Utilizing Vibrotactile Feedback in Heart Rate Monitoring

Juuso Näsi

University of Tampere
Department of Computer Sciences
Computer Science
Master's Thesis
June 2008

University of Tampere
Department of Computer Sciences
Computer Science
Juuso Näsi: Utilizing Vibrotactile Feedback in Heart Rate Monitoring
Advisors: Roope Raisamo, Jukka Raisamo
Master's Thesis
June 2008

Heart rate monitors are becoming increasingly popular among athletes eager to gauge their performance and balance their training. Physical exercise requires concentration and usually puts our sight and hearing in use. Meanwhile, touch remains underutilized. Due to the extensive work-load of sight and hearing, heart rate monitors suffer from usability problems especially in motion.

This thesis reports a study investigating the possibility to utilize touch in conveying information between the device and the user during exercise. This idea brings up three major questions: can people detect vibrations effectively in motion, how does activity level affect recognition performance, and could touch-based interaction improve the usability of heart rate monitors? Since these kinds of questions require empirical research to be answered, a method of utilizing vibrotactile feedback in guiding interval training has been developed in this thesis. The method bases on a 20-minute interval training program implementation. The hardware produces different kinds of vibration sequences which guide the user to catch up certain activity levels.

A user experiment was conducted to reveal how the method performs. The subjects were able to detect 97,1 % of the vibration sequences and recognize 98 % out of the detected ones. The overall performance was 95,2 % which is calculated by dividing the number of correctly decoded vibration sequences with the total number of vibration sequences. The major findings are that the level of user activity affects only minimally the ability to detect vibrotactile stimuli. The users are also able to comprehend the conveyed information without problems. The conclusion derived from the observations and experiences also answers the initial question whether vibrotactile feedback could improve the usability of heart rate monitors? Yes, it could. The results suggest that vibrotactile feedback plays a key role in improving the interactions with heart rate monitors.

Keywords and terms: haptic, tactile, vibration, vibrotactile, heart rate monitoring.

Table of Contents

1. Introduction

Touch is one of the five human senses. Its importance is remarkable in everyday life since it enables us to locate objects in the space around us. Touch is clearly underutilized in human-technology interaction. Sight and hearing play a more fundamental role in the vast majority of user interfaces. Even though touch may seem an alternative modality it has also clear advantages such as discretion. In its ultimate need, touch is an irreplaceable information conveyor for people with seeing or hearing impairments. Touch-based interaction occurs when two or more objects have an effect upon one another. Further explanations regarding the processes behind sensing and particularly touch are presented in Chapter 2.

Heart rate monitors (HRM) are used to detect real time heart rate changes during physical exercise. They can be used to optimize training programs and gauge individual performance. The closely associated physiological processes are introduced at the end of Chapter 2. Moreover, the devices are discussed in greater detail in the same chapter. Most of the current HRMs support only visual and audible feedback. A usability study by Arjanmaa [2006] showed that both visual and audible feedback modalities have severe weaknesses. The problems occur especially in motion when the user attention should be channelled to exercise instead of the device. The visual screen is difficult to view during active sports such as running. Similarly, the audible beeps are not convenient in social situations and evidently troublesome in noisy environments.

How to overcome these usability problems? I argue that HRMs and other handheld devices could benefit more from touch. In my opinion, touch-based interaction can potentially improve the overall usability of HRMs. Currently, the user interfaces are too highly dependant on sight and hearing. The usability is further discussed at the end of Chapter 4.

Before discussing the actual research method a couple of definitions are presented to make sure that the connection between the concepts of touch and vibration is clear. The modality based on touch is called *haptic modality*. Another important term is *tactile feedback* which means that the information is perceived by touch. Most commonly tactile information is delivered by vibrotactile feedback. Even though vibration-based feedback would be more efficient in some cases than its visual or audible counterparts, there are many open questions to be considered. In my own analyses the hardware turned out to be the most crucial limitation when it comes to implementations. Similar and better documented findings were also reported by Poupyrev *et al.* [2002] (see

Chapter 3). Unfortunately, due to the scope of this thesis, the hardware limitations had to be accepted as such, and the best available hardware setup was chosen for the implementation. Instead, the open questions related to the usability of haptics in motion were covered in detail. Based on the initial hypotheses I ended up stating three major research problems. First, I was initially suspicious of whether people can detect vibrations effectively in motion. Second, it was not clear in what extent the activity level affects the recognition performance. Existing studies by Wheeler [2008] and Post *et al.* [1994] have already shown that vibrations can be detected in motion, but it should be taken into account that the results are not always transferable between various technologies. These studies are further reviewed at the beginning of Chapter 4. As a third research problem, aiming to make my results more practical and beneficial, I decided to work up a detailed suggestion regarding the usability issues. Even though Arjanmaa [2006] argued that HRMs would benefit from utilizing vibrations, there were no actual usability results proving that argument. In my opinion, it was initially still unclear whether vibrotactile feedback could improve the usability of HRMs. That is the third question my thesis answers.

A method of utilizing vibrotactile feedback in heart rate monitoring has been developed in this thesis. The method consists of an actual implementation of an interval training program which is guided by vibration-conveyed notification cues. Aiming to evaluate the recognition performance, an HRM was used to record the fluctuations in heart rate. A full description of the used hardware is presented in Chapter 5. The resulting heart rate data was the key indicator in analyzing how the method performs. Unfortunately, as already mentioned, these kinds of methods tend to be quite case specific. Therefore, the achieved results are not directly applicable to applications other than interval training programs, but still well suggestive in various contexts. Even though this particular implementation is quite hardware specific, the general concept can be easily adapted to other scenarios as well. The method leaves open many questions including the human memory capability and learning curve which are common limitations in designing more versatile vibrotactile applications. This brings up a question: what kind of method would be more universally applicable? This question is answered in Chapter 7.

The method was tested by conducting a novel experiment where 12 subjects participated in a 20-minute interval training program. The subjects were guided by vibrotactile notification cues which were produced by transforming audio signals into vibrations. The key indicators of success were the detection and recognition performances which both ended up being excellent. The results were positive across the board and proved the method viable. The experiment including its arrangements, goals, hypotheses, and achieved results is presented in full detail in Chapter 6. The results are discussed in Chapter 7 by comparing them to other related studies and intuitive hypotheses. Finally, the outcome of this thesis is further analyzed in Chapter 8 by making conclusion based on the major findings.

2. Human Physiology and Senses

The basics of human physiology are gone through in this chapter. The five human senses and their special characteristics are introduced in the first section. A deeper look at touch is taken in the second section. Last, heart rate monitoring is discussed in the third section.

2.1. Perceiving the Surrounding World

The human body can be seen as a collection of interacting systems having their own unique purposes. The nervous system is obviously the most interesting of them when it comes to senses. The central nervous system consists of the brain, the spinal cord, and the peripheral nervous system. The brain is the centre of all sensory processing and core activity. A *sense* is often defined as a faculty by which outside stimuli are perceived. The term *perception* means conscious understanding or meaningful interpretation of sensory data. It was already the Greek philosopher Aristotle who classified sight, hearing, touch, smell and taste as the conventional five senses. In addition to Aristotle's classification, humans are argued to have other senses the existence of which is somewhat disputed. The differing definitions of a sense cause that no common agreement on the number of human senses exists. There are also so called non-human senses such as electroception – the ability to detect electric fields – which is a unique property of some species of fish. Another commonly known non-human sense is called magnetoception. Many birds and insects have the ability to observe fluctuations in magnetic fields which is needed for navigational purposes. Similarly, bats are known to have the sense of echolocation to navigate in the dark and track their prey. [Boron and Boulpaep, 2005; Guyton and Hall, 2006]

Sensory receptors are the structures needed for detecting stimuli in the internal or external environment of an organism. Sensory receptors are also responsible for initiating *sensory transduction* which is a process of converting a stimulus into a nerve signal. Thereafter, the nerve signal can be forwarded to the brain by the central nervous system. Each sense has their own brain areas responsible for interpreting the nerve signals. According to Clare [1997] the functional organization of the brain is poorly understood, but the regions involved in sensory interpretation have been identified. The sensory receptors can be classified by their location, morphology, or adequate stimulus.

The various kinds of human-related sensory receptors are presented below [Boron and Boulpaep, 2005; Guyton and Hall, 2006]:

- **Baroreceptors** respond to the short term changes in blood pressure.
- **Chemoreceptors** respond to chemical stimuli. The senses involved include smell and taste.
- **Hydroreceptors** respond to changes in humidity.
- **Mechanoreceptors** respond to mechanical stimuli such as stretching or strain. Mechanoreceptors are mostly related to the skin, but other types such as hair cells exist as well. Mechanoreceptors are the most important type of sensory receptors related to touch.
- **Nocireceptors** (or shortly nociceptors) respond to potentially damaging stimuli and cause the sensation of pain. Nociceptors can be found in all pain-sensitive areas of the body including internal organs. Nociceptors can also trigger reflex actions.
- **Osmoreceptors** respond to the changes in osmotic pressure. They control the fluid balance of the body.
- **Phonoreceptors** respond to sound waves. They are required for detecting audible vibrations.
- **Photoreceptors** respond to light. Photoreceptors are specialized neurons found in the eyes. They are capable of phototransduction which is a process of converting light into neural signals.
- **Proprioceptors** respond to the stimuli regarding the position and movement of the body. They provide feedback only on the internal status of the body. For instance, proprioceptors make it possible to walk in a complete darkness without losing balance. Proprioception is also called the sense of body awareness.
- **Thermoreceptors** respond to temperature including both heat and cold. They are mostly cutaneous. Responses to some tissue-damaging temperatures are handled by nociceptors instead of thermoreceptors.

While most of us can perceive the world using all of the five human senses, some people lack these abilities. That can be because of physical impairments, sensory impairments, or cognitive impairments. Many people with impairments need special attendance depending on the extent of the perception loss and the sense concerned. Usually the remaining senses develop better and compensate for the weak or nonexistent ones. However, several everyday tasks are usually more challenging or even impossible without the applicable sense. Given that the majority of current user interfaces are based on single input and output modalities, many common devices are difficult to use without, for example, sight. Therefore, there should be different options available when it comes to interacting with computers, mobile devices, and other appliances. Unfortunately, these options tend to be less developed, more expensive, and often difficult to use.

2.2. Sense of Touch

Touch can be defined as the sensation produced by the mechanical, thermal, electrical, or chemical stimuli of the skin. Touch is one of the five major human senses. Touch can be felt all over the skin, but the sensitivity varies depending on the area in question. Touch is also a fast sense. Cholewiak and Collins [1991] argued that humans can detect two consecutive haptic stimuli presented within only about 5 ms interval. This record also bears comparison with other senses: touch is approximately five times faster than vision [Geldard, 1960]. All the other major senses have very limited physical areas with perceptual ability. According to Cholewiak and Collins [1991] the average size of the human skin is 1.8 m^2 making it the biggest human organ. The skin is also the heaviest human organ.

The skin consists of three major layers [Boron and Boulpaep, 2005]:

1. **Epidermis** is the outermost layer. It protects the surface by forming a waterproof wrap on the body. It renews quickly and the surface contains plenty of dead cells.
2. **Dermis** is the middle layer which is tightly connected to epidermis. Dermis consists of connective tissue. It contains blood vessels, skin receptors and nerve endings providing the sense of touch.
3. **Subcutis** (or subcutaneous tissue) provides storage for nutrients. It is physically beneath dermis. Thickness varies depending on the individual and the body area.

An overview of the structure of the skin is illustrated in Figure 1. The winding stripes indicate the nerves. Most of the sensory receptors are located just beneath epidermis. These receptors perceive different kinds of stimuli such as heat, cold, light touch and pain. Pain receptors (nociceptors) also perceive heat and cold in case of potentially tissue-damaging temperatures. The sensory receptors perceiving strong pressure and hair movement are located deeper in the tissue. While the skin looks relatively same all over the body, the perceptual abilities differ quite a bit depending on the area in question. Several other factors such as the concerned skin layer and the receptor type affect the sensitivity.

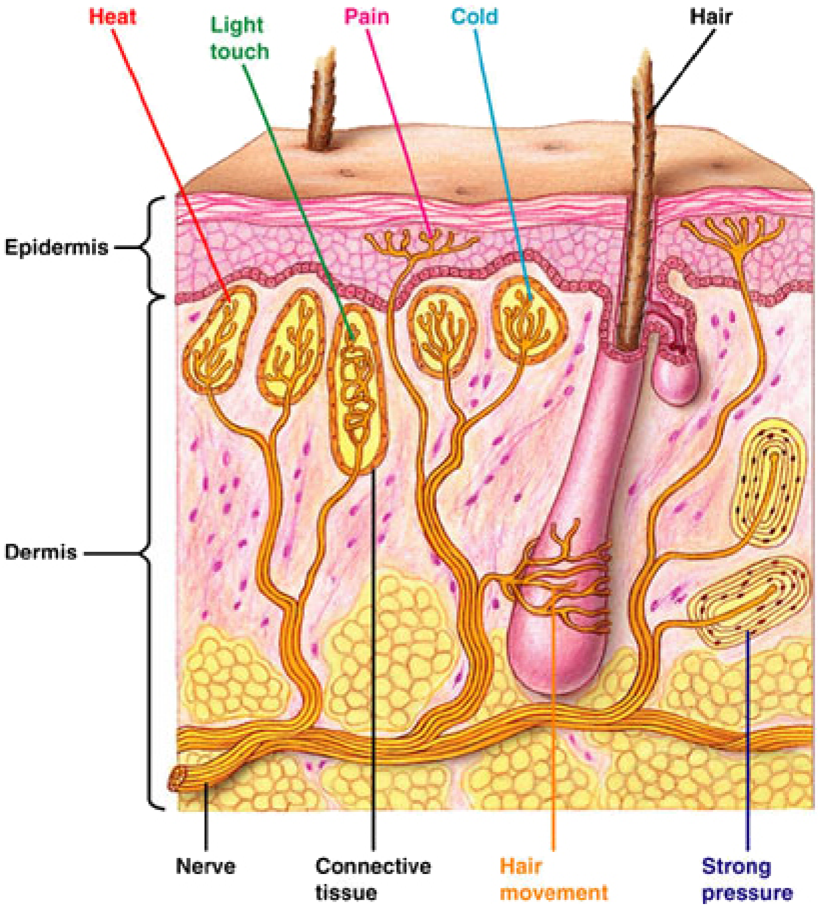

Figure 1 - Receptors of skin [Mallery, 2008].

The term *touch* is often used as a combining term for several actually different senses. Particularly in the field of medicine, touch is often replaced with *somatic senses* in order to better reflect its multiple meanings. Respectively, another term called *somatosensory system* is commonly used to take cognizance of the different senses involved. It can be used to describe the whole sensory system associated with the body. The human skin contains tailored types of sensory receptors capable of perceiving different kinds of stimuli. Particular types of receptors activate themselves depending on the type of stimuli. Somatic modalities can be classified into four types including pain, mechanical touch, thermal sensation, and proprioceptive sensations [Mazzone *et al.*, 2003]. These can be further split in submodalities making it possible to distinguish various kinds of tactile sensation.

Somatic modalities are classified as follows [Mazzone *et al*. 2003]:

1. **Proprioceptive sensations** are elicited by the displacements of joints and muscles. There are two major types of proprioception: the position sense (static) and kinesthesia (dynamic). Kinesthesia is the sensory experience stimulated by bodily movements and tensions. Kinesthesia is sometimes used interchangeably with proprioception, but sometimes it is differentiated by excluding the sense of balance. There is a scientific disagreement regarding the definitions of kinesthesia and proprioception.
2. **Pain** is elicited by tissue-damaging stimuli. Pain receptors – nociceptors – can be divided into three types. Mechanical nociceptors detect strong stimuli caused by mechanical forces such as punching or sharp objects. Thermal nociceptors detect extreme heat or cold. Polymodal nociceptors detect several kinds of damaging stimuli including mechanical, thermal, and chemical ones [Kandel *et al*., 2000].
3. **Thermal sensation** can be caused either by heat or cold stimuli. According to Kandel *et al*. [2000] the receptive fields range from 1 to 2 mm in diameter. Each cm^2 of skin has the density of 1-5 receptors for cold and about 0.4 for warmth. Cold-sensitivity can be considered significantly more accurate compared to warmth-sensitivity.
4. **Mechanical touch** is sensed by mechanoreceptors. Mazzone *et al*. [2003] mentioned five different types of mechanoreceptors for hairless skin and a sixth one for hairy skin. The most important subtypes of mechanoreceptors include Meissner corpuscles, Merkel's disks, Pacinian corpuscles, and Ruffini corpuscles. They all have different functions, frequency ranges, and adaptation rates as presented in Table 1. Their physical locations also range from epidermis to subcutis. Mechanoreceptors have different temporal and spatial resolutions which are the key variables in analyzing the perception capability.

Receptor	**Meissner Corpuscles**	**Merkel's Disks**	**Pacinian Corpuscles**	**Ruffini Corpuscles**
Location	Glabrous skin: Dermis	Glabrous skin: Epidermis	Glabrous and hairy skin: Dermis and subcutaneous	Glabrous and hairy skin: Dermis and subcutaneous
Adaptation rate	Rapid: RA-I	Slow: SA-I	Rapid: RA-II	Slow: SA-II
Receptive field	Small $12mm^2$	Small $12mm^2$	Large $100mm^2$	Large $60mm^2$
Frequency range	20 - 100 Hz	0 - 10 Hz	100 Hz – 1 kHz	0 - 10 Hz
Best response	30 - 40 Hz	0 - 10 Hz	150 - 300 Hz	0 - 10 Hz
Proportion	40 %	25 %	13 %	19 %
Function	Movement, Velocity	Pressure, Vibration	Acceleration, Pressure	Pressure, Skin shear, Thermal Changes

Table 1 - Characteristics of mechanoreceptors, adapted from Mazzone et al. [2003].

Mazzone *et al.* [2003] argued that the common temporal resolution of mechanoreceptors is 5 milliseconds. However, temporal resolution rather depends on the type of stimulus than the type of sensory receptor. Spatial resolution depends on the skin area in question. Certain parts of the skin such as fingertips have more receptors than others. The number and type of receptors are important factors regarding sensitivity. The type of receptors also defines how deep they are located in the tissue. Those receptors located in the outermost parts of the skin have the smallest receptive fields. The size of receptive field determines the ability of mechanoreceptors to resolve the size of an examined object. In other words, it defines the accuracy of tactile perception. Zimbardo *et al.* [1995] claimed that, for example, fingertips sense the location of a stimulus ten times more precisely compared to the human back.

