UNIVERSITY
of
GLASGOW

Tactons: Structured Vibrotactile Messages for Non-Visual Information Display

Lorna Margaret Brown

Submitted for the Degree of Doctor of Philosophy

Department of Computing Science, University of Glasgow

April, 2007

Abstract

This thesis investigates the design of Tactons: structured vibrotactile messages for presenting multidimensional information non-visually. Tactons (tactile icons) are the tactile equivalent of Earcons in audio. While vibrotactile icons have been proposed by other researchers, Tactons are the first to use a structured, abstract approach to encoding multidimensional information. This thesis is the first formal investigation into how Tactons could be designed, and of the identification rates achieved when they are presented to users.

In order to design Tactons and evaluate their effectiveness, two questions need to be answered. Firstly, what are the parameters of vibration that can be manipulated to encode information and, secondly, what levels of performance can be achieved when these parameters are combined to encode multidimensional information? This thesis addresses these questions in two ways. Firstly, the literature on vibrotactile perception is reviewed to identify candidate parameters. Secondly, a series of empirical experiments are conducted to evaluate some of these individual parameters, and to evaluate Tactons which have been created by combining two or more of these parameters.

The review of the literature, along with the results of the experiments on individual parameters, identified a set of parameters which could be used in Tacton design: rhythm, spatial location, roughness, intensity, and intensity change over time. Combinations of these parameters were then used to create Tactons encoding two or three dimensions of information. The results of empirical experiments evaluating two-dimensional Tactons indicated that average identification rates of up to 73% could be achieved when the two dimensions were encoded in rhythm and roughness using a high specification vibrotactile device. Another empirical experiment showed that similar results could be achieved using the low-cost vibration motor featured in mobile phones, by encoding the two dimensions of information in rhythm and intensity. Finally, a series of experiments were conducted to evaluate different combinations of parameters for encoding three dimensions of information in Tactons. These results showed that identification rates ranging from 48% to 54% were achieved when three levels of each parameter (rhythm, spatial location and a third parameter: either roughness, intensity or a combination of both) were used, and that this could be increased to 81% by using just two levels of the roughness parameter.

As this thesis is the first formal investigation into the design of Tactons, the results from this research can be used as a baseline against which future work on Tactons can be compared. In addition, a set of design recommendations has been produced from the results of this research, and can be used by designers or other researchers who wish to use Tactons in their interfaces.

Table of Contents

List of Figures

List of Tables

Acknowledgments

First of all thanks to my supervisor, Stephen Brewster for giving me the opportunity to do this research and for his support throughout my PhD. Thanks also to my second supervisor, Helen Purchase, in particular for her help with experimental methodologies and statistics.

I am very grateful to Roger Cholewiak, who has provided invaluable assistance with many aspects of my work (including proofreading, and providing information on psychophysics, vibrotactile technologies and information transmission). Bruce Mortimer from EAI Inc. has provided excellent support for the C2 Tactors and we have had many useful and interesting discussions about these devices.

I would like to thank Topi Kaaresoja for giving me the opportunity to apply my research in an industrial environment, and for supporting my PhD research while I was an intern at Nokia Research Center.

Thanks to all the members of the Multimodal Interaction Group and my other friends in DCS, for their help with my work, and for being available for coffee or drinks when I needed a break! Special thanks are due to Steve, Andy, Marilyn and Eve for proof-reading chapters of my dissertation. Also, thanks to Jonathan Paisley for his expert help with programming in the early stages of my PhD.

I would like to thank my examiners, Ian Summers and Phil Gray, for their useful feedback and comments on my thesis both during, and following, my viva.

Thanks to my friends for taking my mind off my thesis writing and, especially, to Rebecca Hoyle for making sure that my commas were in all the right places!

My parents know, only too well, what it is like to do a PhD, and I am extremely grateful for their support (technical, emotional and financial!) during mine! They have helped in all manner of ways – piloting experiments, proofreading papers and chapters, and just listening to me.

Last but definitely not least, I would like to thank Steven Wall, for his support and patience which has helped immeasurably throughout the whole of my PhD. In addition, I would like to thank him for the many interesting discussions regarding research ideas, and for an incredible amount of proofreading!

This research was sponsored by EPSRC studentship GR/PO1175/01.

Declaration

Some of the material presented in Chapter 4 has been published in the proceedings of AUIC 2004 [8] and CHI 2004 [13], co-authored with Stephen Brewster.

Experiment 1 in Chapter 5, and Experiment 3(a) in Chapter 6 have been published in World Haptics 2005 [15], co-authored by Stephen A. Brewster and Helen C. Purchase.

Experiment 2 in Chapter 5 has been published in CHI 2006 [16], co-authored by Stephen A. Brewster and Helen C. Purchase.

Experiment 4 in Chapter 7 has been published in CHI 2006 [18], co-authored by Topi Kaaresoja.

Experiment 5(a) and 5(b) in Chapter 8 have been published in Mobile HCI 2006 [17], co-authored by Stephen A. Brewster and Helen C. Purchase.

This thesis only exploits those parts of these papers that are directly attributable to the author. All other referenced material has been given full acknowledgement in the text.

Chapter 1: Introduction

The majority of computer interfaces rely on vision to communicate information to users. However, in complex systems where a large amount of data is presented, the sense of vision can become overloaded. In addition, mobile devices, such as mobile phones and personal digital assistants (PDAs), are becoming increasingly popular. When interacting with these devices, the visual display is often insufficient due to the limited screen space, or inappropriate, such as when the device is in a pocket but requires the user's attention. Therefore, it is important to consider using alternative modalities through which information can be presented. The use of audio for non-visual information display has been widely investigated (e.g. [4, 10, 39]). Another modality which might offer an alternative channel for communication is the sense of touch.

The idea of communicating information through touch is not new. For example, the sense of touch has been used for many years to aid communication for people with sensory impairments. Some deaf-blind people use a technique called Tadoma, or "tactile lipreading" to help them interpret speech. In this technique, a deaf-blind person places his/her hands on the speaker's neck and face and uses the movements of the lips and jaw and the vibrations of the voicebox to identify what is being said [91]. The vibration aspect of the Tadoma method is also used by some people with a hearing impairment to enhance visual lip-reading (http://www.tactaid.com/). Another use of the sense of touch by sensory impaired people is the Braille system, which is used to enable blind and visually impaired people to read text [36]. Technological versions of these techniques also exist, and are known as sensory substitution systems. The tactile sensory substitution systems in existence include tactile hearing aids [74, 101], dynamic Braille displays, and vibrating text displays [5]. Despite the fact that the possibilities of tactile display for communication have been known for some time, the potential of this modality has not yet been fully realised in human computer interaction.

In 1960, Geldard [41] wrote:

> *"For some kinds of messages the skin offers a valuable supplement to ears and eyes."*

In particular, Geldard [40] discussed the potential of mechanical vibrations for communicating information. He proposed that by manipulating different parameters of vibration, vibrotactile stimuli could be used to present a wide variety of information, for example, to advise an airplane pilot to switch fuel tanks, to report weather data, to indicate the state of stocks, or to provide directional information.

Almost 50 years on from Geldard's papers, vibrotactile displays are now common in everyday devices, with mobile phones, handheld computers and games controllers all featuring vibration feedback. However, the vibrations used in such devices are generally very basic and do not fully exploit the potential of vibration as a means of communication. This is now changing with more complex vibration patterns being used (e.g. [25]), but there has still been very little formal work to explore how best to use vibrations in these devices.

Using vibrotactile displays in human computer interaction has some clear advantages. The skin offers a large display space which can be used to display directional information [114] and could, therefore, be useful for navigation systems, for drivers, pilots or visually impaired people. As the skin is less engaged in other tasks than the eyes or ears, it is always ready to receive information [114]. This, along with the fact that vibrations are attention grabbing [41], means that vibrotactile display is a good choice for alerts or warnings. This could be particularly useful in mobile computing, where a user could be alerted to incoming calls or upcoming appointments while mobile, when the other senses would be attending to the sights and sounds around them. Blind people might benefit particularly from this as, unlike audio, it would not block out the environmental noise which they rely on for safe navigation. Another advantage of vibrotactile displays is that they enable discreet, private communication, as they can only be felt by the user. They can, therefore, be used in situations where auditory display would be inappropriate [24], for example in meetings, or while in a quiet environment such as a library. The advantages and disadvantages of vibrotactile display are discussed in detail in Chapter 3.

Researchers in human computer interaction have recently begun to explore the possibilities of vibrotactile display for a range of applications, including sensory substitution (e.g. [42, 101]), alerts (e.g. [57]), navigation (e.g. [45, 111]), communication (e.g. [22, 24]), and enhancing elements of desktop and mobile user interfaces (e.g. [3, 72]). A review of this work can be found in Chapter 3. While the use of vibrotactile display in sensory substitution is quite established, the field of vibrotactile display for human computer interaction is still new, and most of the vibrotactile messages used in the other applications listed above have been quite simple, generally only encoding one dimension of information. Using more complex messages, to encode more dimensions of information, would enable more information to be presented to users, and might be beneficial. For example, instead of simply being alerted to an incoming call by a vibration, users would be able to identify who was calling them and how important the call was from the vibration signal alone. This would allow users to moderate their calls discreetly while in meetings, or other quiet environments.

One new area of research in tactile display is the design of vibrotactile icons [22, 72, 95, 111]: the equivalent of icons in the visual domain and Earcons or Auditory Icons in the audio domain. At the start of this thesis research, there had been no formal research into the best methods of designing vibrotactile icons. While several studies have now been conducted investigating the design and/or identification of vibrotactile icons (e.g. [22, 72]), this area of research is still in its infancy, and there are many different methods of designing vibrotactile icons which could still be explored.

This thesis investigates the design of a specific type of vibrotactile icons, called Tactons. Tactons are defined by Brewster and Brown as “structured, abstract tactile messages that can be used to communicate messages non-visually” [8]. This thesis is the first formal investigation into Tactons, and the first exploration of the use of a structured, abstract approach to design vibrotactile icons.

1.1 Thesis Aims

This thesis aims to investigate how to design Tactons to present multidimensional information, and to understand the identification rates which can be achieved when these Tactons are presented to users. Sherrick and Craig [100] summarise the main problems which need to be addressed when using the sense of touch for communicating information: firstly, the selection and specification of tactile parameters, and values of these which can be readily learned and, secondly, the combination of these values to produce unique clusters. This thesis addresses both of these problems, focusing on the following two research questions.

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

Research Question 1 (RQ1) is addressed through a review of the literature on perception of vibrations, and by experiments investigating perception of individual parameters. Research Question 2 is addressed through a series of empirical experiments investigating absolute identification of Tactons which encode two or three dimensions of information in these vibrotactile parameters. The performance level is measured by recording the percentage correct scores achieved for Tactons and for their individual parameters. An additional measure of performance, which is used in Chapter 8, is the amount of information transmitted through Tactons, which is calculated using Information Transmission [97] [81]. The experiments evaluating Tactons also address Research Question 1 as they provide identification rates for the individual parameters as well as for the complete Tactons.

As there had previously been no formal work on the design of vibrotactile icons using a structured, abstract approach, the results of this work provide a baseline of performance against which the results of future experiments on Tactons can be measured. The other main outcome of this work is a set of design recommendations to help designers and researchers to create effective Tactons for use in computer interfaces.

1.1.1 Thesis statement

This thesis asserts that multi-dimensional information can be encoded in structured vibrotactile messages called Tactons by manipulating parameters of vibration, and that users can learn the mapping between Tactons and their meanings and retrieve information from them.

1.2 Overview of thesis structure

Chapter 2, *Vibrotactile Perception and Technologies*, presents an overview of vibrotactile perception and the technologies available for the presentation of vibrotactile stimuli. Results from psychophysical studies on vibrotactile perception which have been used to inform this research are discussed, and a review of the available vibrotactile technologies is presented to justify the choice of devices used in this research.

Chapter 3, *Related Work on Vibrotactile Display*, reviews related work on the presentation of information through vibrations. This chapter places the work of this thesis in context by summarising related work and identifying an area which has received little attention. In addition, the findings from this related work are considered in terms of how they could be used to inform the design of Tactons.

Chapter 4, *Designing Tactons*, discusses the approach used to design Tactons. This approach is based on the design principles of Earcons. Therefore, this chapter provides an overview of Earcon design and discusses how these principles can be applied to the design of Tactons. This chapter also discusses the different types of Tactons which could be created and the parameters that could be used to encode information in them.

Chapter 5, *Experimental Work on Individual Parameters*, reports two experiments investigating the perception of previously unexplored vibrotactile parameters, to identify whether these could be used in Tacton design. In addition, two different vibrotactile actuators are tested and a decision is made about which to use for the rest of this research.

Chapter 6, *Design and Evaluation of Two-Dimensional Tactons*, presents the first experiments on the absolute identification of Tactons. Two experiments are reported which combine pairs of parameters to create two-dimensional Tactons using a high specification vibrotactile device. A set of recommendations for the design of Two Dimensional Tactons is reported.

Chapter 7, *Design and Evaluation of Two-Dimensional Tactons for a Mobile Phone Vibration Motor*, transfers the Tactons developed in Chapter 6 to a standard mobile phone vibration motor. An experiment is reported which evaluates the effectiveness of these Tactons on this device, and recommendations for the design of Two Dimensional Tactons for phone motors are proposed, based on the findings of this work.

Chapter 8, *Design and Evaluation of Three-Dimensional Tactons*, extends the work reported in Chapter 6 by encoding a third dimension of information in Tactons. This chapter also considers how the design of three-dimensional Tactons could be improved and two further iterative experiments aiming to optimise the design are reported. Finally a set of recommendations for the design of Three Dimensional Tactons are presented.

Chapter 9, *Discussion & Conclusions*, summarises the work presented in this thesis, and relates the findings back to the research questions outlined in Chapter 1. It also sets out the findings in terms of a set of design recommendations, which can be used by researchers or interaction designers wishing to use vibrotactile stimuli. The limitations of this research are identified, and suggestions for future work based on the issues raised by the thesis are proposed.

Chapter 2: Vibrotactile Perception and Technologies

2.1 Introduction

Before designing vibrotactile messages such as Tactons it is necessary to gain an understanding of the capabilities of humans to process vibration stimuli. This chapter begins by providing an overview of the sense of touch and then goes on to present results from the literature regarding perception of the different parameters of vibration, drawing conclusions about the implications of these results for the design of vibrotactile messages. The aspects of vibrotactile perception which can be exploited in such messages are dependent not only on the limits of human perception but also on the capabilities of the vibrotactile actuators used to produce the vibrations. Therefore, the final section of this chapter provides an overview of the most common vibrotactile devices that are currently available, and discusses the choice of actuators for this research.

2.2 The Human Sense of Touch

The sense of touch consists of two main sensory systems – the kinesthetic system, which perceives sensations originating in muscles, tendons and joints such as those caused by movement, and the cutaneous (or tactile) system, which responds to sensations on the surface of the skin. The use of the sense of touch in computing is referred to as Computer Haptics, and two main categories of interfaces exist to stimulate these two sensory systems. Tactile interfaces deform the skin through pressure or vibration, or stimulate the skin using electrical currents or heat and, therefore, stimulate the cutaneous sense. Force-feedback devices, on the other hand, exert forces to oppose movement, thus making use of the kinesthetic system.

This thesis is concerned with the design of vibrotactile messages, which are perceived by the cutaneous system (since vibrations cause deformation of the skin). The kinesthetic system is not involved in the perception of vibrotactile stimuli and, therefore, is not discussed in this chapter. For a review of kinesthetic perception see [106]. The remainder of this chapter focuses on the cutaneous system of human perception.

2.2.1 The Cutaneous System

The cutaneous system responds to stimulation on the skin surface, including skin deformation caused by pressure and vibration, as well as thermal, electrical, chemical and pain stimuli. This thesis is concerned with the perception of skin deformation, rather than with the other stimuli mentioned above, and therefore this review focuses only on these aspects of perception. Skin deformation is sensed by mechanoreceptors (specialised nerve endings) which are located in the skin [64]. The skin is made up of two main layers – the outer layer is known as the epidermis and the inner layer as the dermis. The mechanoreceptors responsible for cutaneous perception are located in, and between, these two layers [64]. In addition, Pacinian corpuscles (the mechanoreceptors responsible for sensing high frequency, low intensity vibration) can be found in the hypodermis, a third layer which is situated below the dermis but above the muscles and bones [64]. Figure 2-1 shows the structure of the skin, with the mechanoreceptors (e.g. Meissner's corpuscle, Merkel's receptors) shown in the epidermis and dermis region. The hypodermis is not labeled but is the region below the dermis, where the Pacinian corpuscles are shown.

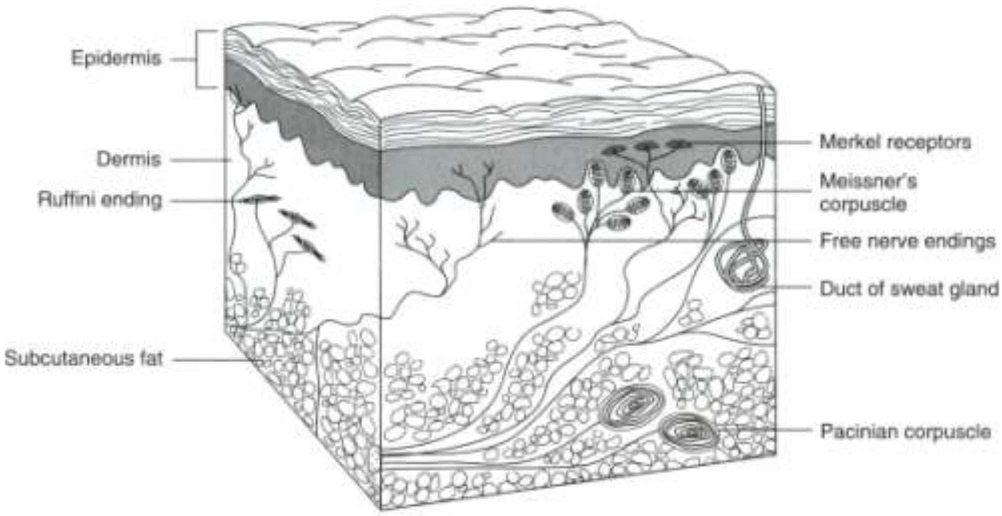

Figure 2-1: A cross section of glabrous skin, showing the layers of the skin and some of its receptors, from [51].

Two types of skin can be found on the human body: glabrous (hairless) skin and hairy skin. These two skin types are populated by different sets of sensory receptors and, therefore, respond differently to touch stimuli. Most research has focused on the glabrous (non-hairy) skin of the human hand. This skin contains four types of mechanoreceptors which respond to skin deformation, with each responding to a different type of stimulation [6, 60]. Table 2-1 summarises the main properties of each of these mechanoreceptors.

Receptor Types

The four mechanoreceptor types can be divided into two categories based on their response to sustained skin deformation [64]. Rapidly adapting (RA) receptors respond quickly to the onset (and in some cases to the offset) of skin deformation, but are not activated during sustained deformation. Slowly adapting (SA) receptors, on the other hand, respond during sustained skin deformation. The receptors can be further categorised in terms of the size of their receptive fields, with Type I receptors having small, well defined receptive fields and Type II receptors having larger, less well defined receptive fields [64].

RAI receptors respond to skin motion and their primary functions are to detect motion and control grip [60]. As rapidly adapting receptors they only detect the onset of skin deformation, and demonstrate no response to sustained skin deformation. The RAII receptors (also known as Pacinian Corpuscles) are the mechanoreceptors responsible for perception of high frequency vibrations, particularly around 200Hz (although they respond to frequencies ranging from 5-1000Hz). The Pacinian Corpuscles have a large receptive field, resulting in poor spatial acuity compared to the other receptor types, but their extreme sensitivity to high frequency vibrations means that they are able to resolve temporal patterns presented to the skin [64]. In real world scenarios, the Pacinian Corpuscles are generally responsible for the perception of vibrations felt through tools resulting from distant events (e.g. the vibrations felt on the hand when a nail is hit with the hammer held in that hand).

Mechanoreceptor type	RAI	RAII (PC)	SAI	SAII
Nerve ending	Meissner	Pacinian	Merkel	Ruffini
Effective Stimulus	Skin motion	High frequency vibration	Edges, points, corners, curvature	Skin stretch
Frequency range	1-300Hz	10-1000Hz	0-100Hz	0-?Hz
Peak Sensitivity	50Hz	200Hz	5Hz	0.5Hz
Spatial acuity	3mm	10+mm	0.5mm	7+mm

Table 2-1: Properties of the Four Mechanoreceptors responsible for cutaneous perception in the glabrous skin of the human hand, adapted from [60].

SAI receptors are responsible for the perception of form and texture as they respond to edges, corners and curvature [60]. The SAI receptors have a very small receptive field and are densely located in the skin, therefore they have extremely high spatial acuity and are able to resolve spatial details extremely well. The SAII receptors, being Type II receptors, have a large receptive field and, therefore, poor spatial acuity. The primary function of the SAII receptors is detection of skin stretch and they are, therefore, used to detect the direction of motion [60].

This dissertation is concerned with perception of vibration. All of the receptors in Table 2-1 respond to vibrations, but each responds to a different frequency range. The Pacinian corpuscles are of greatest relevance to this research since the range of vibrations generated by vibrotactile devices available for human computer interaction use frequencies in the range that is perceived best by these mechanoreceptors.

Less is understood about hairy skin than glabrous skin, but it has been suggested that tactile stimuli presented to hairy skin are processed through three mechanoreceptive channels, rather than the four types of mechanoreceptor in glabrous skin [7]. The three receptor types responsible for processing tactile stimuli in hairy skin are the RAI, RAII and SAII receptors [7]. While the glabrous and hairy skin are populated by different receptors, the important factor for this research is that the Pacinian Corpuscles (RAII receptors), the mechanoreceptors responsible for perceiving high frequency vibrations, are present in both skin types [7], although in hairy skin the Pacinian corpuscles are located in the deep tissue rather than in the subcutaneous regions [73]. Mahns *et al.* [73] report that the detection threshold for vibration on hairy skin is markedly higher than that on glabrous skin, but that vibrotactile frequency discrimination is very similar on both skin types.

2.3 Perception of Vibrotactile Parameters

To use vibrotactile display effectively in computer interfaces it is important to understand the parameters of vibration that can be manipulated to encode information, and the capabilities and limitations of the sense of touch to perceive these parameters. The main parameters of vibration are intensity, frequency, waveform duration, rhythm (temporal pattern), and spatial location [41]. This section provides an overview of each of these parameters in turn, providing information which would be relevant when using each of these parameters in tactile interface design, for example, the perceivable range of each parameter, the number of detectable steps, and the number of levels which can be absolutely identified.

2.3.1 Intensity

The term intensity is used to refer to the strength or magnitude of a vibration. Intensity is defined as the square of the amplitude. However, the terms intensity and amplitude are often used interchangeably when referring to the strength of tactile stimuli.since an increase in amplitude leads to a direct perceived increase in intensity. and is also referred to as the amplitude of vibration. Vibrotactile intensity is expressed in terms of decibels, as is volume in the audio domain. Decibel levels are referred to in terms of their relationship to the threshold of detection, e.g. 28dB SL refers to 28dB above the threshold, or sensation level (SL). This perceptual threshold is dependent on a number of variables, including the person, the tactile actuator used (in particular, the size of the contactor which vibrates against the skin and the presence or absence of a rigid surround), the frequency of the vibration and the location of stimulation.

Figure 2-2 illustrates how the detection threshold is affected by frequency and contactor size. These data could be used to provide an indication of thresholds for guidance when designing vibrotactile interfaces, but do not take account of individual differences in the participants or the technology. Therefore, if an accurate threshold measurement is required it should be determined through psychophysical procedures such as the Method of Constant Stimuli [46] for each user and each experimental setup. In addition to contactor size, the perceptual threshold can also be affected by whether the contactor is surrounded by a rigid surround. The presence of a rigid surround causes a reduction in the threshold at low frequencies and an increase in threshold at high frequencies [49].

Range

When designing a vibrotactile display it is important to ensure that the signal is strong enough to be detected, but not so strong that it causes pain or discomfort [32]. The perception threshold can be established as described above, and intensity levels ranging from this threshold up to around 55dB SL can be used; above 55dB SL stimulation can become painful [115].

The subjective magnitude, as perceived by a user, is a function of the intensity of the input signal [115]. Figure 2-3 shows the relationship between physical intensity and subjective magnitude. The subjective magnitude is also dependent on the frequency of vibration, with higher frequencies requiring lower intensity to equal the subjective magnitude of a low frequency signal [115]. This is also illustrated in

Figure 2-3, which shows the intensity levels required for different frequency levels in order to achieve the same subjective magnitude. In addition to the dependence on frequency, subjective magnitude is also affected by the size of the contactor which vibrates against the skin, with the magnitude of sensation growing as the size of the contactor is increased [115]. Signal duration also affects perceived intensity with intensities feeling stronger as the duration of the signal increases. The final factor affecting subjective magnitude is the body site which is stimulated, with magnitude growing most quickly on areas of the body with lower tactile sensitivity, for example on the hairy skin of the forearm, compared to the glabrous skin on the fingertip [115]. In addition, a fixed intensity can feel very different when applied to different body locations [41] which could cause problems if absolute identification was required.

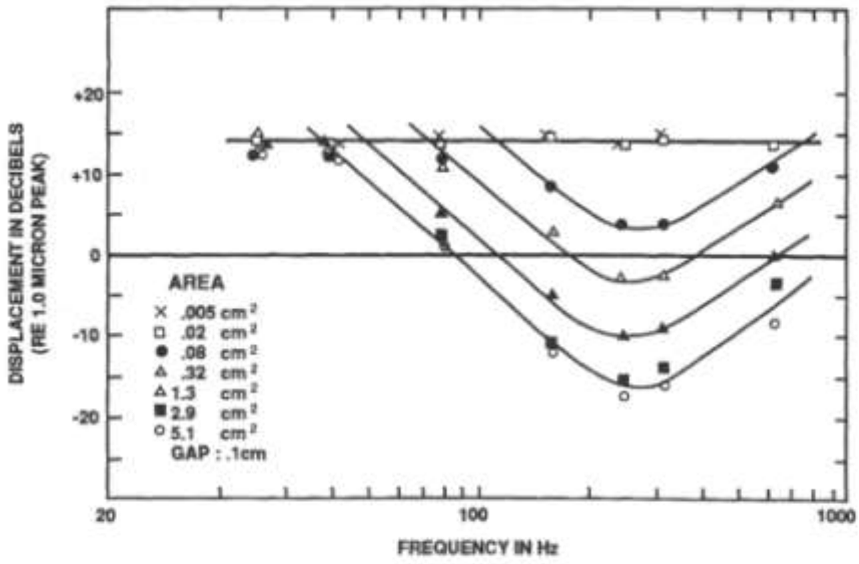

Figure 2-2: Detection thresholds on the thenar eminence (the ball of the thumb) as a function of vibration frequency and contactor size, from [115].

Difference Threshold

The term difference threshold, or difference limen (DL), refers to the amount by which a stimulus has to be altered before a change is perceived. In terms of intensity, the DL will indicate by how much the stimulus amplitude needs to be incremented or decremented for the user to identify that the intensity of the stimulus has changed. Difference thresholds are generally expressed either in terms of the relative amount by which a stimulus should be changed (e.g. by how many decibels), or by a ratio, known as the Weber ratio. The Weber ratio is calculated by dividing the intensity increment required to perceive a change by the existing intensity [100].

A range of psychophysical studies investigating the difference threshold for intensity has been carried out, and a wide range of values has been reported, with the smallest being 0.4dB (Weber ratio: 0.07) and the largest 2.3dB (Weber ratio=0.4) [115]. In the middle of these ranges, Craig [100] found a Weber ratio of 0.2 for intensity when 160Hz vibrations were presented to the index finger over a range of 10-40 dB SL. This result indicates that an increase or decrease of 20% is necessary for a change in amplitude to be perceived. Cholewiak and Collins [28] carried out a similar study using the Audiological Engineering Corp AEC V1440 device (a commercially available vibrotactile device used in tactile hearing aids) using 250Hz vibrations, and found difference thresholds ranging from 0.4 near the perception threshold to 0.2 at 20dB SL.

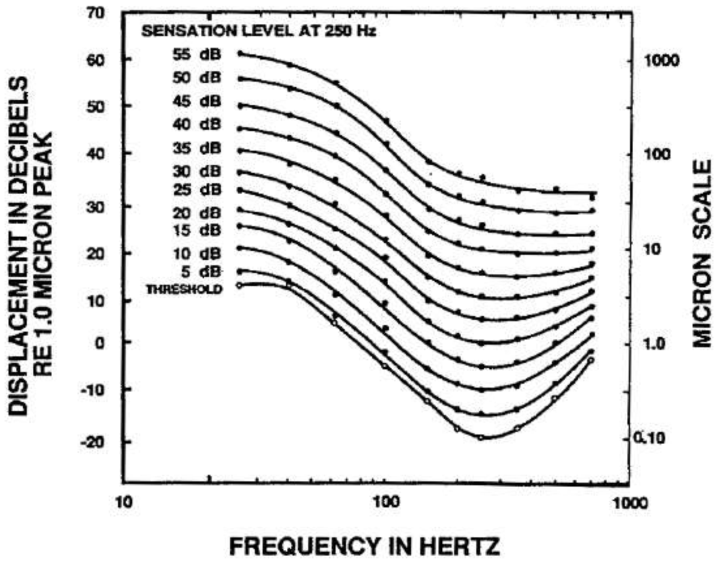

Figure 2-3: Curves showing the relationship between physical intensity (y-axis) and subjective magnitude, and equal sensation magnitude as a function of frequency (x-axis), reproduced from [115].

This variation in results is likely to be due to the use of a range of different experimental methodologies, and different vibrotactile actuators. Therefore, it could be suggested that studies should be carried out to test the DL for every different experimental set-up. However, in terms of tactile interface design, a safe approach would be to use intensity differences greater than 2.3dB (the largest DL identified in previous studies) in order to be sure that stimuli can be distinguished.

Absolute Identification

The difference threshold indicates the amount by which a signal needs to be changed in order for a change to be perceived, but does not indicate how many levels can be uniquely identified. In a tactile communication system, it is likely that the designer will wish to create different levels of a signal which can be absolutely identified by the user. Therefore, it is also necessary to obtain data on absolute identification of different intensity levels. Geldard [41] stated that around 15 levels of intensity can be discriminated, but advised that for absolute identification, especially if users were not to receive a great deal of training, a maximum of three steps widely separated across the perceivable range should be used. Cholewiak and Collins [28] found that, while difference threshold results indicate that up to 20 levels might be distinguishable, users could only absolutely identify three different levels of intensity.

When evaluating a tactile communication system which encoded information in spatial location, intensity and duration, and required absolute identification of each of these parameters, Geldard found intensity to be the least successful parameter and the main cause of errors [41]. He said that this could probably be attributed to two factors. Firstly, fixed amplitudes feel different at different locations due to the differing tactile sensitivity at different locations and, secondly, when the actuators are located on the chest it is difficult to maintain the same level of contact between the vibrotactile actuator and the skin at all times, due to movement caused by breathing. This second constraint should be minimized by using different locations but the first constraint could be problematic for any locations.

Other Factors

Since it may not be possible for people to identify a large number of intensity levels, several suggestions have been made of other ways in which to use intensity in tactile communication systems. Cholewiak and Collins [28] suggested that a wide intensity range should be used to improve the quality of tactile patterns. Another idea, proposed by Geldard [41], was that intensity could be varied over time to create different sensations, e.g. by using abrupt or gradual increases in intensity, much like in music where the attack on a note can be varied from a sharp hit to a gradual slide into a note. He reported Howell's findings that there are six distinguishable steps for the rise time of a vibration, with two levels able to be absolutely identified. It is also postulated that the same findings would apply for decreases in intensity over time.

Another possibility would be to use step changes in intensity, where there is an abrupt jump from a low to a high value (or vice versa) half way through a stimulus. Summers *et al*. [102] carried out experiments using step changes in amplitude and asked users to state whether the amplitude rose, fell or stayed constant. Their results showed that performance improved as the duration of the stimulus increased, with performance levels of around 80% correct achieved for stimuli of 800 milliseconds duration (compared to 40% for stimuli lasting 100 milliseconds).

2.3.2 Frequency

The term frequency refers to the rate of vibration. Frequency is expressed in Hertz (Hz), with 1Hz meaning that an event occurs once every second. There are two main implementations of frequency for vibrotactile stimuli. The first implementation is to use sinusoidal stimuli, in which case the frequency is manipulated by changing the number of cycles of the sinusoid which occur every second. The second implementation is to use pulsed stimuli, with more pulses occurring every second to increase the frequency of vibration.

Range

While humans can hear sounds in the range 20-20,000Hz, the frequency range of the skin is much smaller, ranging from 10Hz to 1,000Hz [100], with the practical, usable range bring further limited to just 10-400Hz [27]. The maximum sensitivity of the Pacinian corpuscles (the mechanoreceptors which respond to high frequency vibration) occurs between 200-300Hz [115]. Figure 2-2 shows the perceptual threshold for vibration on the thenar eminence as a function of frequency to show the differing sensitivity over a range of frequencies; this diagram shows maximum sensitivity at around 250Hz.

Difference Threshold

The interaction between frequency and intensity, as described in Section 2.3.1, means that it can be difficult to measure difference thresholds for frequency. In addition to different frequencies resulting in different subjective magnitudes, changing the intensity of a signal also results in a change in the perceived frequency (also referred to as vibratory pitch). In particular, increasing the intensity leads to a decrease in the perceived pitch [94]. Due to these interactions, it is necessary to match the subjective magnitude of different frequencies in order to obtain reliable and meaningful results for difference thresholds for frequency.

The first study of the difference threshold for frequency in which the subjective magnitude of difference frequencies was matched was carried out by Goff [50] in 1967. Goff [50] presented sinusoidal vibrations over a range of 25-200Hz to the fingertip, and found Weber ratios for frequency ranging from around 0.2 at 50Hz to 0.4 at 200Hz, when the magnitude was matched to a reference intensity level of 35dB SL. When the reference intensity level was reduced to 20dB, the Weber ratios ranged from 0.2 at 50Hz to 0.55 at 200Hz. These results indicate that the difference threshold for frequency is larger when the intensity is lower, and that frequency discriminations are easier at lower frequencies.

Rothenberg *et al*. [94] presented sinusoidal vibrations to the forearm and found Weber ratios for frequency ranging from 0.2 to 0.3 over a range from 10-300Hz, with the smallest ratio of 0.2 around 200Hz. This result is rather surprising when compared to Goff's results since the forearm is less sensitive than the finger and yet lower thresholds are achieved. This result may be due to a difference in experimental methodology. The results of this study do, however, backup Goff's findings that discrimination diminishes as frequency increases, but with the notable exception of improved discrimination around 200Hz (the frequency at which the skin is known to be most sensitive to vibration).

One potential solution to the interaction between frequency and intensity is to use waveforms other than sinusoids. For example, variations in pulse repetition rate, unlike changes in the frequency of a sinusoid, do not produce changes in subjective magnitude [94] and, therefore, pulse stimuli could be used. Rothenberg *et al* [94] presented pulse stimuli (using pulses of 1.1ms) to the fingertip and found that differences as small as 0.09 were perceivable at pulse repetition rates of 10-20Hz, rising rapidly from 0.15 to 0.3 as the rate increased from 100Hz to 300Hz. These results indicate that discrimination of pulse stimuli is best at frequencies below 100Hz and that, as with sinusoidal stimuli, discrimination generally diminishes as the frequency is increased. Comparing these results to the results for sinusoids in the same experiment by Rothenberg *et al*. [94] (as reported above) it can also be seen that frequency discrimination is better with pulse stimuli than with sinusoidal stimuli. Rothenberg *et al.* [94] also report that different sensations can be created by varying the pulse width within these pulse trains and that this might provide another dimension in which to encode information.

The smallest DL reported for frequency was obtained by Franzen and Nordmark [37] who used pulses of 1.5ms for input with a pulse repetition rate ranging from 1Hz to 384Hz. They found the Weber fraction to be as low as 0.03 for frequencies up to 256 pulses per second. This very small DL is surprising compared to the other results achieved, and the validity of the non-standard methodology used has been questioned by other researchers [94].

Absolute Identification

Sherrick [98] carried out rate estimation experiments which indicated that people can absolutely identify three to five different levels of frequency, but that this range can be increased to eight absolutely identifiable levels by varying intensity alongside frequency to encode the same information.

Other Factors

As with intensity, another means of using frequency is to use gradual changes over time, or abrupt step changes. In addition to exploring amplitude step changes as reported in Section 2.3.1, Summers *et al.* [102] also explored perception of step changes in frequency, and step changes when both amplitude and frequency were changed simultaneously. These results showed that identification levels of up to around 85% could be achieved for step changes in frequency (for stimuli lasting 800 milliseconds). The best performance was achieved when both amplitude and frequency were changed simultaneously, where levels of nearly 100% identification were reached for stimuli lasting 800 milliseconds.

2.3.3 Waveform

The term waveform refers to the shape of the vibration wave, for example a sine wave or a square wave. In 1960, Geldard [41] noted that perception of different vibrotactile waveforms was an unexplored area. While it is no longer completely unexplored territory, it is fair to say that waveform has received much less attention than the other parameters of vibration. Where different waveforms have been tested, the majority of this work has investigated how the waveform affects perception of other vibrotactile parameters as opposed to whether people can distinguish or identify different waveforms. For example, Rothenberg *et al.* [94] carried out studies to investigate the difference in frequency recognition when sinusoids, pulsed stimuli and warble tones were used, but did not investigate whether people could distinguish or identify these different waveforms.

Gunther [52] stated that, while people cannot differentiate between sine waves and sawtooth waves, extreme differences such as the difference between a sine wave (perceived as a smooth texture) and a square wave (perceived as a rough texture) can be perceived. However, no experimental results are given to back this up and therefore these findings cannot be validated.

While there has been little work on identification of waveform, a term which occurs often in the literature to describe vibrotactile stimuli is “roughness”. For example, Krueger [66] asked the question “is there a touch analog to the distinction between pure tones and noise?”, and mentioned the example of Katz’s work, where he found that different sensations result from rough vibrations created using a special apparatus, and pure vibrations produced with a tuning fork. Van Doren et al [110] carried out a gap detection experiment using both sinusoids and noise stimuli, and reported that, perceptually, the sinusoid felt smooth while the noise stimuli felt rough. In addition, when Weisenberger [119] presented subjects with pairs of stimuli, one an un-modulated sinusoid and the other an amplitude modulated version of the same sinusoid, subjects reported the task as being a roughness detection task, where modulated sinusoids felt “rougher” than the un-modulated sinusoid. Weisenberger’s results suggest that the sensation of roughness can be simulated by using amplitude modulation. This relationship between amplitude modulation and perceived roughness is supported by research in audio, where the relationship between amplitude modulation and roughness is well established [35, 109].

2.3.4 Duration

The term duration refers to the length of a vibration and can be defined as the time from onset to the termination of the stimulus. Vibrotactile durations are expressed, as in other domains, in terms of seconds (s) and milliseconds (ms).

Range

When using duration in tactile interface design it is important to ensure that stimuli are long enough to be detected by the user, but not so long as to make information transfer too slow. Geldard specified the usable range of durations to be from 0.1 seconds to two seconds. At the outer limits of this range, durations of less than 0.1 seconds can feel like nudges or pokes (which could be undesirable), while durations of over two seconds lead to very slow communication of information [41].

Another important factor when considering duration is the duration of gaps between vibrations, in addition to the duration of the vibrations themselves. Verrillo and Gescheider [115] report Gescheider’s earlier results on tactile gap detection for click stimuli. His results showed that gap detection thresholds were generally around 10ms, but were dependent on the length and intensity of the vibrations on either side of the gap, improving as the both of these variables increased.

Van Doren et al [110] tested gap detection using 256Hz sine waves and 250-500Hz bandpass noise presented to the right thenar eminence (the ball of the thumb). The results for both stimuli showed that the intensity required for detection of the gap decreased as the gap duration increased. For the sinusoid stimuli a gap of 10 milliseconds could be reliably detected (75% of the time) at an intensity of around 25dB SL. Gaps of 100ms or more only required an intensity of around 8dB SL to be detected. For noise stimuli, the intensity required to detect short gaps was significantly higher, at around 40dB SL for gaps of 10ms, reducing to around 8dBSL for gaps of 100ms or more (comparable to that for sinusoids). The significant difference between sine waves and noise stimuli appears to be related to the perceptual differences in these two waveforms. Participants reported that, with the “smooth” sinusoids, they perceived the gap as a click or a pop in the middle of the stimulus, whereas, for the “rough” noise stimuli, it felt like another modulation of the amplitude of the noise.

Difference Threshold

Geldard carried out an experiment to identify duration JNDs for 60Hz vibrations presented to the ventral thorax (the chest). His results showed that the JNDs increased from 0.05 seconds to 0.15 seconds as duration increased from 0.1 to 2 seconds [40]. Within this range he found that users could make 25 distinctions.

Absolute Identification

In Geldard's experiment testing duration perception on the ventral thorax [40] he showed that users could distinguish 25 different durations, but that only four-five levels could be absolutely identified. The result was achieved in a laboratory experiment with extensive training, so he recommended that the number of levels be reduced to three, if such training was not possible.

Other Factors

Interactions between duration and perceived amplitude should be considered when using duration as it has been shown that short intense signals can be confused with longer, lower intensity signals [100].

2.3.5 Rhythms/Temporal Patterns

When designing vibrotactile messages it might be useful to combine vibration bursts and gaps of different durations to form temporal patterns. Such temporal patterns are akin to rhythms in music, and therefore the terms temporal pattern and rhythm can be used interchangeably when referring to them.

When designing temporal patterns, an important consideration will be how many vibration pulses can be perceived in a given period of time. Sherrick and Craig [100] report that the skin is superior to the eye, but inferior to the ear in terms of its ability to count a number of pulses in a given time period. Results indicated that the skin is unable to count accurately if five or more pulses are presented in 700 milliseconds.

A small number of studies have focused specifically on the perception of vibrotactile rhythms. Kosonen and Raisamo [65] investigated the perception of audio, visual and vibrotactile rhythms. Participants were presented with a rhythm in one of these modalities and asked to reproduce it by tapping it out on a mouse button. The audio stimuli were presented via a loudspeaker using a simple tone, the tactile stimuli via the vibration function of a Logitech Wingman mouse (www.logitech.com) and the visual stimuli via a flashing circle on the computer screen. Simple rhythms were created using just two lengths of notes: short (300ms) and long (600ms). The results showed that the audio modality resulted in the best performance, with just 7.8% of rhythms wrongly reproduced in this modality compared to 14% in the tactile modality and 17.5% in the visual modality. Although audio is significantly better than tactile, performance in the tactile modality is also significantly better than that in the visual modality, indicating that the vibrotactile modality may be suitable for display of rhythms.

Work on vibrotactile rhythm perception has also been carried out by Ahmaniemi [2] who presented audio and vibrotactile rhythms to music students and asked them to write down these rhythms in musical notation (a task known as transcription). The audio rhythms were presented via a MIDI bassoon sound, while the vibrotactile rhythms were presented using the vibration motor in a standard mobile phone held in the user's non-dominant hand. The results showed that performance was always better when the rhythms were presented via audio than when they were presented via the tactile modality, although for two of the four rhythms tested these differences were very small. The author suggests that the performance levels may have been lower in the tactile modality since the participants were very experienced in the transcription of audio rhythms, but were completely unfamiliar with the tactile presentation of rhythm. Therefore, he suggests that performance may improve if participants became familiar with this task. No statistical analyses of these results are reported so it is difficult to draw any strong conclusions from this work.

Both of the above studies have asked participants to reproduce rhythms, either by tapping them out or by writing them down. The results of these studies, therefore, indicate whether people are able to reproduce rhythms correctly or not, but do not indicate whether people can later remember these rhythms or absolutely identify them. For rhythm to be used as a parameter in vibrotactile displays it may not be necessary for people to be able to reproduce them, but it would be necessary to ensure that the rhythms could be absolutely identified.

Buttler and Oravainen (unpublished [19]) carried out absolute identification experiments for rhythm presented via both audio and tactile modalities. The audio stimuli were 500Hz sinusoids presented via headphones, while the tactile stimuli were 159Hz vibrations presented via a mobile phone mockup containing a standard mobile phone vibration motor, with headphones worn to block any sound leakage from the vibration motor. The experiment aimed to compare perception of audio and vibrotactile rhythms, and to investigate whether there was any difference in perception between rhythmic and non-rhythmic temporal patterns. The authors define non-rhythmic patterns as temporal patterns which contain beats or pauses of varying length such that the beats in the pattern do not fit to an underlying regular beat. In this experiment the rhythmic patterns all contained six beats of 800ms length, while the lengths varied in the non-rhythmic patterns. The results showed no significant difference in recall performance between the audio and tactile stimuli, and also showed that performance was significantly better in both modalities when rhythmic, rather than non-rhythmic, patterns were used. These results indicate that, when absolute identification is required, the use of rhythm can be as effective in the tactile domain as it is in the audio domain and that using rhythmic patterns rather than non-rhythmic patterns is more effective.

Another study on vibrotactile rhythm perception was conducted by van Erp and Spapé who investigated the creation of "tactile melodies" to represent information in computer interfaces [8]. They transferred pieces of music from the auditory domain to the tactile domain and investigated how people describe and classify certain types of melodies when they are presented via a tactile transducer. The results of this study showed that the two features most important in tactile melody classification are tempo (speed) and intrusiveness. Intrusiveness is not explicitly defined but they state that levels of intrusiveness ranged from soft and polished to loud and aggressive, indicating that users associated intrusiveness with the volume/strength of the pattern and the texture or emotion of the pattern.

2.3.6 Spatial Location (Locus)

The term spatial location, or locus, refers to the location of vibrotactile stimulation on a user's body. There are two main ways in which locus could be utilised in a tactile display; either information could be encoded in the location to which the vibration is presented, or information could be encoded in the movement of a vibration pattern across the body.

Range

The range of possible locations on the body is extremely large, and is only really limited by practicalities of where actuators can be attached and by the difference threshold (as discussed below). In addition, as discussed in Section 2.2.1, tactile sensitivity varies at different locations on the body, particularly between glabrous and hairy skin, and this must be considered carefully when selecting locations. However, the sensitivity of different body locations must be weighed up against the practicalities of using these locations. For example, the fingertips are commonly used due to the fact that they are highly sensitive to small amplitudes and have good spatial acuity [32], but they can be an impractical choice for mobile or wearable computers as users may be required to use their hands for other tasks. Moving the tactile stimuli to the forearm instead leaves the hands free, but the lower tactile sensitivity of the forearm means that greater power requirements are required [115]. Therefore, all of these factors need to be considered when choosing locations.

Difference Threshold

The spatial distance required between two points of stimulation for them to be perceived as different locations is called the two point limen. The two point limen for pressure stimuli has been found to be around 0.9mm on the fingertip [84], varying across the body as shown in Figure 2-4. The large, poorly defined, receptive fields of the Pacinian corpuscles mean that the two point limen for high frequency vibration stimuli would be expected to be much larger for vibrotactile stimuli than for pressure stimuli, where the mechanoreceptors have much smaller, well defined receptive fields [41]. No definitive results appear to exist for two point limens for high frequency vibration stimuli.

Absolute Identification (Localisation)

When designing tactile displays, the ability of people to make absolute identifications of locations is likely to be of greater importance than being able to discriminate two points. If multiple points can be absolutely identified then different information can be encoded in these different locations. Therefore, it is important to be aware of the capabilities of the skin to perform absolute localisation.

Cholewiak and Collins [29] reported early work on tactile perception for pressure stimuli which suggested that tactile localisation was most precise when the stimulus was close to an anatomical reference point, and in particular at points of mobility such as the wrist or elbow. When investigating vibrotactile localisation performance with an array of seven vibrotactile actuators on the forearm, Cholewiak and Collins [29] found that perception at the wrist and elbow was more precise than that for the rest of the array, indicating that these early findings transfer to vibrotactile stimuli. Cholewiak *et al.* [30] also conducted a study on the abdomen, where the main anatomical references are the spine and navel, and found that again localisation was most precise when the stimuli occurred at these reference points. In addition, they found that people were unlikely to mistake stimulation at another point for stimulation at one of the reference points.

Geldard [41] warns against the temptation to use a set of five vibrotactile actuators located on a body site, for example the chest, and encode information by using every on-off combination of the five actuators (31 combinations in total). This approach is unsuccessful because the skin is unable to resolve the difference between one actuator and multiple actuators vibrating once it has adapted to the pressure of the actuators themselves. This, however, only occurs if the stimulation of all actuators is simultaneous and can be overcome by introducing an offset between the stimulation of each actuator. If this offset is introduced it is important to be aware that this can result in the sensation of perceived movement. While this could be undesirable in some situations, it may also be useful to consider utilising perceived movement itself as a parameter, and the following section discusses the possibilities of this parameter.

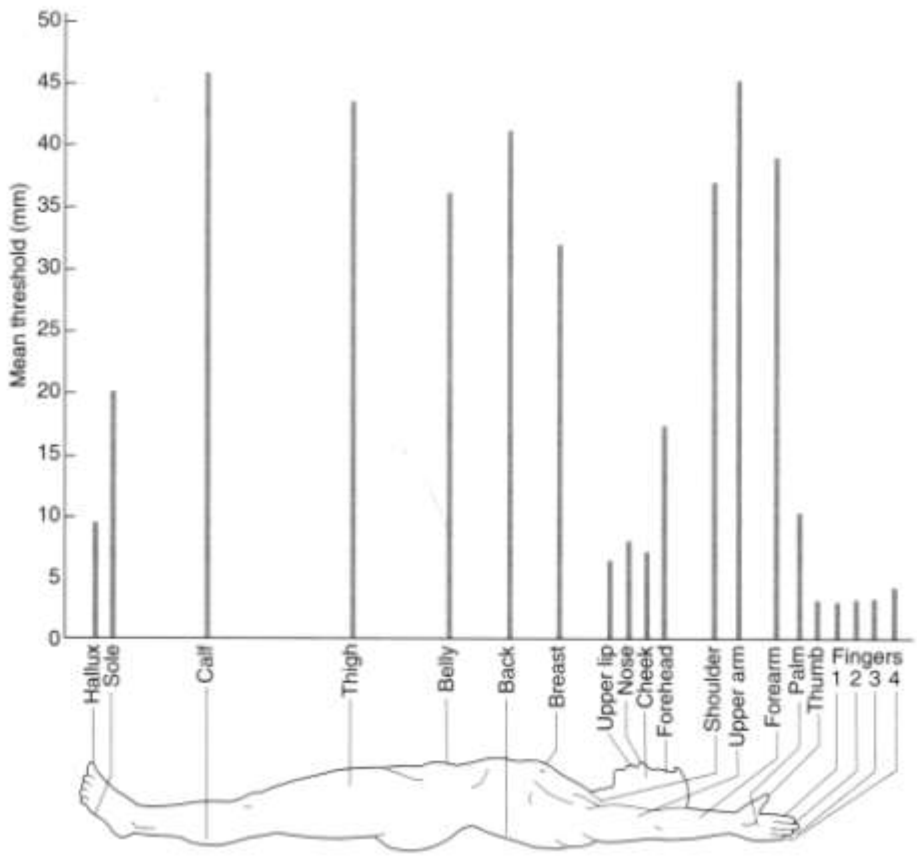

Figure 2-4: Diagram showing two point limens for pressure stimuli at different locations on the male body, from [51], after [118].

Perceived Movement

In addition to using static locations as described above, the use of multiple actuators would enable the use of spatiotemporal patterns. By activating different actuators in turn, the sensation of movement in different directions across the skin could be provided.

One method of inducing the sensation of perceived movement is to use the Sensory Saltation phenomenon, also known as the "Cutaneous Rabbit" [43]. In its simplest form it works as follows. Three vibrotactile actuators are positioned on a user's arm, at equal distances apart. The first actuator pulses three times, followed by three pulses of the second actuator and, finally, one pulse of the last actuator. Instead of feeling like three different actuators pulsing, the user feels a series of taps moving from the first transducer to the last one, as if a small rabbit was hopping along his/her arm [108]. Figure 2-5 shows the actual stimulation and the perceived sensation during sensory saltation. The open circles represent the "phantom" perceived pulses. For the saltation phenomenon to occur on the forearm, the actuators must be no further than 10cm apart [108]. These distances are different at different body sites (for details see [44]). The optimum time between pulses is around 50ms (although it can range from 20ms to 300ms) and between three and six pulses should be delivered to each transducer [108]. The intensity and duration of these pulses are not critical, and this phenomenon can also occur on body locations other than the arm, e.g. on the back or the finger.

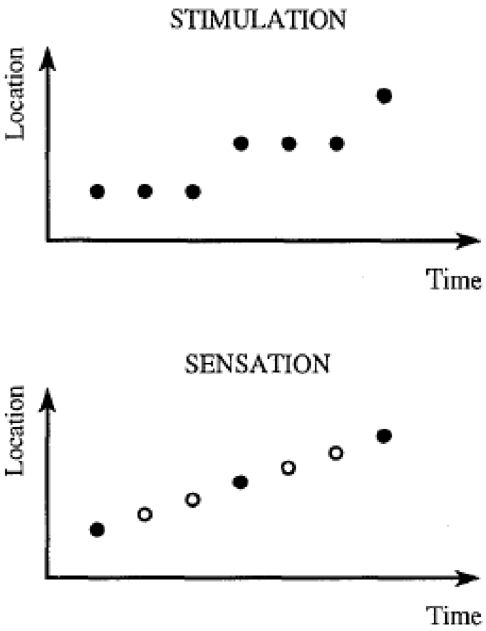

Figure 2-5: The actual stimulation pattern (top) and the perceived sensation pattern (bottom) for sensory saltation. The open circles show the perceived locations where no actual stimulation took place (from [107]).

Cholewiak and Collins [26] found that the perceived quality of lines generated by sensory saltation is similar to that generated by using a row of seven actuators and stimulating each in turn. In addition their results showed that the saltation phenomenon was perceived equally well at three different body sites, namely the volar forearm, the index finger and the back.

Tan et al [108] investigated perception of directional lines presented to a user's back via a 3x3 array of vibrotactile actuators attached to the back of an office chair. Eight possible directions were presented to the user (up, down, left, right and four diagonals) using sensory saltation, and the user was asked to respond via a written description of what they felt or a drawing of the sensation they perceived. These eight directions were implemented in two ways – one a saltatory movement (using sensory saltation) across a single row or column of actuators, and the other, where multiple actuators were stimulated at once to create the sensation of thicker lines. The results showed average percentage correct scores ranging from 79% to 91% for the simple set and from 51% to 87% for the "thick" line set. The authors report that they expect performance to improve further if an absolute identification experiment was carried out, where users chose from a predefined set, rather than having to draw or described the sensation with no prior knowledge about the stimulus set. These results indicate that participants were able to perceive the saltation phenomenon, that the eight directions are identifiable, and that performance is better with simple thin lines than with thick lines.

Cutaneous phi (or apparent motion) is a phenomenon that invokes a feeling of apparent movement [32]. One site is stimulated, followed by another and the user feels as though an object has moved between the two sites. This is the same phenomenon as the visual phi phenomenon, where a dot is lit in one location, then disappears and a dot of the same size is lit in another (nearby) location. In visual phi, the viewer observes apparent motion; they believe that the same dot has moved between the two locations [32]. To create the illusion of cutaneous phi, the duration of the signals, and the length of time between the onsets must be considered, with the time required between the stimuli for optimal perceived movement varying directly with the duration of the stimulus [99]. The longer the duration, the longer the time between the onset of each signal must be. For example, for signal durations of 50ms, the time between onsets should be 80ms, whereas if the signal duration of each stimulus is 400ms, the onsets should differ by 250ms. Apparent motion can be perceived even when the two stimuli are located on different arms; the user will imagine that the object has moved between the two arms [32].

The Phantom of Lateralisation is a phenomenon where, when two tactile stimuli are presented to different positions on the skin within 1 or 2 milliseconds of one another, the user perceives a single tactile sensation at a position between the two positions which were actually stimulated [32, 116]. This phenomenon is less likely to be useful in vibrotactile displays than the apparent movement sensations. In this case, care should be taken to avoid accidentally inducing this phenomenon if multiple body sites are to be stimulated.

2.3.7 Other Considerations when Designing Vibrotactile Stimuli

A number of other factors need to be considered when designing vibrotactile stimuli. Various phenomena, in particular adaptation and masking, can degrade or otherwise affect perception of vibrotactile sensations. Combining several parameters of vibration in a single stimulus can cause interactions between them, and this needs to be considered. Finally, tactile perception can be affected by individual differences, in particular the age of the user. All of these factors are discussed below.

Adaptation

The term adaptation refers to the fact that after extended exposure to a stimulus, the mechanoreceptors' sensitivity to the stimulus decreases. An everyday example of adaptation is that we are not continuously aware of the feel of our clothes against our skin, or the chair we are sitting on even though these things are always there. In terms of vibrotactile stimulation, adaptation causes a decrease in the perceived magnitude of a vibration and, following this stimulation, it has been shown that the absolute threshold for detection increases, making the subject less able to detect low intensity vibrations [48]. Hahn's results [54] show that the effect is greater on detection threshold than on subjective magnitude (e.g. after five minutes of stimulation, the detection threshold increased by 10dB, while the subjective magnitude decreased by just 3.5dB). After stimulation ends, recovery from adaptation occurs, with the length of time for the recovery dependent on both the duration and the intensity of the vibration.

While adaptation could be considered to be a fatigue response to prolonged stimulation, Goble and Hollins [48] have shown that the effects of adaptation are not all negative. Their results show that, while the detection threshold increases following adaptation, there is an improvement in intensity discriminations when the amplitude of the test stimuli are similar to that of the adapting stimulus [47]. In addition, they found an improvement in frequency discrimination following adaptation [48]. These results suggest that adaptation may, in fact, be the cutaneous system's way of tuning itself to optimally receive the types of stimuli to which it is currently being exposed, rather than simply a fatigue response. In terms of designing tactile interfaces it is necessary to be aware of adaptation, particularly when encoding information in the intensity of a stimulus, due to the effect it has on perception threshold and subjective magnitude.

Masking

Masking is a problem which can occur when two or more tactile stimuli are presented close to one another in time or space. This can result in a failure by the user to detect one or more of the stimuli, or in a degradation of quality in the perception of each individual stimulus and care should therefore be taken to avoid this when designing tactile interfaces. The effect of temporal masking is greater when the masker stimulus occurs after the target stimulus (backward masking), than when it occurs before the target (forward masking) [33]. It is possible to reduce the effect of masking by separating the stimuli temporally or physically. The temporal separation can be increased either by slowing down the rate of presentation or by increasing the gap between the offset of one stimulus and the onset of the next. Increasing the physical distance between stimuli, or presenting the stimuli to distinct body sites can also reduce masking, but brings in the added difficulty of requiring users to divide their attention between multiple sites (although people's ability to do this improves with practice) [33].

Combining Parameters

When designing vibrotactile displays it may be necessary to manipulate several vibrotactile parameters at once. Parameters can either be combined redundantly, such that two or more parameters are manipulated together to encode the same information, or orthogonally, where each parameter is manipulated independently to encode a different dimension of information. Rabinowitz *et al.* [89] note that while both of these methods can increase the amount of information transmitted, greater gains are achieved through orthogonal coding. However, significant benefits have also been shown by using redundancy; for example, as discussed in Section 2.3.2 (absolute identification of frequency), Sherrick found that the number of identifiable levels of frequency could be increased from three to eight by varying intensity redundantly alongside frequency. In addition, Summers *et al.* [102] (as reported in Section 2.3.2) found that manipulating frequency and amplitude redundantly resulted in better performance than manipulating either parameter alone.

When manipulating multiple parameters either redundantly or orthogonally it is necessary to consider the interactions between them. In the preceding sections, several of these have been identified. These include the effect of frequency and body site on subjective magnitude, the interaction between intensity and perceived vibratory pitch, and the effect of duration on perceived intensity. Rabinowitz *et al.*'s results [89] for orthogonal coding in one-, two- and three-dimensional vibrotactile messages show that the information transmitted for multidimensional tactile messages is lower than the sum of the information transmitted by each of the individual parameters. This is probably due to interactions between the different parameters, or to the limits of short term memory to process multi-dimensional information. These factors will need to be considered when designing Tactons that encode multidimensional information.

Age of Participants

One important consideration when using vibration stimuli which has not been covered thus far is the age of the people who will be using the system. The cutaneous sense, like the other senses, shows degradation with age and a range of studies has been carried out investigating the effect of age on aspects of vibrotactile perception. These studies have shown that older people are less sensitive to high frequency vibrations and exhibit higher gap detection thresholds, and that subjective magnitude grows more steeply in older people [115]. Controlling for the age of participants was outside the scope of the current research, but these factors should be considered if a system was intended for use by older people.

2.4 Vibrotactile Technologies

A range of vibrotactile actuators is available, with the most common types being inertial actuators (also known as inertial shakers), linear actuators and piezoceramic benders [27, 82].

Inertial actuators are the most common type of vibrotactile actuator. The definition of an inertial actuator is one in which the vibration is caused by the movement of an internal mass, and the entire case in which the mass is housed vibrates [82]. The most common inertial actuator is the standard mobile phone vibration motor, or pager motor. This is a small DC motor featuring an eccentric weight in the shaft. These vibration motors typically vibrate at frequencies around 130Hz. Another typical inertial shaker design is shown in Figure 2-6. In this design a mass is suspended on a spring within a rigid case. Both the mass and the case vibrate when an alternating electro-magnetic force is generated, so the user feels the vibrations through the case itself [27].

The Audiological Engineering Corporation (AEC) VBW32 actuator (Figure 2-7) is a commercially available electromagnetic inertial actuator, used in TACTAID tactile hearing aids (www.tactaid.com), which uses a construction similar to that shown in Figure 2-6. The VBW32 actuators are small and lightweight, weighing 6g and measuring 2.5cm long x 1.9cm wide x 1.1cm thick. They are able to produce amplitudes up to 50dB SL, and have a peak frequency response at 250Hz (www.tactaid.com). The manufacturers state that these actuators are also able to respond over a frequency range of 100-800Hz, but the output is significantly reduced at levels other than 250Hz (which is the resonant frequency of the device).

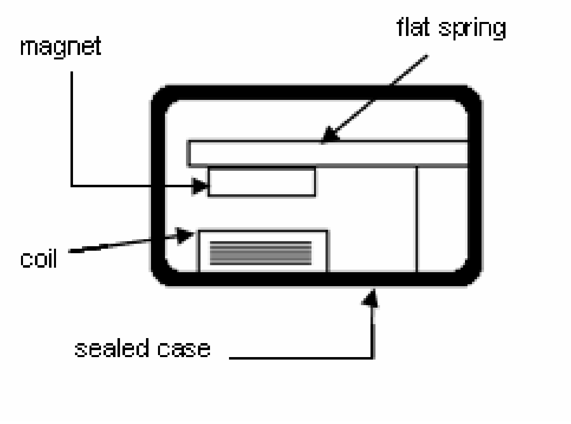

Figure 2-6: Construction of an inertial actuator (adapted from [27])

Figure 2-7: AEC TACTAID VBW32 actuator with 3.5mm jack.

