MudPad

Localized Tactile Feedback on Multi Touch Surfaces

Diploma Thesis at the
Media Computing Group
Prof. Dr. Jan Borchers
Computer Science Department
RWTH Aachen University

by
Yvonne Jansen

Thesis advisor:
Prof. Dr. Jan Borchers

Second examiner:
Prof. Dr. Martina Ziefle

Registration date: July 13th, 2010
Submission date: Mar 15th, 2011

I hereby declare that I have created this work completely on my own and used no other sources or tools than the ones listed, and that I have marked any citations accordingly.

Hiermit versichere ich, dass ich die vorliegende Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt sowie Zitate kenntlich gemacht habe.

Aachen, March 2011
Yvonne Jansen

Contents

List of Figures

List of Tables

Abstract

Touch screens are widespread as of today. They provide a high level of directness as the visual and physical area of interaction fall together. But due to their smooth glass surface they do not provide any tactile feedback upon user actions. We introduce MudPad, a system capable of localized active haptic feedback on multitouch screens. We use an array of electromagnets combined with an overlay containing magnetorheological (MR) fluid to actuate a tablet-sized area. As MudPad has a very low reaction time it is able to produce instant multi-point feedback for multitouch input, ranging from static levels of surface softness to a broad set of dynamically changeable textures. Our system does not only convey global confirmative feedback on user input but allows UI designers to enrich the entire interface with a tactile layer conveying local semantic information. This also allows users to explore the interface by touch.

A quantitative user study suggests that it performs similar to standard touch screens in terms on typing speed. Furthermore, users detect their errors in time and correct them. Users from several exhibitions describe the interaction with MudPad as pleasant and were often engaged with the system for an extended time.

Überblick

Touchscreens sind mittlerweile weit verbreitet. Sie erlauben Benutzern die direkte Interaktion genau dort, wo die manipulierten Information dargestellt werden. Aufgrund ihrer glatten Glasoberflächen bieten sie jedoch keinerlei taktile Rückmeldung über Eingaben. Wir präsentieren mit MudPad ein Touchscreen System, das *lokalisiertes aktives haptisches Feedback* gleichzeitig an mehreren Orten erlaubt. Eine Matrix von Elektromagneten aktuiert eine mit magnetorheologischer Flüssigkeit gefüllte Oberfläche mit der Benutzer interagieren können. Die Reaktionszeit dieser Flüssigkeit ist sehr gering, so dass es für Echtzeitfeedback nutzbar ist. Das Feedback wird durch dynamische Manipulation der Festigkeit der Flüssigkeit durch die Elektromagneten erzielt. Damit sind nicht nur globale Systembestätigungen möglich, sondern die gesamte grafische Oberfläche kann mit taktilen Informationen versehen werden. Dadurch können Interfaces allein durch Berührung erschlossen werden.
Eine quantitative Studie unterstützt die Annahme, dass MudPad in Relation zu herkömmlichen Touchscreens keine signifikant schlechten Einflüsse auf die Eingabegeschwindigkeit hat. Zudem erhöht es die Wahrscheinlichkeit, dass Benutzer frühzeitig bemerken, dass sie Fehler gemacht haben. Während mehrerer Ausstellungen beschrieben zahlreiche Besucher die Interaktion mit MudPad als sehr angenehm und haben den Prototypen ausgiebig getestet.

Acknowledgements

I want to thank everybody who helped me with this thesis. In particular, I thank my supervisors Prof. Jan Borchers for valuable input on my topic, Prof. Martina Ziefle for always taking the time to give some advice, and my advisor Thorsten Karrer for all his help and especially for helping me to get this project published.
I thank the Fablab crew, especially René Bohne and David Peters, for sharing their knowledge.
And finally, I thank Julie Wagner for taking the time to review parts of this thesis and for providing constructive feedback.

Chapter 1

Introduction

Figure 1.1: Ferrofluid exposed to a magnetic field showing spikes in direction of the magnetic flux lines.

"Computer Science is no more about computers than astronomy is about telescopes."

—E. W. Dijkstra

Feedback in interactive systems.

Feedback is vitally important for efficient and pleasant interaction in everyday life and even more so for human computer interaction. With well-designed special purpose devices we use throughout the day it is easy to tell what they do and how to use them. But computer systems are much more complex than that. They can be programmed to do whatever a programmer wants them to do. To interact with them, we almost always use the same indirect physical interface devices such as a display, a keyboard and a mouse. Depending on the task the mappings between this physical devices and the actions they invoke change and can be completely different.