2.3. Heart Rate Monitoring

The human body encounters extreme stress during heavy physical exercise. The body must be able to make rapid organ-system adjustments to meet the demands caused by the various kinds of physical activities. One of the major challenges is to provide enough oxygen-rich blood for the muscles during the exercise. In order to understand the basics of this process a few terms need to be introduced. *Heart rate* indicates the frequency of cardiac cycle which is the event-sequence repeated with every heartbeat. Another important term is *stroke* *volume* which measures how much blood is pumped by each heartbeat. The total volume of blood pumped by the heart – *cardiac output* – can be calculated by multiplying stroke volume by heart rate. When it comes to physical exercise, cardiac output is the most important measurement of oxygen transmission. However, heart rate carries more weight in this thesis as it is the most commonly measured physiological data during exercise. [Boron and Boulpaep, 2005]

Recording heartbeats in order to detect a number of cardiovascular disorders has a long history in medicine. Heart diseases remain a common cause of death. *Electrocardiography* is a non-invasive diagnostic procedure to record electrical changes in the heart and produce an *electrocardiogram* (ECG). ECG is a graphic indicating the overall rhythm of the heart. Heart rate can be calculated as the number of heartbeats per minute (BPM). Typical heart rate ranges between 60 and 100 BPM in case of healthy adults [Boron and Boulpaep, 2005]. The rate can drop to 40 BPM during the sleep. In the other extreme, it can rise as high as 220 BPM during exhausting exercise. In medicine, the body response to exercise is normally measured by changes in cardiac output instead of heart rate. In numbers, exercise can increase cardiac output up to 4-5 times the resting cardiac output. These values should be considered rather suggestive as varying ranges are presented by different medicinal books. Athletes usually have a lower rest heart rate compared to above mentioned averages. According to Guyton and Hall [2006] the effectiveness of each heartbeat is 40-50 % greater in highly trained athletes compared to untrained persons; so that there is a corresponding decrease in heart rate at rest. The differences are presented in Table 1 by comparing nonathletes with marathoners. Training does not increase the maximal heart rate, but instead, it increases the maximal stroke volume which means providing more oxygen-rich blood per heartbeat for the muscles [Boron and Boulpaep, 2005]. While training decreases the lowest achievable heart rate – even though individually – also the maximal heart rate may slightly decrease as shown in Table 2. In addition, the maximal heart rate is also affected by natural factors such as aging [Guyton and Hall, 2006].

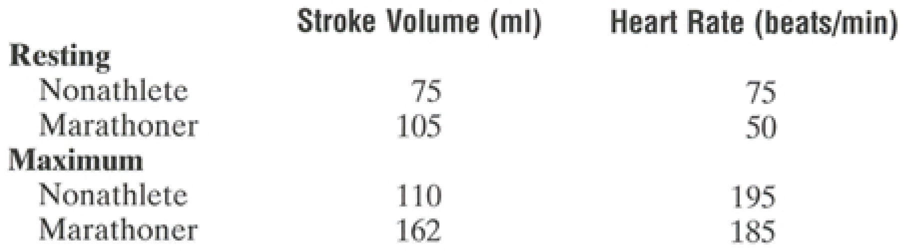

	Stroke Volume (ml)	Heart Rate (beats/min)
Resting		
Nonathlete	75	75
Marathoner	105	50
Maximum		
Nonathlete	110	195
Marathoner	162	185

Table 2 – Comparison of cardiac function between nonathlete and marathoner [Guyton and Hall, 2006].

The term *pulse* is often used as a method of measuring heart rate. To be exact, it is incorrect to denote the frequency of heartbeat by pulse. Some cardiovascular disorders cause pulse to be an inaccurate measurement of heart rate. However, pulse and heart rate are usually equal in case of healthy people.

A heart rate monitor (HRM) is an electrical device which is built to detect real time changes in heart rate. In addition to advanced medical use, HRMs are getting increasingly popular among athletes. The history of HRMs for fitness purposes tracks back to the late 70's when Polar Electro Ltd. pioneered the development of wireless heart rate measurement [Polar Electro, 2008]. In general, monitoring can be conducted during almost any kind of sport as physical exercise affects heart rate rapidly. The final determinants of success in sport are the strength and endurance of the muscles.

Interval training is a widely practiced training technique which has been found effective in, for example, cardiovascular buildup. Typically, interval training consists of repetitions of high intensity activity followed by periods of rest or low activity. Interval training plans are most effective when designed individually. HRMs provide an accurate tool to gauge individual performance and optimize training plans. A modern HRM can be used to determine the safest and most effective heart rate zones for each bout of interval exercise. HRMs can assist in keeping the right balance between light, moderate, and hard exercise. This balance does not only provide optimal results, but also makes training physically easier in the long run.

3. Haptics in Human-Technology Interaction

The term *haptic* means being based on touch. Haptic modality is discussed in this chapter by analyzing its pros and cons regarding various application areas. Special attention is given to the mobile context as the method presented in this thesis studies enabling haptics in motion. This chapter contains three different sections which are summarized below.

First, different kinds of user interfaces are shortly discussed to give an idea of where haptics could potentially be used. The disadvantages regarding the traditional user interfaces are analyzed and ideas are proposed in order to improve the designs.

Second, the existing haptic actuator technologies are reviewed. *Haptic actuators* are the devices producing stimuli on the skin. The stimuli are elicited by physical displacement, temperature-changes, or electricity produced by haptic actuators. In addition to describing the existing technologies, some scenarios of tomorrow's actuators are discussed at the end of the section.

Third, particular attention is given to vibration as an information conveyor. The concept of forming languages based on haptic modality is discussed in this section. In addition, the major parameters affecting the properties of vibrotactile feedback are introduced.

3.1. Beyond the WIMP User Interfaces

Traditional user interfaces – often called as WIMP (Window, Icon, Mouse, Pointer) interfaces – are widely used in various computing devices. The term *WIMP* is often used as a synonym for *graphical user interface* (GUI). In this thesis WIMP is favored due to its better reflection of the variety of interaction elements involved. WIMP interfaces are based on single input and output channels. The input method is haptic. As already mentioned, haptic modality is based on touch. Another related term - *haptics* - means physical interaction via touch including both sensing and manipulation of objects. Haptic modality can be divided into tactile, vibrotactile, force, thermal, and electrical submodalities. The typical input devices in WIMP interfaces include a keyboard and a pointing device. They are reliable input devices from the technical point of view and most problems are related to human factors such as user accuracy [Bradford, 1995]. The reliability is better compared with, for example, speech recognition. In general, haptic input can be considered to be the most efficient option when it comes to desktop environments. While the input method of WIMP interfaces is haptic, the output is mostly provided for eyes only. Vision provides an efficient way to keep track of the constantly changing state of the used application. Visual feedback has some unique properties compared to its auditory and haptic counterparts. The content can be re-examined anytime and therefore lesser short term human memory is needed to get a general insight. Graphical screens can also provide a complex set of data without excessively increasing the cognitive load of users.

Despite having several indisputable benefits, WIMP interfaces have also downsides. The major limitations are related to the size and portability of modern computing devices. As already mentioned before, haptic input and visual output have proven to be efficient in desktop environments. Meanwhile, mobile devices have developed capable of running more applications causing traditional interaction methods to face usability challenges. Applications such as word processors and web browsers require more sophisticated interface design in order to maintain a sufficient level of usability. As a simple example of an output challenge, traditional web pages do not fit on the screens of mobile phones. Similarly, giving input is slower in the mobile world due to the small keypads. Another significant challenge regarding WIMP interaction is related to special user groups [Baca, 1998]. WIMP metaphor has created a huge barrier for people with visual impairments. Visual feedback is the only way to perceive output from WIMP user interfaces making them impossible to be used by blind or visually impaired persons. Inputting information can be considered difficult for people with motor impairments. Because of the above mentioned reasons some people are unable to gain access to the information provided by WIMP user interfaces especially in the mobile context.

Haptics could potentially be used to reduce the heavy reliance on visual feedback and solve some of the above discussed problems. Currently haptics is relatively underutilized as an output modality. Even though the importance of haptics in user interface development is well recognized, there are not many meaningful applications for the average user. The most common haptic application areas include games, simulations, 3D modeling, training of medical operations, and special applications for people with impairments. The current state-of-the-art haptic feedback environments (e.g. Phantom [2008] by SensAble technologies) are complicated systems with customized hardware resulting in that their prices are not reasonable for any private usage. Considering any Phantom-like haptic feedback in the mobile environment is not realistic as the hardware requires grounding. This kind of sophisticated haptic feedback is seldom commercially utilized apart from research purposes and applications targeted to special groups such as visually impaired people. To be exact, even regular keyboards actually provide certain kind of haptic feedback when, for example, a button is pressed. This is called *passive haptics* [Insko et al., 2001]. The problem is that passive haptic feedback is determined only by the mechanical qualities of the input device. Therefore, it does not tell anything about the application concerned. The only advantage is that the user is able to feel that the button really went down. However, from now on, what is called haptic feedback in this thesis should not be connected to passive haptic feedback. As mentioned before, the majority of the current user interfaces still lack of any kind of controlled haptic feedback. While visual and audible interaction is possible over a distance, haptic user interfaces necessitate the presence of the user. Haptic feedback usually requires a direct contact between the device and the user. As an exception, the sensory receptors of the skin respond similarly to strong air pressure as well as to playing loudly low audible tones. As one would expect, there are numerous ways to produce stimuli which are perceivable by the human skin.

The usability of WIMP interfaces could be improved by adding new interaction modalities to replace or support the existing ones. Making use of haptics, gaze, speech, or manual gestures could be beneficial in many applications. In my opinion, these alternative modalities - including also haptics - have a limited capability to manage all interaction in mobile devices. Instead, they could be effectively used as supportive modalities to boost the appropriate tasks. Particularly in the area of mobile computing, multimodal and crossmodal solutions could potentially solve some of the problems mentioned above. The desirability of a certain modality always depends on the context and can be user-specific. Reeves *et al*. [2004] listed age, preferences, skills, and sensory impairments as the affecting factors for individual differences. Providing alternatives, such as haptics, would help especially special user groups. However, it requires some efforts to prepare the way for completely new modalities in the existing systems. It is difficult to justify integrating haptic technologies with mobile devices until there is a broader understanding of the additional value these technologies are able to provide for the user. The trade-off between the cost and the additional value is quite an application-specific question, but the demand should be answered especially in case of the principal applications. However, it must be taken into account that the alternative modalities are still quite challenging to enable in the traditional interfaces. On the other hand, in a well-designed system, all the modalities used share the same back-end functional logic. Preferably all the modalities should have their own tasks and they could be modified and personalized by the user. Reeves *et al*. [2004] recommended that information should never be presented by using two different modalities if the user has to attend both channels at the same time to comprehend the information. Ideally, the multimodal interaction is meant to make things easier without increasing too much the cognitive load of the user.

3.2. Haptic Actuators

Haptic actuators are the actual devices generating stimuli. Their purpose is to deliver information from the device to the skin. The sensory receptors of the skin perceive the physical displacement produced by the actuators. The most common type of tactile feedback is vibration which can be produced by vibrotactile transducers (or shortly tactors). *Vibrotactile transducers* can be considered to be a specified subgroup of haptic actuators. The term *tactor* is used later in this thesis to better reflect the form of the produced feedback. When it comes to this section I refrained from using additional terms besides haptic actuators to keep things as simple as possible.

Before having a look at the actual technologies, a few design issues could be shortly discussed. As mentioned earlier, the current mobile devices are generally able to produce only very simple tactile feedback. The reasons behind that are mostly related to the limitations regarding actuator technologies. These limitations also affect the willingness of manufacturers to enable haptics in mobile devices. The underdeveloped actuator technologies can be considered to be the main bottlenecks in utilizing richer tactile feedback in mobile devices. Poupyrev *et al.* [2002] listed four feasibility requirements for a workable actuator. Similarly, they listed four human factor requirements which an optimal tactile display should satisfy.

> Haptic actuators have following kinds of requirements in the mobile context [Poupyrev *et al.*, 2002]:
> 1. miniature size,
> 2. lightweight,
> 3. low voltage (~5V), and
> 4. easy to customize (platform independent).
>
> Human factors result in the following key requirements [Poupyrev *et al.*, 2002]:
> 1. fast response time,
> 2. variable intensity,
> 3. wide frequency bandwidth (ideally 1 to 1000Hz), and
> 4. multitude of different wave shapes.

Poupyrev *et al*. [2002] mentioned that there is no such actuator which meets all these criteria. Each available technology fails in at least one of the listed requirements. According to Chang and O'Sullivan [2005], there are three major vibration techniques which are commonly used in the commercially available handheld devices. Each of them is capable of producing vibration sequences which are used in the experiment presented in this thesis. In addition to those, there is a great number of other techniques to produce tactile feedback. They are also shortly discussed in this section.

The most common technology is called Rotary Mass Vibrator (RMV) which can be used to produce vibration effect with adjustable duration. However, the lack of finer control makes it difficult to produce more detailed feedback. RMV technology is relatively inexpensive and does not require a lot of computing power. RMV actuators are small and therefore suitable for mobile devices. The basic RMV allows only on-off vibration, but it has been used as a core technology for more sophisticated approaches. A significant and recently released RMV-based technology called VibeTonz by Immersion [2008] provides more advanced vibration features. In comparison with the basic RMV, the biggest differences are the adjustable intensity and the possibility to spin the electric motor backwards. These improvements in the control allow utilizing tactile feedback in completely new ways and developing more diverse applications. VibeTonz can best be put to use in mobile games, messages, call alerts, and ringtones. New breed of touchscreen phones take VibeTonz further by providing navigational aid such as button and scrolling feedback [Immersion, 2008]. Nevertheless, the haptic feedback produced by the VibeTonz can be described versatile only among other mobile-related vibration technologies.

The second notable approach to producing vibration is based on slim piezoelectric actuators [Tikka and Laitinen, 2006]. They can be used to produce very precisely controlled vibrotactile feedback. Poupyrev *et al*. [2002] discussed a novel piezoelectric actuator called TouchEngine. It is built of two thin piezoceramic film layers with printed adhesive electrodes in between. The piezoceramic material either shrinks or expands depending on the polarity of the applied voltage. The top and bottom layers have the opposite polarity with each other causing the whole structure to bend when a voltage is applied. TouchEngine actuators are small, extremely thin, and have also fast response time. However, due to its extremely small size, the produced physical displacement is as low as 0.1 mm. Therefore, in my opinion, TouchEngine needs some further development to become more practical.

The third major approach – Multi Function Transducer (MFT) technology - is founded on a speaker which can be used to produce both vibrotactile and audible output [Chang and O'Sullivan, 2005]. MFT technology is used in several Motorola phones. The haptic effects can be felt when playing low frequencies between 100-300 Hz. The frequencies around the resonant peak are able to produce relatively strong stimuli, whereas frequencies higher than 300 Hz are heard just as normal audio. MFT technology enables audio-haptics which means the synchronicity of sound with vibration. Audio-haptics can be used to enhance ringtones, mobile games, or other applications with haptic effects. Chang and O'Sullivan [2005] claimed that MFT-actuators have several advantages such as low latency, small size, and power efficiency. They conducted a user experiment to compare an MFT-enabled phone with a regular non-haptic phone. The results clearly indicated that the subjects preferred the MFT-phone and liked the vibration effects. In addition, a fair majority of the subjects considered the audio output produced by the MFT-phone better than the audio produced by the normal phone. Chang and O'Sullivan drew the conclusion that the presence of haptics increases perception of sound. This might be true, but the physical qualities of the tested phones were so different that there might be other affecting factors as well. Töyssy [2007] also questioned the results because of the countless differences in the compared devices. However, whether or not the results are skewed by unequal hardware, they still prove that the users feel comfortable using MFT-enabled devices.

While RMV, piezoelectric, and MFT techniques are the leading solutions for producing vibrations in mobile devices, voice coils are an alternative and well-portable solution [Eibeck and Muramatsu, 2007]. According to a technical patent by Cranfill *et al.* [2006] also some MFT-actuators utilize a kind of voice coil technology, but other sources including Chang and O'Sullivan [2005] separate these two techniques. Voice coils are a special form of electric motor which utilizes a coil of wire in a permanent magnetic field [Thibadeau, 2008]. There are linear and rotary voice coils which produce different vibration effects. According to Töyssy [2007] linear voice coil produces a stroking motion comparable to piston movement in a car engine. The displacement of rotary voice coils is similar to RMV actuators. Voice coils are commonly used as positioning devices because of their great accuracy and extremely high acceleration speed. As a simple example, in some modern hard disks the rotary voice coil actuators are used for positioning the disk heads across the platter of the disk [Thibadeau, 2008]. The biggest limitation is that the physical displacement generated by the voice coils is relatively low. Some voice coil actuators can even be used as audio speakers if the signal frequency matches up to the natural frequency of the actuator [Fukumoto and Sugimura, 2001]. In other words, the same actuator can be used to generate both audible and tactile feedback. Technically, the structure of voice coil actuators is quite similar in comparison with regular audio speakers. There is a contactor outside the device case which is actually the only vibrating part of the device (see Figure 2). The coil is located in a magnetic field and pushed along its axis when a current passes through the coil. This procedure causes vibration of the contactor. It is noteworthy that the vibrating skin contactor is relatively small in size, whereas the non-vibrating case is large by contrast. A voice coil actuator was used also in the experiment described in this thesis. That particular actuator is further discussed in the fifth chapter.

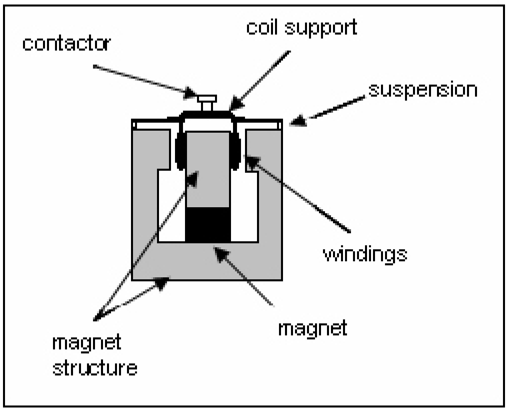

Figure 2 - Typical construct of a voice-coil transducer. Adapted from [Brown et al., 2005] and [Cholewiak and Wollowitz, 1992].

The four introduced haptic actuator technologies are well-portable and therefore the most suitable for producing vibration while in motion. There are also several other types of actuators for varying purposes. In theory, some of them could be used to produce vibrotactile feedback in the mobile environment as well.

Electrostatic displays generate an electrostatic force when an insulated metal plate is energized [Yamamoto *et al.*, 2004]. The surface of the metal plate and the layer of the conductive substance under the skin form a condenser if the actuator is touched. Electrostatic actuators can be very compact and therefore useful in micro scale applications. As a drawback, the skin needs to be dry in order to perceive the electrostatic forces; according to Mallinckrodt *et al.* [1953] a superficial water layer may form a condenser together with the electrode. This effect channels the electrostatic force into the water instead of the skin.

Pneumatic actuators can be used to stimulate the skin by suction or air-pressure. The most common technique is based on compressed air. One of the advantages of pneumatic actuators is that vibratory stimulus can be targeted to a very specific area of the skin. Enriquez *et al*. [2001] developed a tactile alerting system for driving environment. By using pneumatic actuators they were able to localize the stimuli into the palm instead of shaking the whole steering wheel. The produced stimuli vary depending on the configuration and shape of the pneumatic pocket. The pneumatic systems are usually noisy and require lots of equipment such as an air compressor and a pressure regulator. Because of that they are not very convenient in portable devices.

Solenoid actuators are capable of delivering high-energy impulses and producing relatively strong forces. However, solenoids are power wasters compared to most of the other types of actuators. While the poor power consumption is an undisputable downside, Lee *et al*. [2004] found their size, cost, force, reaction speed, and expressive capabilities better in comparison with the traditional linear vibrating motors and piezoelectric actuators.

Shape memory alloys (SMA) are metallic actuators which restore their geometry when an electric current is applied to them. When resistances are encountered during this transformation, they are able to produce strong forces which are useful in, for example, force feedback devices [Mazzone *et al*., 2003]. They require a long time (from seconds up to a minute) for relaxation and therefore they are useless for applications requiring fast response times.

Likewise, voltage can be used to control electrorheological fluids which change their viscosity between liquid and solid gel. Respectively, the applied current defines the shape of electroactive polymers in order to produce tactile feedback. Even plain electricity could be used for stimulation. However, the sensation is usually unpleasant when the mechanoreceptors are fired by an electric current. The electric shocks would be rather irritating and even cause pain.