The frequency response of the VBW32 at two different input levels is shown in Figure 2-8, alongside that for the C2 Tactor, which is discussed below. This diagram shows that, as expected from the specifications, the TACTAID actuator has a peak response at 250Hz and that the response drops off very sharply on both sides of this value. Therefore, to make use of these other frequencies, it would be necessary to boost the amplitude significantly. Another important feature of vibrotactile actuators is their "rise time", which is the amount of time it takes for a signal to reach its full amplitude level. Rise time is inversely related to bandwidth. This is shown by the following equation, which illustrates that the wider the bandwidth of a device, the shorter the rise time will be.

Risetime = 0.35/Bandwidth [1]

---

[1] From http://zone.ni.com/devzone/cda/tut/p/id/2709

The TACTAID actuator has a rise time of 12.5 milliseconds for 250Hz sinusoids (Tan, H., personal communication, 25 August, 2006). TACTAID VBW32 actuators cost approximately $80 US each.

Linear actuators are actuators which move a contactor (contact point with the skin) in a linear motion against the skin. In this type of device the contactor is located outside the case, and therefore the user only feels the vibration through the contactor. This can be contrasted with inertial shakers, where the whole casing vibrates. A typical linear actuator is the Bruel and Kjaer (B&K) 4810 minishaker (http://www.bksv.com/) which has been used in many tactile perception studies, including many of those reported in Section 2.3. The construction of a linear actuator such as the B&K minishaker is shown in Figure 2-9. In such a device, the coil is located in a magnetic field, and when a current passes through the coil, the coil is pushed along its axis, causing the contactor to vibrate in a linear motion against the skin [27].

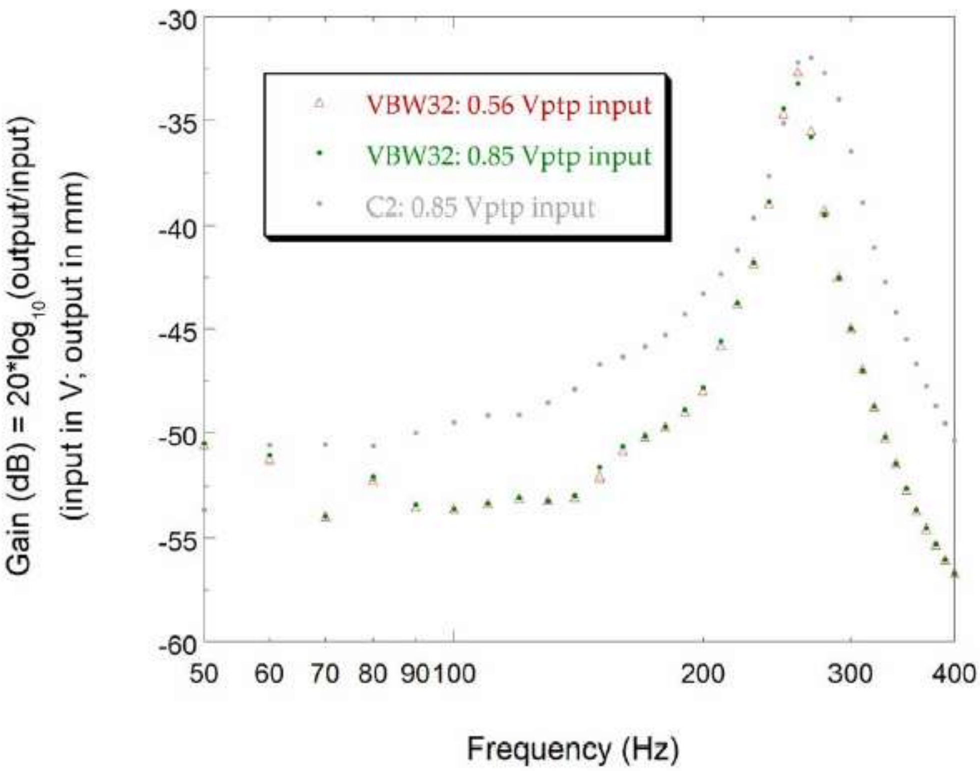

Figure 2-8: Frequency Response of TACTAID VBW32 actuator (at two input levels) and C2 Tactor (Tan, H. personal communication, 25 August 2006).

Linear actuators are able to efficiently produce a wide range of frequencies, from 0.1Hz to 300Hz, and are, therefore, suitable for studies of tactile perception [27]. In fact, most linear actuators are laboratory based devices which have been used for this purpose. However, linear actuators are not easily adapted for use in wearable displays due to their large size and weight (the B&K 4810 minishaker is about 10cm high, 10cm in diameter and weighs 1.1kg [27]), high power consumption, and the fact that, unlike inertial actuators, they require an external support to work against to produce forces against the skin.

The C2 Tactor from Engineering Acoustics Incorporated (www.eaiinfo.com) is a wearable linear vibrotactile actuator, which was designed specifically to provide a lightweight equivalent to the large laboratory-based linear actuators [82]. The contactor in this device is the moving mass itself, which is mounted above the housing and pre-loaded against the skin. This mass is driven perpendicular to the skin, as shown in Figure 2-10. The C2 measures 3cm in diameter by 0.8cm deep, with the contactor being just 0.7cm in diameter. It is heavier than the VBW32 transducers, weighing 17g, but is still light enough for easy use in wearable systems. Like the TACTAID device, the C2 is resonant at 250Hz and therefore exhibits peak frequency at this level, but is designed to be able to produce a wider range of frequencies than other commercially available transducers [82]. Figure 2-8 shows that the C2 exhibits slightly better response at low frequencies than the TACTAID device, and that the response drops off less sharply at either side of the peak frequency. The wider bandwidth of the C2 device also means that the C2 has a shorter rise time than the TACTAID, taking just 5 milliseconds for 250Hz sine waves to reach their full amplitude (Mortimer, B., personal communication, 7 August 2006). A single C2 Tactor costs over $200 US.

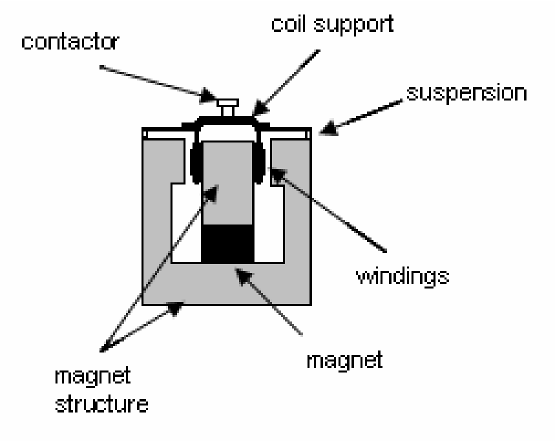

Figure 2-9: Typical construction of a linear voice-coil transducer (adapted from [22]).

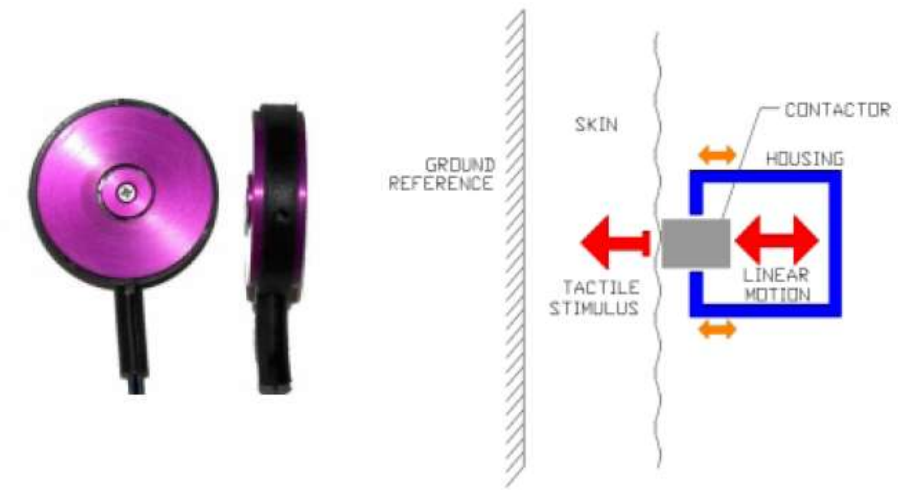

Figure 2-10: The C2 Tactor (left) and the mode of operation of this Tactor (right), from [82]

A third type of actuator is the piezoceramic actuator. Piezoceramic actuators consist of a rectangular plate made up of two layers of piezoceramic material which shrinks or expands dependent on the polarity of the voltage applied to the material. When a voltage is applied to the plate, one layer expands and the other contracts, and so the material bends [27, 87], hence the fact that these actuators are often known as piezoceramic benders. One end of the plate is clamped, and, therefore, only the unclamped end moves, acting as the contactor with the skin. Piezoceramic actuators have been used effectively in devices such as the Optacon [27] (see Chapter 3), Sony TouchEngine [87] and Summers *et al.*'s tactile array [104]. However, to provide an amplitude strong enough to enable these actuators to be used on sites other than the fingertips, voltages of over 100V would be required [27], making these an unsuitable choice for wearable displays.

2.5 Discussion

This chapter has presented a review of vibrotactile perception, in particular the perception of the parameters of vibration. In addition it has reviewed the available technologies which can be used for vibrotactile display. This section discusses how these findings can be applied to the area of vibrotactile display.

2.5.1 Choice of vibrotactile parameter for use in vibrotactile displays

Research Question 1 as set out in Chapter 1 asked:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?

This question will be answered in detail in Chapter 4, where the suitability of the different vibrotactile parameters for use in Tactons is discussed, and in the later chapters where the results of evaluations of these parameters are presented. However, from the information on vibrotactile perception presented above, along with the information on the capabilities of the available devices, it is possible at this stage to draw some preliminary conclusions about the parameters which may be suitable for use in vibrotactile displays. The following sections go through each parameter in turn discussing its feasibility for use in such displays.

Intensity

Intensity appears to be a usable parameter for vibrotactile displays, with around three levels able to be absolutely identified. However, there are a few limitations which might make the use of intensity problematic. Firstly, the subjective magnitude of a signal is dependent on many other variables, such as spatial location, how tightly the actuator is attached to the skin, and frequency of stimulation. The dependence of perceived intensity on these other variables means that its use in a real system could be complicated, as thresholds would need to be measured for each user at each point of stimulation. Since the use of fixed intensities might be problematic, another possible option is to use intensity changes over time, as people may be able to distinguish between increasing and decreasing intensity and between different gradients of each of these.

Frequency

Frequency appears to be problematic for use in vibrotactile displays for a number of reasons. Firstly, the interaction between frequency and amplitude means that controlling the perceived vibratory pitch can be extremely difficult. The results in the literature indicate that frequency should definitely not be manipulated in a tactile communication system if intensity is to be manipulated independently. However, it might be possible to combine intensity and frequency redundantly to improve identification.

The second restriction on the use of frequency in vibrotactile displays is the limitations imposed by the bandwidth of the vibrotactile actuators that are currently available. Many of these devices, as described in Section 2.4, are designed with a resonant frequency of 250Hz and are not engineered to reliably produce a wide range of frequencies. Although devices such as the EAI C2 Tactor claim to be able to produce a wider range of frequencies, the reduced outputs at frequencies other than 250Hz make the use of frequency difficult. For this reason, it is unlikely that frequency could be used reliably with these devices. While devices with a greater bandwidth may be available in future, Sherrick [98] suggests that it might be better to use another parameter such as spatial location in place of frequency in tactile communication systems, since it is easier to reproduce the existing limited-bandwidth transducers than to redesign them as high-fidelity transducers.

One possibility which could enable the use of frequency as a parameter is the implementation of frequency using pulse stimuli rather than sinusoids. This solution appears to overcome the interaction between frequency and intensity, and may also prove to be possible with the low-bandwidth devices.

Waveform

There has been very little work on identification of different waveforms in the literature on tactile perception. One solution which offers promise for the creation of distinguishable tactile waveforms is the use of "rough" and "smooth" waveforms, as several papers have mentioned that distinctions can be made between these two types of sensation. The results in the literature suggest that sinusoids feel "smooth" and that rough waveforms can be created using noise stimuli or by amplitude modulation. In addition, research has looked at peoples' ability to make frequency distinctions when either sinusoidal or pulse stimuli are used, but has not investigated whether people can distinguish between or identify these different types of waveform. The ability of people to make these distinctions may also warrant further investigation.

Duration

Duration seems to be a reliable parameter, with the results from the literature suggesting that at least three different durations can be uniquely identified. The choice of durations is important when designing vibrotactile messages, as vibrations which are too short will be hard to perceive, while using long durations will slow down the rate of information transmission.

Rhythm/Temporal Patterns

In addition to using duration as a parameter, information can be encoded in temporal patterns or rhythms, which can be created by grouping together vibrations and gaps of different durations. Several guidelines on the use of rhythm can be extracted from studies of vibrotactile rhythm perception, which have shown that people are able to identify and reproduce vibrotactile rhythms: musical principles should be applied to the design of temporal patterns, and tempo (speed) can be used as a distinguishing factor.

Spatial Location

The use of spatial location seems very promising. A large number of static locations are able to be identified. The findings suggest that, for accurate identification of three locations, two tactile actuators should be located at anatomical reference points, with the remaining actuator located at a point between these two. Localisation should be accurate at the reference points, and stimulation at the non-reference point is unlikely to be confused with stimulation at reference points.

In addition to the use of static locations, the possibility also exists to make use of perceived movements across the skin. It has been shown that, when using the sensory saltation technique, the perception of movement over many locations can be created using only a few actuators, indistinguishable from using multiple actuator sites. The use of fewer actuators means that the cost of such a system will be lower. The findings have shown that, using sensory saltation with a 3x3 actuator display on the back, at least eight directional signals are identifiable. This could be very effective in a tactile interface where different information could be represented by each of these eight signals.

2.5.2 Choice of Vibrotactile Technologies for use in this Research

The main consideration when selecting vibrotactile actuators for this research is to find devices which are small and light so they can easily be attached to the user’s body and used in wearable displays. In addition, the devices should be commercially available, so that no custom hardware needs to be developed, and should be easily controllable through computer software. These criteria are met by both the TACTAID VBW32 actuators and the EAI C2 Tactors. Both devices are small and lightweight and offer the possibility of using multiple actuators distributed across the body. In addition, they are commercially available, and can be easily controlled using audio outputs if mini jack connectors are attached to the devices.

2.6 Conclusions

This chapter has presented an overview of the aspects of human touch perception relevant to the design of vibrotactile messages such as Tactons, and a review of the available vibrotactile actuators which could be used to present Tactons.

The review of tactile perception allows some conclusions to be drawn about the parameters of vibration which should be used in vibrotactile displays. These findings suggest that the most promising parameters for encoding information in vibrotactile displays are spatial location, duration and rhythm. Spatial location offers a wide range of identifiable values. Duration offers a smaller range of values (around 3 levels), but it may also be possible to increase this range by combining vibrations and gaps of different durations to create temporal patterns/rhythms. Intensity appears to be usable, but needs to be considered carefully as the subjective magnitude perceived by the user is also dependent on a range of other factors. Frequency is likely to be a poor choice due to interactions between frequency and subjective magnitude, and the limited resolution of the available devices, although the implementation using pulsed stimuli instead of sinusoids could be considered. Another conclusion from the review is that further investigation should be carried out into the possibilities of using waveform and intensity changes over time, and Chapter 5 presents two studies investigating perception of these parameters.

The other function of this chapter has been to discuss the possible technologies which could be used to present Tactons, and two actuators have been selected to be used in this work: the C2 Tactor from Engineering Acoustics Inc. and the TACTAID VBW32 actuator.

Chapter 3: Related Work on Vibrotactile Display

3.1 Introduction

The aim of this research is to design vibrotactile messages called Tactons. The purpose of this chapter is to provide an overview of research in the field of vibrotactile display, to place the work of this thesis in context. The term "vibrotactile display" refers to the use of vibrations to present information to a user. This could be as simple as a single vibration from a mobile phone to indicate that a call has been received, or as complex as an entire vibrotactile language, with each character encoded in a different combination of vibrotactile parameters [40]. This is not intended to be an exhaustive review of all research in the field, but rather to provide an overview of the different applications in which vibrotactile feedback has been used and to identify areas which require further investigation. In addition, this review can be used to partially address the research questions posed in the Introduction:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

Research Question 1 is addressed in this review by considering the vibrotactile parameters used in this research, and whether they have been successful. This review does not directly address Research Question 2, which is only addressed in the practical work later in the thesis. However, it does provide a review of different methods of encoding information in vibrotactile messages, and this work can be used to inform the design of Tactons.

The chapter begins by discussing the advantages and disadvantages of vibrotactile display in Sections 3.2.1 and 3.2.2. Section 3.2.3 then discusses the main approaches to encoding information in vibrotactile displays. The remainder of the chapter reviews related research in the field of vibrotactile display, and is structured in terms of its main applications, namely sensory substitution, alerts, navigation, communication and enhancement of user interface elements. Finally, in Section 3.9 the main findings of the chapter are summarised and conclusions are drawn about the vibrotactile parameters used in this research and about how the work in this thesis fits with this related work.

3.2 Vibrotactile Display

Traditionally, most computer systems have relied on visual display to present information to users. However, there are scenarios in which vision is unavailable or inappropriate, and other modalities could be considered. For example, when interacting with mobile devices, such as mobile phones and personal digital assistants (PDAs), the visual display is often insufficient due to the limited screen space, or inappropriate, such as when the device is in a pocket but requires the user's attention. Therefore it is important to consider using alternative modalities through which information can be presented. Another scenario in which non-visual communication is required is for people who do not have access to vision, in particular people with a visual impairment. In addition, sighted people who are engaged in a visually-intensive task could benefit from receiving information non-visually. For example, pilots or car drivers are focused on the view in front of them and on visual information from the dashboard, but might also require navigation information. Providing this through visual displays could overload this sense and therefore presenting information to other senses, such as hearing or touch, might be beneficial. This dissertation focuses on the use of vibrotactile display, to present information to the sense of touch. In the following sections the advantages and disadvantages of vibrotactile display are discussed.

3.2.1 Advantages of Vibrotactile Display

The main advantages of vibrotactile display are as follows:

It is discreet and private to the user.

Vibrotactile displays present information directly to the user's skin and, therefore, the presentation is discreet and private to the user and does not disturb other people [14, 24]. This might be particularly useful when the user wishes to receive information but is in an environment where audio alerts would be unacceptable, for example in a meeting, in a library, or at the cinema. One example of this is given by Chang *et al*. [24], who suggest that tactile messages could be used to give feedback to a politician during a live debate. The advisors could send tactile messages to the politician to indicate how the audience were responding, allowing him to modify his behaviour accordingly [24]. Another example is that a tactile message could be sent by a wife to her husband to express support during a meeting allowing him to know she is thinking of him, without the other meeting attendees being aware of this [24]. The private nature of vibrotactile display can be contrasted with auditory display. Information presented via loudspeakers can be heard by others, which can be annoying to other people and also means that private information cannot be presented. Wearing headphones can overcome this problem but can introduce other problems as discussed below.

It leaves the ears free for other tasks.

Another advantage to using vibrotactile display instead of auditory display is that it leaves the ears free for other tasks. Wearing headphones can overcome some of the problems with audio identified in the previous section but these can block out external noise. This is particularly disadvantageous for blind people who rely on environmental noise for safe navigation, but wearing headphones is also a problem in other situations which require that the ears are free for other tasks, such as in meetings or lectures, or while talking to a friend or colleague. Vibrotactile display is, thus, a more appropriate choice in these situations. In addition, vibrotactile signals can still be felt in noisy environments, when an audio message might not be heard.

Vibrations are attention grabbing.

Van Erp and van Veen [111] report that the tactile sense is always ready to receive information since it is less engaged in other tasks, and that tactile signals are attention grabbing. Geldard [41] also notes that the attention-grabbing nature of tactile sensations means that they would be ideal for presenting emergency warnings and alerts. This can be contrasted with visual alerts, which can often be missed if the user is looking elsewhere. Audio alerts are also attention grabbing but, as discussed above, could be missed in noisy environments.

The skin offers a large, three-dimensional display area.

The large display space on the skin, combined with the three-dimensional nature of the torso means that the body can be used to present directions in three-dimensional space. Van Veen and van Erp [114] suggest that this may make vibrotactile display more suitable than visual display for presenting three-dimensional spatial information, as this kind of information can be hard to interpret when presented visually. In addition it is known that people have some difficulty in accurately locating three-dimensional directional information presented via audio displays: in particular users experience confusion between sounds in front and behind them and find it difficult to localise different elevation levels [120].

3.2.2 Disadvantages of Vibrotactile Display

The main disadvantages of vibrotactile display can be summarised as follows:

The sense of touch has limited resolution compared to the sense of hearing.

As discussed in Chapter 2, the frequency range of the skin is just 10Hz to 1,000Hz [100], compared to the 20-20,000Hz range offered by the ears. This range is restricted further by the limited frequency range offered by currently available vibrotactile devices. O'Sullivan and Chang [83] state that touch is hard to “fake”, meaning that it is difficult, with existing technologies, to generate tactile sensations that feel like real world tactile sensations (e.g. the feel of different fabrics). With audio this is quite the opposite, as real world sounds can be accurately reproduced.

Vibrotactile actuators need to be in contact with the skin.

In order to receive information from a vibrotactile device, the device must be in close contact with the skin. Furthermore, Chapter 2 indicated that one of the most promising parameters for vibrotactile display is spatial location; if this parameter is to be used, multiple vibrotactile actuators must be attached to different locations on the body. This is in contrast to audio where, if a loudspeaker is used, the sounds can be heard without the user being right next to the device (although this results in a loss of privacy), and where regular headphones can be used to present spatialised sound.

Vibrotactile stimuli may be masked by other touch sensations.

While vibrotactile stimuli may not be blocked out by environmental noise there is a chance that vibrotactile stimuli could be masked by other touch sensations, such as clothes rubbing against the skin while walking.

Prolonged use may result in fatigue

Another possible disadvantage of tactile display is that prolonged use might cause fatigue or numbness [14] resulting in reduced sensitivity to the stimuli. However, fatigue may also occur with repeated exposure to audio.

3.2.3 Encoding Strategies for Vibrotactile Display

Geldard, as reported by Craig and Sherrick [32], identified two distinct approaches to presenting information to the skin, namely the Pictorial approach and the Coded approach. The Pictorial approach involves the direct transfer of information to the skin from another sense, usually vision. For example, the character 'A' would be presented to a user's back by stimulating the transducers that make up the shape of the character 'A'. In order to facilitate the possibility of a direct transfer of data from senses other than vision (e.g. audio) to the tactile domain, the term "direct" could be used in place of "pictorial". In the Coded approach, on the other hand, a more abstract tactile representation is used, such that the user needs to learn the relationship between the stimuli and the data it represents. If the coded approach were used, the character 'A' would be presented using a code that the user had learned as representing 'A', for example two short pulses presented to the right hand. Both of these approaches are used in the work reviewed in this chapter, and the encoding strategies employed are discussed throughout the chapter.

3.3 Vibrotactile Displays for Sensory Substitution

One main application of vibrotactile display is in sensory substitution for people who have sensory impairments, e.g. those with sight and/or hearing impairments. Sensory substitution refers to the communication of information through one sense, which would normally have been gained through another sense [117]. This section provides an overview of the work on the use of vibrotactile display in this field.

3.3.1 Tactile Communication of Speech for Hearing Impaired People

The sense of touch has been utilised in natural (non-technology based) methods of communication by deaf people for many years. For example, a technique called Tadoma is used by deaf-blind people to enable them to interpret speech. The deaf-blind person places his/her hands on the speaker's neck and face and interprets what is being said through movements of the lips and jaw, airflow at the mouth and vibrations on the neck [91]. The vibration aspect of Tadoma is of the most interest to this research, and the acquisition of speech information through vibrations from a speaker's neck is also used by deaf people as an aid to lip-reading in a technique called Tactiling (http://www.tactaid.com/).

Since the sense of touch is used naturally by deaf people, it is unsurprising that researchers and product developers have gone on to investigate whether these natural methods of receiving speech information could be replicated using technology. Tactile hearing aids have been developed which receive speech information through a microphone and translate this speech signal into a tactile signal which represents some or all of the speech information. While the work in this thesis is not concerned with the coding of speech information, the strategies used for coding speech information in tactile hearing aids could also be used to encode information in tactile messages.

Two main types of tactile hearing aids have been researched and developed, namely single channel aids (where there is only one tactile actuator) and multi-channel aids (where multiple actuators are used).

Single Channel Tactile Aids

In single channel aids the speech signal is recorded and then presented as a tactile signal by varying a combination of the temporal pattern of the tactile signal, and the intensity and frequency of the vibration. One approach to encoding information in tactile aids has been to transfer the speech signal directly to a tactile signal, varying all three of these parameters according to the speech signal. With this approach (which is an example of the direct approach outlined in Section 3.2.3) the user should be able to feel the frequency change with the pitch of the speaker's voice, the temporal pattern change with the rhythm of the speech, and the intensity of the vibration change according to the stress and emphasis within the speech pattern [101]. While this approach can provide users with a large quantity of information it can be quite difficult for them to extract the relevant information from this complex signal, and the limited frequency range perceivable by the skin, alongside the limited bandwidth of many tactile devices, can make the use of frequency difficult.

An alternative approach has been to present only the most important features of a speech signal to the user, by varying only the parameters that are best perceived through touch. This is called the "processing" approach [101] and is somewhere between the direct approach and the coded approach discussed above. While this approach will transmit a smaller quantity of information than the direct approach, if it is well designed (both in terms of the selection of features, and of the choice of tactile parameters) it should allow more effective information transfer as it is intended to use the sense of touch more effectively. The simplest implementations of this approach maintain a fixed frequency and amplitude and only present the temporal pattern of speech [101]; the commercially available TAM device operates in this manner. More complex devices have retained a fixed frequency but have modulated the amplitude to provide the user with information about the volume of speech, thus giving information about stress and emphasis within speech, as well as the temporal pattern or rhythm of the speech [101].

Multichannel Tactile Aids

The addition of multiple tactile actuators offers a potential solution to the limited ability of the skin to distinguish multiple levels of frequency and intensity [20]. Several common coding strategies have been employed. The most common strategy maps the frequency or pitch of the speaker's voice to spatial location. One example of a frequency-to-place system is the Teletactor system. In this system, the pitch of the speaker's voice was represented by the frequency of the vibration, but different pitch ranges were also presented redundantly to different fingers to aid pitch discrimination [74]. However, the preservation of vibrotactile frequency as well as location in this device made it difficult for users to discriminate the high frequency signals since they were not in the range to which the skin is particularly sensitive. Later frequency-to-place systems such as The Felix system used carrier frequencies of 300Hz (which is in the optimal sensitivity range for the Pacinian corpuscles) to overcome this problem [74]. Another method which has been used in tactile aids is to map both frequency and amplitude of speech patterns to spatial location. In such systems, a two-dimensional area on the skin is used and frequency is presented on one dimension, while amplitude is presented on the other [20, 74].

Tan *et al*. [105] note that the performance achieved with tactile hearing aids is very poor compared to that which can be achieved with Tadoma. They therefore investigated whether a more complex system combining vibrotactile feedback with kinaesthetic feedback might be more successful. A set of stimuli was created by combining kinaesthetic feedback in the form of finger motions with different frequency levels of vibration via a custom built device called the Tactuator. Figure 3-1 shows a schematic drawing of this device showing the placement of the fingers on the device and the available motions. By using either very low (10Hz/30Hz) frequencies or very high frequencies (150Hz/300Hz) extremely different sensations were created, and when a low frequency and a high frequency signal were combined both sensations were discernible. These combinations of frequencies and motions were then combined with four different location cues to create a set of 120 signals. Absolute identification experiments were carried out with these stimuli and the results showed that perfect identification of a set of up to 90 different stimuli was possible with this encoding scheme.

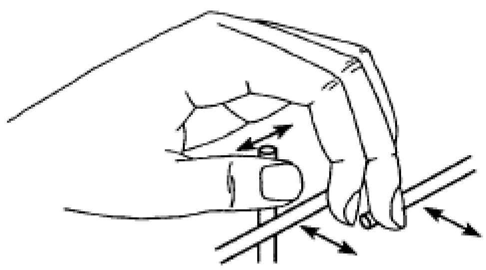

Figure 3-1: Schematic drawing showing the placement of the fingers of the Tactuator device, and the available motions, from [106].

Other approaches

Rabinowitz *et al.* [89] carried out research into the design of multidimensional tactile displays, which could be used for the communication of speech. They created one-dimensional, two-dimensional and three-dimensional tactile messages by manipulating intensity, frequency and contactor area. Their results showed that one dimensional stimuli resulted in information transmission of 1.45 bits (2.7 values perfectly identified), two-dimensional stimuli transmitted around 2.36 bits (5.1 values perfectly identified), and three-dimensional stimuli around 3.99 bits (15.9 values perfectly identified). These results show that increasing the dimensionality of a display increases the amount of information transmitted. This research is of particular interest to the research in this thesis as it is one of very few examples of the use of a multidimensional tactile display. The parameters used, however, may not be suitable. The device used in Rabinowitz *et al.*'s study offered a wide range of intensities and frequencies which cannot be reproduced on most small, commercially available devices. In addition, the devices used in this thesis research do not have adjustable contactor sizes. Although the parameters may not be appropriate, the same approach could be employed to encode multidimensional information in Tactons, by using a set of parameters more suited to the available devices.

3.3.2 Vision Substitution using Vibrotactile Displays

Sensory substitution using the sense of touch has been used by blind people for many years in the form of Braille. In addition, vibrotactile displays have been developed to aid blind people in accessing visual information such as text and pictures. There have been two main approaches to encoding information in these displays. The first two devices described below use a pictorial representation of the data: an image is captured by a camera, and is reproduced directly as a pattern by vibrating the corresponding pins or actuators in a vibrotactile array. The third system described uses a more abstract, coded approach to encode the information, where there is no direct relationship between the vibrotactile stimulus and the data that it represents and, therefore, the mapping between the two has to be learned. The final system uses an encoding scheme which is somewhere between these two approaches; some pictorial elements are retained, but where this is not possible some more abstract coding occurs.

The Optacon

The Optacon was developed in the 1960s as a commercially available reading device for blind people, and was manufactured until production ceased in 1996. A small camera was moved over any material (text or graphics) that the user wished to read. The image captured by the camera was then presented to the user’s fingertips through a 6x24 array of metal pins (Figure 3-2). This system used a direct, pictorial representation of the data, with the pins vibrating (at 230Hz) to create a tactile reproduction of the image captured by the camera [32]. The Optacon was found to be reasonably effective, with reading speeds of around 10 to 12 wpm (words per minute) after the initial 9 day training period, reaching 30 to 50 wpm after further training and experience [32]. Although these reading speeds are significantly slower than Braille, where the average reading speed for an adult is 104 wpm [36], the Optacon was beneficial for blind people as it allowed them to access any text or graphics without having to wait for it to be converted into Braille or tactile diagrams. The only parameter used to encode information in the Optacon display is the spatial pattern created by the pins. The speed of presentation can be controlled by the user by varying the speed at which they move the camera, but no other parameters are varied to encode data.

TVSS (Tactile Vision Substitution System)

The Tactile Vision Substitution System (TVSS), like the Optacon, converted images captured by a camera into vibrotactile patterns presented to the user’s skin, with the aim of substituting vision. The user sat in a chair which had a 20x20 array of vibrotactile tranducers built into the back of it, and objects captured by the camera were presented by stimulating the transducers to make a pattern which represented that object (Figure 3-3). These transducers were either on or off and could, therefore, only represent light and dark (nothing in between). Users were able to distinguish horizontal and vertical lines, and pick which object was being presented from a choice of 25, but had difficulty with the internal details of objects, such as facial features [32].

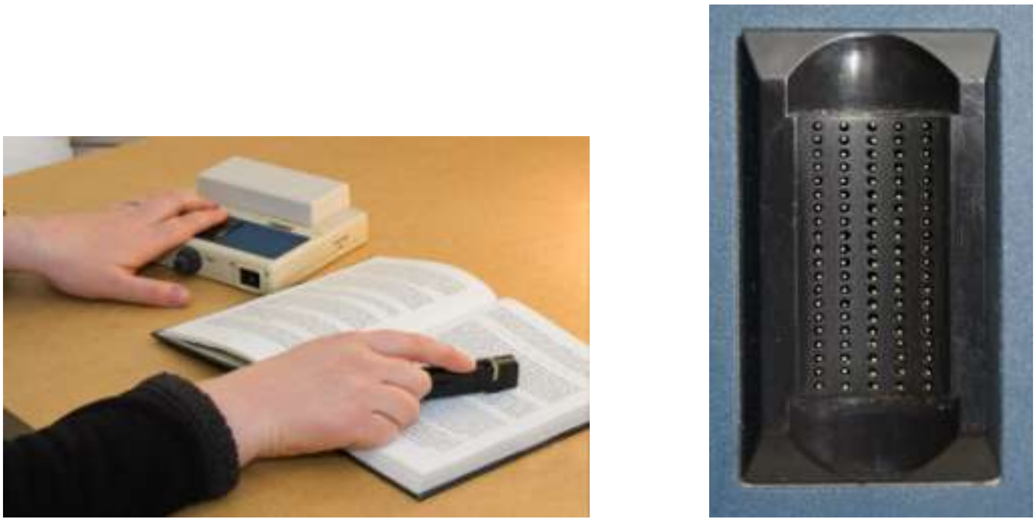

Figure 3-2: Using the Optacon device to read text. The image on the right shows a close-up of the pin array.

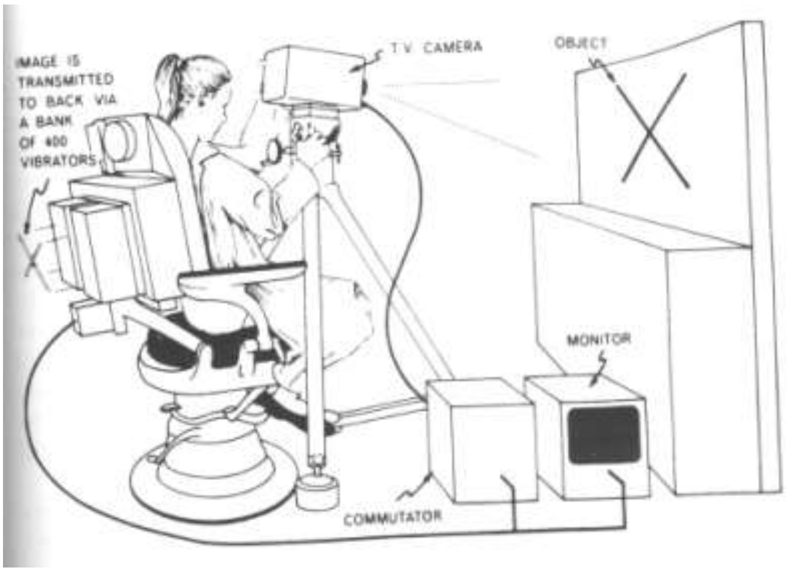

Figure 3-3: Drawing of the TVSS system, from [121] , reproduced in [32].

Vibratese

A study by Geldard [40] showed that people were able to learn a new vibrotactile language, called Vibratese. A tactile language was created, which varied three parameters (duration, intensity and location) in order to assign a unique signal to every letter of the alphabet. Three durations and three intensities were used, and five vibrotactile transducers were positioned in an X-pattern on the user's chest with only one transducer activated at a time (Figure 3-4). These parameters were manipulated orthogonally (independently), resulting in 45 unique combinations of these three parameters (3 intensities x 3 durations x 5 locations). The most commonly occurring letters were represented by the shortest durations to speed up presentation. The encoding scheme used in Vibratese is an example of the coded approach, as the mapping is abstract. There is no direct relationship between the vibrotactile stimulus and the data it represents. This language was found to be quite successful; after several training sessions, one subject managed to interpret this language at a rate of 35 words per minute. As reported in Chapter 2, the results of the evaluation showed that intensity was the least successful parameter, and the main cause of errors in this system [41]. The results achieved with the Vibratese language indicate that vibrotactile display holds promise for communicating complex information, and that using an abstract approach to encoding information in vibrotactile messages can be successful.

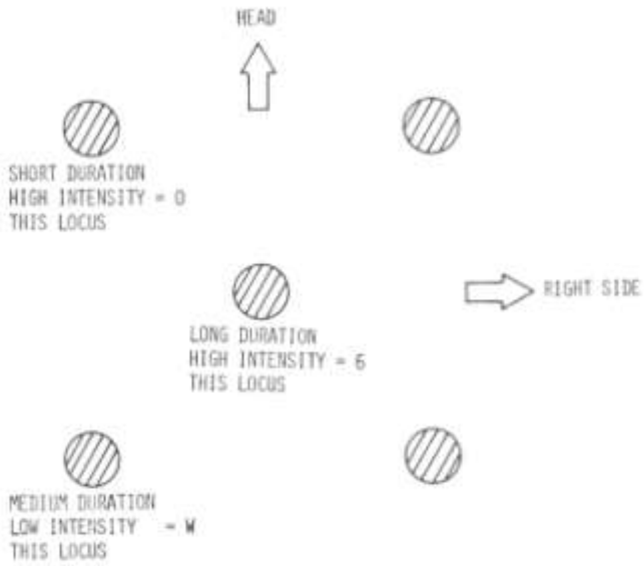

Figure 3-4: The Vibratese Code. The five actuators shown are located in this configuration on the chest. A few examples of letters are shown to illustrate the various possible combinations of locus, intensity and duration of vibration (from [100], after [40]).

Optohapt

Geldard also designed another tactile language called the Optohapt, to turn printed text into a vibrotactile alphabet for visually impaired people. The Optohapt converts printed text into spatio-temporal signals presented to the skin. Nine vibrotactile actuators are positioned on nine different body locations (Figure 3-5 (left)), with these locations carefully chosen to be widely distributed across the body, and to avoid using the same point on two corresponding limbs to avoid confusions (e.g. if an actuator is located on the wrist of the right hand, there would not be an actuator on the wrist of the left hand). In order that the symbols would be more easily distinguished, the alphabetical characters were first converted to more distinguishable symbols before scanning. A vertical array of nine sensors scans horizontally across the written text at a constant speed (Figure 3-5 (right)). When the sensor scans over a black area, the corresponding vibrotactile actuator stimulates the skin. For example, if the symbol were a vertical bar, the vibrotactile actuators would all be "off" at the start, then all come on simultaneously for a short period while the sensors were over the bar, and then all "off" again. This mapping is somewhere between a direct mapping from the data and an abstract approach, as the patterns presented to the skin are a temporal representation of the visual pattern, but the spatial pattern is not preserved. Unfortunately, no results are reported for reading speeds with this system and, therefore, it is not possible to gauge the effectiveness of this language.

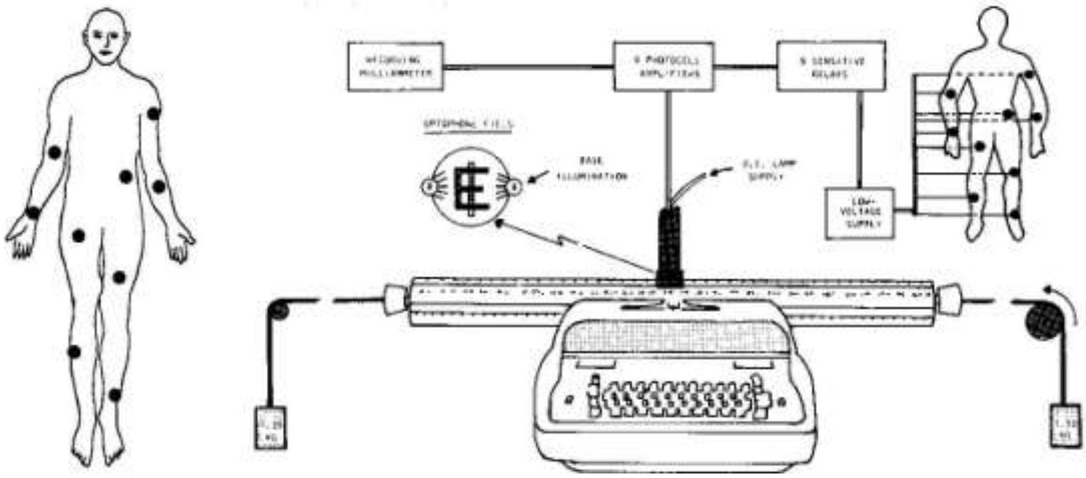

Figure 3-5: (left) The location of the nine actuators in the Optohapt system (from [42]). (right) The Optohapt System: The photocell scans across text at a constant rate. When any of the nine sensors passes over a black area the corresponding vibrotactile actuator on the user's body is activated (from [42]).

3.4 Vibrotactile Displays for Alerts

In Section 3.2.1 it was identified that two of the advantages of vibrotactile display are that it is attention grabbing and is not blocked by environmental noise. Geldard [41] noted that these factors make vibrotactile display an ideal choice for presenting emergency warnings and alerts. Vibration alerts are used by the deaf community to alert them to events such as their alarm clock ringing, or their smoke detector going off (http://www.deafequipment.co.uk/), and are also used in a range of other applications as described below.

3.4.1 Mobile Phones/Mobile Devices

One of the most common uses of vibrotactile display is for alerting users to incoming calls on a mobile phone or upcoming appointments in a personal digital assistant (PDA). These alerts generally consist of a simple vibration to grab the user's attention, but no further information is encoded in the vibration. Recent developments have lead to more complex vibration alerts in commercially available mobile devices. The simplest of these is in some Nokia phone models which enhance the rhythm of ringtones with matching vibration patterns, but it is not possible to use these vibration patterns alone without the audio, and they have not been specially designed to be distinguishable by the vibration alone. The Motorola E398 phone features a multi-function transducer which produces both audio and tactile feedback, with the tactile feedback occurring when frequencies are in the range of 100-300 Hz In this phone, the vibration is used to enhance ringtones with haptic effects [25, 83].

The most extensive addition of vibration feedback to phones has been carried out by Immersion, who developed the "VibeTonz" system (http://www.vibetonz.com), which enables the implementation of more complex vibration effects. This system allows phone manufacturers to create their own tactile effects and add these to phone alerts, user interface elements and messaging applications. In terms of alerts, they suggest that vibration effects could be added to enhance audio ringtones, or could be used independently to allow people to know who is calling them from the vibration alone. However, there has been no research into how to design the most distinguishable messages.

This work on vibrotactile alerts is quite closely related to the work in this thesis, as the aim is to communicate more information through vibrotactile messages. However, most of this work has used the vibrotactile messages simply to enhance the audio ringtones. In cases where the vibrotactile display is intended to be used alone, there appears to have been no formal studies investigating how best to design a set of distinguishable vibrotactile messages that can be understood in isolation, when presented without audio feedback. As discussed in Section 3.2.1, using vibration alone would allow information to be communicated discreetly without disturbing others, and could be used to alert users in noisy environments, therefore this could be beneficial for mobile phone alerts.

3.4.2 In-Car Warning Systems

Vibrotactile feedback has also been used for alerting purposes in cars. One commercial application of this is the Citroen Lane Departure Warning System. In this system the driver's seat is vibrated if the system senses that the car is about to cross white lines without the driver indicating beforehand (http://www.citroen.com/).

Research into the use of vibrotactile alerts for car drivers has also been carried out by Ho *et al*. [57, 58] who investigated the use of vibrotactile cues to alert drivers to potentially dangerous events on the road. Their results showed that drivers responded significantly faster to alerts provided via vibrotactile stimuli than those presented by either audio or visual stimuli [57], and that more rapid responses were facilitated by presenting the vibrotactile stimulus to the side of the body (front or back of torso) that corresponded to the location of the event [58]. These results indicate that vibrotactile messages can be attention grabbing, and that people are able to accurately localise vibrations presented to the front or back of the torso. However, the vibration stimuli are quite simple and only encode one dimension of information (front/back) and, therefore, they are less complex than the type of messages investigated in this thesis.

3.5 Vibrotactile Display for Navigation

One of the advantages of vibrotactile display that was identified in Section 3.2.1 is that the large display space on the body enables the presentation of directional information. This makes vibrotactile display a suitable choice for presenting navigational information. Vibrotactile navigation systems have been developed by a number of researchers, to aid navigation for a range of users including car drivers, pilots, and visually impaired people.

In-Car Navigation Systems

Van Erp and van Veen [111] designed "Vibrocons" (vibrotactile icons) for an in-car navigation system, using vibrotactile devices mounted in a car seat. The Vibrocons encoded the direction the driver should turn using spatial location, with a vibration on the left leg indicating a left turn, and a vibration on the right indicating that the driver should turn right. The distance to the turning was encoded in the temporal pattern, with the direction information presented more frequently as the turning point approached. In his guidelines for the use of vibrotactile display in human computer interaction, van Erp [112] stated that vibrotactile messages should be self-explaining. These Vibrocons are designed according to this principle (e.g. using a vibration on the left to indicate a left-turn), and can therefore be considered to be an example of the direct approach. This vibrotactile navigation system was compared with a visual display where the information was presented using alphanumeric characters, and in a third condition participants used both visual and tactile feedback [111]. The results showed that the tactile feedback resulted in faster reaction times, lower mental effort and lower workload than the visual display. When tactile and visual cues were both available, response times were slower than in the tactile-only condition. This seems to have occurred because, when the visual cues were available, the drivers would check them as well after feeling the tactile feedback, to confirm that they had understood the Vibrocon. This indicates that while users are able to interpret vibrotactile messages, they lack confidence in their interpretation of them. The results of this study indicate that vibrotactile feedback can be effective in improving response times when people are under workload, and that more than one dimension of information can be interpreted from vibrotactile stimuli (in this case, direction and distance were both encoded).

Gemperle *et al*. [45] designed a wearable tactile vest for tactile information display, the main application of which was the display of navigation information. This vest could present information on direction, rotation, speed and acceleration to a range of users including pedestrians and cyclists, as well as car drivers. The stimulation of a single actuator on the appropriate location could indicate direction, while a movement across several actuators could signal rotation. Changing the frequency, intensity or duration of the vibrations could indicate an increase/decrease in speed or acceleration. This system has not been evaluated beyond pilot testing, so it is not possible to assess how successful the method of information presentation is.

Navigation for Pilots

Van Veen and van Erp [114] have investigated the use of a vibrotactile vest to present information, including navigation information, to fighter jet pilots. Tactile display would be particularly useful in this setting as vision can be affected when pilots are in g-load conditions, and tactile information might be more readily received. The tactile vest consists of 128 tactile actuators distributed across the torso. In this display the spatial location of the activated tactile actuators can be used to present a variety of information. For example, activating a single point could indicate the direction of a target or destination. In addition, a line of actuators could be activated to indicate an artificial horizon. These uses are examples of the direct approach and are similar to the Vibrocons described in the previous section. The tactile display could also be used to present coded information such as fuel supply, speed, altitude or time, or to enable discreet communication between pilots or crew members. These examples would use a more coded approach. A study was conducted where users were asked to identify whether a vibration had occurred on their left or their right hand side, both in normal conditions and under g-load. The results showed that performance was consistently high (85-100% correct), and that vibrotactile perception was not significantly impaired under g-load. These results are interesting as they suggest that vibrotactile display can be appropriate in these unusual conditions, and therefore offers more possibilities for this research. However, the stimuli tested were quite simple, and no evaluation has been conducted into perception of the more complex coded information proposed above.

Navigation for Visually Impaired People

In addition to aiding navigation in vehicles, vibrotactile feedback has also been used to help blind or visually impaired people to navigate. For example, Ross and Blasch [93] designed a wearable tactile display, which indicated whether the user was walking in the right direction or if a change of direction was needed. Three vibrotactile actuators were located on the user's upper back. If the user was on target, the centre actuator was activated every two second to indicate this. The left and right actuators were activated to indicate that the user was off target in that direction. When compared with a speech interface and a non-speech audio interface, the results showed that the tactile display resulted in better performance and was preferred by users.

In addition to providing navigation information, tactile displays have also been used to help blind people with obstacle avoidance. For example the UltraCane (www.soundforesight.co.uk) uses ultrasound to detect objects in a user's environment and displays the location and distance to targets by vibrating pads on the handle of the cane. Another similar system is the People Sensor [90] which uses sensors to locate (and distinguish between) people and objects, and indicates their location to the user through vibrotactile output. Two different types of pager motors are attached to the user's waist. As the motors are different types, they produce different sensations so one is used to represent people and the other to represent inanimate objects. The distance to the person or object is indicated to the user by the rate at which vibrotactile pulses are presented, with the pulses becoming more frequent as the person or object gets closer. As no evaluation has been carried out, the success of this system cannot be gauged.

3.6 Vibrotactile Display for Communication

Vibrotactile display is a promising modality for communication as it is discreet and private to the user, and therefore does not disturb other people. In addition, the sense of touch can be very expressive and personal communication may also be enhanced by the addition of this expressive modality. A number of applications of vibrotactile feedback for communication have been investigated and are presented in the following sections.

3.6.1 Tactile Messaging in Mobile Phones

In 2005, the concept of using tactile icons for communication between two people via mobile phones was patented by Nokia Corporation (United States Patent 6963762). This patent is defined as follows (from http://www.freepatentsonline.com/6963762.html):

> *"A mobile phone or telecommunications terminal that sends and receives tactile icons (tactile sensation patterns, including vibration patterns) discernible by feel to a user of the mobile phone or terminal, indicative of a message desired to be communicated between users of such a mobile phone or terminal (not information concerning the operation of the mobile phone or terminal)."*

The idea of sending vibration messages between phones has also been presented by Immersion (http://www.vibetonz.com). One feature of the Immersion VibeTonz system is the ability to send tactile messages along with text messages. They suggest that a user might use this to send a hug sensation or a heartbeat sensation to his girlfriend to let her know he is thinking of her. They describe the hug sensation as follows (from http://www.vibetonz.com):

> *"A warm, soft, purring effect makes you feel like you're being hugged."*

As with the other elements of the VibeTonz system, no formal evaluation of these vibrotactile messages has been reported.

Another example of the use of a vibrotactile display for personal communication is the ComTouch Communication Device designed by Chang *et al*. [24]. This device fits over a mobile phone and users interact with the device by applying pressure to the five sensors positioned under their fingers. The finger pressure is translated into vibrotactile messages which are then sent to another user holding a ComTouch device. The user at the receiving end feels the vibrations presented to the corresponding fingers of their own hand. The five-finger version of this device has not been built, but a prototype version, which allows people to interact using just a single finger to send and receive messages, has been evaluated. In the prototype version of the device the sender applies pressure with their fingertip to a force sensor; the harder they press, the stronger the vibration that is sent to the receiver. In addition to varying the intensity of the vibration the sender could also vary the duration of the vibrations which was dependent on the length of time the finger was in contact with the sensor, and the rhythm of the vibrations by tapping out different temporal patterns on the sensor.

When people were asked to communicate using the device, both to enhance general chatting and to complete a negotiation task where voice chat was intentionally limited, a number of different uses of the tactile communication device were observed. In the general chatting scenario people used the tactile feedback to add emphasis to their speech, both talking and pressing the sensor at the same time to accent certain words or syllables. In addition they used tactile feedback to indicate that they wanted to speak, sending a buzz to interrupt the other speaker. These two behaviours seem to use the tactile feedback to replace non-verbal communication which would occur in face to face interaction, e.g. gesturing to emphasise parts of speech, or making eye contact or gesturing to indicate an intention to speak or a desire to interrupt. People also mimicked the tactile messages sent by the other participant, e.g. participant 1 would tap out a rhythm and participant 2 would repeat it back to them with the same rhythm, duration and intensity. This seems to be a more playful use of the device, just to test out the technology rather than as an aid to communication but indicates the potential of this technique which could be exploited to enable communication. In the negotiation task users had to rely on tactile communication alone, as the use of speech was limited. In this scenario users developed their own encoding schemes, in particular using different durations of vibrations to indicate yes/no, and using multiple taps to indicate a particular number. One problem that was observed in the evaluation of this system was that half the users were unable to perceive the full range of vibrotactile intensities and reported that they could only sense three states of intensity – high, low and off. The authors state that this may be due to the ergonomics of the device, but it may also be due to the problems with absolute identification of intensity identified in Chapter 2.

The results from the evaluation of the ComTouch system indicate that tactile feedback might be a useful enhancement to personal communication devices, enabling people to have a richer remote communication experience that can compensate for the lack of non-verbal cues. In addition, the results showed that people could communicate using vibrotactile feedback alone, and this would allow the communication of information in a discreet manner.

3.6.2 Vibrotactile Communication in Collaborative Applications

Chan *et al.* [22, 23] created a set of "haptic icons" (touch icons) for use in a collaborative application to communicate information such as whether the user is currently in control of the application, waiting for control or observing, and if anyone else has requested control. These haptic icons were presented via the Logitech ifeel mouse (http://www.logitech.com). The parameters of vibration which can be manipulated on the ifeel mouse are frequency, amplitude, and rhythm, so these parameters were used to create the set of icons. The approach to encoding information in these icons was to use common metaphors. For example, the *waiting for control* icons use the metaphor of someone impatiently waiting in line, tapping or drumming his or her fingers. The *gained control* and *lost control* icons use a vibration pattern similar to the audio message received when a USB device is plugged into or removed from a Windows PC, and the *in control* icon uses a heartbeat metaphor where, as a person becomes more anxious, their heart beats harder and faster, so a gentle vibration is felt when no one has requested control, but as requests are made, the intensity of the vibration increases. These haptic icons are grouped into "families" dependent on what they represent. For example, the *gained control* and *lost control* items are from one family, while there are two icons in the *waiting for control* family: *waiting for control after a gentle request*, or *waiting for control following an urgent request*. Chan states that icons belonging to the same family should be perceptually similar, but still distinguishable. In order to ensure that all the haptic icons were distinguishable, and that the icons belonging to the same family were perceived to be similar, a multi-dimensional scaling technique was used. In this technique users are presented with stimuli and asked to categorise and group together stimuli which felt similar. The results showed that subjects categorised first on the number of bursts (rhythm) to distinguish between stimuli, and then categorised equally on the basis of frequency and amplitude. Using these results a set of seven haptic icons were selected for use. An evaluation [22, 23] of these seven haptic icons under various degrees of workload (when the user was also engaged in a visual or auditory task) showed identification rates of 95% on average, with no significant differences in performance with or without distractor tasks, indicating that vibrotactile icons can be an extremely effective means of communicating information. Chan *et al.*'s work is the first formal investigation into the design of vibrotactile icons apart from the work of this thesis, and, therefore, comparisons are drawn between this thesis work and Chan *et al.*'s work in Chapter 6.

Vibrotactile Instant Messaging

Rovers and van Essen [95] proposed the use of haptic icons as an alternative to emoticons in instant messaging applications. An emoticon, also called a "smiley", is defined as follows (from http://en.wikipedia.org/wiki/Emoticon):

> *"a sequence of ordinary printable characters, such as* :-), ;o), t(*_*t), ^_^ *or* :-(, *or a small image, intended to represent a human facial expression and convey an emotion."*

A haptic emoticon is a vibrotactile sensation that can be used to express emotions in place of emoticons. Rovers and van Essen suggest that these haptic icons could be triggered by typing the textual equivalent (just like emoticons), or via special input devices (e.g. by making a gesture on a touchpad). They would then be sent to a vibrating device at the receiver end (which could be a vibration-enabled mouse or joystick, or a small vibrotactile actuator attached to the receiver's clothing or inside a shoe [96]). The proposed haptic icons would be designed using parameters such as frequency, amplitude and duration, and would use metaphorical mappings that were related to what they represented [95]. For example, a big smile emoticon **:D** would be represented by a vibration with increasing amplitude burst that ends abruptly, while a dislike emoticon **:(** would feel more aggressive, consisting of three abrupt pulses of high amplitude. If the touchpad was used for input, a tickling gesture on the touchpad could trigger very weak force pulses, in order to get someone's attention [95]. No evaluation of these haptic icons has been carried out, but this proposal indicates that there is interest in the area of haptic icons for communication.

3.7 Vibrotactile Display for Enhancement of User Interface Elements

One other application in which vibrotactile feedback has been used is to enhance interaction with user interface elements such as menus, buttons, informing the user of system status and confirming user actions. Tactile feedback is present when interacting with real, physical widgets, and therefore may enhance interactions with virtual widgets. In addition tactile feedback can be used to provide information about system state while interacting with user interface elements. The majority of work on enhancing user interface elements using vibrotactile stimuli has focused on mobile devices, since many of these already contain vibrotactile actuators, but some work has also been carried out on the use of vibrotactile feedback to enhance desktop user interface elements.

3.7.1 Mobile Device Interfaces

Poupyrev *et al*. [86] implemented tactile feedback to enhance scrolling interactions in mobile device applications. In this system, the user tilted the device in the direction they wished to scroll and the speed of scrolling was dependent on how far the device was tilted. Tactile feedback (one short vibration) was provided as each line of text scrolled past. Therefore, the user could keep track of the speed of scrolling by monitoring the speed of the tactile pulses. An evaluation of this system [86] showed a 22% improvement in task completion time, with maximum benefit obtained in tasks where the user had to scroll through long distances and very short distances (little benefit was obtained for medium distances). It was concluded that the tactile feedback helped in long distance scrolling as users could scroll at the fastest speed and feel the scrolled amount, and in the one line scrolling because it helped with avoiding overshoots. However, in the medium distance scrolling, users scrolled at a comfortable medium speed and were able to use vision alone for precise selection, so the tactile feedback was not needed.

Poupyrev *et al*. [88] have also carried out research into the use of tactile feedback in touch screens for mobile devices, adding tactile feedback to buttons, sliders, text selection and drawing applications via a piezoceramic actuator located underneath the screen. There has been little formal evaluation of this tactile feedback. Tactile feedback for touch screens has also been explored by Kaaresoja *et al*. [63] for enhancing button presses, scrolling, text selection and drag and drop interactions, but again there has been no formal evaluation of this work.

Luk *et al*. [72] have investigated the presentation of haptic icons via a custom built device using piezoelectric skin stretch technology to provide tactile feedback in mobile devices. The device is mounted on a slider on the side of a phone, and thus tactile feedback can be presented to the finger, or thumb, in contact with the slider. This device is very different to standard vibrotactile devices as it stretches the skin laterally rather than pushing into it. In addition, it requires the device to be in contact with the fingertips unlike some vibrotactile devices which can be mounted on other body locations. Luk *et al*. conducted a perceptual characterisation study and found that a wide variety of distinguishable stimuli could be created by manipulating speed, direction, waveform, duration and amplitude. They suggest that these haptic icons could be beneficial for a range of use-cases within a mobile phone, such as list selection, scrolling, direction signalling and alerts. While scrolling through a list or a webpage, different haptic icons could be presented to represent different list items, or different web page elements. This would allow users to navigate through lists or webpages without looking at their phone. This may be particularly useful for menu navigation as users could discreetly set their phone to silent mode while in a meeting. The actual encoding of information in the vibration parameters established above has not been described.

In their VibeTonz system overview (http://www.vibetonz.com), Immersion propose a number of applications where tactile feedback could enhance interactions with user interface elements in mobile phones. For example, they suggest that tactile feedback could help users to find particular messages when scrolling through a list. Different sensations could represent read or unread messages, or could be assigned to messages from different senders. They also suggest assigning distinctive tactile sensations to different buttons, and providing tactile feedback to indicate that a call has been “dropped” (ie the person on the other end has been cut off).

The multi-function transducer in Motorola phones (also mentioned in Section 3.4) has been used in research evaluations to generate haptic icons to enhance the audio icons in user interface features such as buttons and navigation and for confirming actions [25, 83]. An initial evaluation of this phone enhanced with the tactile feedback compared to one with no tactile feedback showed that people perceived the sound quality of the tactile-enhanced phone to be better than that of the non-enhanced phone, indicating that the tactile feedback led to a richer experience. Unfortunately no information has been provided regarding the types of feedback used.

Finally, the Digital Pen by Nokia (http://www.nokia.co.uk/nokia/0,,18259,00.html) uses simple vibration feedback to enhance user interface features. A vibration pattern is used to confirm when an action has been successful, and also to indicate battery status.

3.7.2 Desktop Interfaces

Commercial products such as the AVB Vmouse (www.avbusa.com) and the (discontinued) Logitech ifeel mouse (www.logitech.com) provide tactile feedback relating to actions in the user interface. The VMouse provides vibrotactile feedback which automatically accompanies any audio feedback. The ifeel mouse provides more complicated tactile feedback in the user interface which is controlled by software called Immersion Touchware desktop (www.immersion.com). The touchware settings can be changed to provide a range of feedback for different interface elements, for example buzzing when the mouse is over an icon or providing tactile clicks when the user moves across menu items. There has been no formal evaluation of the effectiveness of the vibrotactile feedback provided by these devices.

Formal work on the use of vibrotactile feedback for targeting using a vibrating mouse was conducted by Akamatsu *et al.*[3]. They investigated the effect of vibration of the acquisition of single targets, and found that, while it did not significantly reduce targeting time, it did reduce the time the user spent over the target. Cockburn and Brewster [31] recently extended this work to take multiple targets into account. The results of this work showed that the addition of vibrotactile feedback (in the form of a constant 200Hz vibration when the mouse is over a target) can improve targeting speed for small discreetly located targets. However, it also showed that when multiple targets are present vibrotactile feedback can cause confusion and reduce performance as the feedback from different targets interfere with one another. This indicates that while vibrotactile feedback can improve performance, it must be used with caution.

3.8 Other Related Work

Another piece of related work on vibrotactile display, which does not fit into any of the previous categories, is the “Cutaneous Grooves” system developed by Gunther [52, 53]. Gunther developed a language for tactile composition to accompany and enhance musical performances. This application of vibrotactile stimuli is quite different from the others discussed in this chapter as the idea is not to encode specific information, but rather to create an aesthetically pleasing experience. Thus it is not essential that different levels of a parameter can be absolutely identified, rather it is just important that a wide enough range of sensations can be produced to create an interesting and dynamic experience. In the prototype system the “audience member” (the person experiencing the composition) wears a suit containing thirteen vibrotactile actuators. Twelve high frequency actuators are fixed at various points on the user’s body, along with a low frequency actuator (a wearable sub-woofer) attached to his/her lower back. Throughout a composition, vibrotactile stimuli can be presented to different locations on the body, and move across the body using apparent motion (see Chapter 2). In addition variations in intensity, duration, waveform and frequency can be used to enhance the experience. This tactile language was structured in the same manner as music, although the different parameters had different salience from the corresponding parameters in audio. For example, since the skin has low frequency resolution, frequency would be a less important parameter in musical composition than pitch in audio, while spatial location might have a higher prominence in tactile composition than in music. While this research does not investigate the encoding of specific data, and absolute identification of data values, this research is of interest in terms of the selection of parameters for use in vibrotactile display.

3.9 Discussion and Conclusions

This chapter has reviewed the use of vibrotactile display, and has shown that vibrotactile feedback has been used for a wide range of different purposes in many different applications. Five main uses of vibrotactile display have been identified:

1. Sensory Substitution: Vibrotactile display has been used to present speech and printed text to hearing-impaired and visually-impaired people by translating speech signals or alphanumeric-characters into temporal or spatial vibrotactile patterns.

2. Alerts: Vibrations are attention grabbing and are therefore useful for alerts and warnings. Vibrotactile display has been used for alerting purposes in a range of applications including mobile phones, cars, and smoke alarms for hearing-impaired people.

