of the “Forwards N” mapping strategy, the results of Steps 3 and 4 were split into the seven notes of the scale, to calculate the average time it took to play the note after feeling its vibration and the average number of mistakes made. These results are shown in Figure 4.18.

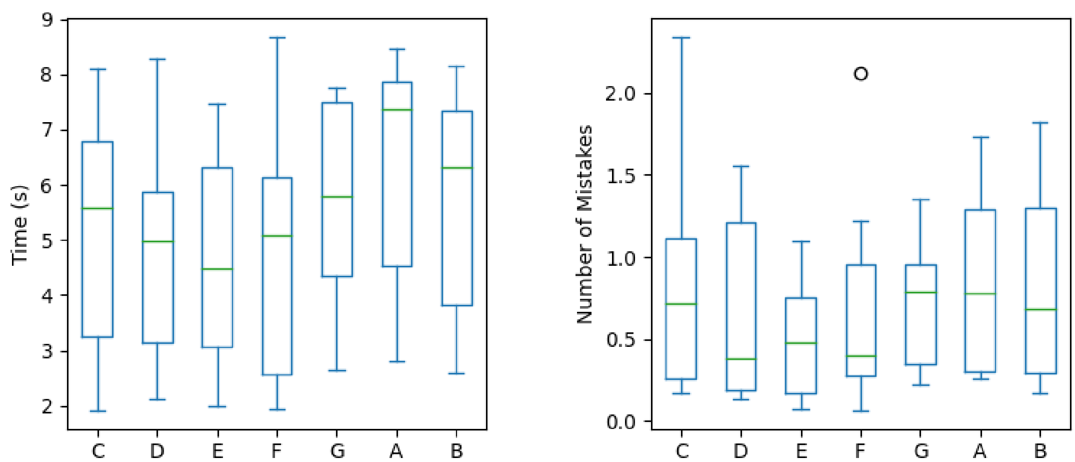

Figure 4.18: Results of individual notes.

Based on the results, there is generally not a huge variation between the notes, although G, A, and B did take slightly longer to play than C, D, and E. During the experiments, participants often counted up from C when determining which note to play, especially when learning, so this might be one of the explanations for this observation.

4.5 Discussion

4.5.1 Choice of Methodology

The methodology for the user studies in this project went through several iterations, particularly regarding the melody-learning portion of the experiments. Initially, the plan was for a user to learn a single melody with the haptic interface, but the audio portion of the procedure was introduced to compare haptics with traditional methods. The audio melody was originally a single chunk, such that participants would have to listen to it from start to finish without interruption and repeat that until they learn it. However, based on initial feedback from colleagues, it was separated into multiple chunks to improve learnability.

During the initial stages of this project, a conversation took place with Stelth Ng, a teacher of blind musicians and the founder of the Toronto Institute of Music for the Blind. He recommended that in the study, participants should learn a piece by both listening *and* using the haptic vest simultaneously. Unfortunately, initial testing with colleagues showed that when both senses were utilized, users tended to ignore the haptics and just focus on the audio, so this procedure was not included in the full methodology.

Nevertheless, pattern association might still have been occurring subconsciously, such that later learning of the haptic mapping would have been accelerated compared to a scenario in which haptics was not included initially. Research by Huang et al. demonstrates the value of passive haptic learning [17], a technique which could also be applied to this project. If this study were to be run again, participants could learn melodies in three ways: through haptics, audio, and both modalities.

4.5.2 Results

The results of the main user study show that participants were able to significantly improve their ability to use the vest over time. However, there was a very wide range of skill levels, as shown in Figure 4.15. By the final session, some participants were able to play the correct note in under two seconds after feeling the vibration, while others took over triple that amount. Musical expertise and age seemed to be influential regarding this, but there are likely other factors at play too. While not included in the pre-test questionnaire, education level and experience with digital technology might also have affected the learning ability of participants.

