Development and Implementation of a Vibrotactile Click Track to Assist Contemporary Music Conducting

Patrick Ignoto

Department of Music Technology
Input Devices and Music Interaction Laboratory
McGill University
Montreal, Canada

June 2018

A thesis submitted to McGill University in partial fulfillment of the requirements for the degree of Master of Arts.

Abstract

This thesis presents the development and implementation of a novel vibrotactile click track system created for use with the McGill Contemporary Music Ensemble (CME), designed in close collaboration with the CME conductor, Prof. Guillaume Bourgogne. Using common HCI techniques, requirements were gathered to better understand how the system should work. These requirements described the need for a system that is different from previous works in haptic click tracks, because it uses a continuous waveform to present the tempo, rather than a discrete series of buzzes. In addition, the system required using existing audio click tracks of pieces in the contemporary repertoire, and presenting them as haptic information instead. Using an iterative design process, prototypes of this system were created and Prof. Bourgogne's comments were used to refine the system to meet all his requirements. The system was designed using the Vibropixels, a wireless, modular tactile display system developed at the IDMIL. Software and firmware were written to interact with the Vibropixels, display the haptic pulses in the way envisioned by the user, and convert existing audio click tracks to this system. The system was evaluated to determine if the requirements were met by measuring vibrations with an oscilloscope, measuring the computational and wireless latency, and by interviewing performers using the system. Results show that the system met the requirements collected, and preparations are underway to use it in a concert setting.

Résumé

Cette thèse présente le développement et l'implémentation d'un système original pour un métronome vibro-tactile, crée pour être utilisée par l'ensemble de musique contemporaine de McGill, dont la conception a été faite en étroite collaboration avec le chef d'orchestre de l'ensemble, le professeur Guillaume Bourgogne. En utilisant des techniques d'interaction homme-machine, les exigences ont été obtenues pour mieux comprendre comment le système devrait fonctionner. Ces exigences ont identifié un système qui est différent des travaux précédents, parce qu'il utilise une forme d'onde continue pour présenter le tempo, au lieu d'une série discrète de vibrations. De plus, le système a besoin de prendre les pistes-métronomes en audio existante pour les pièces du répertoire contemporain et les présenter comme des pistes-métronomes tactiles. En utilisant un processus de conception itératif, des prototypes de ce système ont été créés et les commentaires du professeur Bourgogne ont été utilisés pour améliorer le système et répondre à ses exigences. Le système est conçu avec la plateforme « Vibropixels », un système tactile sans-fil et modulaire créé au laboratoire IDMIL. Des logiciels ont été développés pour intéragir avec les « Vibropixels », afficher les vibrations de façon demandé par le chef d'orchestre et convertir les pistes-métronomes audio à ce système. Le système a été évalue pour déterminer si ces exigences ont été atteintes - en mesurant les vibrations avec un oscilloscope, ainsi que la latence informatique et les délais des transmissions sans-fil, mais aussi avec les commentaires des artistes qui ont utilisé le système. Les résultats montrent que le système a satisfait les exigences requises et des préparations sont en cours pour l'utiliser en concert.

Acknowledgments

Many thanks to my supervisor, Marcelo M. Wanderley, and Guillaume Bourgogne for proposing this project and their guidance and feedback throughout the design process. Thanks to Ian Hattwick for his help with getting me started with the Vibropixels and the initial stages of this project. For their feedback during the evaluation of the design, a big thank you to Suzu Enns and Takuto Fukuda.

A special thanks to my parents for putting up with me, to the Den Hertogs for their help with making nice photos for this document, and to Dirk Dubois for pointing me in this direction and for giving me the motivation to finish this project. I could not have done this without you all.

The research conducted in this thesis was approved by McGill University's Research Ethics Board. REB file number 399-0317.

Contents

List of Figures

List of Tables

List of Acronyms

DMI	Digital Music Instrument
GUI	Graphical User Interface
HCI	Human-Computer Interaction
UCD	User-Centered Design
USB	Universal Serial Bus
IDMIL	Input Devices and Music Interaction Laboratory
CME	Contemporary Music Ensemble

Chapter 1

Introduction

In contemporary music, conductors many times perform pieces of mixed music, requiring synchronizing an orchestra to some pre-recorded piece of music played through speakers in the concert hall. In order to achieve this task, conductors typically use an audio click track as a metronome, played through an ear-piece. A click track is a common tool in assisting a performer with keeping tempo, and is usually used in studio recordings to synchronize musicians to a prominent and consistent rhythmic pulse to simplify later editing (Robertson & Plumbley, 2013). The click track, which the conductor typically receives along with the score, is an audio file containing a series of click sounds played at intervals consistent with the tempi of the pre-recorded audio. This assists the conductor in keeping the tempo of the piece and conducting the orchestra along with the audio.

However, the click track in a live music setting can be quite intrusive to the performer. As the conductor must listen to the orchestra to properly conduct the piece of music; having an ear blocked off to receive information for the piece can interfere with their task or change their overall perception of the music. Using a different modality to transmit this information can be much less intrusive to a conductor in such a setting.

For nearly 60 years, the tactile modality has been explored by many researchers in human-computer interaction (HCI) as a means of sending information (Gallace et al., 2007). This is known as *haptics*. Haptics concerns designing for the sense of touch. In particular, vibrotactile haptics concerns using vibrations as a means of transmitting information tactilely. Devices called vibrotactile displays are used in transmitting information via vibrations (Choi & Kuchenbecker, 2013). Tactons (portmanteau for tactile icons) "are structured, abstract messages that can be used to communicate complex concepts to users non-visually" and have been transmitted with vibrotactile displays by researchers to convey information to users (Brewster & Brown, 2004).

Using such technology and applications, the tactile modality can be used to provide the same tempo information as a click track. This would liberate the auditory modality for the conductor and allow him to fully hear the music being performed by the orchestra during the concert, while still providing the information necessary for him to keep the orchestra in synchrony with the prerecorded audio.

1.1 Project Overview

This document presents a project undertaken to develop a vibrotactile alternative to a click track for contemporary music conductors in live and rehersal settings. It was done in close collaboration with the conductor of the McGill Contemporary Music Ensemble, Professor Guillaume Bourgogne. A series of qualitative interviews were performed with Prof. Bourgogne to gather the requirements for such a device and to design it based on his unique needs. The interviews aimed to highlight what such a system should do, how the vibration patterns should feel, as well as general usability requirements, for instance, where to attach the vibrotactile display. These requirements were used to inform, and further interviews were used to refine, the design of the system. The chosen hardware platform was the Vibropixels, a wireless, scalable vibrotactile display system (Hattwick et al., 2017). The resulting design modified existing firmware and software to meet the requirements gathered. Laboratory tests were performed to determine whether the design met the requirements, and the tests found that these requirements were met. Overall, reception to the design by Prof. Bourgogne was positive and considered a successful haptic alternative to click tracks.

The rest of this document will describe in detail the design process used, the implementation considerations that were necessary to design the system, and the methods used to validate the resulting system and determine its efficiency in transmitting tempo information.

1.2 Thesis Overview

The rest of this document is structured as follows: **Chapter 2** will present previous studies in this field. **Chapter 3** will present in detail the design requirements gathered from initial interviews with Professor Guillaume Bourgogne. **Chapter 4** will present the design implementation considerations in detail. **Chapter 5** will discuss the validations performed on the design and discuss its overall validity. **Chapter 6** will conclude this document and discuss future work.

Chapter 2

Previous Works

This chapter provides a review of existing research and studies that give credence in designing the project discussed in the previous chapter. In particular, it reviews some general concepts regarding the fields of HCI and haptics and their relevance to computer music and music technologies. This provides relevant background to the project and its design for readers not familiar with these fields. The section on haptics presents some of the uses vibrotactile applications have had on the field of music technology and computer music, such as using them for giving feedback to performers using a DMI or notifying performers in a live setting. Previous studies on the use and perception of vibrotactile notifications in lieu of an audio metronome will be presented. Finally, the technology of vibrotactile displays is introduced and two vibrotactile display platforms that were considered for use in this project are described.

2.1 HCI

The field of Human-Computer Interaction (HCI) is the field of research focused on the design of computer systems that can be interacted with in a user-friendly and efficient way (Preece et al., 1994). This can be achieved through interaction design, which focuses on how the user will interact with the product and not just on the product itself (Dix et al., 2003). Contemporary musicians have embraced electroacoustic music into the repertoire and using music technologies on stage is prevalent for these performers (Kimura, 1995), therefore designers of music technologies need to focus on the interaction design of new technologies for these musicians.

A fundamental way to perform interaction design is through user-centered design (UCD), a set of methods used to design systems that focus on how the user will actually use the system in the intended context (Preece et al., 1994). This means using methods such as user requirements analysis to determine the needs of users for the designed system or iterative design to iterate and evaluate multiple prototypes in order to achieve the correct requirements for users. In a survey of UCD practitioners by Vredenburg et al. (2002), it was found that the five most frequently used methods of performing UCD were iterative design, usability evaluation, task analysis, informal expert review, and field studies. In the same survey, the authors found that UCD practitioners believed that field studies, user requirements analysis, and iterative design were among the most important techniques for performing UCD. By using such techniques when designing systems for music technologies, we can develop products that better suit the needs of the musicians who would use such systems.

Once a system has been designed with the requirements found, it is important to evaluate that the system behaves as expected and meets the user's requirements. According to Dix et al. (2003), this can be done through expert analysis (which evaluates based on analytic or review methods) or through user participation (which evaluates based on experimental and observational methods). Ultimately, the authors believe there are no hard and fast rules in evaluation, and suggest choosing between many of the methods they outline based on the design factors.

2.2 Haptics

Haptics is the field of research concerning the sense of touch, typically used for both kinesthetic and tactile sensations (El Saddik, 2007). In the field of HCI, haptics concerns synthesizing tactile stimuli for various types of electronic interfaces (MacLean, 2008). In music technology, the synthesis of tactile stimuli is often used by researchers to notify or give feedback to musicians (Giordano & Wanderley, 2013).

2.2.1 Feedback for DMIs

Although acoustic musical instruments inherently generate haptic sensations due to the resonance and mechanical means in which music is generated, digital musical instruments (DMIs) do not inherently generate haptic sensations, because the gestural controller is mechanically decoupled from the sound synthesizer (Gillespie, 1999; Miranda & Wanderley, 2006). Therefore, incorporating haptic effects in the design of a DMI can be integral to help improve the feel of a DMI, making it play more like a traditional instrument (Bongers, 1997). An early study found that incorporating haptic feedback on DMIs could improve the performance of a trained musician unfamiliar with the instrument, finding that their playing accuracy for simple melodies was improved by an average of 23% (O'Modhrain & Chafe, 2000). The following are some examples of this use of haptic feedback to improve DMIs:

- The Modular feedback keyboard by Cadoz et al. (1990) is a keyboard controller that provides force feedback when striking the keys. Using custom-built motors, they were able to capture the displacement of the keyboard keys and produce a proportionate force. This allowed the displaying of force feedback to the player, similar to striking the keys of an acoustic piano.
- Chafe (1993) used voice coils in a electronic wind controller to simulate haptic feedback in the mouthpiece of traditional brass instruments. The author found this helped performers more accurately play a physical model synthesizer of a french horn.
- The VR/TX system by Rovan & Hayward (2000). This can be either a ring- or foot-mounted vibrotactile display that was used in conjunction with open-air gestural controllers. VR/TX provided vibrotactile feedback to gestural controllers without a physically manipulable controller. It was designed with a specially-made voice coil actuator. A Max/MSP environment controled seven parameters that affected the vibration sensation, including frequency and amplitude.
- Marshall & Wanderley (2006) presented two DMIs, the Vibroltar and the Vibroslide, which benefited from the use of haptic feedback. Both instruments used small speakers embedded in the instrument to generate the sound from the instrument and to provide feedback. The speakers caused the entire instrument to resonate, much like a traditional instrument, simulating the haptic feedback of acoustic instruments.

- Sinclair et al. (2009, 2011) investigated the use of force feedback devices for controlling virtual instruments. The authors sought to simulate the tactile feedback that naturally occurs when a player bows a stringed instrument. The authors used an Ergos TGR (a force feedback haptic device by ACROÉ in France) with a real bow as the end effector of the controller for this virtual instrument. They found bowing this virtual stringed instrument presented a reasonable yet rough representation of bowing a real stringed instrument.

2.2.2 Live Performer Notifications

Other than providing haptic feedback to DMIs, vibrotactile sensations are also used to provide notifications to performers in live music settings. This liberates the auditory and visual modalities of the performers while still providing them information crucial to their performance. Examples of this include:

- The Vibrobyte by Mcdonald et al. (2009) is a vibrotactile display that was used to send vibrotactile notifications from a composer to performers or between performers in a telematic performance. Using a custom built platform, they were able to send messages to cue "the performers with dynamically varying intensity, rhythm and instrumentation combinations".

- Michailidis & Bullock (2011) presents a study in which the authors look at the use of vibrotactile cues in live electronic music performance and determine whether they can help improve the performance. Using a modified prototype of a previous system (Michailidis & Berweck, 2011), they simulated haptic cues when triggering an electronic controller. The authors found that haptic feedback helped the performers control the live electronics.

- The NeVIS system by Hayes & Michalakos (2012) is a networked vibrotactile communication system that can be used by improvisational performers for notification and interaction. Using a custom haptic device with three coin-cell actuators, the system sends cues to a group of improvisational musicians. These cues include denoting sections in a piece, tempo information, and communication between performers.

- Schumacher et al. (2013) incorporated vibrotactile notifications into the CIRMMT Live Electronics Framework (CLEF), a Max-based modular framework for use with live performers in mixed music pieces. Using a hardware prototype with two coin-cell actuators, a module for tactile feedback was added to the CLEF software by the developers. This vibrotactile notification module allowed for two types of modalities, an "individual mode" where each of the actuators were triggered and controlled independently to allow for discrete pulses or continuous motions between the actuators, and a "balanced mode" which is used to display relative values by changing the intensity ratios between the two actuators. The authors evaluated the use of tactile notifications by running tests with a performer. One test evaluated a haptic click-track modality, where discrete haptic pulses were used to convey tempo information. Another test notified performers of the location of a sound source in a sonic spatialization system, which was an irregular and non-deterministic notification. In both cases, the performer said that the tactile notifications were very effective and unobtrusive. A later study by Frid et al. (2014) quantified perceptual information about tactile modalities using the CLEF system. From the data generated by the experiments, the authors provide suggestions for designing tactons in live musical settings; for example, use a duty cycle scale higher than 0.2 and pulses with haptic inter-onset times less than 200ms cannot be perceived as discrete.