Touch screens and direct manipulation.

Over the last couple of years touch screens became very popular. The majority of currently available mobile devices such as phones, PDAs and music players are now equipped with a touch screen. Here, the screen is also the input area, thereby presenting users with a higher level of *directness* (Beaudouin-Lafon [2000]). The concept of *direct manipulation* (Shneiderman [1982], Hutchins et al. [1985]) can almost be taken literally: the graphical representation on the screen and the interaction/manipulation area fall together.
Users usually find touch interaction engaging and intuitive whereas mouse and keyboard imply to navigate through menus and submenus or to remember dozens of shortcuts to invoke more complex actions. Even for rather simple tasks such as navigating a 3D model command keys on the keyboard must be pressed to differentiate between actions such as zooming, panning, and rotating. On touch screens such actions can be performed by simple and easy to learn gestures (Nielsen et al. [2004]). Although these gestures still require a learning phase the memory load is reduced if the gesture is well-designed (Wobbrock et al. [2009]). Additionally, recalling a gesture involves muscle memory and is less abstract than pressing the Ctrl or Alt key on a keyboard.

Touch screens lack tactile feedback.

One drawback of touch screens however is their complete lack of tactile feedback. With keyboard and mouse, a user gets physical, i.e., mechanical feedback from pressing a key, moving the mouse, or pressing a mouse button. Furthermore, these devices provide physical clues such as the small bumps on the "F" and "J" key to allow eyes-free positioning of the fingers on the keyboard. With a touch screen users only get the sense of their finger pressing against a smooth glass surface. There is no telling by touch whether an intended action was recognized by the system. The only way of communicating this to the user is to give salient visual or auditory feedback.

First devices with *global* tactile feedback.

By now, there are a few mobile phones on the market which make use of simple actuators to give tactile feedback, e.g., during text entry. Note though that these devices only use global actuation, i.e., a simple vibrotactile feedback signal slightly displaces the screen. Previous studies point to decreased memory load (Brewster et al. [2007]) and increased speed and accuracy (Hoggan et al. [2008]). For multi touch enabled devices it is not yet possible to create signals for each touch respectively. While this is acceptable for small form factor devices that are mostly only used with one finger, it is preferable to present independent feedback at each screen location in relation to the information or event at that location.

The system we propose is a first attempt at investigating possible interactions for localized active haptic feedback presented simultaneously at multiple points. Apart from the resolution the flexibility of the design is similar to GUI interfaces because the feedback patterns for each screen location can be changed and adapted dynamically. Hence the entire GUI can be enriched with real-time tactile feedback.

1.1 Objective of this thesis

Our contribution.

By adding tactile feedback to a GUI we aim to improve touch screen interactions and to lower the memory load necessary to operate them. An example for memory load reduction would be to reproduce the physical attributes of common GUI objects such as buttons. By presenting a tactile impression similar to the well-known feeling of pressing a physical button, no further evaluation of, e.g., visual feedback is necessary for a user to determine if she actually pressed the button she intended to. This works even better in case of erroneous input where a button gets pressed several times – because of the higher temporal resolution of tactile perceptions such an error is easily detectable through tactile feedback.

Structure of this thesis.

Before we describe the actual system design, we will first introduce the basic principles of human haptic perception. Then we give an overview over touch screen technology, approaches to haptic interaction, and evaluating studies on both. We then describe the design cycle for the hardware development that led to the final design and present different approaches on how such a system could be put to use for common tasks. Before we conclude, we present a pilot study to evaluate the efficiency in terms of speed and error rate in comparison to regular touch screens. We will also outline informal feedback gained from users during various demonstrations during conferences and an art festival. Finally, we will point out possible improvements on current limitations and directions for future studies.

Chapter 2

Human Tactile Perception

"The secret of being a bore is to tell everything."

—Voltaire

In this chapter we give a short introduction into *human tactile perception* to give an idea of its resolution and limitations.

The tactile perception is part of the haptic system that "uses sensory information derived from *mechanoreceptors* and *thermoreceptors* embedded in the skin ("cutaneous" inputs) *together with* mechanoreceptors embedded in muscles, tendons, and joints ("kinesthetic" inputs)" (Lederman and Klatzky [2009]). These receptors are spread over the entire skin area of the human body. We are mostly interested in impressions gained from the fingertips and set the focus for these explanations accordingly.