During the user study, participants still had trouble with the octave notation system, despite it being changed after the pilot study. Some users disliked the system's rules regarding intervals and found them complicated, although the musicians who were familiar with Braille music notation found them easier to understand. One participant suggested that sweeps occur simply when the octave changes, which was the strategy in the pilot study. Another user suggested that instead of sweeps, octaves could be communicated through double taps or long taps, which is a common mechanism in blind technology. Finally, a participant suggested that the octave could be indicated *after* the main vibration on the back instead of beforehand, as the sweeps were often too distracting.

Another common problem was that because participants generally took several seconds to play each note, they had difficulty determining the tempo and rhythm of the notes during the haptic melody portion of the experiment. A few participants suggested that a metronome be used, either before or during the melody playback, to help them understand the musical context. One participant mentioned that he prefers to know the duration of the notes before he plays them, which would require a different method of communicating note length instead of vibration duration.

The vest itself uses ERM actuators that were sometimes noisy, depending on how tight the vest's straps were. This was sometimes distracting from the location of the vibrations. The mapping of the vest, while generally received well, was subjected to some criticism. One participant recommended the original Backwards N design, rather than the Forwards N mapping. Other participants disliked the close proximity of the different vibrations, suggesting that they be more spread out. Designing a system that is appealing to everyone is a common challenge in HCI applications, and this vest was no exception.

4.5.3 Potential Improvements

Participants had several suggestions for ways in which the interface could be improved. First, it was suggested that other areas of the body be used to communicate octaves or more musical information such as accidentals and dynamics. Based on initial testing during this project, communicating additional details does present a greater cognitive challenge, but with sufficient practice, it could be feasible.

The vest's back vibrations were tuned to have the same intensities, although depending on how tight users wore the vest, certain notes could feel weaker than others. Participants suggested that the vibrations be set to the same strength, which would require an additional "calibration" procedure at the beginning of each session.

Finally, several participants said that if this product were to be marketed and distributed, there should be an easy method to translate visual sheet music into vibrations. In this study, the actuator commands were hard-coded for each melody, but in an ideal scenario, users would be able to take a picture of sheet music and receive the vibrations. This would require a computer vision program to detect and translate the notes.

4.5.4 Comparison with Listening

To judge whether participants were better able to learn a melody through the haptic vest or through the traditional method of listening, their haptic and audio scores could be examined, although the formulas for these values are fairly subjective so it is difficult to compare them head-to-head.

One possible way to compare the haptic interface to listening is to simply find the amount of time it took participants to learn the melodies. However, the speed at which the notes are played on the vest is an important factor. If set very slow, participants would be able to easily use the haptic interface to play the melody. The skill level of participants is demonstrated by how fast a melody they can "keep up with" when it is being applied to the vest.

Ultimately, learning a melody through haptics and audio are two different strategies. Participant 10, who had the best haptic score by the final session, put it this way: "Both methods offer a different way of learning music: listening requires a listening phase with playback, whereas we can follow the vest in real-time as it tells us the notes. I found the vest to be an innovative and effective way to learn music."

In the post-session questionnaires, most participants said they prefer to learn a melody by listening because they have much more experience with this method, although many found the vest to be promising. Participants were also asked if they would have used the vest when learning music had they been given the chance. Two said yes, five said they were not sure, and three said probably not. The participants who said yes enjoyed the vest for its ability to stay free from the audio channel. This allowed them to feel more connected to the music they were making, and for those with hearing difficulties, it was even more helpful. The participants who said "probably not" disliked the amount of time and effort it took to read and play notes. Additionally, the vest's inability to communicate multiple notes simultaneously greatly hindered how complex the musical score could be.

Chapter 5

Discussion

5.1 Themes

The chapters in this thesis describe two ways in which haptics can be used to enhance or assist musical performance. Although these projects are quite different and do not share the same audiences, a few overarching themes guide both of them, in addition to the topics of haptics and music.

These projects both apply the principles of user-centered design to the development of haptic musical interfaces. For both the Hapstrument and the vest, multiple rounds of user testing were conducted and the devices went through several iterations based on the user feedback. For the haptic vest, blind musicians were consulted from the beginning of the project regarding its feasibility and the most effective way for users to learn it.