2.2.3 Artistic works with Haptics

Finally, haptics has been used in conjunction with music in several artistic applications. The following show examples of how haptics used in conjunction with sound and lighting can produce a unique sensory experience:

- Cutaneous Grooves was a series of performances by Gunther & O'Modhrain (2003) that performed a piece of music for both the sense of touch and hearing. Users wore a custom-built suit that provided vibration patterns choreographed for the music being played. The authors propose a framework of "tactile composition", with which to use when developing compositions with the tactile sense like Cutaneous Grooves.

- The *Ilinx* garment is a custom-built haptic suit that was designed for a series of multi-sensory art exhibitions (Giordano et al., 2015; Lamontagne et al., 2015). Vibrotactile actuators were sown into a garment worn on the arms, legs, and torso by participants visiting the installation. Participants spent time in a dark room and felt different vibrotactile sensations along their body, as well as choreographed sounds and lights. This arguably provided a unique sensory experience to visitors.

- The *Haptic Fields* art installation used the Vibropixels to perform a multi-sensory art installation much like *Ilinx* (Hattwick et al., 2017). Modular tactile displays, described in 2.4.2, were strapped to participants in lieu of a custom-built garment. The tactile displays also generated light patterns visible to participants.

2.3 Human perception of tactile metronome signals

The sense of touch is governed by four mechanorecptors in the skin providing four channels of information in sensing touch (Bolanowski et al., 1988). Verrillo (1992) later provided a review characterizing human perception of vibrations. The author presents figures demonstrating vibrotactile detection thresholds and contours of equal sensation magnitude related to the sense of touch, gathered from measurement data performed by the author. These data are important to take into account when designing for the sense of touch. Although there is a large body of literature in characterizing vibrotactile sensations, studies characterizing tactile metronome and how well performers are able to follow the beat are few.

Giordano & Wanderley (2015) performed a pilot study that compared guitar performance with both an auditory and tactile metronome. Using a prototype system built with off the shelf parts, they asked four guitar players to play a G major scale at 60 and 120 BPM using both an auditory and tactile metronome. The tactile metronome output discrete haptic pulses much like the auditory metronome outputs clicks for each subdivision of time. The data they gathered from this pilot study preliminarily showed that a tactile metronome can reliably cue a performer to the tempo with accuracy comparable to an auditory metronome.

The rest of the literature involving beat detection with vibrotactile stimulation is with regards to tapping to the beat of rhythms rather than keeping time with a metronome. Ammirante et al. (2016) studied tapping the beat of the rhythms with vibrotactile stimulation. Using voice coils embedded in a chair, they asked participants to tap to the beat of a simple and complex rhythm played using auditory, tactile, and bimodal (auditory-tactile) cues at large and small magnitudes. The inter-tap interval was measured to determine how well they tapped along to the cues. The authors found that people could follow along with the tactile cues as well as the auditory cues for simple rhythms if the tactile cues were sufficiently prominent (large magnitude condition).

2.4 Vibrotactile displays

A vibrotactile display is a technological device that renders and transmits haptic effects to someone holding or wearing the device (Choi & Kuchenbecker, 2013). Such displays can be used to transmit haptic effects or simulations with a wide variety of actuators, such as ERM motors and voice coil actuators. They come in two forms: monolithic, where the entire device vibrates to provide haptic sensations to a wide area, and, localized, where multiple actuators provide sensations across multiple small zones. According to Choi & Kuchenbecker (2013), the vibrotactile displays can be used for a wide variety of applications, such as:

- Simulating the haptic feeling of materials such as floor surfaces (Visell et al., 2009) or fabrics (Huang et al., 2003).
- Communicating with sensory impaired people through tactile sensory substitution, such as helping the hearing impaired to perceive music (Karam et al., 2010; Egloff et al., 2011), or to help the visually impaired perceive their environment (Bach-y Rita & W. Kercel, 2003).
- Communicating navigational assistance to vehicle operators (Erp et al., 2005) or soldiers (Hartnett et al., 2017) whose visual attention is better suited to the task at hand rather than a navigational display.
- Providing tactile stimulation in multi-sensory art installations, such as the Ilinx garment (Lamontagne et al., 2015) or the *Haptic Fields* installation (Hattwick et al., 2017).

Such information can be presented using Tactons. Tactons is a portmanteau of the words tactile icon, and refers to rendered haptic effects that provide a message to a user (Brewster & Brown, 2004). According to Brewster & Brown (2004), five basic parameters can be used to design a tacton. These are: frequency, amplitude, waveform, duration, and rhythm. Additionally, body location can provide an additional parameter if using a distributed set of actuators. Brown et al. (2005) did formal studies in the perception of tactons and found they were as uniquely identifiable as earcons and that rhythm was a very effective way to provide identifiable tactons. Later studies by Giordano (2016) sent easily distinguishable tactons to performers who had to play a melody previously associated with the tacton, and found a high degree of recognition for the tactons.

Using a monolithic vibrotactile display, tactons can be used to display the tempo information of a click track. There were two options readily available at the start of the project: the Soundbrenner Pulse and the Vibropixels. The following subsections will describe the two systems and present some of their features.

2.4.1 Soundbrenner Pulse

The Soundbrenner pulse is a commercial solution of a vibrotactile display that is specifically designed as a watch-like tactile metronome for musicians (Soundbrenner, 2017). It is marketed towards musicians as an alternative to audio click metronomes and click tracks. It can be controlled either via a smartphone application called *The Metronome by Soundbrenner*, or via the MIDI beat clock generated by a supported Digital Audio Workstation (DAW), using the OSX application *Soundbrenner DAW Tools*.

Fig. 2.1 Image of the Soundbrenner Pulse with a Smartphone running *The Metronome by Soundbrenner* application (Soundbrenner, 2017), a watch-like vibrotactile display used for displaying tempo information to the user.

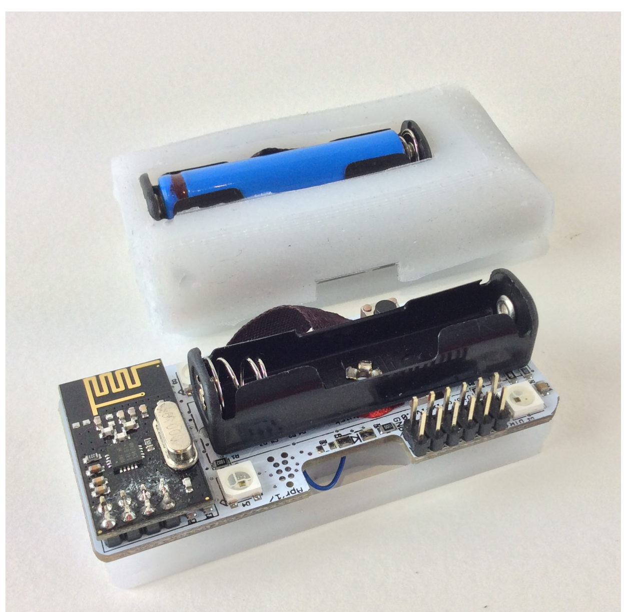

Fig. 2.2 Image of the Vibropixels, one with a silicone cover and battery, the other without. This system is a wireless modular vibrotactile display that can also display RGB colored lights.

2.4.2 Vibropixels

The Vibropixels, developed at the Input Devices for Music Interaction Laboratory (IDMIL) at McGill University, is a reconfigurable, scalable vibrotactile display system (Hattwick et al., 2017). A transmitter connected to a PC wirelessly sends control messages, generated by an interactive software written in Max/MSP, to one or many different Vibropixels in the network, because the Vibropixels are individually or group addressable. The parameters of the control message represent an envelope for the vibrotactile pulse and the receiving Vibropixel outputs a pre-programmed shape based on this envelope. Two motors, a coin cell and a pager motor, are available and individually controlled allowing for different textures of haptic pulses to be generated. The Vibropixels modular design allows them to be reconfigured for use in a wide range of applications. They can be used as a wearable device and are easily programmable and customizable.

Chapter 3

Design Requirements

Before undertaking any design of new technologies, requirements for the design should be gathered. These provide guidelines for what the technology should do and how it would help the user. Since the need for the system was initially introduced by Prof. Guillaume Bourgogne, getting his unique requirements was the best place to begin this project. This was done with semi-structured interviews, which combines focused and open-ended questions, to keep the interviews conversational and allowing unexpected ideas to be brought up that might not have with only focused questions (Hove & Anda, 2005). This chapter will present some background to how the requirements were gathered, a list of the requirements with detailed descriptions, and a description of the system's intended use.

3.1 Interview

An initial interview was the first step done for this project to gather requirements in order to begin designing the system. By using an iterative development process, prototypes would be demoed to Prof. Bourgogne and then he would help refine the design requirements.

An iterative development process is "a cycle of design, test and measure, and redesign, repeated as often as necessary" (Gould & Lewis, 1985). By using an iterative development process, Prof. Bourgogne would help in developing a better design by giving his feedback and helping to further refine the requirements.

The initial interview aimed to understand what Prof. Bourgogne had in mind for the project, and to gather enough requirements to develop a starting prototype. As mentioned previously, the initial interview was semi-structured. Some of the major points that this interview intended to explore were:

- The nature of contemporary music and conducting, in order to establish context.
- The specific problem that required a solution.
- Techniques that have been used to solve this problem and their shortcomings.
- What technique would work better for the conductor.
- A specific practical application in which to use this new technique.

After this interview, enough information was acquired to start developing a prototype based on the understanding of the requirements. It was demoed to Prof. Bourgone, who cleared up any issues this prototype had due to the designer's poor understanding of the requirements. He helped refine the requirements, which allowed planning for a new prototype and the cycle continued until the prototype met his requirements. Five iterations of this cycle were done before a prototype met all the requirements. Fig. 3.1 shows a flowchart describing this process.

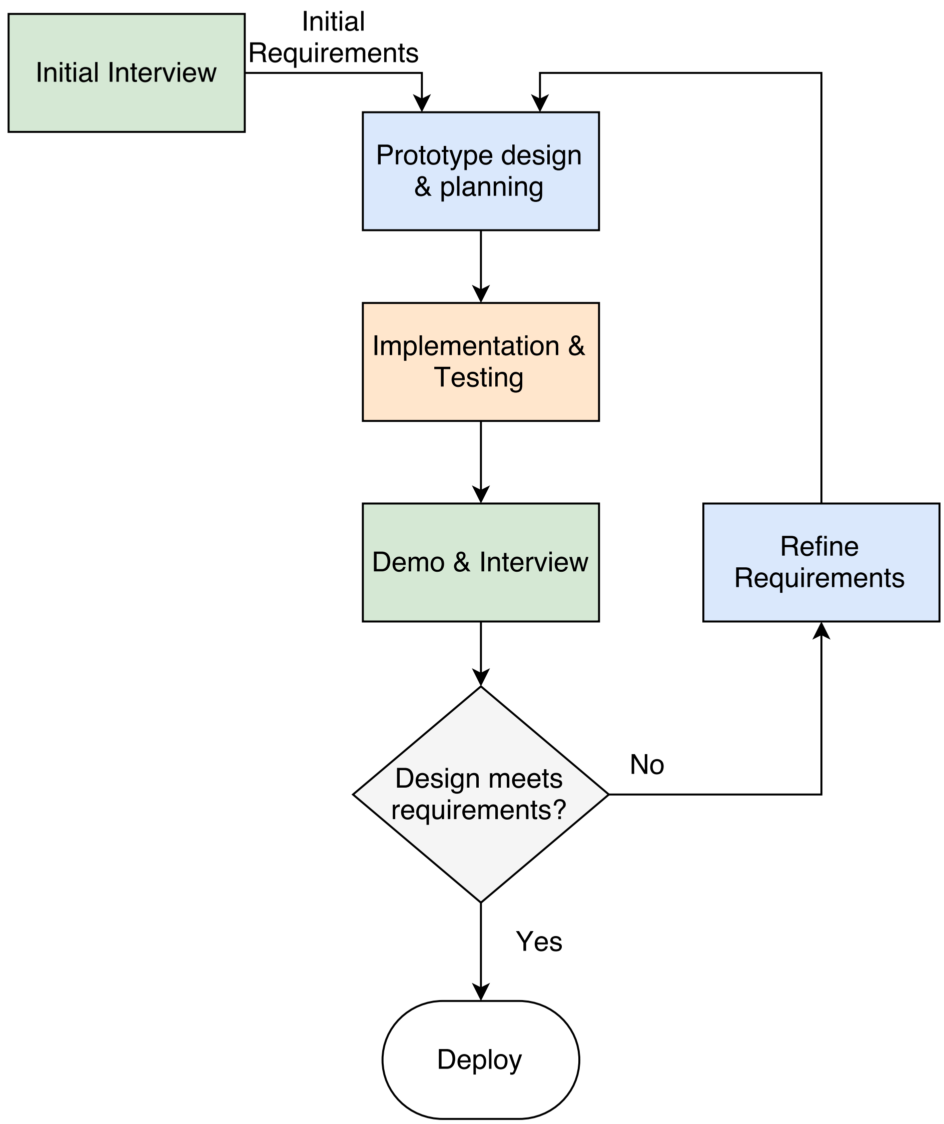

Fig. 3.1 Flowchart detailing the design process used in the creation of the vibrotactile click track. Five iterations of this cycle were done before a prototype met all the requirements.

3.2 Requirements

The following section describes all of the major requirements gathered and refined from the various interviews that took place with Prof. Bourgogne over the course of the project. They are described in detail below:

1. *The system should be able to convert pre-existing audio files with click tracks for musical pieces in the repertoire.* The audio click tracks are used for mixed music pieces in the contemporary music repertoire. When the artistic director chooses one of these pieces to perform, the orchestra is typically given the score, an audio file containing the pre-recorded audio it must synchronize to, and an audio file with the click track that should be played through an earpiece. The technicians that help the orchestra are given the two audio files and tasked with playing them in synchrony during the rehearsals and concerts. Many pieces in the repertoire behave this way.

   Therefore, to reduce the work that would be required in programming an accurate vibrotactile click track from the score, the system should use the audio click track as a basis for generating the haptic click track. This would require processing of the audio file to extract the timing requirements for any given piece and generate a vibrotactile click track from it.

2. *The vibrotactile click track should be able to handle the aesthetic qualities of contemporary music.* Works that are played in the contemporary music repertoire are not always in the same time signature or tempo throughout the entirety of the piece. As an example, the audio click track of the musical work *Charges* by Raphaël Cendo was provided. Fig. 3.2 shows the audio waveform of a ten second sample of the audio click track. As can be seen from this figure, there are two tempo shifts in the span of

ten seconds, one shift to a faster tempo just after the 4 second mark, the other to a slower tempo just before the 6 second mark. Live performances of these works need a system that can seamlessly handle such changes.