3. Providing navigation information: The large three-dimensional display space offered by the skin makes it a suitable choice for displaying directional information. This has been utilised for a range of navigation applications, including systems for car drivers, pilots and visually impaired people.

4. Communication: Vibrotactile display has been used in mobile and desktop interfaces to allow people to express their emotions through vibration enhanced text messages or instant messages. Vibration messages have also been used for a more limited form of communication in collaborative applications, where they can be used to signal the current state of a user or the user's intent. Finally, prototype systems have been developed allowing two-way communication via vibration combined with speech.

5. Enhancing user interface elements: A final use of vibrotactile display has been to enhance interactions with user interface elements in both mobile and desktop interfaces. Vibrotactile display has been used to enhance the feel of widgets such as virtual buttons, but has also been used to display information in the user interface, such as confirming user actions or indicating battery status.

This work reviewed in this chapter can be used to inform the work of this thesis in two ways. Firstly, it is possible to use the related work in this chapter to help to identify suitable parameters for use in vibrotactile display. Secondly, this work can be used to place the work of this thesis into context and to identify an area in the field of vibrotactile display which requires further investigation. The findings from this chapter with regards to these two aspects are discussed in the following sections.

3.9.1 Identification of Suitable Parameters

One aim of this review was to investigate the types of vibrotactile parameters used in other related work on vibrotactile display to inform the choice of parameters for Tactons and help to answer Research Question 1:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?

The main parameter used in this related work was spatial location which was mostly used to indicate direction in navigation applications (Section 3.5). Spatial location was also used in vision substitution systems to map the visual versions of alphanumeric characters to tactile versions, using either a pictorial or an abstract approach (Section 3.3.2). The review of tactile perception in Chapter 2 also indicated that spatial location was one of the most promising parameters for vibrotactile displays, since it offers a wide range of values.

Another parameter which has been used in much of the work in this chapter is temporal pattern or rhythm. Temporal patterns have been speeded up or slowed down to encode information about the distance to a target, or to indicate a change in speed or acceleration (Section 3.5). They have also been used in the design of haptic icons, with a different number of vibration bursts used to distinguish different icons (Section 3.6.2). In addition temporal patterns have been used in tactile hearing aids to encode the temporal patterns used in speech (Section 3.3.1), and in mobile phone alerts to accompany or replace the audio ringtones (Section 3.4.1). Rhythm was also identified to be a promising parameter in the review in Chapter 2.

Intensity has also been used in a number of systems in this chapter, for example in tactile hearing aids to encode the volume of speech (Section 3.3.1), and in Geldard’s Vibratese system [40] to encode an abstract dimension of alphanumeric characters (Section 3.3.2). However, some problems with absolute identification of intensity were identified in several systems, re-iterating the fact stated in Chapter 2 that this parameter should be used with caution.

As expected from the findings of the review in Chapter 2, frequency has been used in very few systems. It was used in early tactile hearing aids but problems were identified and the frequency of speech is now more often encoded in spatial location (Section 3.3.1). Different frequency levels were used by Tan *et al.* [105] in their multi-finger tactual display (Section 3.3.1) but a very high bandwidth device is required to produce this wide range of frequencies. Rabinowitz *et al.*[89] used frequency as a parameter but, again, used a high-bandwidth device (Section 3.3.1). These frequency ranges are beyond the capabilities of most small, commercially available actuators, including those used in this thesis.

Duration has only been used as a parameter by Geldard [40] in his Vibratese system (Section 3.3.2), although it has also been used within temporal patterns or rhythms. It appears that, while duration might be a useful parameter, it will be more useful when used as part of rhythms or temporal patterns. Therefore, from this point forth only rhythm is considered.

As indicated in the review of tactile perception in Chapter 2, waveform and intensity change over time remain largely unexplored parameters. Luk *et al.* [72] (Section 3.7.1) found that distinguishable waveforms could be created, but their system uses piezo-electric stretch technology rather than vibration, therefore the production of these stimuli and the perception of them is very different. Gunther *et al.* [52, 53] proposed the use of waveform and of intensity change over time for use in his tactile composition system (Section 3.8) but no data was provided on the perception of these parameters.

This review of the parameters used in existing systems matches with the review of tactile perception in Chapter 2, indicating that the most promising parameters are spatial location, rhythm, and intensity, and that vibrotactile waveform and intensity change over time are still largely unexplored parameters.

3.9.2 Identifying an Area for Further Exploration

Another aim of this review was to place the work of this thesis in context and to identify an area in the field of tactile display which has not yet been fully explored. This review has shown that there has been a wide variety of research into vibrotactile display. However, much of this research has been concerned with the use of simple vibration messages which usually only encode a single dimension of information. Some of these have only encoded one bit of information (e.g. vibration on or off to indicate whether a mouse is over a target or not [3, 31]), while others have encoded more detailed information (e.g. using a vibration at a specific spatial location to indicate the direction in which a driver should turn [111], or in which direction the driver's attention should be directed [58]).

As noted in Chapter 1, Sherrick and Craig [100] suggested that, to use the sense of touch to communicate information, two main problems needed to be addressed. Firstly, tactile parameters, and values of these which can be readily learned, need to be identified and, secondly, these values need to be combined to produce unique clusters. Much of the work reviewed in this chapter has only concentrated on the selection of values of one parameter, and not on the combination of several values of multiple parameters.

By following the approach suggested by Sherrick and Craig [100], it may be possible to encode more information in vibrotactile messages. For example, each parameter of vibration could encode a different dimension of information and, thus multiple dimensions of information could be encoded. This approach has been used in some sensory substitution systems. In tactile hearing aids different elements of speech signals are encoded in different vibrotactile parameters (e.g. frequency, location, temporal pattern, intensity), while, in the Vibratese language for vision substitution, several vibrotactile parameters are manipulated simultaneously to represent different alphanumeric characters. In addition, this method has been used by Rabinowitz *et al.* [89] to create multidimensional tactile displays for tactile communication of speech. However, outside of the field of sensory substitution the encoding of multiple dimensions of information using different vibrotactile parameters has been largely unexplored. This is, therefore, an area which can be explored in this thesis.

In some of the most recent work on tactile display reviewed in this chapter the term "haptic icon", or "vibrotactile icon" is used to describe the tactile messages [22, 72, 95, 111]. A haptic icon or vibrotactile icon is the touch equivalent of a visual icon. Audio counterparts to icons already exist in the form of both Earcons [4, 9] and Auditory Icons [39] but research into a tactile equivalent of the icon is in its infancy. At the start of this thesis research, there had been very little formal research in this area. Although this area has now received some interest, there has still been little formal work into how best to design vibrotactile icons. Furthermore, there has been little research into how the approach defined above can be applied to the design of such vibrotactile icons. Therefore, this thesis will investigate how vibrotactile icons can be designed by first examining the individual parameters of vibration and then investigating how they can be combined to encode multi-dimensional information. The approach to the design of these vibrotactile icons is presented in the following chapter.

Chapter 4: Designing Tactons

4.1 Introduction

Chapter 3 reviewed related work in the field of vibrotactile display and showed that it has been used for a range of purposes, in particular for sensory substitution, providing alerts, navigation information, communication and enhancing user interface elements. One form of vibrotactile display which has recently received interest is that of vibrotactile icons (also called haptic icons). While vibrotactile icons have been proposed by a number of researchers, until recently there had been little formal investigation into how best to design them. Some formal work has now been carried out [22, 23, 72] but this is still in the early stages, and there are other approaches to vibrotactile icon design which can be considered. One avenue which has received little attention is the use of a structured, abstract approach to encoding information in vibrotactile icons and this approach is explored in this thesis. This chapter discusses what vibrotactile icons are and outlines the approach used in this thesis to design Tactons [8] (a specific form of vibrotactile icons).

4.1.1 Vibrotactile Icons

Shneiderman [5] defines an icon as "an image, picture or symbol representing a concept". Audio counterparts to icons already exist in the form of both Earcons [4, 9] and Auditory Icons [39] but there has been less formal work on how best to design a counterpart of icons and Earcons presented to the sense of touch. The most comprehensive definition of haptic icons is given by Enriquez and MacLean [34] who define haptic icons (or Hapticons) to be "brief computer-generated signals, displayed to a user through force or tactile feedback to convey information such as event notification, identity, content or state." The term Haptic Icon, or Hapticon, can therefore be taken to be the all-encompassing term for any type of force feedback or tactile stimuli conforming to this definition. Within the field of haptic icons there may also be a number of sub-types. In particular vibrotactile icons are haptic icons which make use of vibration, while force feedback icons utilise force-feedback stimuli.

This research aims to investigate formally the design of vibrotactile icons called Tactons (tactile icons): a specific type of tactile icons [8], which are designed using a structured, abstract approach. The term Tacton can also refer to tactile icons presented using different types of tactile technology, e.g. pin arrays [8], but the research in this thesis only focuses on the design of vibrotactile Tactons. Therefore, in the remainder of this thesis the term Tacton only refers to the vibrotactile version.

The definition of a Tacton is discussed in detail in this chapter and the basic design principles that will be used to design these messages are presented and justified. Different encoding strategies are discussed in Section 4.2 while the implementation of the chosen strategy is presented in Section 4.4.1. The selection of vibrotactile parameters with which to encode information is discussed in Section 4.4.2. Finally, possible applications of Tactons are presented in Section 4.5.

The research questions posed in Chapter 1 were:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

Research Question 1 is addressed in the discussion of parameters for Tactons in Section 4.4.2. This chapter does not address Research Question 2 directly, although Section 4.4.1 presents a strategy for how the parameters can be combined to create Tactons.

4.2 Selecting a Strategy for Tacton Design

When designing vibrotactile icons, different approaches to encoding information could be considered. Two main strategies were identified in Chapter 3: the Pictorial/Direct approach and the Coded approach. In addition, from the review of literature in Chapter 3 a third category emerged: the metaphorical approach, where the vibrotactile stimulus is not a direct transfer of information from another sense but uses a metaphor with which the user may be familiar, for example, the heartbeat sensation used by Chan [22] to indicate that the user currently has control of the application, or the hug-like sensations proposed by Immersion in their VibeTonz system (http://www.vibetonz.com). Other approaches to encoding information in tactile messages might be found by investigating the area of semiotics to understand the meanings that people assign to different types of touch. However, this was outside the scope of this thesis and, therefore, is not explored here.

Pictorial/direct mappings can be advantageous as they should be easier for people to learn and remember since they are relatively self-explaining. However, Chang and O'Sullivan [25] state that "touch is hard to fake", meaning that, with the currently available devices, it is difficult to create vibrotactile sensations that feel like real world sensations. In addition, there are few real world vibration sensations with which people are familiar. This, therefore, makes it impractical to create tactile messages that feel like the things they represent. Metaphorical mappings overcome this problem somewhat, and should be easier to learn and remember than abstract mappings. However, they still rely on the existence of sufficient tactile metaphors. Using abstract mappings requires that users learn the mapping between tactile messages and the data they represent, which means that they may be more difficult to learn and remember. However, using these mappings means that it is possible to create tactile messages to represent any data, even when an appropriate metaphor does not exist, and to use the full range of stimuli that the actuators can produce.

A large body of work exists on the design of audio icons and, since sound and vibration are both created from the same source and are both senses which are temporal in nature, it seems likely that work from the audio domain could also be used to inform Tacton design. Two main types of audio icons exist: Auditory Icons [39] and Earcons [4]. Auditory Icons are natural, everyday sounds used to represent events or items within a computer interface. The sounds that are used are semantically linked to the things they represent, and the meanings should therefore be easy to learn and remember. This approach is similar to the pictorial approach. Earcons, on the other hand, are structured, abstract non-speech audio messages which use an approach like the coded approach in the vibrotactile domain. Earcons use musical, rather than natural, sounds and use an abstract mapping that must be learned as there is no semantic link between the sounds and the data they represent.

Earcons have been shown to be a successful means of communicating information in user interfaces (e.g. [9, 10, 70, 76, 77]), but there has been very little investigation into how structured abstract tactile messages could be used to encode information. Chan's work [22, 23] used a structure to link some of the haptic icons used in his work, but there was a metaphorical mapping from the haptic icons to the data they represented. The messages proposed by Rovers and van Essen [95] used a metaphorical mapping and no structure was used in their design. Outside of the domain of tactile icons, there has been some work investigating the use of structured abstract messages for presenting complex information such as language and speech. For example, Geldard's Vibratese language [40] encoded each character of the alphabet in an abstract combination of different tactile parameters. This approach is similar to the approach that will be used in this research, but only explored one set of parameters, so there is still research to be done in this area. In addition, Geldard's work only used the most basic parameters of vibration – location, duration and intensity - whereas Tactons might also make use of parameters borrowed from the audio domain such as rhythm. Tan's multi-finger tactual display [105] also used structured, abstract messages to communicate speech information, but the system used kinesthetic cues as well as vibration cues. The work most similar to the approach used in this thesis is the study by Rabinowitz *et al.* [89] investigating multidimensional tactile displays, where multiple parameters of vibration (intensity, frequency and contactor area) were varied to encode information. This approach is very similar to that used in this thesis, but only explores one set of parameters and, therefore, other parameter sets can still be explored. All three of these systems have been designed to communicate speech and language; there does not appear to have been any research into how structured, abstract vibrotactile messages might be used to communicate information in computer interfaces in the same way that Earcons do in audio.

One novel contribution of this research is the application of Earcon design principles to the problem of vibrotactile icon design, learning from both the structure of Earcons and the parameters used to encode information in them. Earcons have been designed by identifying parameters of audio that users can identify multiple levels of, and then combining these to create unique motifs, or Earcons. This is the approach suggested for tactile communication by Sherrick and Craig [100], which was the basis of this thesis and, therefore, it seems appropriate to model the design of Tactons on the design of Earcons. Tactons can be defined as structured, abstract vibrotactile messages for non-visual information display [8]. Tactons are thus one specific type of vibrotactile icon, with the key aspects being that they are structured and abstract. They use a coded approach as opposed to a pictorial or metaphorical approach.

4.3 Earcon Design Principles

Since the aim of this research is to apply Earcon design principles to the design of Tactons it is important to understand how Earcons are designed. The concept of Earcons was originally conceived by Blattner *et al*. [4], and a large body of work has been carried out on the design and evaluation of Earcons since then (e.g. [9, 10, 70, 76, 77]). This section summarises the basic principles of Earcon design which are of relevance to the design of Tactons.

Blattner *et al.* [4] defined Earcons as “non-verbal audio messages used in the user-computer interface to provide information to the user about some computer object, operation, or interaction”. They described two different types of Earcons: *representational* Earcons and *abstract* Earcons. Representational Earcons are the same as Auditory Icons, and in recent years the term Earcon has only referred to the abstract type. Brewster [10] revised the definition of Earcons to account for this, defining Earcons as “abstract, synthetic tones that can be used in structured combinations to create auditory messages”. The key words in this definition are “abstract” and “structured”. Unlike auditory icons, which use natural everyday sounds which have a semantic link to the data they represent, Earcons use musical sounds, and the mapping between the data and the sounds is abstract. The other key feature of Earcons is the structure which underlies their construction. Section 4.3.1 outlines the three main structures used in Earcon design.

4.3.1 Earcon Structure

Earcons are constructed by combining and manipulating single pitches and motifs (short melodies) to encode complex information. Blattner [4] states that this structured approach to encoding information in audio messages has the following advantages. If the same type of structured approach is applied to the design of tactile icons, the same advantages should apply.

1. “It is systematic, with well-defined building blocks that may be used to construct larger sets of Earcons. A systematic approach is usually straightforward to understand and use. Modularity allows easy modification, future expansion and tailorability.” [4]

2. “Motifs may be transformed, combined, or inherited, thereby creating families of related motifs.” [4]

3. “Motifs may be grouped into families, where Earcons with similar meanings have similar sounds; different families will sound dissimilar.” [4]

A number of different structures has been used for Earcon design, and may also be applicable to Tacton design. These are summarised below.

One-Element Earcons

A one-element Earcon is the simplest type of Earcon, which encodes a single piece of information in a single musical note or a short motif. For example, Figure 4-1 shows four one-element Earcons designed by Blattner *et al.* [4] to present information about objects or actions in a computer environment.

Compound Earcons

A Compound Earcon can be created by presenting two or more one-element Earcons in succession (in the same way that words are combined to create sentences). For example Blattner suggested that Compund Earcons could be created to represent actions in a computer interface. One-element Earcons can be created to represent ‘Create’, ‘Destroy’, and ‘File’ and ‘String’ (Figure 4-1) and then combined to create compound Earcons meaning ‘create file’ or ‘delete string’ as shown in Figure 4-2.

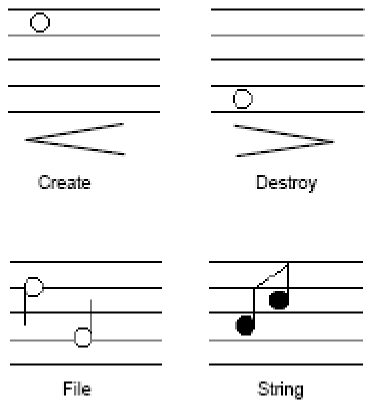

Figure 4-1: Four one-element Earcons representing ‘create’, ‘destroy’, ‘file’ and ‘string’ (from [10], after [4]).

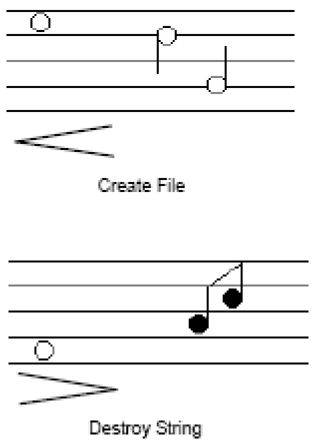

Figure 4-2: The combination of one-element Earcons to form Compound Earcons representing ‘create file’ and ‘destroy string’ (from [10], after [4]).

Inherited Earcons

In addition to Compund Earcons, two types of “family” Earcon exist: Inherited Earcons (also known as Hierarchical Earcons [10]) and Transformational Earcons [4]. A “family” Earcon is one which belongs to a group or family of Earcons which are related by inherited or transformed attributes. This structure should make it easier for users to learn a large set of messages since only a few rules need to be learned and then a relatively large set of Earcons can be understood.

Inherited Earcons can be created using a tree structure, where every node in the tree is an Earcon, and each Earcon inherits from the levels above it. The example in Figure 4-3 shows the inheritance structure for a family of Earcons representing errors in a computer interface. The family is assigned a rhythm which will be used by all Earcons in that family. A different family representing another aspect of the computer interface would be represented by a different rhythm. At the top of the tree the rhythm representing the error family is shown; this rhythm would be presented with an un-pitched click. At the second level the two types of error inherit the family rhythm, and this rhythm is presented, followed by the same rhythm with a different set of pitches (different melody) used to represent the two different errors. In the third level a three-part Earcon is presented: first the family rhythm is presented, then the second level Earcon, and then the melody is presented for a third with the timbre used to present the melody changed to represent the exact error message.

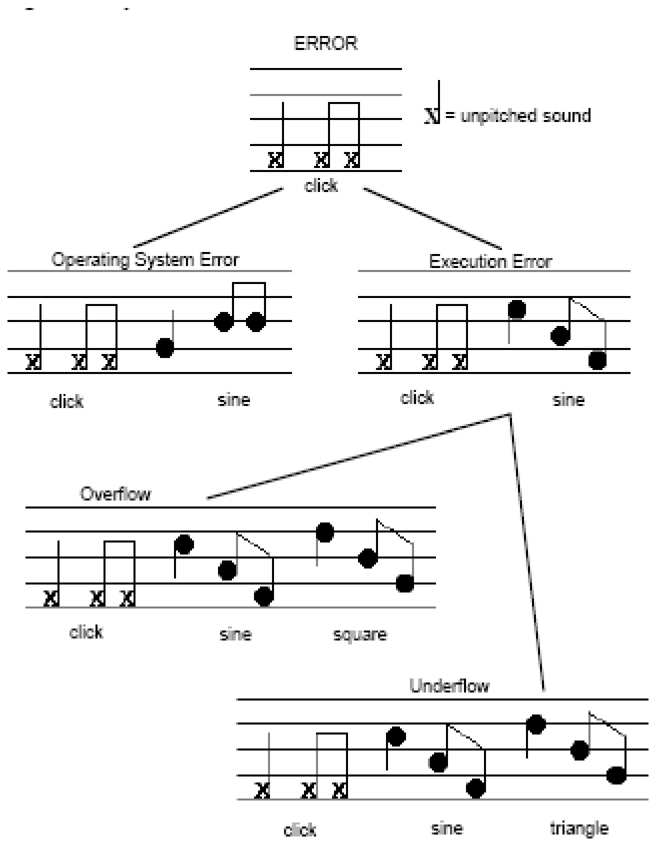

Figure 4-3: A hierarchy of family earcons representing errors (from [10], after [4]).

Transformational Earcons

Transformational Earcons [4, 77] are the other type of "family" Earcons. In Transformational Earcons each data attribute is mapped to a parameter of the Earcon. For example, if Transformational Earcons were used to represent files in a computer interface, the file type could be represented by rhythm, size by the instrument used, and creation date by the register (overall level of pitch). Each file type would be mapped to a unique rhythm. Therefore, two files of the same type and same size, but different creation date, would share the same rhythm and be played by the same instrument, but would be presented in a different register. If two files were of different types, but were of the same size and shared the same creation date, they would be represented by different rhythms but would be played by the same instrument and in the same register.

4.3.2 Earcon Parameter Space

From the results of his PhD research, Brewster produced a set of guidelines for designing Earcons, in particular on the use of different parameters [10, 11]. The following sections briefly summarise his findings and guidelines for each of the parameters in Earcons which may be useful to Tacton design.

Timbre

Brewster [10] showed that using musical instrument timbres (e.g. piano, trumpet) in Earcons was more effective than using simple tones (e.g. sine waves, square waves). He also reported that it is important to select timbres that are subjectively easy to tell apart. McGookin [77] found that users were able to achieve absolute identification rates of over 90% for timbre (piano/violin/trumpet) when it was used in three-dimensional Transformational Earcons in combination with rhythm and register.

Register

The term register refers to the overall pitch level of an Earcon. Brewster [11] suggests that register is a poor choice when absolute recognition is required and therefore it would be better to use it in combination with another parameter. If register must be used then large differences between the different levels will be required. When participants were asked to carry out absolute identification experiments on three-dimensional Transformational Earcons (created by manipulating rhythm, register and timbre), McGookin [77] found that absolute identification rates of around 70% were achieved for register (low/med/high).

Rhythm

Brewster [11] stated that in order to ensure that rhythms can be absolutely identified they should be designed to be as different as possible from one another. One method of doing this is to use a different number of notes in each rhythm. He also notes that to make Earcons sound like a complete rhythmic unit, certain musical rules should be followed. In particular, the first note should be accented (played slightly louder) and the last note should be slightly longer. The other important guideline pertaining to rhythm is that Earcons should be kept as short as possible so they do not slow down interactions when used in human computer interfaces. As discussed in Chapter 2, perception of rhythm in both audio and vibrotactile modalities seems to be relatively similar, therefore it may be possible to learn from these guidelines when designing rhythms for Tactons. McGookin [77] found that users were able to achieve absolute identification rates of over 90% for rhythm (melody) when used in three-dimensional Transformational Earcons in combination with timbre and register.

Intensity

Brewster [11] recommended that intensity should not be used as a parameter in Earcons because users find loud sounds annoying and report annoyance when the volume level is out of their control. Using intensity as a parameter in Tactons could be equally problematic as reducing the intensity could degrade perception of other parameters, or render the signal undetectable, while increasing it too far could cause pain [7]. Therefore, intensity should be used with caution.

Spatial location

Spatial location has not been used much in Earcon design, except to help differentiate multiple Earcons presented simultaneously [75, 79]. Brewster [11] suggested that different families of Earcons could be presented to different locations but this has not been investigated. Spatial location is likely to be a more prominent parameter in Tactons, and could be used either to differentiate multiple Tactons or as a parameter to represent different families of Tactons.

4.3.3 Summary

The research summarised in this section has shown that there are several types of Earcons: *One-element*, *Compound*, *Inherited* and *Transformational*. In addition, it has summarised the parameters which have been used in Earcon design and identified that the most successful of these parameters are timbre and rhythm. Register seems less successful but has also been used in Earcon design.

4.4 Tacton Design Principles

Before designing Tactons a number of decisions need to be made about how they should be constructed. The first decision, whether to use a direct, metaphorical or abstract mapping has been discussed in Section 4.2. In addition, it is necessary to decide how Tactons should be structured and this is discussed in Section 4.4.1. Finally a decision needs to be taken regarding which parameters of vibration should be manipulated to encode information and this is discussed in Section 4.4.2.

4.4.1 Structuring Tactons

Tactons are the vibrotactile equivalent of Earcons, so research on Earcons can be drawn on in order to design Tactons. As discussed in Section 4.3.1 several different structures are used to construct Earcons, and these structures may also be applied to Tacton design. Some possible implementations of Tactons using these structures are discussed below.

One-Element Tactons

A one-element Tacton, like a one-element Earcon, encodes a single piece of information in a short vibration burst or temporal pattern. Building on Blattner's Earcon example from Section 4.3.1, these same items could be represented by one-element Tactons. For example, a vibration that increases in intensity could represent "Create", while a vibration that decreases in intensity could represent "Delete". A temporal pattern consisting of two short vibration bursts could represent a "file", while a temporal pattern consisting of three long vibration bursts could represent a "String". These simple vibration messages only encode one piece of information and are not related by a structure. The benefit of Tactons is that they can encode multiple dimensions of information and are connected by a structure which makes them easy to learn. The following types of Tactons utilise a structure so as to encode multidimensional information.

Compound Tactons

As with Earcons, one-element Tactons could be combined to create Compound Tactons[8]. For example, presenting the “Create” Tacton followed by the “File” Tacton would indicate that a file had been created.

Inherited Tactons

Inherited, or hierarchical, Tactons are created in the same way as inherited Earcons. Each Tacton represents a node in a tree and inherits properties from the levels above it [8]. The top of the tree uses a Tacton which has a basic rhythm played using a sinewave (a different tree would use a different rhythm so it is not confused). The rhythmic structure of Level 2 inherits the Tacton from Level 1 and adds to it, for example by following it with the same rhythm presented using a different waveform. This continues by manipulating a different tactile parameter at each level such that a hierarchical structure is presented.

Transformational Tactons

A Transformational Tacton encodes multiple dimensions of information, with each represented by a different tactile parameter [8]. For example, if Transformational Tactons were used to represent files in a computer interface (as in the example for Transformational Earcons above) the file type could be represented by rhythm, size by frequency, and creation date by body location. Each file type would be mapped to a unique rhythm. Therefore, two files of the same type, and same size, but different creation date would share the same rhythm and frequency, but would be presented to a different body location. If two files were of different types, but were of the same size and shared the same creation date they would be represented by different rhythms but would be played at the same frequency and presented to the same body location. This transformational structure lends itself to orthogonal coding, with each vibrotactile parameter encoding a different dimension of information (See Chapter 1, Section 2.3.7), although redundant coding of one dimension could be used to improve performance if necessary.

4.4.2 Selection of Vibrotactile Parameters

Research question 1 in this thesis asks:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?

In order to address this question, it is important to consider what features are important in a parameter for use in Tactons. The following two factors seem to be the most important in the choice of parameters.

1. The user must be able to distinguish and identify multiple levels of the parameter

2. The user must be able to label the values of each parameter.

Factor 1 is important as multiple levels must be able to be identified so that more than one level of a dimension can be encoded. Therefore, two or more levels must be identifiable for a parameter to be successful. For many vibrotactile parameters, Geldard [41] suggested that three levels is the most that can be expected without extended training and therefore this is the number of levels of each parameter that is aimed for in this research.

In addition to being able to distinguish between different levels of a parameter, it is important that users are able to name the different levels. Being able to label these levels should help users to remember them and also provide a means of explaining what they are to other people. In audio this is made simpler by the existence of music, as labels such as instrument (e.g. trumpet/piano/violin), pitch (high/low), and melody are familiar to most people and can be used in Earcons with little explanation. As reported in Section 4.3.2, Brewster [10] showed that using musical timbres in Earcons was more effective than using simple tones, e.g. sine waves, square waves, and it is likely that this was at least partly due to the difficulty in labeling such abstract tones, compared to the ease of recognising a musical instrument, e.g. trumpet, piano. Many vibrotactile parameters are abstract and unfamiliar, so it is very important to find a way that people can quickly learn to label the values of these parameters.

Vibrotactile Parameters

The findings in the review of perception of vibrotactile parameters reported in Chapter 2 indicated that the most promising parameters for encoding information in vibrotactile messages, such as Tactons, are spatial location, duration and rhythm. This review also indicated that intensity may also be a usable parameter, but needs to be used with care due to interactions with other parameters. Chapter 3 indicated that spatial location and rhythm have been the most used parameters, and intensity has been used but some problems have been reported. In addition to these parameters, the review in Chapter 2 identified two parameters which have received little attention, namely waveform, and intensity change over time. These parameters have not been formally investigated, but may be successful parameters for Tacton design.

One other possibility that can be considered is the use of parameters that have been successful in Earcons. As discussed in Section 4.3.2, the most successful parameters in Earcon design are timbre, rhythm and register. Timbre is the audio equivalent of tactile waveform and an experiment exploring this parameter is reported in the following chapter. Rhythms exist in both domains (although in the tactile domain they are also referred to as temporal patterns) and were used in much of the related work on tactile display reported in Chapter 3. The review of rhythm perception in Chapter 2 indicated that the perception of both audio and vibrotactile rhythms is similar and this parameter may transfer successfully between modalities. It may, therefore, be beneficial to learn from the Earcon guidelines on use of rhythm, e.g. using a different number of notes in each rhythm to make them more distinguishable. Register (or pitch) is the audio equivalent of frequency in the tactile domain and the review in Chapter 2 indicated that frequency would be a poor choice due to the limited frequency range of the skin, interactions between frequency and subjective magnitude, and the limited resolution of the available devices. In addition, frequency was rarely used in any of the systems reviewed in Chapter 3; in fact it was reported that it was avoided in some tactile aids due to the reasons outlined above. Volume (the audio equivalent of intensity) is avoided in Earcon design due to the fact users do not like the volume to be outside of their control, as it can be annoying at high volumes, and hard to hear at low volumes. Using intensity as a parameter in Tactons could be equally problematic as reducing the intensity could degrade perception of other parameters, or render the signal undetectable, while increasing it too far could cause pain. Spatial location has rarely been used in Earcon design, except to differentiate between several Earcons presented simultaneously, but may be more useful in Tacton design, where other parameters are more limited.

4.5 Applications of Tactons

The review in Chapter 3 showed that vibrotactile feedback can be used in a range of applications, and Tactons could be used in many of these applications. As Tactons use abstract rather than direct mappings they are not tied to specific applications and therefore the same Tactons can be used for a range of different applications. The focus of this thesis is on how to design Tactons, rather than on any specific applications. Nonetheless, it is useful to consider some of the applications in which Tactons might be used to place this work in context. Some potential uses of Tactons are discussed below, in relation to the different applications of vibrotactile information display identified in Chapter 2.

Vibrotactile display is useful for alerts as it is attention grabbing. While tactile alerts can be provided through simple vibrations, the use of Tactons would allow more complex information to be presented in these alerts. For example, in a mobile phone a Tacton could indicate who the call was from and how important it was in addition to simply alerting the user to the fact that a call had occurred. In addition, Tactons could be used to provide alerts to people working in loud environments. For example, an initial requirements capture with a local fire brigade indicated that Tactons could be useful for alerting firefighters to various events and warnings, e.g. to indicate that they should evacuate or stop searching, or to enable them to monitor their air supply or current temperature. Some of this information is currently provided by audio alerts, but providing this through vibrations as well would mean that these messages would not be missed in the noise of the fire.

Vibrotactile display has also been used to aid navigation. In general, such applications have used simple vibrations presented to the torso to indicate the direction the user should take. Using Tactons could enable more detailed information to be provided alongside these directional cues. For example it might be possible to provide contextual information to a blind person as they navigate around a town. Tactons could be used to present information such as the type of building (shop, bank, office-block, house), the type of shop (clothes, phones, food, furniture) the price bracket of a shop (budget, mid-range, expensive), or information more related to the concerns of visually impaired people, such as the number of stairs leading up to a shop entrance.

Another application of vibrotactile display is in communication, because it offers a private channel through which information can be displayed to the user without disturbing others. The use of Tactons for communication would enable the communication of a set of pre-defined messages. For example, one suggestion in Chapter 3 was that tactile messages could be used to inform a politician of the reception he was receiving from the audience during a live debate. Tactons could provide an effective means of communicating this information.

Tactons could also be used to enhance interactions with user interface elements. In Chapter 3 it was suggested that different types of vibrotactile feedback could be used to represent different categories of message (e.g. read/unread) when scrolling through a list; Tactons lend themselves well to the encoding of this type of information.

4.6 Conclusions

This chapter has provided an overview of the design of Tactons, discussing both how they could be structured, and the parameters which could be manipulated to encode information in them. While other work has been carried out on the design of vibrotactile icons, this work is still in the early stages and other design approaches can be explored. The approach taken in this thesis is to base the design of Tactons on the design principles employed in the creation of their auditory counterparts, Earcons. In particular, Tactons are a specific type of vibrotactile icon, in which multi-dimensional information is encoded using a structured, abstract approach. This approach has rarely been used in tactile display and has not previously been applied to the design of vibrotactile icons. Using these design principles, along with the understanding of cutaneous perception established in Chapter 2, Tactons can be created and evaluated through empirical studies.

Research Question 1 asked:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?

This chapter has established a set of parameters which might be suitable for use in Tacton design, namely rhythm, waveform, spatial location, and intensity change over time. Chapter 5 explores two of these parameters (waveform and intensity change over time) in more detail through experimental evaluations. The remainder of the work in the thesis is concerned with the design and evaluation of Tactons which have been created based on the design principles established in this chapter.

Chapter 5: Experimental Work on Individual Parameters

5.1 Introduction

The first stage in designing Tactons is to select individual parameters that can be manipulated to encode data. Research Question 1 is concerned with this stage of the design process:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?

Chapters 2-4 identified, from the literature, a set of possible parameters for use in Tactons, namely rhythm, spatial location and intensity. In addition they identified two parameters that require further investigation: waveform and intensity change over time. These two parameters have received little attention and, therefore, require some experimental investigation in order to understand how they are perceived before they are used in Tacton design. This chapter reports two studies investigating the perception of these two parameters, and discusses the implications of these findings for the use of these parameters in Tactons. In addition design recommendations are drawn out from the results of these studies to help designers who wish to use these parameters.

5.2 Vibrotactile Waveform

Very little research has been carried out into perception of different vibrotactile waveforms. When selecting waveforms for use in vibrotactile displays, an obvious first step would be to investigate identification of simple waveforms such as sine waves, square waves and sawtooth waves, and this has been mentioned by Gunther [52]. However, the limited bandwidth of many commercially available vibrotactile actuators means that the differences between these waveforms cannot be accurately reproduced. For example, measurements of the output of the TACTAID VBW2 actuator (see Chapter 2) rendered the outputs shown in Figure 5-1 and Figure 5-2 for sine waves and square waves respectively. The sine wave is accurately reproduced, but the limited frequency range offered by the actuator means that all but two of the frequency components of the square wave are lost, resulting in an output waveform which looks (and feels) very similar to that of a sine wave.

Since the differences between simple waveforms such as sine waves and square waves cannot be reproduced accurately by many transducers (and, therefore, cannot be distinguished) it is necessary to consider different types of waveforms that can be reproduced by the devices and distinguished by users.

In addition to being able to distinguish between different waveforms, another important factor in the selection of parameters, as discussed in Chapter 4, is that users must be able to name the different levels of the parameter, so identifying waveforms that can be labeled by users is also important. In Chapter 2, it was identified that one method of creating distinct waveforms might be the use of amplitude modulation or noise stimuli to create vibrations that feel rough, in contrast to the smooth sensations generated by sine waves. The use of descriptive terms such as “rough” and “smooth” makes such choices practicable for use in Tactons. To a non-scientist the term “amplitude modulation” may be meaningless, and even people familiar with the term would be unlikely to have any idea what such a waveform would feel like. By naming these waveforms in terms of how rough or smooth they feel, people may be able to label them more easily.

When labelling a stimulus in terms of a real world tactile parameter such as roughness, it is important to consider the relationship between the perceived stimulus and the parameter it is meant to represent. This can be considered in terms of the relationship between the proximal and distal stimuli. The proximal stimulus is the sensory information perceived by a person touching an object, while the distal stimulus is the actual object, or the distant source of the sensory information [92]. When modelling a physical stimulus such as roughness it is important that the person receives enough of the important sensory information about the stimulus to be able to correctly infer what the distal stimulus was meant to be. Rosendberg and Adelstein [92] note that this does not need to be an accurate representation of the physical stimulus but, rather, needs to be a perceptually adequate model so that the user can build a mental model of what the stimulus was meant to be. In the case of roughness, the proximal stimulus perceived by the user when amplitude modulated waveforms are presented to them should feel sufficiently like the experience of interacting with a rough object.

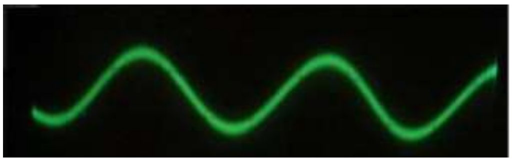

Figure 5-1: Output of TACTAID VBW32 when a sine wave was input.

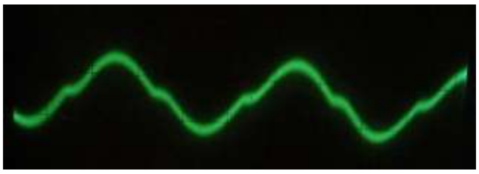

Figure 5-2: Output of TACTAID VBW32 when a square wave was input.

5.2.1 Vibrotactile Roughness

While the perception of roughness is generally associated with sensations resulting from active lateral movement across a surface [21, 66, 69], research has shown that people label certain types of vibrations as “rough” or “smooth” indicating that roughness can be generated through vibration alone [110, 119]. The results in the literature reviewed in Chapter 2 suggest that sinusoids feel “smooth” and that “rough” sensations can be created using amplitude modulated waveforms or noise stimuli [110, 119]. As reported in Chapter 2, Weisenberger [119] found that when subjects were presented with pairs of vibration stimuli, where one was an un-modulated sinusoid and the other an amplitude modulated version of the same sinusoid, they reported that the modulated sinusoids felt “rougher” than the un-modulated sinusoid. Weisenberger’s results on detection of amplitude modulation showed that maximum sensitivity occurred when modulation frequencies were in the range of 30-60Hz, with peak sensitivity at 40Hz, for 250Hz sinusoids.

This relationship between amplitude modulation and perceived roughness is also supported by research in the audio domain [35, 109]. Since both sound and vibration are generated by the same source, and both are temporal in nature, the production of roughness is similar in both modalities and, therefore, it may be beneficial to learn from these results. Terhardt [109] reported that audio roughness can be created by amplitude modulation, and also by frequency modulation or audio beating (the pulsing sound which occurs when two tones, close in frequency, are played simultaneously). He reported that the most important factors affecting the perceived roughness of amplitude modulated signals are modulation frequency and modulation depth. He also noted that, below 20Hz, the listener recognises the individual fluctuations within a signal, whereas above that point the individual fluctuations are no longer perceived as separate events, and the signal sounds “rough” or “harsh”. Fastl [35] reports that above 20Hz the perceived roughness increases as modulation frequency increases, until high frequencies where the ear can no longer detect the fluctuations (around 1000Hz), at which point the roughness disappears. This range is likely to be more limited for tactile roughness due to the limited frequency range perceivable by the skin compared to the ear.

5.2.2 Stimuli Design

While the sensation of vibotactile roughness has been mentioned by researchers, there has been no investigation into how best to use it in vibrotactile display. This section discusses how to create vibrotactile roughness stimuli, and Section 5.3 reports a study investigating perception of these stimuli. These results will indicate how best to use vibrotactile roughness to encode information in vibrotactile messages such as Tactons.

The implementation chosen for vibrotactile roughness was to use amplitude modulation. Noise stimuli could also have been used, but amplitude modulation offers finer control as both modulation frequency and modulation depth can easily be manipulated. As reported above, the two main factors used to manipulate audio roughness are modulation frequency and modulation depth [109]. Changing the modulation frequency changes the speed of the fluctuations, while modulation depth affects the depth of the modulation (the relative amplitude difference). When creating vibrotactile roughness stimuli, informal pilot studies indicated that changes in modulation frequency felt more distinctive than changes in modulation depth, so roughness was controlled by manipulating the modulation frequency.

Amplitude modulation is generated by multiplying a sinusoidal carrier by a sinusoidal modulator with a DC component added to it. This is shown in the following equation, from [119], where m is the modulation depth, f_m is the frequency of the modulator sinusoid, and f_c is the frequency of the carrier sinusoid. The resulting signal is V(t).

$$V(t) = V_0(1 + m\sin 2\pi f_m t)\sin 2\pi f_c t$$

This equation was used in MATLAB (http://www.mathworks.com/) to generate a range of vibrotactile roughnesses by varying the modulation frequency (f_m). The Matlab code is included in Appendix A.1. As discussed above only the modulation frequency was altered to create these signals: the modulation depth (m) was held constant at 1 for all signals. The value of the carrier frequency (f_c) was held constant at 250Hz for all signals. 250Hz was chosen as the carrier frequency for all stimuli because it is the frequency at which the skin is most sensitive, and the devices used in this research exhibit peak response at this frequency (see Chapter 2). 20Hz was chosen as the lowest modulation frequency based on results from the audio domain indicating that this is the minimum frequency at which roughness is perceived (rather than perceiving the individual fluctuations). Informal pilot testing confirmed that a 250Hz vibration signal modulated by 20Hz was considered to feel "rough", rather than like individual pulses. These pilot studies also indicated that the vibration started to feel smooth above 50Hz, so this was chosen as the highest modulation frequency. Within this range, five stimuli were created: sine (an un-modulated 250Hz sine wave), mod20 (the same sine wave modulated by 20Hz), mod30 (modulated by 30Hz), mod40 (modulated by 40Hz), and mod50 (modulated by 50Hz). This range of modulation frequencies from 20-50Hz quite closely matches the range identified by Weisenberger [119] as the most sensitive to amplitude modulation, namely 30-60Hz. The resulting waveforms are shown in Figure 5-3.

These waveforms were converted to audio files in MATLAB using the WAVWRITE command, and could then be presented to users by playing the audio file with the actuator plugged into the headphone socket on a PC. These audio files can be found in Appendix H on the accompanying CD by opening the html file "Exp 1 – Roughness.html".

Amplitude modulation creates different waveforms and therefore vibrotactile roughnesses could be considered to be an analog to timbre (or instrument) in the audio domain. But since the frequency of the modulation changes it could also be considered to be an analog to the pitch. In fact the amplitude modulation implementation is quite similar to the pulse stimuli frequency implementation mentioned in Chapter 2. The use of amplitude modulation to create distinct tactile sensations was also suggested by Gunther [52], who stated that, in music, amplitude modulation is known as vibrato, and that using this in the vibrotactile domain could create a tactile version of vibrato.

5.3 Experiment 1: An investigation into perception of vibrotactile roughness

An experiment was conducted to determine whether people could differentiate between the five different amplitude modulated signals described in 5.2.2 in terms of roughness. The aim was to identify three or more values of roughness that could be used in Tacton design, so it was necessary to find out whether there were three values that people consistently labeled as rougher or smoother than each other, i.e. a stimulus that could be labeled "smooth", one that could be labeled "rough" and another that could be labeled “very rough”. In addition the experiment aimed to find out in which order people placed the stimuli in terms of roughness.

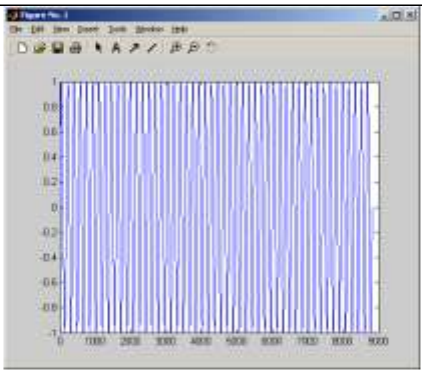

(a) sine: Unmodulated 250Hz sine wave

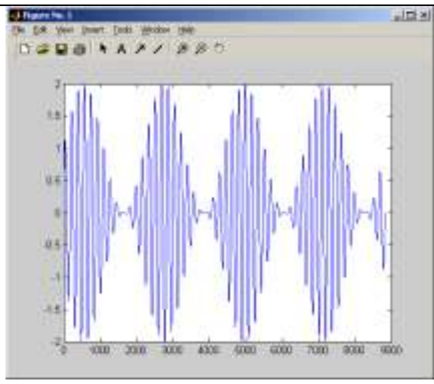

(b) mod20: 250Hz sine wave modulated by 20Hz

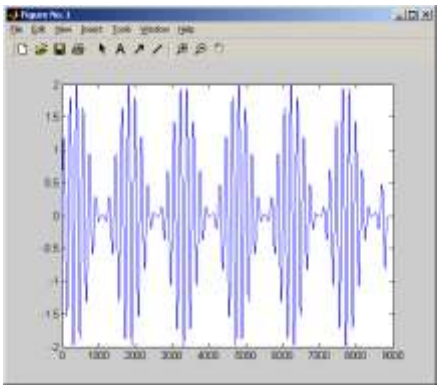

(c) mod30: 250Hz sine wave modulated by 30Hz

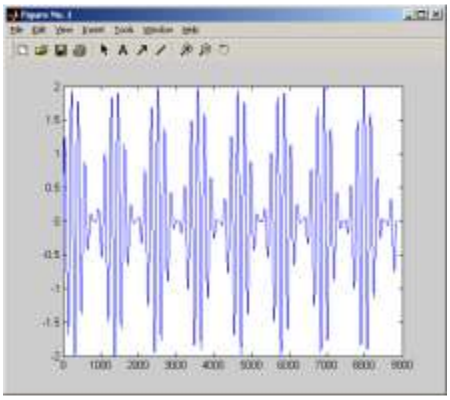

(d) mod40: 250Hz sine wave modulated by 40Hz

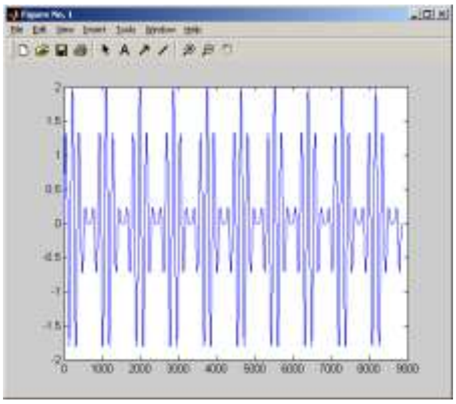

(e) mod50: 250Hz sine wave modulated by 50Hz

Figure 5-3: The five amplitude-modulated stimuli used in the vibrotactile roughness experiment, plotted as amplitude against time.

5.3.1 Hypotheses

Based on the results from the literature the hypotheses were as follows:

1. Participants will be able to distinguish between the stimuli in terms of roughness.

2. Sine wave with no modulation will be perceived as smoother (less rough) than all other stimuli.

3. Perceived roughness will increase as modulation frequency increases.

5.3.2 Hardware Setup

The experiment was run twice, on two different vibrotactile devices – the TACTAID VBW32 (Figure 5-4) and the EAI C2 Tactor (Figure 5-5). These devices were described in Chapter 2. Using both devices allowed a comparison to be made between these two devices, and to see which type of actuator resulted in better perception of vibrotactile roughness.

Figure 5-4: TACTAID VBW32 Actuator

Figure 5-5: EAI C2 Tactor

The actuator was attached to the index finger of the participant's non-dominant hand using double sided sticky tape, and secured using surgical tape (Figure 5-6). The finger-tip was selected as the location for stimulation since it is very sensitive to vibrations and therefore provides a best-case-scenario, and the non-dominant hand was used to leave the dominant hand free to use the mouse to interact with the experiment software. Participants were asked to keep their hand palm-side up as much as possible in order to keep the intensity of the vibrations consistent. They could rest their arm on the arm rest attached to the desk, or on their leg. Participants wore closed-back headphones during the entire experiment to block out any audio leakage from the device. The actuators were controlled via audio files presented via a PC soundcard (with the output boosted by an external amplifier). The vibration output was set to a comfortable level by the experimenter for each device, and was kept constant at that level for all participants using that device.

The C2 Tactor was shown in Chapter 2 with the contactor mounted above the surround and this is the orientation of the device which was intended by the manufacturer. However, when the Tactor was purchased no instructions were supplied and it was unclear which orientation of the device was correct. Pilot testing by the author and her colleagues indicated that sensations on the fingertip felt clearer when the device was used in the other orientation, with the side shown in Figure 5-7 (left) against the fingertip. In this orientation the contactor is mounted below the surround, and starts slightly away from the fingertip, only pressing against the finger when it vibrates. Since this orientation of the device produced clearer sensations, and there was no reason to believe that this was incorrect, the device was used in this orientation for this experiment. In the remainder of this dissertation the two orientations will be labeled as follows. The orientation in which the manufacturers intended the Tactor to be used features the contactor raised above the surround and will, therefore, be labeled as the "raised" orientation (Figure 5 (right)). In the reverse orientation used in the roughness experiment the contactor is indented below the surround, and will, therefore, be referred to as the "indented" orientation (Figure 5 (left)).

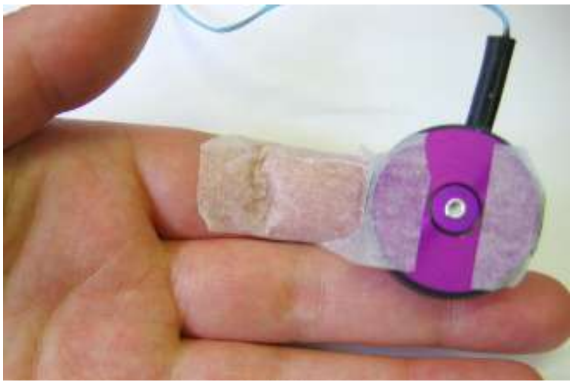

Figure 5-6: Experiment 1: The actuators were attached to the participant's finger using double sided tape and held in position using surgical tape.

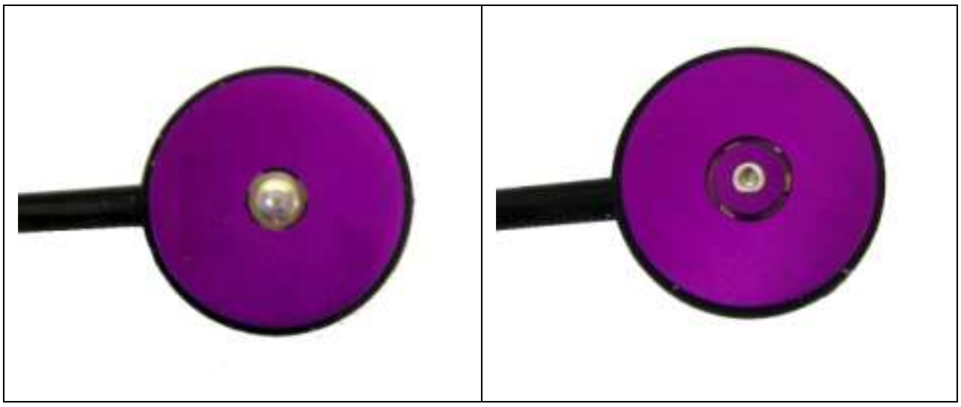

Figure 5-7: The C2 Tactor in the "indented" orientation (left) and in the "raised" orientation (right)

5.3.3 Experimental Method

Participants

Eighteen participants took part in this experiment, with nine different participants carrying out the experiment using each device. All participants were students or staff from the University of Glasgow (5 female, 13 male, all right handed), and were paid £5 for their participation. Participants signed a consent form, which is included in Appendix A.2.

Procedure

The experiment consisted of 50 tasks, and used a forced choice paradigm. In each task participants were asked to compare two stimuli and indicate which stimulus felt "rougher". Within each task participants could choose to feel each of the two stimuli up to four times. Once they had made their decision, they indicated their response by clicking on the corresponding radio button (on the dialogue shown in Figure 5-8). Every possible pairing of stimuli was presented four times, with the order of presentation randomised for each participant.

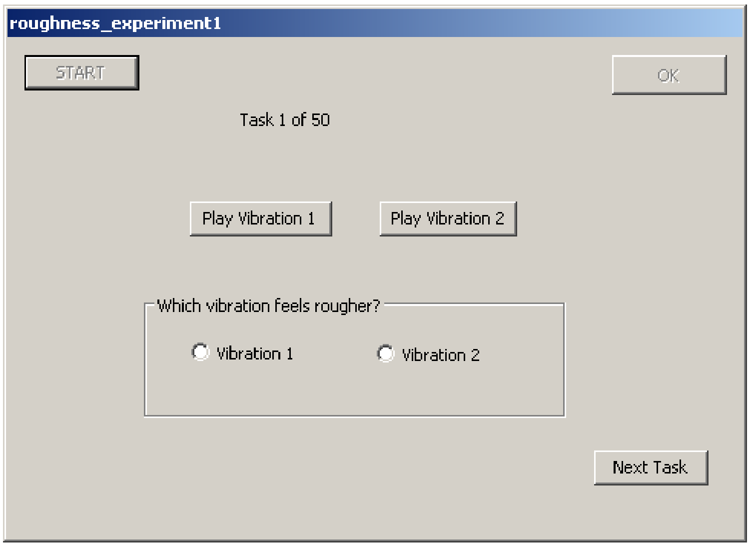

Figure 5-8: Experiment 1: The response dialogue for the Roughness Experiment.

Before starting the experiment, participants were trained to use the response dialogue interface by performing four tasks from the experiment. They received no feedback on their performance during training, and did not receive any training on what was meant by "roughness", as the aim was to find out how people perceive roughness with no training, rather than to pre-condition them to perceive certain stimuli as rougher than others. The instruction sheet that was given to participants for this experiment can be found in Appendix A.3.

It should be noted that this experimental methodology does not provide information on magnitude estimation. The responses made by the users indicate if they perceived one stimulus to be rougher than another but do not indicate how much rougher the stimulus felt. The aim of this experiment was to understand whether people could reliably discriminate these stimuli in terms of roughness, and to place them in an order in terms of roughness, but not to quantify the perceived magnitude differences between them.

5.3.4 Results

During the experiment, data were collected on participants' responses to each pair of stimuli. These data were analysed by considering every pair in turn and looking at how many times each stimulus was considered to be rougher than each other stimulus. This analysis was performed using pair-wise Wilcoxon tests. The raw data from this experiment is included in Appendix G on the accompanying CD. The following sections summarise the main results for both devices.

TACTAID VBW32

The results for the TACTAID showed a significant difference between the un-modulated sine wave and all other stimuli (T=0 (n=9); $p<0.01$), indicating that all other stimuli were perceived to be rougher than the un-modulated sine wave. In addition there was a significant difference between mod40 and mod50 (T=2.5 (n=8); $p<0.05$) with mod40 considered to be rougher than mod50. No other significant differences were shown. After the experiment, four of the nine participants reported confusion with the concept of roughness in terms of how it related to these vibrations. They referred to the "speed", "intensity" and "frequency" of the vibrations, rather than the roughness.

C2 Tactor

As with the TACTAID, the results for the C2 Tactor showed a significant difference between the un-modulated sine wave and all other stimuli. In addition, significant differences were found for all other pairs, except mod20 and mod30. The results can be expressed as followed, with ">" indicating that the stimulus on the left was perceived to be rougher than the stimuli on the right.

- mod20>mod40(T=1.5 (n=8), $p<0.05$)
- mod20>mod50(T=1 (n=9); p=0.01)
- mod20>sine (T=1 (n=8); p=0.01)
- mod30>mod40, mod 50, sine (T=0 (n=9) $p<0.01$)
- mod40>mod50(T=0 (n=8); p=0.01
- mod40>sine (T=0(n=9); $p<0.01$)
- mod50>sine(T=0(n=9); $p<0.01$)

These results indicate that participants felt that roughness increased as modulation frequency decreased, with the exception of the un-modulated sine wave, which felt less rough than all other stimuli. These results are much more consistent than those for the TACTAID. In addition, no participants mentioned a problem with the use of the term "roughness" in relation to the stimuli, unlike when using the TACTAID. However, this is not quantifiable, and may simply be due to the nature of the participants and their willingness to discuss their experience.

5.3.5 Discussion

The results for the TACTAID indicate that, using this device, there are three stimuli which can be distinguished in terms of roughness, namely sine, mod40 and mod50. If using the C2 Tactor, four stimuli can be differentiated in terms of roughness. There are two options for these four stimuli, namely sine, mod20, mod40 and mod50, or sine, mod30, mod40 and mod50. mod20 and mod30 were not distinguishable in terms of roughness, and so should not both be used together.

Hypothesis 1 can be accepted for the C2 Tactor as all but one of the pairs of stimuli were found to be distinguishable. It can also be partially accepted for the TACTAID as several pairs were distinguishable. Hypothesis 2 can be accepted for both devices, as the un-modulated sine wave was always perceived to be smoother (or less rough) than all other stimuli.

Hypothesis 3 can be rejected for both devices. This hypothesis suggested that perceived roughness would increase with increasing modulation frequency. For the C2, the perceived roughness decreased as modulation frequency increased from 20Hz to 50Hz, with the unmodulated sine wave always perceived to be the least rough. This order can be shown as follows with the less than sign (<) indicating that the stimulus on the left is perceived to be less rough than the one on the right. Since mod20 and mod30 could not be distinguished, they are placed together in this order.

Sine<mod50<mod40< mod30/mod20

Hypothesis 3 was based on results from the audio domain, where perceived roughness increases with modulation frequency, however the results from this experiment show that the opposite occurred with these vibration stimuli. This indicates that, while amplitude modulation can be used to create stimuli which sound or feel rough in both modalities, the relationship between amplitude modulation and perceived roughness is quite different.

In the case of the TACTAID actuator, there are insufficient data points to be able to place all of the stimuli on a scale. The only information on order of perceived roughness that can be obtained from the TACTAID results is that the un-modulated sine wave is always perceived to be the least rough and that perceived roughness is higher at mod40 than at mod50. This indicates that perceived roughness decreases with increasing modulation frequency after 40Hz but there is insufficient data to confirm anything about the order of the other stimuli.

The results for the C2 Tactor suggest that, as the modulations became more frequent, the resulting vibration felt smoother. This seems to match with a result observed for the exploration of real, physical, metal gratings explored with a fingertip, where perceived roughness increases with increasing groove width [67, 68]. A vibrotactile equivalent of increased groove width would be an increase in the temporal gaps between fluctuations. Using this analogy the results for the C2 Tactor make sense, as the perceived roughness increased as the amplitude fluctuations got slower, thus leaving bigger gaps between fluctuations (or bigger "grooves" between vibrotactile "ridges", see Figure 5-9). Due to differences in the experimental procedures no conclusions can be drawn about this, but this might warrant further investigation.

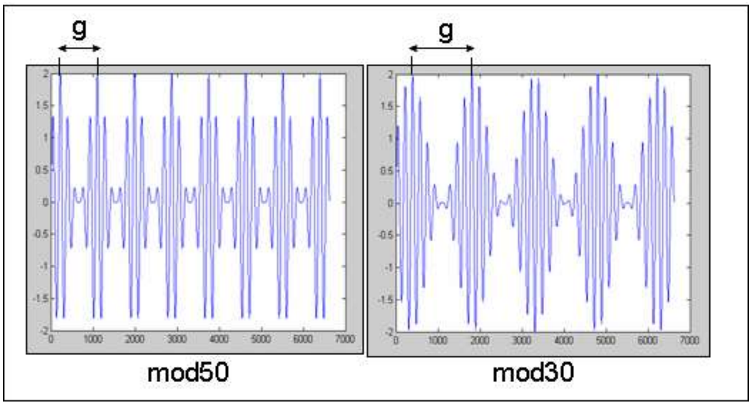

Figure 5-9: Experiment 1: "Groove width" (g) of amplitude modulation increases as modulation frequency decreases. This may, as in perception of real surface textures, be responsible for the increase in perceived roughness

Difference between the devices

The results of this experiment are very different for each device with all but one pair showing a significant difference on the C2 Tactor, and only five of the ten pairs showing significance on the TACTAID (and only one of these which did not include the un-modulated sine wave). With only nine participants in each condition, and a between-groups design, this result could be due to individual differences. However it seems likely that it could be due to the differences in design of the two devices. The most likely reason seems to be the difference in rise times (time taken to reach full amplitude) between the two devices. As reported in Chapter 2, due to the difference in bandwidth between the two devices, the C2 has a rise time of around 5ms (Mortimer, B., personal communication, 7 August 2006), while the TACTAID has a rise time of 12.5ms (Tan, H., personal communication, 25 August, 2006). This difference in rise time means that the C2 will respond more quickly to changes than the TACTAID. Therefore, the TACTAID may have attenuated the intensity and the modulation depth, making the roughnesses less distinguishable. Another possibility is that this difference may be related to the nature of the stimulation provided by the two devices. It is possible that roughness might more distinguishable when a linear vibration occurs against the skin than when the whole device shakes. However, further investigation would be necessary to understand if this was the case.

Limitations of this Study

This study only indicates whether the chosen set of five stimuli can be distinguished in terms of roughness, and allows these stimuli to be placed on a scale in order of perceived roughness. It does not provide any information on the perceived magnitude of roughness of each of these stimuli. This information could have been obtained through using a magnitude estimation study instead of a forced choice paradigm. However, for the purposes of this research it was felt that it was sufficient to identify a set of stimuli that could be differentiated from each other in terms of roughness, and that it was not necessary to understand how much bigger each felt than the other.

The amplitude modulation levels chosen for this experiment were chosen on the basis of informal pilot testing, rather than through a study to choose the optimum levels. A linear relationship between modulation frequency and perceived roughness was assumed, with five stimuli equally spaced 10Hz apart from 20-50Hz. It is possible that more levels could have been attained if different gap spacings had been used, or more levels had been tested.

The results of this experiment only indicate which stimuli can be distinguished from one another; they do not indicate whether each stimulus can be uniquely identified (which will be necessary for encoding information in Tactons). Therefore it will be necessary to test these roughnesses in an absolute identification experiment to understand if unique identification is possible.

Implications for Tacton Design

The results of this study indicate that people are able to distinguish between amplitude modulated sinusoids in terms of roughness. When using the C2 Tactor, it was found that perceived vibrotactile roughness increased as an inverse function of the modulation frequency over a range from 20Hz-50Hz. Over this range, with all stimuli separated by 10Hz, four levels of roughness were able to be distinguished. The only pair that could not be distinguished were mod20 and mod30. When using the TACTAID actuator only three stimuli could be distinguished in terms of roughness (sine, mod40 and mod50), and there were therefore insufficient data points to determine the relationship between modulation frequency and perceived roughness. Since the C2 Tactor allows a greater range of roughnesses to be distinguished, this device should be used in any future experiments that use this parameter.

As reported above, these results indicate that four levels of roughness can be distinguished from one another but do not indicate how many values can be absolutely identified. It was noted earlier that a successful parameter must have multiple levels which can be identified and, therefore, further investigation using an absolute identification paradigm is needed to establish whether roughness is a successful parameter for Tactons.