The hardware in both of these projects is affordable compared to other available alternatives. The Haply 2diy and bHaptics TactSuit X40 both cost USD 500, while force-feedback devices and electronic assistive technologies such as tactile displays and specialized navigation aids usually reach into the thousands of dollars. By using low-cost products, this thesis shows that haptic musical interfaces are not just accessible to the very wealthy.

5.2 Potential Limitations and Future Research

5.2.1 Hapstrument

The Hapstrument, although starting off as just a course project, led to an investigation of the techniques that can be used to create a digital musical instrument using low-cost force-feedback devices. Through an iterative design process, several interfaces were explored and refined to create a fully functional instrument that, to most participants, was musically expressive.

That being said, the Hapstrument was not without its limitations. The pitch selection interface was generally the weakest part of the instrument, as it is difficult to quickly select distinct pitches with just a single end effector. Furthermore, the Haply 2diy is unable to detect the normal force of the handle against the surface, which limits possible features that could be implemented.

Moving forward, there are several ways to improve the Hapstrument. For the pitch selection interface, participants suggested adding a zone in which no pitch is selected, or making it easier to access non-scale notes such as accidentals. Additionally, making each of the discrete segments the same size would allow users to move between notes more easily.

Another potential improvement is to simply replace the pitch selection interface with different hardware, such as a simple MIDI keyboard. Almost every traditional musical instrument uses multiple fingers to select notes because fingers can be easily controlled independently and effectively. Therefore, the Haply 2diy may simply not be the correct physical interface for pitch selection.

If the left-hand Haply 2diy were replaced by a MIDI keyboard, it would not impact the haptic sensations of bowing and plucking, which are the main contributors to musical expression in the Hapstrument. However, it would mean sacrificing the ability to move between discrete and continuous notes as well as the ability to produce vibrato, both of which also help the user to express themselves. An alternative choice could be a Roli Seaboard [49] or Haken Continuum [47], which still allow for continuous pitch selection and vibrato.

Haply Robotics has released a new version of the 2diy, which eliminates much of the internal friction of the device by removing the bottom surface and simply suspending the end effector with the linkages. By using this upgraded version, the rattling issue would be resolved and the user experience would be greatly improved.

Finally, more features could be developed for the Hapstrument. A button or forcesensitive resistor could be installed on the device's handle as another input, or more sound parameters could be included, perhaps by introducing multiple simulated strings of varying thicknesses and timbres.

The user study in this project was able to give valuable insight into the effectiveness of the Hapstrument, although its biggest limitation was its small participant population. Future studies with larger groups may reveal further insights, especially if an inductive thematic analysis is conducted.

5.2.2 Haptic Interface

This thesis also explored how a haptic interface may be used to communicate musical information to blind musicians. Through several iterations, a prototype was designed and tested with the visually impaired community. While this interface was positively received, it also faced some criticism that highlighted a number of challenges and limitations.

The interface's strategy of communicating individual notes can work for simple and slow melodies but would not be possible for very fast or polyphonic melodies. Using a piano keyboard was helpful for the studies but ultimately, monophonic instruments such as trumpet or violin would be better suited for the interface.

Another limitation was with the vest itself. Although bHaptics advertises the vest as one-size-fits-all, the vest was a bit too small for some participants, and if it were to be used with children, it would likely be too big. Therefore, if the interface were to be made publicly available and advertised to the blind community, vests of different sizes should be developed and made accessible.

The user study of the haptic vest, similar to that of the Hapstrument, was limited by its small participant population. However, because the blind music community is a fairly small subset of the general population, a sample size of ten was considered acceptable. Future studies with larger groups could identify additional findings, especially through a thematic analysis.

Moving forward, there are two main paths for this research. First, the interface could be geared toward novice blind musicians who are first learning their instruments. Because beginners do not play fast or complex melodies, the vest would be capable of conveying musical information at a suitable pace. Several participants, in both the pilot and main studies, mentioned that they appreciate learning through tactile feedback, so the haptic vest could make music education more enjoyable.

