TOUCHABLE: ADAPTING A HAPTIC FEEDBACK GLOVE FOR USE IN REHABILITATION CONTEXTS

A THESIS SUBMITTED TO AUCKLAND UNIVERSITY OF TECHNOLOGY
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF CREATIVE TECHNOLOGIES

Supervisors
Dr. Stefan Marks
Assoc. Prof. Frances Joseph

June 2016

By
Jacques Foottit
Colab: Creative Technologies

Copyright

Declaration

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the qualification of any other degree or diploma of a university or other institution of higher learning.

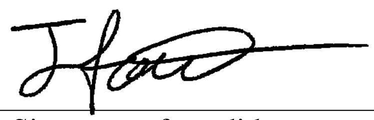

Signature of candidate

Acknowledgements

I would like to acknowledge and thank all those who have supported me during this research journey.

To my supervisors Stefan Marks and Frances Joseph, thank you for your ongoing advice and support.

To my parents, thank you for all your support during the past year.

I would like to acknowledge the Textile Design Lab, and in particular Gordon Fraser. This project could not have been accomplished without your technical expertise.

This research was performed with ethics approval from AUTEC. Application Number 15/328, approved on 7 September 2015.

Abstract

With the increasing miniaturisation of computing and sensor technology, it is becoming common for electronics of all kinds to be integrated into clothing and other wearable items. Motion sensing technologies in particular have been used for a variety of consumer fitness and virtual reality applications for able-bodied people. This research explores the potential for affordable motion capture and haptic feedback technologies to be utilised in a rehabilitation context, with a specific focus on the hand.

An iterative development process was used to adapt and improve an existing prototype haptic feedback glove in response to the unique challenges facing wearable device users in a rehabilitation context. Collaboration with physiotherapists provided valuable feedback throughout the design process. The result is a significantly different prototype device with major design improvements, and insights into how iterative development processes can be utilised for hardware development.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The hand is a fascinating and wondrous piece of human anatomy. It both enables us to change the physical environment around us with incredible freedom, and at the same time defines the manner in which we have shaped our world. Perhaps more than any other part of the human body, our hands represent the human ability to leave a mark on our reality. Simultaneously, it is our sense of touch that gives us a continuous sense of our presence in the world.

Often we are unaware of how significant our sense of touch is, until we consider what happens when it is lost (Robles-De-La-Torre, 2006). It is essential for us to experience and manipulate the world (Augurelle, Smith, Lejeune & Thonnard, 2003), and even influences what we consider part of our body (Tan et al., 2014; Marasco, Kim, Colgate, Peshkin & Kuiken, 2011). Yet despite the entirely new forms of experience computers have enabled, touch remains as one of the last bastions of reality lacking in many of the virtual representations offered by modern technology.

For some this remarkable sense that we often take for granted is absent or limited not in virtual worlds but in the physical world. It is difficult to imagine being able to see but not touch, hear but not feel. Yet this is the reality for many who have experienced traumatic brain or spinal injury. A large body of work has been dedicated to helping those affected to regain control of their movement, granting them the ability to once again interact with the world around them. Yet often sensory therapy is implied but not targeted by contemporary strategies (Bolognini, Russo & Edwards, 2016, Preprint). Introducing a focus on sensory therapy or stimulation may support improved motor rehabilitation outcomes. After all, what is it to move one's hand if you cannot feel that which you are reaching for?

In my final year of undergraduate study, I explored the interaction between human and computer. I was interested in how touch could be brought to the virtual realities modern technology is presenting. In a sense it was about making a virtual reality even more real by giving it tangible substance. The result was a haptic feedback glove which was utilised as a game controller for a flight simulator (Foottit, Brown, Marks & Connor, 2016). Throughout the project I became fascinated with the hand, how we use it, and how as a child it could be a physical manifestation of an imaginary object. It brought up memories of flying around as an imaginary spacecraft, with my hand being the physical representation. My hand made the imaginary become real for me. The haptic feedback glove project explored that idea of bringing reality to the imaginary with the hand as a mediator.

At the conclusion of my final year and the haptic feedback glove project, I thought about those who are without a full sense of touch. I thought about the disconnection I feel towards a virtual reality where I can only see and not move or touch, and wondered if I would feel that same disconnection towards the physical world if I could not move or touch with my own body. I also thought about how so much of today's technology is built for the able-bodied consumer. How someone with limited hand mobility could never wear my haptic feedback glove. I wondered if this thing I had created could be re-purposed, so that instead of bringing reality to the virtual, it could restore a connection with physical reality.

A recent survey of technologies used to measure physiological processes for rehabilitation purposes identified virtual reality as having potential to support motor rehabilitation by allowing a patient to perform exercises in a way that would be easier than interacting with real objects (H. Huang, Wolf & He, 2006). The authors also identified a range of other applications for technology in the rehabilitation context, many of which are yet to be rigorously tested in order to identify how or if they can provide measurable improvement for a patient. This has been the drive for me to develop the device presented in this research. Utilising my previous work as a starting point, I have sought to refine and improve it with the goal of working towards something that may be useful in the area of rehabilitation.

Being an able-bodied person myself, I have relied on the input from physiotherapists to help me understand the challenges associated with creating a wearable device for someone who has limited hand mobility. It has taken me on a journey of discovery. Setting out I had little idea of where I would end up, and so I have utilised design processes that I have come to rely on when exploring the unknown. I have relied on experimentation and incremental improvement. I have investigated a range of different disciplines and technologies as problems arose, rather than attempting to predictively research all the potentially relevant technologies. Sometimes I have relied on my past experience building the previous haptic feedback glove. Other times, I have sought inspiration from the work of others, or advice from experts in the various disciplines this project crosses. In the end, I have created something that differs from any of the ideas I had at the start of this process. I present here one part of the journey that I have undertaken. The outcome of this research is as much a beginning as it is an end. It represents the end of one period in an ongoing process of discovery and innovation.

In order to gain insight into how wearable technologies might be utilised in a rehabilitation context, it is useful to explore the technologies relevant to that area, and how they have been utilised by others in the field. Chapter 2 identifies technologies relevant to this research, and critically assesses how they have been utilised in wearable devices with a similar focus. The methodological approach for this research is then laid out in Chapter 3, followed by a detailed description of the design process in Chapter 4. The final prototype output from the design process is then described in Chapter 5, followed by a reflection on the design process and resulting outcome in Chapter 6. Chapter 7 concludes the research and a looks at areas that may be explored in future works.

Chapter 2

Wearable Technologies for Rehabilitation-Focused Biofeedback

The range of technologies that can potentially be integrated into modern wearable devices is vast and constantly evolving. Understanding the specific goals for a prospective wearable device and investigating the technologies being used in similar devices is therefore of vital importance throughout the design process. This chapter reviews a range of technologies that are capable of providing rehabilitation-focused biofeedback. Section 2.1 provides a basis for the use of biofeedback in the rehabilitation process, and briefly summarises the current understanding on how biofeedback can be most effectively utilised. Sections 2.2 and 2.3 then provide an overview of wearable sensing and haptic feedback technologies respectively, as they are particularly relevant for the development of biofeedback systems oriented around the hand. Finally, Section 2.4 explores some of the currently existing glove-based systems that have potential for use in a rehabilitation context. This provides insights into how particular technologies have been integrated into wearable devices while critically assessing their potential benefits and limitations.

2.1 Task Oriented Biofeedback

H. Huang et al. (2006) provide an overview of the history and implications of task oriented biofeedback. The authors explain that biofeedback has been investigated extensively as a means of supporting and improving rehabilitation processes. While early explorations utilised static biofeedback with limited evidence of benefit and contradictory results, it was soon discovered that dynamic, task oriented biofeedback was much more effective. However this also came with a new set of challenges. One particular challenge identified by the authors is overloading a user with too much information, since biofeedback systems quickly began to include many sensors that each provided a stream of real time data. This led to the development of a technique to deal with information overload known as sensor fusion. The technique involves combining data from multiple sensors to produce a unified output that is relevant and intuitive for the user.

Preventing information overload in modern biofeedback systems with many sensors means that sensor fusion has become a necessity. Combining readings from multiple sensors to produce an intuitive and responsive output to the user requires careful design and attention to the specific goals or tasks the user is attempting to achieve. H. Huang et al. (2006) present a model for modern biofeedback systems where data is acquired from multiple sensors, processed, and then analysed by a central controller with reference to a database of biofeedback rules. The central controller then produces a multimodal biofeedback cue to the user based on the combined sensor data and how it relates to the specific task to be performed as defined by the biofeedback rules.

Use of multimodal cues for biofeedback systems allows a variety of information to be conveyed to the user in a way that is effective. The choice of modality for these cues is driven by the type of information that is to be communicated. This in turn is determined by the task to be performed. For example, Chen et al. (2006) present a multimodal system that incorporates both visual and auditory cues, in order to effectively communicate spatial and temporal information respectively. The choice of modalities was based on knowledge that visual cues are effective at portraying spatial information whereas auditory cues are better at portraying temporal information.

Developments in Virtual Reality technologies present new opportunities for biofeedback by providing flexible, multimodal environments that can accommodate a range of tasks. Creating a biofeedback system that is easily adapted for a range of specific rehabilitation tasks increases the potential benefits of the system by allowing it to be utilised throughout an ongoing recovery process. It is also more likely that a general purpose biofeedback system could be commercialised, due to the increased customer base. Glove based systems already reveal this trend, with general purpose versions being commercialised while special or single purpose gloves have been limited to short term use in research (Dipietro, Sabatini & Dario, 2008).

Despite the benefits a general purpose biofeedback system may offer, it is important that the design does not ignore specific target uses. A device that makes too many compromises for the sake of being a general purpose system may end up being unable to be useful for any application. As a result, careful attention must be paid during the design process to ensure the needs of multiple users and potential use cases are met. This can further complicate the design process and result in conflicting requirements. Resolving these challenges requires investigating a range of potential technologies to identify their relative strengths and weaknesses. The following two chapters summarise technologies that are relevant to the production of a biofeedback glove system. Section 2.2 covers technologies that can be used to sense motion in the hand. Section 2.3 explores technologies that can be used to provide haptic feedback.

2.2 Wearable Sensing Technologies

A range of technologies exist for measuring physiological processes, many of which can be integrated into a wearable device. For the purposes of this research, the relevant technologies are those capable of measuring motion of the hand and wrist. This can be further defined as those technologies capable of measuring either orientation or joint movement. In this area there are a few key categories of sensor. Section 2.2.1 covers sensors that utilise the changing physical attributes of a material as it is deformed in order to measure the deformation. Section 2.2.2 covers sensors that measure myoelectric signals in order to identify muscle contraction and then extrapolates from this data to determine physiological deformations. Section 2.2.3 covers sensors that utilise physical properties such as gravity or magnetic fields in order to measure the orientation of a sensor relative to a fixed reference point such as the Earth. One final category utilising light to track object positions from a distance is covered in Section 2.2.4.

2.2.1 Physical Deformation Sensors

Sensors that utilise changing physical properties under deformation can be used to measure joint angles by attaching the sensor securely over the joint. Movement of the joint will then produce a deformation in the sensor that can be measured. Examples include fibre optic sensors, resistive bend sensors, and piezoelectric bend sensors. Optic fibre sensors shine a light through a fibre optic strand that is designed to leak light as it is deformed. By measuring the intensity of light transmitted through the fibre optic strand, it is possible to determine the amount of deformation. Resistive and piezoelectric bend sensors utilise the changing electrical properties of a conductive material as it is deformed. Resistive bend sensors exhibit a change in resistance under deformation, whereas piezoelectric bend sensors utilise the ability for some materials to produce a voltage when deformed to measure the deformation.

Despite the different physical properties utilised, all deformation sensors share similar challenges when used to measure joint angles. Firstly, the sensor must be securely attached to the joint to be measured, in order to ensure joint deformation translates directly to deformation in the sensor. These sensors also have only 1 degree of freedom, so must be mounted in alignment with the freedom of motion in the joint. If a joint has more than one degree of freedom, multiple sensors must be used. In addition, each time the sensor is attached to a joint it must be calibrated so that readings from the sensor can be translated into corresponding joint deformations. Sensors of this type are also prone to providing output that is not linear or repeatable. If the sensor does not provide linear response, then a proportional change in deformation does not produce a proportional change in sensor output. This can complicate calibration procedures and decrease sensor accuracy. If the output is not repeatable, then applying the same deformation to the sensor will result in a different output. This also reduces the accuracy of the data supplied by the sensor. A key focus for development in this area has been producing sensors that exhibit high linearity and repeatability, with some excellent results (Latessa, Brunetti, Reale, Saggio & Di Carlo, 2009). However often commercially available sensors in this category that exhibit high amounts of linearity and repeatability are also significantly more expensive.

Physical deformation sensors have seen widespread use in data gloves because of their slim and lightweight form factor, low power consumption, and potential for low cost. Cheap piezo-resistive flex sensors have been used in a number of low cost data gloves both in research applications and consumer devices. An array of piezo-resistive flex sensors has also been developed for integration into data gloves (Saggio, 2014). CyberGlove Systems utilise proprietary piezo-resistive flex sensors for their range of data gloves (‘Home’, n.d.). Optical fibre flex sensors are used by 5DT in their range of gloves (‘Data Gloves | 5DT’, n.d.), including a glove optimised for use in MRI environments. The commercial offerings are typically very expensive, while the low cost gloves developed have been limited to research labs and some ventures into consumer gaming markets.

2.2.2 Myoelectric Sensors

Sensors to measure electric signals were some of the earliest wearable sensing devices developed. However recently this technology has been utilised in a unique way. Myoelectric sensors work by measuring electrical signals in muscles to identify muscle contractions. This information can then be used to extrapolate physiological deformation with the use of an appropriate model to map muscle contraction to joint movement. The advantage of these sensors is that the sensor does not need to be attached to the joint under deformation. Instead, it is placed on the muscle group that controls the movement of the joint. Unfortunately this approach also presents a range of new challenges and limitations.

One of the key limitations of myoelectric sensors is that the accuracy of joint deformation is limited by the ability of the system to discriminate between different muscles. To get truly accurate readings, it is likely that electrodes would need to be surgically implanted in muscle groups. Non-invasive sensors that are worn on the arm therefore have limited accuracy. This is further complicated by the requirement for arm-worn sensors to be positioned precisely in order to pick up the relevant signals. This presents challenges to the designer of such a system and may result in tedious calibration requirements or limit the potential users. It is also unclear how effective this sensor would be in a rehabilitation context where users are likely to have abnormal myoelectric signals to the muscles due to injury.

2.2.3 Orientation Sensors

Orientation sensors utilise physical forces such as gravity and the Earth's magnetic field to determine their orientation. The introduction of microelectromechanical system (MEMS) devices has had a great impact in this category of sensor. MEMS devices allow an extremely small mechanical structure to be used in order to measure acceleration, angular velocity, and magnetic fields. In addition to MEMS devices there are a range of other orientation sensors that utilise various mechanisms. Some perform a similar function to MEMS devices but are much larger, heavier and use more power in order to provide more accurate data. Others are simpler and can have a specific purpose. For example a tilt sensor can be a simple switch that is triggered when the sensor is tilted past a certain angle with respect to the ground. The small size, low power consumption and light weight of MEMS devices makes them ideal for portable applications where orientation information is required. In addition MEMS devices are relatively low cost due to their mass production capability and wide adoption in consumer electronics such as smart phones.

An inertial measurement unit (IMU) is a specific type of orientation sensor that incorporates an accelerometer, gyroscope, and in some cases even a magnetometer into a single package. Often the included sensors themselves take measurements on three axes, resulting in up to nine streams of data from a single sensor package. However this data does not directly provide orientation. It must be processed and sensor fusion applied in order to determine the orientation. In addition, when using these sensors for wearable applications they must be attached securely to the part that is to be measured as any movement between the sensor and body to be measured will result in measurement error. Sensor drift is also a common problem with MEMS based IMUs (Cavallo et al., 2014). This phenomenon is a result of accumulating errors in the sensor readings, causing a measured orientation to gradually drift away from the true value.

2.2.4 Vision Based Sensors

Vision based sensors utilise one or more cameras to record images of the object to be tracked, and then apply computer vision algorithms to process the data and output position information about the tracked object. While standard cameras can be used, more recently cameras that capture depth information have been utilised to great effect (Xia, Chen & Aggarwal, 2011). The principle of computer vision algorithms tend to be the same: identify a point of interest and determine the position. One of the most well known vision based systems is the Microsoft Kinect (Microsoft, n.d.), which uses an infra-red depth camera to identify key parts of a person's body and tracks them in three dimensions. More advanced systems such as those produced by the company OptiTrack can use multiple cameras positioned around a room to track a person wearing special markers on their body with sub-millimetre accuracy ('OptiTrack for Virtual Reality', n.d.). While these systems can provide robust tracking solutions, they cannot be considered wearable due to their requirement for cameras to be positioned around a space. The concepts used in these systems has however been explored in wearable devices, for example by using a body worn camera (Mistry, Maes & Chang, 2009; Harrison, Benko & Wilson, 2011).

The main limitation of vision based systems is that they require line of sight to a tracked target in order to work. This is often difficult to achieve in a wearable context, particularly when attempting to measure hand motions. Positioning a camera so that it can see the range of motion of the hand and fingers while still being attached to the body and not limiting mobility is a key challenge for this kind of system.

2.3 Haptic Feedback

Haptic feedback covers a range of devices that can stimulate the sense of touch. They range from simple vibration motors to highly complex mechanical systems that are able to produce realistic forces simulating physical objects. Haptic feedback systems can be roughly divided into two categories: force reflecting and vibration.

2.3.1 Force Reflecting Haptic Feedback

Force reflecting haptic feedback systems produce forces that simulate those produced by real world objects. These systems apply a force against a user that is proportional to the force the user applies to the system, in a sense ’reflecting’ the force applied. In order to achieve this force reflecting behaviour, systems often consist of complex combinations of linkages and actuators. One example of such a system, the CyberGlove Systems Haptic Workstation (‘Haptic Workstation’, n.d.), is shown in Figure 1. Sometimes the actuators used are able to act as sensors for the system to measure the force applied, other times additional sensing systems need to be combined with the feedback interface.

Figure 1. The Haptic Workstation by CyberGlove Systems. One example of a force reflecting haptic interface. Note the complex mechanical systems involved.

The complexity of mechanical systems and power required for driving force reflecting interfaces presents significant challenges for those seeking to design portable and wearable interfaces of this type. The benefit of such a system is that it can not only simulate real world forces, but also is often capable of assisting human movement. This enables such systems to be used in motor rehabilitation for robotic assisted therapies (Heo, Gu, Lee, Rhee & Kim, 2012), or to amplify the force applied by an able-bodied person for tasks such as heavy lifting.

2.3.2 Vibration Feedback

Haptic feedback systems that utilise vibration are generally much simpler than force reflecting interfaces, and are used in a wide range of applications such as mobile phone notifications or communicating facial expressions to visually impaired people (Krishna, Bala, McDaniel, McGuire & Panchanathan, 2010). Vibration is often used as a notification system, but can be used for a variety of purposes, such as conveying temporal information or passive learning of physical skills (K. Huang, Do & Starner, 2008). There are a range of actuators and actuator designs utilised to produce vibration for these systems. Commonly used actuators in vibration based haptic feedback systems include eccentric rotating mass (ERM) motors, linear resonance actuators (LRAs), and ultrasonic transducers.

ERM motors are a simple direct current (DC) motor with a mass attached to the shaft off-centre. When the motor is powered, it spins the mass around and the changing centre of gravity of the device causes it to vibrate. ERM motors come in a range of sizes and shapes. Sometimes the entire unit is enclosed to protect the moving parts, other times the rotating mass is exposed. ERM motors are simple to drive and can produce a range of frequencies based on the rotational speed of the motor.

LRAs also use the principle of a moving mass to produce vibration, but instead of rotating around a central shaft the mass is electromagnetically driven in a linear motion. The mass in an LRA is attached to springs so that it resonates at a particular frequency. The LRA is then driven with an alternating current (AC) waveform that matches the resonant frequency of the mass-spring system in order to produce vibrations in a very efficient manner. LRAs often require more complex driving circuitry since the input waveform must match the resonance of the internal mass-spring system. In addition, LRAs produce a very narrow band of frequencies based on the resonant frequency of the device.

Ultrasonic transducers are speakers that are designed to produce sound at a frequency that is above human hearing. When multiple transducers are combined, the changes in air pressure produced can be used to modulate the texture of a hard surface (Watanabe & Fukui, 1995; Iwamoto, Akaho & Shinoda, 2004), or create points of high pressure air within a three-dimensional space which can be felt by the hand (Hoshi, Iwamoto & Shinoda, 2009). Such systems tend to be fairly complex in comparison to other vibration systems and have limited application in wearable electronics.

Despite the comparative simplicity of vibration feedback systems, it is still important to consider how any form of haptic feedback is integrated into wearable devices. Positioning the same vibration elements in different locations on the body can provide distinctly varied feedback to the user. This is also true for wearable sensing technologies. In the following section, I will explore how various technologies have been used in devices that are relevant to the rehabilitation context.

2.4 Glove-Based Systems in Rehabilitation

This section gives a brief overview of how different technologies are being used in rehabilitation contexts, and identifies common themes and limitations. The focus is on glove based systems that have been or could be used for rehabilitation purposes. Commercially available systems are covered first, followed by some that have not been commercially released.

Glove based systems have a long history of development and an ever increasing range of application areas. Dipietro et al. (2008) defines a glove-based system as:

> a system composed of an array of sensors, electronics for data acquisition/processing and power supply, and a support for the sensors that can be worn on the user's hand.

While this definition is useful in that it is not restricted to the literal interpretation of glove, it does not allow for devices that provide output to the user's hand. For the purposes of this research, the definition is extended to include such devices.