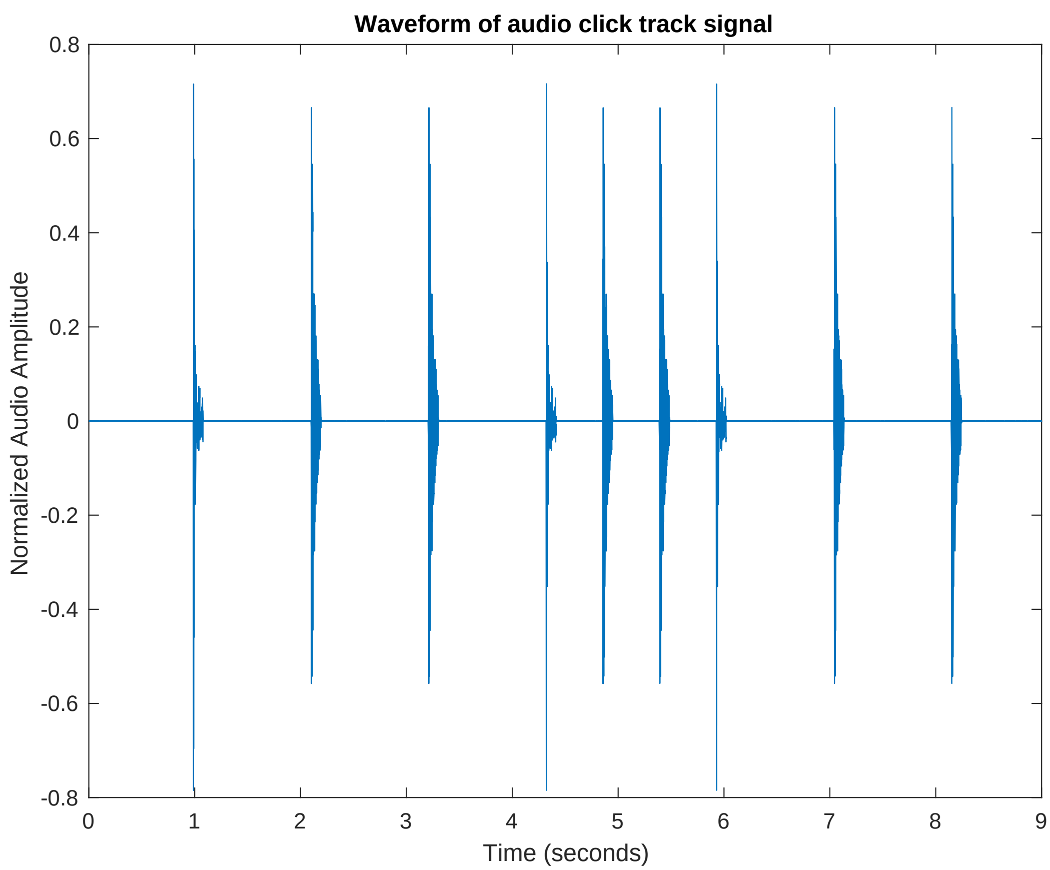

Fig. 3.2 A 10 second waveform sample of an audio click track for the contemporary musical work *Charges* by Raphaël Cendo. This shows the discrete nature of the audio clicks, as well as showing that tempo need not be constant in contemporary music.

3. *Downbeats should be distinct from upbeats.* In the audio click track provided, the downbeats have a distinct sound. The downbeat marks the first beat in a measure of music. This is useful to a conductor for knowing the beginning of a new beat. Having the downbeat feel different from the upbeat would be a nice feature to have for the

conductor using a vibrotactile click track. However, it is not a critical requirement for this system.

4. *Pulses should feel continuous and not like discrete impulses.* Discrete impulses in vibrotactile pulses are when the vibration quickly rises to a maximum vibration, is held there for a specified amount of time and then drops down to no vibration. This is much like an audio click track, where short click-like sounds are used to designate time subdivisions. The works described in the previous chapter use such discrete haptic impulses in order to convey tempo information. However, Prof. Bourgogne required something that allowed him "to feel the variations of tempo with natural motions". He envisioned that "the idea is to have a continuous signal with variations that allows to link the pulses between them and have a better relationship to the tempo". This means that the pulses would feel continuous, never dropping to no vibration. In the interviews, he used the example of a pendulum motion to describe how he believes the pulses should feel, where the pendulum alternates between two positions, but never stops moving. During interviews, visual representations of this were drawn and sung in order to emphasize how it should behave. Fig. 3.3 shows a representation of what was drawn during the interviews.

5. *Pulses should ramp up to a maximum peak* A ramp up to a maximum peak allows for anticipation of the next beat. This gives the conductor a better sense of when the next beat occurs, especially when there is a change in tempo.

6. *The peak of a pulse should be where the click is found in the audio click track.* The maximum peak of the vibration should be located where the click of the audio track is heard. This allows for synchronization of the vibrotactile click track with the audio click track. It is also a good way to evaluate how well the vibrotactile click track

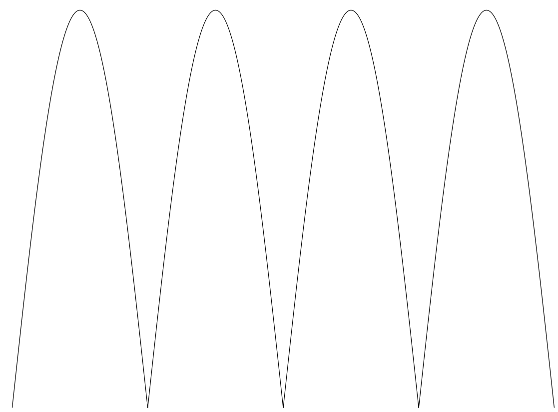

Fig. 3.3 Graphical depiction of vibration pulse intensity as drawn by Prof. Bourgogne. A similar depiction was drawn when discussing how the vibrations should feel and behave. This shows how the haptic pulse should feel to the user, rising up to a specified maximum amplitude and dropping down to a specified minimum (non-zero) amplitude.

outputs the timing information generated using the audio click track. This should be as distinguishable as possible to have a good sense of the pulse in the click track.

7. *Length of a pulse should be dependent on the tempo.* In other words, faster tempos should have shorter pulses than longer tempos. This, much like one of the previous requirements, allows for anticipation of the next beat. At a faster tempo, the ramp up will be much faster than at slower tempos because there is less time between pulses so the overall pulse will be shorter. This, according to Prof. Bourgogne, would allow a conductor to "feel the variations of tempo with natural motions".

8. *Software should allow rewinding/fast-forwarding the vibrotactile click track during rehearsals.* This is of particular importance during rehearsals where the conductor might stop the orchestra to say something and then have them start up at a particular point in the piece. This would require going back a certain number of beats from where the orchestra paused. It must link the click track and the concert audio together,

such that when the vibrotactile click track goes forward or backward the audio will also move to that point in time.

3.3 System use

From the interviews, a general outline of how the system would be used was determined. The system would have two actors involved in its use: the technician and the conductor. It should behave similar to the current system used for audio click tracks. In the audio system, the technician controls a software environment that sends the click track to the conductor's earpiece and plays the audio over the concert speakers, in synchrony. The conductor is wearing an earpiece and conducts the orchestra based on the tempo information in the click track. The vibrotactile click track system should operate in a similar way, where a technician uses a software environment to send vibration pulses to a vibrotactile display, and the conductor wears the vibrotactile display that will be giving him the tempo information as vibration pulses.

Some form of software GUI then needs to be provided to allow the technician to load in the files and play the concert audio and vibrotactile pulses synchronously. The software should transform the tempo information of the audio click track and shape the vibration pulses according to the requirements listed in Section 3.2. The information required to shape the vibration pulses should then be transmitted wirelessly to the vibrotactile display system. The tactile display being worn by the conductor will vibrate in turn, providing the conductor with the required tempo cues. Fig. 3.4 shows a use case diagram of the users and how they would interact with the system.

The next chapter will describe the actual system that was created based on these design requirements in technical detail.

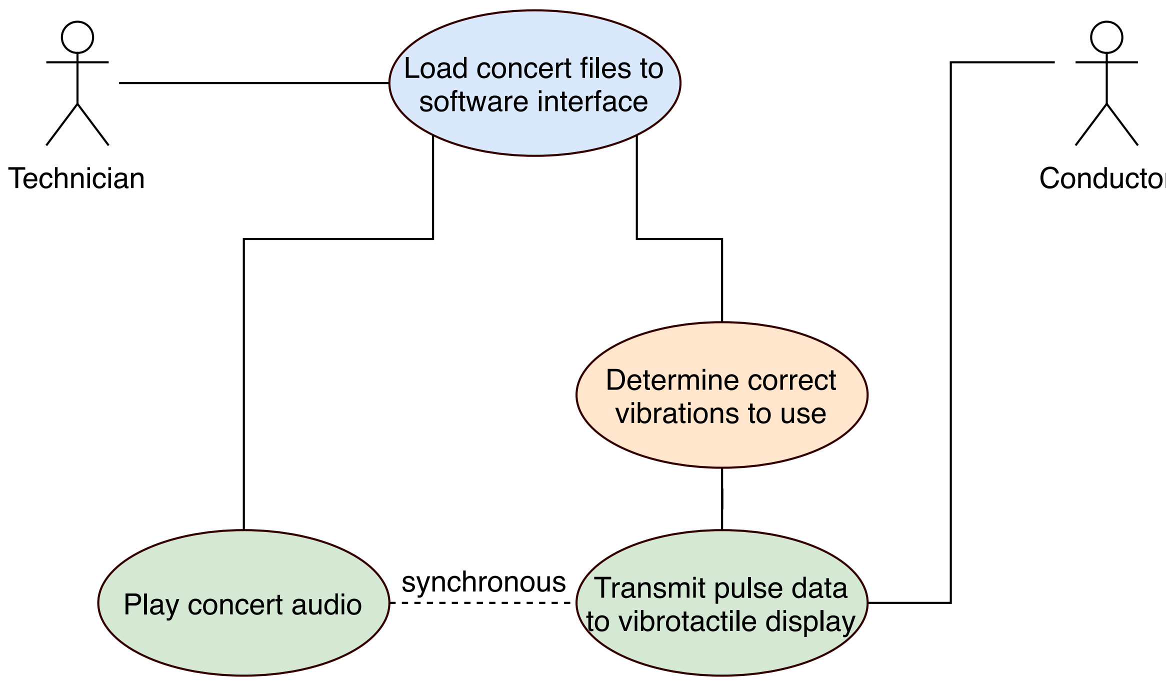

Fig. 3.4 A use case diagram for the vibrotactile click track system, based on the requirements gathered from the interviews.

Chapter 4

Designing the system

With the requirements gathered from the initial interview, prototypes could start being developed for the vibrotactile click track system. An iterative development process was used in designing the system. This helped to improve the overall design and to ensure that the product delivered met Prof. Bourgogne's expectations. This chapter will first present some of the initial explorations of the tactile displays available and developing the tacton for the system. It will then present the final configuration of the system in detail, from aspects of the software algorithms used to the necessary hardware connections.

4.1 Early Design Decisions

Some early design decisions were necessary before starting to prototype in earnest. The first design decision was choosing a tactile display platform to provide the vibrations necessary. Two vibrotactile display platforms were compared and contrasted to determine which would be a better fit for the system. The other design decision was how to present the tactons for the pulse. Several different types were created and presented to Prof. Bourgogne to see which would present the tempo information in the way he described during the initial interview. This section will describe some of the investigations performed for these necessary decisions employed in the final design.

4.1.1 Choosing a tactile display

Before the first prototypes were developed, a tactile display platform was required. Since developing a custom tactile display was beyond the scope of this project, an existing one would be used as the hardware platform for this application. There were ultimately two distinct choices available for use: the Soundbrenner Pulse (a commercial vibrotactile metronome) and the Vibropixels (an open hardware general-purpose vibrotactile display).

(a) Soundbrenner Pulse (Soundbrenner, 2017) (b) Vibropixel

Fig. 4.1 The possible tactile displays for this project. The display in (a) is a commercial device that is sold as a vibrotactile metronome. The display in (b) is an open hardware device that can be customized for use in a wide range of applications.

The Soundbrenner Pulse is a watch-like vibrotactile display specifically used for metronome purposes that was introduced in Sect. 2.4.1. It was designed in Hong Kong and Germany by the Soundbrenner Limited company. It can be controlled either via a smartphone application called *The Metronome by Soundbrenner*, or via the MIDI beat clock generated by a supported Digital Audio Workstation (DAW), using the OSX application *Soundbrenner DAW Tools*. The investigation of the device was done with version 1.0.17 of the firmware and version 1.0.0 of its corresponding software *Soundbrenner DAW Tools*, the latest versions at the time its use was explored.

The platform was explored to see if it could meet the particular requirements for the design. However, there were limitations to customization that made it not ideal for the application. For example, fine control over the vibrotactile pulse of the Soundbrenner was not available to the user. The smartphone application allowed customizing the pulse to only one of the nine available presets and output discrete haptic pulses only. There was no way available to the public to customize the pulses to provide continuous pulses. Therefore it could not meet the requirement of having the continuous feeling pulses rather than discrete ones. Additionally, changing the tempo or time signature with the smartphone application required manual user input, as of the software versions mentioned previously. Changing the tempo via MIDI beat clock was possible with a DAW project that is setup beforehand, but changing the time signature requires manual user input. This is not ideal for an application where pieces with multiple changes in tempo and time signature can occur regularly.

Therefore, the Soundbrenner Pulse is better geared towards beginners and studio musicians who need to be tightly synchronized to a single tempo without time signature variations. The lack of customization meant that this product could not be used as the tactile display platform for this design and still meet the identified requirements.

The Vibropixels did not have many of the features that the Soundbrenner Pulse did, for instance having watch-like straps for easy mounting and providing multiple ways to trigger vibrations, but by being an open design, it was easy to customize and add in more features. The firmware was available for modifications as Arduino sketches and libraries. As well, there was a Max GUI that is used to control the Vibropixels which was open and could be customized to meet the identified requirements. The Vibropixels were also much cheaper than the Soundbrenner Pulse. Furthermore, there were two motors on a Vibropixel: a coin cell motor (seen in Fig. 4.2(a)) and a cylindrical motor (seen in Fig. 4.2(b)). The coin cell is smaller and produces more gentle vibrations. The cylindrical motor is quite powerful and is suited for larger vibrations that can be used to alert users. This feature allowed exploring blending the two motor outputs for the vibrotactile pulses of the system, which could not be done with the Soundbrenner Pulse. For these reasons, the vibrotactile display platform was chosen to be the Vibropixels.

(a) Coin cell motor (b) Cylindrical motor

Fig. 4.2 The two motors available on a Vibropixel. Motor (a) is smaller and produces gentler vibrations. Motor (b) is larger and produces vibrations suitable for alerting users.

4.1.2 Exploring different vibrotactile envelope shapes

Prof. Bourgogne had a particular sensation he had in mind for the vibrotactile pulses. During interviews, he attempted to sing and draw a representation of this sensation. In order to try and emulate these sensations, the firmware of the Vibropixel system was modified to produce different vibration pulse envelope shapes. The original Vibropixels shape pulses for both motors using attack, sustain and decay time parameters; the attack increases linearly up to a specified maximum amplitude, the sustain holds the maximum amplitude for its amount of time, and the decay decreases linearly back down to zero.