Second, the interface could be redesigned to instead communicate chords, which would include the main note (C, C#, ...) and quality (major, minor, dominant 7th, ...). When musicians are playing in an ensemble, particularly with jazz or contemporary music, the individual notes often do not matter, but only whether they fit into the context of the current chord. A foot pedal could be used to trigger the next vibrations so that a musician could feel them at the desired pace. A chord-based system would require a more complex mapping than what was used in this study. One possible solution could be to divide the front of the vest into regions, similar to the approach used on the back of the vest. These regions could be assigned to different chord qualities and accidentals to provide the necessary additional functionality.

Overall, this project was an introductory exploration into whether a haptic interface can effectively communicate specific musical notes. Based on the results and feedback, it has considerable promise for the future. One participant said, “Learning through the vest is totally amazing, and I love this way of learning music. I seriously can’t wait to own one of these vests, can you just imagine how it will shape the way music is learned especially sight reading?” One participant had concerns that the cost of the interface would prove to be a hindrance, although upon being told that the bHaptics vest is only $500, he said that this is “quite affordable compared to other products dedicated to visual impairment.” This interface has significant potential and it will be exciting to see future developments that contribute further to the field of musical assistive technology.

5.3 Conclusion

The first objective of this thesis was to investigate the effect of force feedback on a novel low-cost digital musical instrument and determine its contribution to musical expression. Through the design, creation, and evaluation of the Hapstrument, a DMI was created in which the simulated sensations of bowing and plucking contribute to its expressivity. This research may inspire further exploration into the effectiveness of low-cost force-feedback devices for musical interfaces, and the various techniques they can simulate. This could lead to an increase in affordable haptic digital musical instruments.

The second objective of this thesis was to explore whether a haptic interface may be used to communicate musical notation to blind musicians through vibrations so that they can "read" music and play it in real-time. After several iterations of interfaces and mappings, a design was chosen and a user study was conducted with blind musicians, receiving positive feedback for future developments. This project builds upon previous research to provide another example of how haptics may be used in assistive technology by communicating information without overloading the audiovisual channels.

Moving forward, the field of haptics will keep expanding through future research with more powerful and advanced hardware, leading to further developments in musical haptic interaction. This will yield higher-fidelity digital instruments and additional methods of communicating notation, such that musical experiences will continue to grow in learnability and expressivity.

Bibliography

[1] A. Müller, F. Hemmert, G. Wintergerst, and R. Jagodzinski, "Reflective Haptics: Resistive Force Feedback for Musical Performances with Stylus-Controlled Instruments," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Sydney, Australia, Jun. 2010, pp. 477–478.

[2] M. Giordano, S. Sinclair, and M. M. Wanderley, "Bowing a vibration-enhanced force feedback device," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Ann Arbor, Michigan, USA, May 2012.

[3] B. Verplank, M. Gurevich, and M. Mathews, "THE PLANK: Designing a Simple Haptic Controller," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Dublin, Ireland, May 2002, pp. 177–180.

[4] E. Sheffield, E. Berdahl, and A. Pfalz, "The Haptic Capstans: Rotational Force Feedback for Music using a FireFader Derivative Device," in *Proceedings of the*

International Conference on New Interfaces for Musical Expression, Brisbane, Australia, Jul. 2016, pp. 1–2.

[5] D. M. Howard, S. Rimell, and A. D. Hunt, "Force Feedback Gesture Controlled Physical Modelling Synthesis," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Montreal, Quebec, Canada, May 2003, pp. 95–98.

[6] H.-C. Steiner, "Stickmusic: Using haptic feedback with a phase vocoder," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Hamamatsu, Japan, Jun. 2004, pp. 203–204.

[7] E. Berdahl and A. Kontogeorgakopoulos, "The FireFader: Simple, Open-Source, and Reconfigurable Haptic Force Feedback for Musicians," *Computer Music Journal*, vol. 37, no. 1, pp. 23–34, Mar. 2013.