5.4 Intensity Change Over Time

As identified in Chapter 4, it could be beneficial to draw on the field of audio to inform the design of Tactons. It may also be useful to draw on techniques used in music. It was shown in Chapter 2 that rhythm transfers successfully between the audio and tactile domains, and other techniques may also be transferable. One musical technique which has been mentioned in both Earcon design and tactile interface design is the use of dynamics (changes in volume/intensity). Blattner [4] proposed the use of musical dynamics in Earcons, specifically the use of crescendos (increases in volume) and decrescendos (decreases in volume). Gunther also suggested that these dynamics, and also accents, would be useful in tactile composition [53]. In addition Geldard [41] suggested that intensity changes over time could be a useful way of manipulating intensity within tactile messages, rather than using fixed intensity levels.

While the use of intensity changes over time has been proposed by several researchers there has been little investigation into the effectiveness of this technique for encoding information in vibrotactile messages. The only study of which the author is aware is a study by Howell in his PhD thesis, reported by Geldard [41], where it was shown that there are six distinguishable rise times for intensity increases over time with two levels able to be absolutely identified. Another approach would be to investigate whether people can reliably distinguish between increases in intensity over time, decreases in intensity over time and level stimuli (where the intensity stays constant).

Stimuli Design

Rather than exploring perception of different rates of increase or decrease, this study simply aimed to investigate whether people could reliably identify three different types of intensity change: increases (intensity increases over time), decreases (intensity decreases over time) and level stimuli (no change in intensity over time). If people are able to discriminate between, and uniquely identify, these three types of stimuli, then they could be used to encode information in Tactons.

The stimuli were created by manipulating the amplitude of a 250Hz sine wave over time. All stimuli lasted two seconds, which is at the very upper end of the usable duration range specified by Geldard [41]. These stimuli were created in MATLAB, and saved as audio (.wav) files. The MATLAB code to produce these stimuli is included in Appendix B.1 and the resulting audio files can be found in Appendix H on the accompanying CD by opening the html file “Exp 2 – Intensity Change.html”.

The level stimulus was a 250Hz sine wave in which the amplitude was unchanged over the two second duration (Figure 5-10, note that the figure appears all black as the amplitude is at the maximum level for the full duration). Three different types of increase and decrease (linear, exponential and logarithmic) were created in order to investigate which would be the easiest to identify, and therefore most suitable for use in Tacton design. Figure 5-11 shows the three different types of increase (plotted as intensity against time). The same three types were used for decreases and these are shown in Figure 5-12. It is important to note that the experiment only aimed to discover which of these was most reliably identified as an increase or decrease; users were not asked to distinguish between the different types of increase/decrease.

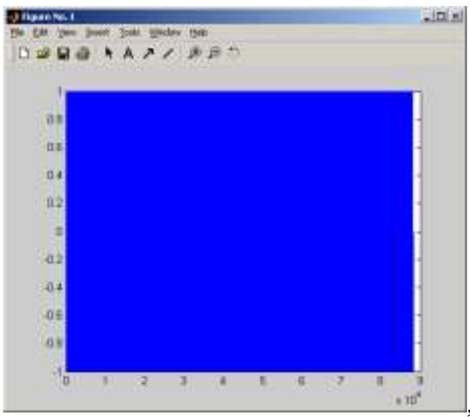

Figure 5-10: The "Level" stimulus.

Pilot testing indicated that it might be difficult to distinguish between increases and level stimuli. This could be due to the fact that the stimulus starts at a low intensity and by the time the user's attention is focused on the stimulus the intensity has leveled off. In music, an increase in volume over time (known as a crescendo [1]) can be made more obvious to listeners by using a technique called sforzando-piano (sfp) before the crescendo (cresc), and this technique could also be applied to tactile crescendos. A sforzando is a strong accent, and piano means soft/quiet [1], so a sforzando-piano has the effect of a sound which starts with a strong attack and high volume (sforzando), and drops quickly to a very low volume (piano). Combining the sfp with a crescendo (to create a sfp-cresc) means that after the sudden drop in volume, the volume then gradually increases over time (crescendo). The high volume accent at the start (sforzando) grabs the listener's attention, making them very aware of the crescendo which follows.

This sfp-cresc effect was applied to the increase stimuli in order to investigate whether this would improve identification of intensity increases (Figure 5-13). Two different intensity levels were used for the sforzando: a high level, which matched the final level of the stimuli (Figure 5-13(a)) and a low level which started at half the level of the final level (Figure 5-13(b)). Both types of sforzando lasted for 250 milliseconds. Pilot testing indicated that the lower intensity sforzando was more subtle, and might improve perception of increases, while the higher intensity one might cause confusion as it felt more like a discreet event of its own.

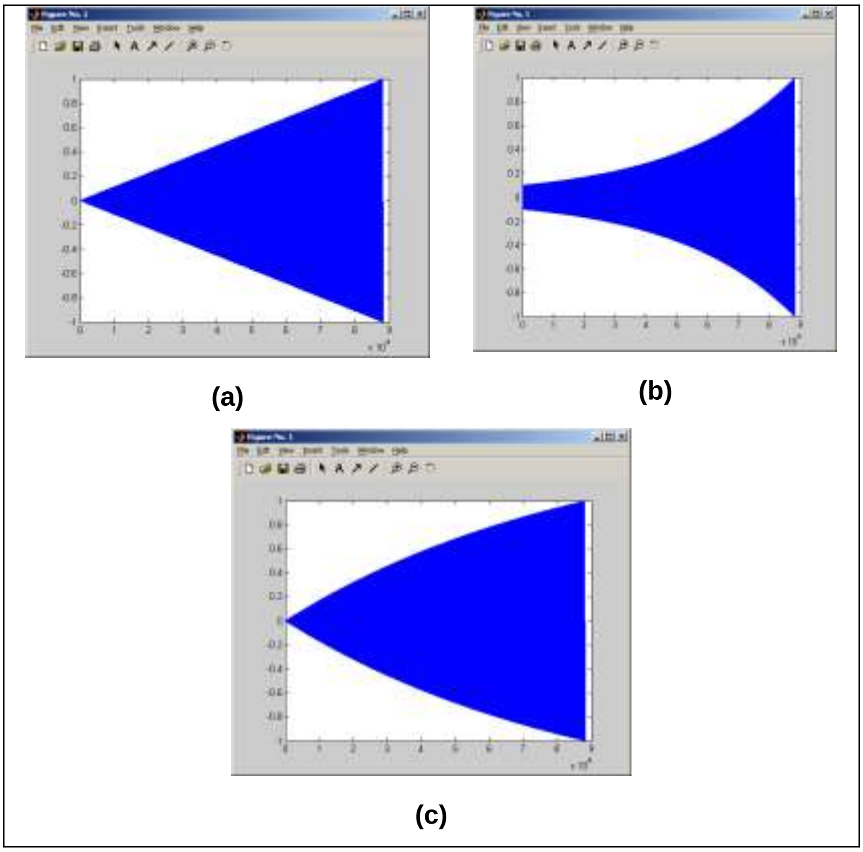

Figure 5-11: The three increase stimuli: (a) Linear, (b) Exponential and (c) Logarithmic (plotted as amplitude against time, with each stimulus lasting 2 seconds).

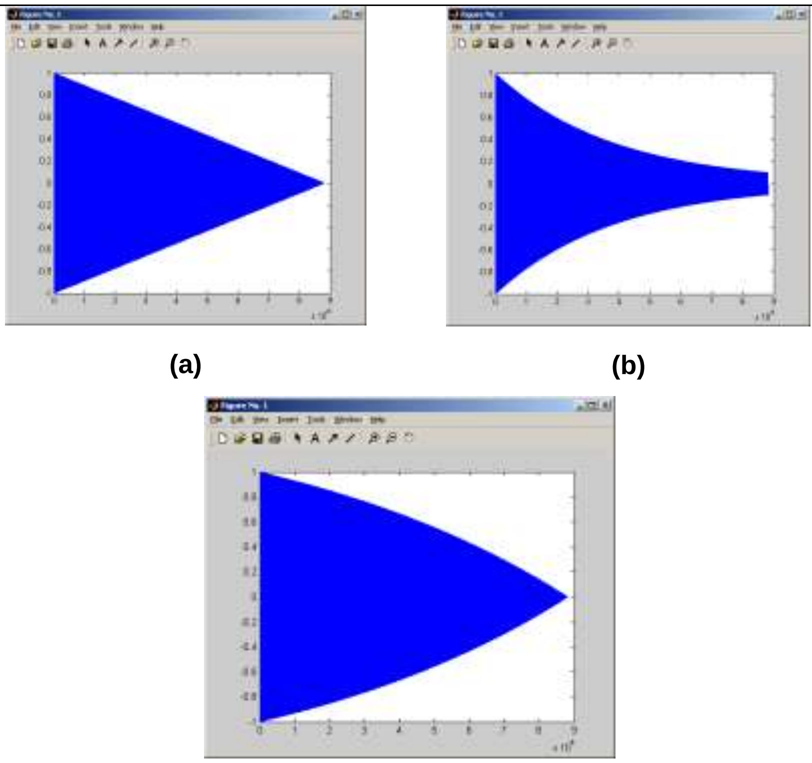

(a) (b)

(c)

Figure 5-12: The three decrease stimuli: (a) Linear, (b) Exponential and (c) Logarithmic, (plotted as amplitude against time, with each stimulus lasting 2 seconds).

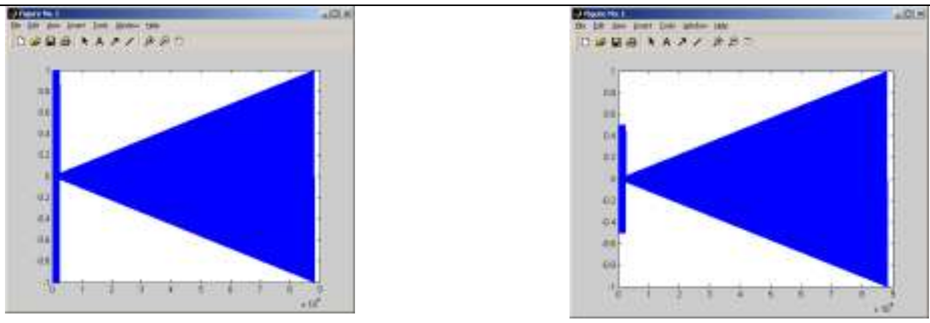

(a) (b)

Figure 5-13: The two types of sforzando applied to a linear increase: (a) high intensity (b) low intensity, (plotted as amplitude against time, with each stimulus lasting 2 seconds). These sforzandos were also applied to exponential and logarithmic increases.

5.5 Experiment 2: An Investigation into Perception of Intensity Change over Time

The aim of this experiment was to discover whether participants could identify and distinguish between intensity increases, decreases and level stimuli, and to identify the most suitable types of increase and decrease for use in Tactons. It also aimed to see whether the use of the musical sforzando-piano-crescendo (sfp-cresc) technique would improve perception of increases.

5.5.1 Hypotheses

The hypotheses were as follows:

1. Participants will be able to distinguish between the three types of tactile dynamics.
2. There will be a difference in the ability to recognise an increase between linear, logarithmic and exponential increases.
3. There will be a difference in the ability to recognise a decrease between linear, logarithmic and exponential decreases.
4. Adding the sfp-cresc effect will improve perception of increases, and avoid confusion with level stimuli.
5. The low intensity sforzando will result in better identification of increases than the higher intensity sforzando as the high intensity sforzando may be perceived as a separate event.

5.5.2 Hardware Setup

The Tactons were presented to users using a TACTAID VBW32 actuator (see Chapter 2) attached to the index finger of their non-dominant hand using sellotape (apart from one participant who had injured his index finger and therefore used the middle finger of his non-dominant hand). The non-dominant hand was used, to leave the dominant hand free to use the mouse. Although Experiment 1 had shown that the C2 is more successful for presenting roughness stimuli than the TACTAID, the TACTAID was used in this experiment since no roughness stimuli were to be presented.

The stimuli, which were saved as audio files, were presented through the TACTAID actuator via a PC soundcard (with the output boosted by an external amplifier). As in the previous experiment, the vibration output was set to a comfortable level by the experimenter and this level was kept unchanged for all participants. Participants wore closed back headphones during the experiment so they did not hear any sounds which might come from the device and would make judgments based on the vibrotactile sensations alone.

5.5.3 Experimental Methodology

Participants

11 right handed participants took part in this experiment. All participants were students or staff of the University of Glasgow. No further demographics were collected for this experiment. All participants were paid £5 for participating.

Procedure

There were 54 tasks in this experiment, consisting of 18 increases (two of each of the nine types), 18 decreases (six of each of the three types), and 18 level stimuli. In each task participants were presented with a two second vibration and asked to identify what happened to the intensity/amplitude of the vibration over time. The term amplitude was used in the questions to users, as this is considered to be interchangeable with the term intensity. The stimulus was presented four times, with a one second pause between presentations, and participants could respond at any time by selecting the corresponding radio button on a dialogue box (Figure 5-14). Participants could choose from four options:

1. Overall Increase: Overall, the amplitude of the vibration increases over time

2. Overall Decrease: Overall, the amplitude of the vibration decreases over time

3. Level: The amplitude of the vibration stays constant over time.

4. Other: The amplitude of the vibration does something other than the three options above.

The "Other" category was included in order to identify any confusion caused by the sfp-cresc. If users selected "other" it would indicate that they had perceived the stimulus as something different, not just as an enhancement of the increase. Qualitative data on users' experiences of the experiment were gathered through an unstructured post-test interview to identify any confusion that may have arisen from the use of the sfp-cresc stimuli.

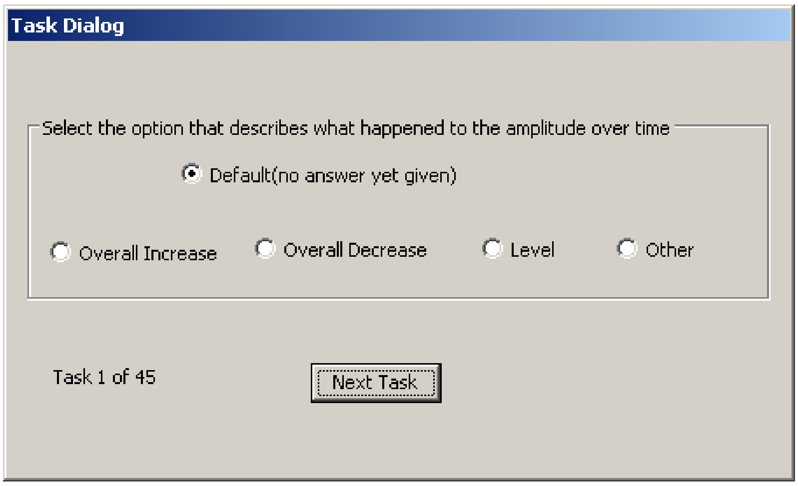

Figure 5-14: Experiment 2: The response dialogue for the Intensity Change experiment

Before starting the experiment participants were trained to use the response dialogue interface by carrying out four tasks from the experiment: identifying a linear increase, linear decrease, linear increase with sforzando-crescendo and a level stimulus. They did not receive any feedback during training or any guidance on what an increase or decrease felt like, as the aim was to identify whether these stimuli were recognisable without pre-conditioning. In particular, they were not trained to recognise the increases containing sforzando-crescendos, as the experiment aimed to identify their instinctive responses to these, rather than to train them to recognise these as increases. The consent form and instruction sheets for this experiment are included in Appendices B2 and B3.

5.5.4 Results

During the experiment, data were collected on participants' responses to each stimulus (see Appendix G on the accompanying CD). From these data, percentage correct scores were calculated for each stimulus, as shown in Figure 5-15. The average recognition rates for all types (standard increase/decrease/level) were over 80%. These rates improve greatly when only the best stimulus from each type is considered, in which case the identification rates are 100% for increases (linear/exponential), 92% for decreases (linear/exponential) and 95% for level stimuli. These results can be seen in Figure 5-15. In this figure, "inc" represents increases, while "dec" represents decreases and "level" represents the stimuli with no intensity change over time. The terms "lin", "exp" and "log" refer to linear, exponential and logarithmic, and the the terms "high" and "low" refer to the intensity level of the sforzando.

Before analysing these data, an Anderson-Darling normality test was run to check if the data were normally distributed. This test indicated that the data did not follow a normal distribution and, therefore, non-parametric analysis was carried out.

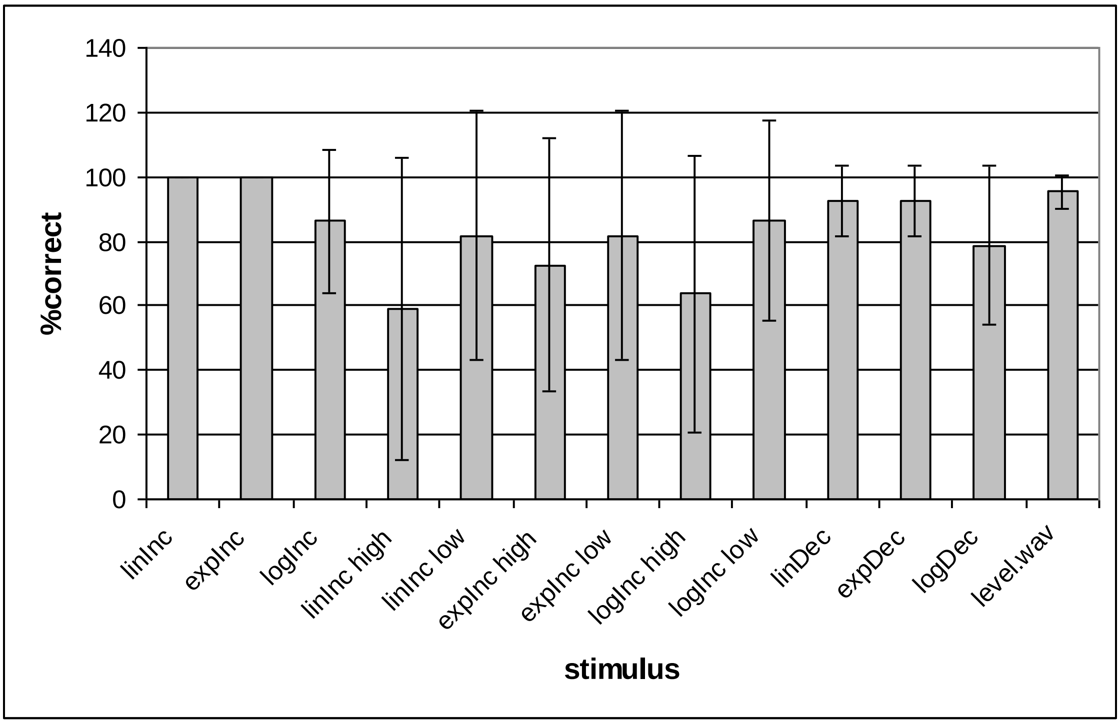

Figure 5-15: Experiment 2: Mean percentage correct scores (with standard deviations) for intensity changes over time.

A Friedman test for within-groups comparisons [56] showed a significant difference overall for the increase stimuli (S = 20.31, DF = 8, p = 0.009). The p-value reported here is adjusted for tied rankings since many of the data had the same percentage correct scores. Post-hoc Wilcoxon tests [56] (with a Bonferroni correction applied to compensate for the fact that multiple pair-wise tests were carried out on the same data) revealed no significant differences between any of the pairs. This lack of significant result may be due to the large number of ties in the data. A Friedman test analysing the results for decrease stimuli also showed no significant results (S = 5.43, DF = 2, p = 0.066).

For the increase stimuli, linear (linInc) and exponential (expInc) increases were correctly identified 100% of the time, while logarithmic (logInc) increases were correctly identified less often: only 85% of the time. Linear (linDec) and exponential (expDec) decreases were identified correctly 92% of the time, while the identification rate for logarithmic (logDec) decreases was 78%. As stated above, none of these differences were found to be statistically significant.

Although the Wilcoxon tests revealed no significant differences between performance on the standard increases and performance on the increases which featured sforzando-pianos, the addition of the sforzando-crescendo seemed to cause confusion, as shown by the high standard deviations for these stimuli (see Figure 5-15). This is backed up by qualitative data which showed that participants were confused by these stimuli. Four subjects thought that the initial pulse was caused by a problem with the transducer or amplifier rather than an intended part of the stimulus. One participant reported that he perceived these stimuli more like parabolas, because they started high, then went low, and then high again. Two participants reported that they knew that these stimuli were increases, but because they were different from standard increases, they selected the "other" option.

5.5.5 Discussion

This study investigated perception of intensity changes over time, in particular the ability of participants to identify three different types of stimuli: increases in intensity, decreases in intensity, and stimuli with no intensity change. The results, with recognition rates of 78%-100% for the three standard stimuli (increase/decrease/level) indicate that these intensity changes can be identified and distinguished from each other, and that intensity change over time could, therefore, be used as a parameter in Tacton design. These results mean that Hypothesis 1 can be accepted.

Three different types of intensity increase and decrease were tested. The linear and exponential increases reached 100% identification, compared to 85% for the logarithmic increase. Similarly, linear and exponential decreases had higher identification rates (92%) than logarithmic decreases (78%). Hypotheses 2 and 3 cannot be accepted as no significant differences were found between the different types of increase stimuli or decrease stimuli. While these differences are not statistically significant it would seem appropriate to use either linear or exponential increases and decreases when designing Tactons since they achieved recognition rates of 92-100%, compared to 78-85% for logarithmic increases and decreases.

Hypothesis 4 suggested that the addition of the sforzando-piano would improve perception of increases. This cannot be accepted as there was no improvement in performance when this technique was used. Although there were no significant differences between performance on these increases and the standard increases, the addition of the sforzando-crescendo seemed to cause confusion rather than improve performance. As there were no significant differences between high and low intensity sforzandos it is not possible to accept Hypothesis 5.

Limitations of this study

The addition of a high intensity accent (sforzando) was intended to improve perception of increase stimuli, however it caused confusion. The decision was taken not to train people to recognise these stimuli but instead to investigate how people perceived these with no pre-conditioning. The results would suggest that sforzandos should not be used in Tacton design and, in fact, with levels of 100% achieved for increases it does not seem necessary to try to enhance them. However, if users were trained to recognise these stimuli performance may improve since this confusion may be avoided. In addition, this attention grabbing technique may be more useful in an interface where the user's full attention is not devoted to the tactile stimuli (e.g. in mobile phone alerts).

The stimuli used in this study were all two seconds long. This is at the upper limit of the range of durations that should be used in vibrotactile display, specified by Geldard [41] (see Chapter 2). If intensity change stimuli were to be used in Tactons, shorter durations might be required to ensure that information transfer was not slowed down too much. Perception may be different for stimuli of shorter durations and therefore if intensity change were to be used, investigation of perception at different durations might be necessary.

Three types of intensity change were investigated in this study, but a wider range may be available. For example, users may be able to tell the difference between e.g. a vibration that increases to a peak intensity and then decreases again, and a vibration with the opposite shape, namely decreasing and then increasing.

This experiment was conducted using the TACTAID VBW32 actuator. It was shown in Experiment 1 that the C2 Tactor enabled better performance for roughness perception. Using the C2 Tactor for this experiment might also have improved performance.

Implications for Tacton Design

The results of this study indicate that people are able to identify intensity increases, intensity decreases and level stimuli. Whereas the roughness experiment only investigated whether people could distinguish between stimuli, this experiment investigated absolute identification of these stimuli. It can therefore be concluded from this study that intensity changes over time could be used effectively in Tacton design, with over 90% recognition possible for these three different types of intensity change.

In order for intensity change over time to be used in Tacton design it will be necessary to combine these stimuli with other vibrotactile parameters such as roughness, rhythm and spatial location. Combining intensity change with these other parameters could change perception of these stimuli, and could also affect perception of the other parameters, therefore if these stimuli were to be used further experiments would be necessary to investigate perception when combined with these other parameters.

5.6 Discussion of Both Experiments

When designing vibrotactile messages such as Tactons it is necessary, first, to identify parameters which can be manipulated to encode information. This chapter reported two studies investigating perception of two vibrotactile parameters which had previously received little attention, namely waveform and intensity change over time. A set of waveforms was created by using amplitude modulation to generate vibrations with varying degrees of roughness, in contrast to the smooth sensations generated by sine waves. It was found that these waveforms could be distinguished in terms of roughness, and therefore it may be possible to use these waveforms to encode information in Tactons. However, further investigation is needed to understand if they can be absolutely identified.

The second study investigated absolute identification of vibrations that increase in intensity, decrease in intensity, or in which the intensity remains constant. It was shown that in an absolute identification experiment, people were able to identify these three types of vibration with 92-100% accuracy. Intensity change over time could, therefore, be used as a parameter in Tactons.

A number of recommendations regarding the creation and use of vibrotactile roughness and intensity change over time stimuli can be extracted from the two studies reported in this chapter. These are reported in Section 5.6.1. Section 5.7 draws conclusions from these studies, in reference to the research questions set at the start of the thesis.

5.6.1 Design Recommendations

Design Recommendations for using Vibrotactile Roughness

Based on the results of Experiment 1, a number of design recommendations for using vibrotactile roughness can be proposed. These recommendations are based on the results of this study, and should be read alongside the results of the study, rather that used in isolation.

1) ***To create a smooth sensation, use an un-modulated s250Hz sine wave.***
   An un-modulated 250Hz sine wave was always perceived to be smooth and was never confused with the “rough” waveforms.

2) ***To create the sensation of vibrotactile roughness, use amplitude modulated sine waves, with modulation frequencies in the range of 20Hz-50Hz. To increase perceived roughness within this range, decrease the modulation frequency.***
   People are able to distinguish between amplitude modulated sinusoids in terms of roughness, indicating that they are happy to label the stimuli in this way. Perceived vibrotactile roughness seems to be dependent on the modulation frequency. For the C2 Tactor, a clear pattern emerged in the results indicating that, apart from the unmodulated sinusoid which always feels smooth, the perceived roughness increased as modulation frequency decreased over a range from 20Hz-50Hz. On the TACTAID actuator, this order was less obvious as less significant differences were found between different levels of roughness.

3) ***If vibrotactile roughness is to be used, use a device with a short response time (e.g. the C2 Tactor responds in 5ms, compared to 12.5ms for the TACTAID).***

   On all the factors above the C2 Tactor obtained better results for roughness perception than the TACTAID. A greater number of levels could be distinguished, a clear order could be established, and participants did not report confusion about the relationship between the vibrations and the concept of roughness. This result is probably due to the faster response time offered by the C2 Tactor.

Design Recommendations for using Intensity Change Over Time

Experiment 2 explored identification of intensity change over time and the following design recommendations can be made based on the results.

1) ***If three types of intensity change are required, use level stimuli (no change in intensity), increases and decreases.***

Users can achieve percentage correct scores of 92-100% when asked to identify whether the intensity of a signal increases, decreases in intensity or does not change (for signals of 2 seconds duration).

2) ***When creating stimuli that increase or decrease in intensity over time, use linear or exponential, rather than logarithmic, increases/decreases.***

In tasks where a user has to identify a signal as either increasing in intensity, decreasing in intensity or level (no change in intensity), linear and exponential increases are correctly identified as increases 100% of the time, compared to 85% for logarithmic increases. In tasks like those described above, linear and exponential decreases are correctly identified as decreases 92% of the time, compared to 78% for logarithmic decreases.

3) ***Do not add a pulse at the start of increases to grab the user's attention.***

Adding a pulse to grab users attention before an increase (based on the sforzando-piano-crescendo technique used in music) did not improve perception of increases, and appeared to cause confusion. With identification rates of 100% for increases, this technique does not seem necessary. However, further training may avoid this confusion, and it may be more useful in a situation where the users full attention is not focused on the stimuli (e.g. in a mobile setting).

5.7 Conclusions

Research Question 1 asked:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons effectively?

This question can be partly answered by some of the results in this chapter. Chapter 4 identified two main conditions which should be met by a parameter for Tactons, and considering these parameters in terms of how they meet these conditions gives an indication of whether they will be successful parameters for encoding information in Tactons:

1. The user must be able to distinguish and identify multiple levels of the parameter

2. The user must be able to label the values of each parameter.

The results from the vibrotactile roughness experiment indicate that up to four levels of roughness can be distinguished from one another but they do not indicate whether these values can also be absolutely identified. Therefore, further study of this parameter is required to understand whether it fully satisfies condition 1. Condition 2 seems to be satisfied by vibrotactile roughness, as users were able to distinguish different levels of amplitude modulation in terms of roughness and should therefore be able to label these in terms of how rough they are. For example, it may be possible to label three values as “smooth”, “rough” and “very rough”. Further study using an absolute identification paradigm, where users are asked to label these parameters in this way, will allow confirmation of this. The results from the roughness perception experiment also indicate that, when presenting vibrotactile roughness, the C2 Tactor should be used, rather than the TACTAID actuator. Perception of roughness was better with the C2 Tactor, probably because of its faster response time. Therefore, the C2 Tactor is used in place of the TACTAID for all further experiments using roughness in this thesis.

The results of the intensity change over time experiment indicate that intensity increases, intensity decreases and level stimuli can be distinguished from one another and can be absolutely identified by users. Therefore, this parameter satisfies condition 1. In order to absolutely identify these stimuli users had to name them as an increase, a decrease or a level stimulus. The results show that people can label these stimuli and, therefore, condition 2 is also satisfied.

The results of the studies in this chapter indicate that both waveform (in the form of vibrotactile roughness) and intensity change over time could be manipulated to encode information in Tactons. Further investigation is required in order to understand whether different levels of vibrotactile roughness can be identified, and further experiments are required to understand how these parameters are perceived when combined with other parameters in Tactons. Therefore, Chapter 6 reports an experiment combining two parameters to encode information in Tactons.

Chapter 6: Design and Evaluation of Two-Dimensional Tactons

6.1 Introduction

Although Tactons can be created by encoding information in a single parameter (one-element Tactons, see Chapter 4) more information can be encoded when several parameters are combined. Therefore, after identifying suitable parameters, the next step in designing Tactons is to combine two or more of these parameters to encode multi-dimensional information. Chapters 2-4 identified a number of parameters (rhythm, spatial location and intensity) which might be suitable parameters for Tactons, and Chapter 5 investigated two further parameters: vibrotactile roughness and intensity change over time. This chapter considers the selection of the two most appropriate parameters and the combination of these parameters to encode two dimensions of information in Tactons. It was identified in Chapter 4 that there are several different types of Tactons; the work in this thesis focuses only on Transformational Tactons.

Section 6.2 discusses the design of these two-dimensional Tactons, and Section 6.3 reports a study evaluating these Tactons. The results of this experiment give a best case scenario for these Tactons as the experiment is conducted using the EAI C2 Tactor, which is one of the best commercially available wearable vibrotactile actuators currently on the market. This experiment also establishes a baseline performance level for Tactons to which any future experiments can be compared. In addition to this experiment a second experiment is reported in Section 6.4 evaluating the same Tactons but with the C2 Tactor used in a different orientation, to compare performance in both orientations.

The experiments reported in this chapter address both of the research questions of this thesis. These research questions were stated in Chapter 1 as follows:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

The experiments in this chapter test absolute identification rates for the individual parameters identified in the previous chapters, when combined with other parameters to create Tactons. This helps to answer Research Question 1. Research Question 2 is also addressed through these experiments as they evaluate Tactons combining pairs of parameters, and therefore provide identification rates for these Tactons.

This chapter concludes by considering how the findings of these experiments answer these two research questions. In addition it draws together a set of recommendations for the design of two-dimensional Tactons based on the results of the experiments reported in this chapter.

6.2 Design of Two-Dimensional Tactons

The following sections report the design of two-dimensional Tactons.

6.2.1 Choice of Vibrotactile Device

The results from the roughness experiment in Chapter 5 showed that the choice of device can have a significant impact on the perception of stimuli and, therefore, the choice of device is an important factor in the design of Tactons. Since the C2 Tactor achieved better performance for roughness perception than the TACTAID actuator, the C2 Tactor was chosen for this study, and was used in the same orientation as that used in the roughness experiment.

6.2.2 Stimuli Design

The aim in creating stimuli for this experiment was to create Transformational Tactons encoding two dimensions of information. As discussed in Chapter 4, Transformational Tactons represent several attributes at once, with each encoded in a different tactile parameter. Chapter 4 also stated that, where possible, three levels should be encoded in each parameter. Therefore, the aim here was to select two parameters and three levels of each of these parameters.

Selection of Parameters

Chapters 2-4 identified three possible parameters: rhythm, spatial location and intensity, and Chapter 5 explored two other parameters in detail: intensity change over time, and vibrotactile roughness. To create Tactons encoding two dimensions of information it was necessary to select two of these parameters.

The parameters chosen to be used for these Tactons were rhythm and roughness. These two parameters also seem promising in terms of tactile perception, with four levels of roughness able to be distinguished (See Chapter 3), and with rhythm having been used effectively by other researchers [19, 65, 113].

While spatial location is perhaps the most promising of the vibrotactile parameters, it was decided that spatial location should not be used in these Tactons, because it would be beneficial to understand what could be achieved with a single actuator before moving on to multiple actuator displays. Other parameters that could have been used in these Tactons were intensity and intensity change over time. The results in the literature indicate that intensity may be a usable parameter, with three levels able to be identified. However, it has also been shown that intensity can be problematic due to interactions with other parameters. In addition, the use of volume (the audio equivalent of intensity) to encode information in Earcons has been shown to be a cause of annoyance to users, who prefer the volume to be under their control [11]. Using intensity as a parameter in Tactons could be equally problematic as reducing the intensity could degrade perception of other parameters, or render the signal undetectable, while increasing it too far could cause pain [32]. The other possibility was to use stimuli which changed in intensity over time, like those evaluated in Chapter 5. These intensity change stimuli were extremely successful and could, therefore, have been used in these Tactons. However these stimuli were very long and using them in combination with the rhythm could have slowed down the rate of communication.

These stimuli might be more effective for Compound Tactons (see Chapter 4). For example, an increase followed by one rhythm could have a different meaning to a decrease followed by the same rhythm.

On the basis of this discussion, it was decided that rhythm and roughness would be more appropriate choices than either intensity or intensity change over time at this stage, and that spatial location should be considered later, once a better understanding of what is possible with a single actuator had been achieved. With rhythm and roughness selected as the parameters for these Tactons, three levels of each of these parameters had to be chosen to encode three levels of each dimension of information. The following sections describe the design of these stimuli.

Roughness

The results from the roughness perception experiment (Experiment 1) reported in Chapter 5 showed that the participants were able to distinguish between four different amplitude modulated signals in terms of roughness. However, they did not indicate whether each stimulus could be uniquely identified. Therefore, it was necessary to select the three stimuli which were most likely to be uniquely identifiable. Based on the results of the vibrotactile roughness experiment, three amplitude-modulated signals were chosen, and mapped to a roughness label. These were sine (smooth), mod50 (rough) and mod30 (very rough). The reasons for these choices are described below.

The un-modulated sine wave (sine) was distinguishable from all other stimuli so was selected. For the other two stimuli it seemed logical that using stimuli that were further apart in frequency would increase the chance of absolute recognition, therefore either mod30 and mod50, mod20 and mod40, or mod20 and mod50 would be suitable choices. Since the difference between mod20 and mod40 had a higher probability ($p<0.05$) of being due to chance than the other two pairs ($p<0.01$), this pair was rejected. Another consideration when selecting roughness stimuli is the speed of the fluctuations of the amplitude modulation. The lower the modulation frequency, the slower the fluctuations will be, and the longer it will take for people to perceive the roughness. Since Tactons should be as short as possible in order to communicate information quickly, it is important to choose higher modulation frequencies, so that shorter pulses can be used. Therefore, mod30 was selected in preference to mod20, so the two stimuli to be used alongside the “smooth” sine wave were mod30 (very rough) and mod50 (rough).

Rhythm

While the roughness stimuli were based on the results of the previous experiment, there was no previous experiment evaluating suitable rhythms for Tactons, so the rhythms were designed based on guidelines from the literature, as discussed in Chapters 2-4. In order to make each rhythm feel as different as possible a different number of notes (vibration bursts) was used in each rhythm, based on Brewster’s guidelines for Earcons [11]. Three rhythms were created: the *7-note rhythm* made up of seven short vibrations, the 4-note rhythm made up of four longer vibrations, and the 2-note rhythm consisting of one short vibration and one very long vibration. These rhythms are shown in Figure 6-1. These rhythms are presented in Figure 6-1 using standard musical notation for rhythm, on a single line since no pitch information is required.

In addition to following Brewster's guidelines, these rhythms also follow advice given by van Erp and Spapé [113] who identified tempo (speed) as an important parameter in the identification of tactile melodies. Although all three rhythms are created using the same tempo, they feel faster or slower due to the use of many short pulses (e.g. 7-note rhythm), or few long pulses (2-note rhythm).

While it might have been possible to use simpler rhythms, there is evidence that people have a good memory for musical attributes such as rhythm [71] and it was felt that people might be able to use this skill to recall tactile rhythms if the rhythms were designed based on musical principles, This is backed up by the work by Buttler and Oravainen ([19], unpublished), who showed that identification of tactile rhythms were best when the rhythms followed musical principles and fitted to an underlying regular beat. Therefore, the three rhythms used here are standard rhythms which might be found in many existing musical works.

7-note rhythm (representing a Voice Call)

4-note rhythm (representing a Text Message)

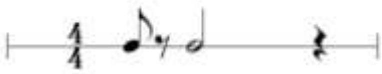

2-note rhythm (representing a Multimedia Message)

Figure 6-1: The three rhythms used in the Two Dimensional Tactons (Experiments 3(a), 3(b) and 4.

The overall durations of these rhythms range from 1.8 seconds for both the 4-note and 2-note rhythms to 2.4 seconds for the 7-note rhythm. The durations of the vibration bursts within these rhythms range from 0.075 seconds for the two shortest notes in the 7-note rhythm, to 1.2 seconds for the longest vibration in the 2-note rhythm. This range more or less complies with Geldard's recommended range from 0.1s to 2s [41] (see Chapter 2). Geldard [41] specified 0.1s as the shortest duration since durations shorter than this might feel like jabs or pokes. However, this recommendation was for single vibrations used in isolation, rather than as part of a rhythm or temporal pattern. Gunther [52] noted that the use of such short vibrations would create an effect much like staccato in music (a short note separated from its neighbouring notes by silence [1]), and this was the effect which was intended to be achieved by the two very short vibrations (0.075s) in the 7-note rhythm. The gaps between the vibrations in these rhythms are all of 75 milliseconds or above. Van Doren's results presented in Chapter 2 indicated that gaps of 100ms or above could be detected even at very low intensity levels, and as these gaps are close to this duration and the intensity is likely to be higher, these gaps should all be detectable.

Although these stimuli are introduced here in terms of the number of notes in each rhythm, they were not introduced to participants in this way. Participants were simply allowed to try the rhythms for themselves, and were free to remember them in any way they wished.

Combining parameters

The Tactons in this experiment were created to represent alerts which might occur when a message or a call arrives on a mobile phone. As the mapping from data to Tactons is abstract, these same Tactons could represent any data, but this particular application domain was chosen here as it is one of the main application areas where vibrotactile feedback might be used. In addition since vibration stimuli are already used in phones this concept would be familiar to users. In a phone alert there are several pieces of information which might be of interest to the user, e.g. the type of alert (new message/incoming call/diary appointment), the priority of the call/message and the caller/sender. As this experiment aimed to investigate Tactons encoding just two parameters, only two of these pieces of information had to be encoded, and the data chosen were the type and priority of the call/message. In a real phone the parameters of these Tactons could encode any data about the phone alerts. In fact these parameters could be used to encode any data for any application; as stated above the choice of application is for experimental purposes only.

To create the two-dimensional Tactons, the two parameters described above were combined, with each parameter encoding a different dimension of information. The type of call/message was encoded by the rhythm. Three types of call/message were represented: voice calls were represented by the 7-note rhythm, text messages by the 4-note rhythm, and multimedia messages by the 2-note rhythm (Figure 6-1). The priority of the alert was encoded in the roughness with the smooth stimulus (unmodulated sinewave) representing low priority, the rough stimulus (sinewave modulated by 50Hz) representing medium priority and the very rough stimulus (sinewave modulated by 30Hz) representing high preiority. Using this mapping, the same rhythm would represent a high priority voice call and a low priority voice call but they would each be presented using a different roughness, whereas a high priority voice call and a high priority text message would share the same roughness, but have different rhythms.

Stimuli Creation

The roughnesses were created in MATLAB (as in the roughness experiment) and saved to audio files. The rhythms were created by importing these files into Cakewalk's music editing software, Sonar (www.cakewalk.com), where the audio files were cut to the length of the notes in each rhythm and arranged temporally to form the rhythm. These rhythms were then saved to audio files which could be presented to users during the experiment.

The audio files for this experiment can be found in Appendix H on the accompanying CD by opening the html file "Exp 3(a) and (b) – 2D Tactons.html".

6.3 Experiment 3(a): Evaluation of Two-Dimensional Tactons

The following sections report the experiment that was carried out to evaluate the Tactons described in Section 6.2.

6.3.1 Aim

The aim of this experiment was to investigate absolute identification of Tactons encoding two dimensions of information. This experiment was the first evaluation of Tactons and would therefore indicate whether Tactons might be successful. Since the experiment had only one condition and no other results existed to which these data could be compared, no hypotheses were set out for this experiment. Instead, the results of this experiment will be used to gauge the level of performance which can be achieved and would also provide a set of results to which future Tacton experiments could be compared.

6.3.2 Hardware Setup

As in all previous experiments, the device was controlled via a PC soundcard (with the output boosted via an external amplifier). The C2 was used in the indented orientation (Figure 6-2), the same orientation as was used in Experiment 1 (roughness perception experiment) and was attached to each participant's index finger with double sided tape and surgical tape in the same way as in Experiment 1. The participant's non dominant hand was used to keep the dominant hand free to use the mouse. During the experiment, participants rested their arm on the desk or on their leg, with the hand facing up so that the pressure between the Tactor and their skin remained constant. Closed-back headphones were worn to block out sounds from the device, to ensure that participants were responding only to the tactile sensations and not to any audio leakage.

Figure 6-2: The C2 Tactor in the "indented" orientation

6.3.3 Experimental Method

Participants

17 participants (staff and students of the University of Glasgow, 4 female, 13 male, all right handed) took part in this experiment and were paid £5 for their participation. No further demographics were collected for this experiment. The consent form signed by participants is included in Appendix C.1.

Training Procedure

As there had been no previous work on Tactons, there was no established methodology for how to train participants to interpret these tactile messages. However, since the design of Tactons is based on that of Earcons, it was possible to learn from the methodologies used for Earcon evaluation. Brewster [12] tested different training methods for Earcons. He found that the most effective approach was for the experimenter to verbally explain the Earcon structure to participants and to present the different Earcons to them during this description, and then to allow the participants to try the Earcons for themselves, with no further experimenter intervention, for up to five minutes. This methodology requires a large amount of experimenter intervention and is unrealistic for real world applications (and for experimental purposes) and Brewster found no significant difference in performance between this method and a training method with less experimenter involvement. In this alternative training method, the participant was provided with a written description of the Earcons, which was not accompanied by any demonstration of the Earcons, and was then given 5 minutes to try the Earcons for themselves. As this "low-cost" training achieved comparable results to those achieved for the "high-cost" training for Earcons, this methodology was used as a basis for the Tactons experiments.

Figure 6-3: Experiment 3(a):The Web Page used for Training.

The participants were given the description of the Tactons used in this experiment, which explained the way in which the data were encoded in the Tactons (see Appendix C.2). After reading this description participants were shown a Web page (Figure 6-3) containing all of the Tactons used in the experiment and were given ten minutes to learn these Tactons through self guided training. Participants were free to end this self-guided training before the ten minutes were up if they felt they had learned all of the Tactons, and most chose to end this training phase after around five minutes.

After the self-guided training phase, participants received instructions on the experimental procedure (see Appendix C.3) and then took part in 18 tasks from the experiment itself as training in how to use the response dialogue. This extended training aimed to reduce the learning effects during the experiment itself. Pilot participants (who did not receive this training) had reported that they became more confident in their responses after around 18 tasks; before this point they still felt they were learning. No feedback was provided during this training, as the aim was to understand how people performed with very little experimenter involvement.

Testing Procedure

The experiment itself consisted of 54 tasks, consisting of six presentations of every Tacton, with the order randomised for every participant. In each task participants were presented with one Tacton, which was repeated four times with a one second pause between repetitions. While the Tacton was being presented, participants had to identify both of the attributes (the type and priority of the call/message) encoded in the Tacton, and indicate their responses by clicking on the corresponding radio buttons (Figure 6-4). They could respond as quickly as they wanted: they did not need to wait until the Tacton had been repeated all four times. Once they had made their response they clicked the “Next Task” button. After a four second break (which was included to allow the skin to recover from adaptation) the message box shown in Figure 6-5 appeared. Participants were encouraged to take a longer break between tasks at any point in the experiment to avoid mental or physical fatigue.

6.3.4 Results

During the experiment, data were collected on the number of correct identifications of rhythm and roughness (see Appendix G on the accompanying CD). Percentage correct scores were calculated for each individual dimension (rhythm and roughness) and for the complete Tactons and these three sets of data were analysed individually. To correctly identify a complete Tacton, both of the individual dimensions had to be correctly identified. Two participants were eliminated from the evaluation as they confused the mapping between rhythms and call/message types, and consistently answered using the wrong mapping (e.g. always selected text message instead of multimedia message). Therefore the results are only reported for the remaining 15 participants. The results for these two participants are also shown in Appendix G so the mis-mapping can be seen.

Statistical analyses were performed on the results for overall Tactons and on the results for individual parameters to understand whether any of the Tactons or their parameters was identified significantly less often or significantly more often than any of the others. This would indicate any particular flaws in the design of these Tactons. Before performing statistical analyses the data were tested using Anderson-Darling normality tests to see if they were normally distributed. As the data did not follow a normal distribution, non-parametric analysis was carried out using the Friedman test for within-groups comparisons [56]. Where significant differences were found, *post-hoc* pair wise comparisons were carried out using Wilcoxon tests [56], with a Bonferroni correction applied to compensate for the fact that multiple pair-wise tests were carried out on the same data. Since there were many tied data (where two or more stimuli had the same percentage correct score) the p-values are adjusted for tied rankings.

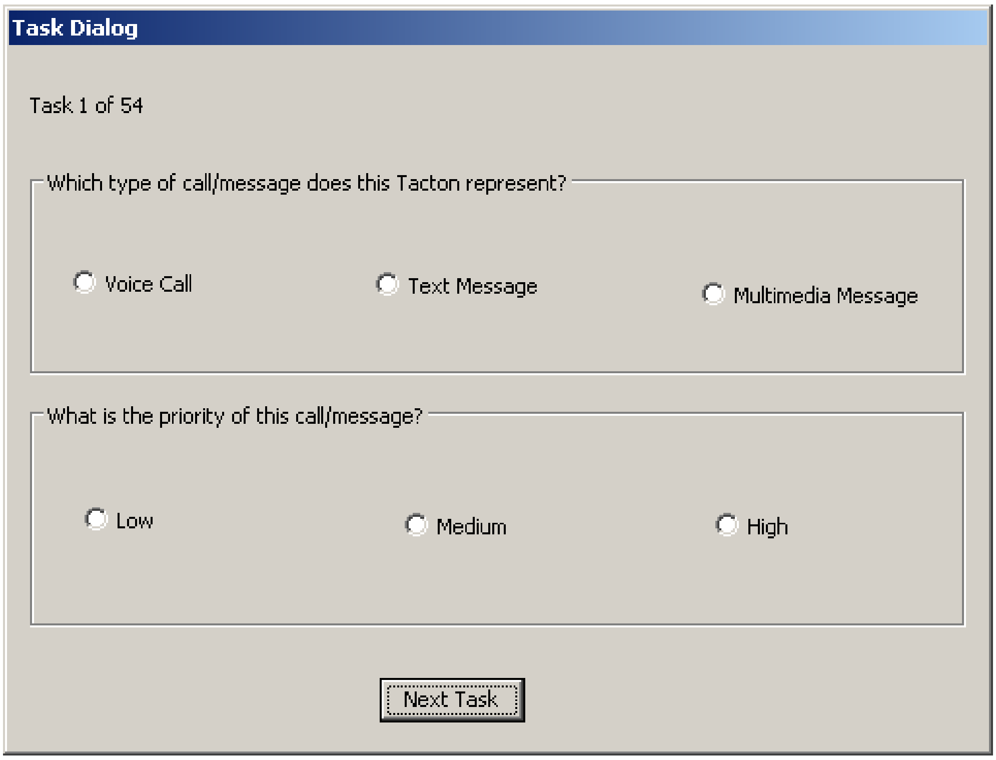

Figure 6-4: Experiment 3(a): Response dialogue for the Two-Dimensional Tactons experiment.

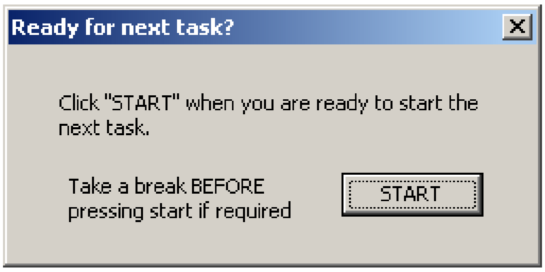

Figure 6-5: Experiment 3(a): The message box that appeared between tasks in the Two-Dimensional Tactons experiment.

Complete Tactons

The results for overall Tacton recognition showed an average recognition rate of 73%, with very rough 2-note rhythms (high priority multimedia messages) having the highest recognition rate (86%), and rough 7-note rhythms (medium priority voice calls) having the lowest (63%). Figure 6-6 shows the results for each Tacton, with the Tactons labeled in terms of the rhythm and roughness used to represent encode the data. The Friedman test analysing differences in the percentage correct score for each Tacton showed a significant difference overall (S = 15.69, DF = 8, p = 0.047). *Post hoc* Wilcoxon tests were carried out to investigate this difference, and a Bonferroni correction was applied, resulting in a threshold p-value of 0.001. These pair-wise tests revealed no significant differences between individual pairs. This is probably due to the fact that the result of the Friedman test was only just significant, and the Bonferroni correction is very conservative. The histogram in Figure 6-7 shows the distribution of scores achieved for overall Tactons. This is included to show the distribution of scores since they did not follow a normal distribution. The histogram also illustrates that although the average score was just 73%, participants scored 80-100% correct on overall Tacton identification in 60% of the tasks (90-100% in around 40% of tasks, and 80-90% correct in 20% of tasks).

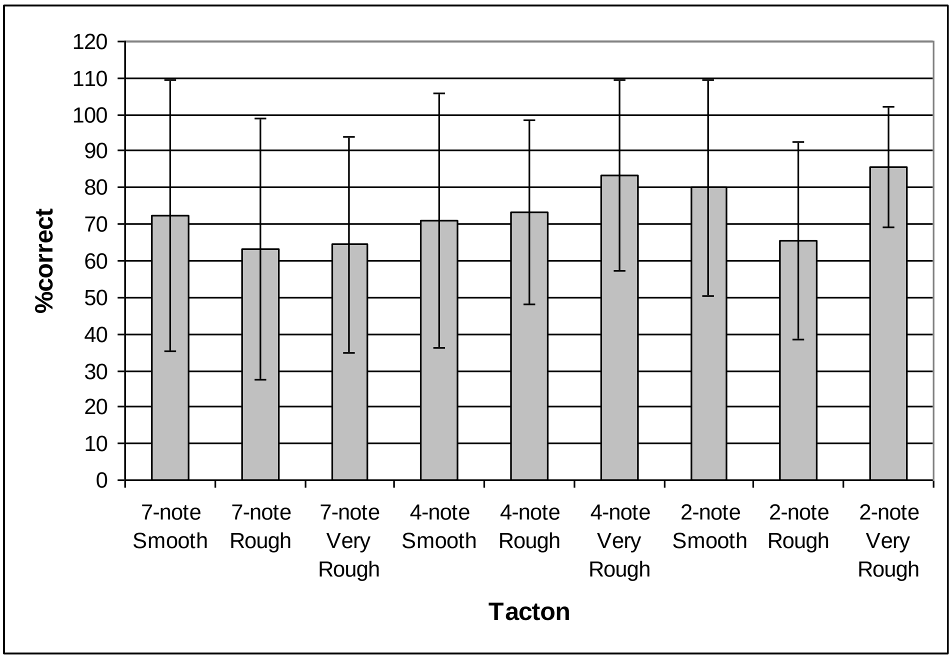

Figure 6-6: Experiment 3(a): Mean percentage correct scores (with standard deviations) for complete Tactons.

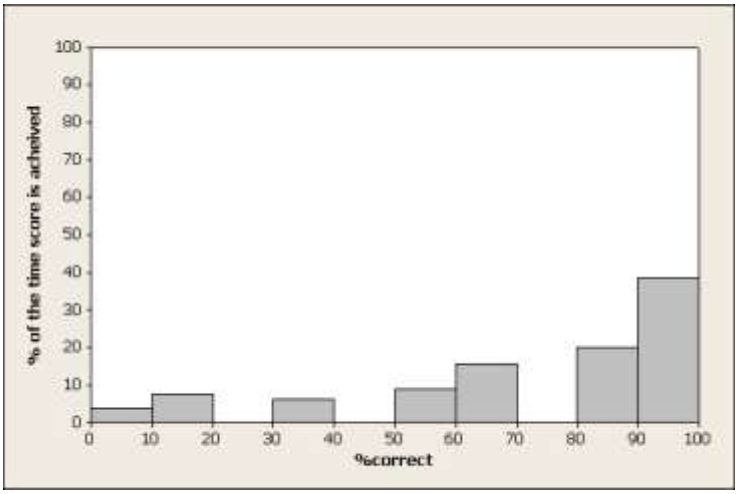

Figure 6-7: Experiment 3(a): Histogram showing distribution of percentage correct scores for overall Tactons

Rhythm

Rhythms (representing call/message types) were correctly recognised on average 96% of the time (Figure 6-8), with the 2-note rhythm (multimedia message) being best recognised (99%) and the 7-note rhythm (voice call) having the lowest recognition rate (91%). No significant differences were found between these three rhythms (S = 3.10, DF = 2, p = 0.212). The histogram in Figure 6-9 shows the distribution of scores achieved by participants for correct rhythm identification. This illustrates that around 90% of time, participants achieved scores of 80-100% for type (rhythm) in 90% of tasks (90-100% in 85% of tasks 80-90% in 5% of tasks).

Informal discussion with participants after the experiment revealed that several felt that they had difficulty distinguishing between the 7-note and 4-note rhythms. Although there is no significant difference, the confusion matrix in Table 6-1 also shows that most confusion occurred between these two rhythms; the 2-note rhythm was very rarely confused with either of the other rhythms.

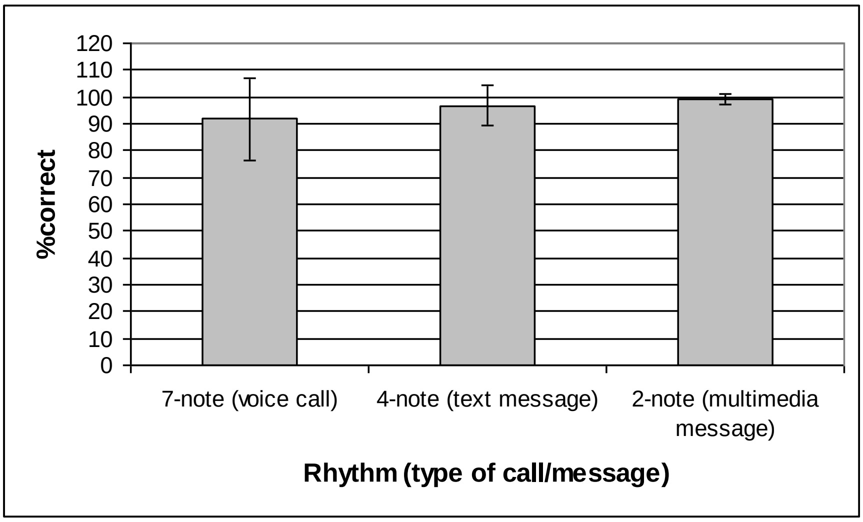

Figure 6-8: Experiment 3(a): Mean percentage correct scores (with standard deviations) for rhythm (representing type of call/message).

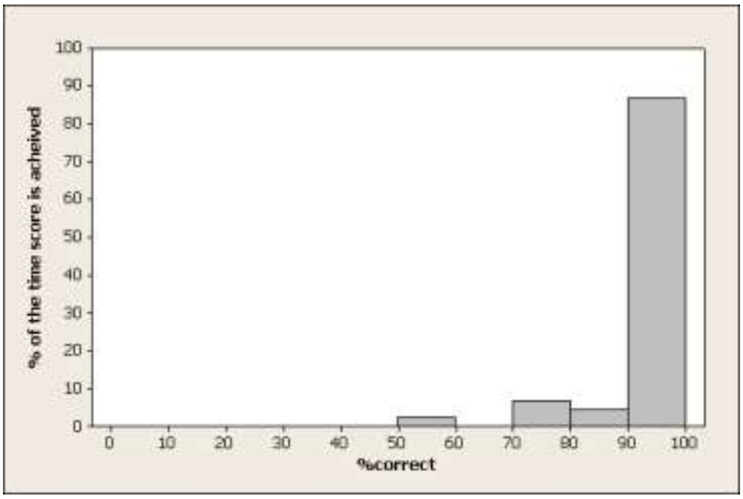

Figure 6-9: Experiment 3(a): Histogram showing distribution of percentage correct scores for rhythm parameter.

		response		
		7-note	**4-note**	**2-note**
stimulus	**7-note**	91.85%	7.41%	0.74%
	4-note	2.59%	96.67%	0.74%
	2-note	0%	0.74%	99.26%

Table 6-1: Experiment 3(a): Stimulus-response confusion matrix for rhythm, showing the percentage of times each response was made, in response to a given stimulus.

Roughness

The results for recognition of roughness (representing priority of call/message) showed an average recognition rate of 79.5% (Figure 6-10), with the smooth stimulus having the highest recognition rate (82%) and the rough stimulus the lowest (76%). A Friedman test showed no significant differences between these three levels of roughness (S = 1.11, DF = 2, p = 0.575). The histogram in Figure 6-11 shows the distribution of scores achieved by participants for correct identification of roughness. This illustrates that although the average score was just 79.5%, participants scored 80-100% on overall Tacton identification in 70% of tasks (90-100% in 40% of tasks, and 80%-100% correct in 30% of tasks).

Informal feedback indicated that many participants felt they often got confused between the rough and very rough stimuli. The confusion matrix in Table 6-2 indicates that participants did experience confusion between adjacent levels of roughness, i.e. between smooth and rough and between rough and very rough, but that there was rarely much confusion between the most widely spaced stimuli (smooth and very rough).

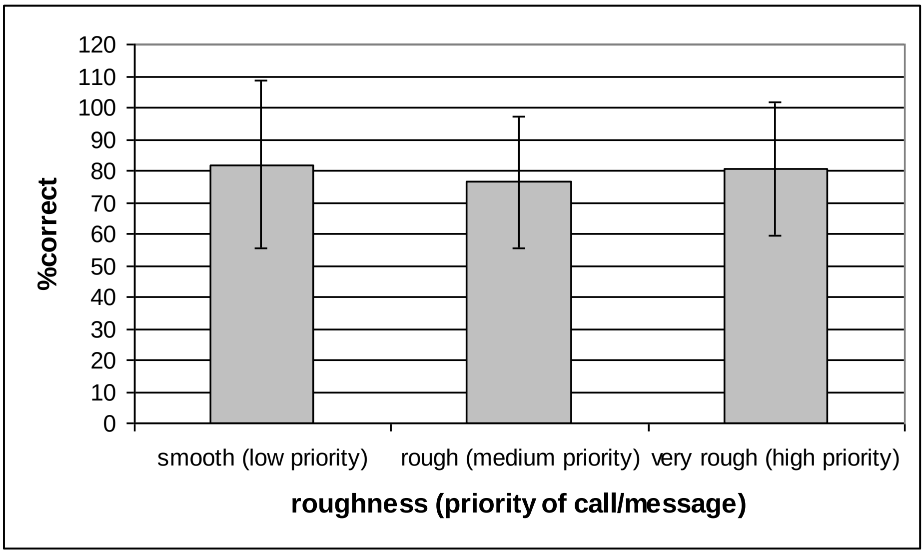

Figure 6-10: Experiment 3(a): Mean percentage correct scores (with standard deviations) for roughness (representing priority of call/message).

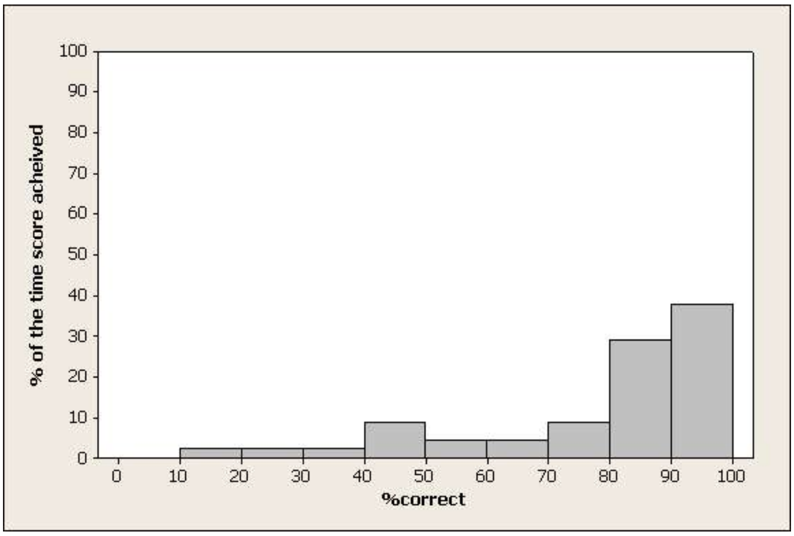

Figure 6-11: Experiment 3(a): Histogram showing distribution of percentage correct scores for roughness parameter.

stimulus \ response	smooth	rough	very rough
smooth	77.41%	15.93%	6.67%
rough	10.37%	71.11%	18.52%
very rough	3.70%	15.93%	80.37%

Table 6-2: Experiment 3(a): Stimulus-response confusion matrix for the roughness parameter.

6.3.5 Discussion

The results of this experiment seem very promising for Tactons. The average result for Tacton identification was 73% with 60% of participants achieving percentage correct scores of over 80%. These results were achieved after less than ten minutes training, which would seem to be a practicable amount of training for many real applications. The average score of 96% for rhythm identification indicates that rhythm is a good choice of parameter for Tactons. Although the average score for roughness identification was lower, at just 79.5%, the results showed that 70% of participants scored between 80-100% correct, indicating that roughness could be a useful parameter for Tactons.

As mentioned in Chapter 5, the C2 Tactor was supplied without instructions as to how it should be oriented and therefore was used in this experiment and the roughness experiment in the orientation that felt best: the "indented" orientation. However, this was not the orientation intended by the manufacturers. After running this experiment, the manufacturers indicated to the author that the device was designed to be used in the "raised" orientation so that the contactor was preloaded on the skin, and that using it in the indented orientation might have significantly reduced the level of the output. However, they did not believe this should have much effect on the fingertip and with the particular stimuli tested. When using the C2 Tactors on other body locations, the preloading of the contactor against the skin is required in order for the vibrations to be strong enough to be perceived, so it would be necessary to use the "raised" orientation for any other location. In order to see whether this indented mounting of the Tactor had affected performance, the experiment was re-run again in the raised orientation.