2.4.1 Commercial Systems

A number of glove based systems are commercially available, with varied capabilities, sensor types, costs, and applications. This section covers five such systems to give a brief overview of the current technologies employed by commercial solutions. The systems covered are the CyberGlove by CyberGlove Systems ('Home', n.d.), the 5DT glove by Fifth Dimension Technologies ('Data Gloves | 5DT', n.d.), the ControlVR system ('Control VR- The Future of Virtual Reality, Animation & more', 2014), the Gloveone system ('Gloveone', n.d.) and the GyroGlove ('Gyroglove | GyroGear', n.d.).

CyberGlove and 5DT

The technology employed by CyberGlove systems for their data glove products is based on proprietary resistive flex sensors ('CyberGlove II', n.d.). These flex sensors are mounted in gloves to capture joint angle data that is then used to control a 'virtual skeleton', providing the user with a virtual representation of their hand movements. The glove itself does not track hand position but rather is limited to finger and wrist flexion. To support applications where hand position is important, the glove allows additional third party motion trackers to be attached. These can be used to provide a more complete motion tracking solution. The number of sensors in the glove can also be varied based on the desired level of detail. More flex sensors means more accurate representation of the hand as less data is interpolated or extrapolated. However this also increases the cost of the glove, as well as requiring more processing power and bandwidth for communication with the computer. The CyberGlove communicates using a custom

2.4GHz band transmitter with a companion receiver, similar in concept to wireless mouse and keyboard devices. The glove also features a removable, rechargeable battery and 3 hours battery life.

The CyberGlove has been used in research both to support the rehabilitation process (A. S. Merians et al., 2002) and as a tool for studying hand movement (Antonin Viau, 2005). It is generally considered to be the most accurate data glove on the market but is also extremely expensive. While prices are not generally published for these gloves, historically they have been priced upwards of USD$10,000 (Karen Moltenbrey, 2006). The high cost of the glove limits potential applications in clinical practice but has made it useful for research areas that require very accurate measurements. While haptic feedback is not integrated into the glove, CyberGlove offers a range of optional haptic feedback add ons.

The 5DT glove range by Fifth Dimension Technologies ('Data Gloves | 5DT', n.d.) utilises optical flex sensors to measure joint angles in a similar way to the flex sensors used in the CyberGlove. The company offers two versions of the glove, which use either 5 or 14 optical flex sensors. Using 14 sensors allows the more expensive glove model to measure 2 joints for each finger and abduction between the fingers. One of the unique offerings from Fifth Dimension Technologies is a glove specifically designed for use in magnetic resonance imaging (MRI) environments (Dipietro et al., 2008). This is achieved by extending the optical fibres from the sensors to a control box that can be up to 7 m away, allowing the glove to be free of any magnetic materials ('5DT Products', n.d.). The ability to operate in an MRI environment enables the glove to be used for capturing hand movement and brain activity simultaneously.

ControlVR and Gloveone

The ControlVR system was first launched as a kickstarter campaign (running from 6th June to 9th July 2014), as an affordable motion capture device for game studios as well as a virtual reality controller for consumers ('Control VR- The Future of Virtual Reality, Animation & more', 2014). By utilising IMUs to measure the orientation of fingers, hands, arms and the chest, it was possible for the system to provide a complete upper body motion capture solution. The ControlVR system was unique in that it was targeted at the consumer market with a low price tag to match. Unfortunately despite successfully achieving their kickstarter campaign funding goal, the product never reached the market.

Gloveone ('Gloveone', n.d.) is a more recent kickstarter funded project (running from 2nd June to 10th July 2015) that has some key similarities to ControlVR. It also utilises IMUs to measure finger flex and hand orientation. Unlike ControlVR it is limited to the hands, so is unable to track hand positions. It also integrates haptic feedback in the fingers and palm, a feature that was not included in ControlVR. The Gloveone is built on a textile glove base, similar to other commercial data gloves available. However it also utilises conductive fabric in the textile design to create conductive zones for command triggering.

GyroGlove

One final upcoming commercial product of note is the GyroGlove. This glove was developed to support patients with Parkinsons by reducing hand tremors (Parkin, 2016). It utilises a custom-designed mechanical gyroscope to dampen hand movement. The glove also records motion data while it is in use, to provide medical professionals with additional data that could help monitor a patient's condition. While the overall objective of the glove is very specific, the concept of recording data at home to support medical professionals in monitoring a condition is something that can be applied to a range of devices.

2.4.2 Other Systems

Systems that have not been released commercially range from simple, low cost devices to highly sophisticated systems that could compete with commercial offerings. In many cases the systems presented are created in response to one or more limitations of the commercially available options. Nathan and Johnson (2007) list a few limitations of commercial gloves that can make them unsuitable for therapy applications. These include: the difficulty in putting on and taking off the commercial devices when a person has limited hand mobility[1], intricate and time-consuming calibration processes, risks of over-extension with mechanical haptic feedback devices, and an inability to integrate the devices with other systems or to record data due to closed source software. The authors also point out that limited sizing options can be an issue with commercial devices, as patients in a rehabilitation context can have unique sizing requirements.[2]

Another limitation identified with commercial offerings that has driven the development of alternative solutions is the high costs involved. A range of low-cost glove systems have been developed for a variety of purposes in response to the high costs associated with commercial systems. One example is presented by Vutinuntakasame, Jaijongrak and Thiemjarus (2011), where a low-cost glove is used to allow a form of sign language to act as an input to a voice box. The glove presented in this research contains only 5 sensors yet is able to accurately identify a wide range of hand gestures. This highlights the fact that even a simple and affordable device can perform a valuable and beneficial function.

At the other end of the scale, sometimes a specific application requires greater capability than what is offered commercially. O'Flynn et al. (2013) developed a glove system with a large number and variety of sensors to allow for highly detailed,

[1]The existing prototype used as a starting point for this research also suffered from this limitation, and responding to this challenge was a primary focus during the development process.

[2]Sizing also emerged as a potential limitation during my development process, and became a key consideration for evaluating the textile elements of the design.

calibration free measurement of the hand to support arthritis rehabilitation. It utilises a modern flexible PCB and surface mount devices (SMDs) to fit a large amount of complexity into a package small enough to fit neatly on a persons hand. This presents a step forward in high quality data capture of the hand.

Integrating more sensors into a smaller space is not the only way that data gloves are moving forward. Different approaches are being explored that utilise new sensor technologies such as IMUs (Moreira et al., 2014), and even utilising new advances in e-textiles to integrate sensors directly into the textile structure of a glove (Åkerfeldt, Lund & Walkenström, 2015). These new approaches reveal an ever expanding range of possibilities for future data glove designs.

This chapter has explored a range of wearable sensing and haptic feedback technologies, and briefly reviewed how they have been integrated into devices relevant to task oriented biofeedback. This provides a foundation for the development of wearable devices within the rehabilitation context. Chapter 3 details the methodological approach for this research.

Chapter 3

Methodological Approach

As the aim of this research was to explore new ways in which technology can be utilised in the rehabilitation context, it was likely that the design process would need to be flexible and adaptable to change. This would be important because of the difficulty in predicting if a given design concept would be suitable without first physically constructing a prototype and testing it. The benefit of a flexible process that incorporates physical prototyping is recognised by the fashion industry, where from the start a designer would work with physical materials to create prototypes. However this project would involve a combination of textile elements and electronic systems, including both hardware and software development. Any method used would therefore need to be suitable for a range of traditionally separate systems. One promising methodological framework comes from the area of software development, and is known as agile.

In the area of software development, agile methods have been utilised so that a project can adapt to meet changing targets. While agile itself has no single or exact definition beyond a simple manifesto, various methods have arisen that share the same philosophy and call themselves agile (Kaisti et al., 2013). Different methods vary based on emphasis of particular practices that respond to specific issues in software development. As a result, situations with different requirements or potential issues may require the adoption of different methods.

The core foundation of any agile method is the Agile Manifesto, which states:

> We are uncovering better ways of developing
> software by doing it and helping others do it.
> Through this work we have come to value:
>
> *Individuals and interactions* over processes and tools
> *Working software* over comprehensive documentation
> *Customer collaboration* over contract negotiation
> *Responding to change* over following a plan
>
> That is, while there is value in the items on
> the right, we value the items on the left more.

Kent Beck	James Grenning	Robert C. Martin
Mike Beedle	Jim Highsmith	Steve Mellor
Arie van Bennekum	Andrew Hunt	Ken Schwaber
Alistair Cockburn	Ron Jeffries	Jeff Sutherland
Ward Cunningham	Jon Kern	Dave Thomas
Martin Fowler	Brian Marick	

©2001, the above authors (Kent Beck et al., 2001)

While emphasising one thing over another, the manifesto also holds that both have value. So it is about shifting focus rather than abandoning everything for a new way of doing things. The unique attribute of agile processes are incremental iterations with an integrated team rather than linear phases with different teams handling each phase.

While a large body of research has investigated agile methods in software development contexts, there has been less investigation into the use of such methods for embedded systems, especially with regard to hardware development. This is partially because embedded systems in general and hardware development particularly have unique challenges that differ from the software development fields agile arose in. This has not prevented some researchers and industry teams from attempting to develop new methods for embedded system development with an agile philosophy at heart.

Kaisti et al. (2013) provides an overview of the literature relevant to applying agile methods to embedded system development and highlights some useful considerations. One of the key aspects of embedded system development highlighted by the author is that the role of architecture and up-front designing cannot be avoided. It is also pointed out that as development continues requirements become more rigid as any changes at later stages of development will have a more significant impact on all the teams involved. Experimentation is another unavoidable element of embedded system development, since hardware constraints may impact code in unpredictable ways. Rather than make an agile method ineffective, these considerations actually align well with the core goal of agile methods - to reduce uncertainty early on by creating complete working systems. The key challenge then is how to balance flexibility with the time and financial resources associated with constructing a complete hardware system.

The relationship between scope of change and cost in software systems is fairly linear. In hardware systems, however, even a small change may require a completely new prototype to be constructed, introducing significant cost. As a result, the idea of small, incremental changes through many iterations is difficult to achieve with a complete embedded system. In order to account for this challenge, the design method for this research utilised a modular development approach. By isolating different components in the system, it would be possible to gain the benefits of incremental improvements through iterative development without the high costs associated with a complete system integration. Different components could also utilise varying iteration cycle lengths based on what would be most effective for that component. These component cycles take place within the prototyping phase of the main cycle, with the start and end of every prototyping phase involving component isolation and system integration respectively. The main cycle is illustrated in Figure 2.

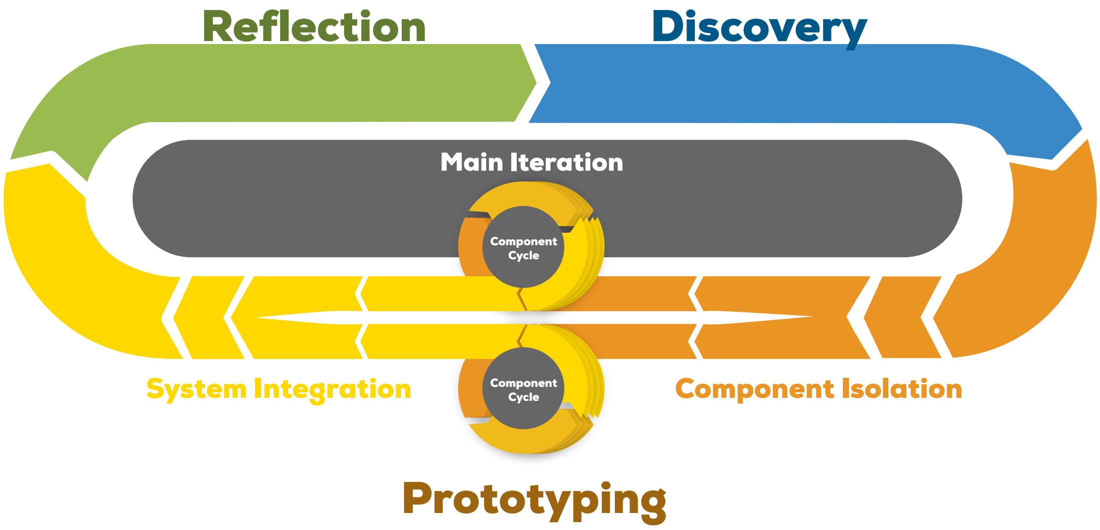

Figure 2. Methodological approach. An iterative method was used, with component isolation for modular development.

During this research, the methodology itself also adapted in response to challenges encountered. This was done through a reflective process that sought to identify what was working and what was not, and then incrementally change the process to improve it. This agile approach to the methodology itself meant that some of the specifics changed. One particular area where this occurred was around the concept of time-boxing.

A key feature of many agile methodologies is time-boxing. The idea behind this is to scale features up or down based on available time and development progress. If features are implemented more easily than anticipated, more can be added. If instead development proceeds slower than expected, features are removed to ensure an on-time release. In a commercial setting this is often an effective way of meeting deadlines for a release. It is useful because generally it is easier to predict the time available to spend on a project than it is to predict the features that could be implemented in the given time frame. Initially for this research the intent was to produce as many as three system integrations with complete functional prototypes. However as the deadline approached for the first system integration, it became evident that the project had not reached a point where it would be beneficial to have a complete prototype. At this point the time-boxing concept was reviewed and altered in two ways. Firstly, the time-box duration of the main cycle was increased significantly in order to allow for more development of individual components. Secondly, the time-boxing was loosened so that the exact deadline for system integration could shift to optimise the balance between delivery time and feature integration. This second consideration is important when there is a long main cycle length, as the delay associated with moving a feature to the next cycle is more significant. If a particular feature is almost complete, it becomes preferable to delay the system integration slightly in order to incorporate the new feature rather than drop it and wait for an entire cycle before it is integrated.

For this research, a single system integration cycle was performed, although multiple iterations were completed for each component of the system. The development cycle covered in this research also continues on from previous work, and at the conclusion of this research another development cycle could begin. Objectives for the next development cycle are noted in Section 7.1.

The main development cycle went through three key phases. These phases were Discovery, Prototyping, and Reflection. The following sections detail the focus of each of these three phases.

3.1 Development Phase 1: Discovery

The primary goal of the Discovery phase was to identify areas for improvement in the current cycle. It had a particular focus on interaction with key stakeholders in the project, and was used to provide scope and direction for the development cycle. User stories were utilised as a tool for facilitating the process.

User stories are a tool commonly utilised in agile methodologies (Cohn, 2004). They are designed to fulfil a similar purpose to system requirements, but promote greater interaction with the users. Instead of setting concrete requirements, a user story provides a description of what a user wants to do with the system. It encourages interaction with the user by removing the specification of how a given task will be achieved. As a result the designer must continue to interact with the user rather than simply check off predefined capabilities.

For this research, a group of physiotherapists at AUT were involved as the user group. While they were not the intended primary end user, they were able to provide expert opinion on what the needs of users with limited hand mobility would be. In addition, they would be the people most likely to put the device on someone, and also could provide expert opinion on what would be required in order for the device to be adopted by the rehabilitation community. Future development cycles would ideally begin to include patients themselves, however for this research the scope was limited to physiotherapists. This was primarily done to ensure the research remained achievable within the time frame.

The outcome of the Discovery phase was a set of user stories that described what the device should be able to do and what some of the key considerations would be when designing how it would be worn. The user stories were also prioritised to facilitate scaling of features during the prototyping phase.

3.2 Development Phase 2: Prototyping

The Prototyping phase began with the identification and prioritisation of key components that needed to be developed. This process involved considering at an abstract level what the system would consist of and isolating the various components that would combine to create this system. Prioritisation of the components was performed by identifying the priority of any user stories that related to a given component. In addition, consideration was given to component dependencies. So if a component needed to be completed in order for development to proceed on other components, it was given a higher priority.

During the prototyping phase identified components were developed individually through iterative processes, with input from physiotherapists occurring at key points in the cycles. The prototyping phase concluded with a system integration process that combined all the component cycle results into a single final prototype. The system integration process itself was also iterative, and can be viewed as another component cycle.

3.2.1 Component Cycles

The component cycles are a key concept for this research method. They formed the basis for providing flexibility and adaptability during the development process. Sufficiently isolated component cycles could also proceed concurrently, reducing down time when one cycle was waiting on delivery of parts or some other external process to complete. While each component was isolated to minimise interactions between them, some dependencies were unavoidable, potentially limiting the ability for affected cycles to proceed concurrently.

Figure 3 depicts a component cycle with multiple iterations. It began with a specification process that identified key requirements and scope for the component. Once this was completed the iteration cycle could commence. Each cycle went through an ideation, testing, and review phase. Once the cycles produced an output that satisfactorily met the specified requirements this became a final output for the cycle which could be utilised for the system integration. Each of the three phases of the component cycles is explained in the following sections.

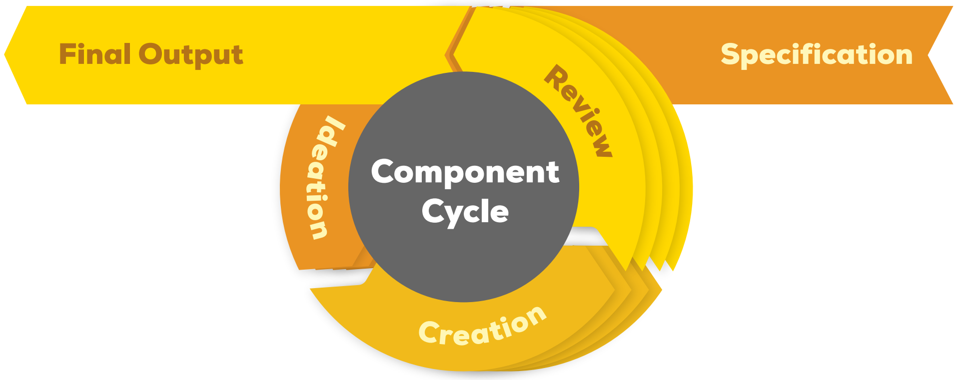

Figure 3. Component iteration cycle. This illustrates the cycle for a single component.

Component Cycle Phase 1: Ideation

The ideation phase involved coming up with ideas based on identified requirements or goals for the component. One example would be if a textile component required more flexibility, then the ideation phase would produce multiple ways of increasing flexibility.

Component Cycle Phase 2: Creation

Following the ideation phase, the creation phase involved rapidly prototyping the most promising ideas from the ideation phase. This would result in one or more physical outcomes that could be analysed to see if they achieved the intended goal. Following the example above, this phase was where multiple textile samples would be created to test their flexibility.

Component Cycle Phase 3: Review

The review phase involved testing the prototypes to see if they achieved the desired goal, and if there were any further limitations or changes that needed to be made. In the textile example, this would involve wearing or feeling the flexibility of textile samples produced during the creation phase. If multiple prototypes were created, they would be compared and the most promising one chosen as a starting point for the next ideation phase.

Component Subdivision

Sometimes during the review phase multiple possible solutions would emerge that could provide a good path for further exploration. In the textile example, there may be two or more samples created that show good signs of flexibility, but with some differing limitations or considerations. In these cases the component cycles could subdivide, and each potential solution explored concurrently. Subdivision meant that less promising solutions that might not normally be considered were explored, rather than allowing predictions of optimal solutions to dictate the design trajectory.

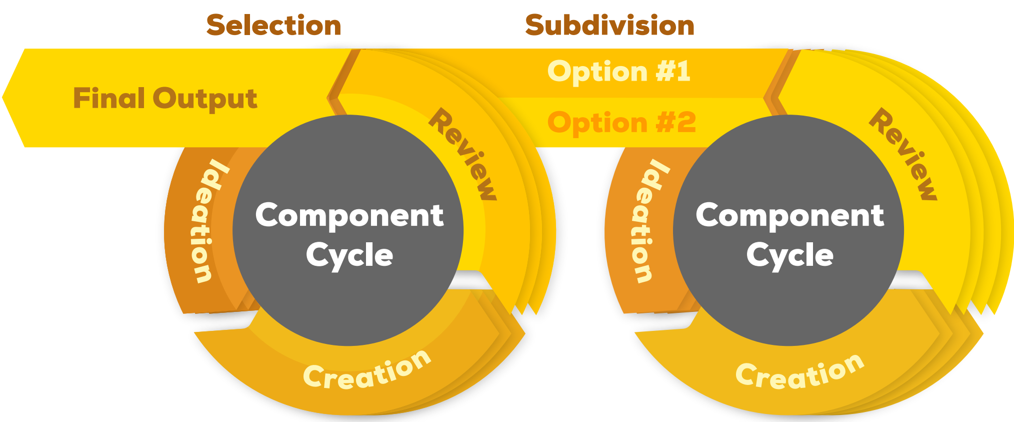

Figure 4. Subdivision of component cycles. An example of one component cycle subdividing into 2 options. Note that more than 2 options may occur.

As the time approached for a system integration to occur, each component cycle needed to conclude so that the outcomes could be combined into a single integrated prototype. If a component cycle had subdivided and still retained multiple options, it is at this point that a selection would need to be made. Where possible, this was done in conjunction with the physiotherapists in order to ensure the option that best met the user expectations was selected. The system integration process involved combining all the final outputs together into a contiguous unit. As component cycles completed at different times, the integration process occurred concurrently to some component development. Once a cycle completed, it was integrated into the final prototype immediately. This allowed component cycles that depended on some integrated systems for development to proceed during the system integration. Once the final component was integrated, the prototyping phase came to an end.

3.3 Development Phase 3: Reflection

Following the prototyping phase was the reflection phase. Reviewing the final prototype and development processes was the primary goal of this phase. The integrated prototype was presented to the physiotherapist group for feedback in preparation for the next development cycle. This gave the physiotherapists an opportunity to review how the various components interacted with each other. This phase also provided an opportunity to review the development process and consider how it might also be altered during future development cycles. Reflecting not only on the outcome but also the design process is an important aspect of this methodology.