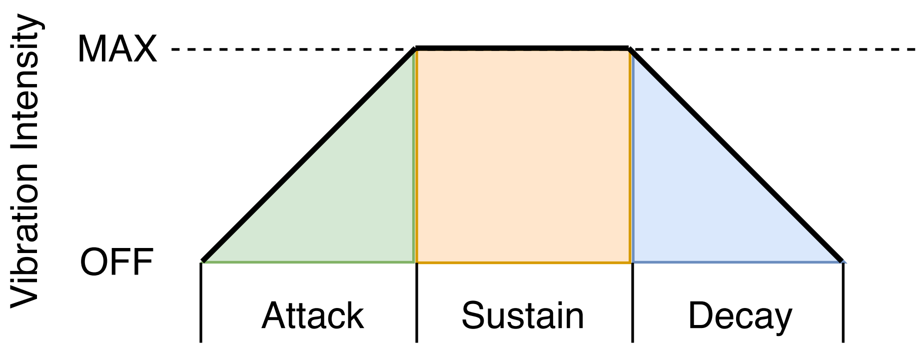

Fig. 4.3 Graphical representations of the three phases of the vibration pulse envelope used in the Vibropixel firmware. The envelope generation is the same for both motors on the Vibropixel.

By modifying the way the vibration envelope is shaped, several different types of pulses were developed to try to emulate the expected vibrations. Four different types of pulses were developed based on what was discussed during interviews:

- **Triangular pulse:** a linear rise and decay model. This produces a triangular pulse, as can be seen in Fig. 4.4(a). This is the simplest and least computationally expensive way of generating the pulse envelope.

- **Exponential pulse:** an exponential rise and decay model. This model exponentially increases up to the specified maximum then exponentially goes back down, as can be seen in Fig. 4.4(b). The idea behind this pulse was to produce well defined peaks while still feeling the rise and decay.

- **Sinusoidal pulse:** based on a rectified sinusoidal shape. This model uses a rectified sine wave to generate the pulse envelope shapes, as can be seen in Fig. 4.4(c). The idea behind this pulse was to generate something akin to a model of pendulum motion, as that was something Prof. Bourgogne mentioned as a "natural" motion that he wanted.

- **Parabolic pulse:** based on a parabolic shape. Uses an inverted parabola to generate the envelope, as can be seen in Fig. 4.4(d). The idea behind this pulse was to produce something similar to the sinusoidal pulse that was less computationally expensive, but still a "natural" motion, since it is based on projectile motion.

These different envelope shapes were performed on the small coin cell motor of the Vibropixel. The cylindrical motor vibration output strength is much less linear than the coin cell motor, therefore differences between the envelopes was not distinguishable. In addition, it is much less power-efficient, leading to the batteries draining faster if it were used continuously (Hattwick et al., 2017). However, the cylindrical motor was used to output a short 25ms rectangular pulse when there was a downbeat to distinguish the downbeats from the other beats.

The different envelope shapes were demoed to Prof. Bourgogne to determine which one was the most similar to his expectations. Although he could distinguish between the different envelopes, none quite met his expectations. The exponential envelope did not provide a strong enough sensation of the beat. Its quick ramp up and ramp down left the entire pulse feeling very weak and therefore not good enough for his needs. The sinusoidal

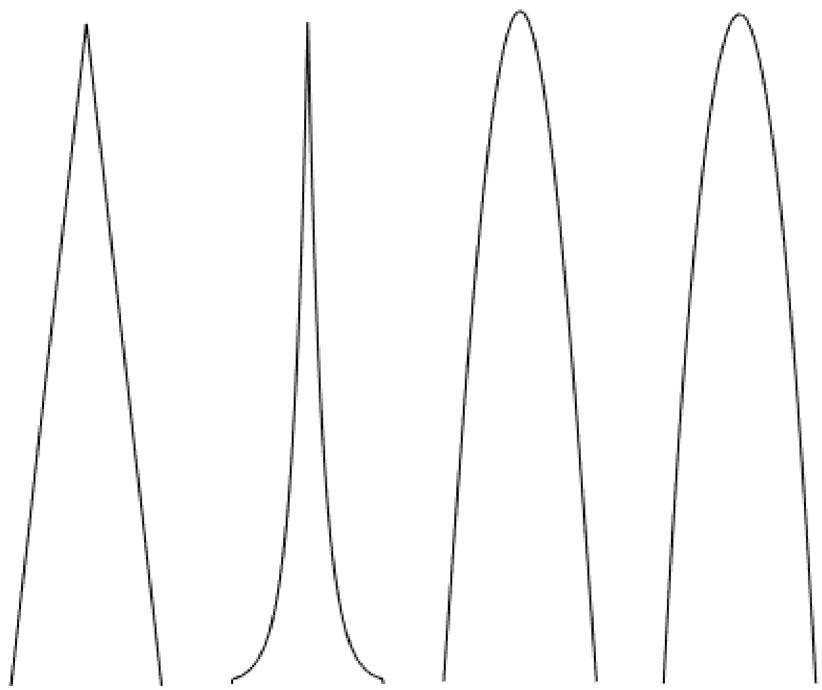

(a) Triangular (b) Exponential (c) Sinusoidal (d) Parabolic

Fig. 4.4 Graphical representation of the four vibration pulse envelopes developed for the system

and parabolic were very similar to him and did not give him a strong enough sensation on where the beat was. He was able to distinguish the "round" peaks and he found that was not good in letting him know exactly where the beat was. The triangular pulse was the best to him in terms of providing anticipation with the ramp up, but he still found it difficult to tell where the beat was. However, after feeling a downbeat with the triangular pulse, including the 25ms rectangular pulse on the cylindrical motor, he thought that was the best way to feel where the beat was while still having the anticipation for the next beat. From this, the final vibration pulses for the system were modified to be triangular envelopes on the coin cell motor with a 25ms vibration from the cylindrical motor to provide a rhythmical stress on the beat.

4.2 Final Design Software

The software for the final design of the vibrotactile click track system consists of three main parts: the pre-processing script in MATLAB, the Arduino firmware of the tactile display system (Vibropixels), and the Max GUI for controlling the tactile display. The pre-processing script was written for the purpose of determining the correct timing and pulse generation parameters for the vibrotactile click track pulses and writing the details in a TXT file that can be read by the Max GUI. The Arduino firmware of the Vibropixel was modified to allow for continuous pulses and to change the control message structure to shape the pulses to meet the requirements for the application. The Max GUI takes the output of the pre-processing script and outputs control messages as USB serial data to the transmitter at the correct time based on the details in the TXT file and the start of the concert audio. Fig. 4.5 presents the components of the software and the various inputs and outputs. The following section will present more details regarding the software design.

4.2.1 Pre-processing Script

From the requirements outlined in the preceding chapter, there should be a ramp up (which will be referred to as the "attack" phase of the pulse) and ramp down (which will be referred to as the "decay" phase of the pulse) to the haptic clicks, to give the conductor a sense of anticipation for each beat. However, the peak of the haptic click should be at the point in time where the audio click occurs. This means that the motor should begin the attack phase before the audio click occurs in time, at the midpoint between two clicks. The motor will start the attack phase when the control message is received by the tactile display. Fig. 4.6 illustrates where the control message should be sent in order to reach the required peak at the correct time.

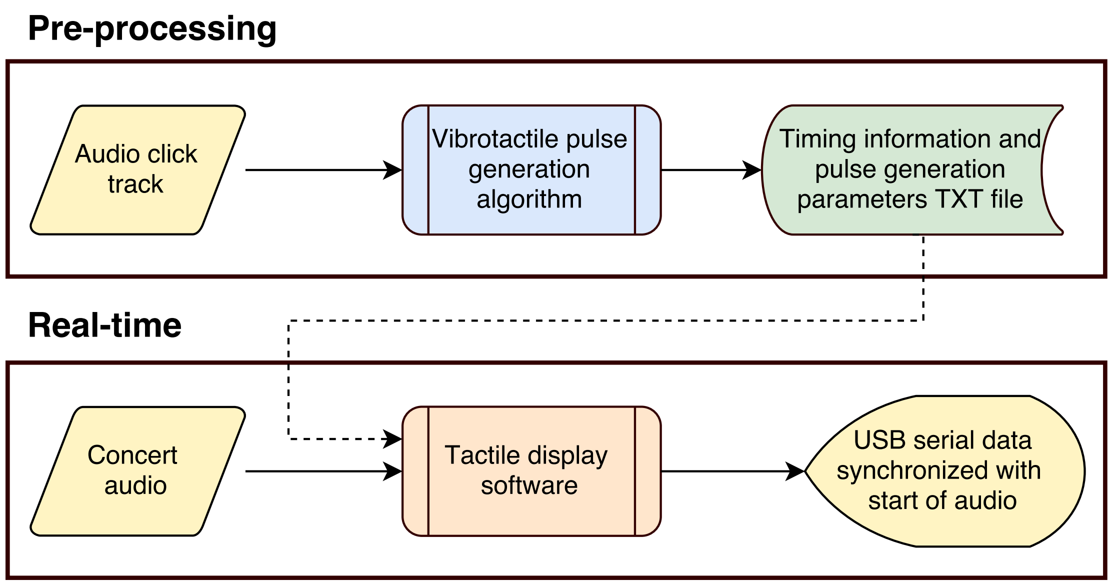

Fig. 4.5 Overview of software inputs and outputs of the final design of the vibrotactile click track system.

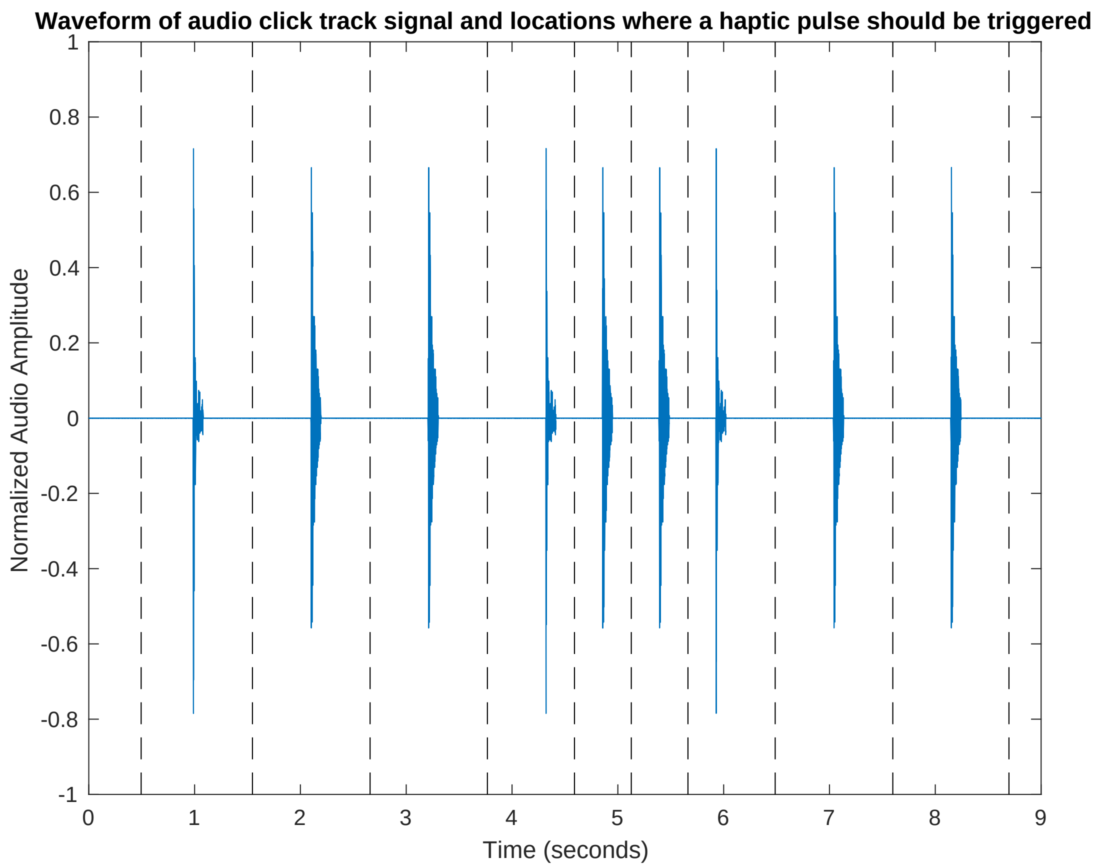

Fig. 4.6 Locations where the motor should begin the attack phase in order for the peak of the pulse to be aligned with the beat. These locations are indicated by the dotted lines. Note that the length of time for attack and decay are not equal when there is a shift in tempo.

We can also see from Fig. 4.6 that the lengths of time for attack and decay are not always equal. If there is a shift in tempo to something faster, the attack will be longer than the decay. Likewise if the tempo shifts to something slower, the attack will be shorter than the decay. The length of time for the attack and decay of a click at time t_i can be determined using the following equations:

$$Length_{attack} = \frac{1}{2}(t_i - t_{i-1}) \tag{4.1}$$

$$Length_{decay} = \frac{1}{2}(t_{i+1} - t_i) \tag{4.2}$$

The pre-processing script therefore determines the locations in time of all the clicks in the audio click track. It calculates the attack and decay for the haptic pulse that will represent each click. Using this it can determine the overall length of time of the envelope of the haptic pulse. It can then generate a control message that the Vibropixel will use to display the haptic pulse. The Vibropixel control message can be seen in the bottom half of Fig. 4.7 and the description of each byte can be seen in Table 4.1.

Having calculated all the parameters for the Vibropixel control message, the script writes it to a TXT file, with the format seen in the top part of Fig. 4.7, with an index and a Timestamp parameter preceding it. It is written in this format to be readable by the Max GUI, which uses the *coll* object to read in the text file and iterate over the pulses. The Timestamp parameter is used to synchronize the vibration pulses to the start of the audio file and provide the Max environment with the accurate time in which to send each pulse message.