In summary, each actuator technique has some downsides at the moment. These downsides are largely the same that Poupyrev *et al*. [2002] listed as the technical feasibility requirements for actuators. Overcoming the remaining problems related to the size, weight, power consumption, and customization would be the next quantum leap in hardware development. In my opinion the size of the actuator is currently the most crucial limitation regarding mobile haptics. I believe that most users prefer the small device size over the quality of haptic features. When it comes to the weight of the actuator it depends on the use case whether it causes problems. Weight can be considered a conditional requirement, whereas power consumption is virtually always critical in the mobile context. It is likely that the battery life is more of a concern than the lack of haptic feedback for most users. There are differences in power consumption between various technologies though. For instance, according to a study by Perez *et al*. [2003], the power consumption of piezoelectric actuators was found significantly lower in comparison with the electromechanical actuators. Perez *et al*. discovered also that electromechanical actuators require two orders of magnitude difference in power to excite the skin at the same level as piezoelectric actuators. Their results are probably well suggestive even though there are certainly differences between the various actuators implementing the same core technology. Many of the related studies estimated the power consumption theoretically instead of conducting actual measurements. Therefore, making any further conclusion based on the existing research is difficult. In addition to the size, weight, and power consumption, also the possibility of customization is a critical requirement. Customization can be seen more of a commercial issue than the other requirements. I believe the best results could be achieved by prioritizing the feasibility requirements depending on the properties of the device. It is most likely possible to find an appropriate actuator for each purpose of use by making some requirement concessions. However, as already mentioned, the requirements are quite device-specific. The most applicable actuator technology usually depends on the context.

It is difficult to predict which actuator technology will be the most successful in the future, but in my opinion piezoelectric technique is the best candidate out of the discussed approaches. At the moment, mobile devices such as mobile phones, PDAs, and mp3 players are already a prominent part of our everyday life. They are constantly developing and new features are introduced. RMV-based vibrating motors are almost a standard feature in mobile phones while more advanced solutions are still scarce. The existing mobile haptic applications are also somewhat underdeveloped. As discussed earlier, desktop computers are clearly ahead in utilizing tactile feedback and, for example, many computer games come with force feedback devices providing fairly versatile feedback. Fortunately, the differences between traditional computers and handheld devices are already quite blurred, and could be almost nonexistent in the future. The support for more detailed tactile feedback is becoming a reality when smartphones develop more powerful. Concurrently, while hardware develops further, user interfaces need to change alongside without question.

3.3. Conveying Information by Vibration

Depending on the use case, vibrotactile feedback can be anything from simple notifications to a complex signal language. At the simplest, a device vibrates for a defined period of time informing the user about a certain state or event. On the contrary, more sophisticated feedback is a stream of information which can be divided into independent signals. Each signal has parameters which can be used to differentiate the signals from each other. These vibrotactile parameters are further discussed later in this section. There are several human factors which affect the actual functionality of vibrations. While small sets of vibration sequences are usually easy to distinguish, more complex vibrotactile languages may require extensive learning. Learning is required in haptic languages in a similar manner as one needs to know the letters of the alphabet in order to read a text. Another affecting human factor is the limited capacity of the human working memory. This was first argued by the cognitive psychologist Miller [1956] who is renowned for his seven-plus-minus-two argument. Miller examined a number of cognitive tasks and each time the effective channel capacity was found to range between 5 and 9. The human working memory sets severe limitations on the complexity of tactile patterns. Vibrotactile feedback is normally presented in a continuous form and there is none or very limited control available to traverse backwards in the stream. Hence, the information disappears right after its initial occurrence. Neither haptic nor audible streams can be re-examined without a particular interface, and the feedback based on these modalities is thereby quite challenging to design.

Despite its potential in improving the usability of user interfaces, vibrotactile feedback has also limitations caused by both technical and human factors. Many of the disadvantages become more crucial in case of mobile devices and concurrent movement. The sensitivity of detecting vibrotactile feedback decreases during physical exercise according to Wheeler's [2008] experiment, which is described in full detail in Chapter 4. Similar findings were also made by Post *et al*. [1994] in their research on human ability to detect, discriminate, and scale the near-threshold vibrotactile stimuli. The results by Wheeler [2008] showed that stimuli can be well perceived also in motion if the signal amplitude is increased. On the other hand, it could make the feedback more uncomfortable and even annoying.

Tactons are structured, abstract tactile messages which are used to convey information non-visually [Brewster and Brown, 2004]. In the visual domain there are icons which are the counterpart of text. In the tactile domain there is Braille which is a tactile code developed by Louis Braille to represent the letters of the alphabet. Brewster and Brown suggested that tactons can be considered the missing iconic counterpart of Braille. Shneiderman [1998] defined an icon as *"an image, picture or symbol representing a concept"*. Tactons can be used to describe interface concepts in a similar manner. They are comparable to visual icons and audio earcons and could be used to replace or support these [Brown *et al.*, 2005]. Tactons can be created using different kinds of tactile displays including vibrotactile, electrotactile, and pin arrays. Tactons are best used for conveying relatively simple messages to the user. In this study, they are used to convey the notification cues between the device and the user.

Tactons can be combined to form more complex information structures. While words are used to form sentences, tactons can be used to form meaningful tacton compounds which are comparable to textual sentences. In theory, these compounds together with agreed rules could form a complete language. Tactons can be differentiated from each other by manipulating certain parameters such as the frequency of vibration (see Figure 3). In the illustration below, Brewster and Brown [2004] indicated 'Create' command by high frequency, and ´Delete´ command by low frequency. Similarly, this example of tactile language has unique tactons representing 'File' and 'Folder'. As mentioned before, combining these tactons (e.g., 'Create' and 'File') makes it possible to produce more sophisticated tactile feedback.

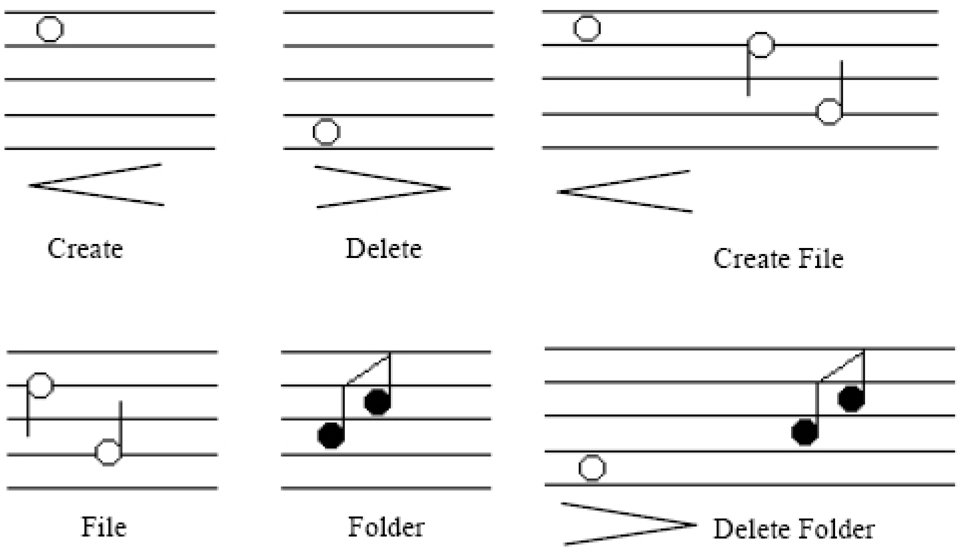

Figure 3 – Compound tactons [Brewster and Brown, 2004].

There are four major parameters in vibrotactile feedback. Brown *et al.* [2005] listed duration, waveform, frequency, and amplitude as the basic parameters for constructing tactons. These parameters are shortly introduced in this section, whereas the optimal values for them are studied in Chapter 5 by prototyping various combinations.

The most important vibrotactile parameter is the duration of vibration. Duration defines how long the haptic actuator vibrates continuously. Combining vibrations of different durations can be used to form rhythms [Brown *et al.*, 2005]. On one hand duration should be long enough to provide the best odds on detecting the produced stimuli, whereas on the other hand it should not irritate the user. The importance of duration is even more critical when stimuli are detected in motion.

Another adjustable vibrotactile parameter is waveform, i.e. the shape of the wave. There are four common periodic waveform types including sine, square, triangle, and sawtooth. The importance of waveform as a vibrotactile parameter is limited due to low resolution of most actuators.

Frequency measures the number of repeating cycles per unit time. As a term, frequency is often used in the context of sound. However, frequency is an equally important parameter also in vibrotactile feedback. Humans are able to hear sounds in the range of 20-20,000 Hz, whereas the practical detection range is only 10-400 Hz in case of the skin [Cholewiak and Wollowitz, 1992]. However, the exact detection range is rather dependant on the used actuator than the human factors. A study of MFT-actuators by Chang and O'Sullivan [2005] claimed that the haptic effects produced with them can be felt when playing the low frequencies between 100-300 Hz. In my own analysis the range 10-400 Hz mentioned by Cholewiak and Wollowitz appeared to be quite realistic when tested using a voice coil actuator.

The fourth meaningful parameter is the amplitude of the wave. Amplitude is a non-negative scalar measuring the wave's magnitude of repetitive variation in time. In other words, amplitude defines the volume of sound, and similarly, it defines the intensity of vibration. Another important term related to amplitude is *root mean square* (RMS) value which measures the mean volume of audible signals. RMS value is calculated by taking a series of equally spaced samples from an audio file. RMS value can be used as an indicator in audio optimization.

There are also more complex parameters which have been developed to produce more sophisticated vibrotactile patterns. Brown *et al.* [2005] evaluated the desirability of using roughness and rhythm as vibrotactile parameters. Their results showed that both of these parameters reached reasonable recognition rates. They used the same voice coil actuator which was used in the experiments presented in this thesis. However, apart from the actuator, their hardware was most likely able to produce stronger audio signal.

In summary, the different combinations of waveform types, frequencies, and amplitudes generate different kinds of vibration effects. Some unusual frequency and amplitude values may render the whole vibration signal undetectable, whereas there is great latitude in manipulating duration. As a common rule, it is vital to take precautions when manipulating vibrotactile parameters.

4. Haptics in Mobile Applications

Prior research on utilizing vibrotactile feedback in heart rate monitoring is barely existent. In an extensive search I found just one study where vibrations are used to monitor patients' heart rates in a clinical environment. In addition, there are a couple of research approaches to pilot vibrotactile feedback during physical exercise. There are also studies examining the effectiveness of vibrotactile feedback in distractive environments. Moreover, there is a variety of other vibrotactile applications which are more distantly related to this study. They are extensively discussed in this chapter to bring up different points of view. A couple of usability studies regarding HRM devices are also reviewed at the end of this chapter.

4.1. Detecting Vibrations in Motion

Physical stress is commonly believed to affect the ability to sense the surrounding world. Wheeler [2008] conducted experiments to determine how physical exercise affects the ability to detect vibrotactile stimuli. The test setup was quite simple consisting of a treadmill and a computer controlling the stimulators. Amplitude was used as a variable while waveform, frequency and duration were fixed values. The users were asked to press a button when a stimulus is detected. The same test was conducted in stationary, walking, and jogging situations. The activity levels were randomized in order to compensate for the learning effects. In addition to determining the recognition rate, the reaction time was also controlled and recorded. The input/output flow outlined below (Figure 4) describes the test setup in its full extent.

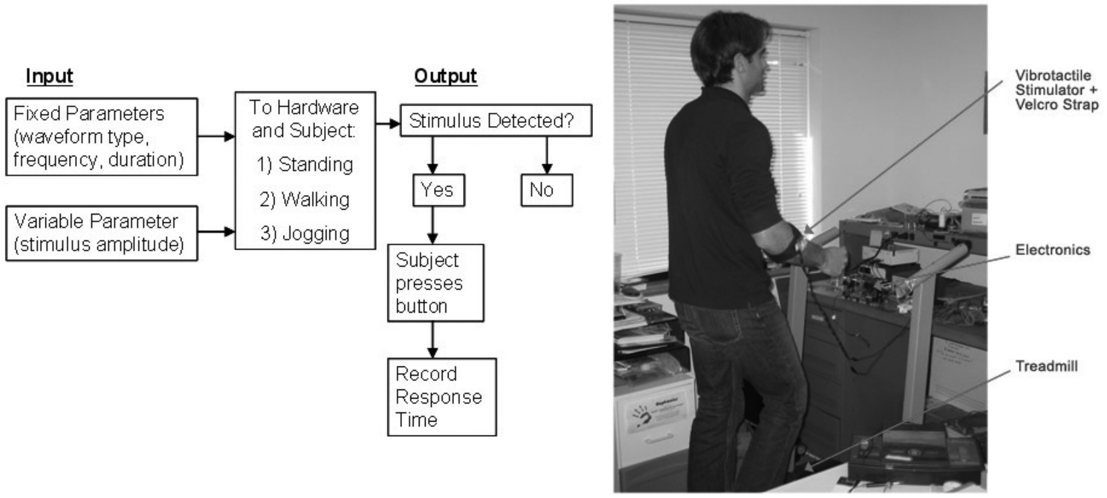

Figure 4 - Mobile haptic display apparatus and experiment design [Wheeler, 2008].

The results of the experiment were almost fully consistent with the hypothesis that stationary subjects have better sensitivity for detecting vibrotactile stimuli. The effect of the activity and stimuli levels is illustrated in Figure 5. From the aspect of this thesis the greatest observation was that the recognition rates can be significantly improved by slightly increasing the intensity of the stimuli.

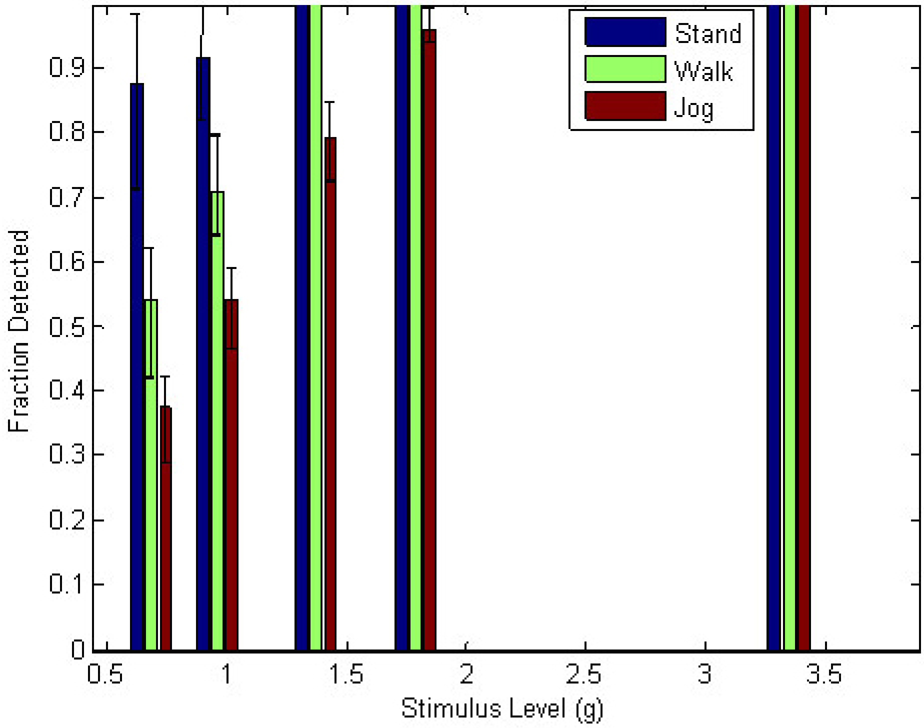

Figure 5 - Fraction of stimuli detected by 8 subjects with a vibrotactile stimulator placed on their arm [Wheeler, 2008].

The results proved that the sensitivity to detect vibrotactile stimuli is proportional to the user activity level. Jogging clearly decreased the perception ability, whereas staying still seemed to be the best situation. It was a very interesting result that the difference was nonexistent when using stronger stimuli. However, higher stimuli level may cause some irritation which sets certain limits on manipulating the amplitude. The optimal stimulus level depends on many factors including the placement of the haptic actuator and the duration of vibration. Certainly there are some natural differences between individuals as well. Regardless, the above mentioned results had a significant influence in designing my own method which is described in Chapter 5.

Post *et al.* [2004] conducted a study to determine if the perception of vibrotactile stimuli is diminished by motor activity. They used three different sites on the operant's arm. Their results showed that motor tasks significantly reduce the ability to detect near-threshold stimuli. Therefore, their findings were in agreement with Wheeler's [2008] results.

In the field of human physiology there are many interesting and more specific studies regarding tactile sensitivity in motion. For instance, a study by Pertovaara *et al*. [1992] showed that the tactile threshold increases especially in the exercising limb during isometric exercise, whereas the perception ability of the resting limbs is not affected. It is an interesting result even though they added that any high-intensity exercise decreases the sensitivity of the skin all over the body, and any prolonged exercise would restore the sensitivity quickly. Their observation could mean - in theory - that actuators are best placed on hands during feet-related exercise, and vice versa. I believe that this could be true especially during extremely hard exercise such as press-ups. However, as already discussed, the sensitivity of the skin varies depending on the area in question resulting that feet probably still have a worse tactile sensitivity in comparison with hands.

An interesting study by Brewster *et al*. [2007] investigated the effectiveness of vibrotactile feedback in handheld devices. They conducted a laboratory study to compare tactile and non-tactile feedback alternatives. Vibrotactile feedback was aimed to spot the possible input errors by providing simple tacton-conveyed cues for the user. The subjects were found to enter more text, make fewer errors, and notice more of the errors they made when vibrotactile feedback was enabled. Another equal experiment was conducted so that the subjects were seated in an underground train. The purpose of these two experiments was to find out whether the positive usability effects were transferable to more realistic scenarios. The underground train can be considered a very distractive environment. According to their results there were very few differences in the total number of errors and the speed of the typing between the laboratory and underground settings. Instead, the subjects noticed more errors in the laboratory setting compared to the underground setting. However, vibrotactile feedback was still beneficial as more errors were made in the visual-only condition. Their study suggested that vibrotactile feedback could improve the usability of handheld devices both in stationary situations and in motion. It has been proved by many other papers that the perception of vibrotactile stimuli is reduced in motion. However, none of the existing studies clearly determines whether that is because of physical exercise or environmental factors such as distraction. Brewster *et al*. proved that perception in motion is reduced partly because of other factors than physical exercise. Even though that result can be considered quite obvious it provides another point of view for analyses regarding the reduced perception capabilities in motion.

4.2. Other Applications

A study by Kaaresoja and Linjama [2005] experimented different durations to find out the optimal length for vibration bursts. They used a Nokia 3310 mobile phone to produce vibrations in three different body locations including the palm, the trouser front pocket, and the belt case. They played the durations of 12.5, 25, 50, 100, 200, and 500 ms five times in each body site making it 30 vibration sequences per site altogether. They used a 5-point Likert scale for the answers which are aggregated in Figure 6 below. There were 18 participants with ages between 19 and 39 years. The results showed that the shortest durations (12.5 and 25 ms) were virtually undetectable which probably resulted from the required spin up time to reach the maximum vibration strength. The technical factors set some limits to the minimum perceivable duration depending on the type of the used actuator. In their conclusion, Kaaresoja and Linjama argued that the duration should be between 50 and 200 ms in this specific case. However, according to the bar chart in Figure 6 it looks like 50 ms provides rather a weak sensation. I would argue that it is better to use such a duration which can be at least moderately detected by the subjects.