Chapter 4

Design Process

Figure 5 shows the existing haptic feedback glove device that was utilised as the starting point for this research (Foottit et al., 2016). It featured six IMUs that measured hand and finger orientation. The angle difference between the fingers and hand were used to provide an indication of finger flex. The device also utilised a Bluetooth connection and was powered by an internal battery. Cylindrical ERM motors provided vibration at the finger tips for haptic feedback. The glove was used to control an aircraft in a basic flight simulator game, where hand orientation controlled the aircraft's movement, while flexing fingers triggered the firing of guns or rockets. Weapons fire produced haptic vibrations that differed based on the type of weapon being fired.

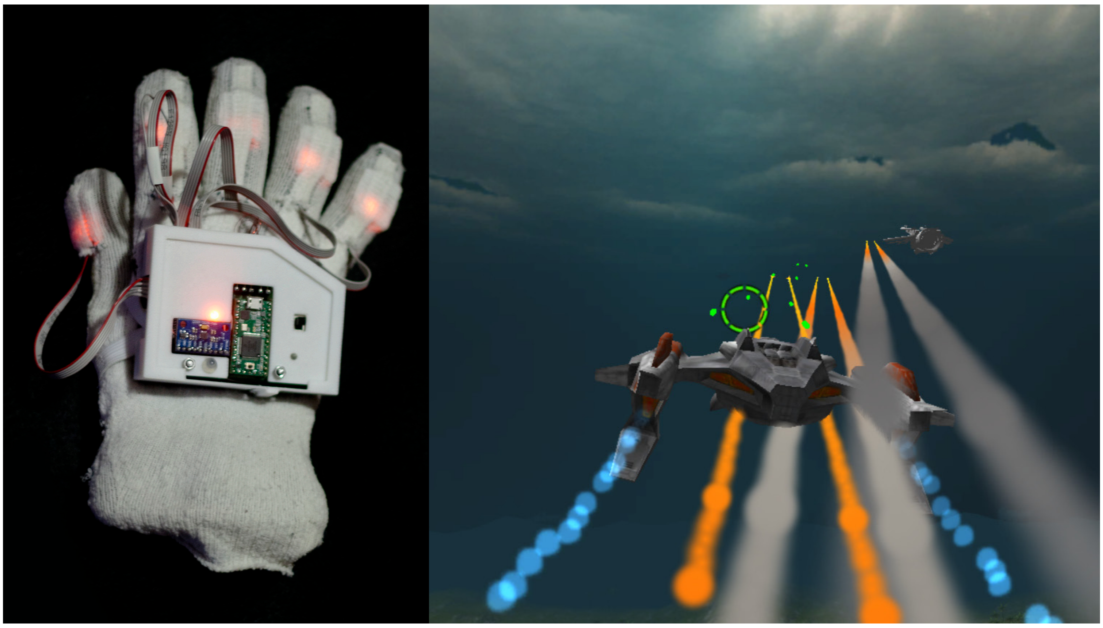

Figure 5. Previous prototype. The previous prototype haptic feedback glove and associated flight simulator game.

Using an existing prototype enabled the research to begin with a review of the current design and identify iterative improvements that needed to be made in order for the glove to better suit rehabilitation purposes. Since it was anticipated that a device used for rehabilitation purposes would most likely be used in a clinical setting, physiotherapists were identified as one of the user groups for the device in relation to the diagnostic quality of the glove. They would also be able to provide professional opinion on behalf of patients.

4.1 Development Phase 1: Discovery

4.1.1 Workshop with Physiotherapists

The first phase of the research involved having a workshop with the Neurological Rehabilitation Team based at the Health and Rehabilitation Research Centre at Auckland

University of Technology (AUT). This team has experience and research interests that align well with this masters research, including recovery of the arm and hand after a stroke. Ethics approval was obtained and all five physiotherapists that were able to attend the workshop consented to being involved in the research, bringing a range of experiences and perspectives. All five physiotherapists attended the first workshop, including the team leader. During this 1 hour session the physiotherapists were acquainted with the current glove prototype, and then an agile user story creation method was used to generate design considerations. The responses from physiotherapists were recorded on cards and prioritised where possible during the session, as seen in Figure 6. User stories emphasise user experience over technical specifications. Using this method it was possible to gain insight into what struggles a user may have with the existing prototype and why. This built a picture of the user that would be very useful during development.

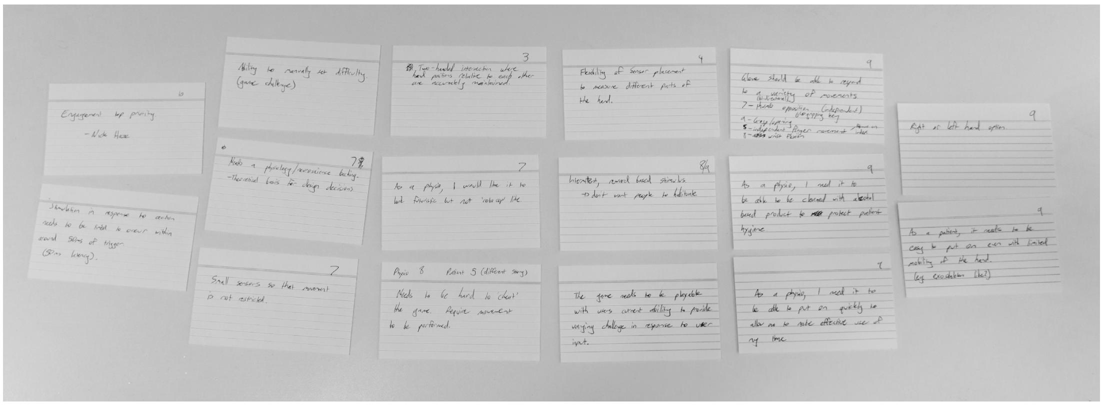

Figure 6. User stories. User story cards from the workshop with physiotherapists.

The user stories revealed three key areas of consideration. The first and most highly rated area was that of wearability. It was extremely important from the physiotherapists perspective that the glove be quick and easy to put on someone with limited hand mobility. This presents a number of significant challenges for the textile design, particularly for glove-based approaches. Responding to these challenges would require the consideration of alternatives to using a glove as the mounting for electronics in addition to exploring glove designs that might respond to the limitations of existing glove designs. One example of an alternative option raised was the possibility of creating a rigid exoskeleton that could be attached to the hand from above. This suggestion also allowed for cleaning the device easily between uses by different people, the second key area of consideration raised. Hygiene was important to the physiotherapists, as they have strict hygiene regulations that must be adhered to in order to prevent the spread of infection when a device is to be used by multiple people. Thirdly, it was important to the physiotherapists that activities performed with the glove be engaging in order to motivate clients to continue using it. While this was a very important area of consideration, it falls outside of the scope of this research.

4.2 Development Phase 2: Prototyping

4.2.1 Component Isolation

In order to break the overall system into components, it was necessary to consider how it might fit together when the phase completed. This helped to define the various components that would be required and the boundaries between them. For example, it was evident early on that the main processing board and battery would be better positioned on the arm rather than the back of the hand in order to reduce fatigue in the user. This meant the module that was previously on the back of the hand would be redesigned and become the arm module as seen in Figure 7. In addition, a new wrist module was designed for the back of the hand. This represented two components of the system.

Figure 7. Repositioning of primary module. The highlighted module on the back of the hand (left) was redesigned to become the highlighted arm module (right).

Textile elements of the glove were another component that needed to be considered. This would serve as a structure for mounting the various electronics and attaching them to the user's body. Since the physiotherapists emphasised that there were significant difficulties associated with putting a fingered glove onto someone with limited hand mobility, it became clear that the fingers would pose a significant design challenge. This meant that the finger modules themselves were considered an independent component, separate to the textile. This enabled a more open exploration of solutions than focusing on the fingers as an integral part of the textile.

The remaining components identified were the interconnections between various modules and the firmware that would drive the system. For this iteration the interconnections component was limited to the wiring options, with connector selection being left as part of the individual module components. The firmware component covered all the programming required for the system to function, including reading from sensors, driving actuators, and communicating with a computer. The firmware component in particular required outcomes from the other components in order to progress.

Due to the interdependent nature of some components, there was a dependency structure where occasionally one component could not progress in development until another component completed an iteration cycle and provided an outcome. For example, firmware could not be written for the sensors until it was determined what devices would be used as sensors. In addition, interdependence meant that changes to one component could have an impact on the design of others, sometimes in unpredictable ways.

A typical challenge when implementing an agile methodology is how to handle dependencies (Brown, 2012), and unfortunately this challenge is even more difficult in hardware development situations due to the interactions between electronics design and programming. Fortunately an agile methodology is also flexible and Brown (2012) offers some ways to handle dependencies in an agile manner. For this research, it was decided to prioritise components based on a combination of known dependencies and predictable delays or lead times. The factors considered while assigning priority are given in table 1.

Table 1

Factors influencing priority of components

Factor	Relative Priority
Sole dependency for another component	Very High
One of multiple dependencies for another component	Medium
Lead Time (e.g. will need to wait on parts)	High

With the various components prioritised, development proceeded with a focus on higher priority components first. When a component was unable to proceed further while waiting for parts or another component to finish, other components could be worked on. The remainder of this section details the process of development for each individual component, and concludes with a description of the final system integration process.

4.2.2 Finger Modules

One of the first objectives for the finger modules was to identify the desired sensor to use. For this purpose the MPU9250 from Invensense (Invensense, n.d.) was chosen, as it offered similar characteristics to the previous sensors used but with the added advantage of an included magnetometer and a choice of either inter-integrated circuit (I^2C) or serial peripheral interface (SPI) buses for communication with the main board. IMUs were chosen instead of alternative sensing technologies such as piezo-resistive flex sensors because they presented the best option for minimising the restriction of joint mobility. Flex sensors require secure attachment over a joint in order to accurately measure the joint angle, whereas IMUs can be positioned away from the interphalangeal spaces and have cables loosely run so as to provide minimal restriction to normal joint motion. In a rehabilitation context this is particularly important as any resistance to joint movement may make it more difficult for a patient to perform a task.

Existing commercially available breakout boards for the chosen sensor were too large for mounting on a finger, so a custom sensor board was developed. Once the board design was completed, parts needed to be ordered and PCBs fabricated. This would entail a long lead time and introduce shipping costs, so the parts and PCBs for the sensor boards were ordered in conjunction with those for the wrist board and main board. Small, coin style vibration motors were also ordered to allow the haptic feedback elements to be encapsulated and reduced in size compared to the cylindrical ERM motors used previously.

Following the receipt of parts and PCBs, the sensor modules were built and each module tested to ensure the sensor was working. This involved connecting each module to a Teensy 3.2 microcontroller (‘Teensy USB Development Board’, n.d.)

via a breadboard connection and running a basic program that would read raw data from the sensor. All sensor modules passed this test.

With the electronics built, a plastic housing was designed to house the sensor modules and vibration motors. This was used as a base to explore a range of mounting options, including stitching to a glove underlay, rigid plastic clips, roll-over knit structure, and self-adhesive bandages. Figure 8 depicts the development stages of the plastic housing for the sensors and clip mountings for the vibration motors.

Figure 8. Iterative refinement of plastic sensor housing and haptic motor clips. In order to obtain a comfortable but firm fit, different clip sizes were made to suit a range of finger sizes.

The challenges associated with using gloves in rehabilitation contexts have been previously identified in existing research (see Section 2.4.2) and during the workshop with physiotherapists as part of the discovery phase of development (see Section 4.1.1). Section 4.2.5 goes into further detail about glove designs that were created in an attempt to develop a glove-based solution. However as development of the textile component progressed, it became evident that a fingered glove solution would not be viable. This redirected attention towards the mounting options that did not require the textile support of glove fingers, such as plastic clips and self-adhesive bandages. Experimentation revealed that clips were likely to become bulky and could not accommodate both a sensor and motor in a single mounting. In contrast, the self-adhesive bandage approach allowed both the sensor and motor to be mounted on the same support as seen in

Figure 10 and was a bit more adaptable to different finger sizes. However the self-adhesive bandage approach was limited in that it would lose adhesion after a few uses and required a high amount of tension to remain securely fastened during finger movement, which had the potential to impede blood flow. In addition the self-adhesive bandage was more difficult to put on due to the wrap-around nature of the mounting.

Figure 9. Initial experiments with self-adhesive bandage.

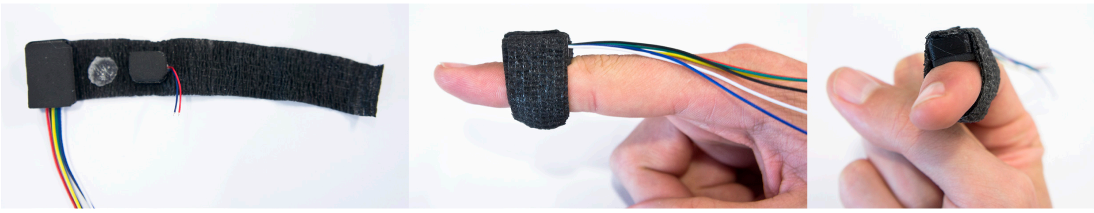

Figure 10. Integrating sensor and motor into a single self-adhesive wrap.

A final solution explored combining the rigid clip idea with the textile nature of the self-adhesive bandage to develop a single-piece elastic ring that could slide onto a finger and have both the sensor and motor attached. This had the advantage of being easy to slide on like the plastic clips, but with the benefit that it more easily adapted to different finger sizes and could hold both the sensor and motor on a single piece. Figure 11 depicts the various stages of development for this approach. Initially plastic casings with wings on either side were utilised to house the sensor and motor modules. Ribbed elastic was then stitched into a loop and attached to the sensor and motor mountings. With these early designs the sensor and motor directly contacted the finger, or had a thin foam padding. The intent of this design was to allow the mounting piece to be pulled apart while being placed on the finger. Unfortunately this also made the mounting pieces wider than the finger and caused them to collide with the mountings on adjacent fingers. For this reason a variation was developed where the sensor and motor modules were stitched to the outside of an elastic loop. Two textile materials were explored for this solution: ribbed elastic and neoprene. While the neoprene offered some advantages in terms of softness and ease of stretching, it was also much more difficult to put on as it gripped the skin more easily causing it to deform rather than slide on smoothly. In contrast the more rigid structure of the ribbed elastic meant it was smooth and easy to slide on.

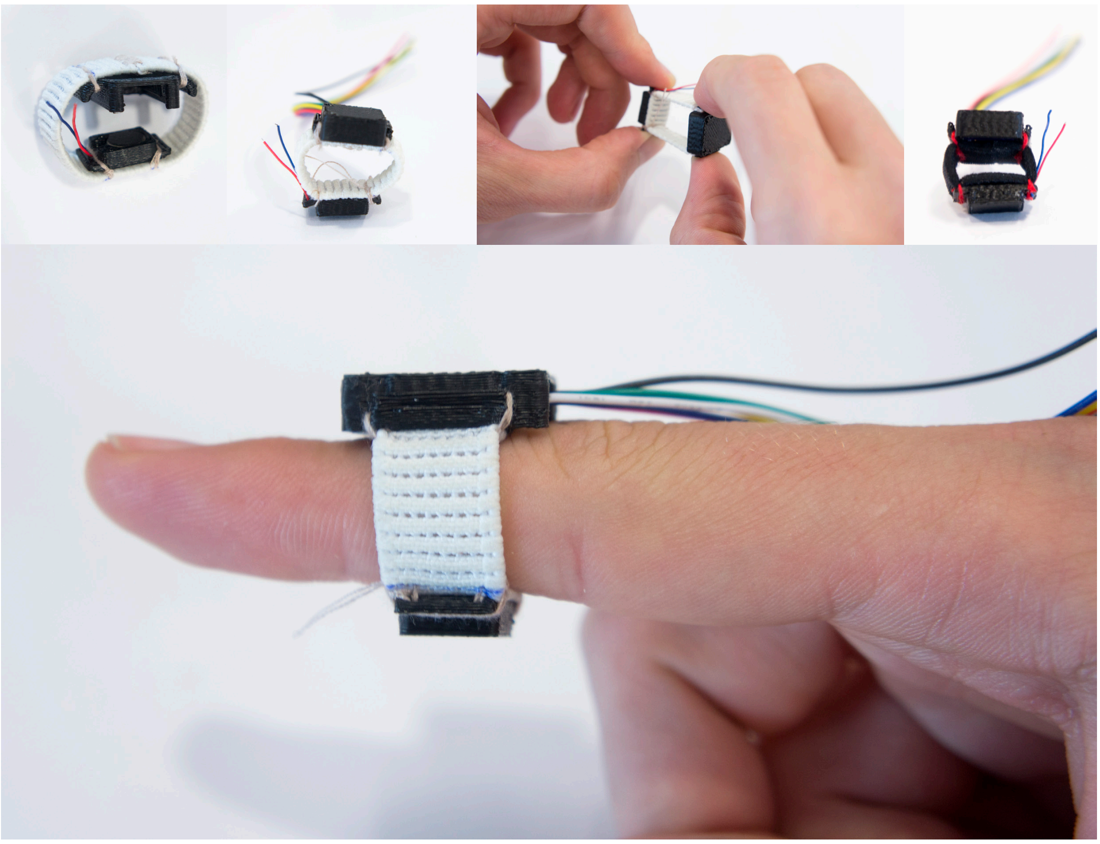

Figure 11. Development of finger module with a motor and sensor mounted on an elastic loop. Early versions with an external textile loop are depicted on the left. An alternative design with the sensor and motor attached to the outside of ribbed elastic (centre) and neoprene (right) is also depicted.

4.2.3 Arm Module

The repositioning of the main board to a position on the arm meant that it needed to be redesigned in order to better suit the mounting position. This presented an opportunity to also revisit some of the electronic component choices that were made for the existing prototype. One of the key decisions was to change from I^2C to an SPI bus for communication with the sensor modules. In addition, a newer motor driver was chosen that allowed communication via an I^2C bus connection. This enabled the drivers to be configured electronically by the main processor, adding flexibility to the system.

A key focus during development of the updated main board was minimising size. However this had to be balanced with the costs in both time and money associated with reducing the size of some elements in the design. A key example of this was the microcontroller board. One way to potentially greatly reduce the main board size would be to integrate a microcontroller directly into the board rather than using a third party board such as the Teensy. However, this would entail potentially spending a disproportionately large amount of time designing the supporting circuitry for the microcontroller. It would also require a more complex and expensive PCB instead of the 2-layer design in use currently. The potential benefits of a smaller board were not great enough to justify the added time and monetary costs, so the Teensy 3.2 board was used for the design as seen in Figure 12.

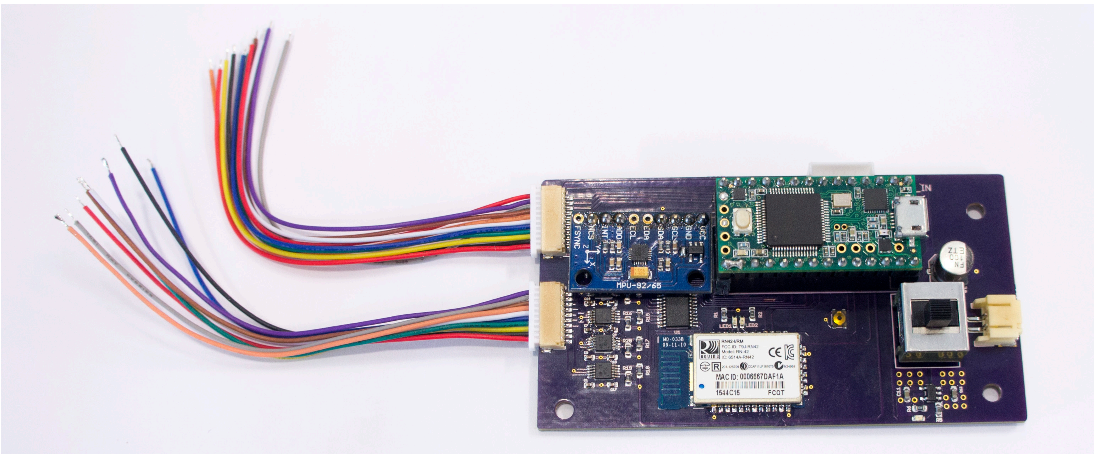

Figure 12. Final produced main board. Note the combination of custom circuits such as the motor drivers on the left, and commercially available breakouts such as the Teensy 3.2 seen at the top right of the board.

One area where size could be reduced was related to the power regulation system. The previous prototype utilised two buck-boost voltage regulators to provide power for the digital electronics and the motor drivers. Buck-boost voltage regulators are a form of switching regulator that can efficiently transform an input voltage into the desired output voltage. They are particularly useful for their ability to increase a low voltage to a higher one, rather than only being able to reduce voltage like simpler linear regulators. This enabled the system to operate even when the battery voltage dropped below 3.3v. However it also required a large amount of space. In addition, lithium polymer (LiPo) batteries can suffer from reduced life if they are discharged too much. After investigating some alternative battery options, it was decided a LiPo battery would be utilised as they offer the greatest power to weight ratio of affordable consumer grade batteries at this time. So the decision was made to simplify the power supply circuitry by using the regulator on the Teensy 3.2 board to supply power for the digital electronics, and to power the motor drivers directly off the battery.

With the main board designed, the casing and battery options were explored further. LiPo batteries come in a range of sizes and shapes, so the first step was to identify an affordable LiPo battery that had a size and shape appropriate to the application while maximising battery capacity. After looking at a range of options, the battery for a Samsung Galaxy Note 1 was chosen. The dimensions for this battery were very similar to that of the main board and it offered a high battery capacity. It was also one of the cheaper options available, and could be removable. The battery was measured and mounting options explored as part of the case design. Figure 13 depicts the battery housing development.