These calculations help determine the triangular pulse generated by the coin cell motor. Between the start of the pulse and the start of the next one, at the exact location of the

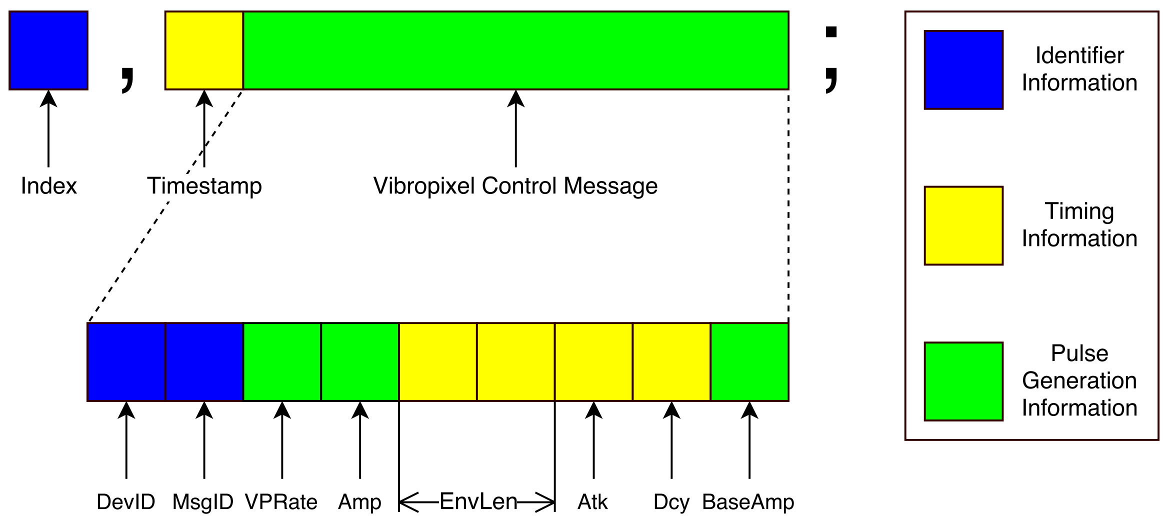

Fig. 4.7 Overview of the different pieces of data for timing and pulse shape generated by the pre-processing script and the Vibropixel control message structures.

Table 4.1 Descriptions of the various control message bytes used for the control message of the tactile display system

Byte	Description
DevID	Device identifier used for addressing multiple receivers
MsgID	Message identifier used for differentiating motors and messages
VPRate	Rate parameter for PWM frequency
Amp	Maximum amplitude of the motor envelope intensity
EnvLen	Length of the envelope expressed as two bytes
Atk	Percentage of envelope length devoted to the attack phase
Dcy	Percentage of envelope length devoted to the decay phase
BaseAmp	Baseline amplitude of motor envelope intensity for continuous vibrations

click in the audio, a message is added to trigger the cylindrical motor with a pre-defined short 25ms pulse. This provides a rhythmical stress on each beat to the conductor, which allows him to understand where the beat would be if there was an audio click track playing.

Since there would be some adjustment of the search parameters of the MATLAB `findpeaks` function to correctly find the peaks of the click track, this script was re-written to take advantage of the MATLAB "live script" format. MATLAB live scripts combine MATLAB code with embedded documentation and allows previewing of code and graphs in-line. This allowed previewing the adjustments in-line before continuing with the rest of the calculations. Fig. 4.8 shows a screen-shot of the live script with the previewed adjusted click track.

4.2.2 Firmware

To meet the requirements outlined in Section 3.2, modifications were made to existing Vibropixel firmware. One of the major changes was to port the existing Arduino methods for controlling the motors to a C++ library, which allowed controlling the motors as software objects and helped make it more intuitive to individually address the two motors.

A new messaging format, as outlined in Fig. 4.7, was developed to apply the haptic click track system to the Vibropixels without interfering with previous functionalities. It allowed removing some unnecessary control bytes, such as LED control, and making modifications to how envelope length was sent, in order to create pulses with more accurate envelope lengths.

This new messaging format was used to create the different vibrotactile pulses outlined in Section 4.1.2, where different message identifiers were used to distinguish between the various shapes. The final tacton described in that section requires two distinct messages to display correctly on the Vibropixel, one for the triangular pulse on the coin cell motor

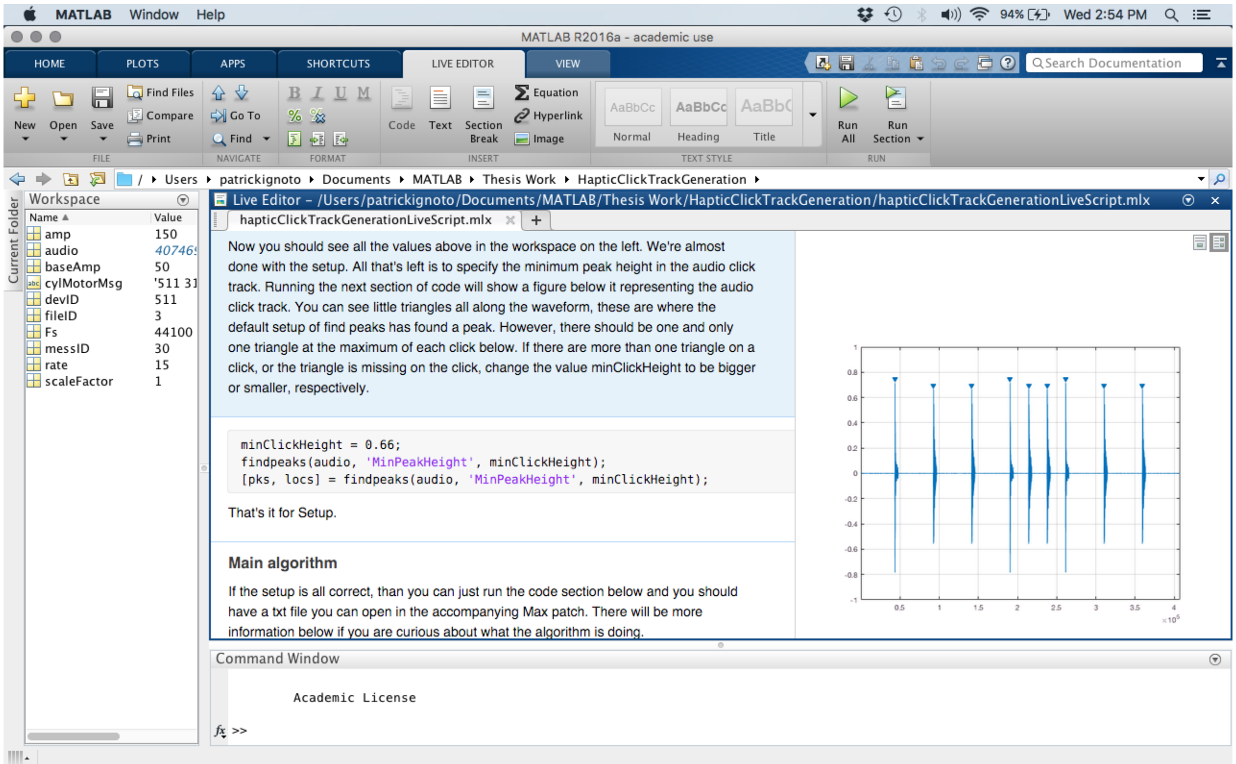

Fig. 4.8 Screen shot of the MATLAB environment running the live script formatted pre-processing script. A technician would use this script once to adjust the search parameters and process the audio click track to generate a file that can display the resulting haptic click track on the Vibropixel.

and one for the short rectangular pulse on the cylindrical motor. The pre-processing script takes this into consideration when performing the calculations and writes both messages for each audio click.

The *BaseAmp* parameter was an artifact from the earliest iterations that was kept to prevent the motor from losing its inertia by turning off. A specific message identifier sets a software flag on the Vibropixel, which prevents the intensity of vibration from dropping below the value specified in *BaseAmp*. This keeps the motor spinning, and mitigates delays caused by the motor starting from a cold start (see Sect. 5.1.2 for more details).

4.2.3 Control GUI

The GUI used to control the system was developed in Max, modifying a previously existing control GUI for the Vibropixels. Max is a visual software programming environment that is commonly used in music and audio processing applications (Cycling '74, 2018). The control GUI for the Vibropixels is developed with Max. The Max control GUI for the Vibropixels was a good starting point for a GUI to control the vibrotactile click track system, because it was more efficient to re-use some existing aspects of the code and it is a familiar environment for musicians and technicians who may need to use the system. In addition, previously existing functionalities, such as setting up communications with the transmitter and a way to test envelope shapes, were maintained in the modified control GUI.

Fig. 4.9 shows the interface that is presented to the user. The left half of the interface controls the vibrotactile click track system. The right half of the interface allows a designer to test the various parameters to create different envelope shapes.

The user is required to load the TXT file output of the pre-processing script and the audio file containing the sound meant to be played for the piece. They can then start the audio and haptic click track together. The haptic click track system uses the `coll` object in Max to read the contents of the TXT file line by line. The *Timestamp* parameter from the TXT file structure tells the system how long to wait before sending the Vibropixel message to the transmitter. Once that time has elapsed, the message is transmitted and the next message is queued. This occurs until there are no more messages to iterate through. In addition, the user can pause both the haptic click track and audio, and can rewind or move the haptic click track forward by a user-specified number of beats. This is linked to the concert audio, so when the user would pause and rewind the haptic click track by 10 beats,

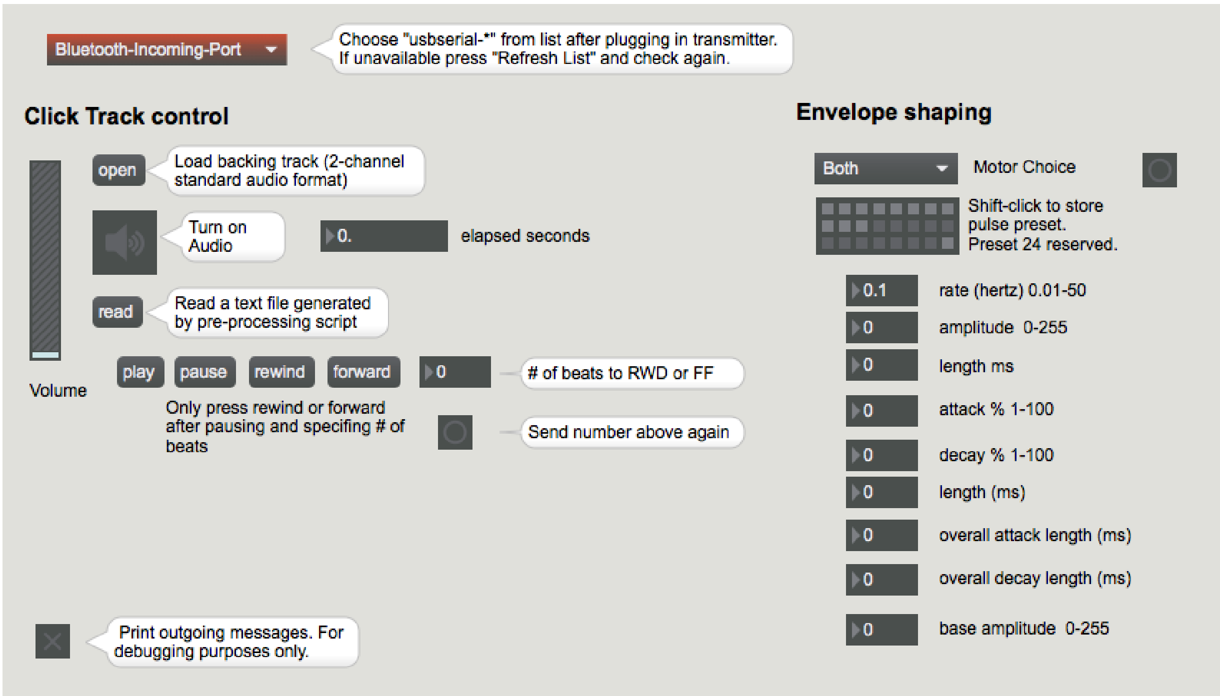

Fig. 4.9 Max user interface to control the output of audio and vibration pulses. The left side of the patch would be used by the technician to display the haptic click track and play the audio at the same time during rehearsals or concerts. The right side is for exploring different vibrotactile envelope shapes and was used to set up the different patterns that were presented to the conductor during the development.

the concert audio will playback at a point in time 10 beats before that.

4.3 Final Design Hardware

The hardware setup for using the vibrotactile metronome is minimal, comprising a computer workstation and a transmitter and receiver for the Vibropixel system. The computer workstation must have the software MATLAB and MAX installed and must have the files for the pre-processing script and the control GUI. The transmitter for the Vibropixel system must have the standard Vibropixel transmitter firmware uploaded. The receiver must have the custom vibrotactile click track firmware uploaded.

In order to use the system in a concert setting, the workstation must be connected to the concert hall's speaker system to play the audio in synchrony with the vibrotactile clicks. The receiver must be strapped onto the conductor's skin, at a location that is comfortable and easily perceivable by the conductor. In the case of Professor Bourgogne, he preferred having the tactile display on the back close to the neck. Fig. 4.10 presents the information displayed by the hardware of the vibrotactile click track system.

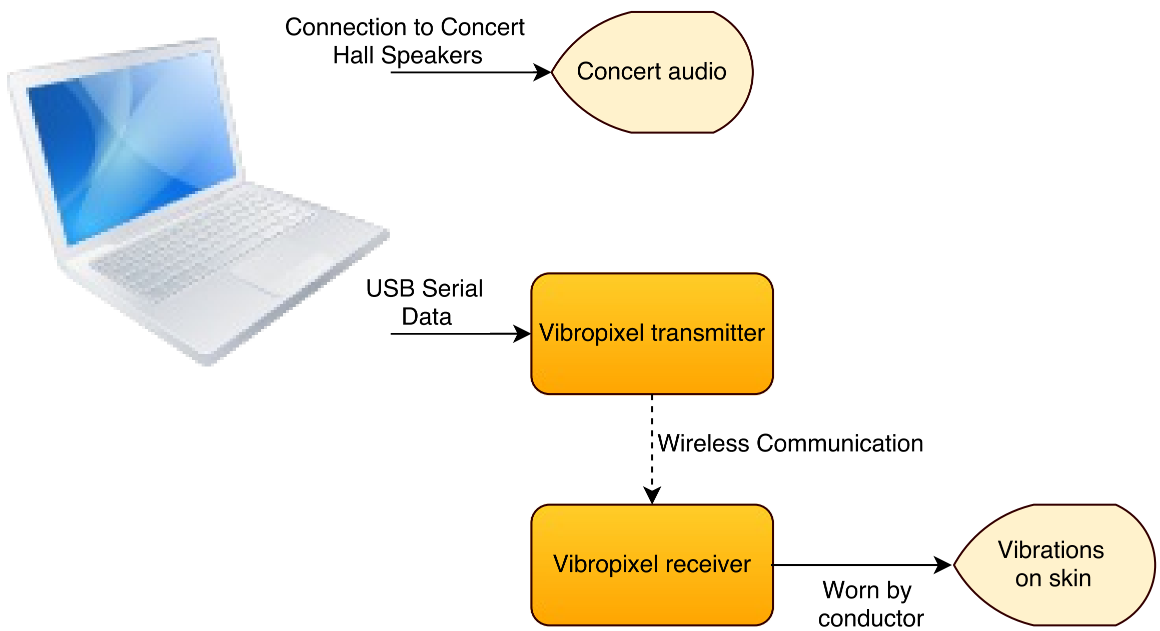

Fig. 4.10 Overview of the required connections and information displayed by the hardware of the vibrotactile click track system.