Tactile perception in relation to other senses.

The temporal and spatial acuity of the tactile sensory system differs immensely from the visual and auditory channels which are prominently employed for human computer interaction (cf. Table 2.1). Tactile impressions from the fingertips are higher grained than those from, e.g., the calves pointing to a huge variety in receptor density. Although the temporal acuity of tactile perceptions is much higher than the eye's, the total information capacity at a fingertip

	Information capacity (bit/s)	Temporal acuity (ms)
Fingertip	10^2	5
Ear	10^4	0.05
Eye	10^6–10^9	25

Table 2.1: Comparison of the three main sensory modalities (Chouvardas et al. [2008]).

is much lower than the ear's or the eye's (see Table 2.1 for a direct comparison).

Four different receptors with individual specializations.

Tactile perception is part of the somatosensory system. A variety of receptors, each one specialized on a certain sensation, detect these signals and feed them up the spinal cord to the somatosensory cortex. Figure 2.1 gives an overview over the distribution of the different receptors within the dermal layers both for hairy and glabrous (hairless) skin. Table 2.2 complements this graphic by supplying further information about the different receptor types and what triggers each of them respectively.

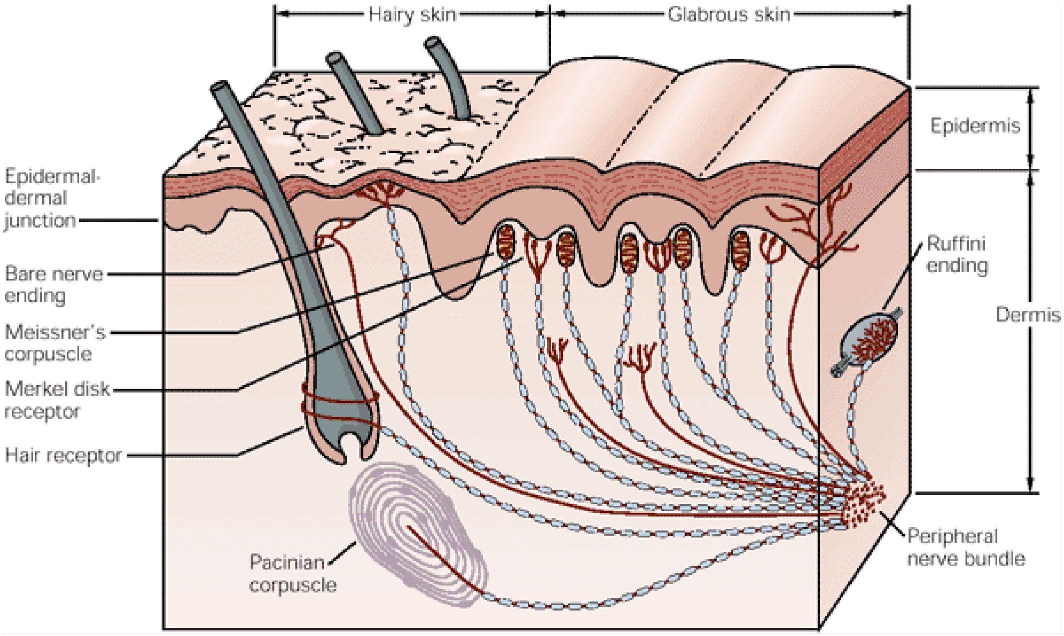

Figure 2.1: Skin cross section showing the four different tactile receptors and their position within the dermal layers (Gardner et al. [2000]).

Receptor	Class	Sense modality	Frequency range (highest sensitivity)	Density at fingertip (per cm^2)
Meissner	RA1	Stroking, fluttering	10–200Hz (200–300Hz)	140
Merkel	SA1	Pressure, texture	0.4–100Hz (7Hz)	70
Pacinian	RA2	Vibration	40–800Hz (200–300Hz)	21
Ruffini	SA2	Skin stretch	7Hz	9
Hair follicle	RA	Stroking, fluttering	?	–
Hair	–	Light stroking	?	–
Field	–	Skin stretch	?	–

Table 2.2: Compilation of the human mechanoreceptors and the stimuli they are receptive for (Chouvardas et al. [2008]).