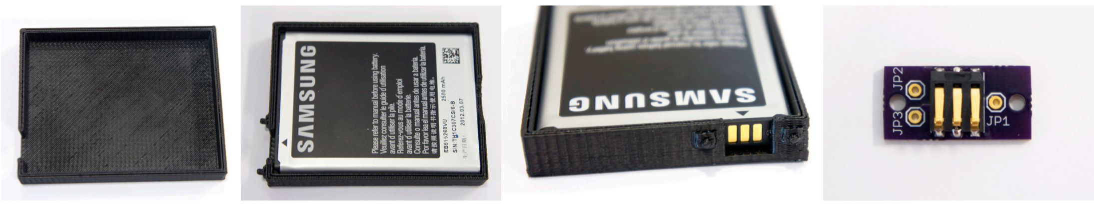

Figure 13. Development of battery housing. Multiple versions were produced, with the final version including a custom made board utilising spring contacts for connecting the battery.

The case design went through a number of iterations as shown in Figure 14. The aim of these developments was to refine the shape in order to minimise size and weight while maintaining an aesthetic look. Early versions of the case had separate mountings for the main board and battery. Later versions mounted the main board directly into the cover and had the battery mount beneath it. Using additive 3D printing, it was possible to verify the positioning of critical parts to ensure that switches, buttons and other elements lined up with the main board. In addition physical models of the case more readily identified potential problems and opportunities for simplification than 3D models alone.

Figure 14. Developments of arm casing utilising affordable and quick 3D printing for rapid prototyping. Note the incremental addition and refinement of mounting holes and supports.

Later stages of development explored how the casing would fit together and attach to the textile support. This too went through a number of iterations. A base plate was developed that would be stitched to the textile support and then would attach to the arm module using 4 machine screws. In this way the entire arm module could be easily assembled and disassembled for maintenance or testing. As the design complexity increased, it became difficult to produce accurate models using affordable 3D printing technologies. As a result assembly simulation was performed using CAD software to identify possible issues with the geometry that could prevent the module from being assembled. For example, the virtual CAD assembly was used to animate the process of assembling the device so that it could be seen if any parts would interfere with each other during the multiple steps involved. Figure 15 depicts some of the CAD visualisations used during the design process.

Figure 15. CAD images at different stages of development. Exploded views and cutaways were utilised to check for interference between different parts of the assembly.

The final casing designs were 3D printed out of a nylon material using a selective laser sintering (SLS) printer. This enabled the designs to have a high degree of precision on fine details and unsupported geometry that would be difficult or impossible to reproduce using the fused deposition modeling (FDM) printer that had previously been used. The outcome of this was a set of high quality parts that fit together without issues. Figure 16 depicts the high quality printed arm casing and push button actuator assembly.

Figure 16. High quality 3D printed arm casing and push button actuator assembly. A piece of carbon fibre rod and a small piece of foam was attached to a 3D printed dome to form the push button actuator assembly (centre image). Note the fine details produced with this rapid prototyping method.

Some of the finer details for the final case design included a push button actuator that was constructed of carbon fibre rod, a small piece of foam, and a 3D printed cap. The assembly slid into a shaft on the main outer shell of the casing and was used to actuate a small tac-switch mounted on the main board. An optical fibre was used to conduct light from SMD light emitting diodes (LEDs) to the outer shell of the case, and a custom switch piece was designed to integrate the power switch into the overall design. Figure 17 depicts the final arm module assembly.

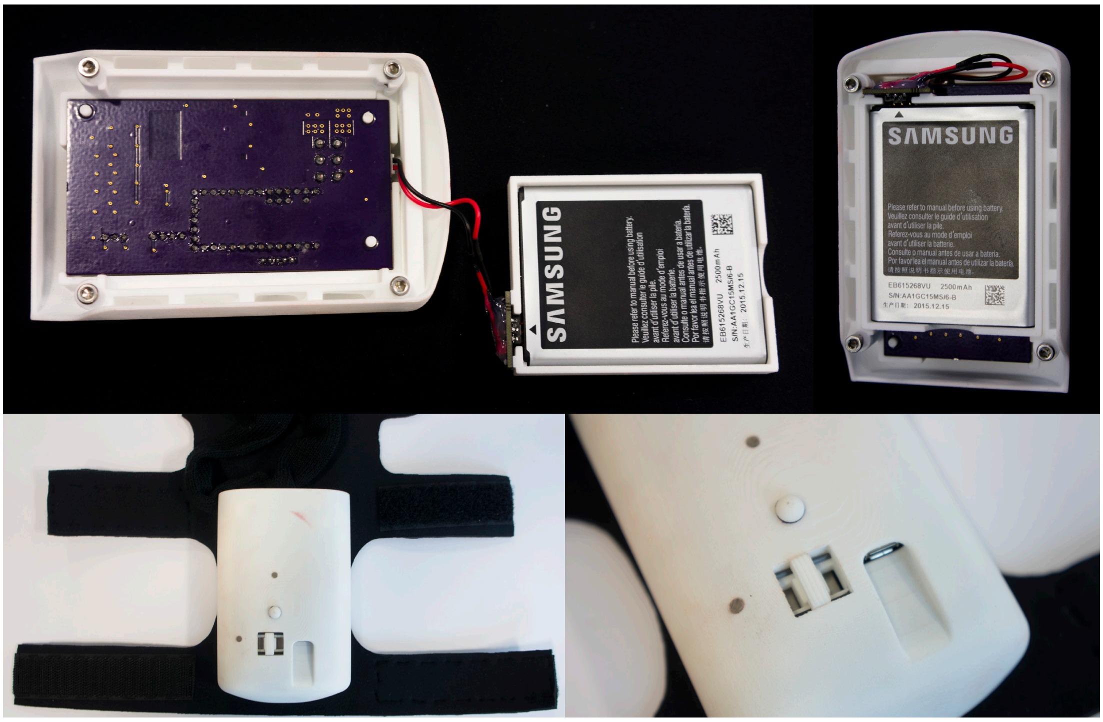

Figure 17. Final arm module assembly. The battery can be seen mounted on the underside before the module is attached to the neoprene support.

4.2.4 Wrist Module

With the main processing board moved to the arm, the wrist module provided a central point where all the finger sensors could be combined into a single bus connection back to the main board. As a result the board design began as an extension of the finger module sensor board with additional connectors added. A single, 11-pin socket was used for the main board connection, while five 6-pin sockets were used for the finger modules. Connector placement was a challenging process. Early designs attempted to position side-entry connectors along edges of the board to facilitate a thin form factor. However it soon became apparent that electrical routing on the board would be extremely complex with this layout, particularly on a 2-layer PCB. As a result, the design was changed so that top-entry connectors were used and grouped in a single line along one side of the board. While this greatly simplified the board design, it increased the height requirements of the casing due to the choice of connector.

Japan solderless terminal (JST) sockets were chosen for the wrist board connectors, as pre-made cables are often readily available for these connectors. They are also very affordable and come in a range of sizes. Connectors with a 1mm pitch spacing were chosen to minimise the footprint requirements of the parts. JST plugs are designed so that the cable comes directly out of the back a connector. This means that top entry connectors have the wires coming up perpendicular to the board. This increased the height requirements of the casing in order to provide sufficient height for the cables to bend from an upwards direction to be parallel with the hand. Figure 18 depicts the wrist board and early casing designs.

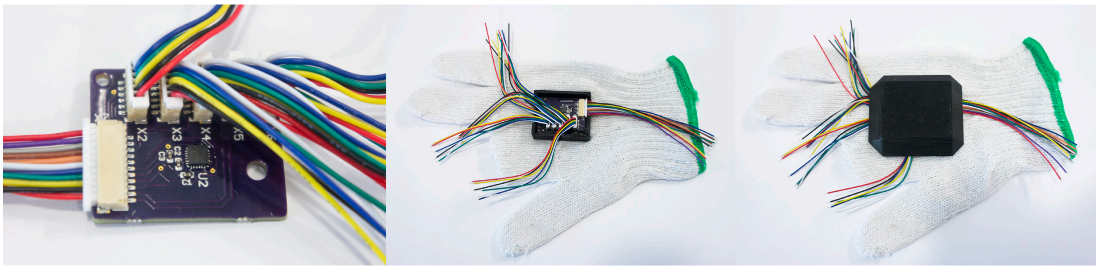

Figure 18. Wrist PCB and housing. Note increased height requirements of the 6-pin top entry connectors compared to the 11-pin side entry connector.

The case design for the wrist module had two main requirements. It had to secure the PCB board with the sensor and connectors to the textile support, and it had to contain the wires in order to provide an aesthetic look while also preventing wires from catching on other objects. Figure 19 depicts the early design iterations for the wrist casing. These remained largely unchanged until the system integration process began.

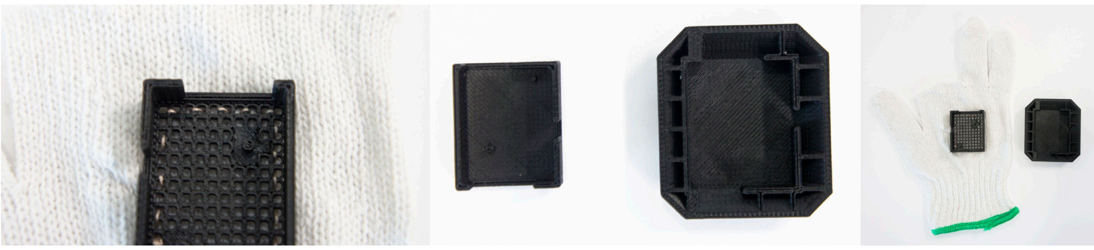

Figure 19. Early iterations of wrist casing. Note the use of a grid pattern in the base to allow for stitching to the textile component.

The wrist module base had a grid pattern on the bottom that enabled it to be stitched to a textile support. Small plastic extrusions allowed the wrist PCB board to be attached after the base was stitched in place. A cover was then pushed on top, using friction forces to hold it in place. The case design went through a number of iterations, with the top piece in particular being redesigned during the system integration process to prevent it from interfering with hand movements.

Figure 20. Final assembly with redesigned top cover. Note the reduced size of the cover.

4.2.5 Textile

The textile support for the electronics was developed so that it would be easy to wear, while also being securely attached to the body. A number of iterations were performed to investigate two competing approaches. The first approach was to adapt a textile glove in order to make it easier to put on someone with limited hand mobility. The second approach was to develop a wrap-around or exoskeleton system that could be attached to the body without needing to put the hand and fingers through openings such as the back of a glove.

The key difficulty in putting a glove on someone with limited hand mobility is that they are unable to hold their hand rigid. This means that the hand itself cannot be used to push apart textile tubes such as glove fingers. It also makes the task of lining up their fingers with the finger openings of a glove very difficult. A number of design elements were explored to alleviate these difficulties. Figure 21 depicts some of the possibilities explored during early experiments.

Figure 21. Initial glove developments. Early developments explored ways of opening up the palm area of the glove and supporting the finger entry points with plastic rings.

The first glove design change involved opening up the wrist part of a glove so that it would wrap around the wrist rather than enclose it. This allowed fingers to be guided into finger holes individually, before securing the glove in place by wrapping the wrist part around and securing with velcro. In addition, small plastic rings were embedded in the finger openings in order to make it easier to guide a finger into the opening. While it was easier to put on the glove with plastic rings embedded, they were tedious to integrate into the glove textile and limited the range of finger sizes that could utilise the glove.

Further glove developments explored using thermoplastic yarns instead of plastic clips to add rigidity at the finger openings. This increased comfort and was much easier to fabricate, but was less effective at keeping the openings wide. Wool and plated lycra yarns were also explored (see Figure 22) to find a balance between firm fit and minimal resistance when putting on. This was particularly necessary since any shifting of sensors relative to the body causes error in measurement. As a result the textile support must be firmly mated to the body.

Figure 22. Glove samples with thermoplastic yarn and sleeves for electronics. The thermoplastic yarn can be seen as white strips near the join between the fingers and hand.

An important consideration was hygiene, so a method of having removable electronics was explored to allow the glove to be washed. This involved knitting a tube attached to the end of a finger that could be folded over to contain the electronics within. Sizing and yarn options were explored, but during tests all the samples tried were difficult to utilise. Constructing a tool to assist the process of rolling the tube over electronics was considered, but not pursued.

Ongoing testing with glove samples indicated that fingered gloves may not be a good solution. Feedback from physiotherapists revealed that even with supported openings it would be difficult to put on a patient. At the same time, a promising alternative solution based on neoprene wrist supports was emerging (see Figure 23). By creating a wrap-around support it could be possible to securely attach the wrist and arm modules without needing to put the fingers through openings such as those in a fingerless glove. While a promising solution, it would require the finger modules to have their own way of mounting. Fortunately an alternative way of mounting the finger modules was devised, making the neoprene wrap an option worth further exploration.

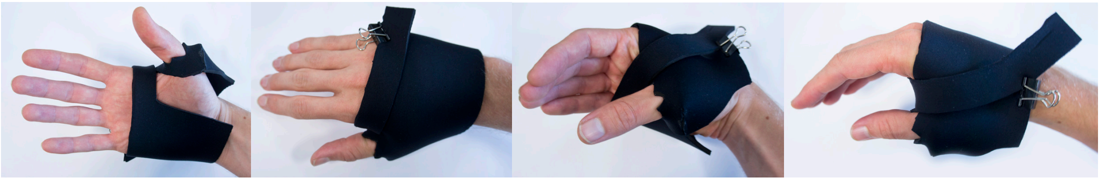

Figure 23. Initial experiments with neoprene as a wraparound method for mounting the wrist module.

Developing the neoprene wrap involved a number of iterations as seen in Figures 24 and 25. A pattern for the textile was developed through this process that allowed a firm attachment to the hand and arm while also being easily adjustable and moulding to the three-dimensional shaping around the thumb.

Figure 24. First complete neoprene support. A first complete neoprene textile support was constructed from offcuts and glued together. It demonstrated a high degree of flexibility and did not impair movement.

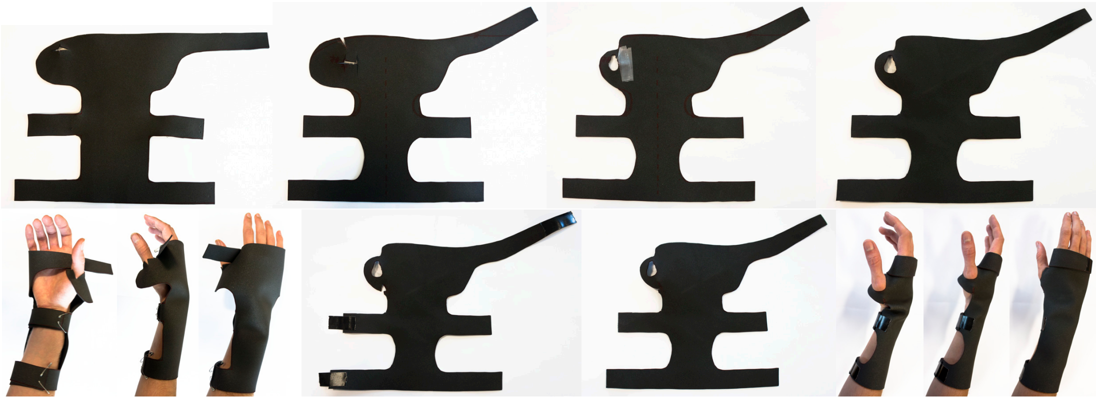

Figure 25. Neoprene support developments. Multiple patterns were developed to improve the fitting.

In addition to the neoprene wrap, a fingerless glove was developed that used a zip along the arm (see Figure 26). This allowed it to be easily put on and then zipped up to provide a firm mating to the arm and wrist. The two alternative designs were presented to physiotherapists and the neoprene wrap was chosen as the preferred design due to the ease with which it could be put on.

Once the neoprene was selected as the ideal option, a final iteration was performed that refined the pattern. High quality 1.5mm neoprene was sourced to fabricate the final piece. Velcro was carefully stitched in place, followed by the final mounting plates for the arm and wrist modules. The final neoprene wrap can also be seen in Figure 26.

Figure 26. Final versions of the knitted glove with zipper and neoprene wrap.

4.2.6 Interconnections

Providing good electrical connections between the various modules was necessary to ensure quality sensor readings, and effective haptic output. The key challenge for this component was to find a way for the interconnections to be flexible and able to stretch. For example, wrist flexion causes a changing distance between the arm and wrist modules. Even more extreme is the requirements for finger modules, which not only need to be capable of accommodating finger flexing but also require the ability to extend past the end of a finger so that the module can be put on and taken off. A number of possible approaches were explored in this area.

One of the first approaches taken was to simply use standard wire and bend it into curved shapes. As the distance increased, the curves would flatten out to provide additional length. The greater challenge for this approach was to get the wire to revert back to the curved shape when the distance decreased. Early experiments with the finger modules explored coiling the wire around a finger in order to provide a structure that would guide the wire back into the original shape when compressed. These initial experiments can be seen in Figure 27.

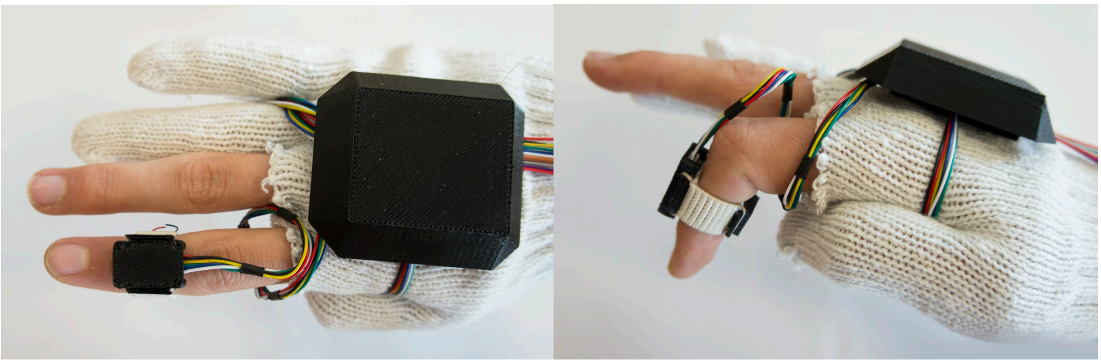

Figure 27. Early development of solid wire connections. Note how the wire is looped around the finger to provide a structure to guide it back to the original shape when compressed.

While the standard wire approach provided excellent electrical characteristics, it could become bulky and interfere with normal hand movement. This was particularly true when the wire was wrapped around a finger. An ideal solution would be some form of stretchable wiring to accommodate hand movements.

No commercial solutions for stretchable wires could be found, so a number of self-made approaches were explored. The first approach investigated was using conductive yarn in a knit structure. The limitation of this approach is the high electrical resistance of conductive yarns. Alternative approaches involved using thin enamelled wire combined with elastic or neoprene structures to provide a stretchable base. Since the wire was not able to stretch, it would be attached to the textile structure in a way that allows it to structurally deform under tension to increase length. Inspired by the ability for a spring to expand and contract many times, most attempts utilised a spiral structure, although a weaved approach was also attempted. Figure 28 depicts some early investigations into possible solutions for flexible wiring.

Figure 28. Early investigations into flexible wiring. Clockwise from top left: conductive yarn, enamelled wire threaded into a knitted textile, enamelled wire coiled around an elastic core, and stitched as a coil into neoprene.

Of the various tests, only two showed particular promise. The first solution was stitching wire into a neoprene base. While early tests showed some promising results (see Appendix C), when more than two wires were required this approach became bulky and tended to remain deformed when stretched. The second solution with promising early results was coiled wire around a central elastic cord. This solution demonstrated an acceptable ability to stretch and return to its original shape, while remaining relatively small compared to the neoprene approach. Further tests investigated how the number of coils per centimetre of elastic cord influenced the stretching characteristic. It was identified that an optimal range of coils per centimetre existed. Below this range, the wire would prevent the elastic from fully stretching as the coils would reach a point of maximum expansion before the elastic. With too many coils, the elastic would remain deformed after stretching due to the stiffness of the wire coil.

Selecting a coil to length ratio within the optimal range, a bundle of six elastic wires was created and attached to JST connectors. This enabled a single sensor to be connected using a stretchable link. Unfortunately weak points became evident, causing either the connector or bonding between standard and stretchy wires to fail (see Figure 29). In addition, with six stretchy wires in a bundle the amount of force required to stretch them becomes quite high (see Figure 30).