6.4 Experiment 3(b) Evaluation of Two-Dimensional Tactons with C2 in "Raised" Orientation

The following sections report the methodology results for the two-dimensional Tactons experiment when it was re-run with the C2 Tactor in the "raised" orientation. The stimuli used were exactly the same as in Experiment 3(a).

6.4.1 Hypothesis

The hypothesis for this experiment was that, with the device used in the raised orientation, as intended by the manufacturers, performance should improve, or at least stay the same, as with the device used in the reverse, "indented" orientation.

6.4.2 Experimental Methodology

Participants

17 participants took part in the study and were paid £5 for their participation. All but two of the participants were staff or students of the University of Glasgow (the other two were friends or relatives of other participants), and the ages of the participants ranged from 22-42 (mean=30). None of the participants had taken part in the previous two-dimensional Tactons experiment.

Procedure

The procedure was exactly the same as that described in Section 6.3.3.

Hardware Setup

The hardware setup was the same as in the last experiment except for the orientation of the C2 Tactor, which was now oriented as shown in Figure 6-12

Figure 6-12: C2 Tactor in "raised" orientation.

6.4.3 Results

The same data were collected as before and analysis was carried out in the same way as for the previous experiment (see Appendix G on the accompanying CD for the raw data). In addition, between groups analysis was carried out to compare the results with those achieved in the previous experiment. As in the previous experiment one participant consistently miss-mapped rhythm to type and was therefore removed from the analysis. The results for this participant are shown in Appendix G, along with the raw data for all participants.

Tactons

The results for overall Tacton recognition (Figure 6-13) showed an average recognition rate of 66%, with smooth 7-note rhythms (low priority voice calls) having the highest recognition rate (76%), and very rough 7-note rhythms (high priority voice calls) having the lowest (54%). A significant difference was found in the Friedman test analysing these data ($S = 28.38$, $DF = 8$, $p < 0.001$). Post hoc analysis with Wilcoxon paired tests (with a Bonferroni correction resulting in a p-threshold for significance of 0.001) found no significant differences between any of the individual pairs.

The histogram in Figure 6-14 shows the distribution of scores for overall Tactons and indicates that in around 25% of tasks participants scored between 90%-100% for correct identification of Tactons, and in 45% of tasks they scored between 80-100%. This is much lower than that achieved with the indented orientation.

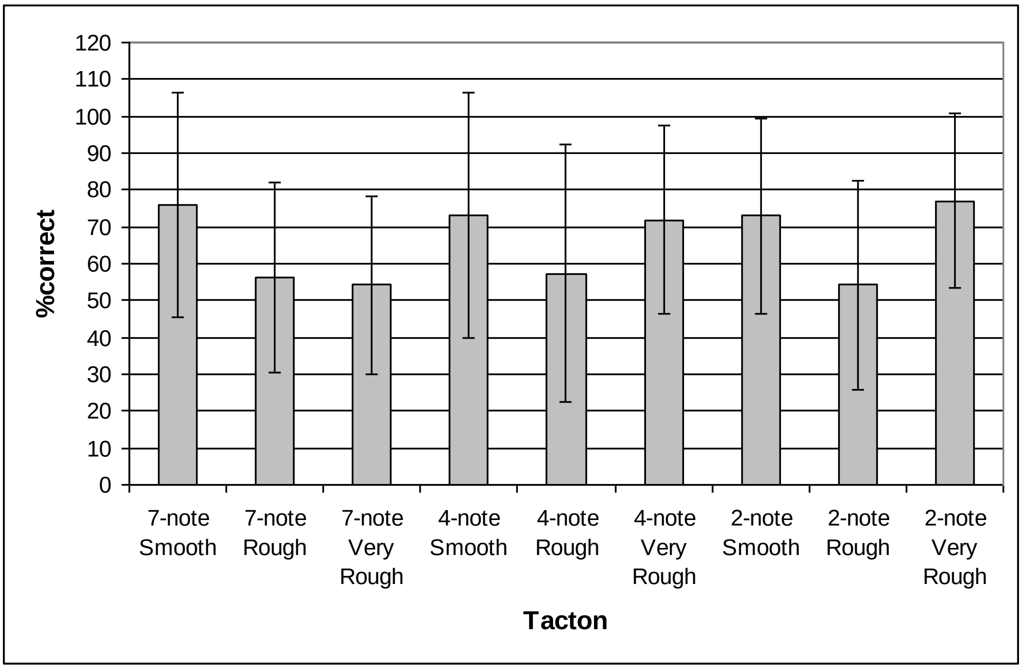

Figure 6-13: Experiment 3(b): Mean percentage correct scores (with standard deviations) for complete Tactons (C2 "raised" orientation).

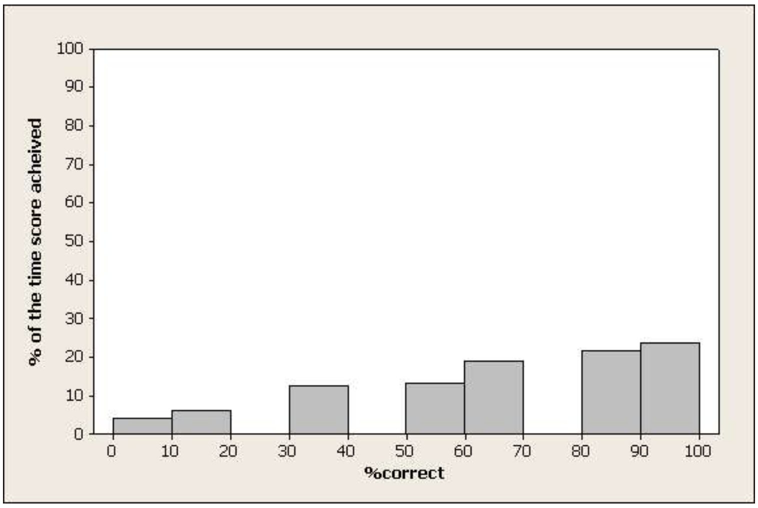

Figure 6-14: Experiment 3(b): Histogram showing distribution of percentage correct scores for overall Tactons.

Rhythm

Rhythms (representing call/message types) were correctly recognised on average 95.5% of the time. As in the previous experiment, 2-note rhythms (multimedia messages) were the best recognised (99%) and 7-note rhythms (voice calls) had the lowest recognition rate (93%). These results are shown in Figure 6-15. No significant differences were found between the results for the three different rhythms (S = 4.06, DF = 2, p = 0.131). The histogram in Figure 6-16 shows that scores of 90-100% were achieved in 90% of tasks for rhythm identification, with scores between 80-100% achieved in 95% of tasks. As would be expected from the results of the statistical analysis, this is a similar distribution to that obtained when using the device in the indented orientation. The confusion matrix in Table 6-3 shows that, as with the indented orientation, most confusion occurs between the 7-note and 4-note rhythms, with the 2-note rhythm rarely confused with the others.

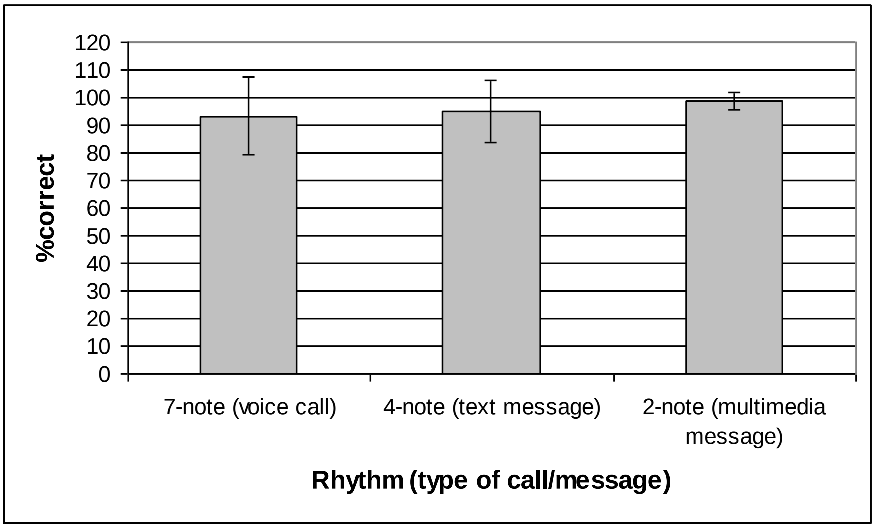

Figure 6-15: Experiment 3(b): Mean percentage correct scores (with standard deviations) for rhythm (C2 "raised" orientation).

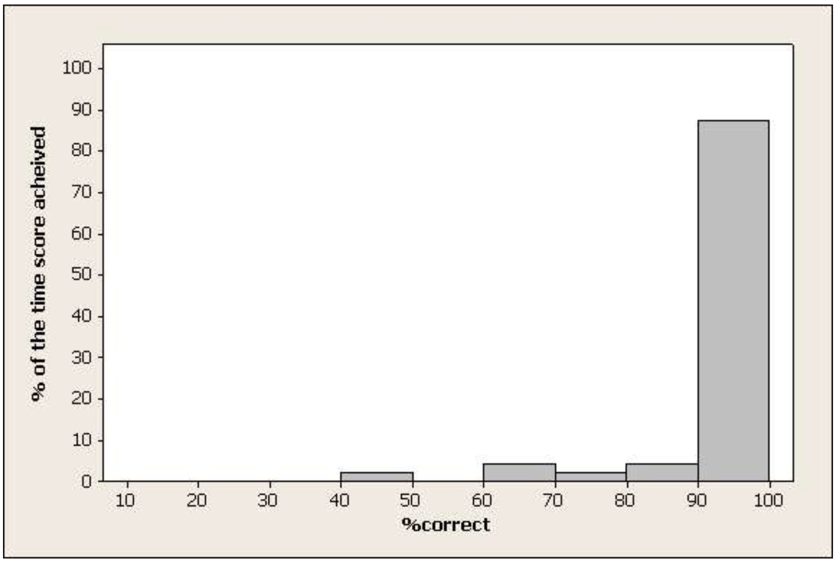

Figure 6-16: Experiment 3(b): Histogram showing distribution of percentage correct scores for rhythm (C2 "raised" orientation).

stimulus \ response	7-note	4-note	2-note
7-note	93.40%	6.60%	0.00%
4-note	5.21%	94.44%	0.35%
2-note	0.35%	1.04%	98.61%

Table 6-3: Experiment 3(b): Stimulus-response confusion matrix for the rhythm parameter.

Roughness

The results for recognition of roughness (representing priority of call/message) showed an average recognition rate of 69%, with the smooth stimulus (representing low priority) having the highest recognition rate (76%) and the rough stimulus (representing medium priority) the lowest (58%). The Friedman test showed a significant difference between these results ($S = 6.41$, $DF = 2$, $p = 0.041$). *Post-hoc* Wilcoxon tests with Bonferroni corrections (resulting in a p-threshold for significance of 0.016), revealed that this difference was between the smooth and rough stimuli ($T=11.5$ ($n=14$), $p=0.011$). There were no significant differences between the other two pairs. The histogram in Figure 6-18 shows that participants obtained scores of 90-100% for roughness identification in 20% of tasks, with scores of 80-100% achieved in 40%. These results are much lower than those achieved with the indented orientation. The confusion matrix in Table 6-4 shows that most of the confusions occurred between the rough and very rough stimuli; the smooth stimulus was less often confused with the other two. This is supported by the statistics which show that performance was significantly worse for the rough stimulus than for the smooth stimulus.

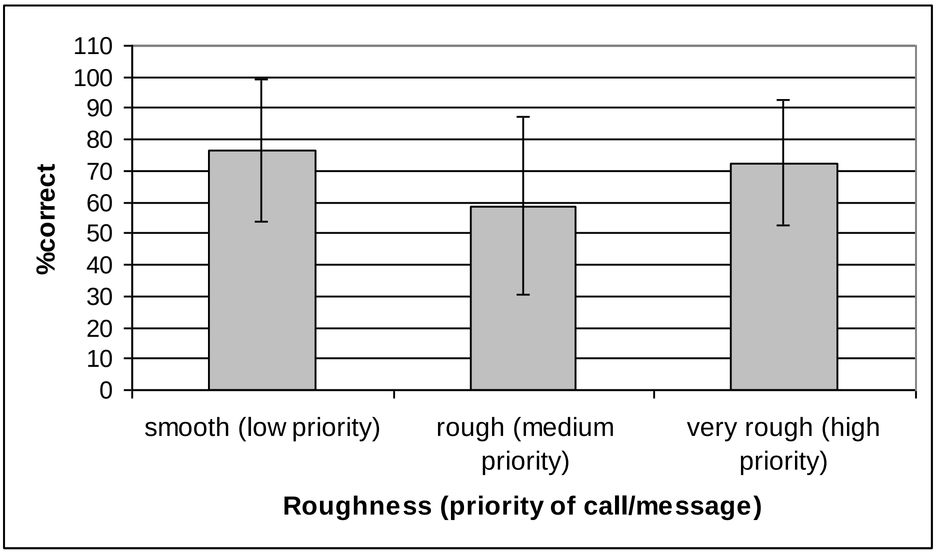

Figure 6-17: Experiment 3(b): Mean percentage correct scores (with standard deviations) for the roughness parameter (C2 "raised" orientation).

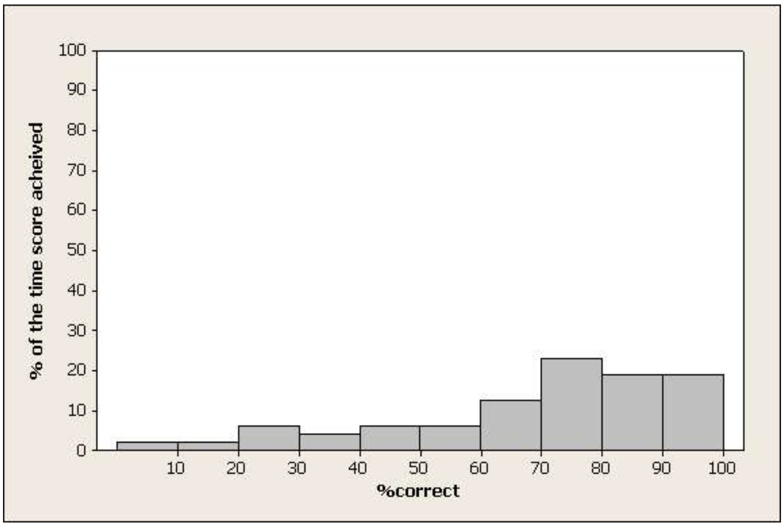

Figure 6-18: Experiment 3(b): Histogram showing distribution of percentage correct scores for the roughness parameter (C2 "raised" orientation).

response

stimulus		smooth	rough	very rough
	smooth	76.39%	19.44%	4.17%
	rough	4.86%	58.68%	36.46%
	very rough	4.86%	22.57%	72.57%

Table 6-4: Experiment 3(b): Stimulus-response confusion matrix for the roughness parameter.

Comparison of results for both device orientations

Statistical analysis was carried out to compare the sets of results for both orientations of the device. As before non-parametric statistics were used because the data were not normally distributed. As this was a between-groups analysis of two conditions, Mann Whitney pair-wise tests were carried out to analyse differences in identification of overall Tactons, and of both the individual parameters. Since multiple comparisons using the same data were not carried out, no Bonferroni correction was applied and the p-value for significance was, therefore, 0.05.

The results of the Mann-Whitney tests showed that the decrease in overall performance when the C2 was used in the raised orientation was significant ($p = 0.01$). In addition there was a significant decrease in priority (roughness) identification when the C2 was used in the raised orientation ($p = 0.019$). There was no significant difference in identification of type (rhythm) between the two orientations ($p=0.523$).

6.4.4 Discussion

The results of this experiment showed an average overall identification rate of 66%, compared to 73% when the device was used in the indented orientation. The identification rate for roughness also decreased, from 79.5% to 69%. There was almost no change in the percentage correct score for rhythm identification which was 95.5% (compared to 96%). These results indicate that, when Tactons are presented to the fingertip, there is a significant decline in performance when the device is used in the "raised" orientation compared to when it is used in the "indented" orientation. As there is no decrease in performance on rhythm, the decline in overall performance appears to be a direct result of a decrease in performance on identification of roughness stimuli.

These results indicate that, when presenting complex waveforms such as these to the fingertip, it is beneficial to use the C2 Tactor in the indented orientation. The hypothesis for this experiment was that performance would stay the same, or improve, when the C2 was used in the raised orientation, since this was the orientation specified by the manufacturers. This hypothesis can be rejected. Discussions with the manufacturers revealed that these actuators are not designed for use on the fingertip and were intended for use on large hairy skin surfaces. On the small surface of the fingertip there may have been insufficient space for the surround to ground itself and drive forces against the fingertip. In the indented orientation the soft skin on the tip of the finger was pressed into the concave opening, and the contactor could press directly against the skin. Another possibility is that the presence of the rigid surround in the indented orientation improved perception of roughness by limiting the propagation of the vibration waves across the skin. Further investigation is necessary to understand if this was the case.

When using the C2 Tactors on body locations other than the fingertip the indented orientation will not work. When used in this orientation the soft skin on the fingertip squashed into the opening enabling the contactor to vibrate against this skin, but the skin on other body sites does not do this, therefore the preloading of the contactor against the skin is required in order for the vibrations to be perceived. Thus, although performance on the fingertip declined with the change in orientation, this "raised" orientation will be used in future experiments on other body locations. It is likely that the decline in performance was due to the fact that the Tactor was not designed for use on the fingertip and, therefore, that these stimuli might be more successful when presented to different body locations.

6.5 Discussion of Both Experiments

The results of the experiments reported in this chapter indicate that two dimensional Tactons can be created by manipulating rhythm and roughness, and that users are able to learn these Tactons with less than ten minutes training and then identify them with above 70% accuracy. Rhythm seems to be very successful as a Tacton parameter, with average identification rates of around 96% achieved in both experiments.

The results also indicate that roughness could be a usable parameter for Tactons, as it achieved a recognition rate of 79.5% in the first experiment, which is comparable to that achieved by register (pitch) in Earcons (around 75%) [77]. Performance on roughness decreased significantly (as did overall identification) when the C2 Tactor was used in the raised orientation rather than in the indented orientation. This decline in performance may be due to the fact that the device was not designed for use on the fingertip. In this case, the excellent results achieved when the device is used in the indented orientation provide a very interesting result: that the C2 Tactor can be used on the fingertip effectively when used in the indented orientation. As discussed in Section 6.4.4, the indented orientation will not work on other body locations and therefore the raised orientation is used for the remaining experiments in this thesis since they do not use the fingertip. However, any further work on the fingertip should continue to use the indented orientation since best performance can be achieved in this orientation.

6.5.1 Comparison to other Non-Visual Icons

As the design of Tactons is based on the design principles used in Earcons, it is beneficial to compare the results for Tactons back to those achieved for Earcons, to understand how these two types of non-visual icons compare. The results from Experiment 3(a) (with C2 in indented orientation) are used for this comparison. The overall recognition rate of 73% compares favourably to results for Earcon recognition, where McGookin’s evaluation of Transformational Earcons [78] showed overall recognition of around 70% (although he used three parameters rather than just two so a direct comparison cannot be made). Looking at the individual parameters, McGookin found a recognition rate of over 90% for melody, which matches the recognition rate for rhythm in Tactons (96%). The roughness parameter in Tactons achieved 79.5% recognition, which is lower than McGookin’s result for timbre (over 90%), but higher than the result for his third parameter, register (75%).

The results can also be compared to those from the only other formal evaluation of vibrotactile icons reported in the literature, namely Chan *et al.*’s work on vibrotactile icons for collaborative applications [22, 23]. In that work, a set of seven vibrotactile icons were created, and people were asked to identify them under various degrees of workload. The results showed identification rates of 95% on average. This level of performance is much higher than that achieved for the Tactons tested in the experiments in this chapter. There are several possible reasons for this high level of accuracy. Firstly, it might result from the use of a metaphorical approach, which may make it easier for people to learn and remember stimuli than an abstract approach (see Chapter 4). In addition, the set of icons used by Chan *et al.* only consisted of seven vibrotactile icons whereas there were nine Tactons in the set used in these experiments. It would be interesting to see how this approach scaled to a larger set of stimuli, as it would require the existence of appropriate metaphors for more components of the system, and would place further memory demands on the user. One other possible explanation for the high performance is that the participants in Chan's study were given the extra motivation that the four people achieving the best scores would receive an extra financial reward, whereas no such bonus was offered in the Tactons experiments. If it was this motivation which improved performance, this is interesting as it indicates that users can achieve higher levels of performance when they have a motivation to do so, and, therefore, that these types of recognition rates may be achievable in real world applications where the users may have an intrinsic motivation to understand the messages.

6.5.2 Limitations of these Studies

One way in which these studies could have been improved would have been to perform experiments testing absolute identification of the individual parameters before running these experiments. The results of these studies would then indicate whether identification of the parameters decreases when they are combined. However, it was decided that it was more important to test how the parameters worked together than to gain a complete understanding of how they functioned independently.

Performance in these experiments might have improved if feedback on performance had been provided during training. As reported in Section 6.3.4, two subjects confused the mapping from rhythm to type; this confusion might have been resolved if the subjects had received feedback about whether their answers were correct during the training phase. In addition, several subjects reported that they found it difficult to judge between medium and high roughness levels. Providing feedback during training might have helped to give users confidence in their responses. However, the training in these experiments was intentionally limited, since in real applications it is unlikely that users would participate in extended training before using such messages in their devices.

In addition there are some factors which may limit the generalisation of these results. For example, the results are only applicable for the three rhythms and three roughness levels used in this study and may not generalise to different stimuli. In addition, only 15-17 users took part in each of these studies; a larger scale study would provide a more accurate representation of the performance levels.

6.5.3 Design Recommendations for Designing Two-Dimensional Tactons

From the results of these experiments it is possible to make a number of recommendations regarding the design of two dimesional Tactons. These recommendations are based on the results of these studies alone, and should be read in combination with the results of the study. This is not a definitive set of guidelines for designing two dimensional Tactons as it is possible that other parameters, which were not tested in these experiments, might yield better performance.

1) ***When designing two-dimensional Tactons, encode one dimension in rhythm and a second dimension in roughness.***

   Average identification rates of 73% were achieved for Tactons which encoded two dimensions of information in rhythm and vibrotactile roughness, with three levels of each parameter encoded in these parameters.

   a) ***Use rhythm to encode the most important dimension.***

      Average identification rates of 96% were achieved for three different rhythms.

   b) ***Use roughness to encode the less important dimension.***

      Average identification rates of 79.5% were achieved when three levels of roughness (an un-modulated sine wave and the same sine wave amplitude modulated by 30Hz and 50Hz) were used to encode information.

2) ***When designing rhythms, using a different number of vibration bursts in each rhythm may help to create absolutely identifiable rhythms.***

   Three rhythms were created with a different number of vibration bursts (2, 4 and 7) used in each rhythm. These were further distinguished by using shorter bursts to create rhythms which feel fast and longer bursts to create rhythms which feel slow.

3) ***When presenting roughness stimuli to the fingertip using a C2 Tactor, use the Tactor in the indented orientation rather than in the raised orientation – the improved perception may be due to the presence of a rigid surround.***

   The above results are those achieved using the C2 Tactor in the indented orientation on the fingertip. When the C2 Tactor was used in the raised orientation there was a significant decrease in identification rates for the roughness parameter and, as a result, for the overall Tactons. This result may have been a result of the presence of the rigid surround in the indented orientation, so may transfer to other devices that feature a rigid surround.

6.6 Conclusions

This chapter reported two experiments investigating identification of two dimensional Tactons, with the two dimensions encoded by manipulating two vibrotactile parameters: rhythm and vibrotactile roughness. These Tactons were presented to the fingertip using the C2 Tactor. The first experiment used the C2 Tactor in the indented orientation, and the second experiment used the device in the raised orientation.

The two research questions posed by this thesis are:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?

RQ2: What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

The experiments in this chapter provide answers to both of these research questions. In answer to Research Question 1, both experiments have shown that rhythm can be manipulated to encode information in Tactons, with average percentage correct rates of around 96% achieved for rhythm in both device orientations. When the C2 Tactor is used in the indented orientation on the fingertip an average percentage correct rate of 79.5% was achieved for roughness. Roughness is a less effective parameter when the C2 is used on the fingertip in the raised orientation, with a recognition rate of 69%.

In answer to Research Question 2, these experiments indicate that performance levels of 73% can be achieved when rhythm and roughness are combined to create Tactons (with scores of 80-100% achieved in 60% of the tasks). The identification rate dropped to 66% when the device was used in the raised orientation rather than the indented orientation. These results were achieved after less than 10 minutes of training, so it may be possible to achieve higher rates if a longer training session took place, or if people were exposed to these stimuli regularly over an extended period of time.

These experiments have shown that it is possible to encode two dimensions of information in Tactons by manipulating rhythm and roughness. Research on Earcons has shown that it is possible to design Earcons which encode three dimensions of information [77] and it would be interesting to see if this can also be achieved by Tactons. Therefore, Chapter 8 investigates the effect of encoding a third dimension of information in Tactons.

These experiments show what can be achieved on a high specification device like the C2 Tactor, but it would also be interesting to investigate whether Tactons can be presented using cheaper, widely available devices such as the vibration motors in mobile phones. Chapter 7, therefore, reports a study evaluating identification of Tactons presented using a standard mobile phone vibration motor.

Chapter 7: Design and Evaluation of Two-Dimensional Tactons for a Mobile Phone Vibration Motor

7.1 Introduction

Having established a baseline for performance on Tacton identification using a specialised vibrotactile actuator (C2 Tactor), as reported in Chapter 6, it is interesting to see how these results compare to the performance achieved using the vibration motors present in devices such as mobile phones. It has been stated that standard phone vibration motors may not be suitable for communicating complex information, as they do not offer subtle control due to their limited bandwidth and high latency [38, 86]. However, if it were possible to present Tactons using such devices this would make the use of such vibrotactile messages more commercially viable as these vibration motors are featured in many common interaction devices (mobile phones, pagers, PDAs, games controllers, and some mice). This chapter presents a design for two-dimensional Tactons for phone vibration motors, and then reports an experiment that was conducted to evaluate these two designs and compare them to the results achieved on the C2 Tactor, reported in Chapter 6. These results will help designers to understand the possibilities offered by standard phone vibration motors for communicating complex information.

This chapter begins, in Section 7.2, with a discussion of the design of Tactons for presentation via phone vibration motors. Section 7.3 then reports an experiment testing absolute identification of these Tactons. This chapter addresses both of the research questions of this thesis, in terms of the presentation of Tactons via a phone vibration motor:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

Possible parameters for use in Tactons presented via a phone vibration motor are discussed in Section 7.2, and absolute identification of these parameters is tested in the experiment reported in Section 7.3 when they are combined to create Tactons. This therefore addresses Research Question 1. This experiment also addresses Research Question 2 which identifies the absolute identification rates for these Tactons. In Section 7.4, the results of the experiment are discussed, and a set of recommendations for designing Tactons for phone vibration motors are presented. Finally, Section 7.5 draws conclusions as to how the results of the experiment reported in this chapter answers the two research questions posed at the start of this thesis.

7.2 Design of Two-Dimensional Tactons for a Mobile Phone Vibration Motor

7.2.1 Device

The device used in this study was a mockup of a widely used mobile phone (Nokia 8210, Figure 7-1), containing a standard phone vibration motor. As discussed in Chapter two, the construction and mode of operation of a mobile phone vibration motor are very different to that of the C2 Tactor. Mobile phone motors are small DC motors featuring an eccentric weight in the shaft. When they are switched on the shaft spins, and the spinning of the eccentric weight causes the sensation of vibration. These vibration motors typically vibrate at frequencies around 130Hz. Unlike the C2 where a small contactor vibrates linearly against the skin, with a phone motor the whole casing in which the motor is housed vibrates. Therefore, when this device is mounted in a phone, the whole body of the phone vibrates.

Figure 7-1: Nokia 8210 Mobile Phone Mockup containing standard phone vibration motor.

Device Control

The motor in this phone was controlled by a piece of hardware, known here as the “Tacton Box”, which was built by staff at Nokia Research Center in Helsinki, Finland. Figure 7-2 shows a block diagram of the Tacton Box. This device is built around a Basic Stamp 2 (BS2) microcontroller by Parallax Inc (http://www.parallax.com/) which can be connected to the serial port of a PC. The vibration motor is driven via an electronic relay connected to the BS2. Three relays are connected in the Tacton box, but only one was required for this study and this one was connected to the motor of the mobile phone mockup.

The BS2 is controlled from the PC by writing code in Basic and sending these commands through the serial port. With this setup it is possible to switch the vibration motor on and off, and control the frequency/intensity of the vibration. A numeric keypad was connected to the Tacton Box (Figure 7-3), and through the Basic software different Tactons could be assigned to each key. Therefore, when a key was pressed the associated Tacton would be presented. This was used during the experiment to present Tactons to the users.

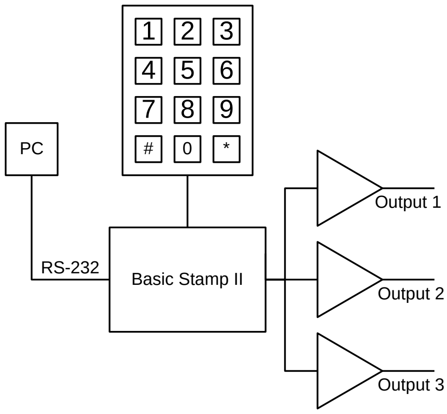

Figure 7-2: Block Diagram showing the design of the Tacton Box used to present Tactons via a phone vibration motor (provided by Topi Kaaresoja).

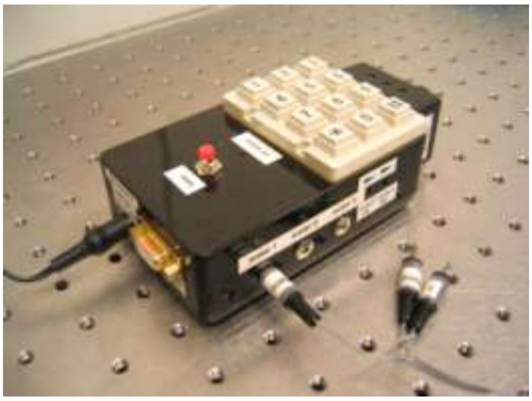

Figure 7-3: "Tacton Box" with numeric keypad: hardware for creating and presenting Tactons via a standard phone vibration motor.

7.2.2 Stimuli Design

As discussed above, the construction and mode of operation of the phone vibration motor is very different to that of a C2 Tactor and, therefore, tactile stimuli will feel quite different when presented using the two different devices. In addition, less control over the stimuli is available when using the phone vibration motor – only the on-off times and the velocity of the rotation can be controlled. Therefore stimuli must be designed within these constraints.

In order to make a direct comparison to the studies on the C2 it was decided to use stimuli that were as similar as possible to those used in the previous experiments. Rhythms are created by controlling only whether the vibration is on or off and therefore the three rhythms used in the previous experiments (7-note, 4-note and 2-note) can be transferred directly to the phone motor.

Roughness, on the other hand, cannot be directly transferred since the amplitude-modulated waveforms used to create the roughnesses used on the C2 could not be reproduced on the phone motor. However, an approximation of roughness was created by using different speeds of on-off pulses, and three levels were selected from the results of informal pilot testing. A "smooth" stimulus was created by using a constant vibration. To create a "rough" sensation, pulses of 10 milliseconds were repeated with 10 millisecond pauses in between, while the "very rough" sensation was created using 30 millisecond pulses, followed by 30 millisecond pauses (Figure 7-4). This implementation exploits the high latency of the vibration motor, as the slow reactions to changes in the input signal result in a signal which resembles those created by amplitude modulation. Figure 7-5 shows the output from the phone measured using a laser vibrometer, using the setup described in [61].

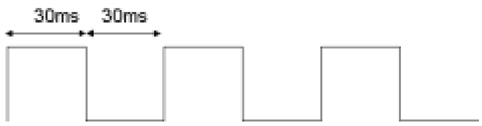

Figure 7-4: Representaton of the on/off input sent to the mobile phone vibration motor to create the "very rough" stimulus.

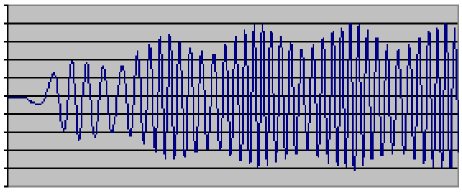

Figure 7-5: The resulting output waveform for the "very rough" stimulus from the mobile phone measured using a laser vibrometer.

In addition to testing roughness it was decided to also test a second parameter that might be better suited to the capabilities of the phone vibration motor, namely intensity. The intensity of the vibration is changed by adjusting the supply voltage to the phone motor. Experimentation indicated that values below 0.93 Volts did not produce a reliable response from the vibration motor, so 0.93V was chosen as the lowest value, 1.38V as the maximum value, and 1.16V as the mid way point. When the supply voltage of a phone motor is adjusted it changes both the intensity and the frequency of the resulting vibration. This, therefore, provides redundant coding of the same information in both of these parameters and may help people to identify multiple levels. Table 7-1 shows the voltage of the three levels of intensity, and also shows the frequency, RPM and velocity values of the outputs as measured using a laser vibrometer.

Creation of Two-Dimensional Tactons

By combining the above parameters, two sets of nine Tactons were created to represent alerts which might occur when a message or a call arrives on a mobile phone. In both sets of Tactons, the rhythms represented the type of alert (voice call, text message, or multimedia message). These rhythms were exactly the same as those used in the C2 experiments in Chapter 6 (Figure 6-1). In one set, the priority of the alert (low, medium or high) was represented by roughness, and in the other set the priority was represented by intensity. The values of both of these parameters were presented above. These two sets of Tactons were then evaluated through an absolute identification experiment, described in Section 7.3.

Intensity Level	Low	Medium	High
Voltage (Volts)	0.93	1.16	1.38
Frequency (Hz)	156	170	180
RPM (1/min)	9360	10200	10800
Velocity (mm/s peak to peak)	19.2	20.8	21.5

Table 7-1: Experiment 4: The voltage, frequency, RPM and velocity of the three levels of intensity (provided by Topi Kaaresoja)

7.3 Experiment 4: Evaluation of Two-Dimensional Tactons presented via a Standard Mobile Phone Vibration Motor

The aim of this experiment was to identify recognition rates for Tactons presented on a phone vibration motor and to compare these to results achieved with the more expensive, specialised C2 Tactor. The two conditions in this experiment are referred to as the Phone Roughness condition and the Phone Intensity condition. Since the experimental design and stimuli used in this study are based on those used in the previous studies on the C2, the results can be compared.

7.3.1 Hypotheses

The hypotheses were as follows:

1. Overall performance will be higher in the Phone Intensity condition than in the Phone Roughness condition (since the intensity parameter is more suited to the capabilities of the phone).

2. Performance on the phone (in both conditions) will be lower than on the C2 due to the fact that it is a less sophisticated device.

3. Rhythm identification will be unaffected by device or by the choice of second parameter.

7.3.2 Experimental Method

Participants

Sixteen participants took part in this study; all were staff members of Nokia Research Center - Helsinki, aged between 23 and 43 (10 male, three female). The experiment consisted of two conditions (*Phone Roughness* and *Phone Intensity*) and the order of these two conditions was counterbalanced, with an equal number of participants taking part in each order of conditions. Participants were given two movie tickets as a reward for participating. The consent form for this experiment is included in Appendix D.1.

Training Procedure

The training procedure was the same as that for the experiments reported in Chapter 6. The participants were given a sheet explaining the Tactons and the information encoded in them (Appendices D2 and D3), and were then allowed to try the Tactons for themselves for up to 10 minutes. The difference was in the experimental setup. Instead of exploring the Tactons via a webpage the participants pressed the buttons on the Tactons box to present the corresponding Tacton. They were provided with a sheet which they could keep alongside the box reminding them which number corresponded to which Tacton. Before starting the experiment, users took part in nine tasks like those in the experiment itself in order to familiarise themselves with the experimental procedure. This procedure is explained below.

Testing Procedure

There were 54 tasks in each condition of this experiment, with each Tacton presented six times during a condition. In each task, a Tacton was presented four times, with a one second gap between presentations. The instruction sheets for this experiment are included in Appendix D.4. While being presented with the Tacton, users made their responses by ticking boxes on a response sheet to indicate the type and priority of the Tacton (see Appendix D.5). The presentation of Tactons was controlled by the experimenter, who pressed the corresponding button on the Tacton box to present the Tacton for each task. The participant indicated when they had completed a task by saying "OK", at which point the experimenter stopped the stimulus. When the participant indicated that they were ready for the next task the experimenter started the next Tacton. The order of presentation was randomised for every participant.

During the experiment the phone was held in the participant's non-dominant hand, and the participant wore headphones to block noise from the device. The participant sat on one side of a screen holding the phone, and the experimenter sat on the other side of the screen with the Tacton box. The participant was able to see the experimenter but not to see which key was being pressed on the Tacton box.

7.3.3 Results

After the experiment the written responses were collated and the number of correct responses to the type (rhythm) and priority (roughness/intensity) of each Tacton were counted. Percentage correct scores were calculated for each individual attribute (type and priority) and for the complete Tactons. These data are included in Appendix G on the accompanying CD. These three sets of data were analysed individually and compared across conditions using Kruskal Wallis tests [56] (as in the previous experiments an Anderson Darling test showed that the data were not normally distributed and therefore non-parametric statistical tests were conducted). Where significant differences were found, *post hoc* Mann-Whitney tests [56] were carried out with a Bonferroni correction applied (resulting in a threshold of 0.008 (0.05/6) for significance) to compensate for carrying out multiple comparisons on the same data. These data sets were also collected in the previous two-dimensional Tactons experiments on the C2 Tactor therefore a comparison can be made between the two conditions in this study and the two versions of the C2 experiment.

Tactons

The results showed overall Tacton identification rates of 52% and 72% for the Phone Roughness and Phone Intensity conditions, respectively, while the C2 experiment had an overall recognition rate of 73% in the indented orientation and 66% in the raised orientation (Figure 7-6). The Kruskal Wallis test for the complete Tactons showed a significant difference in performance between conditions (H=39.44, DF=3, $p<0.001$). Hypothesis 1 can be accepted as *post hoc* Mann Whitney tests revealed that performance was significantly better in the Phone Intensity condition than the Phone Roughness condition ($p<0.001$). Hypothesis 2, however, cannot be accepted as *post hoc* Tukey tests revealed that, while performance was significantly better in both C2 experiments than in the Phone Roughness condition ($p<0.001$), there were no significant differences between performance in the phone intensity condition and either of the C2 experiments (p=0.08 (C2 raised orientation), p=0.3 (C2 indented orientation)).

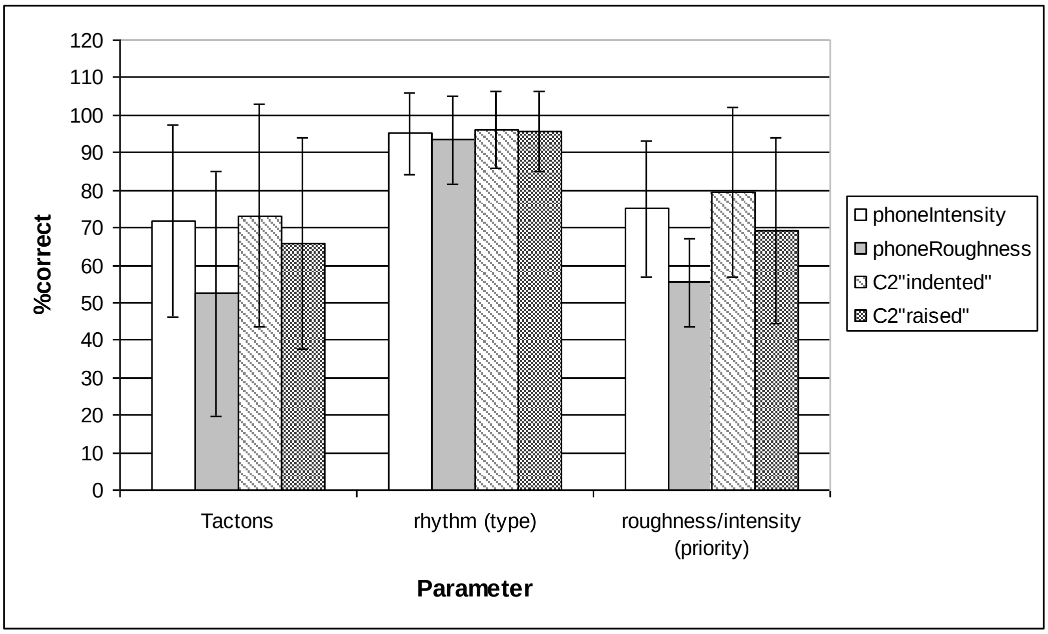

Figure 7-6: Experiment 4: Mean percentage correct scores (with standard deviations) for complete Tactons and individual parameters for the two conditions in Experiment 4 (phoneIntensity, phoneRoughness) and for Experiment 3(a) (C2"indented") and 3(b) (C2"raised").

Rhythm

The results for rhythm identification showed identification rates of 93% in the Phone Roughness condition and 95% in the Phone Intensity condition. Both C2 experiments had found identification rates of around 96% for rhythm. These results are shown in Figure 7-6 while Figure 7-7 shows the percentage correct scores for the individual rhythms in both the conditions using the phone motor. The Kruskal Wallis test showed no significant differences between any of the conditions (H = 4.57, DF = 3, p=1). Since the results are very similar in all conditions, Hypothesis 3 can be accepted.

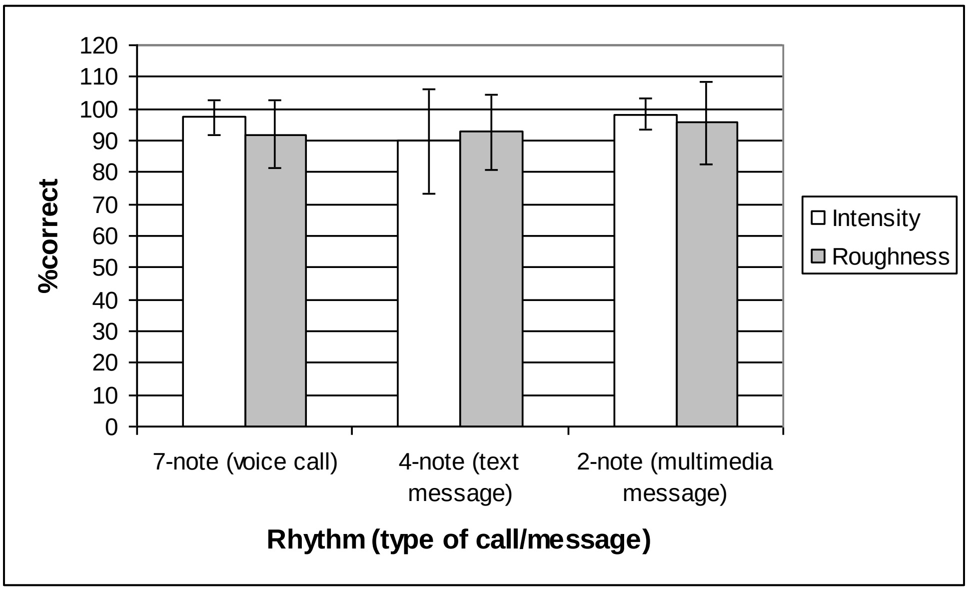

Figure 7-7: Experiment 4: Mean percentage correct scores (with standard deviations) for each of the rhythms in the Phone Intensity and Phone Roughness conditions.

Roughness/Intensity

The results for the second parameter (roughness/intensity) showed identification rates of 55% for the Phone Roughness condition, and 75% for the Phone Intensity condition, while in the C2 experiments the recognition rate for roughness was 79.5% in the indented orientation and 69% in the raised orientation (Figure 7-6). Figure 7-8 shows the individual percentage correct scores for each of the three levels of importance in both of the conditions using the phone motor.

The Kruskall Wallis test showed a significant difference in performance between conditions (H=42.48, DF=3, p<0.001). Hypothesis 1 can, again, be accepted as post-hoc Mann Whitney tests revealed a significant difference between the Phone Roughness condition and the Phone Intensity condition (p<0.001). While there was a significant difference between the Phone Roughness condition and both of the C2 experiments (p<=0.001), there was no significant difference between the Phone Intensity condition and either of the C2 experiments (p=0.09 (C2 in indented orientation), p=0.4 (C2 raised orientation)) and, therefore, Hypothesis 2 cannot be accepted.

By examining the stimulus-response confusion matrices it is possible to gain a better understanding of users' behaviour when identifying the different levels of intensity or roughness. Table 7-2 shows the confusion matrix for the roughness parameter. This reveals that most confusion occurred between the smooth and the very rough stimulus. Qualitative feedback from two participants (see "Qualitative Feedback" section below) indicated that they found the mapping from roughness to priority counter-intuitive as the smooth stimulus felt the strongest. Therefore, it is possible that this confusion resulted from the counter-intuitive mapping from data to stimulus, rather than from an inability to discriminate these different stimuli.

Table 7-3 shows the stimulus-response confusion matrix for the intensity parameter. This shows that there was very little confusion between the medium and high levels of intensity, but that some confusion occurred between the low and high levels, and between the low and medium levels. The confusion between high and low intensities may result from the two-condition experiment. The stimulus used for the high intensity level in the intensity condition was used in the roughness condition to represent low priority, and this may have caused confusion for the participants who took part in the roughness condition first.

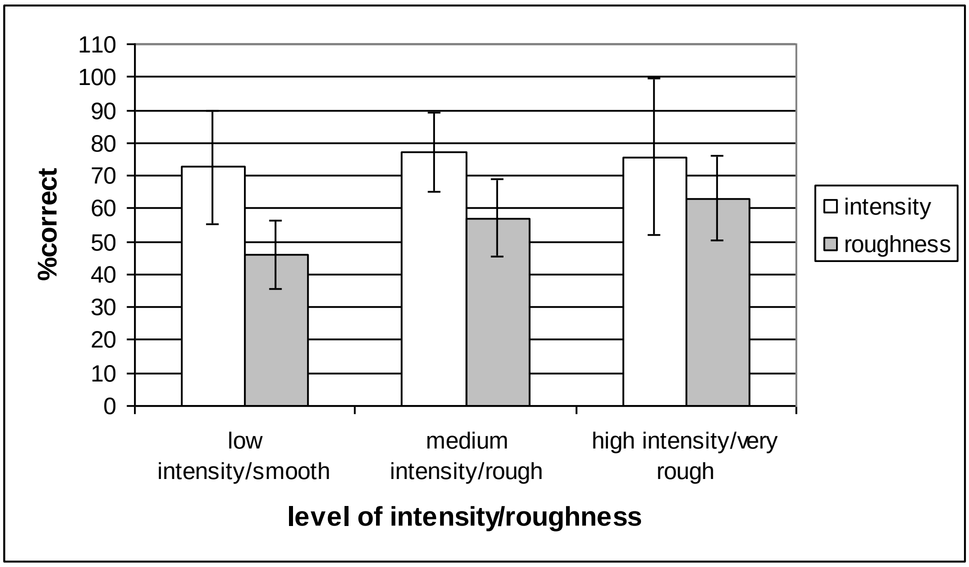

Figure 7-8: Experiment 4: Mean percentage correct scores (with standard deviations) for each of the importance levels (encoded in intensity or roughness) in the Phone Intensity and Phone Roughness conditions.

response

stimulus		smooth	rough	very rough
	smooth	45.83%	11.11%	43.06%
	rough	32.64%	56.94%	10.42%
	very rough	29.17%	7.64%	63.19%

Table 7-2: Experiment 4: Stimulus-response confusion matrix for roughness parameter.

response

stimulus		low	med	high
	low	72.57%	11.46%	15.97%
	med	22.57%	77.08%	0.35%
	high	24.31%	0.00%	75.69%

Table 7-3: Experiment 4: Stimulus-response confusion matrix for intensity parameter

Qualitative Feedback

After each condition, informal interviews were conducted with the participants. They were prompted to discuss how easy or hard they found it to identify the type and priority of a call/message, which condition they preferred, and to give any other comments. The main findings from these comments are discussed in this section.

Thirteen out of the sixteen users expressed a preference for the intensity condition over the roughness condition. This was expected, given that the intensity condition yielded significantly higher performance. Reasons given for preferring this condition were that the concept of intensity was easier to understand than the concept of roughness, or simply that the three levels were easier to distinguish in this condition. Three of the users who expressed a preference for the intensity condition said that they liked the concept of roughness, but that the intensity condition seemed more practical in a real situation (e.g. easier to distinguish levels, particularly if phone was in a pocket). Of the three users who preferred the roughness condition, two said that they found it easier to distinguish the three levels, and one said that they actually found it harder than intensity but they preferred the feeling of it, and would therefore prefer it in a real system.

Eleven of the sixteen users commented that the rhythms were easy to distinguish. Of the users who did not make this comment, four said they found it slightly difficult to distinguish between the 7-note and 4-note rhythms (one of those also found it hard to distinguish between the 4-note and 2-note rhythms). The final user said she found it hard to learn the rhythms as she had no musical skills. Two users commented that the choice of rhythms could be more intuitive. For example, the 2-note rhythm representing multimedia messages was too similar to that used in real phones for a text message alert, and also felt quite urgent, while multimedia messages are not usually urgent at all.

Almost all of the users commented that they found it hard to identify all three levels of priority in either condition, and that it would be much easier if only two levels were used. Furthermore, six participants said that they felt that two levels would be sufficient in a real phone for this application as all they would want to know was whether the call/message was urgent or not. One participant commented that while he liked the rhythms and would use those in a real application, he would not use the priority option at all. Another general comment about the levels of the second parameter was that relative judgements were quite easy but the absolute judgements were very difficult. In addition, two participants commented that the mapping from importance to roughness seemed wrong as the smooth vibration felt strongest but represented the highest priority alert.

One final comment from two users was that, especially in the intensity condition, they could hear some output from the phone, even though their ears were blocked, and could use this to help them judge the intensity level. In a real system it is likely that this sound would be masked by environmental noise and thus the users would have to rely on vibration alone. A follow on study where noise is played through the headphones to completely block the sound from the device might be beneficial in determining the impact of the sound leakage on performance levels.

7.4 Discussion

These results are very promising for presenting Tactons via phone vibration motors. The overall recognition rate of 72% (when intensity is used to encode priority on the phone) compares favorably to the 73% rate achieved on the C2 Tactor (in indented orientation). In addition the results for rhythm (95%) and intensity (75%) are also comparable to those achieved for the individual parameters in both of the C2 experiments. On the other hand, when roughness is used to encode priority the results are significantly worse (52% overall, and 55% for roughness) and, therefore, roughness should not be used as a Tacton parameter when using mobile phone vibration motors. The performance on intensity may have been aided by the fact that changing the intensity level of a phone motor also changes the frequency and thus the same information was encoded redundantly by both of these parameters.

Many participants reported that they found it difficult to identify all three levels of the second parameter in either condition, and several mentioned that for this particular application of Tactons (namely mobile phone alerts), two levels would be sufficient as they are only interested in whether a call or message is important or not; a third level is not required. Reducing the number of levels to two could be considered for this particular application of Tactons, and might improve performance for both roughness and intensity. However there may be other applications in which three levels are required.

7.4.1 Limitations of this Study

During the study users held the phone in their hand, whereas in a real situation the phone would be more likely to be in a pocket. The ability of users to discriminate rhythms and intensity level may be reduced when the phone is not in the hand. In addition in a real application the user's full attention would not be on the tactile alerts. Therefore future work might investigate identification rates when the phone is in a pocket and the user is engaged in another task.

It might be possible to improve intensity recognition. Due to the ramp up time taken to reach the full vibration speed, the motor does not reach full intensity when very short pulses are presented [62]. As a result, in rhythms where pulses are very short (e.g. 7-note rhythm) the intensity will feel lower. Using rhythms with longer pulses would avoid this problem and possibly improve performance. In addition, there was no formal preliminary study to choose the best intensity levels, therefore it may be possible to select more appropriate intensity levels, and thus improve performance.

The roughness stimuli were also designed without any formal preliminary study to select the best levels, therefore it is possible that better results may be possible. However, the response time of the phone motor means that the range of roughness levels that can be used is limited. Two users stated that the mapping from roughness to importance seemed wrong as the smooth signal felt the strongest but represented the least important alert. The relative magnitude of each of the signals may need to be matched to avoid this confusion.

7.4.2 Design Recommendations for Designing Two-Dimensional Tactons for Standard Phone Vibration Motors

Based on the results of this study, a number of recommendations for designing two-dimensional Tactons for phone vibration motors can be made. As in the previous chapter, these are not definitive guidelines, but they do provide guidance based on the results of this particular study.

1. ***When designing two-dimensional Tactons for phone motors, encode one dimension of information in rhythm and another in intensity.***

   An average percentage correct rate of 72% was achieved when two-dimensional Tactons were created, encoding information in three levels of rhythm and three levels of intensity.

   a. ***Use rhythm to encode the most important dimension***

      Percentage correct rates for rhythm of 93-95% were achieved when three different rhythms were used.

   b. ***Use intensity to encode the less important dimension***

      An average identification rate of 75% for intensity was achieved when three levels of intensity were used.

2. ***Do not use roughness to encode information when using standard phone vibration motors.***

   The average identification rate achieved for three levels of roughness on the phone motor was just 55%.

7.5 Conclusions

In chapter 1, the following two research questions were posed:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

Chapter 6 addressed both of these questions based on the results of studies carried out on a high specification vibrotactile device, the C2 Tactor. This chapter reported a study evaluating performance levels for two-dimensional Tactons presented using a standard mobile phone vibration motor. The results of this study can be used to address these questions in relation to the use of a standard phone vibration motor.

In answer to Research Question 1, this experiment showed that the rhythm parameter transferred successfully to the phone vibration motor, with an identification rate of 95% achieved for this parameter. The roughness parameter, however, did not transfer successfully as identification rates of just 55% were achieved for this parameter. Intensity performed well as a Tacton parameter with levels of 75% achieved.

In answer to Research Question 2, when Tactons encoding two pieces of information (type and priority of alert) were created by encoding information in rhythm and intensity recognition rates of 72% were achieved. This dropped to 52% when roughness was used in place of intensity.

The results for two-dimensional Tactons encoding information in rhythm and intensity are comparable to those achieved for two-dimensional Tactons presented on a high-price vibrotactile transducer reported in Chapter 6. It had previously been suggested that mobile phone vibration motors are not suitable for presenting complex tactile messages. These results indicate that it is possible to use such motors to present Tactons, and this finding makes the use of messages such as Tactons more feasible commercially, since these standard vibration motors are extremely cheap and similar motors are already featured in many devices, including games controllers, handheld computers, mice and, of course, mobile phones. In terms of the mobile phone application itself, using Tactons in mobile phone alerts would enable more information to be transmitted in these alerts, and offer personalisation features, such as assigning personalised vibration "ringtones" to different callers.

Chapter 8: Design and Evaluation of Three-Dimensional Tactons

8.1 Introduction

The previous two chapters reported the results of studies investigating identification of Tactons encoding two dimensions of information. When the application in which the Tactons are used requires that more dimensions of information be encoded, it will be necessary to extend the parameter-space of these Tactons. Therefore, a third vibrotactile parameter, in which another dimension of information can be encoded, must be identified. This chapter reports three experiments investigating the design of three-dimensional Tactons.

Section 8.2 describes the initial design of these Tactons, and the decisions made regarding the choice of a third parameter. An absolute identification experiment was conducted to investigate the effect of adding this third parameter on the performance level achieved, and this is reported in Section 8.3. This experiment showed that spatial location was a suitable third parameter but that the use of this parameter degraded performance on identification of roughness. Therefore, the following two experiments aimed to improve performance by changing the "roughness" parameter. Section 8.4 reports an experiment using just two levels of roughness. In Section 8.5 a new design for three-dimensional Tactons is presented, where different parameters are used in place of roughness, and the evaluation of these Tactons is reported in Section 8.6. Section 8.7 discusses the results of these three studies and gives recommendations for the design of three-dimensional Tactons.

This chapter addresses both of the research questions posed at the start of the thesis, in terms of the design of three-dimensional Tactons.

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

Research Question 1 is addressed through a discussion of suitable parameters for three-dimensional Tactons and through the evaluation of individual parameters when used within three-dimensional Tactons. The results of all of the experiments in this chapter provide absolute identification rates for individual vibrotactile parameters when used within three-dimensional Tactons. In addition, the results of these experiments are also used to calculate the information transmitted by each of these parameters. Research question 2 is addressed through the experiments evaluating three-dimensional Tactons, which provide identification rates for these Tactons, and also provide figures for the amount of information transmitted by them.

Section 8.8 concludes this chapter, drawing general conclusions from this work on three-dimensional Tactons and discussing how the findings of these experiments answer the research questions posed in this thesis.

8.2 Design of Three-Dimensional Tactons

8.2.1 Vibrotactile Device

The devices used for this research were the C2 Tactors from EAI Inc. and these were used in the raised configuration, as shown in Figure 8-1 (the same configuration was used in Experiment 3(b)). Therefore, this work builds on the studies using the C2 Tactor from Chapter 6, rather than from the work using the phone vibration motor from Chapter 7.

Figure 8-1: C2 Tactor in "raised" orientation.

8.2.2 Selection of a Third Parameter

The discussion of suitable parameters for encoding information in Tactons in Chapter 4 identified spatial location as one of the most promising parameters. Therefore, it was decided that spatial location should be used as the third parameter for three-dimensional Tactons, along with the two existing parameters established in Chapter 6: rhythm and roughness.

As discussed in Chapter 2, two different ways of using spatial location exist. In the first approach, static locations can be used where the same stimulus can be presented to different locations, with some information encoded in the location. The second approach is to encode information in movements across several spatial locations, such that a movement from right to left might encode different information from a movement from left to right. The static spatial location approach was used here, as it was felt that this would combine more effectively with the existing parameters.

In order to use static spatial location as a parameter in Tactons, it is important to choose the locations carefully. As discussed in Chapter 2, precise localisation can be achieved when stimuli are presented close to anatomical reference points, and in particular at points of mobility, such as the wrist or the elbow. In addition it has been shown that people are unlikely to mistake stimulation at another point for stimulation at one of the reference points. These findings suggest that for accurate localisation of three locations, two actuators should be located at anatomical reference points, with the remaining Tactor located at a point between these two. Based on the above findings, localisation should be accurate at the reference points, and stimulation at the non-reference point is unlikely to be confused with stimulation at reference points.

8.2.3 Stimuli Design

It was discussed in Chapter 6 that, although Tactons are abstract and could therefore represent any data, it is useful to select an application domain with which users are familiar to help with evaluating Tactons. For this study, three-dimensional Transformational Tactons were created to represent alerts which might occur in an electronic diary to remind the user of an upcoming appointment, as this is an application in which simple vibrations are currently used. Three pieces of information were encoded in each Tacton: the type of appointment, the importance of the appointment and the time remaining until the appointment. Each piece of information had three possible values: the *type* of appointment could be a *Meeting*, a *Lecture* or a *Tutorial*, the *importance* could be *Low*, *Medium* or *High* and the *time until the appointment* could be *30 minutes*, *15 minutes* or *5 minutes*. This resulted in a set of 27 Tactons. The following sections describe how these data were encoded in the Tactons.

Rhythm

Rhythm was used to encode the type of appointment: a meeting, a lecture or a tutorial. The rhythms used for these three dimensional Tactons were the same as those used for two dimensional Tactons, but the rhythms represented the type of appointment instead of the type of call/message. The 7-note rhythm now represented a Meeting, the 4-note rhythm represented a Lecture, and the 2-note rhythm represented a Tutorial (see Figure 8-2).

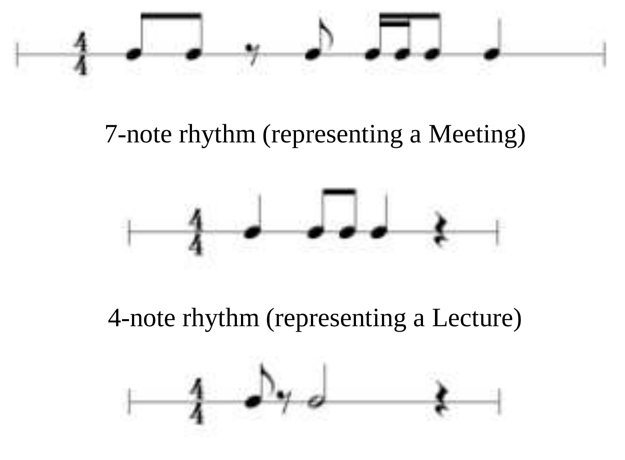

7-note rhythm (representing a Meeting)

4-note rhythm (representing a Lecture)

2-note rhythm (representing a Tutorial)

Figure 8-2: The three rhythms used for Three Dimensional Tactons.

Roughness

The importance of an upcoming appointment was encoded in the roughness parameter. The roughnesses used for three-dimensional Tactons were the same as those used for the two dimensional Tactons, with the smooth stimulus representing low importance, the rough stimulus representing medium importance and the very rough stimulus representing high importance.

Spatial Location

Three locations on the user's volar forearm (the fleshy underside of the forearm) were used to encode information about the time remaining before the appointment – the wrist, the elbow, and the point equidistant between these two reference points. Based on the findings discussed in Section 8.2.2, the reference points (wrist and elbow) should be accurately localised, and the third point should not be confused with either of the two reference points, therefore near-perfect performance should be possible.

C2 Tactors were attached to the user's forearm at these three points using Velcro bands. The "wrist" Tactor was positioned 1cm from the participant's wrist, the "elbow" Tactor 1cm from the participant's elbow and the "mid-point" Tactor at a point equidistant between these two Tactors (Figure 8-3). A vibration at the wrist indicated that the appointment was due to occur in 5 minutes, a vibration at the mid-point that it would occur in 15 minutes, and a vibration at the elbow that the appointment would take place in 30 minutes.

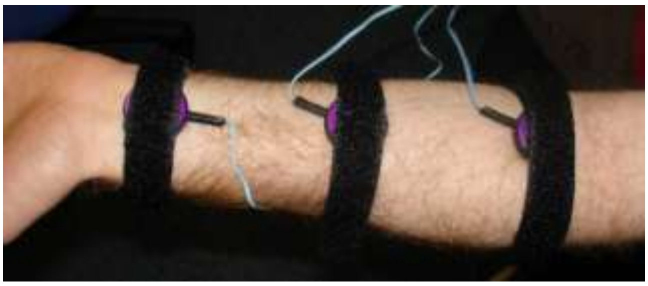

Figure 8-3: The location of the Tactors on the participant's forearm.

Combining parameters

To create the three-dimensional Tactons, the three parameters described above were combined, with each parameter encoding a different dimension of information. As explained above, the Tactons represented alerts for upcoming appointments in an electronic diary, and three pieces of information were encoded regarding these appointments: the type (meeting, lecture or tutorial), the importance (low, medium or high) and the time remaining until the appointment (30 minutes, 15 minutes or 5 minutes). The type of appointment was encoded by the rhythm, the importance by the roughness and the time remaining until the appointment by the spatial location. Three examples will explain this mapping in more detail:

1. A *High Importance Meeting in 5 minutes* would be represented by the 7-note rhythm and a "very rough" vibration and would be presented to the wrist.