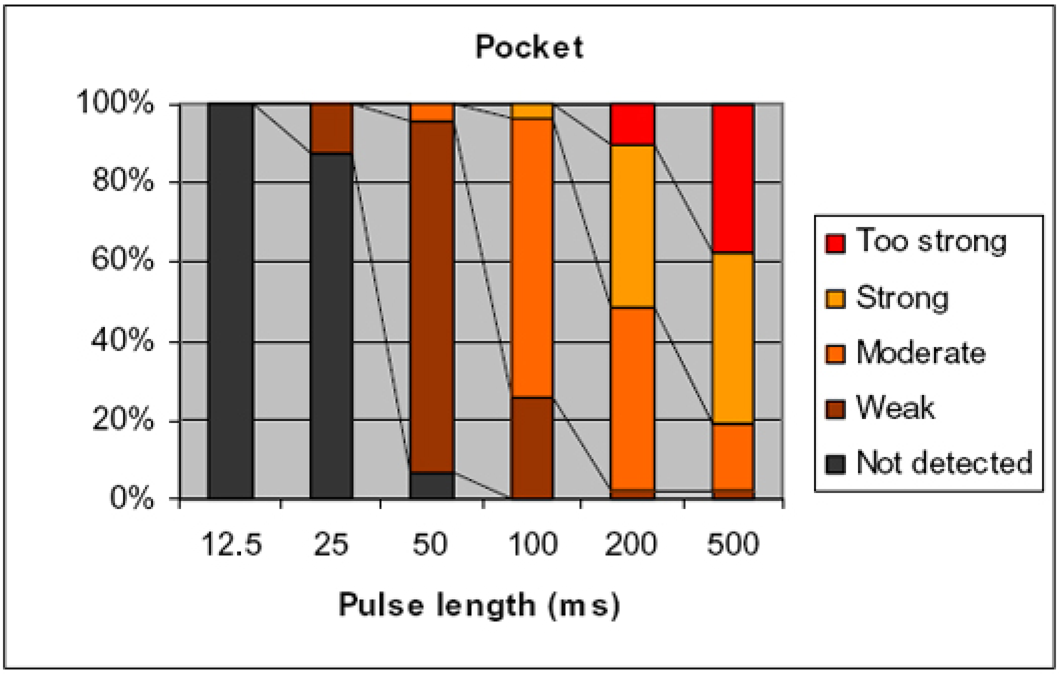

Figure 6 – The effect of duration on the ability to detect vibrations [Kaaresoja and Linjama, 2005].

While Kaaresoja and Linjama [2005] clearly narrowed down the convenient options in finding the optimal duration, these results are quite device-specific. It would be interesting to compare vibrations produced by using Nokia 3310 with vibrations produced by using some newer models such as Nokia N95. The intensity of the stimuli affects directly the obtrusiveness of vibration, and as discussed earlier in this thesis, the different kinds of actuators are capable of producing different kinds of stimuli. Therefore, the results by Kaaresoja and Linjama may not directly transfer to other devices, but they can be seen as suggestive.

Buttussi *et al.* [2006] investigated the use of a mobile guide called MOPET. It was designed to guide the user during outdoor fitness activities. The guide provided navigational, motivational and training support for the user. Buttussi *et al.* argue that MOPET is the first of the kind mobile guide implementation for outdoor fitness activities. They tested the MOPET guide with 12 subjects to evaluate its pros and cons. The results revealed that MOPET was, for example, more useful than traditional fitness trail maps to orientate the subjects during exercise. The percentage of time spent outside the trail was 0,60 % when using MOPET, and 17,80 % when using traditional maps. MOPET was also found to improve the user motivation. While the results proved that mobile guides can be useful for outdoor training purposes, their follow-up research plans are even more interesting from the aspect of this thesis. They are aiming to evaluate users' heart rate and guide the exercise intensity based on heart rate data. That is exactly what was left open for future work by this thesis. In my opinion, their MOPET platform would be an optimal solution when several features are integrated in a single piece of hardware.

A study by Intille *et al.* [2004] explored the algorithms of context-aware applications. They conducted three case studies in order to find the best way to tutor the algorithms to better recognize the various user actions. In the context-aware applications, the functions are triggered based on the data collected by sensors such as accelerometers. Apart from the simple one-to-one sensor-triggered functions, algorithms require extensive training data libraries to detect more complex and user-specific actions. The major goal of their study was to collect user experiences regarding the optimal procedures for the training data collection. There are two major reasons why I found their study interesting from the aspect of this thesis. First, sensor-triggered actions could be very useful in HRMs. Usually athletes are quite busy with their exercise and there is no time for the input. In addition, input methods are not too easy in the current HRM devices. At least the menu navigation can be considered quite challenging during any physical exercise. Context-aware functions could improve the usability by triggering some actions on behalf of the user. Even relatively complex functions could be triggered as some of HRM devices already contain additional sensors such as thermometer, altimeter, or GPS antenna. In addition to the concept of context-aware applications, I found their paper interesting because of one particular case study where real time data was transmitted wireless between a Polar HRM and a PDA. The case study aimed to collect training data which assists the trigger-algorithm in detecting the transitioning between various office activities such as walking, sitting, standing, still, etc. There were 18 subjects wearing a Polar HRM which transmitted the data to a PDA. The subjects were occasionally prompted to self-report their action by making a selection from a short list. In addition, interestingly, the obvious changes in the subject's heart rate also triggered the same PDA prompt. Their goal was to trigger the prompt when it is most likely that some transition has occurred. As one would expect, heart rate is a good indicator to detect activity-related transitions. Their technique could potentially reduce the burden caused by uncalled prompts. The major observations were related to the user comments and complaints. Most of the subjects were frustrated with the random prompts and found the heart rate based prompt-triggering more comfortable. In addition, the users would have preferred to self-report their activities proactively in order to reduce the random prompting. Other notable comments were related to the poor usability of the PDA list selection which was suggested to be replaced by a speech interface. In summary, the most interesting result was that the prompts triggered by heart rate changes reduced user frustration and were found practical. In my opinion, this kind of approach could also improve the usability of HRMs. Even more interestingly, their paper showed that there are existing implementations of wireless real time data transmission between an HRM and a PDA.

Research by Töyssy [2007] experimented telling time by vibration. The success of his method was dependant on both learning abilities and the capacity of human working memory. The main goal of his research was to solve whether people can comprehend time from vibration sequences. He also analyzed the effect of learning on the results. The application was quite a simple J2ME program which produced vibrations using the built-in RMV of the Nokia E50 [2008] smartphone. Vibrotactile feedback was utilized to imitate the digital representation of time in such a way that long vibrations were counted as tens of hours, and short vibrations as single hours. After a short break minutes were represented in the same way. After each recognized time (e.g. 22:14) the subjects were asked to type it using the keypad of Nokia E50. Each tacton contained a single representation of time as illustrated in Figure 7. The long pulses had the duration of 600 ms and the short ones 100 ms. The pause between each pulse was set to 1000 ms, whereas the break between hours and minutes was 2000 ms.

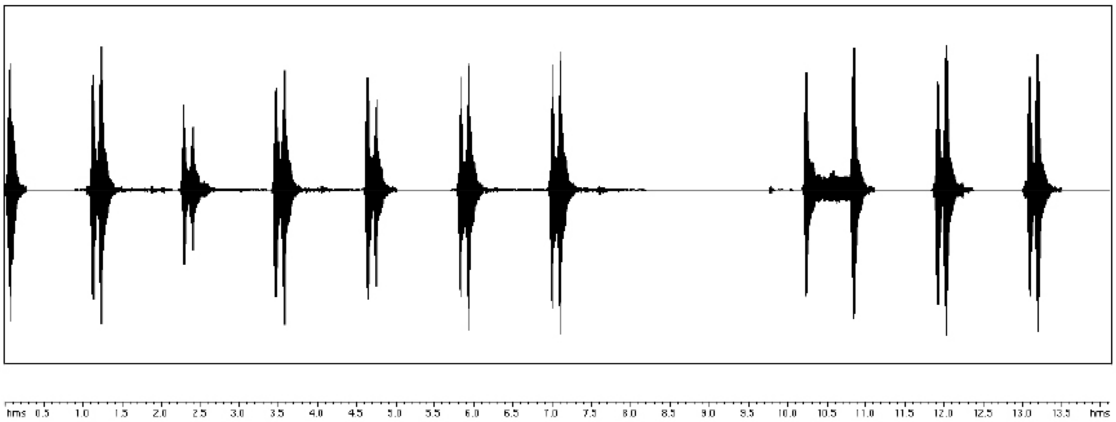

Figure 7 - Pulse sequence recorded with a capacitor microphone from a vibrating mobile phone. The recording represents the time 07:12. [Töyssy, 2007].

His first experiment was meant to prove the method viable. The seven participating subjects were asked to decode 15 different tactons and type the comprehended time accordingly. All the tactons had to be decoded correctly even though it would have required multiple retries. The overall recognition rate was 80 % which determined how many tactons were decoded and typed correctly on the first try. The second experiment studied the effect of learning on the user performance. That experiment had five completely new subjects and five who participated also in the first experiment. There were only 10 different tactons in each segment instead of 15 as in the first experiment. The experiment continued until the subject passed three consequent runs with the recognition rate of 90 % or above. Another terminating condition was the maximum duration of the test which was set to 45 minutes. In this experiment the overall recognition rate reached 88 % which was very close to the initial goal of 90 %. While errors occurred mostly in the minute part (66 %), just 29 % of them were made with the hours, and only 5 % with both. Töyssy believed that it was harder to decode the minutes when the subjects had to keep the hours concurrently in their memory. In addition, minutes could be more difficult to decode because the numbers range between 0-5 for tens of minutes, and 0-9 for single minutes. In case of hours these ranges are just 0-2 and 0-9 respectively. A little surprisingly the effect of learning on the achieved user performance was quite nonexistent. Töyssy concluded that the learning curve was very low and differences in performance emerged rather from individual factors including motivation, concentration, and short term memory. As evidence, two worst performers in the first experiment were also the worst in the second even though they had prior experience on using the application. Töyssy assumed that the recognition rate would improve in case of having the possibility to enter the hours right after recognizing them, and later enter the minutes separately. Of course, it would make the whole procedure of telling time clearly slower. However, in my opinion, the trade-off between cognitive load and quickness should ideally be an individual choice.

Several studies have been conducted aiming to present a variety of physiological data by vibration. Ng and Man [2004] implemented a vibrotactile display which can be used as a silent information display in a medical operation room. Their display was designed to produce vibrations informing doctors about the patient's heart rate status. The current physiological data monitoring systems are based on visual and audible alarms. Visual alarms cannot be effectively observed when some of the attention has to be targeted to the patient at the same time. Similarly, audible alarms cannot be used in an environment which is polluted by other noises. The first experiment by Ng and Man compared vibrotactile alarm and feedback systems against an audio-based system using pseudo-data. The achieved recognition rates of vibrotactile and audible alarms were 97 % and 97,5 % respectively. Both vibrotactile and audible cues were found harder to recognize with distraction. In addition, the recognition process was found to increase the subjects' cognitive load resulting in that they had to stop what they were doing before detecting the alarm. All the subjects preferred vibrotactile alarms to audible alarms. While the results from the alarm test were encouraging, the continuous feedback informing about the actual heart rate was more of a problem. The subjects were able to follow the trend of heart rate changes generally, but they were unable to really recognize the exact heart rate. They also felt the continuous vibration quite annoying. In their second experiment Ng and Man aimed to test vibrotactile alarms and feedback using real-world clinical data. The random subjects were also replaced with an anesthesiologist to improve the reality factor. The results were quite similar to the first experiment. The anesthesiologist subject found vibrotactile alerts more comfortable than audible alerts, and commented that vibrations grab the attention more effectively. The most interesting indicator of success was the average error of 10,44 % in recognizing the actual heart rate. The average error rate shows that it was very difficult to recognize the actual heart rate using the feedback scheme. On the other hand, the general trend was easy to observe. In conclusion, the results proved that vibrotactile alarms can be useful in heart rate monitoring, but the actual heart rate is best presented visually.

4.3. About Usability

There is little prior research on the usability of heart rate monitors. Keinonen [1997] conducted a case study to evaluate the usability of six HRMs to resolve the most affecting factors regarding the product preference. His study focused on analyzing general usability dimensions such as user feeling, ease-of-use, usefulness, presentation, logic, and functionality. The results suggested that the various usability dimensions have only a very little effect on the product preference. A study by Arjanmaa [2006] evaluated more practical usability issues related to HRMs. She found the current feedback modalities impractical especially in noisy usage environments and during sports requiring extensive clothing. Most of the current HRM interfaces provide basic notifications in both visual and audible form while more detailed information is available only visually. Some interfaces let the user select the preferred feedback channel for notifications such as the above/below zone indication. Audible feedback is widely used in most of the commercially available HRMs. Normally HRMs have a certain beep signal which is meant to inform the user of exceeding or going below the intended heart rate zones. According to Arjanmaa this feature is not very functional in the noisy environments. Earplugs could be used to partly solve this problem, but they reduce the ability to hear anything else making, for example, guided fitness training difficult. Another notable disadvantage is that audio signals are omni-directional. Therefore, beeping may annoy other people around if the user is not wearing earplugs. Arjanmaa found this problematic especially in training groups where people are in close proximity to each other. Further disadvantages were found when it comes to the visual feedback of HRMs. Typically, the screens start to blink when a certain heart rate has been exceeded. Blinking can be considered more discreet than beeping, but also more difficult to notice during active exercise. Arjanmaa mentioned skiing as an example of such a sport. Similarly, many other outdoor sports and especially winter sports require extensive clothing which makes visual feedback awkward. Visual screens also suffer from frosting or fogging up under harsh weather conditions. Despite its drawbacks, visual feedback is still usually preferred over audible alerts. According to Arjanmaa’s usability results, most of the users preferred the blinking effect for the notification purposes and disliked all the audible signals. She reckoned that vibrations might be a better way to notify the users. Unfortunately, the HRMs tested in her experiment did not provide vibrations as an option. Arjanmaa argued that it would be meaningful to study the applicability of vibrotactile feedback in the future because the subjects were not completely happy with any of the existing options. She also mentioned that HRMs are usually worn in such a way that there is always a direct contact between the monitor and the skin. Therefore, a haptic actuator could be theoretically easy to integrate with the HRM device without additional wiring. According to Arjanmaa there were no devices with vibrotactile features available at the time of research. However, my repeated search showed that at least Oregon Scientific [2008] manufactures HRMs which are able to produce simple vibrotactile feedback. Their devices have user selectable audible or vibrating heart rate alert for the above/below zone indication. As far as I can judge, they can be used to alert for the changes in heart rate zones in quite the same way as in the experiment presented in this thesis. The only difference is that the HRM by Oregon Scientific alerts for the above/below heart rate zones, whereas this research pilots using four distinguishable tactons to correspond the different activity levels. The difference between these two approaches is quite small, though. However, there are no existing devices which utilize vibrations to guide all-around interval training. Neither is there such a device that would be able to tell the exact heart rate by vibration.

Hansson and Ljungstrand [2000] built a reminder bracelet, a tool to convey event notifications between mobile devices and their users. The aim of the bracelet was to attract the user's attention without disturbing other people around. The subjects wore the bracelet on the wrist and it was connected to a PDA. Light emitting diodes (LEDs) were used to produce the visual cues in order to notify to user of the scheduled PDA events. Their study showed that most of the subjects felt more comfortable using visual cues instead of audible alerts. However, the experiment also showed that blinking can be difficult to notice at times. Hansson and Ljungstrand found the trade-off between the functionality and disturbance challenging. Their study proved that both visual and audible notifications are troublesome at several occasions. It would be interesting to conduct a similar study by using also vibrations for attracting the user's attention. I believe that vibrations would be a more reliable notification method than visual cues and certainly more discreet than audible beeps.

5. Method

A method to encode heart rate monitor notifications into vibration sequences is presented in this chapter. The goal of the method is to study whether vibrations can be efficiently used for feedback purposes in heart rate monitoring. The method includes an actual implementation which is tested by conducting a full-scale user experiment. This research aims to reveal the major pros and cons related to the desirability of vibrotactile feedback in HRMs, as well as in other equivalent portable devices. There are three initially set research problems which are best presented in form of questions. First, can people detect vibrations effectively in motion? Second, how does activity level affect recognition performance? Third, could vibrotactile feedback improve the usability of heart rate monitors? The user experiment regarding my method extensively answers these questions.

The implementation is based on a smartphone playing an audio file. The audio playback is transduced to vibrations by a connected tactor. As discussed earlier, tactors are a specified kind of haptic actuators. The audio file contains encoded notifications whose purpose is to guide the user to catch up certain activity levels during physical training. In other words, the smartphone and the tactor act as a fictional heart rate monitor in this setup and produce notifications for the user. An actual heart rate monitor is also used alongside to record the user data for post-analyses regarding the performance. When the user is requested to catch up, for example, a certain running level, the heart rate would assumedly increase. The behavior of the heart rate can be viewed afterwards by analyzing the HRM data. Of course, it would be preferable to produce the vibrations with an actual HRM, but unfortunately there were no suitable devices available at the time of research. On the other hand, if there had been such a device, I believe this research would have been already conducted by someone else. Anyway, this research relies on using alternative devices as a combination to reproduce the idealistic setup.

In addition to describing the research method, the grounds for paramount decisions are explained in this chapter. As discussed earlier, vibrations can be tricky to detect and distinguish from each other. In search of an optimal setup many different tactor locations and parameter values were tested. Three documented prototypes were implemented during the development period. Their purpose was to straighten out the problems and adjust the values for the key parameters. Most problems were actually only hypothetical by their nature. The prototypes were meant to rule out the possible problems and provide knowledge for designing the final version. The final implementation is based on the results from three major prototypes. Each prototype is extensively described in this chapter. A full description of the final implementation is also presented at the end of this chapter.

5.1. Apparatus

The hardware included a smartphone, a heart rate monitor, and a certain kind of tactor for feedback purposes. These devices are introduced in this section and their co-operation is also shortly discussed.

The smartphone was chosen to be Nokia N95 [2008]. This particular model was selected because it was one of the most advanced smartphones available at the moment of writing. It was used to control the vibrations and provide output for the actual feedback device.

Choosing a haptic actuator was a more difficult task. It required lots of efforts to get familiar with the available options. Because of very limited time I was forced to decide quite early on. I ended up choosing a C2 type vibrotactile transducer by Engineering Acoustics [2008]. It is relatively small in size (3.05 cm in diameter by 0.79 cm high) and weights only 17 g. The skin contactor is 0.76 cm in diameter, whereas the actual raised area is 0.064 cm. The C2 vibrotactile transducer – or shortly tactor – was largely chosen based on the recommendations from the advisors of this thesis.

The connectivity between Nokia N95 and the C2 tactor is excellent. Nokia N95 contains 3.5mm stereo headphone plug. Therefore, the C2 tactor can be connected with its pre-mounted cable. Some older phones without 3.5mm plug could also be connected up by using a suitable adapter. However, their headphone amplifiers may be incapable of producing strong enough output signal for vibration purposes. Even when playing the audio files using the nominal center frequency of the C2 tactor, the vibration effect was not strong by any means. In my early trial runs, Nokia N95 was briefly replaced with a USB sound card [Audiotrak Maya, 2008]. The results showed that the generated vibration effect was significantly stronger. Thereby, the audio amplifier of Nokia N95 can be seen as the bottleneck in producing stronger vibration. According to my quick technical reviews it still seemed to be one of the best options among the present mobile devices. The upcoming smartphones are hopefully able to produce stronger audible output.