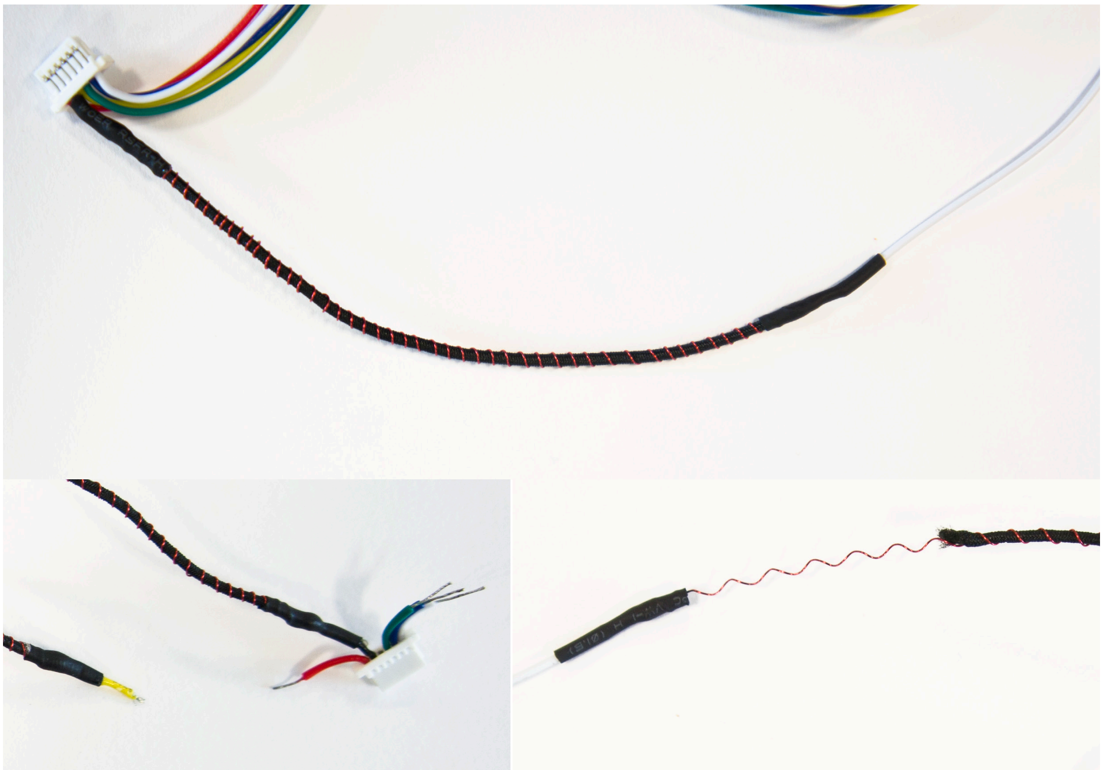

Figure 29. Stretchy wire failures. When integrated into connectors, stretchy wires were prone to failure at the connector or joint between normal and stretchy wire.

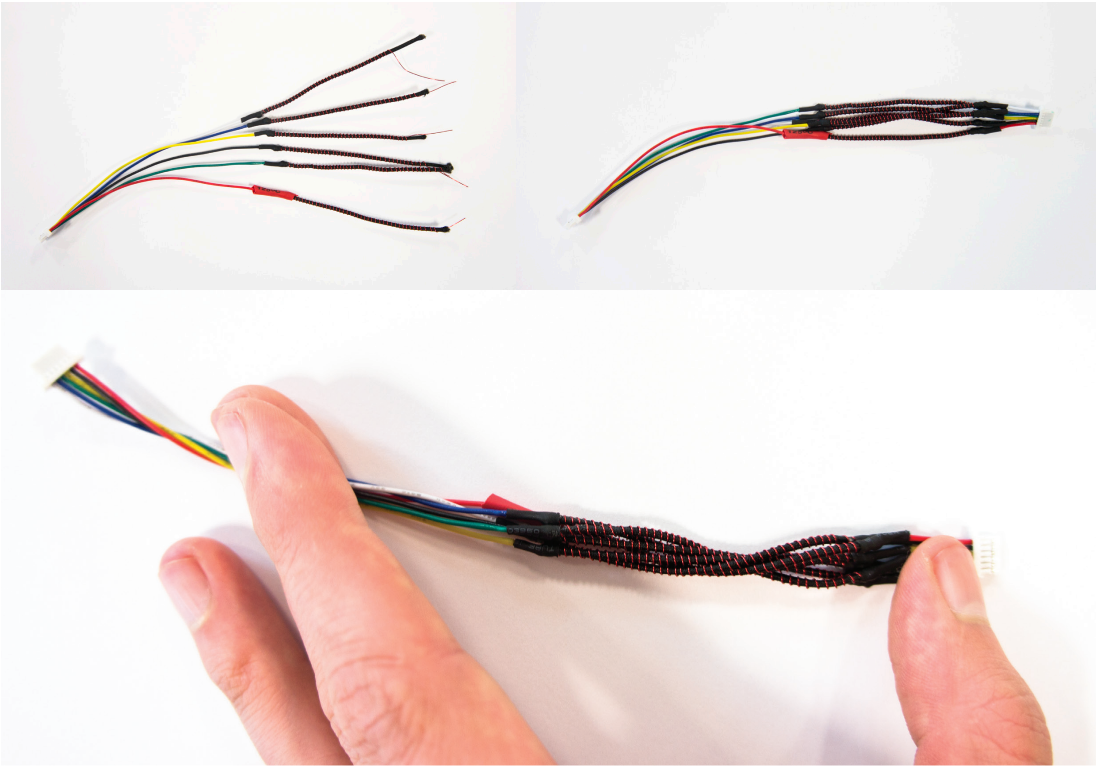

Figure 30. Stretchy wire bundles. Bundles of stretchy wires proved to be too difficult to stretch, resulting in high risk of failure.

With stretchy wires proving to be fragile and difficult to utilise, the original standard cabling was used. In order to improve the aesthetic appearance of the wires they were covered in a knitted tube made of black wool. In addition, enamelled wire was used for the haptic motor connections. This was coiled through the elastic loop of each finger module in order to provide a stretchable connection from the main wire bundle to the motor. Figure 31 depicts a single complete finger module connection.

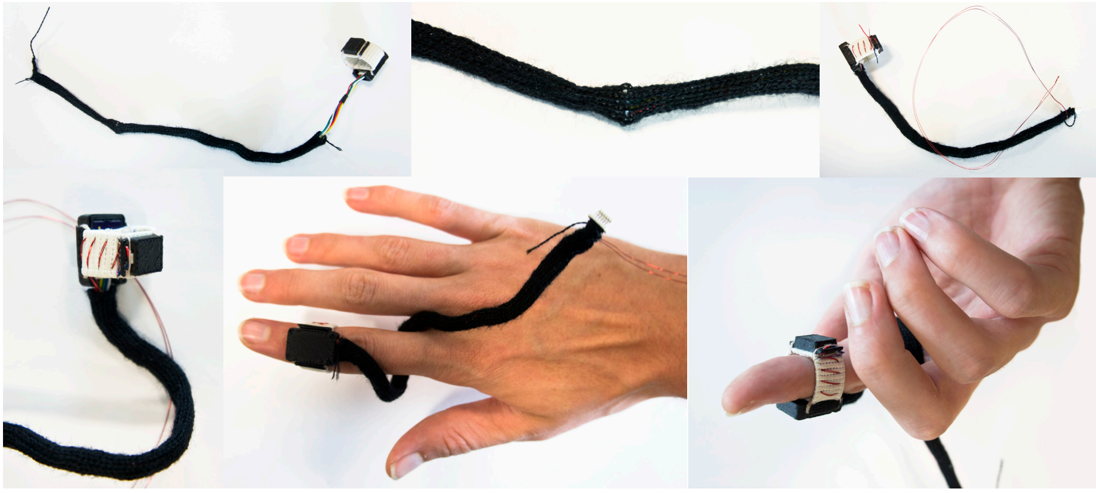

Figure 31. Connection from finger module to wrist board. Enamelled wire was used for the motor, and was stitched into the elastic ring before being coiled around the wire bundle from the sensor. Black knitted tubes were then put over the wire bundle for aesthetic appearance.

Once all the finger module wiring was completed, the final element to resolve for the interconnections was connecting the haptic motors to the motor drivers on the main board. Figure 32 depicts the process used for connecting the thin enamelled wires used for the haptic motors to the thicker wires used for the JST connector needed to plug them into the main board. This connection was done by soldering the enamelled wire to the thicker wire, and then looping it around the insulation before covering the joint with heat shrink. By looping the enamelled wire around the insulation of the thicker wire it was possible for the heat shrink to secure it in place, providing strain relief to the solder joint.

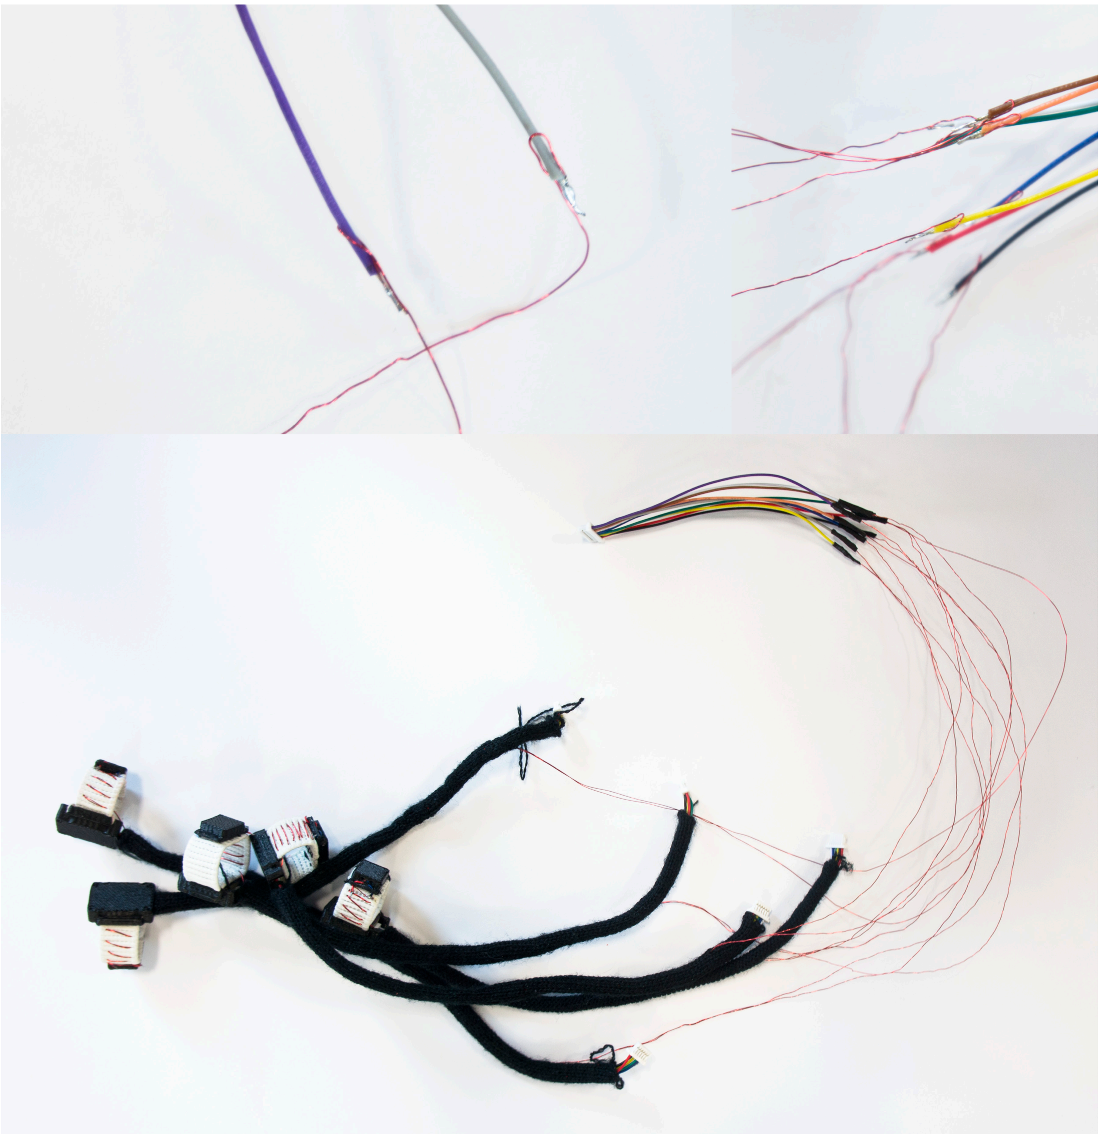

Figure 32. Connecting haptic motor wires to the connector plug for the main board. Thin enamelled wires were soldered to thicker wires and looped around the insulation before being covered by heat shrink.

4.2.7 Firmware

The existing prototype used as a starting point for this research included firmware for the device. However due to some key design changes in the hardware for the new prototype developed in this research, the previous firmware had to be significantly modified to accommodate the new hardware.

The first process that needed to be undertaken was to write a library that would interface with the MPU9250 sensors. The existing open-source library for MPU6050/MPU9150 sensors utilised by the previous firmware was used as a starting reference, and adapted to use the SPI bus instead of the I^2C bus. Modifications were also made to the library in order to improve support for the first in first out (FIFO) buffer feature of the MPU9250 sensor. Early versions of the adapted library were used when the sensor modules were built in order to test each module individually and confirm it was functioning, but did not support sensor fusion or FIFO buffering of data.

Sensor fusion for each of the MPU9250 sensors was performed using one of two well-known algorithms. Both the Madgwick algorithm (Madgwick, 2010) and the Mahoney algorithm (Mahony, Hamel & Pflimlin, 2008) were implemented, so each sensor could be configured to use either one. This provided flexibility in configuring the sensor fusion in order to allow further optimisation. The Madgwick and Mahoney algorithms are well documented (Munoz Diaz, de Ponte Muller, Jiménez & Zampella, 2015) and open source code for their implementation is readily available (x-io Technologies, 2012).

The selection of a new haptic motor driver required an entirely new library to be written for the Texas Instruments DRV2605 device. A library for the PCA9547 I^2C bus multiplexer was also required in order to allow communication with multiple motor drivers on the same I^2C bus. Both of these libraries were developed throughout the development process and tested as hardware was assembled.

Once the libraries were written, the firmware to control the system was developed. The communication protocol for the previously existing firmware was revised to improve data transmission speeds and add new features. The resulting communication protocol is documented in Appendix B. During the development multiple changes were made to the libraries as optimisations and bugs were identified. The device has four possible states for operation: standard, configuration, standby, and version. When powered up, the device initialises all sensors, motors, and input controls, then goes into the standard state.

In the standard state, sensor fusion is continuously occurring. When readings are requested, the most up to date values available are reported. Motor outputs can also be driven while the device is in the standard state. Different states can also be selected by sending the appropriate instruction to the device over Bluetooth or through a universal serial bus (USB) connection.

The configuration state allows the Bluetooth module to be configured. It does this by passing through serial commands from the USB port to the Bluetooth module, and responses from the Bluetooth module back to the USB port. In this state the device does not respond to commands transmitted to the Bluetooth module. Sending a ‘q ’character over the USB connection returns the device to the standard state.

The standby state places the device in a low power mode to conserve battery during inactive times. While in this state the device does not respond to commands and does not read from the sensors or drive the motors. To exit the standby state the device must be turned off and then on again.

The version state outputs the firmware version and then reverts to the standard state.

4.2.8 System Integration

The final system integration began once hardware from the various components had reached the point where a single solution could be produced. A feature lock was placed on the hardware components of the system, meaning that any new features or alterations that would require additional component iterations were pushed to the next development cycle. The exception to this was the firmware component, which relied on integrated components in order to progress. As a result, the firmware component underwent multiple iterations during the system integration process, with the output of each cycle being integrated into the final prototype immediately.

Since the various component cycles were not synchronised to end at the same time, the system integration process occurred in stages, with components being integrated into the final prototype as they were completed. The finger, wrist, and arm modules were the first to be completed. The next component to produce a final output was the interconnections, allowing the various electronic modules to be connected together as shown in Figure 33.

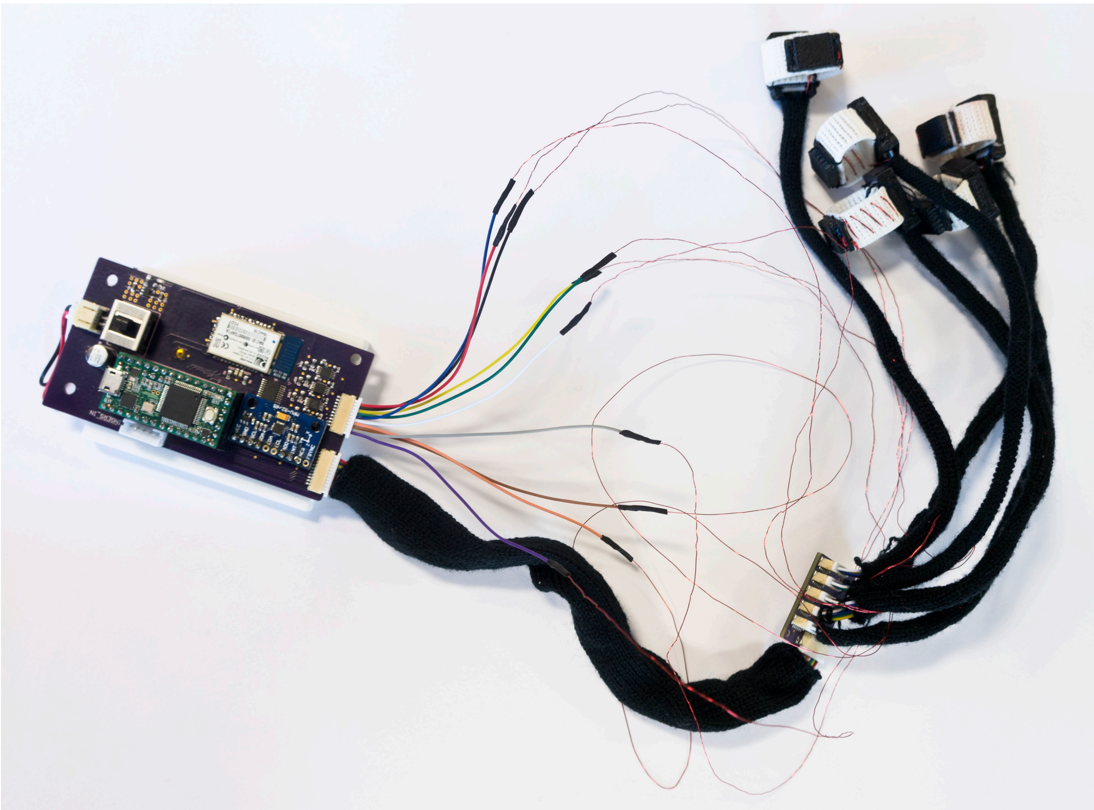

Figure 33. Electronics connected together for attaching to the textile base. Note the thin enamelled wires for the haptic motors.

The textile component was the next to complete, however there were still two competing solutions after the final component cycle completed. The two options were presented to the physiotherapist group, and the neoprene-based wrap method was chosen as the preferred option. During this session suggestions were also made that would involve further iteration of multiple components. These were excluded from the final prototype in order to maintain the feature lock and conclude the textile component iterations, but are discussed in Section 7.1.

With the textile solution finalised, it was possible to integrate the connected electronic modules into the final prototype as shown in Figure 34. During this process the cable bundles from the finger modules were stitched to the neoprene textile in order to secure them in place and prevent any stretching of the cables putting strain on the wrist module connectors. The enamelled wire from the haptic motors were also routed at this point. Unlike the finger modules, the haptic motor wiring went directly from the each finger module to the main board located on the arm. This introduced unexpected difficulties during the system integration. For example, the thumb and finger modules were routed on opposite sides of the neoprene, complicating the routing of haptic motor wires.

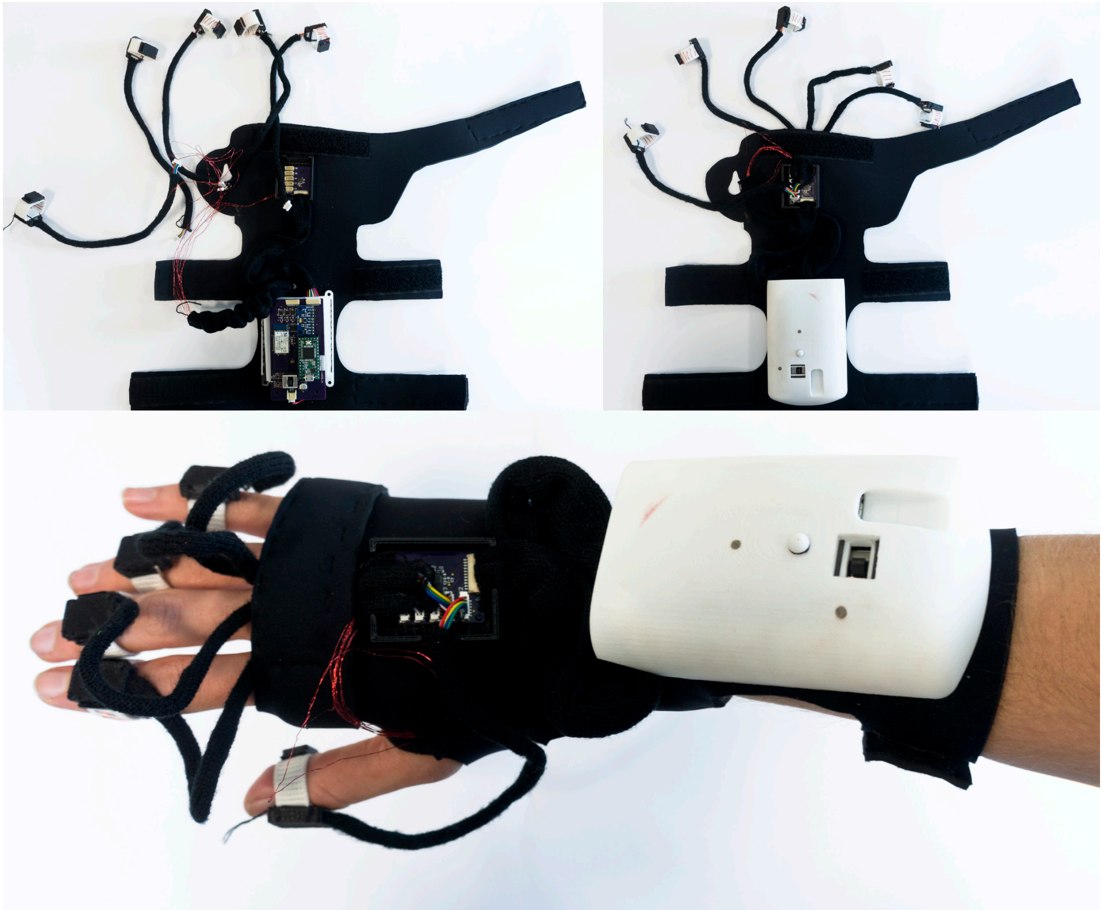

Figure 34. Mounting electronics to the neoprene base. Note how the finger module cables were secured to the neoprene base so that they roughly lined up with the fingers.