4.4 Discussion

The design outlined in this chapter meets all but one of the requirements discussed in the previous chapter. The requirement for downbeats to be different than the others was not met as the method used to distinguish the downbeats was used on all the beats in the end. Initially, distinguishing downbeats was achieved by using the *bonk~* extension for Max. *bonk~* is a percussive envelope detector that takes in an audio file and distinguishes attacks of different instruments or sounds (Puckette et al., 1998). This object was found as a way to distinguish downbeats from upbeats, as they are distinct sounds in the audio file. Since the *bonk~* object reads the audio file in real-time, the control GUI was programmed to take in the audio click track, and output the ictus whenever the *bonk~* object recognizes one of the downbeat clicks. This was moved to the pre-processing script after Prof. Bourgogne asked that the ictus be on each beat instead. He claimed that, for his purposes, having distinguishable downbeats was low priority and could be left for future work. Therefore, this particular requirement was left as future work, which includes determining a good tacton to represent it and to add this to the pre-processing script.

The use of this system does differ from the use case diagram initially developed in Fig. 3.4. The use of the pre-processing script adds an additional step that the technician must do but because pieces are comprehensive and static, this is a one time operation. However, other than this additional step, the use case does not change greatly. Fig. 4.11 updates the use case diagram based on these changes.

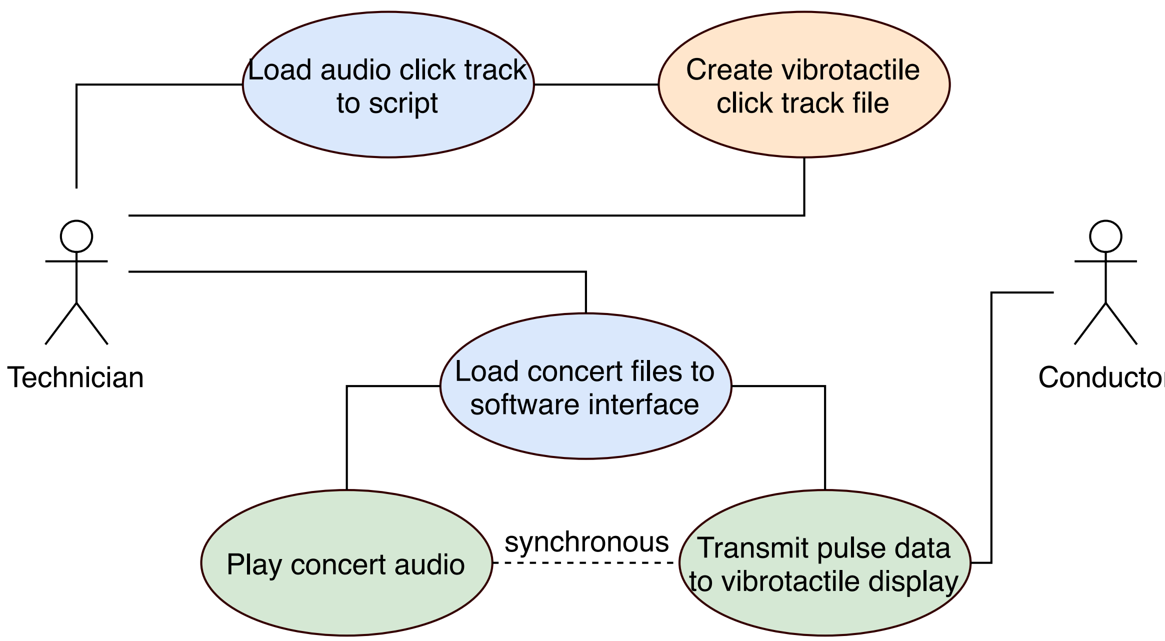

Fig. 4.11 A use case diagram for the vibrotactile click track system, based on the final software design.

Chapter 5

Evaluation

The previous chapters outlined the requirements for the vibrotactile click track and the specific design decisions made to achieve them. In order to evaluate whether those requirements were met, tests were run in the laboratory to visualize the pulses and verify what effect latency might have on the system. As well, interviews with Prof. Bourgogne and other musicians were done to get qualitative feedback on the design. This chapter outlines the tests performed and presents the data collected. Finally, it discusses some of the qualitative feedback gained from interviews with performers.

5.1 Visualizing the pulses

In order to quantitatively determine if the pulses created by the system met the requirements of providing continuous pulses to the user, an experiment was run to visualize the pulses. The experimental results were used to extract the envelope of the pulses and plot them along with the audio click track, allowing the verification of whether the envelope pattern coded was similar to what was discussed in requirements 4-7 of Section 3.2.

5.1.1 Experimental Setup

Using an oscilloscope, we can easily measure and visualize analog electrical signals, however we require a means to convert the vibrations to an analog electric signal. Since the motors are controlled via a PWM signal, it is not possible to visualize the vibrations by measuring the control signal. Therefore, an analog accelerometer was clamped to the vibrating device in order to capture the data. As well, the audio click track data was played in synchrony with the vibrations and captured with the oscilloscope. Fig. 5.1 shows the interconnections between devices for the experiment.

5.1.2 Results and Analysis

Fig. 5.2 shows the amplitude envelope of the pulses generated by the coin cell motor and the audio click track, as the orange line and the blue line respectively. The amplitude envelope of the pulses were shown rather than the raw data to remove the spurious changes caused by the vibrations and provide more clarity to the graph. As well only the continuous pulses generated by the coin cell motor were plotted; it does not include the ictus created by the cylindrical motor because the cylindrical motor creates a much stronger vibration which effectively masked the envelope of the vibrations on the coin cell motor on the accelerometer.

One can see that the amplitude envelope of the vibrations is continuous, having a similar outline to the graphical depiction in Fig. 3.3, therefore satisfying requirement 4 of Section 3.2. Additionally, the pulses ramp up and down and the peak of this pulse is at the location of each peak, which satisfies requirements 5 and 6. Finally, one can see that the lengths of the pulses are dependent on the tempo and are smaller during the faster tempo between 4 and 6 seconds, satisfying requirement 7.

One issue is that the first pulse reaches its peak considerably later than the others. This

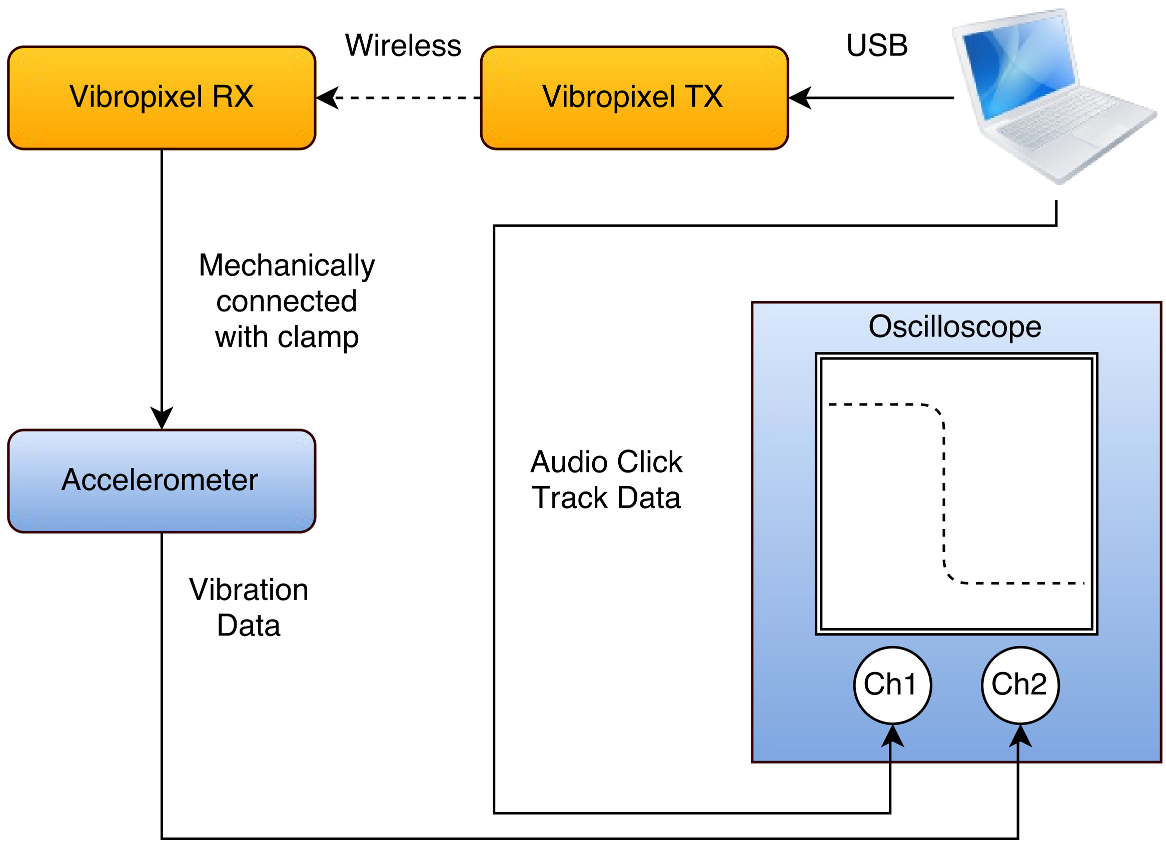

Fig. 5.1 Experimental setup to measure and visualize the vibrations and audio. The oscilloscope measures the audio level from the computer on one channel and the intensity of vibration on the accelerometer on the other channel. This provides a visualization of the synchrony between haptic and audio pulses in time.

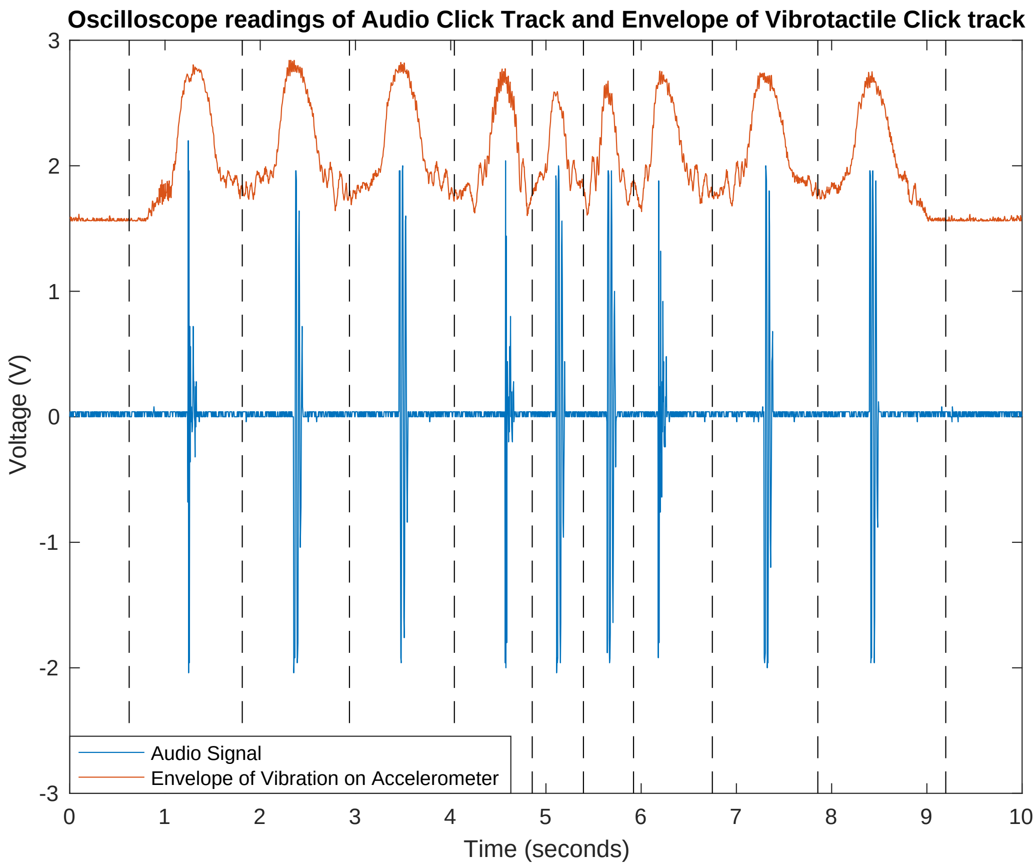

Fig. 5.2 Oscilloscope data showing the audio click track and the envelope of the vibration pulses on the Vibropixel. This demonstrates that the haptic and audio pulses are synchronous in time and that the haptic pulses behave as expected. The dotted lines show where the haptic pulses should begin the attack phase and are provided as a reference.

issue is not apparent with the other vibration pulses, likely because they are prevented from dropping to no amplitude (through the use of the *BaseAmp* parameter). Therefore this issue is likely due to the fact that the motor is starting from a complete stop, and requires more momentum to get up to the peak. This can be easily resolved with some kind of start-up pulse, which could inform the conductor that the piece is about to begin, and would allow the motor to gain the extra momentum needed to be aligned like the others. The option of adding this is possible in the pre-processing script, if the click track for the piece does not already start with a count-in.

5.2 Timing measurements

Latencies in the system would greatly reduce the efficacy of the system if they were to be inconsistent within a playthrough of the piece. The plots in Fig. 5.3 show why this is the case. If a pulse is consistently late, then the tempo would remain the same, and the entire signal would just be delayed a certain amount. However, an inconsistency in the latency would cause the tempo to change and therefore not be an adequate tool in providing tempo information. For this reason, it is important to know if the latencies in the system are random or only systematic. As well it is important to note whether the measured latencies are repeatable; if it is not, it could imply that the latency is not systematic.

In a second test, two measurements were collected, the time when the message is sent to the transmitter in the Max patch, and the time the message is received on the receiving Vibropixel. The measured values were compared to the value given by the pre-processing script to determine the computational latency caused by Max and the operating system and transmission latencies in the wireless transmission through the hardware and the firmware. Since the pre-processing script determines the time messages are meant to be transmitted exactly, we can use this value as the correct theoretical value and compare it to the experimental ones gathered. This would provide information not only to determine how accurate the system is, but whether any systematic errors caused by latencies in the system affect the system in any significant way.