Choosing the heart rate monitor was a relatively easy task as the decision did not directly affect the results of this research. According to my original plans the heart rate data was meant to be analyzed in real time. The plan was to process the ECG information in Nokia N95 and produce the vibrations based on the resulting data. However, the technology required for wireless communication and real time processing turned out to be troublesome with Nokia N95. It would have been time consuming to solve the problems in the scope of this thesis. This was a real backlash for this research given that such a wireless setup has been previously employed at least by Intille *et al.* [2004]. As discussed earlier in this thesis, their study involved wireless real time data transmission (HRM-PDA). Buttussi et al. [2006] also described similar setup in which the wireless range was reported to be 9 meters. Instead of changing the already chosen hardware to overcome these problems I decided to keep the focus more precisely on analyses of user performance. Implementing a more sophisticated application was left for follow-up research. Of course, a PC based setup would have been easier to implement, but I preferred an outdoor experiment over a lab setup to maintain a reasonable reality factor. Due to the above mentioned reasons I decided to use a heart rate monitor separately to view the resulting changes in heart rates. The chosen model was Polar Electro 810i [2008] which is one of the most modern HRM devices available. The wireless communication between the transmitter belt and the wrist unit made it well wearable. There would have been other quite equal options, but I cannot see that any other existing option would be better than the chosen device for this particular purpose. Ideally, vibrotactile feedback would have been produced by the HRM itself without any external actuators. However, no existing device is capable of that. Even though the earlier discussed HRM by Oregon Scientific was found capable of producing basic vibrotactile feedback, the features are almost certainly tightly tied to the manufacturer provided software. Because of focusing on the primary research goals, the HRM was not used for controlling the vibrations in any way. Instead, the data provided by the HRM was used to analyze whether the test subjects followed the vibrotactile instructions given by Nokia N95 and the C2 tactor.

5.2. Controlling the Vibrations

The vibration control ended up being relatively simple. In the beginning I planned to write a J2ME program to control the vibrations. Later in the technical analysis it turned out that it is more straightforward to use a regular audio file in which the notification cues are encoded in form of tactons. This audio file was created beforehand and Nokia N95 was only used for playing it. This approach can be considered the easiest as the experiments were designed so that there was no real need for any dynamic vibration sequences. However, it would be easy to change over to a J2ME or a native Symbian solution later on if necessary. Either one of these technologies could be used to manipulate the audio output of the device. They could be used even without external tactor by enabling the integrated RMV of the smartphone. According to Töyssy [2008] out of those two approaches the native Symbian provides finer control over the RMVs of smartphones. Töyssy also mentioned that J2ME implementations may not work on all J2ME phones as vibration controls and their extensions vary. While J2ME and Symbian would enable dynamic feedback, they will likely suffer from upcoming deprecations as this kind of APIs are developing rapidly. However, in my opinion dynamic feedback would have not increased the scientific value of this research. Therefore, I decided to use just a regular audio file to keep the setup simple.

Quite early on I decided to use an open source cross-platform sound editor called Audacity [2008] to create the customized audio files. The file format was chosen to be mp3 because it tends to be the most popular audio format in the mobile environment. In order to find an optimal bit rate for the experiment, a few tests were conducted by using the 32, 128, 192, and 320 kbit/s bit rates. According to the results there was no big difference regarding how many kilobits the file used per second to encode audio. The same audio file encoded and exported using any of the above bit rates was found well perceptible with a C2 tactor. The vibration sequences produced using the different bit rates seemed to be all the same. I argue that it is not even possible to distinguish the vibrations by hand. However, the visual spectrum analyzer of Winamp media player [2008] produced slightly different graphic representation in case of each tested bit rate. Otherwise, the differences were barely existent except for the file sizes. After all I ended up choosing the 192 kbit/s which is a pretty common bit rate of encoded music. The additional kilobits were not an issue as Nokia N95 has enough space for the file in any case. It is worth noting that besides the bit rate also the encoding algorithm affects the quality of mp3 files. In this experiment all files were encoded using an encoder called *lame_enc.dll* [Lame, 2008] which was recommended by the Audacity editor.

In order to perceive any worthwhile vibrotactile feedback, there must be an agreement on the meaning of vibrations. In this experiment the information was encoded in tactons. Each tacton contained information in the form of temporal vibration sequences. As discussed earlier in this thesis, tactons are the tactile counterpart of visual icons. Each tacton has a unique meaning which has been agreed on beforehand. As the purpose of this research is to guide the users during interval training, there was no need for too many different tactons. In addition, the tactons used in this research are quite different from each other and therefore extensive learning should not be required to distinguish them. The expected bottleneck of this experiment was rather related to the sensitivity to detect vibrations during exercise. Another limitation was related to the ability of the used hardware to produce sufficient stimuli. The key requirement for worthwhile vibrotactile feedback is that the users are able to detect the stimuli, decode the tactons and react accordingly.

5.3. Prototype I

The first prototype was designed based on the varied dummy runs conducted during the hardware integration. In the first prototype, there were only two different vibration sequences indicating the below and above zone alerts during interval training. The major goal of the first prototype was to find out whether it was possible to perceive any stimuli in motion using the chosen hardware setup. The tacton representing the below zone alert contained a single continuous 10 second vibration. The above zone tacton contained one second vibrations followed by pauses of one second. The rotation between them was continued until a total of 9 seconds was reached. In total, there were five separate pulses in each above zone alert sequence. The pause between the separate tactons was approximately 50 seconds, whereas the total duration of the audio file was 20 minutes. Thus, there were 20 separate tactons altogether. The prototype was tested by two female test subjects (ages 21 and 25) resulting in that each error hit meant 2.5% in the overall results. The first prototype had fixed values for waveform, frequency, and amplitude. The changing variables were the location of the tactor and the type of activity. The test environment was an outdoor setting with natural background noise including some passing traffic. While Nokia N95 was placed in the pocket, the C2 tactor was attached by turns to the front wrist, the palm, and the back. A simple cloth strap was used to keep the tactor stable on the back. Similarly, tight woolen gloves were worn to keep the tactor balanced on the wrist and the palm. Extensive clothing was put on due to cold winter conditions. Therefore, both devices were completely out of sight. Naturally, the output volume of Nokia N95 was set to maximum to get the best odds on detecting the stimuli. A comprehensive description of the different test combinations is presented in Table 3. The first two columns indicate the changing variables and the last column reveals the results of each combination.

Location	**Activity**	**Duration**	**Waveform**	**Frequency**	**Amplitude**	**Performance**
front wrist	stationary	1s rot. / 10 s	square	250 Hz	1 (max)	**100%**
front wrist	walking	1s rot. / 10 s	square	250 Hz	1 (max)	**100%**
front wrist	jogging	1s rot. / 10 s	square	250 Hz	1 (max)	**97,5%**
palm	stationary	1s rot. / 10 s	square	250 Hz	1 (max)	**100%**
palm	walking	1s rot. / 10 s	square	250 Hz	1 (max)	**100%**
palm	jogging	1s rot. / 10 s	square	250 Hz	1 (max)	**100%**
back	stationary	1s rot. / 10 s	square	250 Hz	1 (max)	~**50%** guessing
back	walking	1s rot. / 10 s	square	250 Hz	1 (max)	**0%** no sensation
back	jogging	1s rot. / 10 s	square	250 Hz	1 (max)	**0%** no sensation

Table 3 – Results from Prototype I.

The surprising difference between the indoor and outdoor environments was one of the first findings. Despite the fact that the C2 tactor is relatively noiseless, it gives some audible cues in the indoor setting. Any indoor experiment should be ideally conducted using hearing protectors to block all audible signals. Outdoors it was notably more difficult to notice when the tactor starts vibrating especially if the tactor was placed on the back. In addition to the missing audio cues, another affecting factor was the constantly changing contact between the actuator and the skin.

As assumed, the level of activity also affected the ability to detect the stimuli, but the difference between the stationary and the jogging case was less of a problem than Wheeler [2008] discovered in his research. The effect is not viewable in the performance rate, but the subjects reported verbally that detecting the vibrations was slightly more difficult in motion. Prior to implementing this prototype I assumed that Wheeler's error rates would be lower than mine as his user experiment was conducted in a lab setting without external disturbance. Probably this result is caused by the fact that the used hardware was too different to make a valid conclusion. Any two different experiments would be comparable only if stimuli intensities were measured somehow.

This prototype proved that detecting stimuli all over the body is by no means self-evident even when playing the 10-second continuous vibration using the nominal center frequency (250 Hz). At least the differences between the palm and the back were huge. It appeared to be almost impossible to perceive any stimuli when the tactor was placed on the back. Some vibration was sensed in the back in the stationary position, but it was still quite hit-or-miss to recognize whether it was continuous or rotational vibration. The back was initially tested also indoors prior to this prototype and I found it slightly easier to recognize the vibrations. However, the unintended audible cues may have affected perceptiveness. While the back was out of the question as a tactor location for further experiments, both the palm and the wrist fulfilled the requirements. Both locations achieved excellent recognition rates, but the sensation was still reported stronger in the palm by the subjects. Despite its poorer perceptual capability, I decided to choose the wrist because the palm is often needed for other purposes. In addition, most HRMs are normally placed on the wrist making it a more natural choice. In case a haptic actuator and an HRM should be physically integrated some day, it should anyway be placed either in the HRM wrist unit or in the transmitter belt.

Based on the results regarding user performance, using only two different tactons appeared to be too simple as 100% success rates were recorded in the experiment. Apart from using the back as a tactor location there was only a single error hit which resulted more or less from poor concentration during the long-winded 60 minutes (3 x 20 minutes). In conclusion, both the palm and the front wrist were found excellent locations to perceive vibrotactile stimuli. The trade-off between the optimal sensitivity and the minimal disturbance can be considered a case-specific issue. The only design flaw was probably the length of the audio file. It was very time consuming as well as physically demanding to conduct this kind of user experiment. However, this prototype made it possible to alter the plans and design more versatile tactons for the final version. On the other hand, it is a remarkable result in practice that two tactons can be distinguished with a 100% success rate. However, as mentioned before, the first prototype was just meant to prove that the basic goals are within reach.

5.4. Prototype II

The second prototype was implemented to find an optimal duration and waveform for the vibration. In advance to this prototype it was already quite clear that one of the tactons is going to be a single continuous vibration in the final version. Therefore, the main goal of this prototype was to find the best distinguishable options in contrast to the continuous vibration. Several different vibration sequences were tested.

This prototype sampled the durations of 100, 200, 400, 600 and 800 milliseconds. These durations were chosen based on the results from the first prototype. The subjects of the first prototype believed that they would be able to detect shorter signals as effectively as they detected the one second vibrations. Therefore, the duration of one second was left out of this prototype as it was found too long without a question. Kaaresoja and Linjama [2005] experimented different durations using a mobile phone RMV and found even 500 milliseconds too long. According to their results the ideal duration ranges from 50 to 200 milliseconds. They also argued that the subjects found longer durations irritating. However, in the first prototype, even the 10-seconds continuous vibration was quite pleasant or at least acceptable. I would assume that the hardware used by Kaaresoja and Linjama produced stronger stimuli which naturally irritate sooner. However, I decided to try slightly longer vibrations compared to the 50-200 milliseconds range suggested by Kaaresoja and Linjama because of environmental factors such as outdoor setting and concurrent motion. This decision was made taking also into account Wheeler's [2008] results which proved that the higher activity level remarkably reduces the ability to perceive vibrotactile stimuli. Actually I could have kept the initial duration of one second as it was not found irritating, but both subjects claimed that they would be able to detect shorter vibrations as well. My hypothesis was that the optimal value for the duration parameter lies in range from 100 to 800 milliseconds in this particular scenario.

The selected five durations were tested so that a pause of one second was followed after each vibration. The rotation between the vibrations and the pauses was continued until a total of 1-4 vibrations were played. The number of vibrations was randomized in order to find out whether it is more difficult to detect fewer vibrations in a row. After each sequence the subject was asked the number of vibrations. The answers were written down by the supervisor of the experiment. The total length of the audio file was again 20 minutes and the tactons were separated by 15-19 seconds breaks. In total there were 20 rotational tactons for each parameter combination. In contrast to the first prototype, the breaks were shorter to save some time. Therefore, each 20-minute file had capacity for three different parameter combinations. Even then, the experiment consisted of 100 minutes of jogging per subject. There were two subjects: one male (24 years) and one female (23 years).

In addition to the various durations, three major waveform types were tested to find out their differences. The subjects were not asked to report the waveform type they detected, but the performance rate was calculated separately for each waveform. The order of waveform occurrences was randomized. The tested combinations of varying duration and waveform are illustrated in Table 4. In this prototype the activity level was fixed to be jogging in all combinations. Similarly, the location, the frequency and the amplitude were fixed values in this prototype.

Location	**Activity**	**Duration**	**Waveform**	**Frequency**	**Amplitude**	**Performance**
front wrist	jogging	100 ms rot.	square	250 Hz	1 (max)	**77,5%**
front wrist	jogging	100 ms rot.	sine	250 Hz	1 (max)	**72,5%**
front wrist	jogging	100 ms rot.	saw tooth	250 Hz	1 (max)	**65,0%**
front wrist	jogging	200 ms rot.	square	250 Hz	1 (max)	**82,5%**
front wrist	jogging	200 ms rot.	sine	250 Hz	1 (max)	**95%**
front wrist	jogging	200 ms rot.	saw tooth	250 Hz	1 (max)	**87,5%**
front wrist	jogging	400 ms rot.	square	250 Hz	1 (max)	**90%**
front wrist	jogging	400 ms rot.	sine	250 Hz	1 (max)	**92,5%**
front wrist	jogging	400 ms rot.	saw tooth	250 Hz	1 (max)	**90%**
front wrist	jogging	600 ms rot.	square	250 Hz	1 (max)	**100%**
front wrist	jogging	600 ms rot.	sine	250 Hz	1 (max)	**95%**
front wrist	jogging	600 ms rot.	saw tooth	250 Hz	1 (max)	**95%**
front wrist	jogging	800 ms rot.	square	250 Hz	1 (max)	**95%**
front wrist	jogging	800 ms rot.	sine	250 Hz	1 (max)	**100%**
front wrist	jogging	800 ms rot.	saw tooth	250 Hz	1 (max)	**97,5%**

Table 4 – Results from Prototype II.

The results of the second prototype were quite straightforward. Duration had a clear effect on the recognition rate, whereas waveform was less influential. The results also emphasized the initial hypotheses very well. While 578 out of 600 (2 subjects x 20 tactons x 15 combinations) tactons were correctly reported by the subjects at the right time, some additional errors occurred in the answers regarding the number of vibrations. As mentioned earlier, the subjects were asked how many vibrations there were in each tacton. There were only a very few errors with 600 and 800 milliseconds. Also, the duration of 400 ms was satisfactory with only 11 errors, whereas a total of 14 error hits were recorded with 200 milliseconds. The worst performing duration was 100 milliseconds with 34 errors in 120 tactons. There was probably some guessing involved in the results regarding the duration of 100 milliseconds. Especially the shortest duration caused problems in calculating the number of consecutive vibrations.

In conclusion, when the duration parameter was set to 400 milliseconds or higher the recognition performance was excellent as in the first prototype.

All tested waveform types produced sufficient stimuli when the duration was long enough. It was really difficult to identify the waveform as all types produced quite equal stimuli. The shortest duration of 100 milliseconds provided no chance at all to identify which waveform was used. The duration of 200 milliseconds made waveform distinguishing also extremely challenging. According to Brown *et al.* [2005], the differences in waveform are lost in case of low bandwidths. Therefore, it was not surprising that the waveform types were difficult to distinguish in this prototype. Brown *et al.* also claimed that the subtle differences in waveform are not perceivable by the skin. In this prototype, the differences were found virtually nonexistent when examined in motion. In the stationary position it was possible to distinguish the saw tooth waveform from the other two types. The saw tooth produced slightly weaker stimuli compared to the sine and square waves. Apparently because of that the recognition rate was the worst in case of saw tooth waves.

The results clearly proved that the duration affected the ability to calculate the number of vibrations in a tacton. In addition, the majority of non-reported tactons contained only a single vibration, which raises the question of whether tactons should always contain at least two separate vibrations. The detection performance was 96,3 %, whereas the overall performance ended up being 90,5 %. The results were excellent taking into account that most of the errors occurred with certain durations. The results regarding the final version may further improve as only the best possible duration is used.

5.5. Prototype III

Original plans included an implementation of a third full-scale prototype to test the different frequencies and amplitudes. However, the third prototype was simplified as the results were almost certainly known beforehand. This prototype was rather meant to support the possible future work. It was found out during the previous prototypes that all combinations appeared to produce weaker stimuli than the 250 Hz frequency and the max amplitude setting. Similar findings were reported by Brown *et al.* [2005]. They were using exactly the same kind of tactor in their experiment and claimed that C2 tactors have a reduced output at any other frequency than 250 Hz. They also argued that reducing the amplitude could potentially degrade the perception of other parameters, or even render the signal undetectable. The 250 Hz frequency was also recommended by the tactor manufacturer for producing the strongest possible stimuli [Engineering Acoustics, 2008].

Regardless of the above mentioned facts I decided to test different frequencies including 1, 10, 50, 100, 150, 200, 250, 300, 350, 400, 450, 500, and 1000 Hz in order to clarify the actual perceivable frequency range. Duration was fixed to 2 seconds, waveform was chosen to be square, and the amplitude was set to maximum. All combinations were tested in the stationary position without any hearing protection. As there was no multi-user test plan regarding this prototype the perceptiveness could not be expressed using percentages. Instead, the perceptiveness was rated using a five-point Likert-scale. There were five grades: very good, good, acceptable, poor, and very poor.

The results indicated that any frequency between 100 and 400 Hz would produce sufficient stimuli for feedback purposes (see Table 5). Against my initial hypothesis all the frequencies between roughly 150 and 350 Hz produced quite reasonable stimuli. As expected, the strongest one was 250 Hz, but the differences in comparison with other frequencies were not found that radical as described by Brown *et al*. [2005]. Even the 50 Hz frequency was still detectable without any problems in the stationary position. The frequencies lower than 50 Hz required quite a bit of concentration to be detected. Similarly, anything higher than 500 Hz was virtually just audible. The difference between 400 and 500 Hz frequencies was found surprisingly remarkable.

Location	**Activity**	**Duration**	**Waveform**	**Frequency**	**Amplitude**	**Perceptiveness**
front wrist	stationary	2 seconds	square	1 Hz	1 (max)	**Very Poor**
front wrist	stationary	2 seconds	square	10 Hz	1 (max)	**Very Poor**
front wrist	stationary	2 seconds	square	50 Hz	1 (max)	**Poor**
front wrist	stationary	2 seconds	square	100 Hz	1 (max)	**Acceptable**
front wrist	stationary	2 seconds	square	150 Hz	1 (max)	**Good**
front wrist	stationary	2 seconds	square	200 Hz	1 (max)	**Good**
front wrist	stationary	2 seconds	square	250 Hz	1 (max)	**Very good**
front wrist	stationary	2 seconds	square	300 Hz	1 (max)	**Very good**
front wrist	stationary	2 seconds	square	350 Hz	1 (max)	**Good**
front wrist	stationary	2 seconds	square	400 Hz	1 (max)	**Acceptable**
front wrist	stationary	2 seconds	square	450 Hz	1 (max)	**Poor**
front wrist	stationary	2 seconds	square	500 Hz	1 (max)	**Very Poor**
front wrist	stationary	2 seconds	square	1000 Hz	1 (max)	**Very Poor**

Table 5 – Vibration perceptiveness using different frequencies in Prototype III.

The results indicated that any frequency between 100 and 400 Hz is able to produce perceivable stimuli. However, it must be taken into account that these values were tested in the stationary position, whereas any other activity level would probably narrow down the sufficient Hz-scale. In addition, the duration parameter was set to two seconds which is longer than normally. In practical usage shorter durations would definitely be needed especially when designing more complex tactile patterns. In some cases, increasing the activity level and shortening the duration could possibly render the signal completely undetectable. In conclusion, 250 Hz was quite a self-evident choice for the value of the frequency parameter. Any manipulation of frequency would probably make the detection more difficult for the user.