The next stage of the system integration involved wearing the device and testing that it did not restrict hand and finger motion. Cable lengths were tested to ensure they were not being overstretched during movements. Extremes of wrist flexion were also tested to ensure the arm and wrist modules did not collide and restrict the range of motion. Figure 35 depicts the wrist motions used to check for potential restrictions to range of movement.

Figure 35. Testing to ensure the final prototype did not restrict range of motion of the hand and wrist.

With the final prototype assembled, it became evident that there was excess wiring for the haptic motors that needed to be secured in order to prevent it getting in the way. A small patch of fabric was glued on the underside of the neoprene base to hold the excess wire as seen in Figure 36.

Figure 36. Final tweak. A patch was glued over excess enamelled wire. The wire bundles from the finger modules can be seen stitched in place at the top of the neoprene in the bottom image.

Once the hardware integrations were complete, the device firmware went through a testing process and was iteratively improved in response to failed tests. This approach allowed the code to be written during the development and integration of hardware components before it could be tested with the final hardware. New features were not added at this stage, development was limited to fixing bugs in the code that prevented the existing features from working. The system integration was concluded once a fully functional prototype was produced, also concluding the prototyping development phase.

4.3 Development Phase 3: Reflection

A final review session was held with the physiotherapists to conclude the development iteration cycle. This provided an opportunity for final feedback on the design in order to assess the outcome and provide input for a future development cycle. In addition, time was spent reflecting on the design process and how it responded to the challenges faced during development. The final outcome is detailed in Chapter 5, while Chapter 6 presents reflections on the design process.

Chapter 5

Outcome

The final prototype features 5 vibration motors and 7 IMUs. There are 5 finger modules, each with a single sensor and motor. These can be positioned on the thumb and fingers, as seen in Figure 37. Additional IMU sensors are mounted on the back of the hand and on the arm. With these sensors the glove is able to measure finger flexion, wrist flexion, and arm orientation. Each of the 5 vibration motors is connected to a haptic motor driver, capable of producing a range of haptic outputs. The drivers include a set of licensed libraries with 123 different haptic effects such as clicks and pulses. A full listing of the effects contained within the libraries is included in Appendix D. Custom waveforms are also possible, but would require additional firmware development to implement.

Figure 37. The final prototype. Note the positioning of sensor and motor modules on the middle phalanges of the fingers and the distal phalange of the thumb.

A Bluetooth module allows the device to connect wirelessly with any Bluetooth enabled device. Currently only PC software has been developed, but this could be easily expanded to other platforms as required. The application side has not been a part of this research, and is considered part of the future work.

Power is supplied by a $2500\,\mathrm{mA\,h}$ Lithium Polymer battery. Battery life measurements are difficult to predict as it can vary widely depending on the intensity and frequency of haptic vibrations. However a stress test of the battery indicated that it can last up to $17.5\,\mathrm{h}$ when only sensors are in use. For this test, the glove was connected to a laptop computer using Bluetooth and continuously transmitted raw data and quaternion orientation information for each of the 7 sensors. Data were recorded to the hard drive in comma separated value (CSV) format. Included in the data was a recorded timestamp that was used to calculate the operating time of the device. During recording the device connection dropped out 2 times, and was restarted. The timestamp was not interrupted during this time, so that the missing portions of data during signal drop-out did not influence the duration calculation.

A three-way toggle switch allows the device to be placed into three different modes: ON, OFF, and CHARGE. In the ON mode, the device calibrates itself and waits for a computer connection. The USB connection can also be used to update the device firmware while it is in the ON mode. In the CHARGE mode, the USB connection acts as a charging input for the battery. The OFF mode disables all circuits. The device cannot be programmed in either CHARGE or OFF mode.

LED indicators are used to communicate the device state. There are two LEDs used for this purpose, one indicates the device state in the ON mode, the other indicates CHARGE mode states. Figure 38 illustrates the different LED indicators.

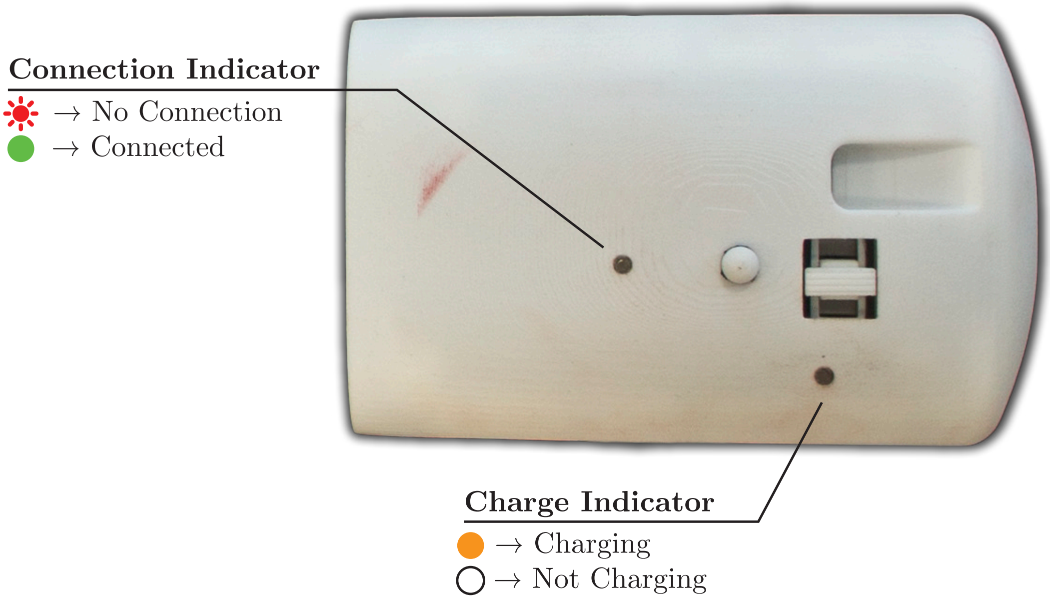

Figure 38. LED indicators on the arm module.

A push button on the arm module allows the user to interact with the firmware. A single, short push will record the current orientation of each sensor and use them as offsets in order to calibrate the sensor-body offsets. A long push will restart the device.

The device weighs a total of 252 g. The arm module accounts for a significant proportion of this weight at 185 g, while the finger modules only weigh 7 g each and the entire weight on the hand is only 70 g. The outer shell of the arm module weighs 51 g, and is the most likely source of potential weight reduction for future work. While this prototype positions the bulk weight in a better position than the original prototype, during the final reflection session feedback from the physiotherapists indicated there was still room for improvement in this area.

Subjectively, the weight distribution and comfort is significantly improved in the new prototype. Moving the majority of weight to the arm instead of the hand makes it less fatiguing during use, and places less strain on the wrist. By utilising finger modules and a wrap-around design for the hand and wrist, the device also feels less bulky. The new design also allows greater ventilation around the hand, which makes it less sweaty when being worn for extended periods of time. The finger modules are easily pushed on with one hand, and can stretch to fit a range of finger sizes. When being put on another person, it is possible to hold their finger in place while pushing the finger modules into position. The cable sleeving and custom printed casing gives the device a more refined and robust appearance, potentially increasing user confidence that the device will not break.

The IMUs used offered a significant benefit over those present in the previous prototype since each of the sensors incorporated a magnetometer. This is a key improvement to orientation accuracy as it allows the sensor fusion algorithms to account for accumulating errors in the yaw axis.

The accuracy of sensor readings remains to be tested, but previous research has demonstrated that the Madgwick and Mahoney algorithms are capable of producing results with an accuracy ranging from $<0.03°$ under ideal conditions to $<7.10°$ in the presence of magnetic distortions or rapid movements (Cavallo et al., 2014). Yaw rotations fare the worst, since the accelerometer cannot be used to correct errors from the gyroscope readings on this axis. Ongoing research to improve the sensor fusion algorithms provides a basis for further improvement in this area. Device usage can also be tailored to avoid rapid movements and magnetic disturbances in order to maximise the accuracy of sensor readings. Moreira et al. (2014) also utilises IMU sensors, detailing calibration and algorithm considerations. Given the existence of current research in this area, sensor calibration and algorithms has not been a focus for this research.

Feedback from physiotherapists during the final reflection session was generally positive, with all physiotherapists agreeing the device was a significant improvement over the previous prototype. In particular it was noted that it would likely be possible for the device to be worn by a patient with limited hand mobility. Some potential improvements were also noted, which could be integrated in future work. These are covered in Section 7.1.

Currently the prototype takes 27 s after powering up to calibrate the sensors, followed by another 10 s to calibrate motors. Once these calibrations are complete, the glove can be put on by an able-bodied person in approximately 40 s. This amounts to a total time of around 77 s. With further work, it is possible to reduce the time required to calibrate the sensors and completely eliminate the need for repeated motor calibrations. This could easily reduce the total time required to put the device on below 60 s. This is important as the user story for wearability included being able to put on and take off the device quickly. See Appendix E.1 for video documentation of the start up process, as well as putting on and taking off the device.

Another benefit of the final design is that it accommodates a range of hand and arm sizes. The pattern for the neoprene base is also able to be easily adjusted for arm and hand sizes that may be too large for the current prototype. Smaller sizes are also potentially possible, but may require further refinements to the arm module. This could be a consideration for future work if the device is to be used with children.

Chapter 6

Reflection

A key part of the methodology used in this research was reflection on the design process. This occurred throughout the development process, but was a specific focus during the reflection development phase. This chapter presents the key insights gained through this reflective process.

6.1 Methodological Flexibility

Utilising agile methods for the design process presented a range of challenges. In many ways the methodology for this project was being iteratively improved just as the final artefact itself was. Initially I had planned to perform three system integrations and maintain frequent contact with the physiotherapists. Originally inspired by the Scrum method (Rubin, 2013), I also had planned on utilising time-boxing and scaling features in order to allow the frequent system integrations to occur. However it became evident that some of the processes I was planning to utilise were ill-suited to the project. I had overestimated the amount of production I could achieve in a given time-frame, and when the time came for the first system integration I was not at a point where it made sense to do one. This highlighted the first key observation regarding hardware development: rigid time-boxing is far less effective. Due to the cost in both time and money associated with performing a system integration, shifting features from one development cycle to the next is costly. If a feature is close to completion, it is better to delay the system integration and allow it to be included rather than delay the feature for an entire cycle.

In addition to noting the need for flexible time-boxing constraints, I also observed that I had underestimated the amount of time to devote towards each development cycle. It became clear that I would need a large amount of time to resolve the multiple component cycles to a point where it would be worthwhile integrating them. This is where working as an individual was limiting. In a structured work environment with a team of people it is likely that shorter development cycles would be achievable since components could be worked on concurrently. I also noted that working as an individual made the project particularly susceptible to delays.

6.2 Responding to Delays and Uncertainty

During the first six months of the project, I found that my plans for development were delayed due to an inconsistent schedule. Some weeks I had a large amount of time to devote to the project, while other weeks were taken up with other commitments. This further exacerbated the issues I had identified with time-boxing, since it became difficult to predict when the majority of progress would be made. While this would be somewhat mitigated with a more stable time schedule, even in the context of a business it is likely that unexpected events could cause delays to the development of a project. Even though agile is generally very good at handling these delays, the challenges of hardware design requires that it be given priority where possible. This is particularly true when a solution is not clearly evident at the beginning.

While many elements of the project were well defined and I had an existing prototype as a starting point, there were still many elements of the project that had no firm solutions at the outset. One particular example was the finger modules. It was clear that the existing solution of mounting electronics to a regular glove was not suitable for rehabilitation purposes, however there was no clear way forward from the literature.

As Section 2.4 reveals, virtually every device covered in the literature utilises some kind of textile glove as a support for sensors. The exceptions seemed to be complex, articulated exoskeletons that would likely be time consuming to attach and adjust for different hands. Some solutions utilised plastic rings to hold sensors instead of textile gloves, but these also did not provide an ideal solution. So a number of avenues were explored to search for a better solution.

Exploring multiple avenues for a solution introduced massive uncertainty as it was impossible to predict when an idea might emerge that would solve the problem. It required constant experimentation and time spent investigating a range of possibilities. In the end it was not until two ideas were combined that a promising solution emerged. Up to that point, it was not worthwhile performing a complete system integration as it would not have produced a usable output.

In order to account for the uncertainty inherent in the design process, the time-boxing constraints needed to be loosened and the duration of a single development iteration significantly increased. However this had the potential to greatly reduce the level of user involvement in the project. A solution was necessary that could enable longer development cycles without limiting user involvement.

6.3 Maintaining User Involvement

The increase in development cycle length meant that only a single system integration would be performed, rather than the multiple system integrations originally planned. So that I could maintain user involvement without the need for frequent system integrations, I decided to present the developments of different components for feedback from the physiotherapists. In particular, the textile and finger modules were demonstrated to get ongoing feedback and inform further development. However, once the final system integration was performed a problem with this approach emerged.

While individual components remained the same during the system integration, integrating them together identified additional considerations. This was particularly evident for the arm module. While the arm module and the textile independently worked well, when the two were integrated together the position and weight of the arm module was identified as an area that could be improved. Unfortunately changing the arm module position would require significant reworking of the electronics.

What would have worked better here would have been to use a low-fidelity substitute for the arm module that could approximate the size and weight. This would allow interactions between the arm module and the textile to be estimated before the final system integration. Using low cost approximations of electronics and the associated cases or mounting may allow an agile method to be even more effective in embedded systems development and particularly for developing wearable devices. This observation also highlighted a key difference in the design processes used when developing electronic devices compared to those used for textile design.

6.4 Integrating Different Design Processes

Reflecting on my own design processes, I observed a key difference between the development process used for textiles and that used for electronics. While textile development follows a typically physical process, electronics tend to be simulated virtually before fabrication. This creates a disconnect between the two processes that must be overcome if wearable technologies are to be developed effectively.

I observed that in my own process in particular, when designing electronics and casings I was unable to simulate how the textile would attach. At the same time, when developing the textile element it was necessary to create physical prototypes of the casing for electronics in order to get an idea of how it could attach to the textile. In order to achieve this I utilised 3D printing in order to rapidly produce prototypes of the casing. I used a physical process for the casing development as well, which required me to assemble the electronics first.

In hindsight, it may have been more useful to produce an approximation of the electronics casing to provide a rough idea of the size and weight the final module might have. This could then be used to inform the design process for the textile when deciding where to mount the electronics. This is an important consideration because changing the electronics mounting location can influence the circuit design. Once fabricated making changes to the electronics is much more costly, so utilising physical approximations early on may be a very useful technique for iteratively improving the integration of electronics and textile. It also provides a technique for including users in the design process.

Collaboration between the designer and user is one of the key goals of an agile methodology. Frequent interaction with users helps to reduce uncertainty and increases the likelihood of a final product meeting user expectations. However with hardware development and the simulation processes often utilised it is very difficult to involve the user. Physical prototypes provide a very useful tool for helping a user make sense of the development process.

Creating prototypes that are not fully functional but rather approximate the look and feel of a final device can help to identify design considerations which may have an impact on the electronics components before they are fabricated. When a user can see and feel what a device would be like, they can point out things that may not be apparent from computer renderings and simulations. It is the physical equivalent of a sketch or mock up.

During my development process I utilised some sketched pieces, but neglected to sketch the physical characteristics of the electronics I was designing. The result was receiving feedback from users at a point in development where I could not easily respond to it. This issue is particularly relevant to wearable technologies, where the interface between device and body is much more important than typical electronic devices.

Chapter 7

Conclusion

Starting from an existing haptic feedback glove designed to explore human-computer interaction in a gaming environment, this research has explored how modern technologies can be utilised to support rehabilitation of the hand. An iterative development process was utilised to allow the design to change as the challenges associated with developing wearable devices for a rehabilitation context became clearer. Isolating different device modules allowed rapid prototyping and iterative development of various elements to proceed without generating significant time and monetary costs associated with producing fully functional prototypes. This development process also adapted in response to challenges encountered with iterative hardware development.

An investigation into how biofeedback can be utilised to support rehabilitation highlighted the importance of combining sensor data to produce meaningful and intuitive feedback to the user so that they would not be overloaded with information. Using multimodal systems tailored to communicate information using a relevant modality was one solution that has been utilised to great effect in existing research.

Wearable sensing and haptic feedback technologies relevant to rehabilitation of the hand were then explored. Existing glove-based solutions provided insights into how different technologies could be integrated into wearable devices. The limitations of these systems in a rehabilitation context were also considered.

Collaboration with physiotherapists provided valuable information throughout the design process. An initial workshop revealed key elements of the previously existing prototype that would need to be adapted for rehabilitation use, in particular the wearability of the device. This became a focus for development throughout this research. The sensing and haptic feedback modules for the fingers were redesigned so each finger module could be put on individually using only one hand. The main processing module and battery was integrated into a unit mounted on the arm, while a smaller hand module contained a sensor and combined connections from the fingers into a single link to the arm unit. A variety of textile options were also explored as a base for the hand and arm units. The final solution utilised neoprene and wrapped around the hand and arm, allowing it to fit a variety of hand and arm sizes while also being easy to put on.

The final prototype produced represents a significant development towards a wearable haptic feedback device that could be utilised in a rehabilitation context, responding primarily to two limitations of existing commercial solutions highlighted in Section 2.4.2. Firstly, by designing the device to be put on in stages, the final prototype responds to the challenge of putting on and taking off gloves when a person has limited hand mobility. Secondly, it responds to the issue of limited sizing options by utilising a design that is easily adapted for different hand sizes. The research was limited to physical device development, and has not detailed specific rehabilitation exercises or interactive environments that might utilise the device. While future work could further improve the device based on physiotherapist feedback, the research is now at a stage where it would be beneficial to include patients in the design process and to further explore exactly how the device could be used in a rehabilitation context.

7.1 Future Work

Future work could address a number of limitations both with regard to the device itself and the supporting software. Potential improvements can be broken down into three categories: hardware, embedded firmware, and supporting software. During the later stages of the project, there were two particular improvements suggested that could not be achieved in the development cycle for this research.

Firstly, it was recommended to reposition the arm module to a position higher up the arm in order to further reduce the perceived weight of the device. This was considered at the time it was proposed, however doing so would require changes to the main board and the addition of another sensor near the wrist. It would also increase the connection lengths between the main board and the other parts of the device, which could introduce issues for the communication and power supply between different parts of the device. As a result, this was determined to be too significant to include in the current development cycle, and is considered future work.

Secondly, a concern was raised regarding the ability for the finger modules to accommodate different sized hands. While the neoprene textile is designed to be easily adjusted so that a range of sizing options could be made available, responding to the need for different sizes of finger modules presents additional challenges. One ideal solution could be to have plugs for the modules so different module sizes can be used interchangeably, however this would require a change to the connections between the finger and wrist modules. Although a potential connection solution was identified (see Figure 39), the changes required to implement the different connector options placed it outside of the scope of this research.

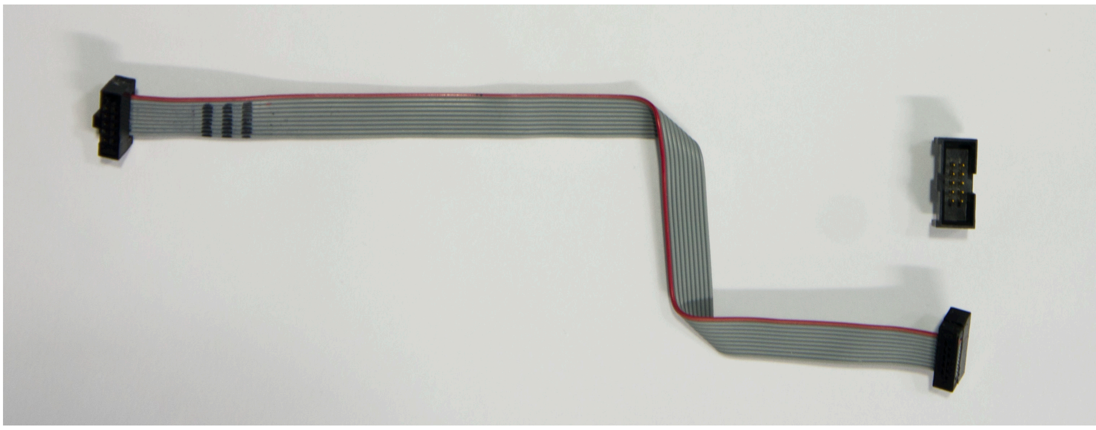

Figure 39. Pluggable wiring solution. A potential wiring option that would allow interchangeable finger modules by enabling them to be plugged in by the user.

In addition to the potential improvements highlighted by the physiotherapists, hardware features such as on-device storage for independent data recording could be added, enabling the device to store motion data even when not connected to a computer. Firmware improvements could also be made that improve the sensor fusion and calibration algorithms, and add additional functionality such as custom haptic waveforms. Supporting software also needs to be developed for specific applications such as rehabilitation games or controlling virtual environments.

Future use of the device in a rehabilitation context would require some form of clinical trials to establish if and how it could be used to support the rehabilitation process. This would extend existing research that explores the application of various technologies including haptic feedback and data gloves to support rehabilitation goals (Boian et al., 2002; Timmermans, Seelen, Willmann & Kingma, 2009; Beursgens, Timmermans & Markopoulos, 2012; A.r et al., 2014; Jacobs, Timmermans, Michielsen, Vander Plaetse & Markopoulos, 2013). Continued research exploring how the device could be used would also inform the development process, which would need to respond to specific user needs and device limitations identified in clinical research.