2. A *Low Importance Meeting in 5 minutes* would be represented by the 7-note rhythm and presented to the wrist (as above), but would be represented by a "smooth" vibration.

3. A *Low Importance Lecture in 30 minutes* would be represented by the 4-note rhythm and a "smooth" vibration, and would be presented to the elbow.

Stimuli Creation

The wavefiles used in Experiments 3(a) and 3(b) were used in this experiment, but instead of being presented to a single point, they were presented via three different Tactors to add the third dimension of information. Since multiple Tactors had to be controlled, the audio files were sent via a multi-output soundcard from MAudio (http://www.maudio.co.uk/) via three external amplifiers to the three Tactors[2]. The audio files can be found in Appendix H on the accompanying CD by opening the html file "Exp 5(a) and (b) – 3D Tactons.html".

8.3 Experiment 5(a): Evaluation of Three Dimensional Tactons

An experiment was conducted to investigate absolute identification of Tactons encoding three dimensions of information. The Tactons described in Section 8.2.3 were presented to users and data were recorded on the identification of complete Tactons (correct identification of all three parameters – rhythm, roughness and spatial location), in addition to the recognition of the individual parameters.

The aim of this experiment was to investigate identification rates for three-dimensional Tactons in an absolute identification experiment, and to understand the effect on the results of adding a third parameter, compared to the two dimensional Tactons experiments reported in Chapter 6. Adding a third parameter to Tactons is likely to degrade perception of the existing parameters. Therefore, it was expected that the overall identification performance on Tactons would decrease. In particular, the fingertip (which was used in the previous study) is one of the most sensitive parts of the body for perceiving vibrotactile stimuli [32] and is therefore likely to be the most sensitive to subtle variations in vibration such as roughness. The forearm contains fewer receptors and is therefore less likely to be as sensitive to these subtleties.

Rhythm is less likely to be affected by the change in location as it does not require subtle discriminations to be made. Therefore, comparable performance to that achieved in the previous study (i.e. over 90%) was expected. Performance for the spatial location parameter was expected to be close to perfect (over 90%), due to the careful choice of locations as discussed previously.

8.3.1 Hypotheses

Based on the information presented above, the hypotheses for this experiment were as follows:

1. The identification of complete Tactons will be lower for three-dimensional Tactons than for two-dimensional Tactons.

2. Rhythms will be correctly identified over 90% of the time.

---

[2] The code to control the multiple outputs of the soundcard was adapted from code by Toby Opferman, available from http://web2.codeproject.com/audio/wavefiles.asp

3. Perception of roughness will be lower for three-dimensional Tactons than for two-dimensional Tactons.

4. Locations will be correctly identified over 90% of the time.

8.3.2 Hardware setup

Three C2 Tactors were attached to the participant's non-dominant forearm using Velcro bands (Figure 8-3) at the positions described in Section 8.2.3. The participant's dominant arm was left free to control the mouse, which was used to make responses. Participants wore closed-back headphones to block out any noise from the Tactors.

Tactile sensitivity varies across the length of the forearm, therefore the intensity of the vibrations could feel very different at different points on the arm, even if the actual output levels were carefully controlled [29]. These differences in intensity could be used by participants to aid in localising the vibrations, rather than just using the spatial location. Therefore, before starting the experiment it was necessary for each participant to set the levels of the three Tactors so that they all felt perceptually equivalent. The intensity of the mid-point transducer was fixed, but participants were able to alter the strength of the vibrations at the elbow and wrist Tactors. They could alter the strength by turning the volume control on an amplifier up and down for each Tactor until they felt that the strength perceptually matched the strength of the vibration at the mid-point Tactor. Instructions were provided to explain this procedure to the participant and these are included in Appendix E.2.

8.3.3 Experimental Method

The following sections describe the methodology used for this experiment.

Participants

16 participants took part in this experiment and were paid £6. The participants were all students or staff of the University of Glasgow and were aged 17 to 36 (mean=23). Six of the participants were male, while 10 were female. Two were left handed and the rest were right handed. The consent form, which was signed by all participants, is included in Appendix E.1.

Training Procedure

As in the Two-Dimensional Tactons experiments, participants were first provided with a written description of the Tactons (Appendix E.3), and were then allowed to try the Tactons for themselves for up to ten minutes. This self-guided training was carried out using the interface shown in Figure 8-4. When the participants clicked on a button they were provided with the corresponding Tacton. Participants were advised to use this time to learn the mapping between the Tactons and their meanings before being tested on these in the experiment.

After familiarising themselves with the Tactons, participants carried out 27 practise tasks so as to learn how to use the response dialogue. In this training phase, each of the Tactons was presented in one of the tasks, so that users had equal exposure to each of the stimuli. As in the previous experiments, this extended training also aimed to reduce the learning effect during the experiment itself. To remain consistent with the previous experiments, no feedback on performance was provided during any stage of the training.

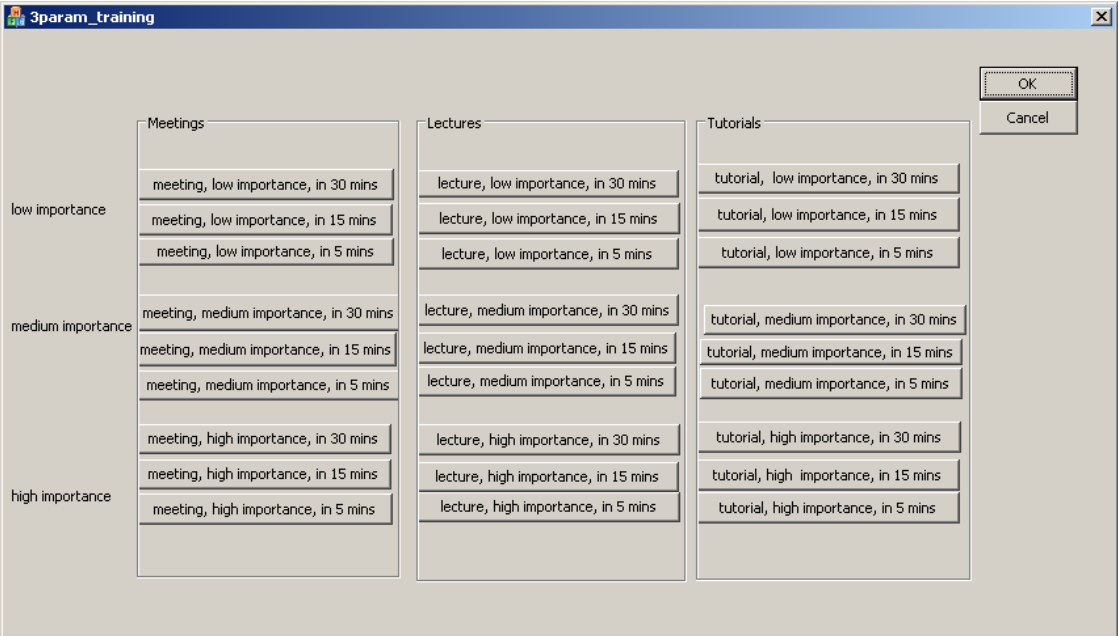

Figure 8-4: Experiment 5(a): Training interface.

Testing Procedure

There were 81 tasks in this experiment, with each Tacton being presented three times during the experiment. The order of these tasks was randomised for every participant to avoid learning effects. The testing procedure was the same as that used in the two-dimensional Tactons experiments. In each task a Tacton was presented to the user and was repeated four times with a one second pause between repetitions. During this time the participant had to identify all three attributes encoded in the Tacton by clicking the corresponding radio buttons on a dialogue box (Figure 8-5). It was not necessary to wait until the Tacton had been presented all four times before responding if a decision was made before then. When participants had made their decision, clicking the "Next Task" button entered their response and popped up a dialogue box prompting them to click "OK" when they were ready for the next task. The instruction sheet given to participants is included in Appendix E.4.

8.3.4 Results

During the experiment data were collected on the number of correct responses to the parameters of each Tacton, and percentage correct scores were calculated for each individual parameter and for the complete Tactons. These data can be found in Appendix G on the accompanying CD. In addition the information transmitted for the complete Tactons and for the individual parameters was calculated from stimulus-response confusion matrices [97].

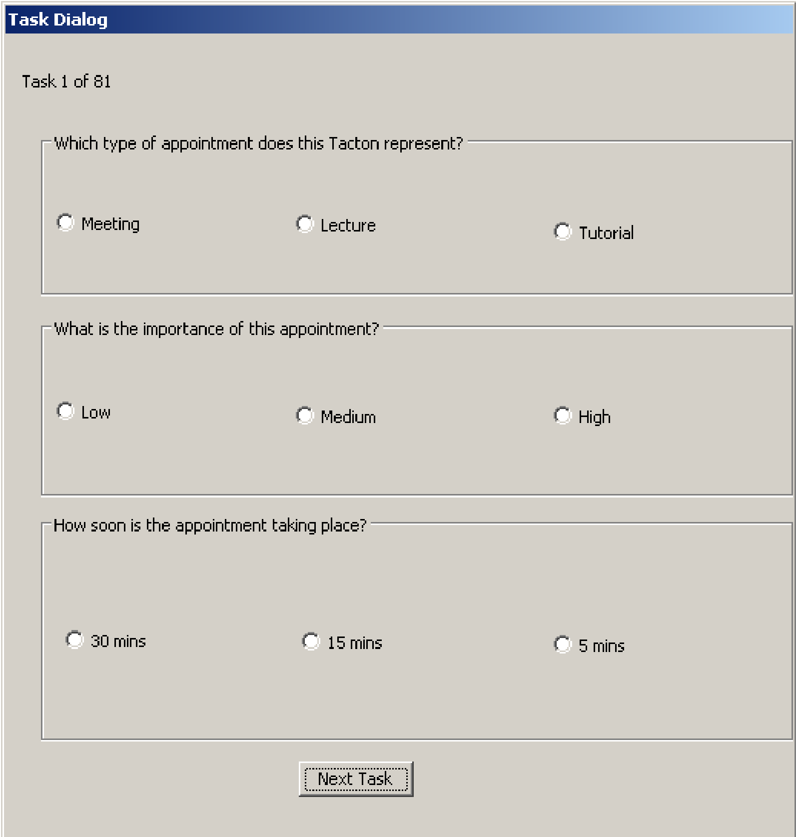

Figure 8-5: Experiment 5(a): Response Dialogue for the Three Dimensional Tactons experiment.

Percentage Correct Scores

The results for overall Tacton recognition showed an average recognition rate of 48%. The best performance was achieved for the High Importance Lecture in 15 Minutes (7-note rhythm, very rough, presented to elbow), which had a recognition rate of 71%, while the Low Importance Tutorial in 5 Minutes (2-note rhythm, smooth, presented to wrist) had the lowest recognition at 21%. The scores for all Tactons are shown in Figure 8-6, where each Tacton is labeled in terms of the rhythm, roughness and location used to encode the data in it. This graph shows a large range of scores with the highest scores generally achieved for the very rough stimuli, and the lowest scores for smooth stimuli.

The rhythms (representing appointment types) were correctly recognised on average 90% of the time. This was calculated by averaging the scores for each individual rhythm and the identification rates for each rhythm are shown in Figure 8-7. The results for recognition of location (representing time until appointment) showed an average recognition rate of 89%: the average scores for each individual location are shown in Figure 8-8.

The results for roughness identification (representing importance of appointment) showed an average recognition rate of 50%. Figure 8-9 shows the individual recognition rates for each level of roughness, and indicates that, as shown in the overall Tactons results, performance is best on high priority stimuli and worst on low priority stimuli. The confusion matrix in Table 8-1 shows that most confusion occurred between adjacent stimuli (smooth-rough or rough-very rough); less confusion occurred between the smooth and very rough stimuli.

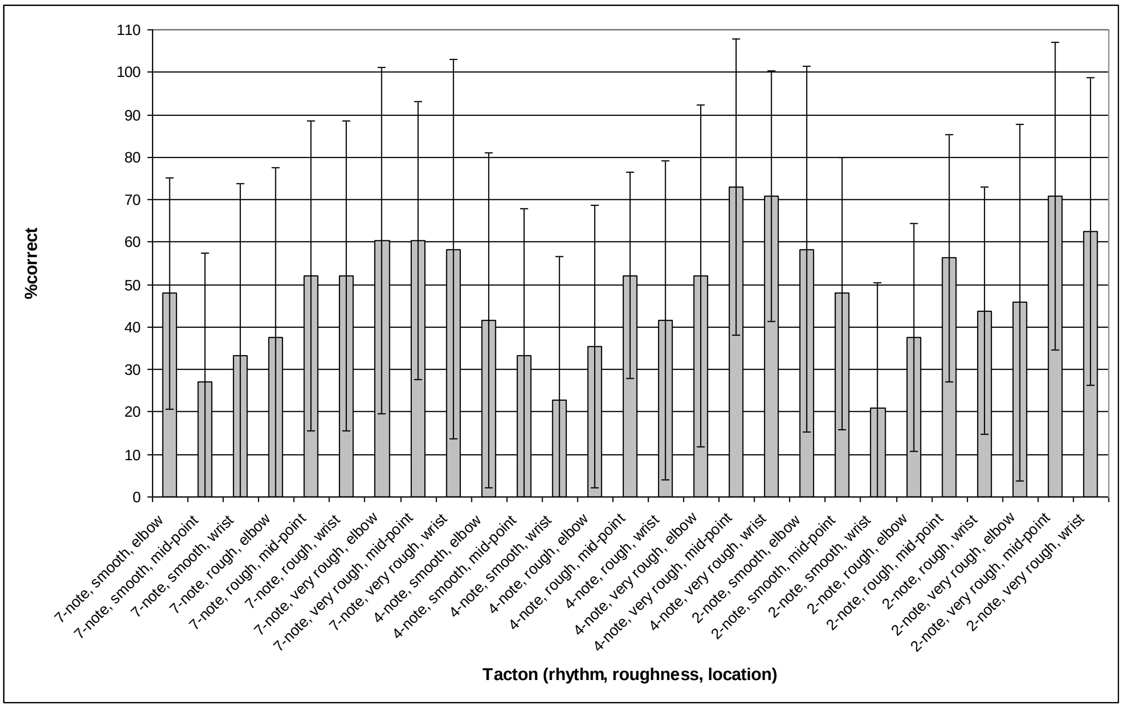

Figure 8-6: Experiment 5(a): Mean percentage correct scores (with standard deviations) for complete Tactons.

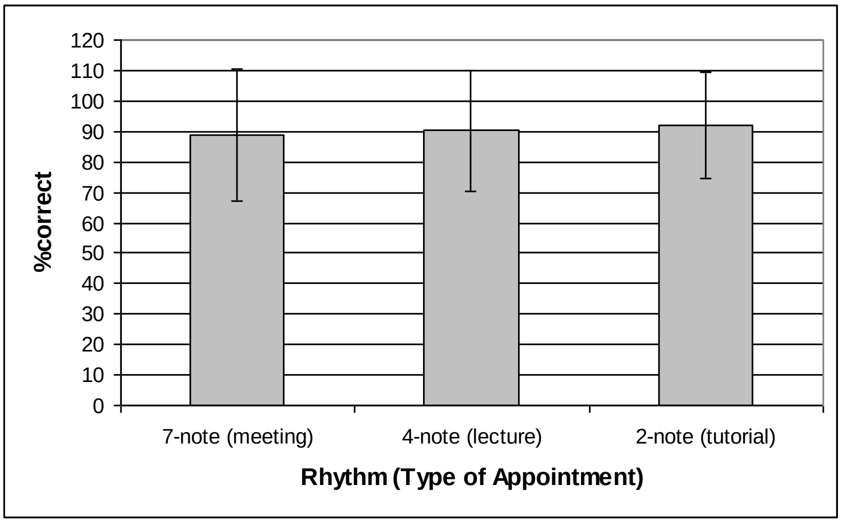

Figure 8-7: Experiment 5(a): Mean percentage correct scores (with standard deviations) for the rhythm parameter (representing type of appointment).

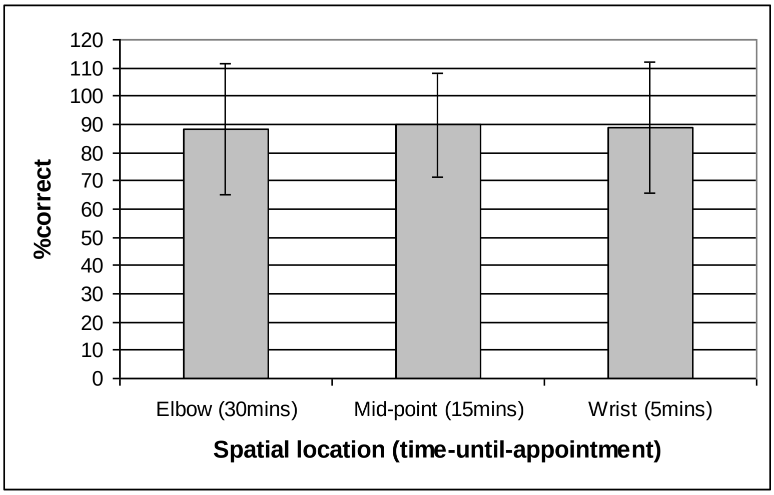

Figure 8-8: Experiment 5(a): Mean percentage correct scores (with standard deviations) for the spatial location parameter (representing time-until-appointment).

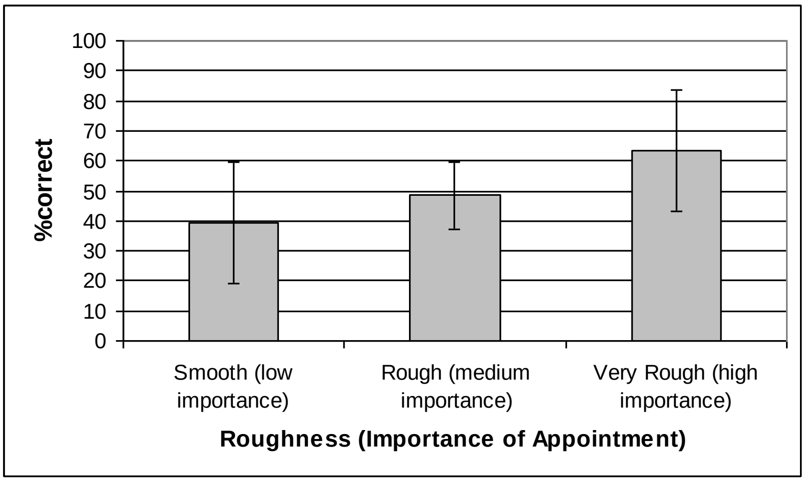

Figure 8-9: Experiment 5(a): Mean percentage correct scores (with standard deviations) for the roughness parameter (representing importance of appointment).

stimulus \ *response*	**smooth**	**rough**	**very rough**
smooth	39.35%	37.27%	23.38%
rough	24.31%	48.38%	27.31%
very rough	10.88%	25.69%	63.43%

Table 8-1: Experiment 5(a): Stimulus-response confusion matrix for the roughness parameter.

Comparison to Two-Dimensional Tactons results

A Kruskal-Wallis test was carried out to compare the results of this experiment with those for both of the two-dimensional Tactons experiments reported in Chapter 6 (Experiments 3(a) and 3(b)). This test showed a significant difference in the percentage correct scores for overall Tactons ($H = 13.35$, $DF = 2$, $p = 0.001$) between the three experiments. Post hoc Mann Whitney tests with Bonferroni correction (resulting in a threshold of $p=0.016$ for significance) indicated that there was a significant difference between the results in the three-dimensional Tactons experiment and the results in both the indented orientation (p=0.0005) and the raised orientation (p=0.01) versions of the Two-Dimensional Tactons experiment.

The Kruskall-Wallis analysis of percentage correct scores for roughness showed a significant difference (H = 16.50, DF = 2, P = 0.000). *Post hoc* Mann Whitney tests with Bonferroni corrections indicated that, as for complete Tactons, there were significant differences between the results from the three-dimensional Tactons experiment and the results from both the indented orientation (p=0.0005) and the raised orientation (p=0.007) versions of the two-dimensional Tactons experiment.

No significant differences were found in the percentage correct scores for type (rhythm) between the three-dimensional Tactons experiment and either of the two-dimensional Tactons experiments (H = 1.23, DF = 2, p = 0.540). Spatial location was not used as a parameter in two-dimensional Tactons therefore no comparison can be made.

Discussion of Percentage Correct Scores

The percentage correct score for complete Tactons (48%) is significantly lower than that achieved for two-dimensional Tactons (66-73%). In addition, the percentage correct score for roughness (50%) is also significantly lower than that achieved for two-dimensional Tactons (69-79.5%). The scores for rhythm and spatial location were around 90% and no significant difference was found between the scores for rhythm in this experiment and in the two-dimensional Tactons experiments.

Since the performance level was significantly lower for both overall Tacton identification and for roughness identification, Hypotheses 1 and 3 can be accepted. Hypothesis 2 stated that the identification rate for rhythm would be over 90%. The rate achieved was 90% and, therefore, Hypothesis 2 can be accepted. Hypothesis 4 was that the identification rate for spatial location would be over 90%. This recognition rate was 89% which is very close and, therefore, Hypothesis 4 can be accepted.

These results suggest that perception of roughness degrades when spatial location is added to encode a third dimension of information. This could be due to two factors: firstly that the use of locations other than the finger degrades perception of roughness since the sensitivity of the skin on the arm is lower that that on the fingertip, or, secondly, that the brain is overloaded with too much information. By examining the information transmission results it is possible to identify how many values of each parameter should be used to achieve near-perfect performance.

8.3.5 Analysing Effectiveness using Information Transmission

Another metric which can be used to measure the effectiveness of Tactons is information transmission. Information transmission (also known as "uncertainty reduction" [97]) is a measure of the correlation between the amount of information in a stimulus set and the amount of information in the responses given by the participant [81]. Miller [81] suggests that this can be considered as two overlapping circles (Figure 8-10), with the left circle showing the information input (the amount of information in the stimuli), the right circle showing the information output (the amount of information in the participant's responses), and the overlapping section the information transmitted (the information in the user's response which matches that in the stimuli). Information transmission has previously been used to evaluate the effectiveness of vibrotactile stimuli and has been shown to be an effective measure [29, 30, 89, 103, 106].

Calculating Information Transmission

A full description of how information transmission is calculated can be obtained from [97], but the following describes the calculations briefly.

Information transmission is also known as uncertainty reduction [97], which refers to the fact that our uncertainty about the content of a stimulus should reduce once we gain some knowledge about it. Before receiving a stimulus the uncertainty score will be high as the user has no knowledge of its content. The original uncertainty (or the information in the stimulus set) is calculated by taking the logarithm (to base 2) of the number of alternatives available for the user to choose from. For example there are 27 Tactons in the set used in this experiment and therefore the original uncertainty score for this set of Tactons is 4.75 bits.

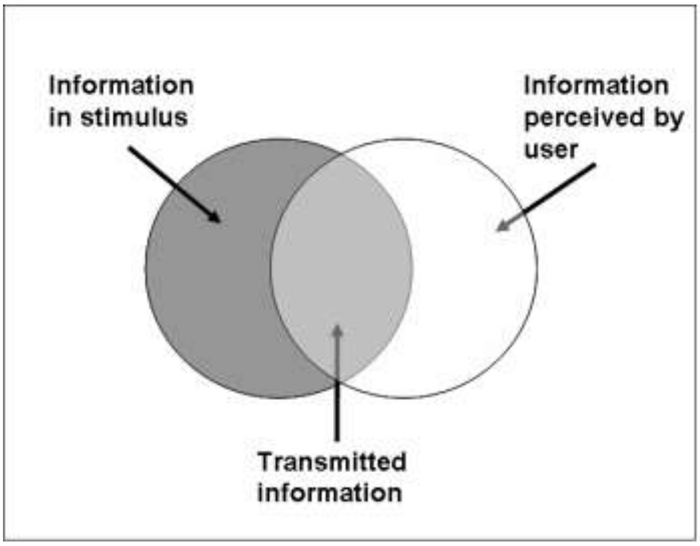

Figure 8-10: Information Transmission Diagram adapted from a description by Miller [81].

The uncertainty in the responses is calculated from the stimulus-response confusion matrices resulting from an absolute identification experiment. The following equation (from [97]) is used to calculate the response uncertainty, where p(i) is the proportion of times each stimulus is selected, and the result is the sum of p(i) log2p(i) for every stimulus.

$$H = \sum p(i) \log_2 p(i)$$

The reduction in uncertainty (or the amount of information transmitted) is then calculated by subtracting the response uncertainty from the original uncertainty. This provides a value for the amount of information transmitted in bits.

An alternative to expressing the amount of information transmitted as a number of bits is to convert this number to a value which indicates the number of stimulus categories that can be correctly identified [97, 106]. This is calculated by raising two to the power of the number of bits (e.g. if the information transmitted via the rhythm parameter was 1 bit, the number of rhythms that should be used if near perfect performance was to be achieved would be 2). In this paper every time an amount of information is reported in bits, the corresponding number of stimulus categories that this refers to is also reported. These are referred to as "tokens".

The above calculations depend on a large number of samples. When the number of samples is low the stimulus-response confusion matrix is sparsely populated and the above calculations will result in an overestimate of the information transmitted [59, 80]. A number of corrections have been proposed for small samples [80], but these can be unreliable when the samples are very small (n<5 for each cell in the confusion matrix) [59]. However, it is possible to estimate the overall information transmitted by a multidimensional data set by calculating the information transmitted by each individual dimension, and then summing these value, and this can be used for small samples as the number of samples for each dimension is sufficiently large [59]. This method of estimating overall information transmission for a multidimensional data set has also been used by Pollack and Ficks [85] and Houtsama [59], both of which are reported in [59].

Information Transmission Results for Experiment 5(a)

The maximum information available from the set of Tactons was 4.75 bits (27 tokens) while, for each individual parameter, a maximum of 1.58 bits (3 tokens) was available. The actual information transmitted in this experiment was calculated for the complete set of Tactons and for the individual parameters using the stimulus-response confusion matrices (the calculations can be seen in Appendix G on the accompanying CD).

As the number of samples was low for complete Tactons (less than 5 samples per cell), the calculations outlined above could not be used to calculate the information transmitted by complete Tactons. Instead, the information transmitted was calculated for each individual parameter (as there were sufficient samples for each individual parameter), and then these were summed to estimate the total information transmitted.

For the individual parameters, the information transmitted through rhythm was 1.03 bits (2.05 tokens), through location was 0.98 bits (1.99 tokens), and through roughness was 0.12 bits (1.08 tokens). Summing the information transmitted by the individual parameters results in a total of 2.13 bits (4.38 tokens) transmitted by complete Tactons. The overestimate that would have occurred if the information transmission had been calculated directly from the confusion matrix can be seen in Appendix G on the accompanying CD, where the total information transmitted in Experiment 5(a) is calculated to be 2.55 bits (5.87 tokens).

The information transmitted through complete Tactons indicates that perfect performance could have been achieved if only four Tactons had been used rather than the 27 in this experiment, as 4.38 tokens

(2.13 bits) were transmitted. The information transmission results for rhythm and location suggest that near-perfect performance could have been achieved for both of these parameters if only two levels of each had been used. The number of bits of information transmitted through roughness was 0.12 bits (1.08 tokens). Since the corresponding number of tokens is less than two this indicates that it may not be possible to achieve near-perfect performance with two or more roughnesses. However, reducing the number of roughnesses to two should help to improve performance. This is explored in Section 8.4.

Comparison to Two-Dimensional Tactons Information Transmission results

After calculating the information transmitted by three-dimensional Tactons, the results for two-dimensional Tactons were re-examined and the information transmitted by two-dimensional Tactons was also calculated. The information transmitted by the two-dimensional Tactons with the C2 in the indented orientation was 2.52 tokens (1.34 bits) for rhythm, and 1.51 tokens (0.6 bits) for roughness, resulting in a total of 3.84 tokens (1.94 bits). In the raised orientation, these values were 2.5 tokens (1.32 bits) for rhythm, and 1.43 tokens (0.52 bits) for roughness, resulting in a total of 3.58 tokens (1.84 bits).

By comparing the results for three-dimensional Tactons to those for two-dimensional Tactons it can be seen that the overall information transmitted increased when a third dimension was encoded. Therefore, it can be concluded that it is beneficial to add a third parameter to Tactons. However, these results also show that the information transmitted through rhythm and roughness decreased, indicating that perception of these parameters is impaired by using a location other than the fingertip, or that the addition of a third dimension overloads the user’s memory. The decrease in performance on roughness is particularly noticeable, again indicating that it might be beneficial to reduce the number of levels of this parameter.

Analysis with Outliers Removed

The results of this experiment showed quite poor identification performance. However, closer inspection of the results indicated that this low average performance might be due, in part, to some outlier behaviour by several participants. In order to understand what performance levels might be achieved if such outliers were not present, the results for complete Tactons and individual parameters were examined, and any participant whose results on any of these were more than two standard deviations away from the mean was removed from the analysis. On this basis, three participants were removed.

With these three participants removed, the results for overall Tacton recognition showed an average recognition rate of 48%. This overall score is the same as when all participants were included, as is the score for roughness which is, again, 50%. However, the results for rhythm and spatial location improved. *Rhythms* were correctly recognised on average 96.7% of the time (compared to 90%) while location showed an average recognition rate of 95.5% (compared to 89%).

These outliers will also have affected the information transmission results, since it is calculated from confusion matrices. Therefore the information transmission scores were recalculated. For the individual parameters, the information transmitted through rhythm was 1.36 bits (2.6 tokens), through roughness was 0.12 bits (1.09 tokens), and through location was 1.28 bits (2.43 tokens).Summing these values results in overall information transmission for the complete set of Tactons of 2.76 bits (6.77 tokens).

With the outliers removed, the information transmitted through complete Tactons indicates that perfect performance could have been achieved if only 6 or 7 Tactons had been used rather than the 27 in this experiment. For the individual parameters, the results for rhythm and location suggest that near-perfect performance could have been achieved for both of these parameters if only two levels of each had been used, while the roughness results again indicate that it may not be possible to achieve near-perfect performance with two or more roughnesses. However, reducing the number of roughnesses to two should help to improve performance and this is explored in the following sections.

8.3.6 Discussion

The overall identification rate of 48% and the information transmission level of just 4.38tokens out of a possible 27 tokens is very low. This was improved somewhat by removing outliers from the analysis: the identification rate was still 48% but the information transmitted increased to 6.77 tokens as there was less confusion in the identification of locations and rhythms. However, the performance level is still low and, therefore, ways in which performance can be improved must be considered.

Reducing the amount of information encoded in the Tactons should reduce confusion and, therefore, increase the identification rate and the amount of information transmitted. Although the information transmission results for rhythm and spatial location when all participants are included would indicate that to achieve perfect performance only two levels could be used, the results with outliers removed indicate that more than two levels might be possible. The performance on roughness, however, indicates that three levels cannot be distinguished. Although the information transmission results indicate that it might not be possible to achieve perfect performance using two levels of roughness, reducing the number of roughnesses to two is a reasonable step towards improving performance whilst retaining a parameter with some use. Therefore, the next experiment aimed to investigate the effect of reducing the number of roughnesses to two.

Limitations of this study

The choice of roughnesses for this study was informed by studies conducted on the fingertip and therefore may not have been the optimal set for use on the forearm. However, by using the same roughnesses, the stimuli are consistent across all of the experiments and a direct comparison can be made between performance on the arm and the fingertip stimuli; changing the roughness parameter at the same time as adding an additional parameter would have introduced an additional confound. However better performance levels might have been achieved if roughness had been optimised for presentation on the arm. In addition, only one set of locations were tested. Therefore, it can be concluded that these locations are successful, but no firm conclusions can be drawn about the use of other locations.

8.4 Experiment 5(b): Reducing the complexity of Three-Dimensional Tactons

Based on the results of the experiment reported in Section 8.3, the number of roughnesses used in the three-dimensional Tactons was reduced from three to two, and a second experiment was carried out to investigate whether performance improved as a result.

8.4.1 Stimuli

The rhythms and spatial locations used in this experiment were exactly the same as those used in the previous experiment. The only change was to the roughness parameter, where only two levels were used instead of the three in the previous experiment. The smooth (un-modulated sine wave) and very rough (sine wave modulated by 30Hz) stimuli were used in this experiment, to represent low and high importance appointments. These two roughnesses were selected as they were at two ends of the scale – one smooth and one very rough – and should, therefore, be the two most perceptually different stimuli. In addition, the confusion matrix in Table 8-1 showed that there was less confusion between these two stimuli than between either of the adjacent pairs (smooth-rough or rough-very rough). Reducing the number of roughnesses to two resulted in a set of 18 possible Tactons.

The audio files for this experiment can be found in Appendix H on the accompanying CD by opening the html file “Exp 5(a) and (b) – 3D Tactons.html”.

8.4.2 Aim and Hypotheses

This experiment investigated whether performance improved when the number of roughnesses was reduced from three to two. In addition to improving the percentage correct score, it is possible that reducing the number of roughnesses could reduce confusion and, therefore, increase the amount of information transmitted, even though less information was input.

The hypotheses were that reducing the number of roughnesses from three to two would result in:

1. a significant improvement in the percentage correct score for complete Tactons.
2. a significant improvement in the percentage correct score for the roughness parameter.
3. an increase in the amount of information transmitted.
4. comparable performance on identification of rhythm and location to that found in the previous experiment.

8.4.3 Hardware Setup

The hardware setup was exactly the same as in the previous experiment.

8.4.4 Experimental Method

The experimental methodology for this experiment was replicated from the previous experiment. Any differences are outlined in the following sections.

Participants

Six participants who had taken part in the three-parameter experiment were asked to return to take part in this follow-on study. The participants were self selecting: they were the first six participants to respond to an email request for participants to take part in a follow-on study. This experiment took place within the two weeks following the original experiment. Participants were, once again, paid £6 for their participation. Five of the six participants were male, and one was female, and they were aged between 19 and 32 (mean=24). All were right handed. The results in this follow-on experiment can be compared back to the results of the same six participants in the original study for a within-groups comparison.

Training Procedure

The training procedure was replicated exactly from the previous experiment. Participants were again given ten minutes to familiarise themselves with the set of Tactons, but most participants chose to move on from this training after around two minutes as they were already familiar with the Tactons from the previous experiment. They then took part in 18 tasks like those in the experiment itself to refamiliarise themselves with the interface. The training instructions for this experiment are included in Appendix E.5.

Testing Procedure

The testing procedure was again identical to that used in the previous experiment. Since there were only 18 Tactons in this set, there were 54 tasks in the experiment (three repetitions of each Tacton). As in the previous experiment the order of presentation of the Tactons was randomised for each participant. The training instructions for this experimental procedure are included in Appendix E.6.

8.4.5 Results

Before reporting the results of this experiment it is important to examine the results of these six participants in the original experiment. In the previous experiment these six participants achieved a recognition rate of 50% for complete Tactons, 96.9% for rhythms, 51.4% for roughness and 99.8% for spatial location. The information transmitted was 1.38 bits (2.6 tokens) for rhythm, 0.15 bits (1.1 tokens) for roughness and 1.57 bits (2.96 tokens) for spatial location, resulting in overall information transmission of 3.1 bits (8.57 tokens) for complete Tactons The percentage correct scores are shown in Figure 8-11, alongside the results when the number of roughnesses was reduced.

It should be noted that the results of these six participants in the previous experiment are higher than the average scores achieved for that experiment. However, this is mainly due to the fact that none of these participants exhibited outlier behaviour in that experiment. When compared to the results with no outliers included the difference is smaller. These participants were not selected for their good performance, rather they were self-selecting. It is possible that the participants most motivated to take part in an additional experiment would be those who found the original experiment easy rather than those who struggled.

Complete Tactons

When the number of roughnesses was reduced to two the percentage correct score for complete Tactons increased from 50% to 80.56. These scores are shown in Figure 8-11 along with the percentage correct scores for the individual parameters. In addition, Figure 8-12 shows the scores for each participant, showing that performance improved for every participant, with subject 5 reaching a score of 100% correct in the 2-roughness condition. A one factor Wilcoxon test showed a significant improvement in performance (T=0 (n=6), p=0.036), therefore Hypothesis 1 can be accepted.

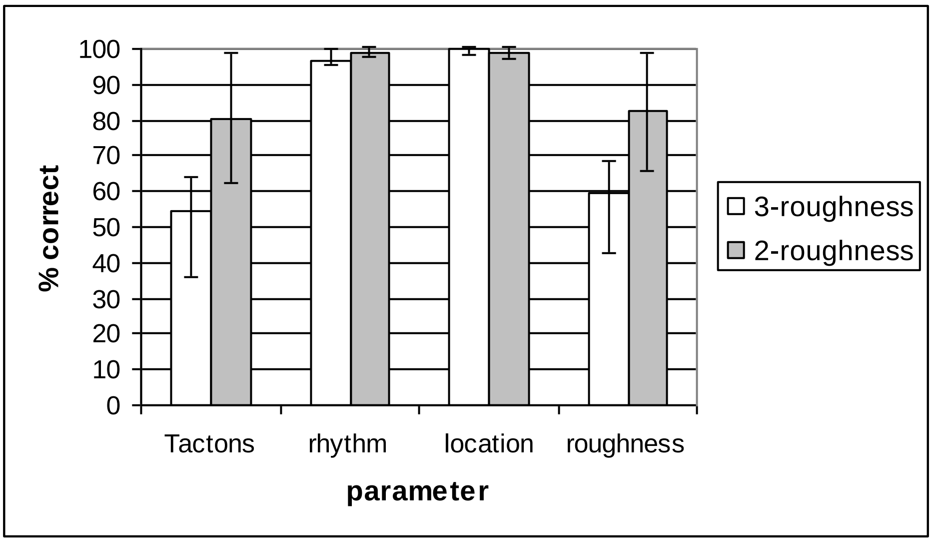

Figure 8-11: Experiment 5(b): Mean percentage correct scores (with standard deviations) for complete Tactons and individual parameters for the 6 participants in the 3-roughness (3rough) and 2-roughness (2rough) conditions.

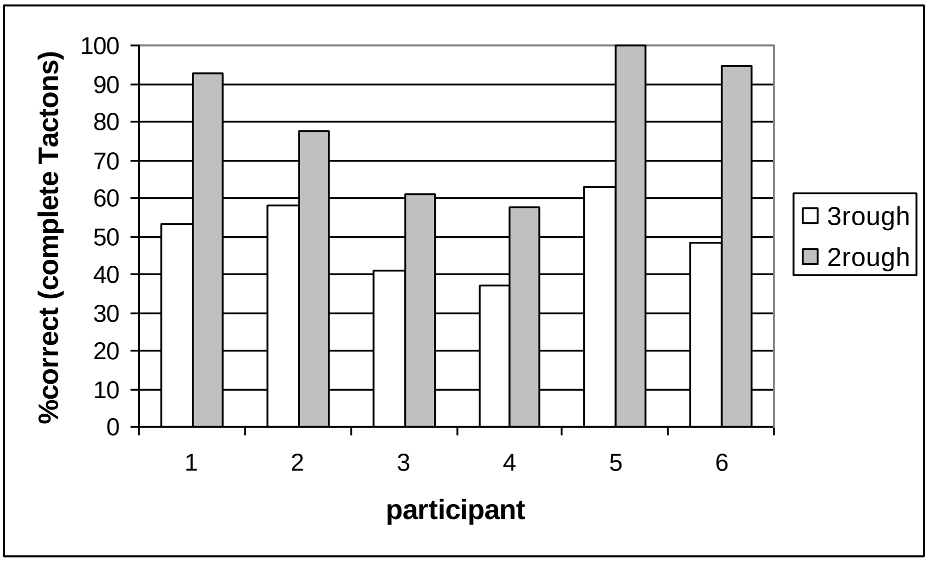

Figure 8-12: Experiment 5(b): Mean percentage correct scores (with standard deviations) for complete Tactons, for each individual participant in the 3-roughness (3rough) and 2-roughness (2rough) conditions.

Summing the information transmission results for the individual parameters (reported below) showed that the information transmitted by complete Tactons increased from 3.1 bits (8.57 tokens) to 3.35 bits (10.2 tokens) when the number of roughnesses was reduced to two. This result indicates that perfect performance could have been achieved with around 10 Tactons. It also shows that decreasing the number of Tactons from 27 to 18 actually increases the amount of information transmitted (from 3.1 bits (8.57 tokens) to 3.35 bits (10.2 tokens)) even though the amount of information input has been decreased (from 4.75 bits (27 tokens) to 4.17 bits (18 tokens)). This increase in information transmission indicates that Hypothesis 3 can be accepted.

Rhythm

Rhythms were correctly recognised on average 96.9% of the time in the three-roughness experiment and 99.1% of the time in the two-roughness experiment (Figure 8-11). A one-factor Wilcoxon test showed no significant difference in performance (T=5 (n=5), p=0.5). This result suggests that performance on rhythm identification is comparable in both experiments, and indicates that Hypothesis 4 can be accepted for rhythm. In the three-roughness experiment, the information transmitted by the rhythm parameter was 1.38 bits (2.6 tokens), while in the two-roughness experiment this increased slightly to 1.52 bits, corresponding to 2.86 tokens. This indicates that perfect performance could have been achieved if only two rhythms had been used, but that this result was close to perfect.

Roughness

The results for recognition of roughness showed an average recognition rate of 51.4% in the three-roughness experiment and 82.4% in the two-roughness experiment (Figure 8-11). A one-factor Wilcoxon test showed a significant difference in performance (T=0 (n=6), p=0.036) therefore Hypothesis 2 can be accepted. The graph in Figure 8-13 illustrates the performance in both conditions for each individual participant and shows that performance on roughness identification improved for every participant when the number of roughnesses was reduced to two.

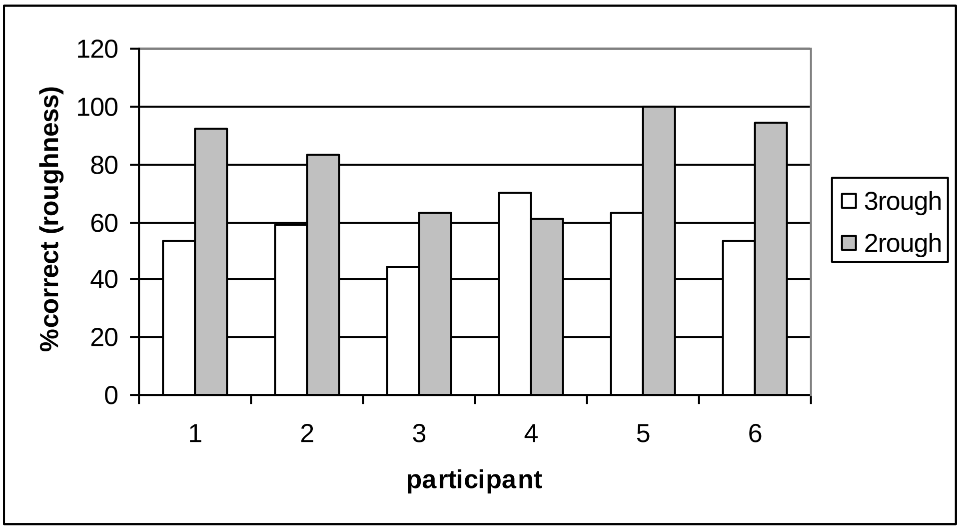

Figure 8-13: Experiment 5(b): Mean percentage correct scores (with standard deviations) for roughness, for each individual participant in the 3-roughness (3rough) and 2-roughness (2rough) conditions.

Information transmission was calculated for roughness at 0.15 bits in the three-roughness experiment (1.1 tokens) and increased to 0.33 bits (1.26 tokens) when the number of roughnesses was reduced to two, indicating that it is beneficial to reduce the number of roughnesses to two. As the number of tokens is still less than two, the information transmission results indicate that it may not be possible for users to reliably identify two or more levels of roughness and, therefore, roughness may not be a suitable parameter for Tactons if perfect performance is required. However, it may be suitable if recognition of around 80% is acceptable.

Spatial Location

The results for recognition of location showed an average recognition rate of 99.8% in the three-roughness experiment and 98.8% in the two-roughness experiment (Figure 8-11). A one-factor Wilcoxon test showed no significant difference between the results of these two experiments ($T=1.5$ ($n=3$), $p=0.59$), suggesting that the performance on location identification is comparable in both. This indicates that Hypothesis 4 can be accepted for location. Information transmission was calculated for this parameter at 1.57 bits (2.96 tokens) in the 3-roughness experiment and 1.5 bits (2.8 tokens) in the two-roughness experiment. This indicates that perfect performance could have been achieved if only 2 locations had been used but that, as with the rhythm parameter, this result was very close to perfect.

8.4.6 Discussion

The results of this follow-on study indicate that decreasing the number of roughnesses from three to two significantly improves performance for absolute identification of Tactons encoding three dimensions of information. Although less information is encoded in the Tactons, more information is actually transmitted and, therefore, it is beneficial to use this smaller set of Tactons.

The percentage correct rate for roughness increased significantly from 57.2% to 82.4% when the number of roughnesses was reduced to two, indicating that it is beneficial to use two, rather than three, levels of roughness in Tactons when combined with spatial location and rhythm. However, the performance for roughness is still much lower than that for the other two parameters therefore it might be worth considering whether a different parameter such as intensity (or a combination of roughness and intensity to provide redundancy) could achieve better performance. This is explored in Section 8.6.

Limitations of this study

One limitation of this study is that the same participants were used for both experiments, and the order of the experiments was not counterbalanced, and it is therefore possible that the increase in performance was due in part to the increased training received. In addition only six participants took part in this study so the results can only be used as an indication of what might be possible; a larger study is necessary to confirm these findings.

8.5 Redesigning the Parameter Space for Three-Dimensional Tactons

Another method of improving performance on three dimensional Tactons might be to use a parameter other than roughness, since this is the parameter that degrades performance. The discussion in Chapter 4 identified that intensity could be a usable parameter for Tacton design, and this parameter was used effectively in Chapter 7 for Tactons presented using a phone vibration motor. Therefore, the Tactons were redesigned using intensity in place of roughness, and this is discussed in Section 8.5.2. In addition, previous research has shown that when few values of a parameter can be identified, manipulating another parameter redundantly to encode the same information can increase the number of values that can be encoded. In fact for the experiment reported in Chapter 7 this occurred to some extent, as changing the intensity of the phone motor vibration also changed the frequency and thus both parameters encoded the same information. Section 8.5.3 reports the design of stimuli combining roughness and intensity redundantly to encode information.

Before either of these parameters could be designed it was necessary to establish the detection threshold for vibrations presented to the forearm via a C2 Tactor, to provide a level against which the intensity of the other stimuli could be measured. Section 8.5.1 reports an experiment which was conducted to establish this threshold.

8.5.1 Measurement of Detection Threshold

Before creating stimuli that used intensity as a parameter it was necessary to find the perceptual threshold for the C2 Tactors on the forearm. Identification of this level would enable suitable values of intensity to be chosen. In addition, since intensity levels are reported as decibels relative to the perceptual threshold, it is necessary to identify this threshold in order to report the levels such that these stimuli could be reproduced.

Several procedures exist for identifying detection thresholds; the procedure used here was the Method of Constant Stimuli [46]. In this procedure, a set of stimuli of different intensities are chosen, ranging from those which are seldom perceived to those which are almost always perceived. These stimuli are then repeatedly presented to a user, in a random order, and the user responds with "yes" if they perceived a stimulus or "no" if they did not. The point at which a stimulus is perceived 50% of the time is considered to be the detection threshold.

To obtain an accurate measure of detection threshold, this procedure should be carried out for every participant. However, in a real application this would not be feasible and, in fact, for the purposes of this experiment a precise value is not necessary. The main reason for establishing this threshold was to have an approximation of this level to provide a value against which the other stimuli could be measured. Therefore, this threshold was established through an experiment with six users.

To establish the threshold for this experimental setup a range of intensities were chosen, and were labeled accordingly to an arbitrarily chosen value, called zero. The intensity of the stimuli were changed relative to this zero value and ranged from -36dB to -14dB, with -36dB chosen as the lowest value as it was undetectable by the author, and -14dB chosen as the highest as it was always clearly detected. Within this range, there were 12 stimuli, each of them 2dB apart. All of these stimuli were 250Hz sine waves.

An experiment using the method of constant stimuli procedure was carried out with six participants (staff and students of the University of Glasgow), who were unpaid. A C2 Tactor was attached using a Velcro band (as in the previous experiments) to the mid-point of the participant's non-dominant forearm. This point on the arm was used since this is the point to which participants compared the intensity of the other two Tactors during the hardware set-up phase in both of the previous experiments described in this chapter (see Section 8.3.2). For every participant, each stimulus was presented 20 times, so the experiment consisted of 240 tasks. In each task, the screen shown in Figure 8-14 appeared and the participants could press the "Play Stimulus" button up to two times. Then they responded as to whether they had perceived the stimulus by clicking on the corresponding radio button. The instructions for this experiment are included in Appendix F.1.

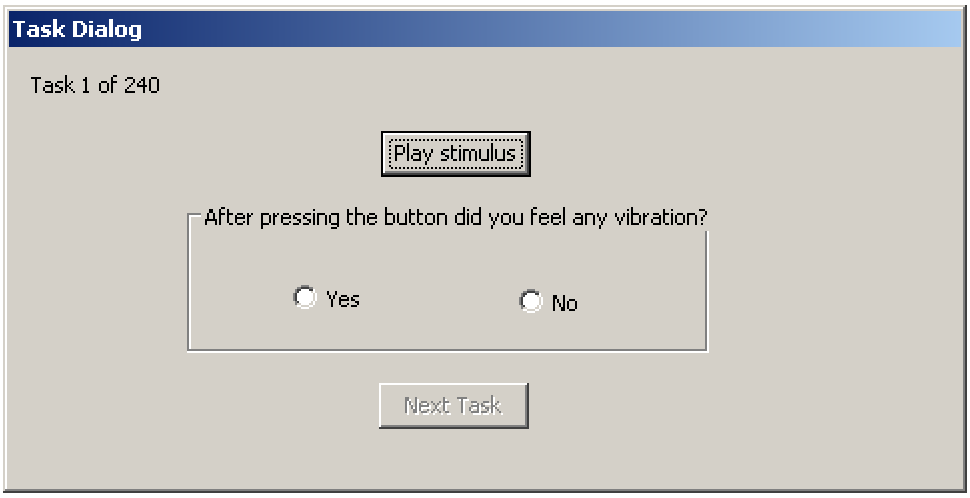

Figure 8-14: Measurement of Detection Threshold: response dialogue.

The results, showing how many times the users selected "Yes" for each stimulus, are shown in Table 8-2. It can be seen that the 50% detection threshold occurs between -26 dB and -24 dB. All that is required for this experiment is an estimate of the detection threshold, not an exact value, as the important factor was to ensure that all stimuli were detectable. Therefore, a conservative estimate of the threshold would be "-24 dB", and this was considered to be the threshold for the following studies. It should be noted again that "-24 dB" is a name given to this stimulus relative to an arbitrary zero value, and in itself has no meaning. From this point on the "-24 dB" stimulus will be renamed as "zero" and will be the threshold against which other stimuli are measured, thus 12 dB SL will mean that the intensity is 12 dB above this threshold value.

8.5.2 Design of Intensity Stimuli

Three intensity levels were required for this experiment. The important factors in selecting these intensity levels was that they were strong enough to be perceived, not so strong as to cause discomfort, and could be discriminated from each other. Pilot testing by the author and her colleagues enabled the selection of three values for intensity based on these criteria, and the levels selected to be used for these stimuli were 12dB SL, 16dB SL and 21dB SL. The lowest of these is well above the perception threshold and was felt to be clearly perceivable, while the highest was well below the limit (55dBSL) at which tactile stimuli become painful or unpleasant. In addition, the differences between them are all greater than the most conservative finding for difference thresholds (namely 2.3dB), and pilot tests indicated that these stimuli felt as if they were approximately equally spaced in terms of subjective magnitude (even though the actual spacings are not exactly equal). The audio files for these stimuli can be found on the accompanying CD by opening the html file “Exp 6 – Redesigned 3D Tactons.html”.

Level relative to arbitrary zero value(dB)	Yes-Score	proportion	Level relative to detection threshold (dB SL)
-36	7	0.06	-12
-34	14	0.12	-10
-32	21	0.18	-8
-30	21	0.18	-6
-28	31	0.26	-4
-26	48	0.40	-2
-24	81	0.68	0
-22	110	0.92	2
-20	119	0.99	4
-18	119	0.99	6
-16	120	1.00	8
-14	120	1.00	10

Table 8-2: Measurement of Detection Threshold: "Yes" responses to whether a stimulus could be perceived.

8.5.3 Design of “Roughness + Intensity” Stimuli

Redundancy is used in real life, for example when we process speech or recognise a face, many pieces of information are used redundantly to help us to interpret meaning. Research in tactile display has also shown that gains can be achieved through using redundancy. In particular, Sherrick [98] showed that while varying frequency alone only created three to five distinguishable stimuli, varying both intensity and frequency redundantly meant that eight levels could be distinguished. Geldard [41] also suggested that, where people are having difficulty interpreting vibrotactile signals, using redundancy may result in improved performance. Therefore, it may be beneficial to vary both intensity and roughness redundantly to encode the same information. These stimuli were created by adjusting the intensity of the existing roughness stimuli, as used in the previous experiment to the levels established for the intensity stimuli in Section 8.5.2. The smooth stimuli were presented at a level of 12dB SL, the rough stimulus at 16dB SL, and the very rough stimuli at 21dB SL. The audio files for these stimuli can be found in Appendix H on the accompanying CD by opening the html file “Exp 6 – Redesigned 3D Tactons.html”.

8.6 Experiment 6: Evaluation of Redesigned Three-Dimensional Tactons

The aim of this experiment was to investigate the effect on absolute identification of three-dimensional Tactons when the third parameter was changed from roughness to intensity, or to roughness plus intensity, and to investigate the effect of this change on identification of individual parameters.

8.6.1 Hypotheses

The hypotheses for this experiment were as follows:

1. Performance on the third parameter will improve when intensity is used in place of roughness.

2. Performance on the third parameter will improve when roughness+ intensity stimuli are used in place of roughness alone.

3. Performance on rhythm and spatial location identification will not be affected by the change in the third parameter.

4. Overall performance will improve when intensity or roughness+intensity are used in place of roughess.

8.6.2 Hardware setup

As in the previous experiments three C2 Tactors were attached to the participant's non-dominant forearm using Velcro bands (Figure 8-3) at the positions described in Section 8.2.3. The participant wore closed-back headphones to block out any noise from the Tactors and used their dominant hand to control the mouse. As in the previous experiments, participants were asked to adjust the levels of the three Tactors until the subjective magnitude felt equivalent on all three (see Section 8.3.2). This was particularly important for this experiment, where intensity was a parameter.

8.6.3 Experimental Method

The following sections describe the methodology used for this experiment.

Participants

Seventeen new participants took part in this study, aged 18 to 38 (mean=23). All were staff or students of the University of Glasgow and were paid £10 for their participation (the money was increased since this experiment consisted of two conditions and took more time). Nine of the participants were male and eight were female. All, but one, of the participants were right handed. The remaining participant said he was ambidextrous, but used the mouse in his right hand and, therefore, the Tactors were attached to his left arm (as they were for all the right handed participants). The consent form for this experiment is included in Appendix F.2.

This was a within-groups study with two conditions: the "intensity" condition and the "roughness plus intensity" condition. All participants took part in both conditions and the order was counterbalanced to take account of learning effects. In the analysis only 16 participants' results were analysed, with eight for each order of conditions, as one participant was removed from the analysis (see Section 8.6.4).

Training Procedure

The training procedure for each condition was exactly the same as for the previous three-dimensional Tactons experiments. Participants received a sheet explaining the mapping from stimuli to the Tactons they represented (see Appendices F3 and F4) and could then spend up to ten minutes trying out these Tactons for themselves. They then carried out 27 tasks from the experiment itself to get used to using the interface. This training occurred before both conditions to familiarise the participants with the set of

Tactons used in that condition and to give them practice using the interface for that particular set of Tactons.

Testing Procedure

The testing procedure was exactly the same as for the previous experiment. The only difference was that, in order to keep the experiment to a manageable length, the number of times each stimulus occurred was reduced to two, meaning that each condition consisted of 54 tasks. Although this is a small number of repetitions, each individual parameter occurred 27 times, so there are a large number of data points for each of these. The instruction sheet is included in Appendix F.5.

8.6.4 Results

During the experiment data were collected on the number of correct responses to the three parameters of each Tacton. These data are included in Appendix G on the accompanying CD. Percentage correct scores were calculated for each individual parameter and for the complete Tactons, and information transmission was also calculated from the stimulus-response confusion matrices [97]. One participant consistently confused the mapping from time to spatial location (consistently selected wrist for 30mins and elbow for 15mins instead of *vice versa*) and was removed from the analyses. Therefore, the results reported below are for the remaining 16 participants. The results for the removed participant can also be seen in Appendix G.

Percentage Correct Scores

For complete Tactons, the average percentage correct score was 54% in the Intensity condition and 49% in the Roughness+Intensity condition. The average percentage correct score for rhythm was 89% in the Intensity condition and 82% in the Roughness+Intensity condition. For identification of spatial location, average scores of 93% and 84% were achieved in the Intensity condition and the Roughness+Intensity condition respectively. These scores are shown in Figure 8-15 along with those from Experiment 5(a) (labeled “roughness”). In addition, Figure 8-16 shows the individual scores for each participant for the intensity condition and the roughness+intensity conditions. This graph shows that there were very large differences in individual performance, and for some participants there was a large difference in performance between conditions. Closer examination of the results for these participants show that this difference was not only due to a difference in performance between intensity and roughness+intensity but also due to a difference in performance on spatial location identification, indicating that there may be some interaction between these parameters.

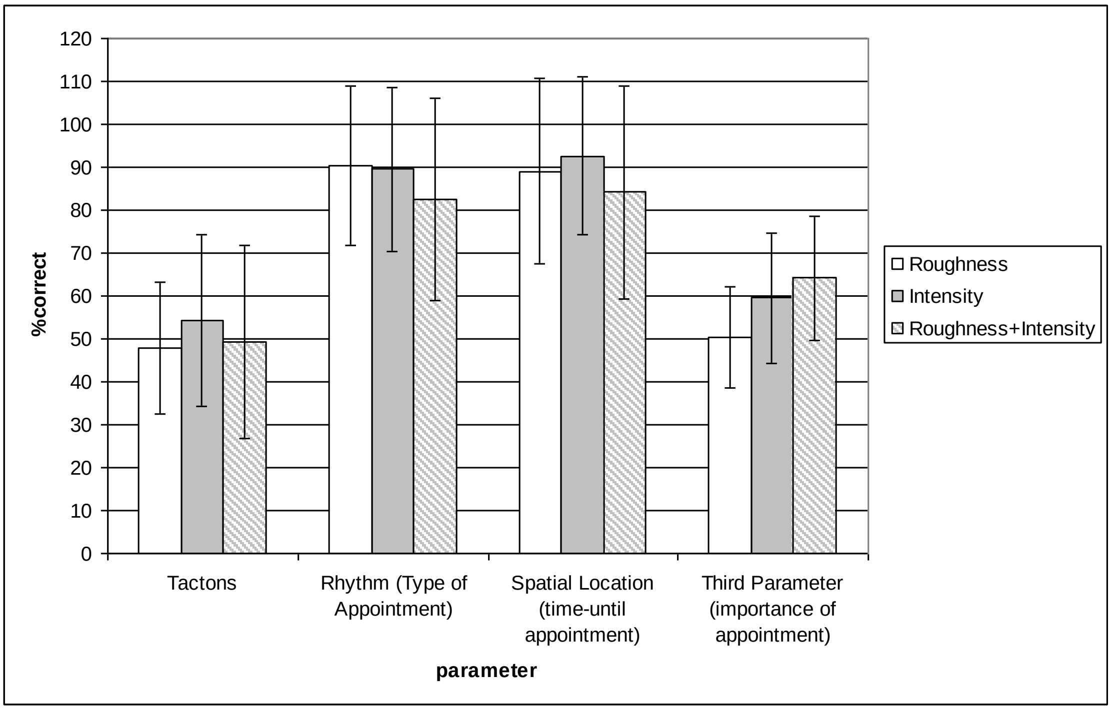

Figure 8-15: Experiment 6: Mean percentage correct scores (with standard deviations) for Tactons and individual parameters from the two conditions in this experiment (*Intensity*, *Roughness+Intensity*) and from Experiment 5(a)(*Roughness*).

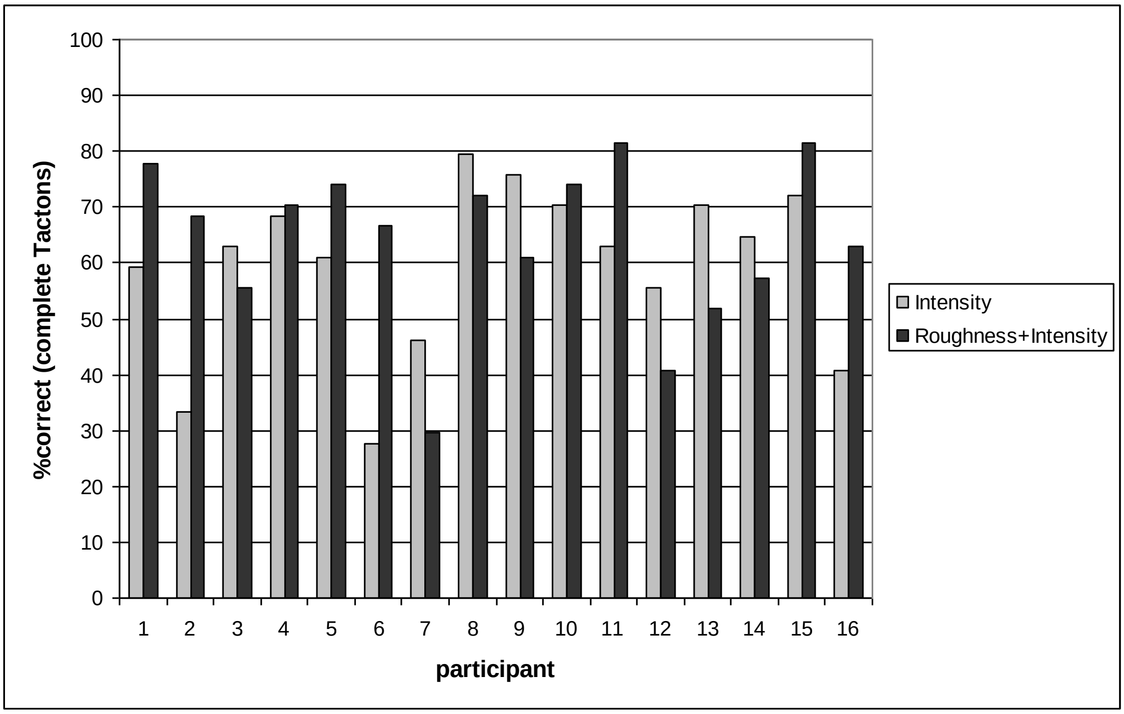

Figure 8-16: Experiment 6: Mean percentage correct scores for complete Tactons, by condition (Intensity, Roughness+Intensity), for each individual participant.