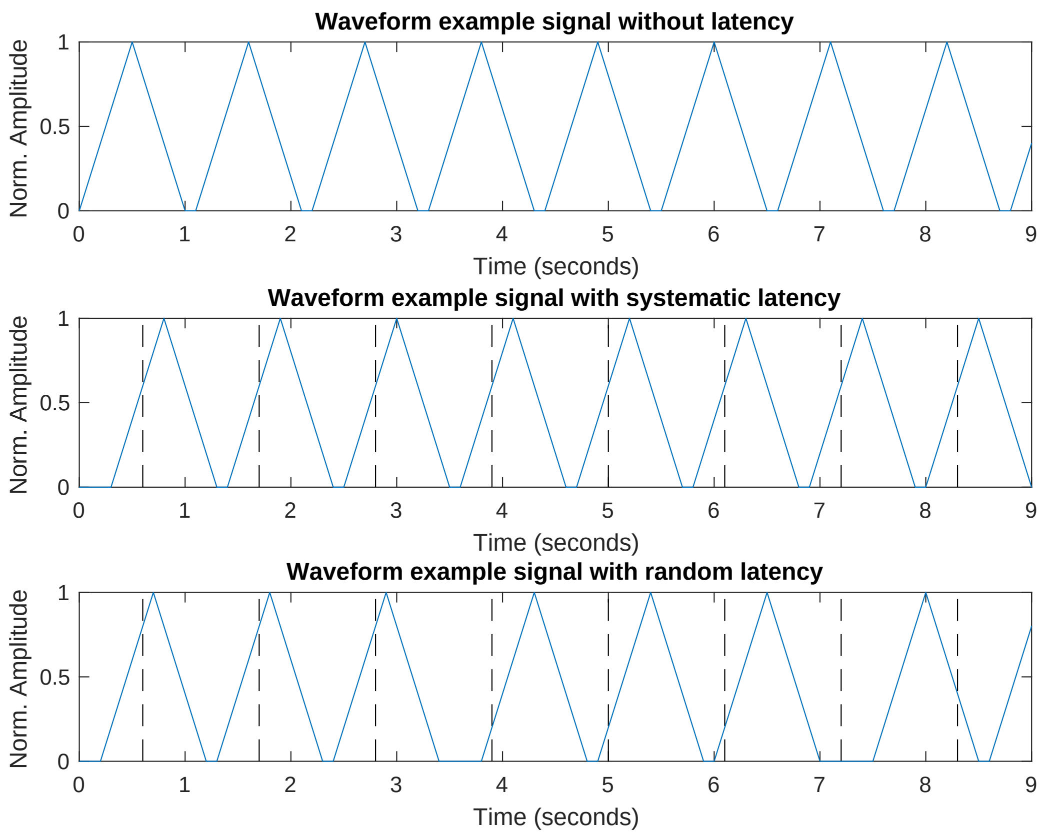

Fig. 5.3 Plots illustrating why systematic latency is not an issue and why random latency is very undesirable. Note that these are not actual measurements but an illustration meant for demonstrative purposes. The dotted line illustrates where the audio clicks should be in time. The audio clicks with the systematic latency keep their intervals intact with, but with random latency this is not the case.

5.2.1 Experimental Setup

The experiment was done in software on a PC and the firmware of the Vibropixel. Both receiver and transmitter were connected to the PC via the USB FTDI connectors. The Max GUI would take a timestamp with the `cpuclock` object at the point in time in the software where the message was sent to the transmitter's serial port for wireless transmission. This was then written out to a `coll` object collection and saved at the end of the test. The receiver would take a timestamp (using the Arduino `micros()` function) of when the message was received and then subtract the time from the timestamp taken upon reception of the first message. It would then write the timing data to the serial port where a script in the Processing environment would read the serial port and write the data to a CSV file. Once the test was done, the data in the CSV file was compared with the theoretical value in MATLAB. Fig. 5.4 shows the flow of the experimental data.

The entirety of the vibrotactile click track for Cendo's *Charge* was played on the Vibropixel and the timing data for each of the 1240 clicks in the piece were recorded (2480 messages/data points in total). This was performed four times to produce four data sets for both transmission and reception.

5.2.2 Results and Analysis

In Fig. 5.5, the experimental error in the time the message is transmitted is shown as box plots for each of the four data sets gathered. It shows a very low median error in milliseconds. Additionally, the distributions for each test are similar, which shows that the latencies caused by the Max GUI and MacOS are systematic and not random. The data also shows that the Max GUI is uniform in timing and any computational latencies caused by it are negligible, as the maximum amount determined was five milliseconds. Previous

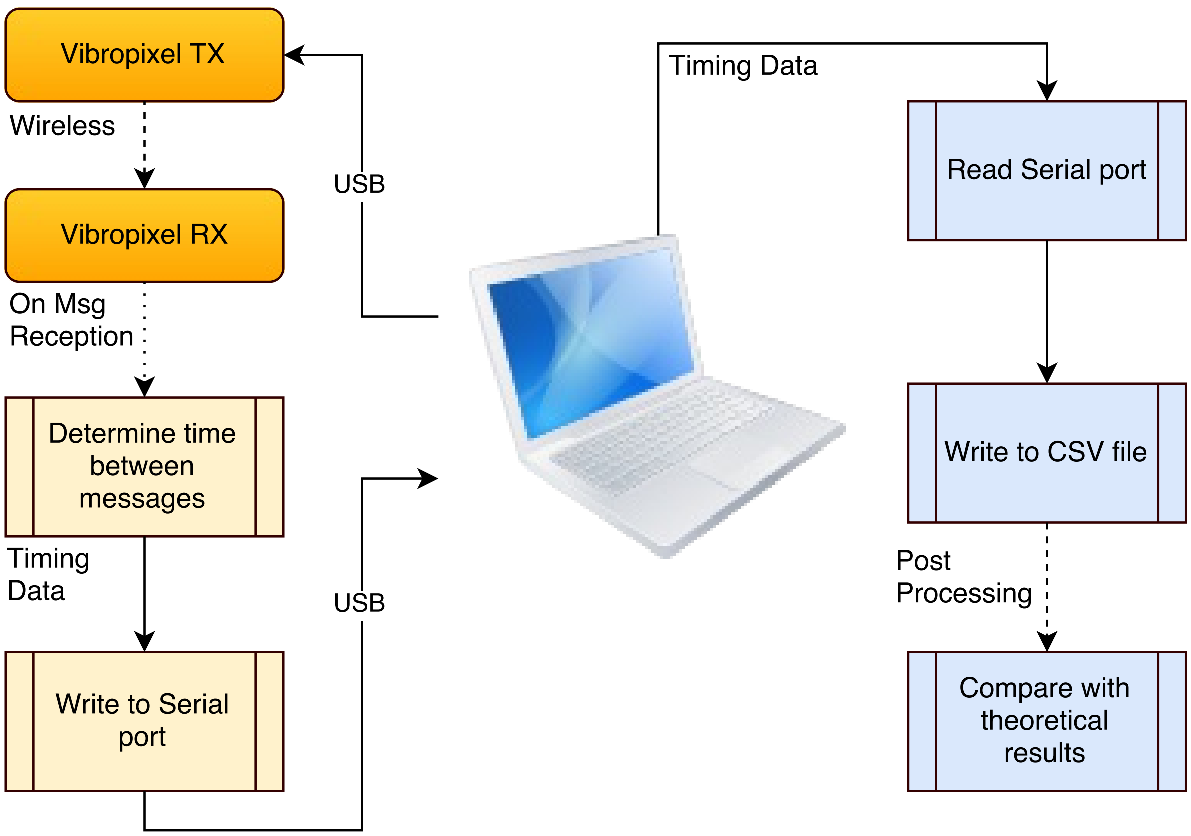

Fig. 5.4 Experimental setup to measure pulse lengths and determine the computational and wireless latencies involved in the system.

studies showed that a deviation from the beat of ten to twenty milliseconds was natural for musicians in a musical performance (Hennig et al., 2011). In addition, at a fast pace of 300 beats per minute, a 16th note would be approximately fifty milliseconds, ten times more than the maximum latency output measured.

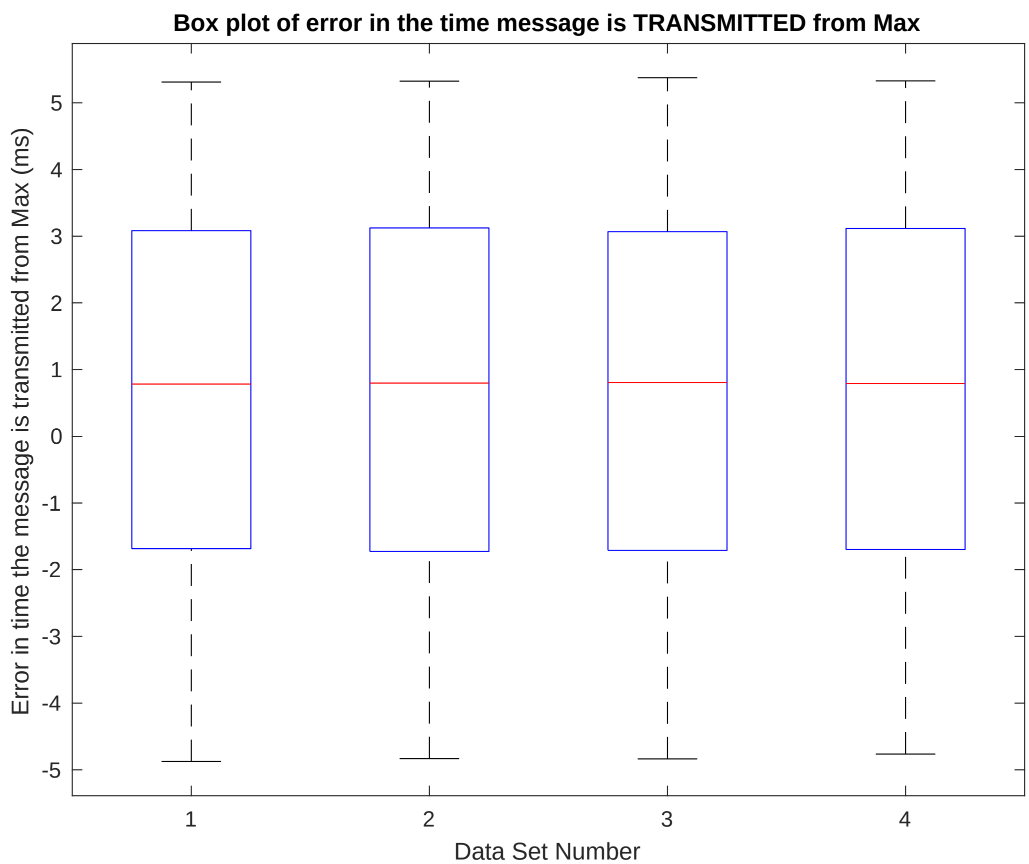

Fig. 5.5 Box plots of the error measured in the Max GUI for the timing experiment. The timestamps taken in the Max GUI were compared to the actual value determined by the pre-processing script. Each data set represents the number of pulses in the musical piece *Charges* by Cendo played in its entirety. The data sets show consistency in the computational latency in Max and can be considered negligible in musical performance.

In Fig. 5.6, the experimental error in the time the message is received is shown as box plots for each of the four data sets gathered. It again shows a low median error in milliseconds, although higher than the median of the computational latency. The variability between tests is more pronounced than in the previous plot. The values have a similar range to the previous plot, therefore it would not cause a perceptible amount of error to a musical performance. These figures show that this system meets the timing requirements.

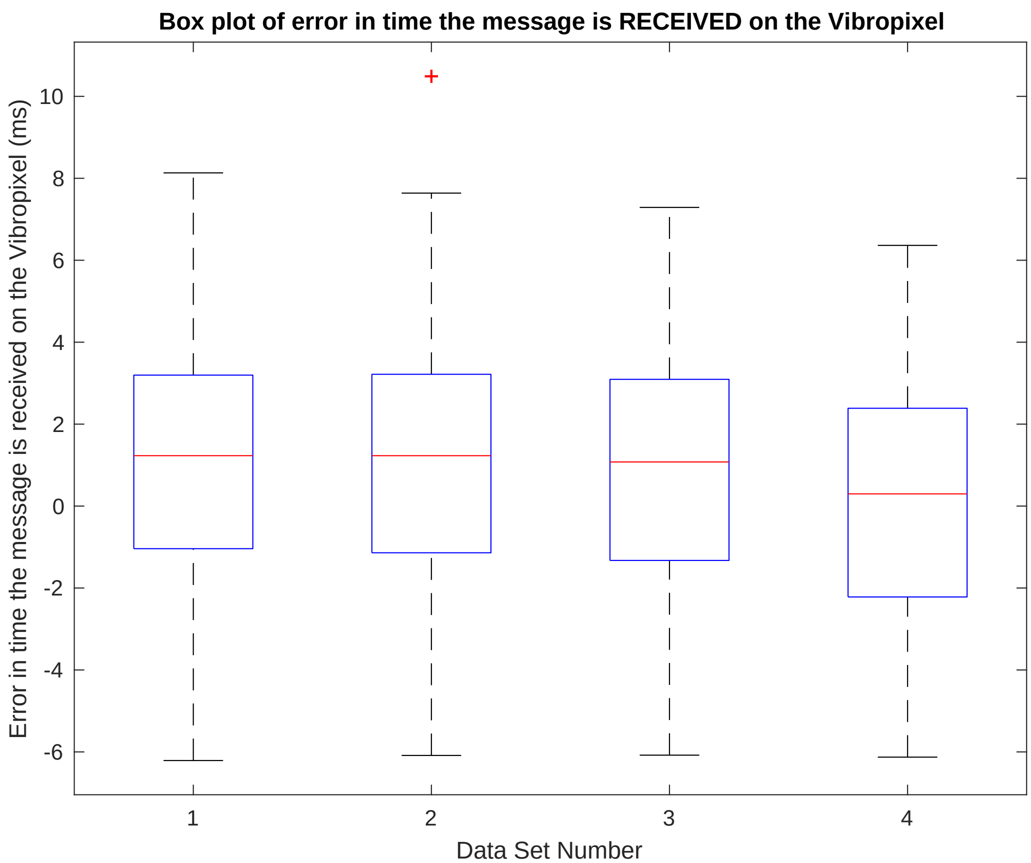

Fig. 5.6 Box plots of the error in the time the message is received on the Vibropixel. The timestamps taken on the Vibropixel were compared to the actual value determined by the pre-processing script. Each data set represents the number of pulses in the musical piece *Charges* by Cendo played in its entirety. Although the wireless latency is more variable than computational latency, it is still negligible in musical performance.

5.3 Performer Comments

The final design was demoed to Prof. Bourgogne and an interview was carried out to get a better idea of his reaction to the system. He was quite pleased with the system, saying "It's really close to what I'm trying to get, because we have this curve which is exactly what I want plus the accuracy of the ictus." He was able to feel the differences in the ramp ups between fast and slow tempos, which he had said was not possible when it was the continuous pulses alone. He preferred when the vibrations were on his core, either the torso or the back of the neck, rather than on arms or legs. This provided bone conduction resonance that allowed him to better perceive the haptic pulses. Overall, these comments suggest that the design was done to specification, and can be used to provide the information necessary to this performer.