References

Åkerfeldt, M., Lund, A. & Walkenström, P. (2015). Textile sensing glove with piezoelectric PVDF fibers and printed electrodes of PEDOT:PSS. *Textile Research Journal*, *85*(17), 1789–1799. doi:10.1177/0040517515578333

Antonin Viau, A. G. F. (2005). Reaching in reality and virtual reality: a comparison of movement kinematics in healthy subjects and in adults with hemiparesis. *Journal of neuroengineering and rehabilitation*, *1*(1), 1–7. doi:10.1186/1743-0003-1-11

A.r, A., Wan Khairunizam, W. A., Nazrul Hamizi, A., Shahriman, A. B., Hazry, D., Zuradzman, M. R., ... M. Hazwan, A. (2014). Glove based virtual reality (VR) interaction for the purpose of rehabilitation. *Australian Journal of Basic and Applied Sciences*, *8*(4), 170–175. Retrieved from http://dspace.unimap.edu.my:80/xmlui/handle/123456789/35407

Augurelle, A.-S., Smith, A. M., Lejeune, T. & Thonnard, J.-L. (2003). Importance of cutaneous feedback in maintaining a secure grip during manipulation of hand-held objects. *Journal of Neurophysiology*, *89*(2), 665–671. doi:10.1152/jn.00249.2002

Beursgens, L., Timmermans, A. & Markopoulos, P. (2012). Playful ARM hand training after stroke. *CHI '12 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '12, 2399–2404. doi:10.1145/2212776.2223809

Boian, R., Sharma, A., Han, C., Merians, A., Burdea, G., Adamovich, S., ... Poizner, H. et al. (2002). Virtual reality-based post-stroke hand rehabilitation. *Studies in health technology and informatics*, 64–70. Retrieved from http://ti.rutgers.edu/publications/papers/2002_mmvr_boian.pdf

Bolognini, N., Russo, C. & Edwards, D. J. (2016). The sensory side of post-stroke motor rehabilitation. *Restorative Neurology and Neuroscience*, 1–16. Retrieved from http://content.iospress.com/articles/restorative-neurology-and-neuroscience/rnn150606

Brown, R. (2012, April 26). Managing technical dependencies [Agile coach journal]. Retrieved May 11, 2016, from http://www.agilecoachjournal.com/2012-04-26/managing-technical-dependencies

Cavallo, A., Cirillo, A., Cirillo, P., De Maria, G., Falco, P., Natale, C. & Pirozzi, S. (2014). Experimental comparison of sensor fusion algorithms for attitude estimation. *Proceedings of the 19th IFAC World Congress*, *19*, 7585–7591. doi:10.3182/20140824-6-ZA-1003.01173

Chen, Y., Huang, H., Xu, W., Wallis, R. I., Sundaram, H., Rikakis, T., ... He, J. (2006). The design of a real-time, multimodal biofeedback system for stroke patient rehabilitation. *Proceedings of the 14th ACM International Conference on Multimedia*. MM '06, 763–772. doi:10.1145/1180639.1180804

Cohn, M. (2004). *User stories applied: for agile software development*. Boston: Addison-Wesley.

Dipietro, L., Sabatini, A. & Dario, P. (2008). A survey of glove-based systems and their applications. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, *38*(4), 461–482. doi:10.1109/TSMCC.2008.923862

Foottit, J., Brown, D., Marks, S. & Connor, A. M. (2016). Development of a wearable haptic game interface. *EAI Endorsed Transactions on Creative Technologies*, *3*(6), 151165. doi:10.4108/eai.25-4-2016.151165

Harrison, C., Benko, H. & Wilson, A. D. (2011). OmniTouch: wearable multitouch interaction everywhere. *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*. UIST '11, 441–450. doi:10.1145/2047196.2047255

Heo, P., Gu, G. M., Lee, S.-j., Rhee, K. & Kim, J. (2012). Current hand exoskeleton technologies for rehabilitation and assistive engineering. *International Journal of Precision Engineering and Manufacturing*, *13*(5), 807–824. doi:10.1007/s12541-012-0107-2

Hoshi, T., Iwamoto, T. & Shinoda, H. (2009). Non-contact tactile sensation synthesized by ultrasound transducers. *Proceedings of EuroHaptics conference, 2009 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2009. Third Joint*, 256–260. doi:10.1109/WHC.2009.4810900

Huang, H., Wolf, S. L. & He, J. (2006). Recent developments in biofeedback for neuromotor rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, *3*, 1–12. doi:10.1186/1743-0003-3-11

Huang, K., Do, E. Y. L. & Starner, T. (2008). PianoTouch: a wearable haptic piano instruction system for passive learning of piano skills. *Proceedings of 12th IEEE International Symposium on Wearable Computers*, 41–44. doi:10.1109/ISWC.2008.4911582

Invensense. (n.d.). MPU-9250 | InvenSense [MPU9250]. Retrieved May 15, 2016, from http://www.invensense.com/products/motion-tracking/9-axis/mpu-9250/

Iwamoto, T., Akaho, D. & Shinoda, H. (2004). High resolution tactile display using acoustic radiation pressure. *Proceedings of SICE 2004 Annual Conference*, *2*, 1239–1244. Retrieved from http://www.hapis.k.u-tokyo.ac.jp/public/hiroyuki_shinoda/research/pdf/04SICE_iwamoto.pdf

Jacobs, A., Timmermans, A., Michielsen, M., Vander Plaetse, M. & Markopoulos, P. (2013). CONTRAST: gamification of arm-hand training for stroke survivors. *CHI '13 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '13, 415–420. doi:10.1145/2468356.2468430

Kaisti, M., Rantala, V., Mujunen, T., Hyrynsalmi, S., Könnölä, K., Mäkilä, T. & Lehtonen, T. (2013). Agile methods for embedded systems development - a literature review and a mapping study. *EURASIP Journal on Embedded Systems*, *2013*(1), 1–16. doi:10.1186/1687-3963-2013-15

Karen Moltenbrey. (2006). Tools of the trade [Computer graphics world]. Retrieved from http://www.cgw.com/Publications/CGW/2006/Volume-29-Issue-11-Nov-2006-/Tools-of-the-Trade.aspx

Kent Beck, Mike Beedle, Arie van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, ... Dave Thomas. (2001). Manifesto for agile software development. Retrieved from http://www.agilemanifesto.org/

Krishna, S., Bala, S., McDaniel, T., McGuire, S. & Panchanathan, S. (2010). VibroGlove: an assistive technology aid for conveying facial expressions. *CHI '10 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '10, 3637–3642. doi:10.1145/1753846.1754031

Latessa, G., Brunetti, F., Reale, A., Saggio, G. & Di Carlo, A. (2009, June 4). Piezoresistive behaviour of flexible PEDOT:PSS based sensors. *Sensors and Actuators B: Chemical*, *139*(2), 304–309. doi:10.1016/j.snb.2009.03.063

Madgwick, S. O. (2010). An efficient orientation filter for inertial and inertial/magnetic sensor arrays. *Report x-io and University of Bristol (UK).*

Mahony, R., Hamel, T. & Pflimlin, J. M. (2008, June). Nonlinear complementary filters on the special orthogonal group. *IEEE Transactions on Automatic Control*, *53*(5), 1203–1218. doi:10.1109/TAC.2008.923738

Marasco, P. D., Kim, K., Colgate, J. E., Peshkin, M. A. & Kuiken, T. A. (2011). Robotic touch shifts perception of embodiment to a prosthesis in targeted reinnervation amputees. *Brain*, *134*(3), 747–758. doi:10.1093/brain/awq361

Merians, A. S., Jack, D., Boian, R., Tremaine, M., Burdea, G. C., Adamovich, S. V., ... Poizner, H. (2002). Virtual reality–augmented rehabilitation for patients following stroke. *Physical Therapy*, *82*(9), 898–915. Retrieved from http://ptjournal.apta.org/content/82/9/898

Microsoft. (n.d.). Kinect sensor. Retrieved May 31, 2016, from https://msdn.microsoft.com/en-us/library/hh438998.aspx

Mistry, P., Maes, P. & Chang, L. (2009). WUW - wear ur world: a wearable gestural interface. *CHI '09 Extended Abstracts on Human Factors in Computing Systems*. CHI EA '09, 4111–4116. doi:10.1145/1520340.1520626

Moreira, A. H. J., Queirós, S., Fonseca, J., Rodrigues, P. L., Rodrigues, N. F. & Vilaca, J. L. (2014). Real-time hand tracking for rehabilitation and character animation. *Proceedings of 2014 IEEE 3rd International Conference on Serious Games and Applications for Health (SeGAH)*, 1–8. doi:10.1109/SeGAH.2014.7067086

Munoz Diaz, E., de Ponte Muller, F., Jiménez, A. R. & Zampella, F. (2015). Evaluation of AHRS algorithms for inertial personal localization in industrial environments. *Proceedings of 2015 IEEE International Conference on Industrial Technology (ICIT)*, 3412–3417. doi:10.1109/ICIT.2015.7125605

Nathan, D. & Johnson, M. (2007). Design of a grasp assistive glove for ADL-focused, robotic assisted therapy after stroke. *Proceedings of IEEE 10th International Conference on Rehabilitation Robotics, 2007.* 943–950. doi:10.1109/ICORR.2007.4428537

O'Flynn, B., Sanchez, J. T., Angove, P., Connolly, J., Condell, J., Curran, K. & Gardiner, P. (2013). Novel smart sensor glove for arthritis rehabiliation. *Proceedings of*

2013 IEEE International Conference on Body Sensor Networks (BSN), 1–6. doi:10.1109/BSN.2013.6575482

5DT Products. (n.d.). Retrieved August 30, 2016, from http://www.5dt.com/products/pdataglovemri.html

CyberGlove II. (n.d.). Retrieved August 30, 2016, from http://www.cyberglovesystems.com/cyberglove-ii/

Data Gloves | 5DT. (n.d.). Retrieved May 11, 2016, from http://www.5dt.com/?page_id=34

Gloveone: The haptic glove to feel & touch real virtual reality. (n.d.). Retrieved June 29, 2016, from https://www.gloveonevr.com/

Gyroglove | GyroGear. (n.d.). Retrieved March 16, 2016, from http://gyrogear.co/gyroglove/

Haptic Workstation. (n.d.). Retrieved May 31, 2016, from http://www.cyberglovesystems.com/haptic-workstation/

Home. (n.d.). Retrieved May 11, 2016, from http://www.cyberglovesystems.com/

OptiTrack for Virtual Reality. (n.d.). Retrieved September 8, 2016, from https://www.optitrack.com/motion-capture-virtual-reality/

Teensy USB Development Board. (n.d.). Retrieved May 16, 2016, from https://www.pjrc.com/store/teensy32.html

Control VR- The Future of Virtual Reality, Animation & more. (2014). Retrieved May 25, 2016, from https://www.kickstarter.com/projects/controlvr/control-vr-motion-capture-for-vr-animation-and-mor

Parkin, S. (2016). Hope in a glove for parkinson's patients [MIT technology review]. Retrieved March 16, 2016, from https://www.technologyreview.com/s/545456/hope-in-a-glove-for-parkinsons-patients/

Robles-De-La-Torre, G. (2006). The importance of the sense of touch in virtual and real environments. *IEEE MultiMedia*, *13*(3), 24–30. doi:10.1109/MMUL.2006.69

Rubin, K. S. (2013). *Essential scrum: a practical guide to the most popular agile process*. Upper Saddle River, NJ: Addison-Wesley.

Saggio, G. (2014). A novel array of flex sensors for a goniometric glove. *Sensors and Actuators A: Physical*, *205*, 119–125. doi:10.1016/j.sna.2013.10.030

Tan, D. W., Schiefer, M. A., Keith, M. W., Anderson, J. R., Tyler, J. & Tyler, D. J. (2014). A neural interface provides long-term stable natural touch perception. *Science Translational Medicine*, *6*(257), 138–257. doi:10.1126/scitranslmed.3008669

Texas Instruments Incorporated. (2014). DRV2605l 2 to 5.2 v haptic driver for LRA and ERM with effect library and smart-loop architecture (rev. c). Texas Instruments. Retrieved September 14, 2015, from http://www.ti.com/lit/gpn/drv2605l

Timmermans, A. A., Seelen, H. A., Willmann, R. D. & Kingma, H. (2009). Technology-assisted training of arm-hand skills in stroke: concepts on reacquisition of motor control and therapist guidelines for rehabilitation technology design. *Journal of neuroengineering and rehabilitation*, *6*, 18. doi:10.1186/1743-0003-6-1

Vutinuntakasame, S., Jaijongrak, V.-R. & Thiemjarus, S. (2011). An assistive body sensor network glove for speech- and hearing-impaired disabilities. *Proceedings* *of 2011 International Conference on Body Sensor Networks (BSN)*, 7–12. doi:10.1109/BSN.2011.13

Watanabe, T. & Fukui, S. (1995). A method for controlling tactile sensation of surface roughness using ultrasonic vibration. *Proceedings of 1995 IEEE International Conference on Robotics and Automation. 1*, 1134–1139. doi:10.1109/ROBOT.1995.525433

Xia, L., Chen, C. C. & Aggarwal, J. K. (2011). Human detection using depth information by kinect. *CVPR 2011 workshops*, 15–22. doi:10.1109/CVPRW.2011.5981811

x-io Technologies. (2012). Open source IMU and AHRS algorithms | x-io technologies. Retrieved from http://www.x-io.co.uk/open-source-imu-and-ahrs-algorithms/

Glossary

AC alternating current. 27

AUT Auckland University of Technology. 45

biofeedback the use of instrumentation to make covert physiological processes overt. 17–19, 33, 92

Bluetooth a wireless communication protocol commonly used in portable consumer electronics. 44, 73, 81

CSV comma separated value. 82

DC direct current. 27

embedded system a computer system with a dedicated function within a larger mechanical or electrical system, often with real-time computing constraints. 36

ERM eccentric rotating mass. 27, 44, 50

FDM fused deposition modeling. 58

FIFO first in first out. 72

haptic feedback any form of output that is designed to interact with the human sense of touch. 5, 14, 17, 19, 24, 25, 27, 28, 30, 33, 114

I^2C inter-integrated circuit. 50, 54, 72

IMU inertial measurement unit. 23, 31, 33, 44, 50, 80, 84

JST Japan solderless terminal. 61, 69, 70

kickstarter a crowd funding platform that connects financial backers with creative projects looking for funding. 30, 31

LED light emitting diode. 60, 82

LiPo lithium polymer. 56

LRA linear resonance actuator. 27

MEMS microelectromechanical system. 23

MRI magnetic resonance imaging. 30

myoelectric signal an electrical impulse that produces contraction of muscle fibres in the body. Also known as a motor action potential. 20, 22

PCB printed circuit board. 11, 33, 50, 62, 104, 106, 107

SLS selective laser sintering. 58

SMD surface mount device. 33, 60

SPI serial peripheral interface. 50, 54, 72

USB universal serial bus. 73, 82

Appendix A

Electronic Circuit Schematics and Board Layouts

A.1 Sensor Board

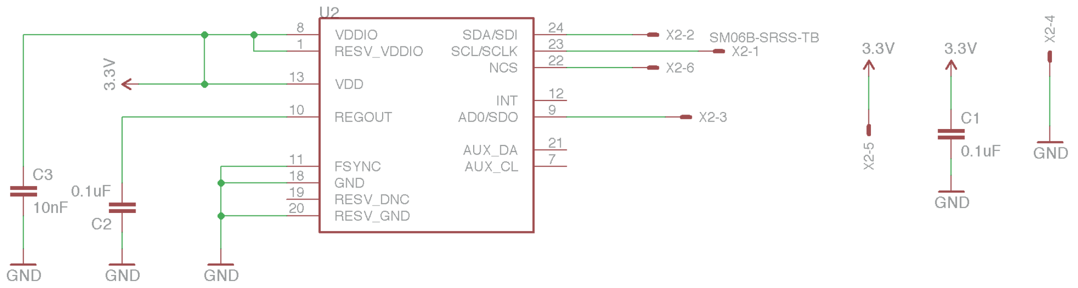

Figure A1. Sensor Board Schematic. This shows the electrical schematic for a single sensor board (used for the finger and thumb sensors).

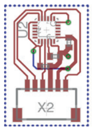

Figure A2. Sensor Board PCB Layout. 1:1 scale image of PCB layout for a single sensor board.

A.2 Main Board

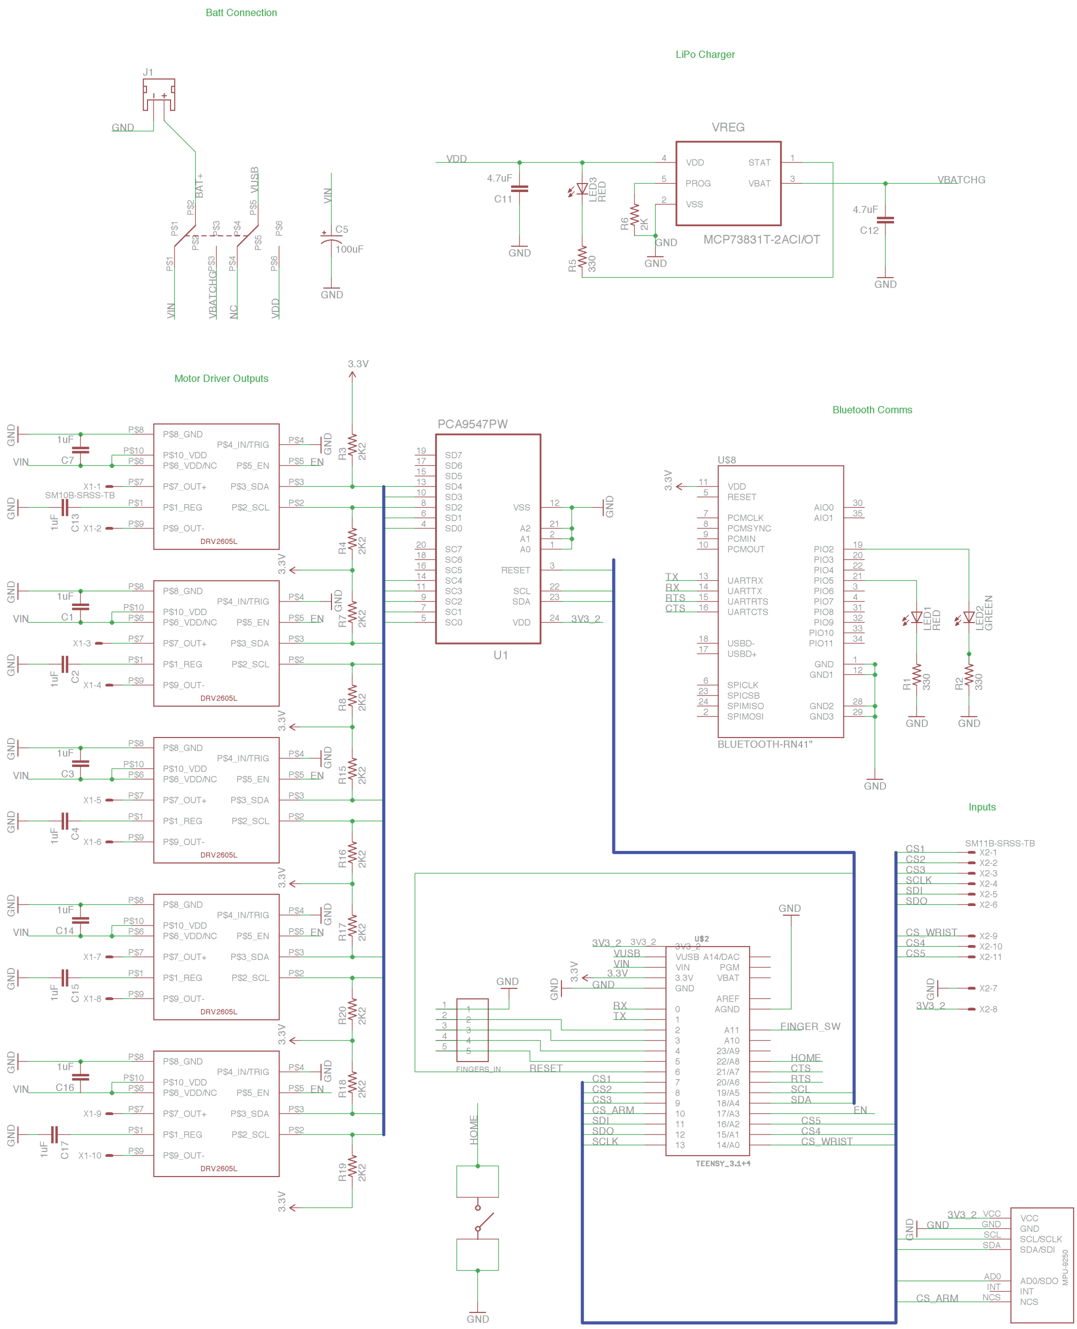

Figure A3. Main Board Schematic. The electrical schematic for the main board, that was mounted on the arm.

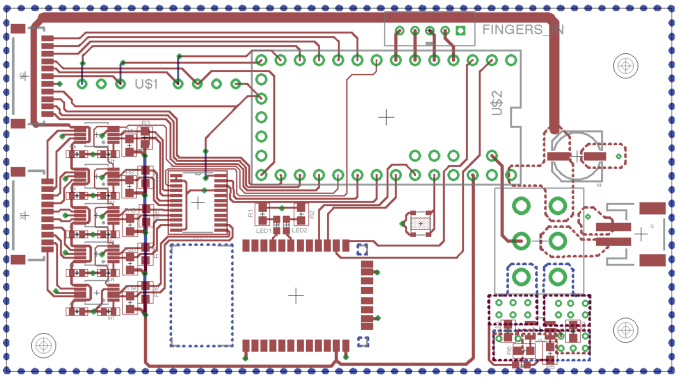

Figure A4. Main Board PCB Layout. 1:1 scale image of the PCB layout for the main board.