As the focus of this study was on improving performance on the third parameter, the results for this parameter (roughness/intensity/roughness+intensity) are examined in greater detail. The average percentage correct score for intensity was 59% while the average score for roughness+intensity was 64%.

Figure 8-17 and Figure 8-18 show the average scores for each of the three levels of each of these parameters. In addition, Table 8-3 Table 8-3: Experiment 6: Stimulus-response confusion matrix for intensity levels in the *Intensity* condition.and Table 8-4 show the confusion matrices for intensity and roughness+intensity respectively. In both conditions it can be seen that most confusions occur between adjacent levels with less confusions between the two extremes. This was also the case when roughness was used alone in the previous experiment, and indicates that just using low and high values of each of these parameters might also improve performance.

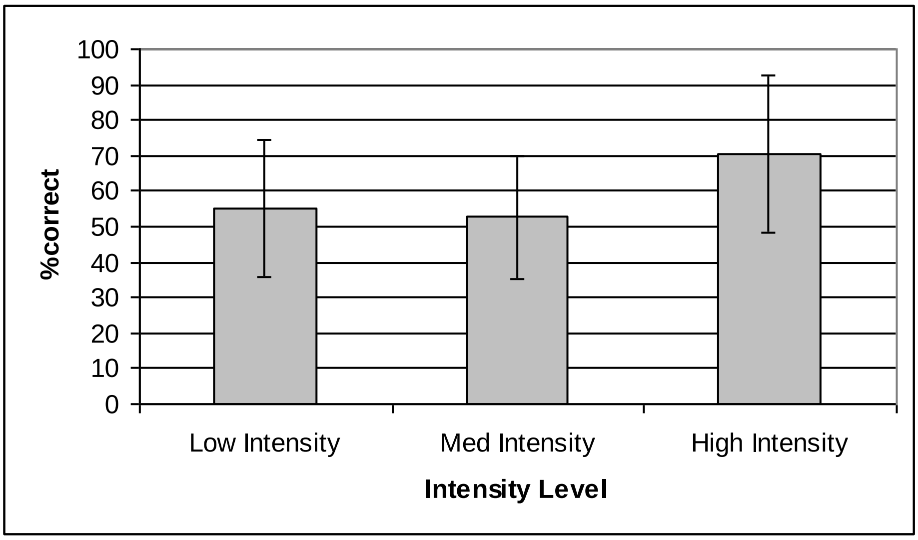

Figure 8-17: Experiment 6: Mean percentage correct scores (with standard deviations) for intensity identification in the *Intensity* condition.

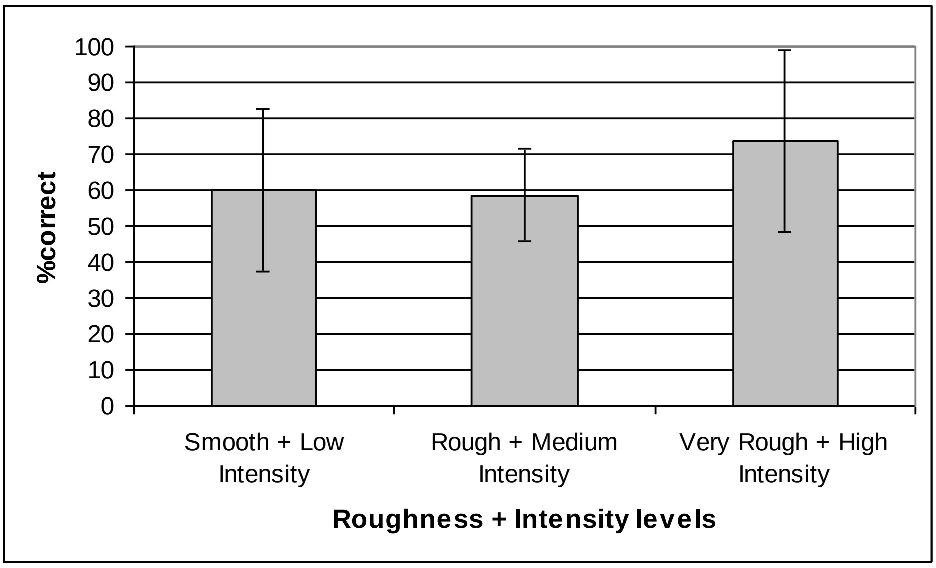

Figure 8-18: Experiment 6: Mean percentage correct scores (with standard deviations) for roughness+intensity identification in the Roughness+Intensity Condition.

Information Transmission Results

As in the previous experiment, the maximum information available from each of these sets of Tactons was 4.75 bits (27 tokens) while, for each individual parameter, a maximum of 1.58 bits (3 tokens) was available.

The results showed that the information transmitted through rhythm was 0.99 bits (2 tokens) in the intensity condition and 0.78 bits (1.71 tokens) in the roughness+ intensity condition. The information transmitted by spatial location was 1.13 bits (2.19 tokens) in the intensity condition and 0.80 bits (1.74 tokens) in the roughness+intensity condition. Finally, for the third parameter (roughness/roughness+intensity) the information transmitted in the intensity condition was 0.30 bits (1.23 tokens) and in the roughness+intensity condition was 0.35 bits (1.27 tokens).

Summing the information transmitted by the three parameters estimates the information transmission for complete Tactons to be 2.42 bits (5.35 tokens) in the intensity condition, and 1.93 bits (3.81 tokens) in the roughness+intensity condition.

response

stimulus		**low**	**med**	**high**
	low	55.21%	36.11%	8.68%
	med	14.93%	52.78%	32.29%
	high	3.82%	25.69%	70.49%

Table 8-3: Experiment 6: Stimulus-response confusion matrix for intensity levels in the *Intensity* condition.

response

stimulus		**smooth**	**rough**	**very rough**
	smooth	60.07%	32.99%	6.94%
	rough	20.14%	58.68%	21.18%
	very rough	7.29%	19.10%	73.61%

Table 8-4: Experiment 6: Stimulus-response confusion matrix for roughness+intensity levels in the *Roughness+Intensity* condition.

The information transmitted through complete Tactons indicates that, as in the previous experiment, perfect performance could have been achieved if only 4 or 5 Tactons had been used rather than the 27 in this experiment, as 3.81 - 5.35 tokens were transmitted. Again, as in the previous experiment the information transmission results for rhythm and location suggest that near-perfect performance could have been achieved for both of these parameters if only two levels of each had been used. The number of tokens transmitted through intensity was 1.23, and through roughness+intensity was 1.27. This again indicates that it may not be possible to achieve near-perfect performance with two or more levels of either of these parameters.

Comparison across conditions

When considering the results of this experiment it is possible to compare these two sets of data with the data from the experiment reported in Section 8.3. That experiment can be considered to be a third condition of the current experiment, and the conditions called be labeled as the *Roughness* condition (from Experiment 5(a)), the *Intensity* condition and the *Roughness+Intensity* condition. The average identification rates for complete Tactons and for the individual parameter for all three conditions are shown in Figure 8-15. It should be noted that all participants (including outliers) are included in the analysis for all conditions.

Statistical analyses across the data sets from these three conditions were carried out using Kruskal Wallis tests, since some of the comparisons were between-groups comparisons. Non-parametric statistics were used, as before, because the data were not normally distributed. Where statistically significant differences were found, post hoc Mann-Whitney tests were carried out, using a Bonferroni correction to compensate for carrying out repeated analyses of the same data. As there were tied ranks in these data the p-values reported are those which are adjusted for tied ranks.

No significant differences were found in the percentage correct scores for complete Tactons ($H = 3.09$, $DF = 2$, $p = 0.214$) between the three conditions. In addition there were no significant differences in the scores for rhythm ($H = 3.09$, $DF = 2$, $p = 0.214$) or spatial location ($H = 1.23$, $DF = 2$, $p = 0.542$). A significant difference was found for the third parameter between conditions ($H = 7.97$, $DF = 2$, $p = 0.019$). Post hoc Mann-Whitney tests with a Bonferroni correction (resulting in a threshold p-value of 0.016 for significance) showed that this difference in performance was due to a statistically significant difference between the roughness condition and the roughness+intensity condition ($p=0.007$). No significant differences were found between the roughness condition and the intensity condition (0.052), or between the intensity condition and the roughness+intensity condition ($p=0.37$).

8.6.5 Discussion

The aim of this experiment was to investigate whether performance on the third parameter could be improved by changing the vibrotactile parameter used to encode this information. Hypothesis 1 stated that performance on this parameter would improve when intensity was used in place of roughness. This hypothesis cannot be accepted as the results showed no significant difference in performance between these two conditions (48% in roughness condition, 59% in intensity condition). Hypothesis 2 stated that performance on this parameter would improve when roughness + intensity stimuli were used in place of roughness alone. This hypothesis can be accepted as identification was significantly better in the roughness + intensity condition (64%) than in the roughness condition (48%). Hypothesis 3 stated that performance on spatial location and rhythm would be unaffected by the change in the third parameter. While Figure 8-15 shows that performance on both of these parameters was lowest in the roughness + intensity condition, no significant differences were found and therefore this hypothesis can be accepted.

Although performance on the importance parameter improved as a result of changing the parameter used to encode this, this improvement did not result in any significant increase in the overall identification rates for Tactons. Hypothesis 4 must, therefore, be rejected. Figure 8-15 shows that performance on rhythm and spatial location was lower in the roughness+intensity condition than in either of the other conditions. While this difference was not significant, this decrease in performance may have been sufficient to offset the improvement in performance on the third parameter.

The average percentage correct score for complete Tactons was 54% in the Intensity condition and 49% in the roughness+intensity condition. These scores are very low, and the information transmission results indicate that to achieve perfect performance the number of Tactons would need to be reduced to just four or five. The information transmission results for rhythm and spatial location indicate that perfect performance could be achieved with around two levels of each of these parameters, while the results for intensity and roughness+intensity indicate that it may not be possible to achieve near-perfect performance with even two levels. However, the experiment reported in Section 8.4 showed that performance could be improved significantly by reducing the number of values of roughness from three to two. As the current experiment has shown that performance on roughness+intensity is significantly better than performance on roughness alone, it is likely that even better performance could be achieved if the experiment was re-run using just two levels of the roughness+intensity parameter. However, the impact of this parameter on the performance on complete Tacton identification might need to be further investigated.

Limitations of this study

One limitation of this study is that the levels of intensity were selected from the results of informal pilot studies, rather from the results of formal studies. It is therefore possible that a more optimal set of intensity levels might have been available and that, by using these, the performance level may have improved. However, it had been reported in related work on vibrotactile display that intensity judgements can be very difficult [41] and that people are often only able to distinguish two levels: low and high intensity [24]. Therefore, it is unsurprising that the identification rate was quite low. Reducing the number of levels of intensity to two should improve this. The results for roughness+intensity showed higher performance levels, but the same problems may apply to this and it is likely, once again, that performance would be improved by reducing the number of levels to two.

Geldard [41] noted that a fixed amplitude applied to different locations can feel very different at the different locations. This could also account for the poor performance levels. This should have been overcome by the level setting at the start of the experiment as participants should have set all of the outputs to feel the same. However, accurate level setting can be difficult for novice users, and therefore a more formal procedure for this setting might be required to ensure that the stimuli feel the same at all locations. Another possibility is that the subjective magnitude of the three intensity levels could have been affected by adaptation (see Chapter 2), and that this may have affected identification. Further investigation is necessary to understand if this was the case.

8.7 Discussion of Three-Dimensional Tactons results

Three experiments investigating the identification of three-dimensional Tactons have been reported in this chapter. These Tactons were been created by adding spatial location as a parameter in addition to the two parameters, roughness and rhythm, used in the Two-Dimensional Tactons reported in Chapter 6. In addition, the use of different parameters in place of roughness has been explored. The results of these experiments have shown that the overall identification rates for three-dimensional Tactons are generally much lower than those for two-dimensional Tactons when three values of each parameter are used. The identification rates ranged from 48% to 54%, depending on the parameters used. This identification rate increased to 81% when the number of levels of roughness was reduced to two, and doing this also increased the amount of information transmitted, indicating that it is beneficial to use this smaller set of Tactons.

In terms of individual parameters, the results of these experiments have shown that spatial location is successful as a Tacton parameter, with identification rates of 84%-99% achieved for this parameter, depending on the other parameters used in the Tactons. Rhythm transferred successfully to being used in three-dimensional Tactons with the results for this parameter ranging from 82%-99%, again depending on the other parameters. Roughness did not transfer so successfully to three-dimensional Tactons, with an identification rate of just 50% achieved when three levels were used. The experiments showed that this rate can be significantly improved by reducing the number of roughnesses to two (81%) or by manipulating intensity at the same time to encode the same information redundantly (62%).

The maximum information transmitted by three-dimensional Tactons was 3.37 bits (10.3 tokens). In Rabinowitz *et al.*'s work on multidimensional tactile display [89], around 3.9 bits (15.9 tokens) were transmitted by three-dimensional displays. While these results are much higher than those achieved for Tactons, their research used specialised technology in which the contactor size could be altered, and a high level of control over intensity and frequency could be achieved. This research on Tactons uses small, commercially available actuators which could be used in wearable devices and, given these limitations, the decrease in performance seems understandable. In addition, future work will investigate how the design of Tactons could be improved and the amount of information transmitted may then increase.

8.7.1 Comparison to Earcons

One of the motivations for investigating three-dimensional Tactons was that research on Earcons has shown that Earcons encoding three dimensions of information can be identified, and this research aimed to make a comparison to those results. McGookin's evaluation of Three-Dimensional Transformational Earcons [78] (with three levels of each parameter) showed overall recognition of around 70%, with individual recognition rates of over 90% for melody, over 90% for timbre and 75% for register. With the same number of levels of each parameter, the overall recognition rate for Three-Dimensional Tactons is much lower, with the highest level achieved being 54%. There is less difference in the individual parameters, where rhythm achieved scores of 82-99% and spatial location achieved scores of 84-99%. The parameter which reduces performance is the third parameter: roughness, intensity or roughness+intensity. The highest level achieved for this parameter was 64% when both roughness and intensity were manipulated. This level is a lot lower than the 75% achieved for register in Earcons.

While the results for Tactons are much lower than for Earcons this result is actually unsurprising, given that the resolution of the tactile sense is much lower than the available resolution in audio. By reducing the number of levels of the importance parameter, an overall identification rate of 81% can be achieved, with 99% identification of rhythm and spatial location, and 82% identification of roughness. Although a smaller set of messages are available, this result compares favourably to the results achieved for Earcons, and perhaps it is unreasonable to expect the same level of performance to be achievable when a more limited sensory modality, namely touch, is used.

8.7.2 Design Recommendations for Designing Three-Dimensional Tactons

From the three studies reported in this chapter it is possible to make a number of design recommendations for designing three dimensional Tactons. As in the previous chapters, it should be noted that these are recommendations rather than definitive guidelines, since these recommendations are based only on these studies, and better performance may be achieved if different parameters, which were not explored in these studies, were to be used.

1) ***When designing three-dimensional Tactons, encode the two most important dimensions in rhythm and spatial location.***

   When three rhythms were used, average identification rates of up to 99% were achieved for this parameter. When three Tactors were located on the volar forearm average identification rates of up to 99% were achieved for spatial location.

2) ***When using spatial location, locating two Tactors at anatomical landmarks, and one between them can be successful.***

   When three Tactors were located on the volar forearm with one at the wrist, one at the elbow and one equidistant between these two, average identification rates of up to 99% were achieved.

3) ***A less important dimension can be encoded in roughness, intensity or a combination of both.***

   a) ***If three levels of a third parameter are needed, roughness and intensity should be combined redundantly to encode this information, however doing so may impair recognition og other parameters.***

      The results showed that varying roughness and intensity redundantly resulted in significantly better performance than varying roughness or intensity alone. However, the identification rate for this parameter was still low, at just 64% and the use of this parameter also seemed to impair identification of the other parameters, indicating that there is some interaction between them.

   b) ***Where possible only two levels should be used.***

      The results showed that identification rates of 81% can be achieved when two levels of roughness are used. This may be improved by using two levels of roughness+intensity rather than just roughness alone.

8.8 Conclusions

This chapter has reported three experiments investigating the identification of Tactons encoding three dimensions of information. The results from these three experiments can be used to answer the two research questions posed at the start of this thesis:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

In relation to Research Question 1, both experiments have indicated that rhythm and spatial location can be used to encode information in Tactons. Identification rates of 82-99% have been achieved for these parameters. In addition, the maximum information transmitted by each of these parameters was around 1.5 bits, indicating that around 2.8 of the 3 levels of these parameters could be perfectly identified. In addition the experiments have investigated the use of various third parameters: roughness alone, intensity alone and roughness and intensity combined. The results have shown that the most successful of these parameters is the combination of both, but that the identification rate for this parameter is still low, at just 64%. The results of Experiment 5(b), which reduced the number of roughnesses to two, showed that performance levels of around 82.4% can be achieved. If the roughness+intensity parameter was used in place of roughness alone this result might be even better.

In relation to Research Question 2, these experiments have shown that an identification rate of 81% can be achieved, and 3.35 bits (10.2 tokens) of information transmitted, when three dimensions of information are encoded in Tactons, on the condition that only two levels of roughness are used. When three levels of the importance parameter are used the best result achieved was 54% correct, when the importance was encoded in the intensity (although there were no significant differences between the scores for any of the conditions). In terms of information transmission, 2.42 bits (5.35 tokens) were transmitted in this condition. Since better results on identification of importance levels were achieved for both intensity and roughness+intensity than for roughness alone, reducing the number of levels of either of these parameters to just two might result in higher performance than two levels of roughness alone. The 81% performance level achieved with two levels of roughness exceeds the 73% identification rate achieved for Tactons encoding two dimensions of information (see Chapter 6), indicating that encoding three dimensions of information in Tactons can be successful.

These experiments have shown that it is possible to encode three dimensions of information in Tactons by manipulating rhythm, spatial location and a third parameter: roughness, intensity or a combination of both, with varied results. The results ,when three levels of the third parameter are used, are quite low and, therefore, it could be difficult to use these Tactons in applications requiring high identification rates. Reducing the number of levels of the third parameter improves performance significantly, increasing both the percentage correct rate and the amount of information transmitted. There may also be other solutions: a number of alternatives are proposed in the following chapter, along with a summary of the work in this dissertation and conclusions regarding this research as a whole.

Chapter 9: Discussion & Conclusions

9.1 Introduction

This thesis has investigated the design of vibrotactile icons, called Tactons. While vibrotactile icons have been proposed, or studied, by other researchers [22, 23, 72, 95, 111], there have been very few formal studies of how best to design them. This research is the first formal investigation into the design of Tactons: a specific type of vibrotactile icon, designed using a structured, abstract approach [8].

The aim of this research was to investigate how to design Tactons to encode multidimensional information, and to understand the identification rates which can be achieved when these Tactons were presented to users. In Chapter 1, the thesis statement was as follows:

> *"This thesis asserts that multi-dimensional information can be encoded in structured vibrotactile messages called Tactons by manipulating parameters of vibration, and that users can learn the mapping between Tactons and their meanings and retrieve information from them."*

This statement has been defended by answering the following two research questions:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?
>
> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

These two questions have been addressed through a review of the literature on tactile perception and vibrotactile display, and by a series of empirical experiments evaluating individual parameters, and Tactons created by combining these parameters. This chapter summarises the work reported in this thesis and discusses how the findings answer the two research questions above. It then sets out a series of design recommendations derived from this research, which could be employed by designers or researchers who wish to use Tactons in their work. Limitations of this research are then identified and areas for future work are described. Finally, some general conclusions are drawn from this research, and the main contributions of this work are stated.

9.2 Summary

Chapter 2 provided a review of tactile perception, concentrating on the perception of the parameters of vibration which could be manipulated to encode information in Tactons. This chapter identified a number of candidate parameters which might be able to be used in Tactons: spatial location, rhythm, duration and intensity. In addition, this chapter discussed the devices which could be used to present Tactons and selected two for use in this research.

Chapter 3 reviewed related work on vibrotactile display. This review showed that, while tactile display has been investigated for many years, much of this work has used simple vibration stimuli, rather than vibration stimuli encoding multi-dimensional information. In addition, people have only recently started to investigate how to design vibrotactile icons. Therefore, this thesis focused on the design of vibrotactile icons encoding multidimensional information.

Chapter 4 discussed different strategies for encoding information in vibrotactile icons. One strategy, which had received little attention in previous work, was the use of a structured, abstract approach to encode information in vibrotactile icons, like the approach used in Earcons in audio. This approach was, therefore, used in this thesis to design Tactons, which were defined as structured, abstract vibrotactile messages which can be used to present multidimensional information.

Chapter 5 reported two experiments investigating the perception of vibrotactile parameters which had not previously been formally investigated: waveform and intensity change over time. The results of Experiment 1 showed that different vibrotactile waveforms could be created by using amplitude modulation to create the sensation of roughness, and that up to four levels of this parameter could be distinguished from one another. In Experiment 2, it was shown that people are able to absolutely identify at least three different types of intensity change: increases, decreases, and level stimuli.

In Chapter 6, two-dimensional Tactons were created using three rhythms and three levels of roughness to encode information and were presented to the fingertip. Two experiments were conducted evaluating these Tactons presented via the C2 Tactor used in two different orientations: "indented" and "raised". The results of these experiments indicated that Tactons could be identified by users, and that this new method of interaction showed promise. These evaluations of two-dimensional Tactons were the first experiments investigating identification of Tactons and, therefore, provided a baseline against which future research on Tactons could be compared.

Having established a baseline identification rate for two-dimensional Tactons on a high specification vibrotactile actuator, Chapter 7 investigated how these Tactons could be transferred to a standard phone vibration motor. Two sets of two-dimensional Tactons were created: both sets used rhythm to encode one dimension of information while the other dimension was encoded either in roughness, or in intensity of vibration. When intensity was used to encode the second dimension, the results were comparable to those achieved with the C2 Tactor. It had previously been suggested that these devices were not suitable for the presentation of complex vibrotactile messages, but the results of this experiment indicate that, with the selection of appropriate parameters, it is possible to create effective Tactons for presentation on phone motors.

Some applications might require that more than two dimensions of information be presented via Tactons. Chapter 8, therefore, extended the work reported in Chapter 6 by encoding a third dimension of information in Tactons using spatial location. The results showed that adding spatial location degraded perception of roughness and, therefore, two further experiments were carried out to investigate whether performance could be improved. These experiments showed that the greatest improvement in performance could be achieved by reducing the number of levels of roughness from three to two. Replacing the roughness parameter with intensity, or manipulating both roughness and intensity redundantly, improved identification on that parameter, but did not result in any significant improvement in performance overall.

The following sections consider how the results from the practical work in this thesis can answer the two research questions posed by this thesis.

9.2.1 Research Question 1

Research question 1 asked:

> **RQ1.** What are the parameters of vibration that can be manipulated to encode information in Tactons?

Research Question 1 (RQ1) was addressed through a review of the literature on perception of vibrations, and of related work on vibrotactile display in Chapters 2-4, and by experiments investigating perception of individual parameters reported in Chapter 5. In addition, the experiments in Chapters 6, 7 and 8 evaluated identification of Tactons, which were created by combining two or three of these individual parameters. The results of these experiments provide identification rates for the individual parameters and, therefore, can also be used to answer this question. It should be noted, however, that these identification rates are somewhat conservative since the parameters are combined with others rather than assessed in isolation (thus placing greater memory demands on users and possibly introducing interactions between parameters). However, they do provide realistic figures for what can be achieved when these parameters are used in Tactons.

The candidate parameters established in the literature reviews in Chapters 2-4 were rhythm, spatial location and intensity, and it was also found that waveform and intensity change over time had not been formally investigated. The findings from the practical work investigating each of these parameters are reported below.

Rhythm

Tactile rhythm perception had previously been studied by a few researchers [2, 19, 65], but had mainly been concerned with the reproduction of these rhythms rather than the absolute identification of them. The research in this thesis has shown that rhythm can be used successfully as a parameter for Tactons. Identification rates of over 90% were achieved for a set of three rhythms presented via the C2 Tactor to the fingertip (Experiments 3(a) and (b)) and to the forearm (Experiments 5(a) and (b)), and for the same three rhythms presented via a phone motor (Experiment 4). In addition, information transmission rates of up to 1.52 bits (2.86 tokens) have been achieved for this parameter (Experiment 5(b)). This result indicates that reducing the number of rhythms to two should result in perfect identification, but that identification is already close to perfect (since there were three rhythms in the set). All of these results indicate that rhythm is a successful parameter for Tactons which can be transferred successfully between different body locations and different devices.

Spatial Location

Spatial location had previously been used by many researchers [45, 58, 111, 114], but was generally used in one-dimensional tactile displays, rather than in combination with other parameters. The results of this work have shown that spatial location can be used successfully as a parameter for Tactons. The average identification rates for spatial location (when three locations on the volar forearm were used) ranged from 84%-99%, depending on the other parameters used along with it (Experiments 5(a), 5(b) and 6). The maximum information transmitted by this parameter (in Experiment 5(b)) was 1.5 bits (2.8 tokens). As with the rhythm parameter, this result indicates that performance is near perfect, and that reducing the number of values of this parameter to two should result in perfect performance. These results suggest that the use of three spatial locations on the forearm is a successful means of encoding information in Tactons.

Waveform

Vibrotactile waveform had previously received little investigation. The work in this thesis has shown that distinguishable waveforms can be created by using amplitude modulation to create the sensation of roughness and that up to four different levels of roughness can be distinguished. When three levels of roughness were used, in combination with rhythm, to create two-dimensional Tactons presented via the C2 Tactor, average identification rates of up to 79.5% were achieved (Experiment 3(a)). However, when the C2 was used in the “raised” orientation rather than the “indented” orientation, the identification rate dropped to 69% (Experiment 3(b)), indicating that roughness perception is affected by the nature of the vibrotactile stimulation. Performance is lower when roughness is used as a parameter in 3D Tactons (50%), but can be improved by using just two levels of roughness (82%) or combining roughness and intensity to encode the same information (64%). In Experiment 4, roughness was also simulated on the phone vibration motor, but performance was significantly lower on the phone motor than on the C2 Tactor (55%). These results indicate that roughness can be used in Tactons but that it is less successful than rhythm and spatial location and does not transfer as easily between devices.

Intensity

The results from the review of tactile perception indicated that intensity could be a successful parameter, but that it should be used with care as it could interact with other parameters. In Experiment 4, intensity was used to encode one dimension of the two-dimensional Tactons for presentation via a phone vibration motor. The results showed that intensity was more suitable than roughness for use on phone vibration motors, with an identification rate of 75% achieved for this parameter, compared to 55% for roughness. Changing the intensity on a phone motor also changes the frequency of the vibration and this, therefore, provides redundant coding which may have improved identification. Intensity was also used in Experiment 6, to encode one dimension of the three-dimensional Tactons, in place of roughness. The results showed that using three levels of intensity resulted in slightly better performance than using the same number of roughnesses (59% compared to 55%) but there was no statistically significant difference between them. The best performance was achieved when roughness and intensity were combined redundantly to encode the same information (64%). These results suggest that intensity can be used in Tacton design, but is most successful when combined redundantly with another parameter.

Intensity Change Over Time

Intensity change over time was investigated and it was shown that very high identification rates can be achieved. The results showed that, when users were asked to identify whether the intensity of a stimulus increased, decreased or did not change, average identification rates of up to 100% could be achieved for increases, 92% for decreases and 95% for level stimuli. This parameter was not tested in combination with other parameters to create Tactons and, therefore, its actual success as a Tacton parameter cannot be confirmed. However, the high identification rates indicate that this is a promising parameter, which should be investigated further.

Conclusions regarding Research Question 1

From these results, it can be concluded that, when using a device like the C2 Tactor, the most successful parameters for encoding information in Tactons are rhythm and spatial location. Roughness or intensity can also be used, but it might be more beneficial to use just two, rather than three, levels of these parameters. Combining these two parameters redundantly also improves identification over using roughness alone. Intensity change over time appears to be very successful, but has not been tested in combination with other parameters and should, therefore, be investigated in future work. When using a standard mobile phone vibration motor, the most successful parameters appear to be rhythm and intensity.

9.2.2 Research Question 2

Research question 2 asked

> **RQ2:** What levels of performance can be achieved when these parameters are combined to create structured vibrotactile messages (Tactons) to present multi-dimensional information?

Research Question 2 (RQ2) has been addressed through the experiments in Chapters 6, 7 and 8. These chapters reported experiments investigating the performance levels achieved for absolute identification of Tactons encoding two or three dimensions of information. The following sections discuss how the results for two-dimensional and three-dimensional Tactons can answer this question.

Two-Dimensional Tactons

The results for the experiments investigating identification of two-dimensional (2D) Tactons are shown in Table 9-1. These results show that the best performance achieved for 2D Tactons was 73%, which was achieved in Experiment 3(a), when the C2 Tactor was used in the indented orientation with the two dimensions encoded in rhythm and roughness. The result was significantly lower, at just 66%, when the C2 was used in the raised orientation in Experiment 3(b). The information transmission results for the two dimensional Tactons presented via the C2 Tactor are also presented in this table. These results indicate that perfect performance could be achieved if a set of just three or four Tactons were used in place of the nine used here.

It is interesting to note that the best performance achieved in Experiment 4, using a phone motor (when rhythm and intensity are used) is almost identical to the best performance on the C2 Tactor at 72%. The performance on the phone motor when roughness is used, however, is significantly lower, at just 52%. This shows that the transfer of Tactons between different devices is possible, as long as appropriate parameters are chosen. This is a very promising result in terms of the commercial feasibility of Tactons, since these motors are common in many devices (e.g. mobile phones, PDAs, games controllers).

Experiment	**3(a)**	**3(b)**	**4** **(condition 1)**	**4** **(condition 2)**
Type of Tacton	**2D**	**2D**	**2D**	**2D**
Device	*C2 (indented)* *on fingertip*	*C2 (raised)* *on fingertip*	*phone motor* *in hand*	*phone motor in hand*
Parameters (number of levels)	*rhythm (3), roughness (3)*	*rhythm (3), roughness (3)*	*rhythm (3), roughness (3)*	*rhythm (3), intensity (3)*
%correct	73%	66%	52%	72%
IT (bits)	1.94	1.84		
IT (tokens)	3.84	3.58		

Table 9-1: Overall Tacton identification results from all two-dimensional (2D) Tactons experiments.

Three-Dimensional Tactons

The results for the three experiments investigating identification of three-dimensional (3D) Tactons are shown in Table 9-1. These results show that, while more information is transmitted, identification rates for 3D Tactons are generally much lower than those achieved for 2D Tactons. When three levels of each parameter (rhythm, spatial location and a third parameter: either roughness, intensity or a combination of both) were used, identification rates ranged from 48% to 54% (Experiments 5(a) and 6). The best performance was achieved when intensity was used as the third parameter (Experiment 6, condition 1), but the difference was not statistically significant. Experiment 5(b) indicated that reducing the number of levels of roughness to two can result in a significant increase in performance. This experiment resulted in an identification rate of 81%: the highest identification rate achieved for 3D Tactons. In addition, the amount of information transmitted by the complete Tactons increased from 2.13 bits (4.38 tokens) to 3.35 bits (10.2 tokens), even though the amount of information in the stimulus set had been decreased (reduced from 27 Tactons to 18 Tactons).

The maximum information transmitted by 3D Tactons was 3.35 bits (10.2 tokens) in Experiment 5(b). This result indicates that perfect performance could be achieved if a set of just 10 Tactons were used. Further investigation is needed to evaluate a set of this size, and to see whether this number can be improved through further refinements of the Tacton design.

Experiment	**5(a)**	**5(b)**	**6 (condition 1)**	**6 (condition 2)**
Parameters (number of levels)	*rhythm (3), roughness (3), location (3)*	*rhythm (3), roughness (2), location (3)*	*rhythm (3), intensity (3), location (3)*	*rhythm (3), roughness+intensity (3), location (3)*
%correct	48%	81%	54%	49%
IT (bits)	2.13	3.35	2.42	1.93
IT (tokens)	4.38	10.2	5.35	3.81

Table 9-2: Overall Tacton identification results from all three-dimensional (3D) Tactons

9.3 Design Recommendations for Tactons

In addition to answering the two research questions posed in the introduction, another contribution of this thesis is the production of a set of recommendations to aid designers who wish to use Tactons in their interfaces. Recommendations have been extracted from the results of the experiments at the end of every chapter and these are listed again below. The relevant chapters should be consulted for more detail on each of these recommendations. As stated in each chapter, these are reported here as recommendations rather than as definitive guidelines, as they are based on the results of the studies in this thesis alone, and better performance may be achieved using sets of parameters that have not been explored in this work. In addition, it should be noted that these recommendations are based on the results of controlled, lab based studies, and further research is required to investigate the ecological validity of these results (e.g. to test whether perception of Tactons is affected when the user is engaged in other tasks).

9.3.1 Design Recommendations for Individual Parameters

Design Recommendations for using Vibrotactile Roughness (from Chapter 5 and Chapter 6)

1) To create a smooth sensation, use an un-modulated 250Hz sine wave.
2) To create the sensation of vibrotactile roughness, use amplitude modulated sine waves, with modulation frequencies in the range of 20Hz-50Hz.
    a) To increase perceived roughness within this range, decrease the modulation frequency.
    b) If distinguishable levels of roughness are required, use up to three levels within this range, plus the smooth un-modulated sine wave.
3) For accurate reproduction of roughness stimuli, the actuator should be selected carefully:
    a) Use a device with a short response time (e.g. the C2 Tactor responds in 5ms, compared to 12.5ms for the TACTAID)
    b) When presenting roughness stimuli to the fingertip using a C2 Tactor, use the Tactor in the indented orientation rather than in the raised orientation – the improved perception may be due to the presence of a rigid surround (from Chapter 6).

Design Recommendations for using Intensity Change Over Time (from Chapter 5)

4) If three types of intensity change are required, use level stimuli (no change in intensity), increases and decreases.
5) When creating stimuli that increase or decrease in intensity over time, use linear or exponential, rather than logarithmic, increases/decreases.
6) Do not add a pulse at the start of increases to grab the user's attention.

9.3.2 Design Recommendations for Tactons

Design Recommendations for Two-Dimensional Tactons (from Chapter 6)

7) When designing two-dimensional Tactons, encode one dimension in rhythm and a second dimension in roughness.
    a) Use rhythm to encode the most important dimension.
    b) Use roughness to encode the less important dimension.
8) When designing rhythms, using a different number of vibration bursts in each rhythm may help to create absolutely identifiable rhythms.

Design Recommendations for Two-Dimensional Tactons for Phone Motors (from Chapter 7)

9) When designing two-dimensional Tactons for phone motors, encode one dimension of information in rhythm and another in intensity.
    a) Use rhythm to encode the most important dimension
    b) Use intensity to encode the less important dimension
10) Do not use roughness to encode information when using standard phone vibration motors.

Design Recommendations for Three-Dimensional Tactons (from Chapter 8)

11) When designing three-dimensional Tactons, encode the two most important dimensions in rhythm and spatial location.
12) When using spatial location, locating two Tactors at anatomical landmarks (e.g. wrist, elbow), and one between them can be successful.
13) A less important dimension can be encoded in roughness, intensity or a combination of both.
    a) If three levels of a third parameter are needed, roughness and intensity should be combined redundantly to encode this information, however doing so may impair recognition og other parameters.
    b) Where possible only two levels should be used.

9.4 Limitations of this research

The work reported in this thesis used a systematic approach to exploring the parameters that can be used to encode information in vibrotactile messages called Tactons. While this work has made a significant contribution to knowledge in the field of vibrotactile icon design, and vibrotactile display in general, there are some limitations which must be identified. The limitations of each individual experiment are summarised in the corresponding chapters, and, therefore, this section only presents the general limitations.

There are many interesting factors that could be considered in a dissertation on the design of tactile messages, but the scope of the thesis meant that this dissertation could only focus on a small number of these. This thesis focused on using the tactile modality alone, rather than in combining vibrations with other modalities, and only one approach to encoding information in Tactons is explored. A number of other areas might be of interest in this thesis, but are outside the scope of the current work. These could be explored in future work. One area of interest might be how tactile messages could be used alongside other modalities such as audio, to create multimodal messages. However, before exploring multimodality it is important to gain a good understanding of each modality alone. A large body of literature already exists on the design of audio messages (e.g. [9, 39, 78]) but very little exists on the design of tactile messages. Therefore, this thesis provides an investigation into the design of tactile messages, the results of which could be used to design multimodal messages in future.

This thesis only investigated one approach to Tacton design, namely a structured, abstract approach, based on the techniques used in the design of Earcons. In particular this thesis focused on the design of the design of Transformational Tactons, based on one form on Earcons. Several other types of Tactons were also proposed, namely Compound Tactons and Inherited Tactons and these could also be investigated. These three types of Tactons were defined based on the existence of the same three types in Earcon design. It may be possible to consider other types of Tactons which do not have an equivalent type in the domain of Earcons. One particularly interesting approach would be to investigate the semiotics of touch to understand the meanings which users assign to different types of touch, or to explore cultural differences in how these meanings are assigned. These approaches were beyond the scope of this thesis, but could be of interest in future work.

The findings reported in this thesis may be specific to the devices used to present the Tactons. As was shown in the roughness perception experiment in Chapter 5, quite different results can be achieved when two different devices (TACTAID VBW32 actuator and C2 Tactor) are used to present the same stimuli. In addition, the experiment in Chapter 7 showed that, while rhythm transferred successfully from the high specification C2 device to a standard mobile phone motor, roughness did not. The results indicate that some parameters, like rhythm, may be device independent, but that others are not and, therefore, it is important to consider which device is being used when selecting parameters, rather than just assuming that the results can be directly transferred from device to device. The majority of work in this thesis was conducted with the EAI C2 Tactor, which seems to be the best small vibrotactile actuator that is currently commercially available. These results, therefore, indicate what can be achieved with the best available device, but the experiments in Chapters 2 and 7 also provide some recommendations which could be followed when two lower specification devices are used.

One of the main applications of Tactons is in mobile computing and, therefore, any system to present Tactons must be wearable. The results of the experiments in this thesis have shown that one of the most effective parameters of Tactons is spatial location. To use this parameter, multiple Tactors must be attached to the user's body. With the C2 Tactors it is technically possible to make a wearable tactile display system as the Tactors are small and lightweight enough to be attached to the body, and a Bluetooth control board is now available (Mortimer, B. Personal Communication, 12 November 2005), or the devices can be controlled via the audio port on a mobile computer. However, there are some practical issues in terms of intergrating the actuators into clothing or jewellery in such a way that they will be appealing to users.

The mapping between Tactons and the data they represent is abstract and a set of Tactons could, therefore, be used to represent any multi-dimensional information. However, in each experiment reported in this thesis, only one set of data was ever tested for each set of Tactons. While the mapping is abstract it is likely that users will develop their own methods of remembering the mapping between the Tactons and the data. Therefore, it is not known to what extent the performance in these experiments was dependent on the particular data that was represented. However, since the mapping was abstract and the Tactons were not designed specifically with those application domains in mind, it seems likely that performance would be similar if the same Tactons were used to encode different data.

The work in this thesis has only investigated the identification of Tactons in a lab setting with the user's full attention on the identification task. There has been no evaluation of Tactons while users are engaged in other tasks, such as walking, talking, or reading. If Tactons are to be used in human computer interaction it will be necessary to establish whether they can be identified in these types of conditions, therefore further work is required to investigate this.

The work in this thesis concentrated on the use of quantitative measures to investigate identification of different tactile parameters. Another important measure in HCI is the subjective user experience. This is not measured in this thesis, as the aim of the work was to understand how tactile stimuli were perceived and identified in lab studies, in order to get a best-case scenario measurement, rather than using them within applications. If the Tactons were to be evaluated within applications, as proposed above, the subjective user experience would be an extremely important measure which could be gathered through standard questionnaires e.g. NASA TLX [55]or through purpose-built questionnaires.

9.5 Future Work

The work in this thesis is the first formal investigation into the design of vibrotactile icons using a structured, abstract approach. It has, therefore, opened up many more questions and research ideas which have been outside the scope of this thesis. These ideas are summarised in this section.

9.5.1 Evaluation of Tactons in real world situations

As identified in Section 9.4, one limitation of this work is that the Tacton identification experiments have only been carried out in a lab setting, with the user's full attention on the task. To establish whether Tactons could be effective in real world applications it will be necessary to carry out identification experiments in different contexts and under different degrees of workload. Therefore, experiments should be carried out to investigate how Tacton identification is affected by other tasks, such as walking, talking or reading.

9.5.2 Assessing the Impact of Training

The experiments in this thesis tested identification of Tactons after a very short training period of around 10 minutes. It would be interesting to see how performance is affected by different types of training. In particular, no feedback was provided to users during the training for these experiments. It is possible that providing feedback during the training session might give users more confidence in their responses, and could also correct any miss-mappings they may have made.

Another possibility is to investigate how performance changes after people have been exposed to them regularly over an extended period of time. This could be achieved by conducting a longitudinal study where the Tactons are implemented in a real device and used in peoples' everyday life. The users could then be brought back to the lab to carry out identification experiments at different stages of the study to see how performance had changed.

9.5.3 Improving Tacton Design

While this thesis has reported a detailed investigation into the design of Tactons, it is still possible that the design of Tactons could be improved. The overall levels of identification, particularly for three-dimensional Tactons, were quite low, and it may be possible to improve this. One solution to improve identification rates for Tactons is to use the information transmission results to identify the number of Tactons that could be perfectly identified in any one set, and reduce the number of Tactons to that number. For example, the results from this thesis have indicated that a maximum of 10 Tactons could be perfectly identified (Section 9.2.2). Another possibility is to redesign some of the parameters used in Tactons, and some proposals for how this could be done are outlined in this section.

Improving the Roughness/Intensity Parameter

The work in this thesis has only investigated the use of several fixed levels of each parameter and has relied on the ability of users to uniquely identify these values. In the case of roughness and intensity, many users reported that they found this absolute identification extremely difficult, and felt that making relative judgments was easier. Therefore, it might be possible to design stimuli which rely on relative judgments instead of absolute judgments. This approach was used successfully by Summers *et al.* [102], as reported in Chapter 2, where they found that users could achieve extremely good recognition of step changes in amplitude and frequency.

This approach could be applied to make use of roughness or intensity in a similar manner to how pitch is used in music. In music, different melodies can be created with the same rhythm by changing the pitches used for the notes in the rhythm. For example, the three melodies shown in Figure 9-1 use the same rhythm: two notes lasting 1 beat followed by one note lasting 2 beats. This rhythm is presented using three different pitch structures. In the first melody the pitch rises over time (from low to medium to high), in the second it falls (from high to medium to low), and in the third it goes from high to low to high again.

This same technique could be applied to Tacton design, with the roughness or intensity changing within a rhythm to create tactile “melodies”. One dimension of information could be encoded in the rhythm, while a second dimension could be encoded in the pattern of roughnesses or intensities. This is shown in Figure 9-2, where standard musical notation is used to show the rhythms and the letters L, M, and H represent low, medium and high levels of roughness or intensity.

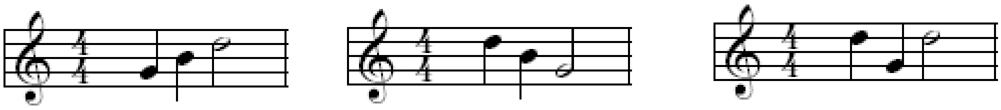

Figure 9-1: In music, different melodies can be created with the same rhythm by changing the pitch structure, as shown in this diagram.

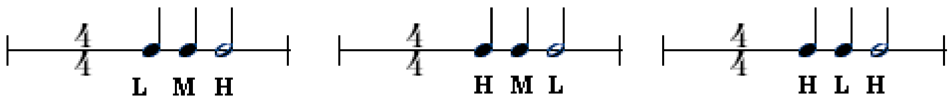

Figure 9-2: By changing roughness or intensity over time, tactile “melodies” could be created. In this diagram: L=low roughness/intensity, M=medium roughness/intensity and H=high roughness/intensity.

Another option which might improve identification of roughness or intensity would be to precede every Tacton with a reference vibration from which users could judge the level of the roughness or intensity used in the Tacton itself. For example, if roughness were used, a single “rough” vibration would be presented before every Tacton. The Tacton itself would be smooth, rough or very rough. Having felt the rough vibration first, the user might find it easier to judge if the Tacton was of the same roughness, or if it was smoother or rougher, than they would have found it to identify the level of roughness if it had been presented in isolation.

Expanding the Dimensionality of Rhythm

In the experiments reported in this thesis, rhythm has been shown to be a very successful parameter. The identification results for rhythm have been consistently good, across different devices and different body locations. However, only three rhythms have been used, and it may be possible to get more from this parameter. In particular, the main differentiating factor in the three rhythms used in this thesis was the number of notes in each rhythm. There may be other factors on which rhythms could be distinguished, such as tempo (speed), time signature (number of beats in a bar of music), and articulation (e.g. short, sharp bursts compared to long, flowing phrases). It is possible that each of these could be considered to be a parameter of rhythm and that a different dimension of information could be encoded in each of these parameters. For example, one dimension of information could be encoded in the basic rhythm (as it was in the experiments in this thesis), while another could be encoded in the speed at which the rhythm is presented (fast/slow). Alternatively, one dimension could be encoded in the basic rhythm and another in the articulation of the rhythm, such that a rhythm presented with a staccato articulation (short, sharp bursts of vibration separated by gaps) could represent something different from the same rhythm presented with a legato articulation (smoothly flowing phrases with little or no gap between notes). An example of this can be seen in Figure 9-3 where the dots represent staccato articulation and the curve represents legato articulation.

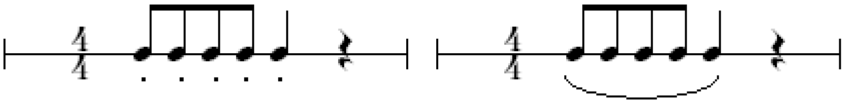

Figure 9-3: The same rhythm articulated in two different ways could be used to represent two different piece of information. This diagram shows the same rhythm with two different articulations: staccato (left) and legato (right).

9.6 Final Remarks

This thesis has provided the first detailed experimental investigation into the design of structured vibrotactile messages called Tactons. Tactons are a specific type of vibrotactile icon, in which information is encoded using a structured, abstract approach, and this is the first time that this approach has been applied to the design of vibrotactile icons. The results from this thesis research, therefore, provide a benchmark against which the results of future research on Tactons, or other tactile icons, can be measured.

The results have shown that it is possible to encode multi-dimensional information in Tactons by manipulating parameters of vibration. In addition, although the mapping between Tactons and the data they represent is abstract, the results have shown that users are able to learn these mappings and, therefore, retrieve information from the Tactons. The amount of information that can be retrieved from Tactons is dependent on the parameters of vibration used to encode it. Therefore, it is important that these parameters are selected carefully.

While a range of studies on vibrotactile perception exist in the domain of psychophysics, there has been little work on how these results can be applied in the area of tactile display for human computer interaction (HCI). Much of the information in the literature on perception of tactile parameters was established on high bandwidth lab based devices, and is not directly transferable to the commercially available transducers that are used by HCI researchers. This thesis has provided data on perception of vibrotactile parameters presented via commercially available actuators. In addition it has provided data on some vibrotactile parameters which had previously been largely unexplored. Furthermore, a set of design recommendations has been produced from the results of the empirical experiments reported in this thesis to aid other researchers or interface designers to design Tactons for use in their own interfaces.

Tactons offer a new means of displaying information non-visually in computer interfaces. As Tactons use an abstract mapping, they can be used to encode any data, and can, therefore, be used in a wide range of applications, e.g. mobile phone alerts, navigation systems for blind people, enhancement of mobile or desktop computer interfaces, etc. This research has provided a set of recommendations for designers and other researchers who might wish to use Tactons, and has indicated what levels of performance can be achieved. With this information, designers can now design Tactons for use in their own interfaces.

References

[1] Grove Music Online: http://www.grovemusic.com/index.html

[2] Ahmaniemi, T., "Vibration as a Rhythm Excitation in Music", in *Proceedings of the Twelfth International Congress on Sound and Vibration*, (2005), pp.

[3] Akamatsu, M., Mackenzie, I.S., and Hasbrouc, T., "A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device", *Ergonomics 38*, 1995, pp 816–827.

[4] Blattner, M.M., Sumikawa, D.A., and Greenberg, R.M., "Earcons and Icons: Their Structure and Common Design Principles", *Human Computer Interaction 4* (1), 1989, pp 11-44.

[5] Bliss, J.C., "Summary of Optacon-related cutaneous experiments", in *Cutaneous Communication Systems and Devices.*, F.A. Geldard (ed). Psychonomic Society: Austin, Texas, 1974, pp 84-94.

[6] Bolanowski, S.J., Gescheider, G.A., Verillo, R.T., and Checkosky, C.M., "Four channels mediate the mechanical aspects of touch", *Journal of the Acoustical Society of America 84* (5), 1988, pp 1680-1694.

[7] Bolanowski, S.J., Gescheider, G.A., and Verrillo, R.T., "Hairy skin: psychophysical channels and their physiological substrates", *Somatosensory and Motor Research 11* (3), 1994, pp 279-290.

[8] Brewster, S. and Brown, L.M., "Tactons: Structured Tactile Messages for Non-Visual Information Display", in *Proceedings of AUIC 2004*, Australian Computer Society (2004), pp 15 - 23.

[9] Brewster, S.A., Wright, P.C., and Edwards, A.D.N., "A detailed investigation into the effectiveness of earcons. " in *Proceedings of ICAD 92*, Addison-Wesley (1992), pp 471-498.

[10] Brewster, S.A., "Providing a Structured Method for Integrating Non-Speech Audio into Human-Computer Interfaces", PhD Thesis, University of York, UK, 1994.

[11] Brewster, S.A., Wright, P.C., and Edwards, A.D.N., "Experimentally derived guidelines for the creation of earcons", in *Adjunct Proceedings of HCI'95*, Springer (1995), pp 155-159.

[12] Brewster, S.A., "Using Non-Speech Sounds to Provide Navigation Cues", *ACM Transactions on Computer-Human Interaction 5* (2), 1998, pp 224-259.

[13] Brewster, S.A. and Brown, L.M., "Non-Visual Information Display Using Tactons", in *Extended Abstracts of CHI 2004*, ACM Press (2004), pp 787-788.

[14] Brewster, S.A. and King, A., "An Investigation into the Use of Tactons to Present Progress Information." in *Proceedings of Interact 2005*, Springer (2005), pp 6-17.

[15] Brown, L.M., Brewster, S.A., and Purchase, H.C., "A First Investigation into the Effectiveness of Tactons", in *Proceedings of World Haptics 2005*, IEEE Press (2005), pp 167-176.

[16] Brown, L.M., Brewster, S.A., and Purchase, H.C., "Tactile Crescendos and Sforzandos: Applying Musical Techniques to Tactile Icon Design", in *Extended Abstracts of CHI 2006*, ACM Press (2006), pp 610-615.

[17] Brown, L.M., Brewster, S.A., and Purchase, H.C., "Multidimensional Tactons for Non-Visual Infomation Display in Mobile Devices", in *Proceedings of Mobile HCI*, ACM Press (2006), pp 231 - 238.

[18] Brown, L.M. and Kaaresoja, T., "Feel Who's Talking: Using Tactons for Mobile Phone Alerts", in *Extended Abstracts of CHI 2006*, ACM Press (2006), pp 604-609.

[19] Buttler, A.J. and Oravainen, K.S., "The effect of rhythm on the memorability of auditive and vibrotactile patterns" (unpublished), Perceptional and Cognitive Psychology: Course Assignment Report, University of Helsinki, Finland (2002).

[20] Carney, A.E. and Beachler, C.R., "Vibrotactile perception of suprasegmental features of speech: A comparison of single-channel and multichannel instruments", *Journal of the Acoustical Society of America 79* (1), 1986, pp 131-140.

[21] Cascio, C.J. and Sathian, "Temporal Cues Contribute to Tactile Perception of Roughness", *Journal of Neuroscience 21* (14), 2001, pp 5289-5296.

[22] Chan, A., "Designing Haptic Icons to Support an Urgency-Based Turn-Taking Protocol", M.Sc. Thesis, Computing Science, University of British Columbia, 2004.

[23] Chan, A., MacLean, K.E., and McGrenere, J., "Learning and Identifying Haptic Icons under Workload", in *Proceedings of World Haptics 2005*, IEEE Press (2005), pp 432-439.

[24] Chang, A., *et al.*, "ComTouch: Design of a Vibrotactile Communication Device", in *Proceedings of DIS 2002*, ACM Press (2002), pp 312-320.

[25] Chang, A. and O'Sullivan, C., "Audio-Haptic Feedback in Mobile Phones", in *Proceedings of CHI 2005*, ACM Press (2005), pp 1264 - 1267.

[26] Cholewiak, R.W. and Collins, A.A., "The generation of vibrotactile patterns on a linear array: Influences of body site, time, and presentation mode", *Perception and Psychophysics 62* (6), pp 1220-1235.

[27] Cholewiak, R.W. and Wollowitz, M., "The design of vibrotactile transducers", in *Tactile Aids for the Hearing Impaired*, I. Summers (ed). Whurr Publishers Ltd: London, 1992, pp 57-82.

[28] Cholewiak, R.W. and Collins, A.A., "Human Testing Program for the Sensor Electronics ELan-01 Tactor, Under Subcontract 255-4118, 1 Sept 99 - 31 Aug 01, Final Technical Report". 2001, The Cutaneous Communication Laboratory at Princeton University.

[29] Cholewiak, R.W. and Collins, A.A., "Vibrotactile localization on the arm: Effects of place, space and age", *Perception and Psychophysics 65* (7), 2003, pp 1058-1077.

[30] Cholewiak, R.W., Brill, C.J., and Schwab, A., "Vibrotactile localization on the abdomen: Effects of place and space", *Perception and Psychophysics 66* (6), 2004, pp 970-987.

[31] Cockburn, A. and Brewster, S.A., " Multimodal feedback for the acquisition of small targets", *Ergonomics 48* (9), 2005, pp 1129-1150.

[32] Craig, J.C. and Sherrick, C.E., "Dynamic Tactile Displays", in *Tactual Perception: A Sourcebook*, W. Schiff and E. Foulke (ed). Cambridge University Press, 1982, pp 209-233.

[33] Craig, J.C., "Tactile pattern recognition and its perturbations", *Journal of the Acoustical Society of America 77* (1), 1984, pp 238-246.

[34] Enriquez, M.J. and MacLean, K.E., "The Hapticon Editor: A Tool in Support of Haptic Communication Research", in *Proceedings of the 11th Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE (2003), pp 356-362.

[35] Fastl, H., "Roughness and temporal masking patterns of sinusiodally amplitude modulated broadband noise", in *Psychophysics and Physiology of Hearing*, E.F. Evans and J.P. Wilson (ed). Academic: London, 1977, pp 403-414.

[36] Foulke, E., "Reading Braille", in *Tactual Perception: A Sourcebook*, W. Schiff and E. Foulke (ed). Cambridge University Press, 1982, pp 168-208.

[37] Franzen, O. and Nordmark, J., "Vibrotactile Frequency Discrimination", *Perception and Psychophysics 17*, 1975, pp 480-484.

[38] Fukumoto, M. and Sugimura, T., "Active Click: Tactile Feedback for Touch Panels", in *Extended Abstracts of CHI 2001*, ACM Press (2001), pp 121-122.

[39] Gaver, W., "The SonicFinder: an interface that uses auditory icons." *Human Computer Interaction 4* (1), 1989, pp 67-94.

[40] Geldard, F.A., "Adventures in tactile literacy." *The American Psychologist 12*, 1957, pp 115–124.

[41] Geldard, F.A., "Some Neglected Possibilities of Communication", *Science 131* (3413), 1960, pp 1583-1588.

[42] Geldard, F.A., "Cutaneous coding of optical signals: The optohapt", *Perception and Psychophysics 1*, 1966, pp 377-381.

[43] Geldard, F.A. and Sherrick, C.E., "The Cutaneous Rabbit: A Perceptual Illusion", *Science 178* (4057), 1972, pp 178-179

[44] Geldard, F.A. and Sherrick, C.E., "The cutaneous saltatory area and its presumed neural basis", *Perception & Psychophysics 33*, 1983, pp 299-304.

[45] Gemperle, F., Ota, N., and Siewiorek, D., "Design of a Wearable Tactile Display", in *Proceedings of the Fifth International Symposium on Wearable Computers*, (2001), pp 5-12.

[46] Gescheider, G.A., "Psychophysics: The Fundamentals", 3rd ed., Lawrence Erlbaum Associates: Mahwah, New Jersey, 1997.

[47] Goble, A.K. and Hollins, M., "Vibrotactile adaptation enhances amplitude discrimination", *Journal of the Acoustical Society of America 93*, 1993, pp 418-424.

[48] Goble, A.K. and Hollins, M., "Vibrotactile adaptation enhances frequency discrimination", *Journal of the Acoustical Society of America 96* (2), 1994, pp 771-780.

[49] Goble, A.K., Collins, A.A., and Cholewiak, R.W., "Vibrotactile threshold in young and old observors: The effects of spatial summation and the presence of a rigid surround", *Journal of the Acoustical Society of North America 99* (4), 1996, pp 2256-2269.

[50] Goff, G.D., "Differential discrimination of frequency of cutaneous mechanical vibration", *Journal of Experimental Psychology 74* (2), 1967, pp 294-299.

[51] Goldstein, E.B., "Sensation and Perception", 5th ed., Brooks/Cole Publishing Company: Pacific Grove, CA, 1999.

[52] Gunther, E., "Skinscape: A Tool for Composition in the Tactile Modality", M.Eng. Thesis, Department of Electrical Engineering and Computer Science, MIT, 2001.
[53] Gunther, E., Davenport, G., and O'Modhrain, S., "Cutaneous Grooves: Composing for the Sense of Touch", in *Proceedings of the Conference on New Instruments for Musical Expression 2002*, (2002), pp 37-42.
[54] Hahn, J.F., "Vibrotactile adaptation and recovery measured by two methods", *Journal of Experimental Pscyhology 71* (5), 1966, pp 655-658.
[55] Hart, S.G. and Wickens, C., "Workload assessment and prediction", in *MANPRINT: an approach to systems integration*, H.R. Booher (ed). Van Nostrand Reinhold: New York, 1990, pp 257 - 296.
[56] Hinton, P.R., "Statistics Explained: A Guide for Social Science Students", Routledge, 1995.
[57] Ho, C., Spence, C., and Tan, H.Z., "Warning Signals Go Multisensory", in *Proceedings of HCI International 2005*, Lawrence Erlbaum Associates (2005), pp 1-10.
[58] Ho, C., Tan, H.Z., and Spence, C., "Using spatial vibrotactile cues to direct visual attention in driving scenes", *Transportation Research Part F: Traffic Psychology and Behaviour 8*, 2005 pp 397-412.
[59] Houtsama, A.J.M., "Estimation of mutual information from limited experimental data", *Journal of the Acoustical Society of North America 74* (5), 1983, pp 1626-1629.
[60] Johnson, K.O., "Neural Basis of Haptic Perception", in *Handbook of Experimental Psychology, Volume 1: Sensation and Perception*, H. Pashler and S. Yantis (ed). Wiley: New York, 2002, pp 537-583.
[61] Jyrinki, T., "Perception of Small Device Vibration Characteristics - Test Facilities and Setup", MSc. Thesis, Department of Electrical and Communications Engineering, Helsinki University of Technology, 2004.
[62] Kaaresoja, T. and Linjama, J., "Perception of Short Tactile Pulses Generated by a Vibration Motor in a Mobile Phone", in *Proceedings of World Haptics 2005*, IEEE Press (2005), pp 471-472.
[63] Kaaresoja, T., Brown, L.M., and Linjama, J., "Snap-Crackle-Pop: Tactile Feedback for Mobile Touch Screens", in *Proceedings of Eurohaptics 2006*, (2006), pp 565-566.
[64] Klatzky, R.L. and Lederman, S.J., "Touch", in *Handbook of Psychology: Volume 4 Experimental Psychology*, A.F. Healy and R.W. Proctor (ed). John Wiley & Sons, 2003, pp 147-176.
[65] Kosonen, K. and Raisamo, R., "Rhythm Perception through different modalities", in *Proceedings of Eurohaptics 2006*, (2006), pp 365-369.
[66] Krueger, L.E., "Tactual Perception in Historical Perspective: David Katz's World of Touch", in *Tactual Perception: A Sourcebook*, W. Schiff and E. Foulke (ed). Cambridge University Press, 1982, pp 1-54.
[67] Lederman, S.J., "Tactile roughness of grooved surfaces: The touching process and effects of macro- and microsurface structure." *Perception and Psychophysics, 16*, 1974, pp 385-395.
[68] Lederman, S.J., "Auditory Texture Perception", *Perception 8*, 1979, pp 93-103.
[69] Lederman, S.J. and Klatzky, R.L., "Hand movements: A window into haptic object recognition", *Cognitive Psychology 19* (3), 1987, pp 342-368.
[70] Leplatre, G. and Brewster, S.A., "Designing Non-Speech Sounds to Support Navigation in Mobile Phone Menus." in *Proceedings of ICAD 2000*, ICAD (2000), pp 190-199.
[71] Levitin, D. and Cook, P.R., "Memory for musical tempo: Additional evidence that auditory memory is absolute", *Perception & Psychophysics, 58*, 1996, pp 927-935.
[72] Luk, J., *et al.*, "A role for haptics in mobile interaction: initial design using a handheld tactile display prototype", in *Proceedings of CHI 2006*, ACM Press (2006), pp 171-180.
[73] Mahns, D.A., *et al.*, "Vibrotactile Frequency Discrimination in Human Hairy Skin", *Journal of Neurophysiology 95*, 2006, pp 1442-1450.
[74] Mason, J.L. and Frost, B.J., "Signal Processing Strategies for Multichannel Systems", in *Tactile Aids for the Hearing Impaired*, I.R. Summers (ed). Whurr Publishers: London., 1992, pp 128-145.
[75] McGookin, D.K., "The Presentation of Multiple Earcons in a Spatialised Audio Space", in *Volume II Proceedings of British HCI 2002*, BCS (2002), pp 228-229.
[76] McGookin, D.K. and Brewster, S.A., "An Investigation into the Identification of Concurrently Presented Earcons", in *Proceedings of ICAD 2003*, ICAD (2003), pp 42-46.
[77] McGookin, D.K., "Understanding and Improving the Identification of Concurrently Presented Earcons", PhD Thesis, Department of Computing Science, University of Glasgow, 2004.

[78] McGookin, D.K. and Brewster, S.A., "Understanding Concurrent Earcons: Applying Auditory Scene Analysis Principles to Concurrent Earcon Recognition", *ACM Transactions on Applied Perception 1* (2), 2004, pp 130-155.

[79] McGookin, D.K. and Brewster, S.A., "Space, The Final Frontearcon: The Identification Of Concurrently Presented Earcons In A Synthetic Spatialised Auditory Environment", in *Proceedings of ICAD 2004*, ICAD (2004), pp 1-8.

[80] Miller, G.A., "Note on the bias of information estimates", in *Information Theory in Psychology: Problems and Methods*, H. Quastler (ed). Free Press: Glencoe, Illinois, 1955, pp 95-100.