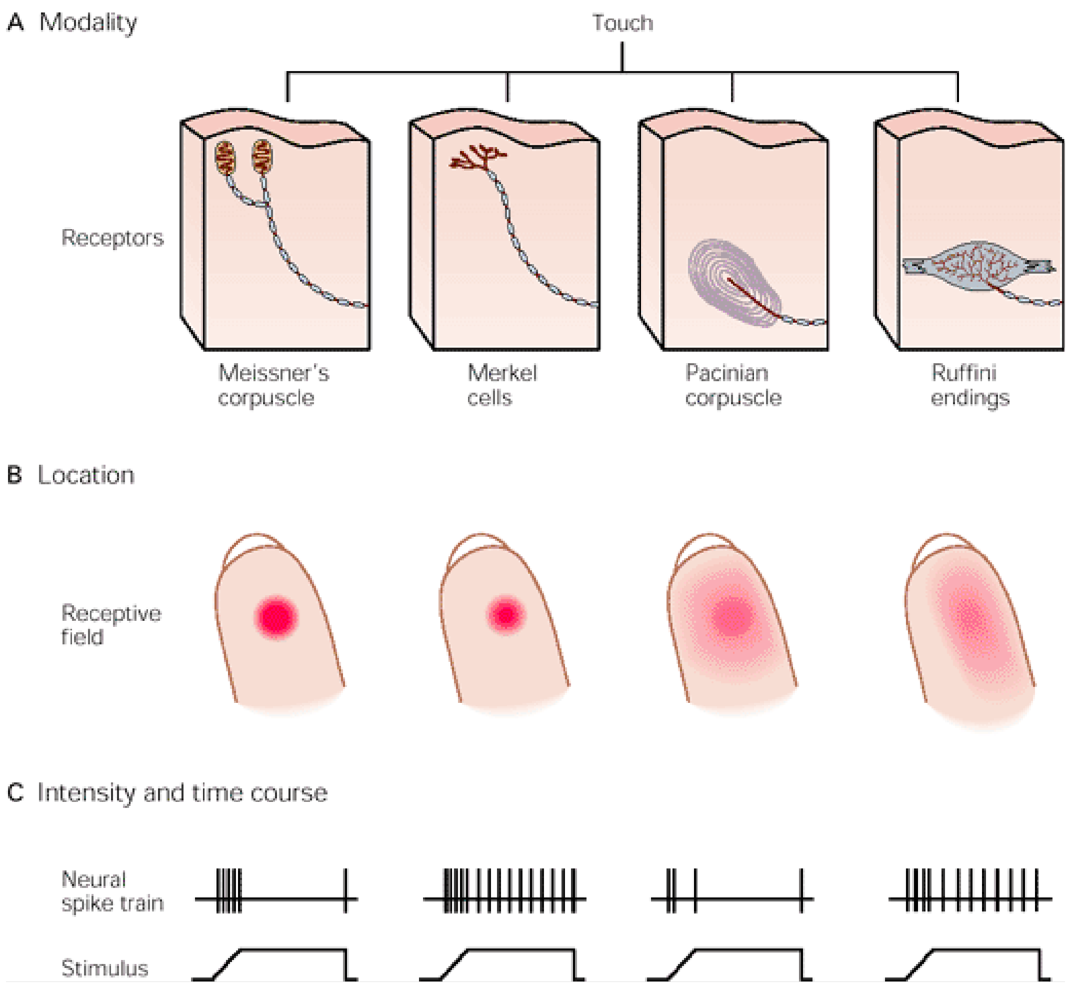

Figure 2.2: Illustration of the receptive field sizes and adaptation rates for the four different tactile receptors (Gardner and Martin [2000]).

Receptors adapt to stimuli.

Receptor adaptation. The receptor types differ in their adaptation behavior. Adaptation means in this context that a receptor adapts to a sensation and stops firing or reduces its fire rate even though the stimulus is still present. Table 2.2 lists two different sensor classes: slowly-adapting (SA) and rapidly-adapting (RA) receptors. The SA type is, e.g., necessary to estimate the amount of pressure one needs to carry a glass without breaking or dropping it. It would be impossible to hold a glass for an extended time if those receptors would stop firing after a while. The RA receptors are most susceptible to vibration. If those would not adapt to a rumpling car or train, a longer trip would be almost unbearable. Figure 2.2c illustrates the neural firing behaviors for the respective receptors to the same stimulus. The SA receptors keep firing while the RA types only react to changes in stimulus intensity, i.e., in its onset and offset phase.