[8] E. Berdahl and D. Huber, "The Haptic Hand," in *Proceedings of the International Conference on New Interfaces for Musical Expression*, Baton Rouge, Louisiana, USA, May 2015, pp. 303–306.

[9] E. Berdahl, P. Vasil, and A. Pfalz, "Automatic Visualization and Graphical Editing of Virtual Modeling Networks for the Open-Source Synth-A-Modeler Compiler," in *Haptics: Perception, Devices, Control, and Applications*, London, UK, Jul. 2016, pp. 490–500.

[10] E. Berdahl and Z. Berkowitz, "Portable Modeling of Virtual Physics for Audio and Haptic Interaction Design," in *Proceedings of the 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, New Orleans, Louisiana, USA, Mar. 2017, pp. 6568–6569.

[11] L. Turchet, "Musical Haptics for the Listener: Reflections of a Design Practitioner," *Multimodality & Society*, vol. 2, no. 4, pp. 379–385, 2022.

[12] L. Turchet, T. West, and M. Wanderley, "Touching the Audience: Musical Haptic Wearables for Augmented and Participatory Live Music Performances," *Personal and Ubiquitous Computing*, vol. 25, pp. 749–769, Aug. 2021.

[13] A. Mazzoni and N. Bryan-Kinns, "Mood Glove: A Haptic Wearable Prototype System to Enhance Mood Music in Film," *Entertainment Computing*, vol. 17, pp. 9–17, Jun. 2016.

[14] M. S. O'Modhrain, "Playing by Feel: Incorporating Haptic Feedback into Computer-Based musical Instruments," Ph.D. dissertation, Stanford University, Nov. 2000.

[15] I. Hwang, H. Son, and J. R. Kim, "AirPiano: Enhancing Music Playing Experience in Virtual Reality with Mid-Air Haptic Feedback," in *2017 IEEE World Haptics Conference (WHC)*, Jun. 2017, pp. 213–218.

[16] A. Bouwer, S. Holland, and M. Dalgleish, "The Haptic Bracelets: Learning Multi-Limb Rhythm Skills from Haptic Stimuli While Reading," in *Music and Human-Computer Interaction*. London: Springer, Feb. 2013, pp. 101–122.

[17] K. Huang, E. Y.-L. Do, and T. Starner, "PianoTouch: A Wearable Haptic Piano Instruction System for Passive Learning of Piano Skills," in *Proceedings of the 2008 12th IEEE International Symposium on Wearable Computers*, Washington, D.C., USA, Sep. 2008, pp. 41–44.

[18] T. Grosshauser and T. Hermann, "Augmented Haptics – An Interactive Feedback System for Musicians," in *Haptic and Audio Interaction Design*. Dresden, Germany: Springer, Sep. 2009, pp. 100–108.

[19] F. A. Geldard, "Adventures in Tactile Literacy," *American Psychologist*, vol. 12, no. 3, p. 115, Mar. 1957.

[20] W. M. Rabinowitz, A. J. M. Houtsma, N. I. Durlach, and L. Delhorne, "Multidimensional Tactile Displays: Identification of Vibratory Intensity, Frequency, and Contactor Area." *The Journal of the Acoustical Society of America*, vol. 82, no. 4, pp. 1243–1252, Oct. 1987.

[21] G. Luzhnica and E. Veas, "Optimising Encoding for Vibrotactile Skin Reading," in *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, ser.

CHI '19. Glasgow, Scotland, UK: Association for Computing Machinery, May 2019, pp. 1–14.

[22] S. Brewster and L. M. Brown, "Tactons: Structured Tactile Messages for Non-Visual Information Display," in *Proceedings of the Fifth Conference on Australasian User Interface - Volume 28*, ser. AUIC '04. Dunedin, New Zealand: Australian Computer Society, Inc., Jan. 2004, pp. 15–23.

[23] L. A. Jones, B. Lockyer, and E. Piateski, "Tactile Display and Vibrotactile Pattern Recognition on the Torso," *Advanced Robotics*, vol. 20, no. 12, pp. 1359–1374, Jan. 2006.