The system was also informally experimented with a clarinetist as part of a pilot study to determine whether this design also worked as intended for other performers. Three types of metronomes were demoed to the clarinetist using the Vibropixels, a discrete haptic metronome, the continuous haptic pulses alone, and the continuous haptic pulses with the ictus at the centre. Of haptic metronomes in general, the clarinetist said "Yes it's definitely an aid and it's better than in the ear. I know how much of a problem that can be." Of the three haptic metronomes demoed to her, the continuous haptic pulses with an ictus was found to be the most useful. On the method of continuous haptic pulses with the ictus, she said "it's clear but not as dry like when it's just the click, that's so dry just like a metronome... That's why people say it's so unmusical to play with metronomes. This *[type]* gives the very precise pulse and seemed stronger". The continuous pulses did help with following along when a change in tempo occurs: "you'd need to practice, but the fact that I was able to even get a few is a good sign...".

In addition, the system was presented to a composer whose piece would be used with the system for an upcoming CME concert. On the system, he said "It's amazing because synchronizing the computer and performance is always hard". He also envisioned many more applications that could use vibrotactile cues for other cases.

The comments from Guillaume and the other performers were promising and show that this system can be usable to performers as a haptic click track.

Chapter 6

Conclusion

This chapter summarizes the findings in this thesis, discusses the overall design, and presents other observations made from this process. As well, it details other possible areas of development and future research for this work.

6.1 Summary and Discussion

This thesis presented work done in developing a vibrotactile click track system for the conductor of the McGill Contemporary Music Ensemble. The conductor was interviewed to determine the requirements such a device would require. An iterative design approach was used, where a cycle of gathering requirements, prototyping, and evaluation were used to develop the product. Feedback from the conductor helped improve the design of not just the system, but of the tactons that were used to send tempo information to the conductor. This helped to shape the design based on the vision this user had for the system, as well as provided a better understanding for the use case of such a system. The system was evaluated by looking at the tactons with an oscilloscope, verifying latencies in the system, and through performer feedback using interviews. The evaluation has been positive in demonstrating that the system met the requirements set out by the user.

The next logical step for this system is to use the prototype in a live concert setting. At the time of this writing, development is underway in order to use this system for a CME concert in April 2018. This will provide useful data on how effective this system is in its intended environment, a live concert setting. With the results from this trial run, we will evaluate how well this system behaves in real-world conditions, and will help inform us how robust this system is and whether changes must be made for other concerts.

6.2 Future Work

Current development includes finalizing the location where the conductor will wear the Vibropixel and how the Vibropixel can be placed on the conductor in a way that does not distract the audience and can be easily put on and taken off. As well, ensuring that the audience will not hear the vibrations is paramount to the conductor, and so some investigation of whether this is the case and how to mitigate it is underway.

Future iterations could benefit from designing or modifying the haptic signals according to the standards outlined in the standards document ISO 9241-920:2009 (International Organization for Standardization, 2009). Additionally, future usability studies could make use of standardized assessment tools such as the NASA Task Load Index (TLX) to provide a better assessment of the performance of the system and the workload for users (Hart & Staveland, 1988).

There can also be work done to improve usability of the software parts, add additional features and further improve on the quality of the design. For example, the pre-processing script was written in MATLAB, using the "live script" format, that allowed documentation along with executable blocks of code. This was useful in better documenting how the pre-processing script worked and to give some context for the adjustments that may need to be done by the user to get this to work for different pieces. However, someone completely unfamiliar with MATLAB might have a difficult time understanding how to use it. In addition, MATLAB is an expensive and not commonly available program outside academia, which can be problematic for music technicians in non-academic areas. The pre-processing script can easily be ported to Python, with the NumPy library, or with Octave. This would allow applying the algorithm in an open-source environment. However, this would eliminate the ease offered by the MATLAB live script formatting. A further look at the options available and possible user studies with concert technicians would likely help perfect this aspect of the product.

There are many other possible areas for future research that can be explored from this project. For design researchers, an interesting study based on this project is to determine whether a product designed based on the requirements of a single user is as effective as a product designed by collating the requirements of several potential users. The idea is to perform more formal testing on whether designs created in such a way are as user-friendly and practical as a design created in the conventional way. Similarly, a study could be done on how tactons are perceived and recognized by others when they are designed using UCD techniques.

For haptic perception researchers, an interesting area of research would be to determine how well the continuous pulses designed for this system help musicians keep their tempo and compare these results with discrete haptic pulses. This would provide quantitative data about how useful the continuous pulses are in helping keep tempo and having a better sense of the length of a beat.

References

Ammirante, P., Patel, A. D., & Russo, F. A. (2016). Synchronizing to auditory and tactile metronomes: a test of the auditory-motor enhancement hypothesis. *Psychonomic Bulletin & Review*, 1–9.

Bach-y Rita, P., & W. Kercel, S. (2003, December). Sensory substitution and the human–machine interface. *Trends in Cognitive Sciences*, *7*(12), 541–546.

Bolanowski, S. J., Gescheider, G. A., Verrillo, R. T., & Checkosky, C. M. (1988, November). Four channels mediate the mechanical aspects of touch. *The Journal of the Acoustical Society of America*, *84*(5), 1680–1694.

Bongers, B. (1997, January). Tactile display in electronic musical instruments. *Proceedings of IEEE Colloquium Developments in Tactile Displays*, 7/1 – 7/3.

Brewster, S., & Brown, L. M. (2004). Tactons: structured tactile messages for non-visual information display. In *Proceedings of the fifth conference on australasian user interface* (pp. 15–23).

Brown, L. M., Brewster, S. A., & Purchase, H. C. (2005). A first investigation into the effectiveness of tactons. In *First joint eurohaptics conference and symposium on haptic interfaces for virtual environment and teleoperator systems. world haptics conference* (pp. 167–176).

Cadoz, C., Lisowski, L., & Florens, J.-L. (1990). A Modular Feedback Keyboard Design. *Computer Music Journal*, *14*(2), 47–51.

Chafe, C. (1993). Tactile audio feedback. In *Proceedings of the International Computer Music Conference* (pp. 76–79).

Choi, S., & Kuchenbecker, K. J. (2013). Vibrotactile display: Perception, technology, and applications. *Proceedings of the IEEE*, *101*(9), 2093–2104.

Cycling '74. (2018). *Max Software Tools for Media.* Retrieved 2018-03-14, from `https://cycling74.com/products/max`

Dix, A., Finlay, J., Abowd, G. D., & Beale, R. (2003). *Human-computer interaction.* Harlow, England; New York: Pearson/Prentice-Hall.

Egloff, D., Braasch, J., Robinson, P., Van Nort, D., & Krueger, T. (2011, April). A vibrotactile music system based on sensory substitution. *The Journal of the Acoustical Society of America*, *129*(4), 2582–2582.

El Saddik, A. (2007). The potential of haptics technologies. *IEEE Instrumentation & Measurement Magazine*, *10*(1), 10–17.

Erp, J. B. F. V., Veen, H. A. H. C. V., Jansen, C., & Dobbins, T. (2005, April). Waypoint Navigation with a Vibrotactile Waist Belt. *ACM Trans. Appl. Percept.*, *2*(2), 106–117.

Frid, E., Giordano, M., Schumacher, M. M., & Wanderley, M. M. (2014). Physical and perceptual characterization of a tactile display for a live-electronics notification system. In *Proceedings of the Joint International Computer Music Conference (ICMC) and Sound and Music Computing Conference (SMC)* (pp. 954–961).

Gallace, A., Tan, H. Z., & Spence, C. (2007). The body surface as a communication system: The state of the art after 50 years. *Presence: Teleoperators and Virtual Environments*, *16*(6), 655–676.

Gillespie, B. (1999). Haptics. In P. R. Cook (Ed.), *Music, cognition, and computerized sound : an introduction to psychoacoustics* (pp. 229–260). Cambridge, Mass.: MIT Press.

Giordano, M. (2016). *Vibrotactile feedback and stimulation in music performance* (PhD Thesis). McGill University.

Giordano, M., Hattwick, I., Franco, I., Egloff, D., Frid, E., Lamontagne, V., … Wanderley, M. (2015). Design and implementation of a whole-body haptic suit for “ilinx”, a multisensory art installation. In *12th international conference on sound and music computing (smc-15)* (Vol. 1, pp. 169–175).

Giordano, M., & Wanderley, M. M. (2013). Perceptual and technological issues in the design of vibrotactile-augmented interfaces for music technology and media. In *International workshop on haptic and audio interaction design* (pp. 89–98).

Giordano, M., & Wanderley, M. M. (2015). Follow the tactile metronome: Vibrotactile stimulation for tempo synchronization in music performance. In *Proceedings of the Sound and Music Computing Conference, Maynooth, Ireland.*

Gould, J. D., & Lewis, C. (1985, March). Designing for Usability: Key Principles and What Designers Think. *Commun. ACM*, *28*(3), 300–311.

Gunther, E., & O'Modhrain, S. (2003). Cutaneous grooves: composing for the sense of touch. *Journal of New Music Research*, *32*(4), 369–381.

Hart, S. G., & Staveland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In P. A. Hancock & N. Meshkati (Eds.), *Advances in Psychology* (Vol. 52, pp. 139–183).

Hartnett, G., Elliott, L., Baraniecki, L., Skinner, A., Riddle, K., & Pettitt, R. (2017, July). Can You Feel Me Now? Wearable Concept for Soldier Communications. In *Advances in Human Factors in Robots and Unmanned Systems* (pp. 236–247).

Hattwick, I., Franco, I., & Wanderley, M. M. (2017). The vibropixels: A scalable wireless tactile display system. In S. Yamamoto (Ed.), *Human interface and the management of information: Information, knowledge and interaction design: 19th international conference, hci international 2017, vancouver, bc, canada, july 9–14, 2017, proceedings, part i* (pp. 517–528).

Hayes, L., & Michalakos, C. (2012). Imposing a networked vibrotactile communication system for improvisational suggestion. *Organised Sound*, *17*(01), 36–44.

Hennig, H., Fleischmann, R., Fredebohm, A., Hagmayer, Y., Nagler, J., Witt, A., ... Geisel, T. (2011). The Nature and Perception of Fluctuations in Human Musical Rhythms. *PLOS ONE*, *6*(10), e26457.

Hove, S. E., & Anda, B. (2005). Experiences from conducting semi-structured interviews in empirical software engineering research. In *11th IEEE International Software Metrics Symposium (METRICS'05)* (pp. 23–32).

Huang, G., Metaxas, D., & Govindaraj, M. (2003). Feel the "Fabric": An Audio-haptic Interface. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (pp. 52–61).

International Organization for Standardization. (2009). *ISO 9241-920:2009 - Ergonomics of human-system interaction – Part 920: Guidance on tactile and haptic interactions.* Retrieved 2018-06-13, from `https://www.iso.org/standard/42904.html`

Karam, M., Branje, C., Nespoli, G., Thompson, N., Russo, F. A., & Fels, D. I. (2010). The emoti-chair: an interactive tactile music exhibit. In *Chi'10 extended abstracts on human factors in computing systems* (pp. 3069–3074).

Kimura, M. (1995). Performance Practice in Computer Music. *Computer Music Journal*, *19*(1), 64–75.

Lamontagne, V., Martinucci, M., Salter, C., Hattwick, I., Franco, I., Giordano, M., ... W, M. M. (2015). The Ilinx Garment: Whole-body tactile experience in a multisensorial art installation. In *in Proceedings of the International Symposium on Electronic Arts.*

MacLean, K. E. (2008). Haptic interaction design for everyday interfaces. *Reviews of Human Factors and Ergonomics*, *4*(1), 149–194.

Marshall, M. T., & Wanderley, M. M. (2006). Vibrotactile feedback in digital musical instruments. In *Proceedings of the 2006 conference on new interfaces for musical expression* (pp. 226–229).

Mcdonald, K., Kouttron, D., Bahn, C., Braasch, J., & Oliveros, P. (2009). The Vibrobyte: A Haptic Interface for Co-Located Performance. In *Proceedings of the New Interfaces for Musical Expression Conference (NIME)* (pp. 41–42).

Michailidis, T., & Berweck, S. (2011). Tactile feedback tool: approaching the foot pedal problem in live electronic music. In *Proceedings of the International Computer Music Conference (ICMC)* (pp. 661 – 664).

Michailidis, T., & Bullock, J. (2011). Improving performers' musicality through live interaction with haptic feedback: a case study. *Proceedings of the Sound and Music Computing Conferece (SMC).*

Miranda, E. R., & Wanderley, M. M. (2006). *New digital musical instruments: control and interaction beyond the keyboard.* AR Editions, Inc.

O'Modhrain, S., & Chafe, C. (2000). Incorporating haptic feedback into interfaces for music applications. In *Proceedings of the International Symposium on Robotics with Applications, World Automation Conference.*

Preece, J., Rogers, Y., Sharp, H., Benyon, D., Holland, S., & Carey, T. (1994). *Human-computer interaction.* Wokingham, England; Reading, Mass.: Addison-Wesley Pub. Co.

Puckette, M. S., Apel, T., & Zicarelli, D. D. (1998). Real-time audio analysis tools for Pd and MSP. In *Proceedings of the 1998 International Computer Music Conference.* (pp. 109–112).

Robertson, A., & Plumbley, M. D. (2013). Synchronizing sequencing software to a live drummer. *Computer Music Journal*, *37*(2), 46–60.

Rovan, J., & Hayward, V. (2000). Typology of tactile sounds and their synthesis in gesture-driven computer music performance. *Trends in gestural control of music*, 297–320.

Schumacher, M., Giordano, M., Wanderley, M. M., & Ferguson, S. (2013). Vibrotactile notification for live electronics performance: A prototype system. In *Proceedings of the International Symposium on Computer Music Multidisciplinary Research (CMMR)* (pp. 516–525).

Sinclair, S., Scavone, G. P., & Wanderley, M. M. (2009). Audio-Haptic Interaction with the Digital Waveguide Bowed String. In *International Computer Music Conference (ICMC)* (pp. 275–278).

Sinclair, S., Wanderley, M. M., Hayward, V., & Scavone, G. (2011). Noise-free haptic interaction with a bowed-string acoustic model. In *IEEE World Haptics Conference (WHC)* (pp. 463–468).

Soundbrenner. (2017). *Soundbrenner Press Kit.* Retrieved 2017-12-11, from `http://press.soundbrenner.com/photos`

Verrillo, R. T. (1992). Vibration Sensation in Humans. *Music Perception: An Interdisciplinary Journal*, *9*(3), 281–302.

Visell, Y., Law, A., & Cooperstock, J. R. (2009, July). Touch Is Everywhere: Floor Surfaces as Ambient Haptic Interfaces. *IEEE Transactions on Haptics*, *2*(3), 148–159.

Vredenburg, K., Mao, J.-Y., Smith, P. W., & Carey, T. (2002). A Survey of User-centered Design Practice. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 471–478).