[81] Miller, G.A., "The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information", *The Psychological Review, 63*, 1956, pp 81-97.

[82] Mortimer, B., Zets, G., and Cholewiak, R., "Vibrotactile Transduction", *submitted to Journal of the Acoustic Society of America*, 2006, pp.

[83] O'Sullivan, C. and Chang, A., "Dimensional Design; Explorations of the Auditory and Haptic Correlate for the Mobile Device", in *Proceedings of ICAD 2005*, ICAD (2005), pp 212-217.

[84] Phillips, J.R. and Johnson, K.O., "Neural Mechanisms of Scanned and Stationary Touch", *Journal of the Acoustical Society of America 77* (1), 1984, pp 220-224.

[85] Pollack, I. and Ficks, L., "Information of elementary multidimensional auditory displays", *Journal of the Acoustical Society of North America 25*, 1954, pp 765-769

[86] Poupyrev, I., Maruyama, S., and Rekimoto, J., "Ambient Touch: Designing Tactile Interfaces for Handheld Devices", in *Proceedings of UIST 2002*, ACM Press (2002), pp 51 - 60.

[87] Poupyrev, I., Rekimoto, J., and Maruyama, S., "TouchEngine: A Tactile Display for Handheld Devices", in *Extended Abstracts of CHI 2002*, ACM Press (2002), pp 644 - 645.

[88] Poupyrev, I., Okabe, M., and Maruyama, S., " Haptic Feedback for Pen Computing: Directions and Strategies", in *Extended Abstracts of CHI 2004*, ACM Press (2004), pp 1309-1312.

[89] Rabinowitz, W.M., Houtsma, A.J.M., Durlach, N.I., and Delhorn, L.A., "Multidimensional tactile displays: Identification of vibratory intensity, frequency, and contactor area", *Journal of the Acoustical Society of America 82* (4), 1987, pp 1243-1252.

[90] Ram, S. and Sharf, J., "The People Sensor: A Mobility Aid for the Visually Impaired", in *2nd. International Symposium of Wearable Computers*, IEEE (1998), pp 166-167.

[91] Reed, C.M., Durlach, N.I., and Delhorn, L.A., "Natural Methods of Speech Communication", in *Tactile Aids for the Hearing Impaired*, I.R. Summers (ed). Whurr Publishers, 1992.

[92] Rosenberg, L.B. and Adelstein, B.D., "Perceptual Decomposition of Virtual Haptic Surfaces", in *IEEE Virtual Reality 1993*, IEEE Press (1993), pp 46-53.

[93] Ross, D.A. and Blasch, B.B., "Wearable interfaces for orientation and wayfinding", in *Proceedings of ASSETS 2000*, ACM Press (2000), pp 193-200.

[94] Rothenberg, M., *et al.*, "Vibrotactile frequency for encoding a speech parameter", *The Journal of the Acoustical Society of America 62*, 1977, pp 1003–1012.

[95] Rovers, A.F. and Van Essen, H.A., "HIM: A Framework for Haptic Instant Messaging", in *Extended Abstracts of CHI 2004*, ACM Press (2004), pp 1313-1316.

[96] Rovers, A.F. and van Essen, H.A., "FootIO – Design and evaluation of a device to enable foot interaction over a computer network", in *Proceedings of World Haptics 2005*, IEEE Press (2005), pp 521-522.

[97] Senders, V.L., "Measurement and Statistics", Oxford University Press: New York, 1958.

[98] Sherrick, C., "A Scale for Rate of Tactual Vibration", *Journal of the Acoustical Society of America 78*, 1985, pp 78-83.

[99] Sherrick, C.E. and Rogers, R., "Apparent Haptic Movement", *Perception and Psychophysics 1*, 1966, pp 175-180.

[100] Sherrick, C.E. and Craig, J.C., "The Psychophysics of Touch", in *Tactual Perception: A Sourcebook*, W. Schiff and E. Foulke (ed). Cambridge University Press, 1982, pp 55-81.

[101] Summers, I.R., "Signal Processing Strategies for Single-Channel Systems", in *Tactile Aids for the Hearing Impaired*, I.R. Summers (ed). Whurr Publishers: London., 1992, pp 110-127.

[102] Summers, I.R., *et al.*, "Vibrotactile and electrotactile perception of time-varying pulse trains", *Journal of the Acoustical Society of America 95* (3), 1994, pp 1548-1558.

[103] Summers, I.R., Cooper, P.G., and Wright, P., "Information from time-varying vibrotactile stimuli", *Journal of the Acoustical Society of America 102* (6), 1997, pp 3686-3696.

[104] Summers, I.R., Chanter, C.M., Southall, A.L., and Brady, A.C., "Results from a Tactile Array on the Fingertip", in *Eurohaptics 2001*, (2001), pp 26-28.

[105] Tan, H., Durlach, N.I., Reed, C.M., and Rabinowitz, W.M., "Information transmission with a multifinger tactual display", *Perception & Psychophysics 61* (6), 1999, pp 993-1008.

[106] Tan, H.Z., "Information Transmission with a Multi-Finger Tactual Display", PhD Thesis, Electrical Engineering and Computer Science, MIT, 1996.
[107] Tan, H.Z. and Pentland, A., "Tactual displays for wearable computing", in *Proceedings of the International Symposium on Wearable Computing 1997*, IEEE Press (1997), pp 84-89.
[108] Tan, H.Z., Lim, A., and Traylor, R., "A psychophysical study of sensory saltation with an open response paradigm", in *ASME Dynamic Systems and Control Division*, (2000), pp 1109-1115.
[109] Terhardt, E., "On the perception of periodic sound fluctuations (roughness)", *Acustica 30*, 1974, pp 201-213.
[110] Van Doren, C.L., Gescheider, G.A., and Verillo, R.T., "Vibrotactile temporal gap detection as a function of age", *Journal of the Acoustical Society of America 87* (5), 1990, pp 2201-2206.
[111] van Erp, J.B.F. and van Veen, H.A.H.C., "Vibro-Tactile Information Presentation in Automobiles", in *Proceedings of Eurohaptics 2001*, (2001), pp 99-104.
[112] van Erp, J.B.F., "Guidelines for the Use of Vibro-Tactile Displays in Human Computer Interaction", in *Proceedings of Eurohaptics 2002*, (2002), pp 18-22.
[113] van Erp, J.B.F. and Spapé, M.M.A., "Distilling the Underlying Dimensions of Tactile Melodies", in *Proceedings of Eurohaptics 2003*, (2003), pp 111-120.
[114] van Veen, H.A.H.C. and van Erp, J.B.F., "Tactile Information Presentation in the Cockpit." in *Proceedings of the First International Workshop on Haptic Human-Computer Interaction*, Springer LNCS (2000), pp 174-181.
[115] Verrillo, R.T. and Gescheider, G.A., "Perception via the sense of touch", in *Tactile Aids for the Hearing Impaired*, I.R. Summers (ed). Whurr Publishers: London., 1992, pp 1–36.
[116] von Békésy, G., "Neural Funneling along the Skin and between the Inner and Outer Hair Cells of the Cochlea", *Journal of the Acoustical Society of America 31* (9), 1959, pp 1236-1249.
[117] Wall, S.A. and Brewster, S.A., "Sensory Substitution Using Tactile Pin Arrays: Human Factors, Technology and Applications", *To appear in Special Issue of Signal Processing on Multimodal Interfaces*, 2006, pp.
[118] Weinstein, S., "Intensive and extensive aspects of tactile sensitivity as a function of body part, sex and laterality, " in *The Skin Senses*, , D.R.K. (ed). (ed). Charles C. Thomas: Springfield, IL, , 1968, pp 195-218.
[119] Weisenberger, J.M., "Sensitivity to amplitude-modulated vibrotactile signals", *Journal of the Acoustical Society of America 80* (6), 1986, pp 1707-1715.
[120] Wenzel, E.M., Arruda, M., Kistler, D.J., and Wightman, F.L., "Localization using nonindividualized head-related transfer functions. " *Journal of the acoustical society of America 94* (1), 1993, pp 111–123.
[121] White, B.W., *et al.*, "Seeing with the skin", *Perception and Psychophysics 7*, 1970, pp 23-27.

Appendix A: Experiment 1 Files

A.1 Experiment 1: Roughness Experiment, Matlab code

Matlab code to produce a sine wave

```
function s = make_sine(f,fs,d)
%f is frequency, fs is sampling frequency, d is duration

T = 1/fs;
t=0:T:d;
s = sin(2*pi*f*t);
```

Matlab code for roughness stimuli (amplitude modulation)

```
function f = make_ampMod_waveform(fcarrier,fmod,modDepth, fs, d)
%fcarrier is carrier frequency, fmod is modulation frequency, modDepth is modulation depth
%fs is sampling frequency, d is duration.

carrier = make_sine2(fcarrier,fs,d);
mod = make_sine2(fmod, fs,d);
f = (1+modDepth.*mod).*carrier;
```

A.2 Experiment 1: Roughness Experiment, Consent Form

This experiment aims to discover how people perceive roughness through vibrations. These vibrations will be presented to the forefinger of your non-dominant hand (left hand if you are right-handed), via a small transducer that vibrates (the experimenter will show this to you). The transducer will be held onto your finger by double sided sticky tape and surgical tape. If you have any allergies to adhesives such as sticky tape please advise the experimenter immediately.

In this experiment you will be presented with 50 pairs of vibrations and for each pair you will be asked to make a decision on which vibration feels "rougher". Roughness is a subjective concept, so there is no right or wrong answer; we just want to know which vibration feels rougher to you.

Some of the vibrations may feel very similar; if you find a pair which you feel unable to tell apart please just make a guess. You should not worry if you find it difficult to tell some pairs apart.

Since your skin is probably not accustomed to this type of vibrations you may experience slight tingling sensations in your finger during the experiment. If you feel any discomfort, pain or fatigue at any time, please inform the experimenter immediately. We strongly recommend that you take regular breaks during the experiment. You can stop and take a break at any time during tasks. The length of time taken to respond is not recorded, therefore you should feel free to take a break for as long as required. You will receive a reward of £5 upon completion of the experiment.

During the experiment the computer will record data on your responses. This data will be stored anonymously using an ID number, rather than your name or any number that could identify you. All results will be held in strict confidence, ensuring the privacy of all participants. There will be no record kept that would allow your results to be traced back to you. All data will also be held securely in a password protected computer system (for electronic data) or a filing cabinet in a locked office (for paper based data). The data (including that from other participants) will be analysed to identify trends which may be generalised over the larger population. The results of this analysis may be published in appropriate scientific journals and conferences.

A feedback sheet (containing the summaries of the data mentioned above) will be sent to all participants who request it, after the data has been analysed.

Your participation in this experiment will have no effect on your marks for any subject at this, or any other university.

You may withdraw from the experiment at anytime without prejudice, and any data already recorded will be overlooked.

A.3 Experiment 1: Roughness Experiment, Instructions

In this experiment you will be presented with pairs of vibrations and asked to make a decision on which vibration feels "rougher". Roughness is a subjective concept, so there is no right or wrong answer; we just want to know which vibration feels rougher to you. Some of the vibrations may feel very similar; if you find a pair which you feel unable to tell apart please just make a guess. The experiment will not allow you to carry on if you do not make a selection.

When you start the experiment you will see the dialogue shown in Figure 1. You will be able to feel the vibrations by clicking on the buttons "Play Vibration 1" and "Play Vibration 2". You can play each vibration up to four times in each task. After four presses, the button will be disabled and will appear greyed out. Once you have made a decision about which you feel is rougher, make your selection by clicking on the radio button corresponding to your decision, e.g. in Figure 2 the user has selected Vibration 2.

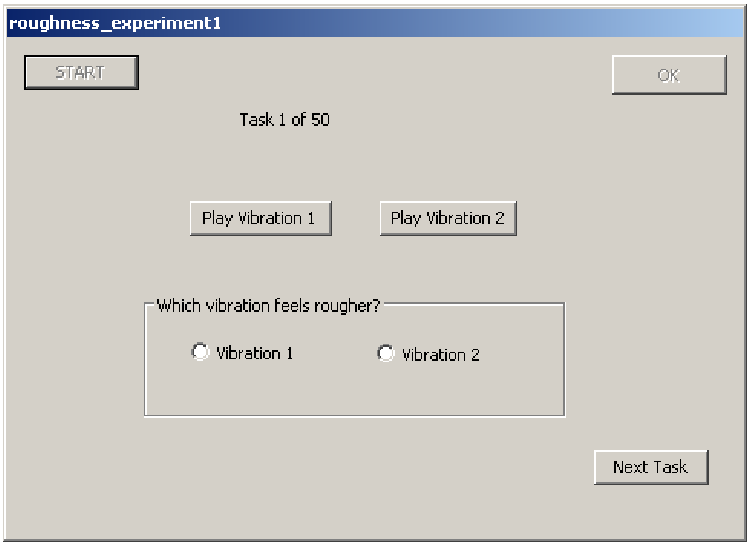

Figure 1: The experiment dialogue

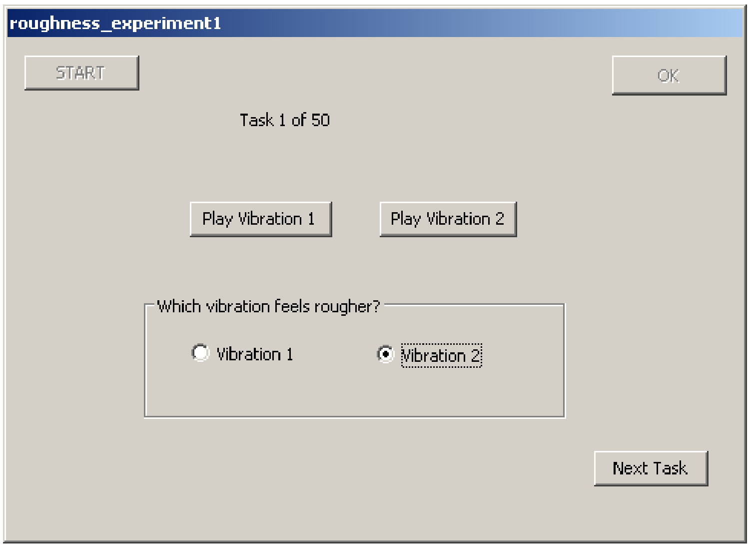

Figure 2: The user has selected "Vibration 2" as the rougher vibration

Once you have made your selection, press the “Next Task” button to move onto the next task. You will see the Task Number display change to indicate which task you are currently on, e.g. “Task 2 of 50”.

When you press next after the final task (Task 50), you will see the message shown in Figure 3. At this point, please inform the experimenter that you are finished.

Figure 3: message indicating end of experiment

Appendix B: Experiment 2 Files

B.1 Experiment 2: Intensity Change Over Time Experiment, Matlab Code

Linear Increase

```
function a = linear_increase(wave, fs, d)
% wave is the waveform (already created) to be increased in amplitude
%fs is the sampling frequency, d is the duration.

% create array that increases in amplitude over time
increasingArray = [0:(fs*d)];
size(increasingArray)

% scale so it goes from 0 to 1
scaledIncrease = increasingArray/((fs*d));

% then multiply the original wave by this so that the original wave increases in amplitude over time.
size(scaledIncrease)
a = scaledIncrease.*wave;
```

Linear Decrease

```
function a = linear_decrease(wave, fs, d)
% wave is the waveform (already created) to be decreased in amplitude
% fs is the sampling frequency.
% d is the duration.

% create array that decreases in amplitude over time
decreasingArray = [(fs*d):-1: 0];
size(decreasingArray)

% scale so it goes from 0 to 1
scaledDecrease = decreasingArray/((fs*d));

% then multiply the original wave by this so that the original wave decreases in amplitude over time.
a = scaledDecrease.*wave;
```

Linear Increase with “sforzando-piano”

```
function a = linear_increase_sfz(wave, fs, d, maxTime, maxAmpFactor)
% wave is the waveform (already created) to be increase in amplitude
%fs is the sampling frequency, d is the duration
% maxTime is the length of the sforzando, maxAmpFactor is the intensity of the sforzando.

% create array that increases in amplitude over time
increasingArray = [0:(fs*d)];
size(increasingArray)

% scale so it goes from 0 to 1
scaledIncrease = increasingArray/((fs*d));

% then multiply the original wave by this so that the original wave
% increases in amplitude over time.
size(scaledIncrease)
scaledIncrease(1:maxTime) = max(scaledIncrease)*maxAmpFactor;

a = scaledIncrease.*wave;
```

Exponential Increase

```
function a = exponential_increase(wave, fs, duration)
s = logspace(0,1,(fs*duration)+1);
 a = (s.*wave)/10;
```

Exponential Decrease

```
function a = exponential_decrease(wave, fs, duration)
s = logspace(1,0,(fs*duration)+1);
a = (s.*wave)/10;
```

Exponential Increase with “Sforzando piano”

```
function a = exp_increase_sfz(wave, fs, duration, maxTime)
s = logspace(0,1,(fs*duration)+1);
s(1:maxTime) = max(s);
res= s.*wave;
a = res/10;
```

Logarithmic Increase

```
function a = linspace_increase(wave, fs, duration,multFactor)
lin = linspace(0,1,(fs*duration) +1);

for a = 1:(fs*duration) + 1 % 1 for array starting at 0
   result(a) = 1.0 - exp(-multFactor * lin(a));
end

result2 = result*1.58;% to scale it up to 1
size(result2)
a = result2.*wave;
```

Logarithmic Decrease

```
function a = linspace_decrease(wave, fs, duration,multFactor)
lin = linspace(1,0,(fs*duration) +1);

for a = 1:(fs*duration) + 1 % 1 for array starting at 0
   result(a) = 1.0 - exp(-multFactor * lin(a));
end

result2 = result*1.58;
size(result2)
a = result2.*wave;
```

Logarithmic Increase with "Sforzando piano"

```
function a = linspace_increase_sfz_maxVal(wave, fs, duration, maxTime,
multFactor,maxValFactor)
lin = linspace(0,1,(fs*duration) +1);

for a = 1:(fs*duration) + 1 % 1 for array starting at 0
   result(a) = 1.0 - exp(-multFactor * lin(a));
end

result2 = result*1.58;
result2(1:maxTime) = maxValFactor*max(result2);
size(result2)
a = result2.*wave;
```

B.2 Experiment 2: Intensity Change over Time Experiment, Consent Form

This experiment aims to discover whether people can identify changes in the amplitude of a vibration over time. These vibrations will be presented to the forefinger of your non-dominant hand, via a small transducer that vibrates (the experimenter will show this to you). The transducer will be held onto your finger by a piece of sticky tape. If you have any allergies to adhesives such as sticky tape please advise the experimenter immediately.

There are 45 tasks in the experiment. The experiment uses multiple choice for responses, so for each task you will be asked to pick from four options. At the end of the experiment you will be asked to complete a short questionnaire to identify how difficult you found the experiment and how you felt about your performance. You will receive a reward of £5 upon completion of the experiment.

It is important to remember that we are not testing you; we are testing the effectiveness of these vibrations as tactile cues. We want to find out which cues are effective, so if you find a task difficult it is indicative that the cue would be ineffective, and not of any inability on your part.

During the experiment the computer will record data on your responses and the time taken to respond. This data will be stored anonymously using an ID number, rather than your name or any number that could identify you. All results will be held in strict confidence, ensuring the privacy of all participants. There will be no record kept that would allow your results to be traced back to you. All data will also be held securely in a password protected computer system (for electronic data) or a filing cabinet in a locked office (for paper based data). The data (including that from other participants) will be analysed to identify trends which may be generalised over the larger population. The results of this analysis may be published in appropriate scientific journals and conferences.

Since your skin is probably not accustomed to this type of vibrations you may experience slight tingling sensations in your finger during the experiment. If you feel any discomfort, pain or fatigue at any time please inform the experimenter immediately. We strongly recommend that you take regular breaks during the experiment. You will have you will have the opportunity to take a break after responding to each task, before starting the next task.

A feedback sheet (containing the summaries of the data mentioned above) will be sent to all participants who request it, after the data has been analysed.

Your participation in this experiment will have no effect on your marks for any subject at this, or any other university.

You may withdraw from the experiment at anytime without prejudice, and any data already recorded will be overlooked.

B.3 Experiment 2: Intensity Change over Time Experiment, Instructions

In this experiment you will be presented with vibrations and asked to choose, from four options, what happens to the amplitude (intensity) of the vibration over time. The four options are:

1. *Overall Increase*: Overall, the amplitude of the vibration increases over time
2. *Overall Decrease:* Overall, the amplitude of the vibration decreases over time
3. *Level*: The amplitude of the vibration stays constant over time.
4. *Other*: The amplitude of the vibration does something other than the three options above

Before starting the experiment the experimenter will let you try a few tasks like those in the experiment in order to familiarise you with the software and the types of stimuli you will encounter.

There are 45 tasks in the experiment, and the experiment will run as follows:

Step 1:

The experiment will start with the dialog box shown in Figure 1. When you are ready to start the experiment, press the START button.

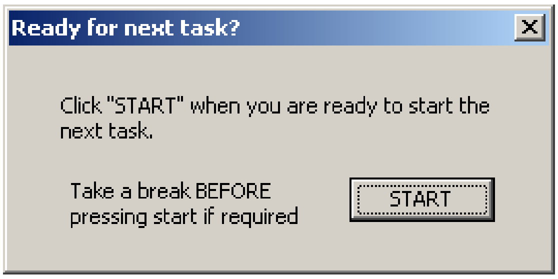

Figure 1: Prompt to start next task

Step 2:

When you press START you will see the screen shown in Figure 2 (on next page of these instructions). After a few seconds a 90 second vibration will be presented to your finger. This will be repeated four times with a one second pause between presentations. As soon as you have decided what happens to the amplitude of the vibration over time, please select an answer by clicking on the corresponding radio button. . There is no need to wait until the vibration has been repeated all four times before choosing your answer. As soon as you have made your selection, and are happy with it, press the NEXT button.

Step 3:

When you press NEXT there will be a short pause and then the dialog shown in Figure 1 will reappear. If you feel tired at all please take this opportunity to rest – once you press NEXT the next task will begin immediately. When you are ready to start the next task, press the START button.

Steps 2 and 3 will be repeated for all 45 tasks. As the tasks progress you will see the Task Number display on the dialog box change, e.g. to “Task 2 of 45” so you can keep track of your progress.

Once the final task is complete a dialog box will appear, which says “experiment ended, please see experimenter”. At this point please remove the transducer from your finger and tell the experimenter that you are finished. They will then ask you to fill in a short questionnaire about this condition.

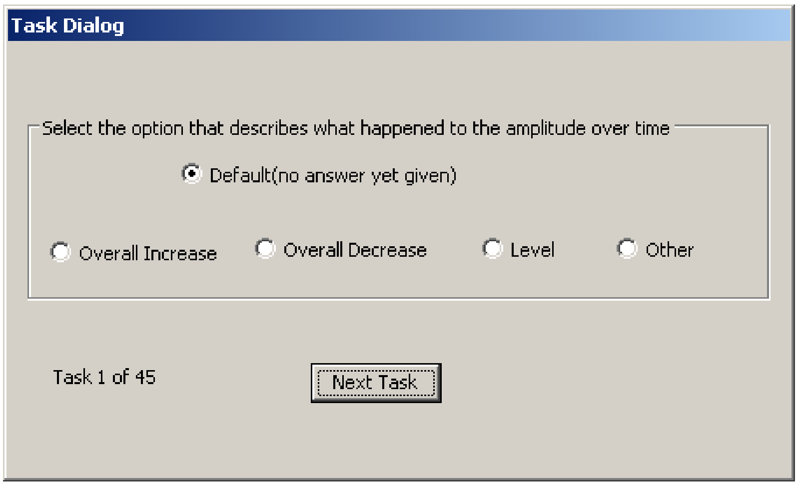

Figure 2: The Task Dialog

Appendix C: Experiment 3(a) & 3(b) Files

C.1 Experiment 3(a) and 3(b): Two-Dimensional Tactons Experiment, Consent Form

This experiment aims to discover whether people are able to identify parameters of vibro-tactile messages called Tactons. These vibrations will be presented to the forefinger of your non-dominant hand (left hand if you are right-handed), via a small transducer that vibrates (the experimenter will show this to you). This will be attached to your finger with surgical tape.

This experiment consists of 54 tasks. In each task you will be presented with one Tacton asked to identify two parameters of the vibro-tactile message.

Some of the tasks are very hard so if you find it difficult to identify a parameter please just make a guess. You should not worry if you find this experiment difficult. We are trying to identify parameters which people are able to identify so if you are unable to identify them it will be good for us to learn this at this stage.

Since your skin is probably not accustomed to this type of vibrations you may experience slight tingling sensations in your finger during the experiment. If you feel any discomfort, pain or fatigue at any time please inform the experimenter immediately. We strongly recommend that you take regular breaks during the experiment. You can stop and take a break at any time during tasks. The length of time taken to respond is not recorded, therefore you should feel free to take a break for as long as required. You will receive a reward of £5 upon completion of the experiment.

During the experiment the computer will record data on your responses. This data will be stored anonymously using an ID number, rather than your name or any number that could identify you. All results will be held in strict confidence, ensuring the privacy of all participants. There will be no record kept that would allow your results to be traced back to you. All data will also be held securely in a password protected computer system (for electronic data) or a filing cabinet in a locked office (for paper based data). The data (including that from other participants) will be analysed to identify trends which may be generalised over the larger population. The results of this analysis may be published in appropriate scientific journals and conferences.

A feedback sheet (containing the summaries of the data mentioned above) will be sent to all participants who request it, after the data has been analysed.

Your participation in this experiment will have no effect on your marks for any subject at this, or any other university.

You may withdraw from the experiment at anytime without prejudice, and any data already recorded will be overlooked.

C.2 Experiment 3(a) and 3(b): Two-Dimensional Tactons Experiment, Training Sheet

Each of the vibrations present in the experiment is called a Tacton, and represents an alert which might occur when a message or a call arrives on your mobile phone. There are two pieces of information in the Tacton – the type of call/ message, and the priority of the call/message.

In order to successfully perform the experiment you must understand the vibrations used, and how to retrieve the two pieces of information that are encoded in a Tacton.

The Tactons are structured so that the *rhythm* encodes the *type* of call/message. The type of call/message can either be a **voice call,** a **text message** or a **multimedia message**.

The *roughness* of the vibration used to produce the Tacton encodes the priority of the call/message. There are 3 distinct roughnesses for **low** (smooth), **medium** (rough) and **high** (very rough) priority calls/messages.

In order to ensure that you understand the Tactons, all of the Tactons used in the experiment, along with their attributes are provided on a web page. You will have 10 minutes to ensure you thoroughly understand the Tactons and their structures.

Please see the Experimenter.

C.3 Experiment 3(a) and 3(b): Two-Dimensional Tactons Experiment, Instructions

This experiment consists of 54 tasks. Each task will consist of you feeling one Tacton. You will feel the Tacton repeat 4 times (with a one second pause between repetitions). During this time you will need to identify both of the attributes (the **type** and **priority** of the call/message) encoded in the Tacton. As soon as you have identified these attributes, please select your answers by clicking on the corresponding radio buttons. There is no need to wait until the Tacton has been repeated all seven times before choosing your answer. As soon as you have made your selection, and are happy with it, press the NEXT button.

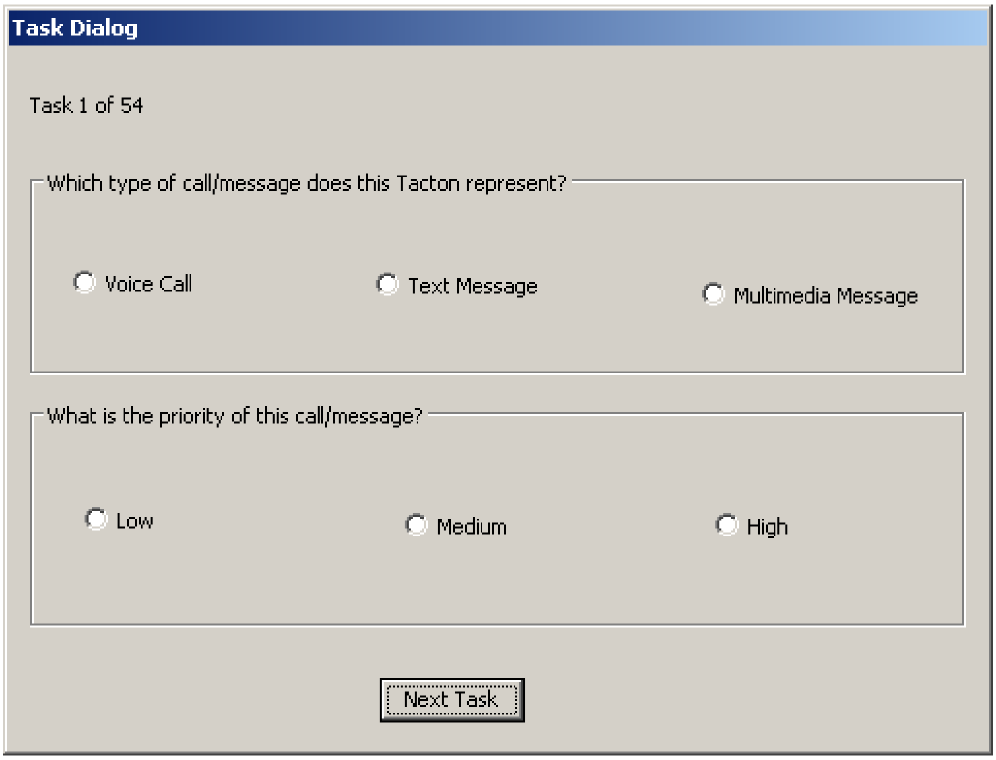

When you press NEXT there will be a short pause and then the dialog shown in below will appear. If you wish to pause for a short break, **you should do so before clicking this dialog.** When you are ready to start the next task, press the START button.

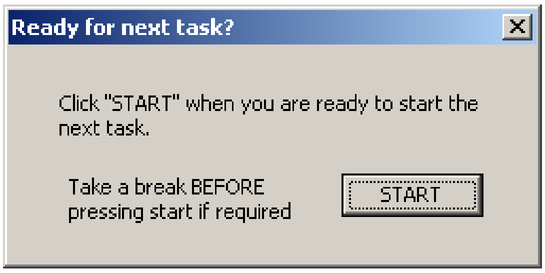

Before starting the experiment itself, you will do 18 of the above tasks as practise for the experiment.

Please see the Experimenter.

Appendix D: Experiment 4 Files

D.1 Experiment 4: 2D Tactons for Mobile Phone Motors Experiment, Consent Form

This experiment aims to discover whether people are able to identify parameters of vibra messages called Tactons. These vibrations will be presented to the your hand

This experiment consists of 2 conditions, with 54 tasks (and 10 training tasks) in each condition. In each task you will be presented with one Tacton asked to identify two parameters of the vibro-tactile message.

Some of the tasks are very hard so if you find it difficult to identify a parameter please just make a guess. You should not worry if you find this experiment difficult. We are trying to identify parameters which people are able to identify so if you are unable to identify them it will be good for us to learn this at this stage.

Since your skin is probably not accustomed to repeated vibrations you may experience slight tingling sensations in your hand during the experiment. If you feel any discomfort, pain or fatigue at any time please inform me immediately. You can stop and take a break at any time between tasks, and we recommend that you do so – please just me if you would like a break.

You will receive a reward of two movie tickets upon completion of the experiment.

During the experiment I will record data on your responses. This data will be stored anonymously using an ID number, rather than your name or any number that could identify you. All results will be held in strict confidence, ensuring the privacy of all participants. There will be no record kept that would allow your results to be traced back to you. The data (including that from other participants) will be analysed to identify trends, which may be generalised over the larger population. The results of this analysis may be published in appropriate scientific journals and conferences.

You may withdraw from the experiment at anytime without prejudice, and any data already recorded will be overlooked.

D.2 Experiment 4: 2D Tactons for Mobile Phone Motors, Training Sheet – Roughness Condition

Each of the vibrations present in the experiment is called a Tacton, and represents an alert which might occur when a message or a call arrives on your mobile phone. There are two pieces of information in the Tacton – the type of call/ message, and the priority of the call/message.

In order to successfully perform the experiment you must understand the vibrations used, and how to retrieve the two pieces of information that are encoded in a Tacton.

The Tactons are structured so that the *rhythm* encodes the *type* of call/message. The type of call/message can either be a **voice call,** a **text message** or a **multimedia message**.

The *roughness* of the vibration used to produce the Tacton encodes the priority of the call/message. There are 3 distinct roughnesses for **low** (smooth), **medium** (rough) and **high** (very rough) priority calls/messages.

In order to ensure that you understand the Tactons, all of the Tactons can be tested by pressing the numbers on the numerical keypad provided. The chart below shows how the numbers on the box map to the Tactons. You will have 10 minutes to try these and to ensure you thoroughly understand the Tactons and their structures.

Please see the Experimenter.

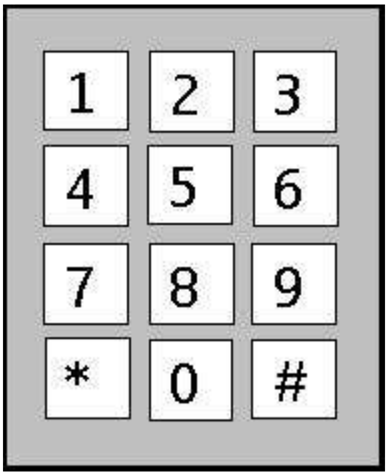

1. Voice Call, Low Priority
2. Voice Call, Medium Priority
3. Voice Call, High Priority
4. Text Message, Low Priority
5. Text Message, Medium Priority
6. Text Message, High Priority
7. Multimedia Message, Low Priority
8. Multimedia Message, Medium Priority
9. Multimedia Message, High Priority

*, 0 and # have no functionality

D.3 Experiment 4: 2D Tactons for Mobile Phone Motors, Training Sheet – Intensity Condition

Each of the vibrations present in the experiment is called a Tacton, and represents an alert which might occur when a message or a call arrives on your mobile phone. There are two pieces of information in the Tacton – the type of call/ message, and the priority of the call/message.

In order to successfully perform the experiment you must understand the vibrations used, and how to retrieve the two pieces of information that are encoded in a Tacton.

The Tactons are structured so that the *rhythm* encodes the *type* of call/message. The type of call/message can either be a **voice call,** a **text message** or a **multimedia message**.

The *intensity* of the vibration encodes the priority of the call/message. There are 3 distinct intensity levels for **low** (low intensity), **medium** (medium intensity) and **high** (high intensity) priority calls/messages.

In order to ensure that you understand the Tactons, all of the Tactons can be tested by pressing the numbers on the numerical keypad provided. The chart below shows how the numbers on the box map to the Tactons. You will have 10 minutes to try these and to ensure you thoroughly understand the Tactons and their structures.

Please see the Experimenter.

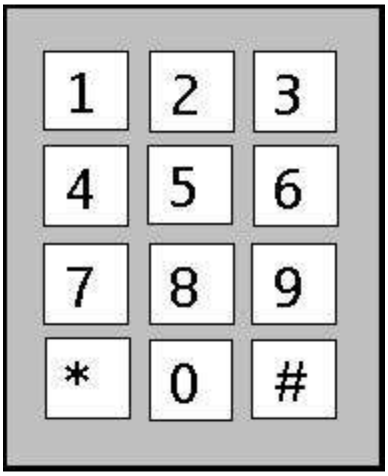

1. Voice Call, Low Priority
2. Voice Call, Medium Priority
3. Voice Call, High Priority
4. Text Message, Low Priority
5. Text Message, Medium Priority
6. Text Message, High Priority

7. Multimedia Message, Low Priority
8. Multimedia Message, Medium Priority
9. Multimedia Message, High Priority

*, 0 and # have no functionality

Experiment 4: 2D Tactons for Mobile Phone Motors, Experiment Instructions

This experiment consists of 54 tasks. Each task will consist of you feeling one Tacton. You will feel the Tacton repeat 4 times (with a one second pause between repetitions). During this time you will need to identify both of the attributes (the **type** and **priority** of the call/message) encoded in the Tacton. As soon as you have identified these attributes, please mark you answer on the answer sheet by putting a cross in the boxes corresponding to the type and priority of the message.

Once you have made your response, please indicate to the experimenter that you are ready for the next Tacton. You are welcome to take a break between tasks at any time. If you wish to do so, please tell the experimenter you would like a short break before continuing.

Before starting the experiment itself, you will do 10 of the above tasks as practise for the experiment.

Please see the Experimenter.

D.4 Experiment 4: 2D Tactons for Mobile Phone Motors, Response Sheet

Participant ID ________

Task	Type		Priority	
1	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐
2	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐
3	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐
4	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐
5	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐

6	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐
7	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐
8	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐
9	Voice	☐	Low	☐
	Text	☐	Medium	☐
	Multimedia	☐	High	☐

Appendix E: Experiment 5(a) and 5 (b) Files

E.1 Experiment 5(a) and 5(b), Three-Dimensional Tactons Experiment, Consent Form

This experiment aims to discover whether people are able to identify parameters of vibro-tactile messages called Tactons. These vibrations will be presented to three points on your forearm via small transducers that vibrate (the experimenter will show these to you). These will be held against your arm by Velcro bands.

This experiment consists of 81 tasks. In each task you will be presented with one Tacton asked to identify three parameters of the vibro-tactile message.

Some of the tasks are very hard so if you find it difficult to identify a parameter please just make a guess. You should not worry if you find this experiment difficult. We are trying to identify parameters which people are able to identify so if you are unable to identify them it will be good for us to learn this at this stage.

Since your skin is probably not accustomed to this type of vibrations you may experience slight tingling sensations in your arm during the experiment. If you feel any discomfort, pain or fatigue at any time please inform the experimenter immediately. We strongly recommend that you take regular breaks during the experiment, between tasks. You will receive a reward of £6 upon completion of the experiment.

During the experiment the computer will record data on your responses. This data will be stored anonymously using an ID number, rather than your name or any number that could identify you. All results will be held in strict confidence, ensuring the privacy of all participants. There will be no record kept that would allow your results to be traced back to you. All data will also be held securely in a password protected computer system (for electronic data) or a filing cabinet in a locked office (for paper based data). The data (including that from other participants) will be analysed to identify trends which may be generalised over the larger population. The results of this analysis may be published in appropriate scientific journals and conferences.

A feedback sheet (containing the summaries of the data mentioned above) will be sent to all participants who request it, after the data has been analysed.

Your participation in this experiment will have no effect on your marks for any subject at this, or any other university.

You may withdraw from the experiment at anytime without prejudice, and any data already recorded will be overlooked.

E.2 Experiment 5(a) and 5(b), Three-Dimensional Tactons Experiment, Level Setting Instructions

Before you start we need to make sure that the strength of the vibration from all three devices feels equal to you. We will now ask you to set the level of the devices by turning the volume dials on the amplifiers until the levels of all three devices feel the same. You should not adjust the mid-point level, only the level of the elbow and the wrist devices. The elbow device is controlled using the top amplifier, the wrist by the bottom amplifier. You can feel a vibration on each device by pressing the corresponding buttons on the dialog.

You can take up to 5 minutes to set these levels. It is difficult to get them to feel exactly the same but do your best to make them feel as similar as possible.

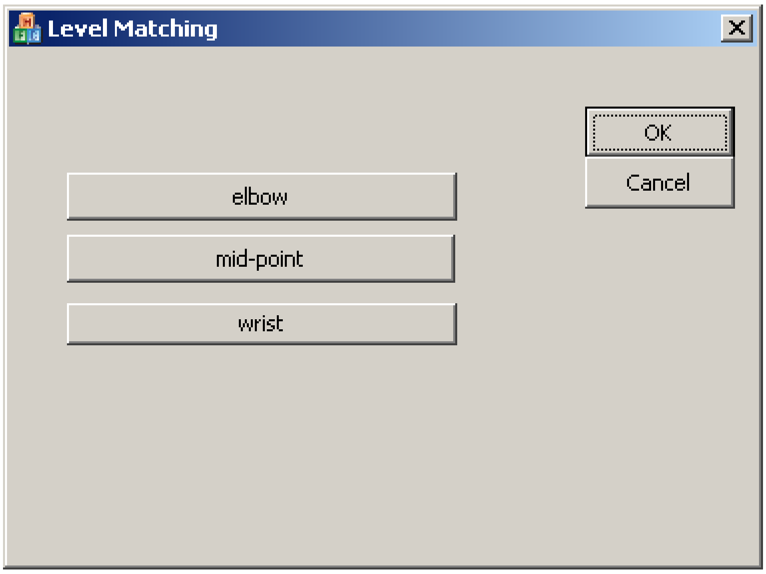

E.3 Experiment 5(a), Three-Dimensional Tactons Experiment, Training Sheet

Each of the vibrations present in the experiment is called a Tacton, and represents an alert which might occur when on your phone or handheld computer to alert you to an upcoming event. Three pieces of information are encoded in each Tacton – the type of appointment, the importance of the appointment, and the time remaining before the appointment.

In order to successfully perform the experiment you must understand the vibrations used, and how to retrieve the two pieces of information that are encoded in a Tacton.

The Tactons are structured so that the *rhythm* encodes the *type* of appointment. The type of appointment can either be a **meeting,** a **lecture** or a **tutorial**.

The *roughness* of the vibration used to produce the Tacton encodes the importance of the appointment. There are 3 distinct roughnesses for **low** (smooth), **medium** (rough) and **high** (very rough) importance appointments.

The *location* of the vibration on your arm encodes the time remaining until the appointment. A vibration presented near your **elbow** means that the appointment will takes place in **30 mins**, a vibration in the **middle of your arm** indicates that the appointment will take place in **15 minutes**, and a vibration at your **wrist** means that the appointment will occur in **5 minutes**.

In order to ensure that you understand the Tactons, all of the Tactons used in the experiment, along with their attributes are provided on a dialog box (see below). You can start each Tacton by clicking on the corresponding button. You will have 10 minutes to ensure you thoroughly understand the Tactons and their structures.

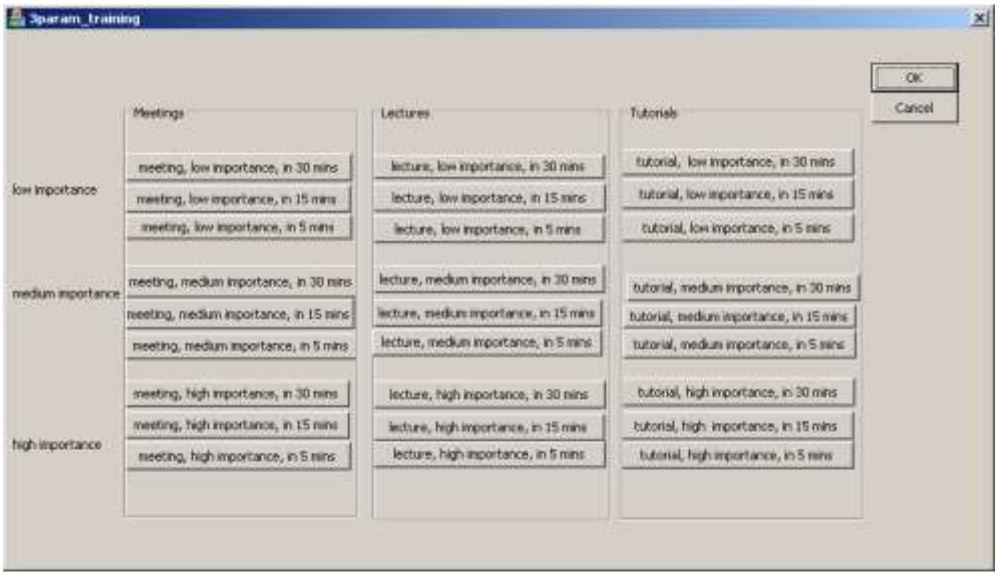

Please see the Experimenter.

E.4 Experiment 5(a), Three-Dimensional Tactons Experiment, Instructions

This experiment consists of 81 tasks. Each task will consist of you feeling one Tacton. You will feel the Tacton repeat 4 times (with a one second pause between repetitions). During this time you will need to identify all three attributes (the **type** of appointment, the **importance** of the appointment and the **time until** the appointment) encoded in the Tacton. As soon as you have identified these attributes, please select your answers by clicking on the corresponding radio buttons. There is no need to wait until the Tacton has been repeated all four times before choosing your answer. As soon as you have made your selection, and are happy with it, press the NEXT button.

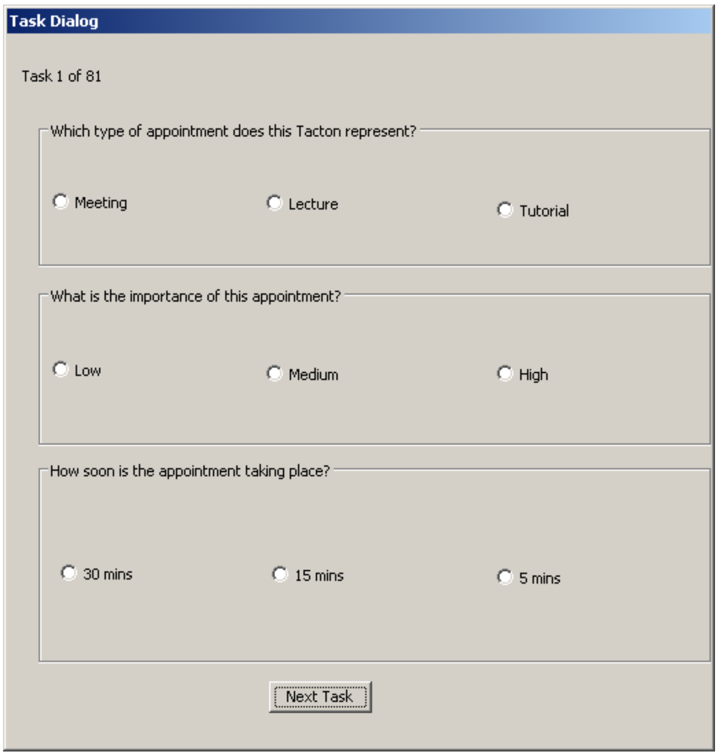

When you press NEXT there will be a short pause and then the dialog shown in below will appear. If you wish to pause for a short break, **you should do so before clicking this dialog.** When you are ready to start the next task, press the START button.

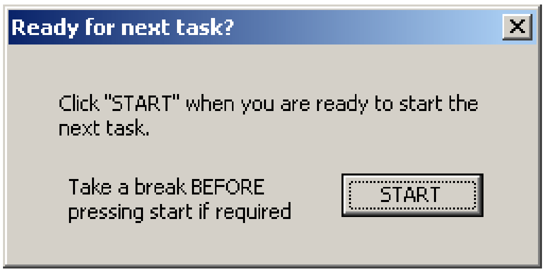

Before starting the experiment itself, you will do 27 of the above tasks as practise for the experiment. Please see the Experimenter.

E.5 Experiment 5(b), Three-Dimensional Tactons (2 levels of roughness) Experiment, Training Sheet

Each of the vibrations present in the experiment is called a Tacton, and represents an alert which might occur when on your phone or handheld computer to alert you to an upcoming event. Three pieces of information are encoded in each Tacton – the type of appointment, the importance of the appointment, and the time remaining before the appointment.

In order to successfully perform the experiment you must understand the vibrations used, and how to retrieve the three pieces of information that are encoded in each Tacton.

The Tactons are structured so that the *rhythm* encodes the *type* of appointment. The type of appointment can either be a **meeting,** a **lecture** or a **tutorial**.

The *roughness* of the vibration used to produce the Tacton encodes the importance of the appointment. There are **two** distinct roughnesses for **low** (smooth) and **high** (rough) importance appointments.

The *location* of the vibration on your arm encodes the time remaining until the appointment. A vibration presented near your **elbow** means that the appointment will takes place in **30 mins**, a vibration in the **middle of your arm** indicates that the appointment will take place in **15 minutes**, and a vibration at your **wrist** means that the appointment will occur in **5 minutes**.

In order to ensure that you understand the Tactons, all of the Tactons used in the experiment, along with their attributes are provided on a dialog box (see below). You can start each Tacton by clicking on the corresponding button. You will have 10 minutes to ensure you thoroughly understand the Tactons and their structures.

Please see the Experimenter.

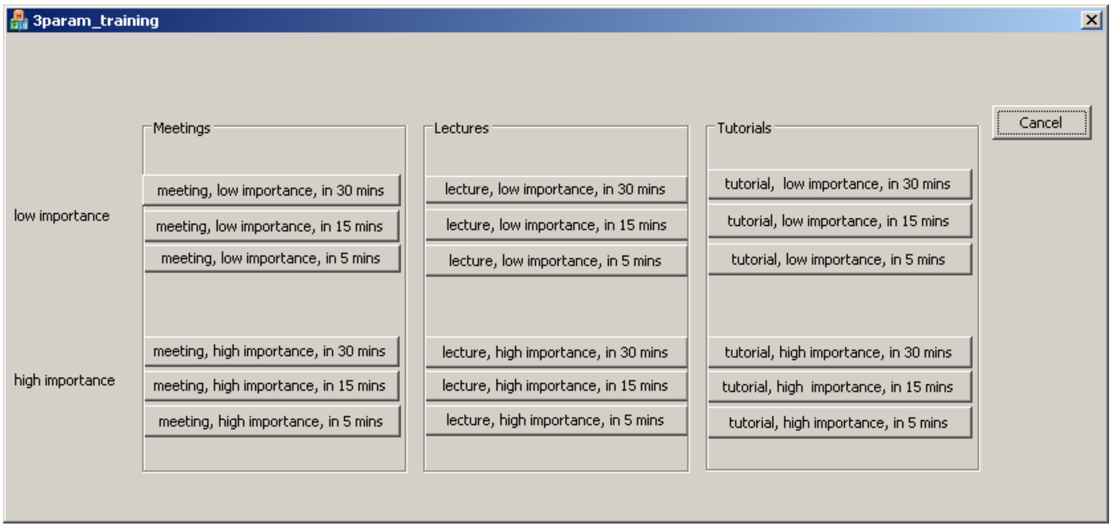

E.6 Experiment 5(b), Three-Dimensional Tactons (2 levels of roughness) Experiment, Instructions

This experiment consists of 54 tasks. Each task will consist of you feeling one Tacton. You will feel the Tacton repeat 4 times (with a one second pause between repetitions). During this time you will need to identify all three attributes (the **type** of appointment, the **importance** of the appointment and the **time until** the appointment) encoded in the Tacton. As soon as you have identified these attributes, please select your answers by clicking on the corresponding radio buttons. There is no need to wait until the Tacton has been repeated all four times before choosing your answer. As soon as you have made your selection, and are happy with it, press the NEXT button.

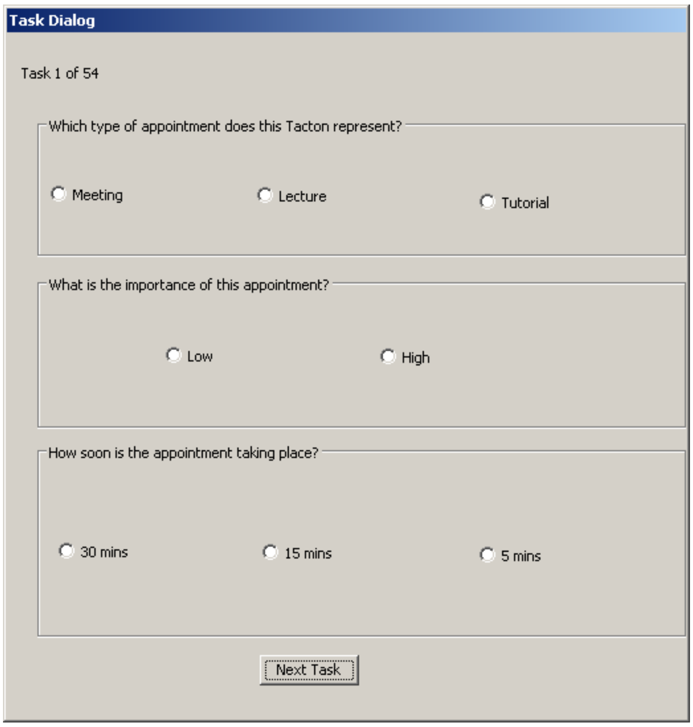

When you press NEXT there will be a short pause and then the dialog shown in below will appear. If you wish to pause for a short break, **you should do so before clicking this dialog.** When you are ready to start the next task, press the START button.

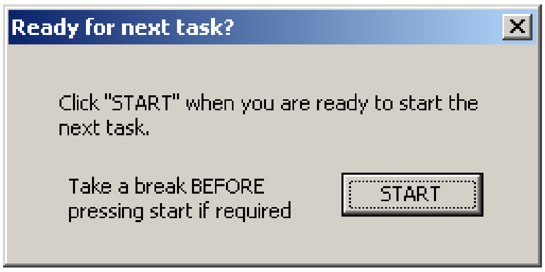

Before starting the experiment itself, you will do 18 of the above tasks as practise for the experiment.

Please see the Experimenter.

Appendix F: Experiment 6 Files

F.1 Experiment 6: Redesigned 3D Tactons Experiment, Detection Threshold Instructions

This experiment aims to discover how strong a vibration has to be before it is perceived. Therefore in this experiment you will be presented with vibrations and will be asked to respond with a yes or no answer to indicate whether you felt anything.

The vibrations will be presented to the forearm of your non-dominant arm (the arm you do not use to control the mouse) using a small vibration device.

You will be presented with the following dialog box. You should click the "Play Stimulus button) and then concentrate to see if you feel any vibrations on your arm (where the vibration device is attached). If you feel any vibration, select the "yes" radio button; if you do not then select the "no" radio button. Once you have responded with your decision, click "Next Task" to move on to the next task.

You can press "Play Stimulus" up to two times for each task. You will not be able to move onto the next task if you have not played the stimulus at least once, or if you have not made a response.

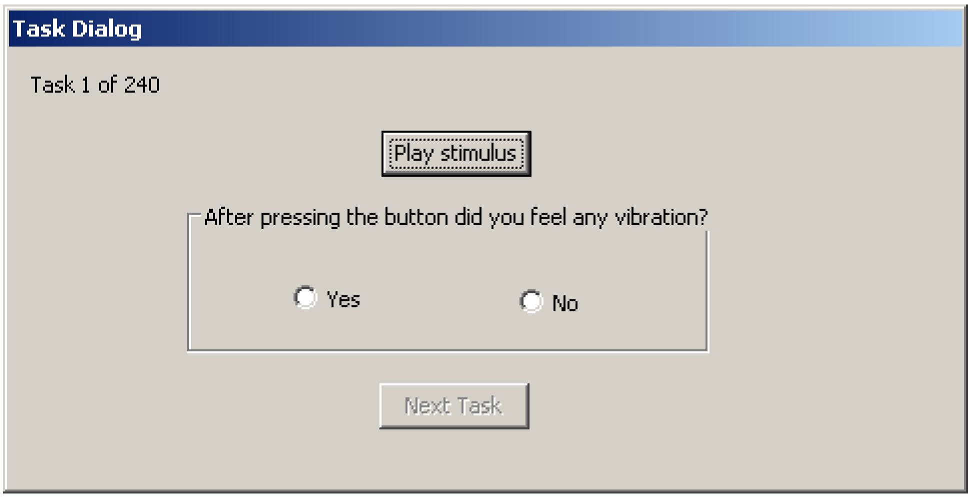

Please note that as this experiment is testing the limits of perception you may find that you do not feel anything for many of the stimuli. Do not worry about this, this is perfectly normal! Please concentrate as hard as you can during each task to ensure that you make an accurate judgement of whether or not you have felt anything.

F.2 Experiment 6: Redesigned 3D Tactons Experiment, Consent Form

This experiment aims to discover whether people are able to identify parameters of vibro-tactile messages called Tactons. These vibrations will be presented to three points on your forearm via small transducers that vibrate (the experimenter will show these to you). These will be held against your arm by Velcro bands.

This experiment consists of two sets of 54 tasks. In each task you will be presented with one Tacton asked to identify three parameters of the vibro-tactile message.

Some of the tasks are very hard so if you find it difficult to identify a parameter please just make a guess. You should not worry if you find this experiment difficult. We are trying to identify parameters which people are able to identify so if you are unable to identify them it will be good for us to learn this at this stage.

Since your skin is probably not accustomed to this type of vibrations you may experience slight tingling sensations in your arm during the experiment. If you feel any discomfort, pain or fatigue at any time please inform the experimenter immediately. We strongly recommend that you take regular breaks during the experiment, between tasks. You will receive a reward of £10 upon completion of the experiment.

During the experiment the computer will record data on your responses. This data will be stored anonymously using an ID number, rather than your name or any number that could identify you. All results will be held in strict confidence, ensuring the privacy of all participants. There will be no record kept that would allow your results to be traced back to you. All data will also be held securely in a password protected computer system (for electronic data) or a filing cabinet in a locked office (for paper based data). The data (including that from other participants) will be analysed to identify trends which may be generalised over the larger population. The results of this analysis may be published in appropriate scientific journals and conferences.

A feedback sheet (containing the summaries of the data mentioned above) will be sent to all participants who request it, after the data has been analysed.

Your participation in this experiment will have no effect on your marks for any subject at this, or any other university.

You may withdraw from the experiment at anytime without prejudice, and any data already recorded will be overlooked.

F.3 Experiment 6: Redesigned 3D Tactons Experiment, Training Sheet - Intensity Condition

Each of the vibrations present in the experiment is called a Tacton, and represents an alert which might occur when on your phone or handheld computer to alert you to an upcoming event. Three pieces of information are encoded in each Tacton – the type of appointment, the importance of the appointment, and the time remaining before the appointment.

In order to successfully perform the experiment you must understand the vibrations used, and how to retrieve the two pieces of information that are encoded in a Tacton.

The Tactons are structured so that the *rhythm* encodes the *type* of appointment. The type of appointment can either be a **meeting,** a **lecture** or a **tutorial**.

The *intensity (or strength)* of the vibration encodes the importance of the appointment. There are 3 distinct intensity levels for **low** (low intensity), **medium** (medium intensity) and **high** (high intensity) importance appointments.

The *location* of the vibration on your arm encodes the time remaining until the appointment. A vibration presented near your **elbow** means that the appointment will takes place in **30 minutes**, a vibration in the **middle of your arm** indicates that the appointment will take place in **15 minutes**, and a vibration at your **wrist** means that the appointment will occur in **5 minutes**.

In order to ensure that you understand the Tactons, all of the Tactons used in the experiment, along with their attributes are provided on a dialog box (see below). You can start each Tacton by clicking on the corresponding button. You will have 10 minutes to ensure you thoroughly understand the Tactons and their structures.

Please see the Experimenter.

F.4 Experiment 6: Redesigned 3D Tactons Experiment, Training Sheet - Roughness + Intensity Condition

Each of the vibrations present in the experiment is called a Tacton, and represents an alert which might occur when on your phone or handheld computer to alert you to an upcoming event. Three pieces of information are encoded in each Tacton – the type of appointment, the importance of the appointment, and the time remaining before the appointment.

In order to successfully perform the experiment you must understand the vibrations used, and how to retrieve the two pieces of information that are encoded in a Tacton.

The Tactons are structured so that the *rhythm* encodes the *type* of appointment. The type of appointment can either be a **meeting,** a **lecture** or a **tutorial**.

The importance of an appointment is encoded in the *intensity (or strength)* and the *roughness* of the vibration. There are 3 distinct intensity levels for **low** (low intensity, smooth), **medium** (medium intensity, rough) and **high** (high intensity, very rough) importance appointments.

The *location* of the vibration on your arm encodes the time remaining until the appointment. A vibration presented near your **elbow** means that the appointment will takes place in **30 minutes**, a vibration in the **middle of your arm** indicates that the appointment will take place in **15 minutes**, and a vibration at your **wrist** means that the appointment will occur in **5 minutes**.

In order to ensure that you understand the Tactons, all of the Tactons used in the experiment, along with their attributes are provided on a dialog box (see below). You can start each Tacton by clicking on the corresponding button. You will have 10 minutes to ensure you thoroughly understand the Tactons and their structures.

Please see the Experimenter.

F.5 Experiment 6: Redesigned 3D Tactons Experiment, Instructions

Each condition in this experiment consists of 54 tasks. Each task will consist of you feeling one Tacton. You will feel the Tacton repeat 4 times (with a one second pause between repetitions). During this time you will need to identify all three attributes (the **type** of appointment, **the importance** of the appointment and the **time until** the appointment**)** encoded in the Tacton. As soon as you have identified these attributes, please select your answers by clicking on the corresponding radio buttons. There is no need to wait until the Tacton has been repeated all four times before choosing your answer. As soon as you have made your selection, and are happy with it, press the NEXT button.

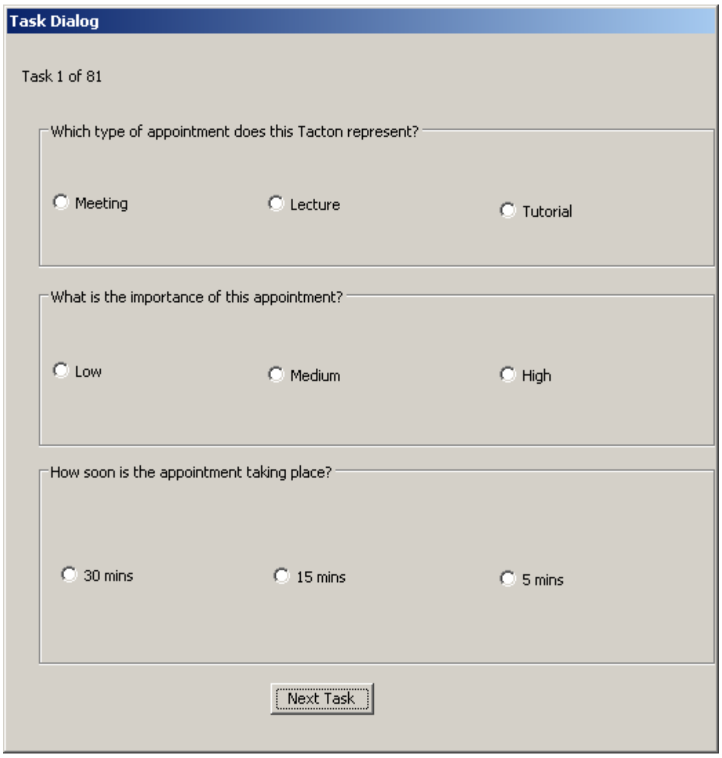

When you press NEXT there will be a short pause and then the dialog shown in below will appear. If you wish to pause for a short break, **you should do so before clicking this dialog.** When you are ready to start the next task, press the START button.

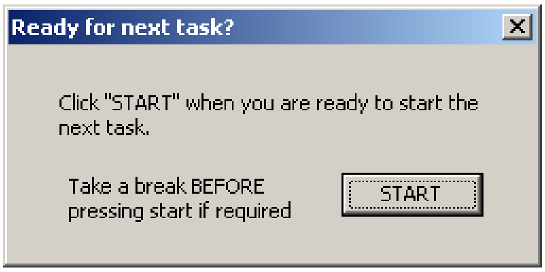

Before starting the experiment itself, you will do 27 of the above tasks as practise for the experiment. Please see the Experimenter.

Appendix G: Experiment Raw Data (on Accompanying CD)

Appendix H: Tactons Audio Files (on accompanying CD)