[24] M. F. de Vargas, A. Weill-Duflos, and J. R. Cooperstock, "Haptic Speech Communication Using Stimuli Evocative of Phoneme Production," in *2019 IEEE World Haptics Conference (WHC)*, Tokyo, Japan, Jul. 2019, pp. 610–615.

[25] Y. Golan, B. Serota, A. Shapiro, O. Shriki, and I. Nisky, "A Vibrotactile Vest for Remote Human-Dog Communication," in *2019 IEEE World Haptics Conference (WHC)*, Tokyo, Japan, Jul. 2019, pp. 556–561.

[26] L. Turchet and M. Barthet, "Co-Design of Musical Haptic Wearables for Electronic Music Performer's Communication," *IEEE Transactions on Human-Machine Systems*, vol. 49, no. 2, pp. 183–193, Dec. 2019.

[27] L. Turchet, D. Baker, and T. Stockman, "Musical Haptic Wearables for Synchronisation of Visually-Impaired Performers: A Co-Design Approach," in *ACM International Conference on Interactive Media Experiences*, ser. IMX '21. Virtual Event, USA: Association for Computing Machinery, Jun. 2021, pp. 20—-27.

[28] M. Giordano, J. Sullivan, and M. M. Wanderley, "Design of Vibrotactile Feedback and Stimulation for Music Performance," in *Musical Haptics*. Cham, Switzerland: Springer International Publishing, May 2018, pp. 193–214.

[29] T. J. West, A. Bachmayer, S. Bhagwati, J. Berzowska, and M. M. Wanderley, "The Design of the Body:Suit:Score, a Full-Body Vibrotactile Musical Score," in *Human Interface and the Management of Information. Information in Intelligent Systems.* Cham, Switzerland: Springer International Publishing, May 2019, pp. 70–89.

[30] T. McDaniel, S. Krishna, V. Balasubramanian, D. Colbry, and S. Panchanathan, "Using a Haptic Belt to Convey Non-Verbal Communication Cues during Social Interactions to Individuals who are Blind," in *2008 IEEE International Workshop on Haptic Audio visual Environments and Games*, Lecco, Italy, Nov. 2008, pp. 13–18.

[31] H. Nicolau, J. a. Guerreiro, T. Guerreiro, and L. Carriço, "UbiBraille: Designing and Evaluating a Vibrotactile Braille-Reading Device," in *Proceedings of the 15th International ACM SIGACCESS Conference on Computers and Accessibility*, ser.

ASSETS '13, no. 23. Bellevue, Washington, USA: Association for Computing Machinery, Oct. 2013.

[32] B. Duvernoy, S. Topp, and V. Hayward, "HaptiComm, a Haptic Communicator Device for Deafblind Communication," in *Haptic Interaction. AsiaHaptics 2018*, ser. Lecture Notes in Electrical Engineering. Singapore: Springer Singapore, May 2019, pp. 112–115.

[33] J. G. J. Vuijk, J. Gay, M. A. Plaisier, A. M. L. Kappers, and A. Theil, "PatRec: A Mobile Game for Learning Social Haptic Communication," in *Proceedings of the 23rd International ACM SIGACCESS Conference on Computers and Accessibility*, ser. ASSETS '21, no. 81. Virtual Event, USA: Association for Computing Machinery, Oct. 2021.

[34] S. D. Novich and D. M. Eagleman, "[D-79] A Vibrotactile Sensory Substitution Device for the Deaf and Profoundly Hearing Impaired," in *2014 IEEE Haptics Symposium (HAPTICS)*, Houston, Texas, USA, Feb. 2014, pp. 1–1.

[35] S. Nanayakkara, E. Taylor, L. Wyse, and S. H. Ong, "An Enhanced Musical Experience for the Deaf: Design and Evaluation of a Music Display and a Haptic Chair," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '09. Boston, MA, USA: Association for Computing Machinery, Apr. 2009, pp. 337–346.