5.6. Final Implementation

As an afterthought, all prototypes were worthwhile to implement and helped in designing the final version. They cleared out early on what is possible and what is not. Actually they set more limits than I assumed beforehand, but they also proved that the planned goals of this research are realistic and achievable. The outcome of the prototype phase was a set of optimal parameter values for the tactons. However, there were so many affecting variables that all combinations could not be tested thoroughly.

The location of the tactor was chosen to be the front wrist. Results from the prototypes indicated that the sensitivity of the wrist was good enough even though it was worse than in the palm. In my opinion the tactor is easier to wear on the front wrist than on the palm.

The duration of a single vibration was set to 600 milliseconds. In case of the first prototype both subjects complained that the pulse duration of one second was too long. On the other hand, it was revealed by the second prototype that the performance started decreasing when reducing duration below 400 milliseconds. In case of 400 milliseconds, the recognition rate was slightly over 90 %, but very interestingly, there were two tactons with only a single vibration which were not detected at all by the subjects. It comes to a conclusion that single vibrations are not necessarily detectable in motion even if the duration is relatively long. Because of that I designed the tactons in such a manner that there were always 2-4 vibrations in each sequence. Just in case, I also decided to choose the duration of 600 milliseconds instead of 400 milliseconds to reduce the risk of having imperceptible tactons.

The waveform types were tested in the second prototype. The differences were so marginal that there was no wrong option to be chosen. Both sine and square waves produced almost identical stimuli which were very difficult to discriminate even in the stationary position. The saw tooth waves produced slightly weaker stimuli even though the difference was virtually meaningless. Finally, I decided to use the sine waves because it was recommended by the tactor manufacturer for producing the strongest stimuli [Engineering Acoustics, 2008].

Different frequencies were sampled during the whole development phase. The third prototype further examined the possible values and provided a more detailed documentation of the results. The results indicated that the frequencies between 100 Hz and 400 Hz would be suitable for producing strong enough stimuli. However, as the stimuli were rather weak anyway, there was no reason to use any other value than the nominal center frequency (250 Hz) of the C2 tactor.

Amplitude was the least examined variable as the maximum value clearly produced the strongest stimuli. All of the briefly tested alternative values appeared to weaken the stimuli. Theoretically, alternative values would be useful in designing more sophisticated tactile patterns on the condition of having more advanced hardware. However, it was justified to keep using the maximum amplitude.

The above discussed parameters were used to design the four well distinguishable tactons. Each tacton contained a request to catch up a certain activity level (stationary, walking, jogging, or running). Tactons were differentiated by manipulating the duration of the stimuli and the number of consecutive vibrations, whereas other parameters were not very meaningful having this particular hardware in use. The following design was considered to be the best possible setup for this particular scenario: a tacton containing two consecutive short vibrations suggested walking, three short vibrations suggested jogging, whereas four suggested running. The fourth kind – continuous 5 second vibration – suggested stopping any ongoing activity and keeping the standing position until the next tacton is recognized.

As a quick recap, the tacton parameter values for the final user experiment were chosen as follows:

- Duration: five seconds continuous for stationary, 600 milliseconds rotational for others,
- Waveform: sine,
- Frequency: 250 Hz, and
- Amplitude: maximum (1/1).

These parameter values provided simple vibration effects without any multi-dimensional mappings. This design was considered the most effective for this application and any further manipulation would have most likely just hindered the detection. The temporal presentation of each tacton is shown in Figure 8. The Y-axis indicates the elapsed time, while the X-axis indicates the tacton in question.

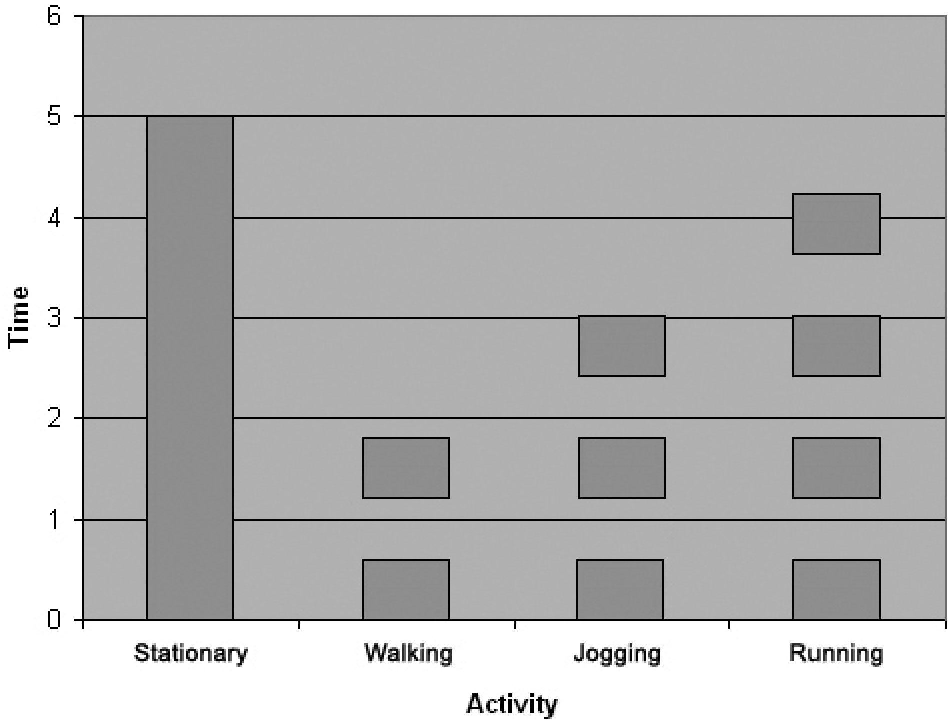

Figure 8 – Tacton design as in the final version.

The total duration of the audio file was chosen to be 20 minutes. The file consisted of 21 tactons which were surrounded by silence. Technically, the audio file was optimized to best correspond the recommendations by device manufacturer. The recommended drive was as follows: sine wave tone bursts 250 Hz at 0.25 A RMS nominal, and 0.5 A RMS max for short durations [Engineering Acoustics, 2008]. The Audacity [2008] editor was found incapable to effectively adjust RMS values. Therefore, another music editor called Cubase SX3 [2008] was used for fine-tuning the audio file. The resulting RMS value was quite close to the recommended 0.5 A. The volume was normalized to maximum peak without crossing 0 db. The original db value produced by Audacity editor was approximately -3 db which was already quite good. The 0 dB level is the highest level that sound can hit in the digital world without peaking the meters and causing massive distortion in the sound.

The file contained an encoded interval training plan which was illustrated in form of vibrotactile cues. The plan involved brief intervals at near-maximum exertion interspersed with periods of lower-intensity activity. The near-maximum bouts were assumed to be well viewable in the data provided by the HRM. The purpose of the lower-intensity activity was to provide enough recovery and let the heart rate settle. The lower-intensity intervals were expected to be recognizable by lower BPM rates in the heart rate curves. The intervals were designed in such a manner that the HRM data would include as little fluctuation as possible. In other words, there were no sharp changes in the activity level during the experiment. The entire interval plan is illustrated in Figure 9. There were two major near-maximum intervals which were surrounded by lower-intensity activity. This particular interval plan is not probably the most effective from the aspect of fitness training. Instead, it was planned to be simple in order to channel more attention to detecting and recognizing the tactons.

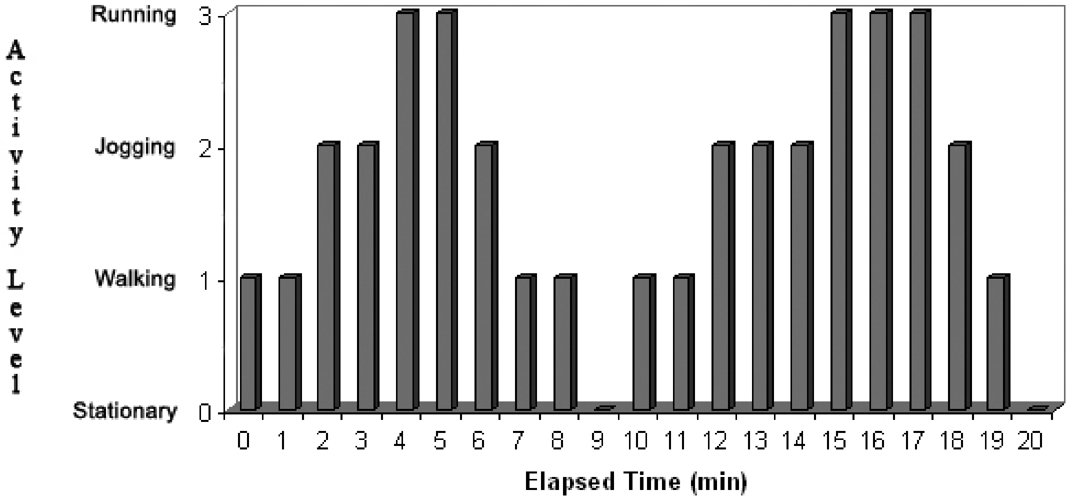

Figure 9 – Requested activity levels in the final audio file implementation.

As illustrated in Figure 11, each tacton contained a suggestion to the user to either increase, decrease, or keep the current intensity of training. After nine minutes there was one tacton suggesting stopping the activity for a period of one minute. The break was put there partly because of testing the associated tacton, and partly because of providing some rest in order to let the heart rate settle. Same kind of tacton was also placed at the end of the experiment.

In the real world the user might be able to estimate the next tacton-suggestion during the interval training. The audio file created for this experiment was meant to be unpredictable, but also understandable. The subjects knew the total duration of the interval plan, whereas all other details about different bouts were kept secret. During normal interval training athletes usually set their plans themselves and know roughly how they progress. Therefore, at least in theory, the error rate should be higher in this experiment compared to a real world scenario. When having no good guesses about the type of the next tacton, the subjects really need to rely on their recognition abilities. Töyssy [2007] also argued that the possibility of estimation would narrow down the options when decoding vibration sequences.

6. User Experiment

There are four sections in this chapter. The arrangements regarding the conducted user experiment are described in the first section. The goals of the user experiment are discussed in the second section, whereas the hypotheses are presented in the third section. Last, the actual results are revealed in the fourth section.

6.1. Arrangements

There were quite a few arrangements regarding the user experiment. A detailed description of the setup is presented in this section. First, the basic information regarding the test subjects is presented. Second, the test environment including location, weather, and clothing issues is discussed. Third, the software settings are shortly mentioned and the hardware setup is analyzed in greater detail. Last, the preparation procedures such as user briefing are further discussed.

There were 12 participating subjects including five females and seven males. Their ages ranged from 21 to 28 years. Even though the interval training was not very grueling some basic precautions were still taken. The audio file was designed in such a way that anybody relatively fit should be able to participate in the user experiment. All subjects had a normal sense of touch by their own report. Two of the subjects had prior experience of HRMs, but they had never used any interval programs. One of them had also conducted his own research on tactile feedback. However, that piece of research was not related to heart rate monitoring in any way and it was not expected to give any advantage in the experiment. Other subjects had no notable experience of vibrotactile feedback apart from typical mobile phone RMVs and gaming devices.

The user experiments took place on a straight stretch of a jogging path. The path was located approximately at an elevation of 90 meters. There were no noteworthy differences in altitude on the path. All subjects could not attend the test during the same day and therefore the weather conditions were not exactly equal for all of them. All experiments were conducted during the winter, but the path was not slippery at all. The temperature ranged between -5 and 5 °C, whereas the wind was not notable on any of the testing days. Due to the cold temperature all subjects wore suitable tracksuits so that all the hardware was totally out of sight.

An HRM was used for recording the heart rate during the experiment. As mentioned earlier, the used model was Polar Electro 810i [2008]. The heart rate was recorded by a transmitter belt which communicates wireless with the wrist unit. The default settings were used except for the recording frequency which was set to 15 seconds. This frequency was chosen because the heart rate curve would have become somewhat incoherent in case of recording all heartbeats. As far as I can judge, the 15 seconds frequency actually denotes the calculated average of all occurred heartbeats during that period.

The tactor was connected to the front wrist of the right hand. As in the prototype experiments, tight woolen gloves were worn to keep the tactor balanced on the wrist. Nokia N95 was put into the pocket of the tracksuit, whereas the connecting wire was placed inside the sleeve. The output volume of Nokia N95 was set to maximum and other audio playback settings were kept as default.

In summary of the used hardware, there were four separate items in the setup. The transmitter belt and the wrist unit were used for recording the heart rate, whereas Nokia N95 and the C2 Tactor were used for producing vibrations. The hardware setup is illustrated in its full extent in Figure 10. As shown in the main image, the tactor was placed inside a glove in order to keep it well balanced. The smaller image expands upon the actual tactor placement on the front wrist. Even though the tight woolen gloves were worn, the tactor could potentially move slightly during exercise. Initially I considered using some kind of sticky tape for keeping the tactor stable on the wrist. However, it might have been overkill as the devices are not taped up to the skin during normal training either. There are also some natural differences regarding hand shapes making it very difficult to place the tactor on exactly the same location for all subjects. In theory, this kind of differences may affect the ability to detect stimuli. Double-checking the tactor location after each test was a good way to make sure the tactor has not moved excessively. That way, possible tactor movement could be taken into account in the analyses regarding the results.

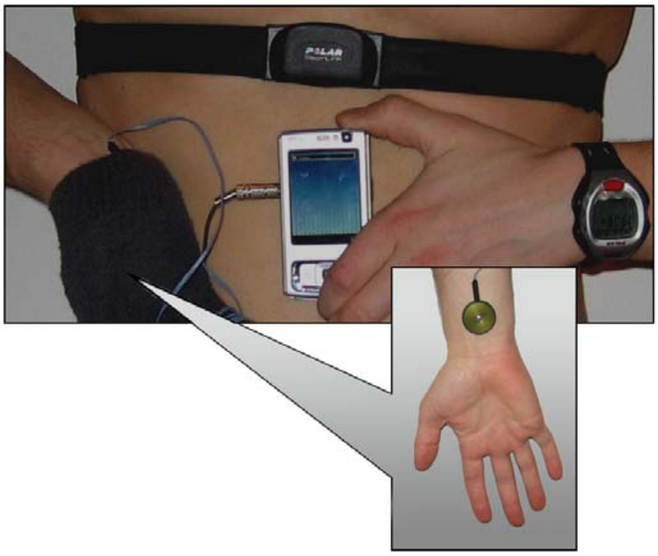

Figure 10 – Overview of hardware setup

A technical cross-check list was used to ensure that the settings would be equal for all subjects. There were quite a few settings regarding the HRM and Nokia N95 so the list made it easier to remember them all. The same cross-check was conducted both before starting each test and after finishing each test. For example, the list reminded of setting the volume level to maximum which was a key requirement regarding this user experiment. I was also slightly afraid that the buttons of Nokia N95 could be accidentally pressed during the experiment because of body movement. In my opinion the control buttons of Nokia N95 are slightly over-responsive at times. Therefore, the list reminded of locking the keypad to deactivate the audio controls. Another important task was to set the phone to offline mode in order to prevent it from receiving incoming calls during the experiment. When it comes to the HRM there were settings such as recording frequency and exercise set which were double-checked.

All the subjects were given only a very short briefing about the test. They were shown a written description of the experiment. In short, they were informed about the purpose of the experiment, about the duration of the experiment, and about the possibility to decide their walking, jogging, and running speeds themselves. In addition, the basic information needed for distinguishing the tactons was presented in the written instructions. The subjects were also allowed to familiarize themselves with the four distinguishable tactons in practice. An illustrative audio file was played preparatory to the actual experiment. It contained a single occurrence of each tacton type. None of the subjects requested to repeat the illustrative file. Neither had any of them attended any of the prototype experiments so they all had a quite equal starting point. Even the person with the research background was not familiar with the particular tactor type used in the experiment. All the subjects were asked to report real time when they detected any vibration during the experiment. In addition, they were asked to report verbally the number of consecutive vibrations in each tacton. This procedure was planned out in order to make the performance rates as reliable as possible.

6.2. Goals

There were three major questions this thesis was meant to answer. The first question: can people detect vibrations effectively in motion? The second question: how does activity level affect recognition performance? The third question: could vibrotactile feedback improve the usability of heart rate monitors? These questions pretty much defined the three major goals for this user experiment.

The question about vibrotactile feedback in motion is a key issue in this thesis. As mentioned earlier, interval training is such an exercise that consists of repetitions of high intensity activity followed by periods of low activity. Since there are several levels of motion, it can be seen as an optimal technique for testing whether vibrations can be effectively detected and recognized in motion. In my opinion, this could be proven by achieving sufficient performance levels and by getting heart rate curves which fluctuate in relation to the tacton-conveyed requests. When it comes to user performance, the minimum goal was set to 90 %. The error quota of 10 % should include both detection and recognition errors. Since the results of the prototypes were quite encouraging, this goal can be considered more than realistic. Achieving 90 % performance appears to be quite a common goal in this kind of experiments. For instance, the same threshold was used by Töyssy [2007] in his research, which was discussed earlier.

The second question was designed to evaluate the effects of exertion and activity level on the recognition performance. These effects can be optimally studied during interval training as the activity level can be increased step by step. It is worth noting that the second question cannot be examined at all if the answer regarding the first question is unfavorable. Assuming that the second question can be examined, the answer is probably directly readable from the user performance results. In case the performance is worse regarding the tactons conveyed during the higher activity bouts, it would quite obviously indicate that the activity level affects the performance. Of course, the differences should be remarkable to make this conclusion on the grounds of just a single user experiment. The resulting findings can be compared with the results by Wheeler [2008]. As mentioned earlier, in my opinion Wheeler's results can only be seen as suggestive because of the differences in the context.

When it comes to the third question some existing studies can be used to support the findings. There is prior research on the usability of traditional modalities in HRM devices. These modalities have been found problematic in many ways. Arjanmaa [2006] found it difficult to pay attention to visual feedback during intensive training sessions. According to the same study, audible feedback was also found irritating by many subjects. My research was not aimed to compare vibrotactile feedback with these other modalities, but the purpose was to examine whether vibrations are utilizable for conveying the HRM notifications. In case of receiving positive results from this experiment, it would be a logical next step to conduct further experiments where visual, audible, and vibrotactile notification cues are compared with each other. This comparison could be meaningful even though audible and visual cues have already been proven troublesome for HRM notification purposes. However, the third major goal of this experiment was to solve whether vibrotactile feedback could improve the usability of heart rate monitors. On the grounds of the prior research only the weaknesses of the current modalities are known so far.

In addition to the three major goals, this user experiment also provided a good opportunity to gather up general ideas and suggestions. All the test subjects were asked to say a few words about the concept of having vibrotactile feedback in HRM devices. Even though the raw numbers such as the performance rates receive most of the attention in this kind of studies, one of the equally important research goals was to get various user views and opinions. All subjects were asked for their opinion whether vibrotactile feedback has some downsides in this particular purpose of use. I believe that the general feedback can be very useful when it comes to possible future research.

6.3. Hypotheses

This section gives some initial hypotheses regarding the three major questions this thesis sets out to answer. In addition to foreseeing those answers, this section presents some general hypotheses as well.

Can people detect vibrations effectively in motion? In my opinion, it is quite expected that the subjects are able to detect and recognize the vibrotactile cues without major problems. It can be considered probable when taking into account the excellent recognition performance rates achieved in the prototype versions. Likewise, the calculation of the consecutive vibrations can be hypothesized to go smoothly. Even though some error hits were recorded in the second prototype the recognition performance was still quite good. As mentioned already in the previous section, the performance goal was set to 90% which is realistic in my opinion.