A.3 Wrist Board

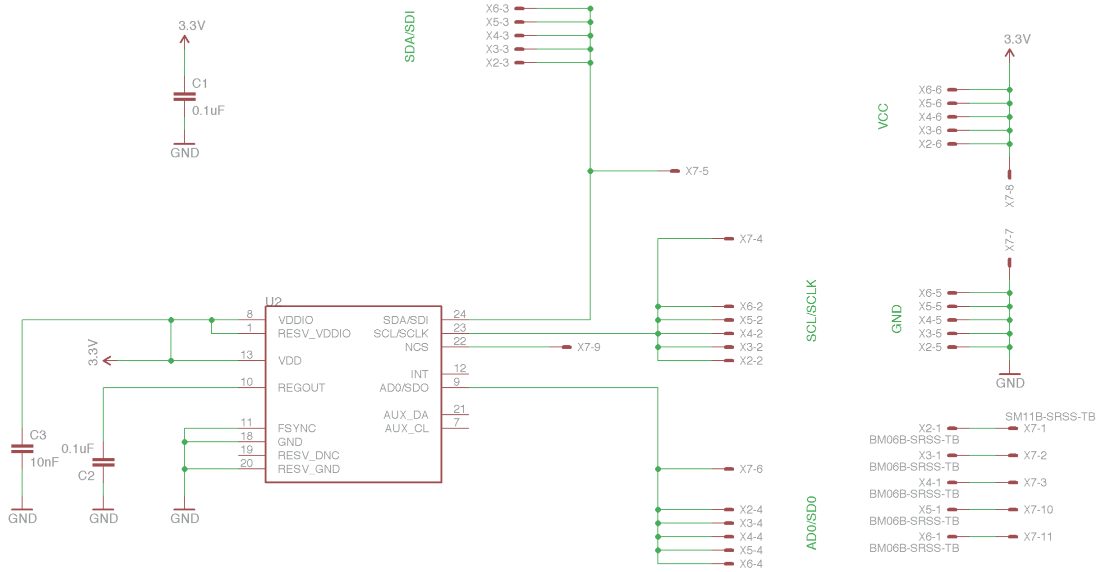

Figure A5. Wrist Board Schematic. The electrical schematic for the sensor board mounted on the back of the hand. Note this board also combines the connections from each finger module into a single data connection back to the main board.

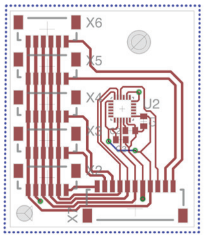

Figure A6. Wrist Board PCB Layout. 1:1 scale image of the PCB layout for the sensor board for the back of the hand.

A.4 Battery Connector Board

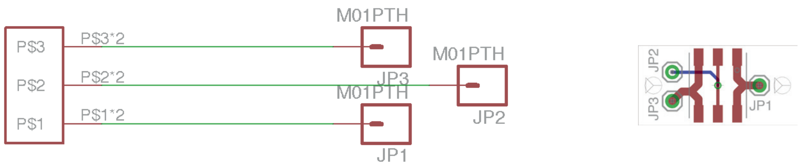

Figure A7. Battery Connector Board Schematic and PCB Layout. Electrical schematic and 1:1 scale image of the PCB layout for the battery connector board.

Appendix B

Serial Communication Protocol

Quaternion Data Packet									
Length: 22 Bytes									
Packet Header				**Data Payload**				**Packet End**	
				Quaternion (float)					
Open	*Type*	*Id*	*Close*					*CR*	*NL*
$	q	i	:	Qw[0..3]	Qx[0..3]	Qy[0..3]	Qz[0..3]	\r	\n
0	1	2	3	4..7	8..11	12..15	16..19	20	21

Figure B1. Quaternion Data Packet. Details the packet structure for quaternion data that is transmitted from the device to a paired computer.

Raw Data Packet								
Length: 36 Bytes								
Packet Header				**Data Payload**				
Open	*Type*	*Id*	*Close*	*Timest (long)*	*Temp (int)*	*Accelerometer (int)*		
$	r	i	:	Ts[0..3]	T[0..1]	Ax[0..1]	Ay[0..1]	Az[0..1]
0	1	2	3	4..7	8..9	10..11	12..13	14..15

Data Payload						**Packet End**	
Gyroscope (int)			*Magnetometer (float)*			*CR*	*NL*
Gx[0..1]	Gy[0..1]	Gz[0..1]	Mx[0..3]	My[0..3]	Mz[0..3]	\r	\n
16..17	18..19	20..21	22..25	26..29	30..33	34	35

Figure B2. Raw Data Packet. Details the packet structure for when raw sensor data is transmitted from the device to a paired computer.

Calibrated Data Packet

Length: 42 Bytes

Packet Header				Data Payload		
Open	*Type*	*Id*	*Close*	*Accelerometer (float)*		
$	c	i	:	Ax[0..3]	Ay[0..3]	Az[0..3]
0	1	2	3	4..7	8..11	12..15

Data Payload						Packet End	
Gyroscope (float)			*Magnetometer (float)*			*CR*	*NL*
Gx[0..3]	Gy[0..3]	Gz[0..3]	Mx[0..3]	My[0..3]	Mz[0..3]	\r	\n
16..19	20..23	24..27	28..31	32..35	36..39	40	41

Figure B3. Calibrated Data Packet. Details the packet structure for when calibrated sensor data is transmitted from the device to a paired computer.

Request Data Packet

Length: 7 Bytes

Packet Header				Data Payload		Packet End	
Open	*Type*	*Id M*	*Close*	*Data to Send*	*Qty**	*CR*	*NL*
$	d	i	:	Type Mask	uint8	\r	\n
0	1	2	3	4	5	6	7

***Note: Qty = 0 means toggle stream on/off**

ID Masks

Hex	*Label*
0x01	Thumb
0x02	Index Finger
0x04	Middle Finger
0x08	Ring Finger
0x10	Little Finger
0x20	Hand
0x40	Arm

Type Masks

Hex	*Label*
0x01	Quaternion
0x02	Raw Data
0x04	Calibrated Data

Figure B4. Request Data Packet. Details the packet structure for requesting sensor data from the device.

Device Mode Packet

Length: 7 Bytes

Packet Header				Data Payload	Packet End	
Open	*Type*	*Resv*	*Close*	*Mode*	*CR*	*NL*
$	c	0	:	Mode Mask	\r	\n
0	1	2	3	4	5	6

Mode Masks

Hex	*Label*
0x01	Standard
0x02	Configuration
0x03	Standby
0x04	Version

Figure B5. Device Mode Packet. Details the packet structure for setting the device mode.

Motor Command Packet

Length: 7 – 9 Bytes

Packet Header				Data Payload			Packet End	
Open	*Type*	*Id M*	*Close*	*Mode*	*Intensity**	*Duration**	*CR*	*NL*
$	m	i	:	Mode Mask	Int Value	Milliseconds	\r	\n
0	1	2	3	4	5	6	7	8
					*Custom Mode Only			

ID Masks

Hex	*Label*
0x01	Thumb
0x02	Index Finger
0x04	Middle Finger
0x08	Ring Finger
0x10	Little Finger

Mode Masks

Hex	*Label*
0x01	Stop
0x02	Pulse
0x04	On
0x08	Custom*

Figure B6. Motor Command Packet. Details the packet structure for instructing the device to output a haptic vibration.

Appendix C

Wire Stretch Tests

While testing options for stretchable wiring solutions, a simple experiment was performed to evaluate the potential of stitching enamelled wire into a neoprene substrate. The neoprene provided an elastic, stretchable base, while the stitch structure of the enamelled wire allowed it to deform so that it would increase in length when pulled. The idea was that as the neoprene base was stretched the width would decrease, resulting in an increased length due to the zig-zag stitch structure.

A series of 5 neoprene strips were hand stitched with enamelled wire in different patterns and spacings to test their ability to stretch and return to the original shape. A strand of enamelled wire was threaded through a sewing needle and used to make hand stitches at spacings of $1\,\mathrm{cm}$ and $0.5\,\mathrm{cm}$. A single and double stitch structure was used. The single stitch consisted of a single, zig-zag stitch of wire. The double-stitch combined two single-stitches started at opposite sides of the strip to form an 'x' pattern. Each stitch spacing was tested with both stitch structures. In addition, the $0.5\mathrm{cm}$ stitch spacing was tested with a double density single stitch, where two enamelled wires were threaded through the sewing needle instead of one.

Once the enamelled wires were stitched into the neoprene substrate, they were stretched by hand until they were subjectively assessed to be at their limit. Measurements were taken with the strip stretched and after it was released. The rigidity of enamelled wire meant that the strips did not return to their original length after being stretched, so only the length achieved after stretching was recorded. Approximately $5.8\,\mathrm{N}$ of force was applied to reach maximum stretch.

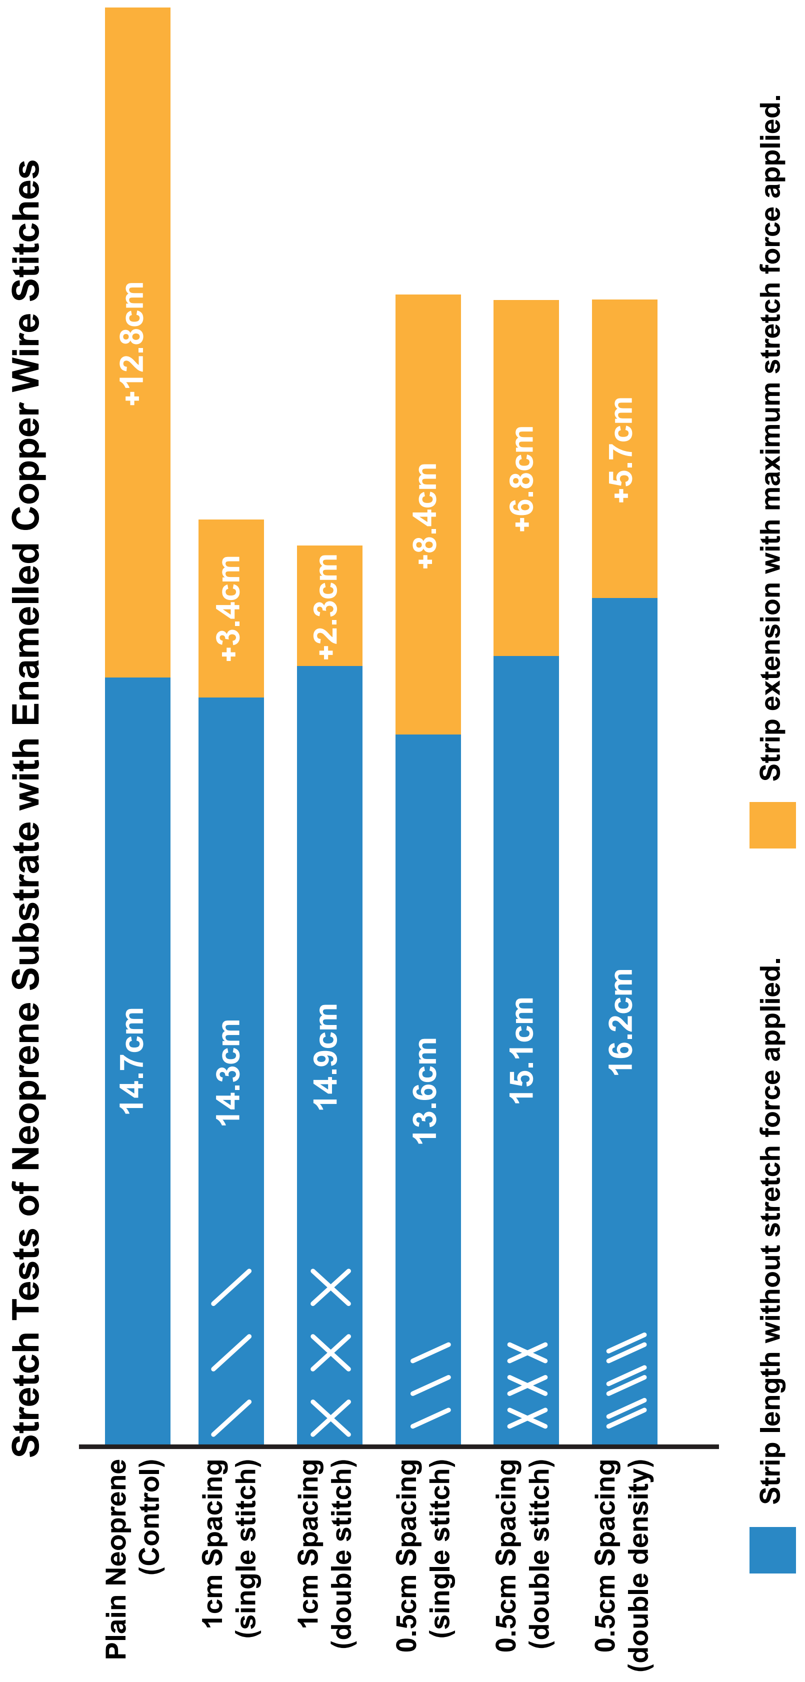
Stretch Tests of Neoprene Substrate with Enamelled Copper Wire Stitches
Plain Neoprene (Control)
1cm Spacing (single stitch)
1cm Spacing (double stitch)
0.5cm Spacing (single stitch)
0.5cm Spacing (double stitch)
0.5cm Spacing (double density)
14.7cm
+12.8cm
14.3cm
+3.4cm
14.9cm
+2.3cm
13.6cm
+8.4cm
15.1cm
+6.8cm
16.2cm
+5.7cm
Strip length without stretch force applied.
Strip extension with maximum stretch force applied.

Appendix D

Haptic Effects Libraries

The datasheet supplied by Texas Instruments Incorporated for the DRV2605L Haptic Motor Driver documents the haptic feedback effects included in the device (Texas Instruments Incorporated, 2014). The device contains 7 different libraries, each of which contain the same effects but with differing output waveforms to suit a range of vibration motors. The 123 different effects possible are listed below.

EFFECT ID NO.	WAVEFORM NAME	EFFECT ID NO.	WAVEFORM NAME
1	Strong Click - 100%	64	Transition Hum 1 - 100%
2	Strong Click - 60%	65	Transition Hum 2 - 80%
3	Strong Click - 30%	66	Transition Hum 3 - 60%
4	Sharp Click - 100%	67	Transition Hum 4 - 40%
5	Sharp Click - 60%	68	Transition Hum 5 - 20%
6	Sharp Click - 30%	69	Transition Hum 6 - 10%
7	Soft Bump - 100%	70	Transition Ramp Down Long Smooth 1 - 100 to 0%
8	Soft Bump - 60%	71	Transition Ramp Down Long Smooth 2 - 100 to 0%
9	Soft Bump - 30%	72	Transition Ramp Down Medium Smooth 1 - 100 to 0%
10	Double Click - 100%	73	Transition Ramp Down Medium Smooth 2 - 100 to 0%
11	Double Click - 60%	74	Transition Ramp Down Short Smooth 1 - 100 to 0%

EFFECT ID NO.	WAVEFORM NAME	EFFECT ID NO.	WAVEFORM NAME
12	Triple Click - 100%	75	Transition Ramp Down Short Smooth 2 - 100 to 0%
13	Soft Fuzz - 60%	76	Transition Ramp Down Long Sharp 1 - 100 to 0%
14	Strong Buzz - 100%	77	Transition Ramp Down Long Sharp 2 - 100 to 0%
15	750 ms Alert 100%	78	Transition Ramp Down Medium Sharp 1 - 100 to 0%
16	1000 ms Alert 100%	79	Transition Ramp Down Medium Sharp 2 - 100 to 0%
17	Strong Click 1 - 100%	80	Transition Ramp Down Short Sharp 1 - 100 to 0%
18	Strong Click 2 - 80%	81	Transition Ramp Down Short Sharp 2 - 100 to 0%
19	Strong Click 3 - 60%	82	Transition Ramp Up Long Smooth 1 - 0 to 100%
20	Strong Click 4 - 30%	83	Transition Ramp Up Long Smooth 2 - 0 to 100%
21	Medium Click 1 - 100%	84	Transition Ramp Up Medium Smooth 1 - 0 to 100%
22	Medium Click 2 - 80%	85	Transition Ramp Up Medium Smooth 2 - 0 to 100%
23	Medium Click 3 - 60%	86	Transition Ramp Up Short Smooth 1 - 0 to 100%
24	Sharp Tick 1 - 100%	87	Transition Ramp Up Short Smooth 2 - 0 to 100%
25	Sharp Tick 2 - 80%	88	Transition Ramp Up Long Sharp 1 - 0 to 100%
26	Sharp Tick 3 - 60%	89	Transition Ramp Up Long Sharp 2 - 0 to 100%
27	Short Double Click Strong 1 - 100%	90	Transition Ramp Up Medium Sharp 1 - 0 to 100%
28	Short Double Click Strong 2 - 80%	91	Transition Ramp Up Medium Sharp 2 - 0 to 100%
29	Short Double Click Strong 3 - 60%	92	Transition Ramp Up Short Sharp 1 - 0 to 100%
30	Short Double Click Strong 4 - 30%	93	Transition Ramp Up Short Sharp 2 - 0 to 100%
31	Short Double Click Medium 1 - 100%	94	Transition Ramp Down Long Smooth 1 - 50 to 0%

EFFECT ID NO.	WAVEFORM NAME	EFFECT ID NO.	WAVEFORM NAME
32	Short Double Click Medium 2 - 80%	95	Transition Ramp Down Long Smooth 2 - 50 to 0%
33	Short Double Click Medium 3 - 60%	96	Transition Ramp Down Medium Smooth 1 - 50 to 0%
34	Short Double Sharp Tick 1 - 100%	97	Transition Ramp Down Medium Smooth 2 - 50 to 0%
35	Short Double Sharp Tick 2 - 80%	98	Transition Ramp Down Short Smooth 1 - 50 to 0%
36	Short Double Sharp Tick 3 - 60%	99	Transition Ramp Down Short Smooth 2 - 50 to 0%
37	Long Double Sharp Click Strong 1 - 100%	100	Transition Ramp Down Long Sharp 1 - 50 to 0%
38	Long Double Sharp Click Strong 2 - 80%	101	Transition Ramp Down Long Sharp 2 - 50 to 0%
39	Long Double Sharp Click Strong 3 - 60%	102	Transition Ramp Down Medium Sharp 1 - 50 to 0%
40	Long Double Sharp Click Strong 4 - 30%	103	Transition Ramp Down Medium Sharp 2 - 50 to 0%
41	Long Double Sharp Click Medium 1 - 100%	104	Transition Ramp Down Short Sharp 1 - 50 to 0%
42	Long Double Sharp Click Medium 2 - 80%	105	Transition Ramp Down Short Sharp 2 - 50 to 0%
43	Long Double Sharp Click Medium 3 - 60%	106	Transition Ramp Up Long Smooth 1 - 0 to 50%
44	Long Double Sharp Tick 1 - 100%	107	Transition Ramp Up Long Smooth 2 - 0 to 50%
45	Long Double Sharp Tick 2 - 80%	108	Transition Ramp Up Medium Smooth 1 - 0 to 50%
46	Long Double Sharp Tick 3 - 60%	109	Transition Ramp Up Medium Smooth 2 - 0 to 50%
47	Buzz 1 - 100%	110	Transition Ramp Up Short Smooth 1 - 0 to 50%
48	Buzz 2 - 80%	111	Transition Ramp Up Short Smooth 2 - 0 to 50%
49	Buzz 3 - 60%	112	Transition Ramp Up Long Sharp 1 - 0 to 50%
50	Buzz 4 - 40%	113	Transition Ramp Up Long Sharp 2 - 0 to 50%
51	Buzz 5 - 20%	114	Transition Ramp Up Medium Sharp 1 - 0 to 50%

EFFECT ID NO.	WAVEFORM NAME	EFFECT ID NO.	WAVEFORM NAME
52	Pulsing Strong 1 - 100%	115	Transition Ramp Up Medium Sharp 2 - 0 to 50%
53	Pulsing Strong 2 - 60%	116	Transition Ramp Up Short Sharp 1 - 0 to 50%
54	Pulsing Medium 1 - 100%	117	Transition Ramp Up Short Sharp 2 - 0 to 50%
55	Pulsing Medium 2 - 60%	118	Long buzz for programmatic stopping - 100%
56	Pulsing Sharp 1 - 100%	119	Smooth Hum 1 (No kick or brake pulse) - 50%
57	Pulsing Sharp 2 - 60%	120	Smooth Hum 2 (No kick or brake pulse) - 40%
58	Transition Click 1 - 100%	121	Smooth Hum 3 (No kick or brake pulse) - 30%
59	Transition Click 2 - 80%	122	Smooth Hum 4 (No kick or brake pulse) - 20%
60	Transition Click 3 - 60%	123	Smooth Hum 5 (No kick or brake pulse) - 10%
61	Transition Click 4 - 40%		
62	Transition Click 5 - 20%		
63	Transition Click 6 - 10%		

Appendix E

Supplementary Videos

E.1 Touchable: Start up and Wearability Demo

This video depicts the device start up process, putting it on, and taking it off.
Link: https://vimeo.com/172637167

E.2 Touchable: Operation Demo

This video depicts the device in operation controlling a flight simulator game.
Link: https://vimeo.com/172535712