[36] U. Trivedi, R. Alqasemi, and R. Dubey, “Wearable Musical Haptic Sleeves for People with Hearing Impairment,” in *Proceedings of the 12th ACM International Conference on Pervasive Technologies Related to Assistive Environments*, ser. PETRA ’19. Rhodes, Greece: Association for Computing Machinery, Jun. 2019, pp. 146–151.

[37] A. Sharp, B. A. Bacon, and F. Champoux, “Enhanced Tactile Identification of Musical Emotion in the Deaf,” *Experimental Brain Research*, vol. 238, no. 5, pp. 1229–1236, May 2020.

[38] B. Remache-Vinueza, A. Trujillo-León, M. Zapata, F. Sarmiento-Ortiz, and F. Vidal-Verdú, “Audio-Tactile Rendering: A Review on Technology and Methods to Convey Musical Information through the Sense of Touch,” *Sensors*, vol. 21, no. 19, p. 6575, Sep. 2021.

[39] B. Petry, T. Illandara, D. S. Elvitigala, and S. Nanayakkara, “Supporting Rhythm Activities of Deaf Children Using Music-Sensory-Substitution Systems,” in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. CHI ’18. Montreal, Quebec, Canada: Association for Computing Machinery, Apr. 2018, pp. 1–10.

[40] P. Tranchant, M. M. Shiell, M. Giordano, A. Nadeau, I. Peretz, and R. J. Zatorre, “Feeling the Beat: Bouncing Synchronization to Vibrotactile Music in Hearing and Early Deaf People,” *Frontiers in Neuroscience*, vol. 11, Sep. 2017.

[41] B. Krolick, *New International Manual of Braille Music Notation.* Braille Press Zurich, 1996.

[42] W. C. Payne, F. Ahmed, M. Zachor, M. Gardell, I. Huey, A. Hurst, and R. L. Dubois, “Empowering Blind Musicians to Compose and Notate Music with SoundCells,” in *Proceedings of the 24th International ACM SIGACCESS Conference on Computers and Accessibility*, ser. ASSETS ’22, no. 17. Athens, Greece: Association for Computing Machinery, Oct. 2022.

[43] J. A. Borges and D. Tomė, “Teaching Music to Blind Children: New Strategies for Teaching through Interactive Use of Musibraille Software,” *Procedia Computer Science*, vol. 27, pp. 19–27, 2014.

[44] B. Lussier-Dalpé, C. Houtekier, J. Duquette, M.-C. Wanet-Defalque, and W. Wittich, “The Challenge of Reading Music Notation for Pianists with Low Vision: An Exploratory Qualitative Study Using a Head-Mounted Display,” *Assistive Technology*, vol. 34, no. 1, pp. 2–10, 2022.

[45] K. E. MacLean, O. Schneider, A. Weill-Duflos, V. Levesque, P. Irani, and J. R. Cooperstock, “CanHap 501: Learning Haptic UX Design in Remote Teams,” in *2021 IEEE World Haptics Conference (WHC)*, Montreal, Quebec, Canada, Jul. 2021, pp. 348–348.

[46] A. Weill-Duflos, N. Ong, F. Desourdy, B. Delbos, S. Ding, and C. Gallacher, "Haply 2diy: An Accessible Haptic Platform Suitable for Remote Learning," in *Proceedings of the 2021 International Conference on Multimodal Interaction*, Montreal, Quebec, Canada, Oct. 2021, pp. 839–840.

[47] L. Haken and E. Eagan, "Playing the Continuum Fingerboard: New Performance Techniques for a Continuous Instrument." *The Journal of the Acoustical Society of America*, vol. 128, no. 4, pp. 2345–2345, Oct. 2010.

[48] G. W. Young and K. Crowley, "The Design of Tangible Digital Musical Instruments," in *Proceedings of MusTWork16*, Dublin, Ireland, Jun. 2016.

[49] R. O. Lamb, "The Seaboard: Discreteness and Continuity in Musical Interface Design," Ph.D. dissertation, Royal College of Art, Jun. 2014.