How does activity level affect recognition performance? According to the results acquired from the prototypes, the differences should be insignificant regarding the various activity levels. However, Wheeler [2008] made opposite findings when it comes to different activity levels. The recognition performance was found better at lower activity levels, whereas more errors occurred at higher levels. He also found the higher stimuli intensity to shorten the gap between the various activity levels. All his results were quite rational, but for some reason the same behavior was not repeated in my own prototype results even though the used stimuli intensity was quite low. It raises some questions: are there affecting differences between the indoor and outdoor environments (treadmill vs. jogging path)? Could Wheeler's results be repeated in this experiment by further reducing the stimulus level? The latter question will probably never be answered in this thesis as it is not wise to reduce the stimuli level. Instead, in my opinion it is a significant finding if the activity level does not affect too much the ability to detect stimuli. However, any further conclusion cannot be made based on prototypes which were tested only by two subjects. Therefore, my hypothesis for the final user experiment is a trade-off between my own findings and Wheeler's results. I believe that activity level has a very slight effect on user performance.

Could vibrotactile feedback improve the usability of heart rate monitors? This is probably the most difficult question to answer without any practical experience. My hypothesis is that usability could be improved at least in scenarios where the existing modalities suffer from unfavorable conditions. As a simple example, the usage of audible cues can be considered problematic in noisy environments as well as in socially discreet situations.

Apart from the major questions, there are also general hypotheses regarding the experiment. The resulting heart rate curves can be expected to be quite individualistic as the subjects are of different ages, shapes, and genders. As mentioned earlier in this thesis, the normal heart rate ranges between 60 and 100 BPM for a healthy adult. The heart rates of the various test subjects should increase and decrease most likely in relation to each other. However, each subject is asked to define themselves what is a suitable running speed for them. These differences may cause the heart rate curves being slightly skewed and not fully comparable to each other. A treadmill experiment would probably result in better comparable heart rate curves. Even then, the further analysis of the heart rate behavior would be non-professional as this is not a medicine study. In the real world differing running speeds can be considered natural and the interval training should preferably be planned individually anyway. One of the most desired findings would be that the heart rate of each subject starts instantly increasing or decreasing after recognizing a tacton which suggests changing the ongoing activity level. There is probably some natural delay which varies depending on the physical condition and the level of exertion. However, the HRM data can be assumed to clearly show the difference between the near-maximum bouts and the lower-intensity activity. A hypothesis regarding the user comments is that some subjects may complain about the intensity of the stimuli. As mentioned already several times, the intensity of the stimuli cannot be further increased with the current hardware. Despite the fact that there are no ways to solve this issue it is important to be well aware of it.

6.4. Results

There were two major ways to analyze the results. First, the heart rate curves provided by the HRM were analyzed in detail. Second, the answers provided by the subjects were examined. There are three major terms which are used to report the results. *Detection performance* indicates how many tactons were detected out of their total number. *Recognition performance* indicates how many tactons out of the detected were also correctly decoded. *Overall performance* is calculated by dividing the number of correctly decoded tactons with the total number of tactons.

The detection performance was 97,1 % which is better than in the prototype experiments. This was quite hypothesized because the parameter values were optimized for the final experiment. As many as five subjects detected all the tactons without missing any. Interestingly, many users missed the second tacton requesting to continue walking. I believe that this was partly caused by the users' initial confusion regarding the test. It was also interesting that half of the missed tactons represented walking. This result is in total disagreement with Wheeler's [2008] findings which suggest that the detection should be easier during walking. On the grounds of that it could be argued that two consecutive vibrations are more difficult to detect than three or four. This may seem almost self-evident, but in the prototype versions only tactons containing just a single vibration suffered from poorer detection performance, whereas those with two or more were detected with quite an equal success. Taking all the above discussed views into account the initial confusion could also be a potential reason for poorer performance in detecting tactons representing walking.

The recognition performance was 98% which is even better than the detection performance. It seems that it was not very difficult to calculate the number of vibrations in a tacton even though this caused some problems in the prototype versions. On the grounds of the improved recognition performance, the applied alterations in implementation can be seen successful. There was one especially interesting recognition error though. One of the subjects reported the long vibration as five short ones even though there was no such an option at all. After discussing with the subject I came to the conclusion that this recognition error was caused by hand movement which occurs roughly in one second cycles during walking. Natural hand movement could have caused the feeling of rotational vibrations by loosening the contact between the tactor and the skin in cycles.

The overall performance was 95,2 % meaning that only less than 5 % of the tactons were missed or incorrectly recognized. While detection and recognition performance only indicate the results related to a certain process, the overall performance measures the real applicability of the method. In my opinion, overall performance can be seen as the most important measurement of usability in this research. Having an overall performance rate of 95,2 % proves that the research method is viable. All the performance rates are presented in detail in Table 6.

While the user performance was excellent there were also some unfortunate backlashes. The test results regarding two subjects could not be included because of either flawed setup or technical problems. The first discarded test proceeded without problems until the audio playback was accidentally stopped. It turned out that the keypad of Nokia N95 was left unlocked by accident. That was totally my own fault as the locking was mentioned also on the used cross-check list. For some reason this task was forgotten when doing other preparations. It is quite obvious that body movement caused some pressure on the pocket and the button got pressed accidentally. As discussed earlier, there were some slight hypothetical worries regarding the over-responsive keys of Nokia N95. If the keypad is unlocked, the audio control buttons can be pressed by accident as they are located in the main navigation wheel of the phone. However, the test had to be discarded because the erroneous setup was not caused by the test subject. In hindsight, it would have been better to place the smartphone in a belt case instead of the regular tracksuit pocket. The second discarded test was more of a mystery. The tactor just finished producing any vibrations after the fifth tacton in the experiment. The test was continued for another five minutes before coming suspicious of whether there was something wrong with the hardware. First, it was double-checked that Nokia N95 keeps playing the audio file. Second, the wire connection was checked to ensure the devices are properly connected to each other. Third, the tactor was placed on a fingertip to test whether it works correctly. It turned out that the tactor produced no vibration at all for some reason. The audio file was certainly fine as it produced sound when the tactor wire was disconnected. The test had to be terminated due to these technical problems. Surprisingly, the tactor strength was normalized by degrees when tried again back home. The fault could have been caused by excessive humidity or cold weather conditions. Still, it was very strange and the breakdown could not be repeated again when trying the tactor in extremely bad weather by design. It seems that this fault remains a big mystery. However, I decided to leave both of the above discussed tests out of the performance calculations as they were clearly technical problems. Counting them in would have caused skewed performance results as the problems were not related in anyway to the ability to detect or recognize tactons. In summary, it seems that the hardware setup is quite vulnerable to different kinds of problems.

Subject	**Detection Performance**	**Recognition Performance**	**Overall Performance**
#1	20/21	20/20	20/21
#2	20/21	19/20	19/21
#3	**N/A** (flawed setup)	**N/A** (flawed setup)	**N/A** (flawed setup)
#4	21/21	21/21	21/21
#5	21/21	21/21	21/21
#6	19/21	19/19	19/21
#7	20/21	19/20	19/21
#8	21/21	19/21	19/21
#9	**N/A** (technical problems)	**N/A** (technical problems)	**N/A** (technical problems)
#10	21/21	21/21	21/21
#11	21/21	21/21	21/21
#12	20/21	20/20	20/21
All	204/210 (97,1 %)	200/204 (98 %)	200/210 (95,2 %)

Tacton	**Detection Performance**	**Recognition Performance**	**Overall Performance**
stationary	20/20 (100 %)	19/20 (95 %)	19/20 (95 %)
walking	67/70 (95,7 %)	66/67 (98,5 %)	66/70 (94,3 %)
jogging	69/70 (98,6 %)	68/69 (98,6 %)	68/70 (97,1 %)
running	48/50 (96 %)	47/48 (97,9 %)	47/50 (94 %)

Table 6 – User performance rates regarding the user experiment.

All the activity levels were well viewable in the heart rate curves produced by the Polar HRM. Actually the behavior was more punctual than assumed beforehand. As illustrated in Figure 11, the changes in heart rate were visible virtually without delay. It was quite surprising because some delays were noticed when analyzing the prototype results. However, every single heartbeat was recorded in the prototype versions, whereas the average value of 15 seconds was used in the final version. The sparser frequency probably hid the delays in the graphical presentation. Each time after recognizing a new tacton the heart rate started changing almost immediately. Of course, there were individual differences and some subjects had more fluctuating curves than others, but the differences were quite irrelevant.

Figure 11 - The heart rate behavior during the interval training session.

When changed to higher activity level the heart rate curve ascended steeply during the first 30-60 seconds. Thereafter the ascent was gentler for the next minute. It seems that the heart rate became relatively even at a certain BPM level after about two minutes of continuous activity. There were no longer bouts than three minutes of the same continuous activity so it is hard to estimate how long the heart rate would have remained at the same level.

When changed to easier activity, it was highly dependant on the individual factors how long it took for the heart rate to settle down. The heart rate descent was faster with seemingly fit subjects, whereas it took even double the time with a couple of unfit and overweight subjects. Of course, the selfdecided jogging and running speeds also affected the required recovery time. Some subjects run faster than others and naturally needed more time for recovery.

The heart rate curves clearly revealed the recognition errors, which fulfilled one of the major goals regarding the research method. Generally all the occurred errors were found in the graphical presentations, but some cases seemed to be more obvious than others. In most cases the errors were more visible when changing between the walking and jogging levels in comparison with the change between jogging and running. Some subjects had only a very small difference between their jogging and running speeds which caused that the heart rates were quite equal during both activity levels. Therefore, some recognition errors caused the heart rate curve to diverge from its hypothetical path even less than 10 BPM. Despite having a few hard-to-detect errors the majority were actually clearer than assumed beforehand. A good example of spotting a recognition error from the graphical heart rate curve is illustrated in Figure 12. In this particular case the curve behaved extraordinary evenly for some reason.

Tacton Log:
13:00 Jogging
14:00 Jogging
15:00 Running (not perceived)
16:00 Running
17:00 Running
18:00 Jogging
19:00 Walking
20:00 Stationary

Figure 12 – An example of spotting a recognition error from a graphical heart rate curve.

The general feedback regarding the experiment was also positive and some subjects came up with creative ideas. The most common remark was that the produced stimuli should be stronger to make the application more workable.

This complaint was also among the initial hypotheses regarding the usability of the final implementation. There were only two subjects who would have rather kept the current intensity of the stimuli. Some subjects seemingly concentrated more than others on the detection process. I believe that having a stronger stimuli level could have reduced the need for concentration.

The subjects also gave some suggestions and comments regarding the setup. A couple of subjects would have preferred having the tactor placed on the palm instead of the front wrist. As discussed earlier, the perceived sensation would be stronger in the palm, but it could be more irritating as well. In addition, the front wrist is a more natural choice when considering that this research could form the basis for integrating a suitable tactor with an HRM device in the future.

Another proposal for improving the test setup was to provide a possibility to re-examine the latest tacton. This could be conducted for example by pressing a particular button in the Nokia N95 keypad. As the tactons were encoded in an audio file, it would not be a major problem to traverse backwards in the audio stream. The control interfaces for this purpose are already existent in most mp3 players and smartphones. To take the idea of re-examination further, one subject suggested enabling multimodal output in such a manner that the user could confirm an unclear vibrotactile sensation by looking at a visual HRM screen.

When it comes to hardware there was an interesting suggestion to replace Nokia N95 with an mp3 player such as Apple iPod [2008]. There are probably some differences in amplifiers regarding the various portable devices. At least in theory, the stimuli could be stronger when playing the audio file with some other device than Nokia N95. Another hardware-related idea was to connect an external headphone amplifier to Nokia N95 to strengthen the stimuli. These kinds of solutions are commonly used in, for example, sound recording. Some portable headphone amplifiers were found in a quick search, but their functionality was not tested in practice. The products by Boostaroo [2008] promise an impressive 100% - 400% increase in audio output depending on the model in question. Their flagship model is powered by two small AAAA-batteries providing up to 40 hours of battery life. When it comes to scientific purposes, this kind of devices could potentially be useful in strengthening vibrotactile stimuli. Ideally, the future HRM devices could have this technology built-in alongside an appropriate haptic actuator.

7. Discussion and Future Work

The results were positive across the board: the detection performance was 97,1 %, the recognition performance was 98 %, and the overall performance 95,2 %. Even though the overall performance goal was very carefully set to 90 %, the achieved results surpassed the expectations. The best performers detected and recognized all the tactons without any errors. Promisingly, even the worst performers exceeded the goal of 90 % overall performance.

It appeared to be surprisingly easy to detect vibrotactile stimuli in motion. The performance rates alone answered the initially set question whether vibrotactile feedback can be effectively perceived in motion. The overall performance rate was 94 % when calculating only the tactons conveyed during the running bouts. The difference to the other activity levels was virtually nonexistent. When analyzing the results from the aspect of the first research problem they seemed to be well in agreement with the existing studies [Post *et al.*, 1994; Wheeler, 2008]. In my opinion the most important remark was that the detection results appear to be almost entirely proportional to the intensity of the stimuli.

This study also showed that the recognition of vibrotactile stimuli is – even though surprisingly little - diminished by physical exercise. One might ask why surprisingly? Many subjects complained that the stimuli were rather too weak. Because of that, in my opinion, the stimuli level can be considered to be near-threshold using the same definition as Post *et al.* [1994] used to describe stimuli which are just perceivable. Therefore, the achieved results are not in agreement with other studies [Post *et al.*, 1994; Wheeler, 2008] which argued that the near-threshold perception ability is significantly decreased by motion. Neither their studies nor my hypotheses support the achieved results. Running-tactons achieved the overall performance rate of 94 %, jogging-tactons 97,1 %, walking-tactons 94,3 %, and stationary ones 95 %. Even though running-tactons were the worst performers, the difference was too small to correspond with the results presented in other studies. I believe that the detection performance was affected by the differences in tactons. Walking-tactons were possibly more difficult to detect because they contained only two separate 600 ms vibrations. In contrast, running-tactons contained four vibrations which may have advanced the perception. However, some further experiments would be needed to make any valid conclusion.

The heart rate curve behavior was quite surprising in my opinion. According to the hypothesis the curve path should have ascended quite gently in the beginning and steeply after some minutes of continuous activity.

However, the behavior seemed to be almost opposite as it was illustrated by the graphical presentations (Figures 11 and 12) in the previous chapter. The heart rate appeared to become quite even after a couple of minutes of continuous activity. This behavior was slightly surprising as the textbook of medical physiology by Guyton and Hall [2005] argued that the heart rate should increase slowly at the beginning of the exercise. Instead, the stroke volume should increase more rapidly until reaching its maximum by the time the cardiac output has increased only halfway. After that any further increase in cardiac output should occur by increasing the heart rate. In my opinion, the curve behavior is somewhat unexpected when taking into account the high reliability of the textbook of medical physiology by Guyton and Hall. It might be that the interval bouts are too short in this experiment to get a better idea of the heart rate behavior. In theory, another possible reason for the strange behavior - even though unlikely - might be the inaccuracy of the used heart rate monitor. Many of the subjects were relatively fit so I would rule out the possibility that their stroke volumes would have reached the maximum capacity so quickly after the start.

Even though vibrations would be theoretically an optimal way to convey information, there are many open questions to be solved. The major headache is that several subjects would have preferred having stronger vibrations. I am afraid it is not possible using this particular hardware combination of Nokia N95 and the C2 tactor. Another remarkable limitation is the low displacement produced by the C2 tactor. Therefore, it is not possible to encode more complex information patterns to the tactons. According to Poupyrev *et al.* [2002] this problem is associated with all voice coil type actuators. Their findings were confirmed by my prototype results which proved that it was virtually waste of time to manipulate other parameters than duration. To address this problem, other types of haptic actuators could be tried out to produce stronger stimuli, but as discussed before, basically all current devices have at least some downsides. When having a look at the used C2 tactor I found three major weaknesses in addition to the limited capability to produce stimuli. At least the price of the tactor should be a lot lower before considering any commercial usage (currently over $200). Also, the power consumption should be brought further down as the HRM wrist unit uses a watch-type battery which is changeable only at authorized service centres. According to Polar Electro [2008] the estimated battery life of an HRM is 1-2 years depending on the model and the frequency of use. They also reminded that the battery will run out more rapidly if the backlight and the beep signal are used excessively. Most likely the C2 tactor would consume too much power to keep the battery alive for even just a few months. The third major problem is related to the portability. The tactor should be smaller before it could be elegantly integrated with an HRM wrist unit. In summary, the above discussed notes are well comparable with the four actuator feasibility requirements listed by Poupyrev *et al.* [2002]. I argue that the C2 tactor fails to meet the criteria at least regarding the size and customization potential. The size is not far from being reasonable, but it is still quite far from being suitable for a manufactured commercial model in which the haptics is not the main thing. Of course, the requirements listed by Poupyrev *et al.* [2002] can be interpreted in many ways. My opinion is based on the fact that the tactor should be integrated with an actual HRM wrist unit such as Polar Electro 810i. As mentioned several times, the hardware is usually the major headache in implementing tactile interfaces, and so it was in this research as well. Another hardware related drawback – from the scientific aspect – is that the results reported by this thesis are specific to the used hardware combination and may not fully transfer to other devices. These particular devices were chosen since they are portable, easily available, and fully compatible without additional adapters.

There were also a few design flaws which caused some headache. The subjects were not informed whether they should guess the answers if they were unsure of them. In hindsight, the written instructions should have contained some guideline about this issue as well. In a couple of cases the subject ended up guessing the answer. In my opinion, it would have been better to inform the subjects to keep the current activity level in case of being unsure of the tacton type. However, the effect of guessing was quite insignificant when it comes to the performance rates. Another design flaw was related to the hypothesis that the perception ability is affected by the activity level. In case of missing several tactons in a row during the running bout, should the running be continued for the whole test? This would have been a problem because some of the subjects were not so fit that they could have kept the full running speed for the whole test. This design flaw fortunately never came true in practice, but there was no plan of action regarding this issue. In addition to the already discussed design flaws there were some debatable decisions regarding the hardware. It would have been meaningful to test the output capability of a few different amplifiers before deciding on Nokia N95. On the other hand, the scope of this thesis set some limitations on the number of prototypes. It would be interesting to conduct a similar experiment in the future using Apple iPod [2008] or an equivalent mp3 player as suggested by one of the test subjects. Nokia N95 has definitely better programmability, but iPods are smaller in size and more purpose-built for playing audio files. There are definitely differences in audio output between the different devices. However, the European regulations limit the maximum allowed amplifier output of mp3-players to protect against hearing defects. This limitation is the reason why mp3 players purchased in the member countries of European Union cannot be optimally used together with most of the voice coil actuators. Even if this was possible, mp3 players would be useful mostly in prototyping. In my opinion, one of the future goals of the possible follow-up research would be to integrate a tactor physically with an HRM wrist unit, or even further integrate an HRM with a smartphone.

The third research problem was more abstract than the others. Therefore, the obtained results cannot directly answer the initial question of whether vibrotactile feedback could improve the usability of HRMs. On the grounds of the user feedback and other findings I argue that vibrotactile feedback could be used effectively for guiding the interval training at least. However, the findings are not necessarily valid if producing feedback relating to other HRM features. The major question is the overall desirability of utilizing vibrotactile feedback in HRMs. As discussed earlier, the research by Arjanmaa [2006] argued that the potential of vibrotactile feedback in HRMs should be extensively studied. However, Arjanmaa concentrated on the whole user interface, whereas this research examined only a very specific HRM program. For instance, a more complex design would be needed in order to convey actual heart rate by means of vibrotactile feedback. So far at least Ng and Man [2004] have experimented conveying actual heart rate by vibration. They argued that the general heart rate trend is conveyable, but the exact heart rate is difficult to recognize by the user. However, in this kind of studies the performance is highly dependant on the data presentation method. In addition, equivalent tasks requiring complex vibration patterns suffer from learning requirements which was not an issue in this experiment. I believe that vibrotactile feedback could potentially handle various tasks in HRMs and equivalent devices, but the limits will be faced when the amount of conveyable information is increased too much. It was also pointed out by Töyssy [2007] that the human short term memory sets some limits on the possibilities. However, in my research that limitation was bypassed altogether as the vibration sequences were very simple in the final user experiment. Human factors are not any problem in my opinion as long as the information is kept simple enough. Even though the subjects' comments regarding the usability were encouraging and the method was proven viable there is a long way to go before vibrotactile feedback would be ubiquitous in heart rate monitoring.

The users' motivation to vibrotactile interaction depends probably mostly on the signal quality. When asked whether the subjects liked vibrotactile feedback, most answers were a conditional yes. All subjects except for two complained about the intensity of the signal. Otherwise there were no major complaints. Most subjects preferred the vibrotactile alternative over the visual channel. On the contrary, there were quite conflicting thoughts when comparing vibrotactile cues with audible beeps. A couple of subjects strongly preferred having audible beeps even though they were not even tested during the experiment. It is probably a more natural way of interaction if it is socially accepted. As a downside, in an outdoor setting the background noise may distract hearing the audible beeps which was also mentioned by Arjanmaa [2006]. Similarly, audible feedback would be less preferred in social situations. On the other hand, Töyssy [2007] believed that distractive environments might be problematic also for vibrotactile applications as the human memory capacity for vibration sequences is small. This issue is not critical in case of this research, but it can be considered a remarkable limitation when designing more complex applications.

As an afterthought, the initially set research problems were probably too heavily dependant on the subjects' ability to perceive stimuli. The recognition performance was fully dependant on the detection performance. Similarly, the third research question about the usability improvements was fully dependant on the detection ability. If the first question had been answered unfavorably the remaining questions would have been completely meaningless. The reason behind selecting these research problems was the lack of existing studies in this particular application area. While the research problems regarding the possible follow-up studies could be more specific, this thesis offered just one view into the world of vibrotactile interaction.

How about the future? When it comes to heart rate monitoring, the paramount future goal would be an actual integrated device which is capable of displaying varieties of tactile feeling. I believe that vibrotactile feedback could improve the quality of feedback not only in HRMs, but also in a wide variety of other portable devices. Even though the possibilities seem clear I would argue that vibrotactile feedback is best adapted very carefully to unexplored purposes of use as its functionality heavily depends on the context. The future research is only limited by the bounds of human imagination. This thesis left many open questions which could be potentially studied in the future. For instance, it would be interesting to implement a more complex application with an extensive set of distinguishable tactons. That would mean exploring the limits of human short term memory and learning curve. Another interesting future study would be to conduct an experiment with visually impaired subjects. I believe that they – if some users – are the right ones to evaluate applications like this. The perspective would be probably completely different because of their everyday experiences with the alternative interaction modalities such as haptics. This kind of experiment would result in very valuable feedback at least. However, in case of visually impaired subjects the experiment would be probably better conducted using a treadmill instead of the outdoor setting.

8. Conclusion

This thesis investigated the effectiveness of vibrotactile cues in guiding interval training. On the grounds of the achieved 95,2 % overall performance rate the method seems to be viable. The three major questions were answered. First question: can people detect vibrations effectively in motion? Yes, they can, on the condition that the intensity of the stimuli is reasonable. The detection performance of 97,1 % left no questions whether vibrotactile stimuli can be detected in motion. Second question: how does activity level affect the recognition performance? It seems that the activity level affects only slightly the recognition performance which ended up being 98 %. Neither perception nor recognition performances were found poorer at higher activity levels. In my opinion, this result is significant as activity level – at least when intuitively thinking – should affect the detection performance in some extent. Third question: could vibrotactile feedback improve the usability of heart rate monitors? Yes, it could, but more research is definitely needed before making valid conclusion. This research examined usability just by analyzing raw performance numbers rather than conducting real usability tests.

While this study answered the initially set questions, it also raised a bunch of new ones for the follow-up research. In addition to them, many other existing questions remain unanswered. How would additional user distraction affect the results? How would the perception ability of the front wrist change if the hand muscles were exercised instead of leg muscles? What would be the highest number of different tactons to be remembered without problems? How would multimodal feedback affect performance rates? Can this method be applied to other applications than interval training programs?

The limitations regarding the hardware left too few adjustable parameters making it impossible to produce the ideal stimuli. The user reception could have been warmer if the vibrations had been clearer. Even though the problems were mostly foreseen, this research widened the perspective and hopefully gave some contribution to the future research. At least it was proven that vibrotactile feedback is not necessarily an inferior notification method compared to vision and sound. Taking into account that the explored application area is very limited the main contribution of this study might be the ice-breaking for the future work. Despite the limits set by the technical and human factors there are still many unexplored and potential purposes of use. While the traditional vibration may seem the most stable of concepts, yet it continues to evolve as the devices develop further. Meanwhile, the competing notification conveyors such as audible alerts will most likely remain quite the same causing the technological gap to shorten. The realms of possibility in tactile interaction are still unknown.

References

[Arjanmaa, 2006] Arjanmaa, T., Sykemittarit liikunnan apuna – hyöty vai haitta? (In Finnish). Department of Computer Sciences, University of Tampere, Master's Thesis, June 2006. Also available as: http://tutkielmat.uta.fi/pdf/gradu01139.pdf. (Checked on April 5th, 2008)

[Apple iPod, 2008] Apple iPod nano, http://www.apple.com/ipodnano/. (Checked on March 15th, 2008)

[Audacity, 2008] The Audacity cross-platform sound editor, http://audacity.sourceforge.net/. (Checked on February 16th, 2008)

[Audiotrak Maya, 2008] Audiotrak Maya EX5 CE External 5.1 USB Surround Audio Solution (Sound Card), http://www.digit-life.com/articles3/multimedia/audiotrak-maya-ex5.html. (Checked on May 29th, 2008)

[Baca, 1998] Baca, J. 1998. Comparing effects of navigational interface modalities on speaker prosodics. In *Proceedings of the Third international ACM Conference on Assistive Technologies* (Marina del Rey, California, United States, April 15 - 17, 1998). Assets '98. ACM Press, New York, NY, 3-10.

[Boostaroo, 2008] Boostaroo Portable Headphone Amplifiers, http://www.boostaroo.com/. (Checked on March 18th, 2008)

[Boron and Boulpaep, 2005] Boron, W. F., Boulpaep, E. L., *Medical Physiology, Updated Edition*. Elsevier Saunders, 2005.

[Bradford, 1995] Bradford, J.H. 1995. The human factors of speech-based interfaces: a research agenda. In *SIGCHI Bull.* **27**, 2 (Apr. 1995), 61-67.

[Brewster and Brown, 2004] Brewster, S. and Brown, L. M. 2004. Tactons: structured tactile messages for non-visual information display. In *Proceedings of the Fifth Conference on Australasian User interface - Volume 28* (Dunedin, New Zealand). A. Cockburn, Ed. ACM International Conference Proceeding Series, vol. **53**. Australian Computer Society, Darlinghurst, Australia, 15-23.

[Brewster et al., 2007] Brewster, S., Chohan, F., and Brown, L. 2007. Tactile feedback for mobile interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (San Jose, California, USA, April 28 - May 03, 2007). CHI '07. ACM, New York, NY, 159-162.

[Brown and Kaaresoja, 2006] Brown, L. M. and Kaaresoja, T. 2006. Feel who's talking: using tactons for mobile phone alerts. In *CHI '06 Extended*

Abstracts on Human Factors in Computing Systems (Montréal, Québec, Canada, April 22 - 27, 2006). CHI '06. ACM, New York, NY, 604-609.

[Brown et al., 2005] Brown, L. M., Brewster, S. A., and Purchase, H. C., A First Investigation into the Effectiveness of Tactons, in *Proceedings of the WorldHaptics 2005*, IEEE (2005), 167 - 176.

[Buttussi et al., 2006] Buttussi, F., Chittaro, L., and Nadalutti, D. 2006. Bringing mobile guides and fitness activities together: a solution based on an embodied virtual trainer. In *Proceedings of the 8th Conference on Human-Computer interaction with Mobile Devices and Services* (Helsinki, Finland, September 12 - 15, 2006). MobileHCI '06, vol. **159**. ACM, New York, NY, 29-36.

[Chang and O'Sullivan, 2005] Chang, A. and O'Sullivan, C. 2005. Audio-haptic feedback in mobile phones. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (Portland, OR, USA, April 02 - 07, 2005). CHI '05. ACM, New York, NY, 1264-1267.

[Cholewiak and Collins, 1991] Cholewiak, R.W. and Collins, A.A. (1991). Sensory and physiological bases of touch. In M. A. Heller & W. R. Schiff (Eds.), *The Psychology of Touch* (pp. 23-60). Hillsdale, N. J.: Lawrence Erlbaum Associates.

[Cholewiak and Wollowitz, 1992] Cholewiak, R.W., Wollowitz, M., "The design of vibrotactile transducers", Tactile Aids for the Hearing Impaired, I. Summers, ed, Whurr Publishers Ltd: London, 1992. pp 57-82.

[Clare, 1997] Clare, S., Functional MRI: Methods and Applications. University of Nottingham, Doctor's Thesis, October 1997. Also available as: http://users.fmrib.ox.ac.uk/~stuart/thesis/index.html. (Checked on April 5th, 2008)

[Cubase SX3, 2008] Cubase SX3 – The Music Creation and Production System, Steinberg Media Technologies GmbH., http://www.steinberg.net/35_1.html. (Checked on March 10th, 2008)

[Engineering Acoustics, 2008] The C2 Tactor, Engineering Acoustics, http://www.eaiinfo.com/. (Checked on February 16th, 2008)

[Enriquez et al., 2001] Enriquez, M., Afonin, O., Yager, B., and Maclean, K. 2001. A pneumatic tactile alerting system for the driving environment. In *Proceedings of the 2001 Workshop on Perceptive User interfaces* (Orlando, Florida, November 15 - 16, 2001). PUI '01, vol. **15**. ACM, New York, NY, 1-7.

[Fukumoto and Sugimura, 2001] Fukumoto, M., and Sugimura, T. 2001. Active click: tactile feedback for touch panels. In *CHI '01 Extended Abstracts on*

Human Factors in Computing Systems (Seattle, Washington, March 31 - April 05, 2001). CHI '01. ACM, New York, NY, 121-122.

[Guyton and Hall, 2006] Guyton, A. C., Hall, J. E., *Textbook of Medical Physiology, Eleventh Edition*. Elsevier Saunders, 2006.

[Hansson and Ljungstrand, 2000] Hansson, R. and Ljungstrand, P. 2000. The reminder bracelet: subtle notification cues for mobile devices. In *CHI '00 Extended Abstracts on Human Factors in Computing Systems* (The Hague, The Netherlands, April 01 - 06, 2000). CHI '00. ACM, New York, NY, 323-324.

[Immersion, 2008] Immersion Corporation, http://www.immersion.com. (Checked on January 29th, 2008)

[Insko et al., 2001] Insko, B. E., Meehan, M. J., Whitton, M. C., and Brooks, F. P. 2001 *Passive Haptics Significantly Enhances Virtual Environments*. Technical Report. UMI Order Number: TR01-010., University of North Carolina at Chapel Hill.

[Intille et al., 2004] Intille, S. S., Bao, L., Tapia, E. M., and Rondoni, J. 2004. Acquiring in situ training data for context-aware ubiquitous computing applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Vienna, Austria, April 24 - 29, 2004). CHI '04. ACM, New York, NY, 1-8.

[Kaaresoja and Linjama, 2005] Kaaresoja, T. and Linjama, J. 2005. Perception of Short Tactile Pulses Generated by a Vibration Motor in a Mobile Phone. In *Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic interfaces For Virtual Environment and Teleoperator Systems - Volume **00*** (March 18 - 20, 2005). WHC. IEEE Computer Society, Washington, DC, 471-472.

[Kandel et al., 2000] Kandel E. R., Schwartz J. H., Jessell T. M., *Principles of Neural Science, Fourth Edition*. McGraw-Hill Medical, 2000.

[Keinonen, 1997] Keinonen, T. 1997. Expected usability and product preference. In *Proceedings of the 2nd Conference on Designing interactive Systems: Processes, Practices, Methods, and Techniques* (Amsterdam, The Netherlands, August 18 - 20, 1997). S. Coles, Ed. DIS '97. ACM, New York, NY, 197-204.

[Lame, 2008] Lame mp3 encoder, http://lame.sourceforge.net/. (Checked on May 29th, 2008)

[Lee et al., 2004] Lee, J. C., Dietz, P. H., Leigh, D., Yerazunis, W. S., and Hudson, S. E. 2004. Haptic pen: a tactile feedback stylus for touch screens. In *Proceedings of the 17th Annual ACM Symposium on User interface Software and Technology* (Santa Fe, NM, USA, October 24 - 27, 2004). UIST '04. ACM, New York, NY, 291-294.

[Mallery, 2008] Mallery, C., Principles of biology at cellular, genetic, organismal levels. Department of Biology, University of Miami. Available as: http://henge.bio.miami.edu/mallery/150/. (Checked on February 4st, 2008)

[Mallinckrodt et al., 1953] Mallinckrodt, E., Hughes, A. L., and Sleator, J. W. Perception by the skin of electrically induced vibrations. In *Science* **118**, 3062, (1953), 277-278.

[Mazzone et al., 2003] Mazzone, A., Zhang, R., and Kunz, A. 2003. Novel actuators for haptic displays based on electroactive polymers. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (Osaka, Japan, October 01 - 03, 2003). VRST '03. ACM, New York, NY, 196-204.

[Miller, 1956] Miller, G. A., The Magical Number Seven, Plus or Minus Two: Some Limits on Our Capacity for Processing Information. In *The Psychological Review* **63** (1956), 81-97.

[Ng and Man, 2004] Ng, J., Man, J. 2004. Vibro-Monitor: A Vibro-tactile display for physiological data monitoring. In *Proceedings of the Human Interface Technologies Conference*, University of British Columbia, (Apr. 2004), 1–8.

[Nokia E50, 2008] Nokia E50 Smartphone, Nokia Corporation, http://business.nokia.fi/A4272027. (Checked on March 7th, 2007)

[Nokia N95, 2008] Nokia N95 Smartphone, Nokia Corporation, http://nseries.com/nseries/index.html#l=products,n95. (Checked on February 16th, 2008)

[Oregon Scientific, 2008] Oregon Scientific Speed & Distance HRM, http://www.oregonscientific.co.uk/prod_speedanddistancehrm.htm. (Checked on February 1st, 2008)

[Cranfill et al., 2006] Cranfill, D. B., Guzman, S. J., Isabelle, S. K., Saliba, R., Tactile transducers and method of operating, Patent assignee: Motorola Inc. (Schaumburg, IL), Issued on October 31st, 2006. Also available as: http://www.patentgenius.com/patent/7129824.html#show-page2. (Checked on April 22nd, 2008)

[Phantom, 2008] SensAble Phantom, http://www.sensable.com. (Checked on January 29th, 2008)

[Perez et al., 2003] Perez, C. A., Santibañez, A. J., Holzmann, C. A., Estévez, P. A., Held, C. M., Power requirements for vibrotactile piezo-electric and electromechanical transducers. In *Medical and Biological Engineering and Computing* **41**, 5 (Nov. 2003), 718-726.

[Pertovaara et al., 1992] Pertovaara, A., Kemppainen, P., Leppänen, H., Lowered cutaneous sensitivity to nonpainful electrical stimulation during isometric exercise in humans. In *Experimental Brain Research* **89**, 2 (May. 1992), 447-452.

[Polar Electro, 2008] Polar Electro Heart Rate Monitors, http://www.polarusa.com. (Checked on March 15th, 2008)

[Post et al., 1994] Post, L. J, Zompa, I. C., Chapman, C. E., Perception of vibrotactile stimuli during motor activity in human subjects. In *Experimental Brain Research* **100**, 1 (Jul. 1994), 107-120.

[Poupyrev et al., 2002] Poupyrev, I., Maruyama, S., and Rekimoto, J., Ambient touch: designing tactile interfaces for handheld devices. In *Proceedings of the 15th Annual ACM Symposium on User interface Software and Technology* (Paris, France, October 27 - 30, 2002). UIST '02. ACM, New York, NY, 51-60.

[Reeves et al., 2004] Reeves, L. M., Lai, J., Larson, J. A., Oviatt, S., Balaji, T. S., Buisine, S., Collings, P., Cohen, P., Kraal, B., Martin, J., McTear, M., Raman, T., Stanney, K. M., Su, H., and Wang, Q. Y., Guidelines for multimodal user interface design. *Commun. ACM* **47**, 1 (Jan. 2004), 57-59.

[Shneiderman, 1998] Shneiderman, B., *Designing the User Interface: Strategies for Effective Human-Computer Interaction, Third Edition*. Addison-Wesley, Reading (MA), 1998.

[Thibadeau, 2008] Thibadeau, R., Tutorial on the rotary or swing arm voice coil actuator used on modern hard disk (disc) drives. Carnegie Mellon University. Available as: http://www.internetlab.ri.cmu.edu/rotaryvoicecoil/. (Checked on February 28th, 2008)

[Tikka and Laitinen, 2006] Tikka, V., Laitinen, P., Designing Haptic Feedback for Touch Display: Experimental Study of Perceived Intensity and Integration of Haptic and Audio. In *Haptic and Audio Interaction Design*. Springer Berlin / Heidelberg, 2006, 36-44.

[Töyssy, 2007] Töyssy, S., Telling time by vibration. Department of Computer Sciences, University of Tampere, Master's Thesis, November 2007. Also available as: http://www.cs.uta.fi/research/theses/masters/Toyssy_Sampo.pdf. (Checked on April 5th, 2008)

[Vadas et al., 2006] Vadas, K., Patel, N., Lyons, K., Starner, T., and Jacko, J. 2006. Reading on-the-go: a comparison of audio and hand-held displays. In *Proceedings of the 8th Conference on Human-Computer interaction with Mobile Devices and Services* (Helsinki, Finland, September 12 - 15, 2006). MobileHCI '06, **159**. ACM, New York, NY, 219-226.

[Eibeck and Muramatsu, 2007] Eibeck P. A., Muramatsu, B., Voice coil theory of operation. Available as: http://bits.me.berkeley.edu/beam/vc_2.html. (checked on February 25th, 2008)

[Wheeler, 2008] Wheeler, J., Haptic Display for Active People. Mobile Haptics Project. University of Tampere, Finland & Stanford University, USA. http://bdml.stanford.edu/twiki/bin/view/Haptics/HapticsInMotion. (Checked on May 29th, 2008)

[Winamp, 2008] Winamp Media Player, http://www.winamp.com/. (Checked on April 5th, 2008)

[Yamamoto et al., 2004] Yamamoto, A., Nagasawa, S., Yamamoto, H., and Higuchitokyo, T., Electrostatic tactile display with thin film slider and its application to tactile tele-presentation systems. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (Hong Kong, November 10 - 12, 2004). VRST '04. ACM, New York, NY, 209-216.

[Zimbardo et al., 1995] Zimbardo, P., McDermott, M., Jansz, J., Metaal, N., *Psychology: A European Text*, Harper & Collins, 1995.