Receptive fields. The receptors can furthermore be differentiated by the size of their *receptive fields*. Again there are two different types: type 1 with small and type 2 with large receptive fields. The size of the field correlates to the density of the receptors as the last column of table 2.2 shows for the fingertip. Consequently the spatial acuity correlates to the density – two locations on the skin can only be distinguished as such if different receptors can pick up the stimuli. The receptive field of a receptor can be measured by identifying a nerve for a certain skin area and by recording its signals. Fingertips have a very high density of receptors with small fields and therefore both the localization and the differentiation is very good, i.e., in the area of 1–2mm. In general, the lower skin layers contain receptors with larger receptive fields whereas the receptors in the epidermis have small receptive fields.

Different sized receptive fields result in variable spatial acuity.

We can conclude that the spatial resolving power of the skin is poorer than the eye's but better than the ear's. The temporal resolving capacity is much better than the eye's where a succession of still images higher than 20Hz is already perceived as moving. But it is much worse than the ear's as successive taps on the skin with a gap of less than 5ms can hardly be distinguished as such.

Vision-Touch Interactions. The sensory information gained from visual and tactile receptors needs to be integrated to form one coherent perception. The intersensory interactions during this integration process can accelerate and, in case of conflicting information, decelerate the process (e.g., Butter et al. [1989]). These interactions have been investigated by perception researchers by comparing data from unimodal with that of bimodal or multimodal conditions. By presenting intersensory conflicts, i.e., non-consistent information from different sensory inputs,

Intersensory integration.

the weighting of the modalities can be inferred when compared to unimodal reaction times (Ernst and Banks [2002]).

Chapter 3

Related work

"Great things are not done by impulse, but by a series of small things brought together."

—Vincent van Gogh

This body of related work is an intersection of different areas of previous research projects. We roughly follow the line of argumentation from the introduction. First we present popular technologies for (multi-) touch input and their implications and limitations for user interaction. This is necessary as most of these approaches are not compatible with our system design. Next, we will present some studies evaluating touch screens in comparison to mouse interaction for direct manipulation tasks. We then give a more focused survey of research on haptic feedback for touch screens and advance to evaluating studies on the effectiveness of tactile feedback.

3.1 Touch Technologies

By now there are many different touch sensing technologies available. Each only works in a certain environment and most are not suitable for our approach. We will give an overview here for a later discussion on the most suitable technology for our setup.

Capacitive sensing. Most of the mobile touch screen devices available today, e.g., the Apple iPad, use capacitive sensing. It is a very mature and reliable technology. Lee et al. [1985] introduced the first multi touch tablet which even included relative pressure sensing. Hardware-based sensing of pressure is not possible with capacitive sensing, but Lee et al. interpreted changes of touch sizes as changes in pressure. That way the tablet could be operated in a similar way as a mouse – navigate with light pressure to a target and push harder to select. Later, Rekimoto [2002] created a similar system to sense multi-finger gestures and tangible objects placed on a surface. Due to the capacitive sensing, the objects needed to be touched to become active and invoke changes in the capacitive field.

Capacitive sensing is mature and reliable.

Zimmerman et al. [1995] also measured changes of small electric fields caused by hands in an interaction space or entire persons in an interactive room. Dietz and Leigh [2001] (DiamondTouch) integrated an emitting electrode in a chair to sense multiple user inputs on a tabletop surface. A major benefit of this technique is data about ownership of the sensed touches. As of today, it is still the only technology that can reliably relate touches to users.

DiamondTouch can distinguish between users.

The major drawback of capacitive sensing is the need to operate the surface with bare fingers or special styli. As soon as a non-conductive layer such as gloves is between surface and finger the surface any input attempt fails.

Capacitive sensing needs bare finger contact.

Resistive sensing. With resistive touch screens users have a wider choice how to enter data. The base technology used is called *force sensitive resistor (FSR)* (Eventoff [1984]). As the name already implies, it is pressure based and can be used with arbitrary objects to induce that pressure on the screen. The simplest implementations are called 4-wire sensors and are common, e.g, in touch screen ticket machines. They depend on two flexible layers of resistive material that are separated by spacer dots and connected to wire electrodes at the edges. One of these layers is connected to a 5V power source and the voltage induced on the other layer is measured. The position of a touch is calculated by interpreting these two layers as a voltage divider.

Resistive sensors detect arbitrary objects by pressure.

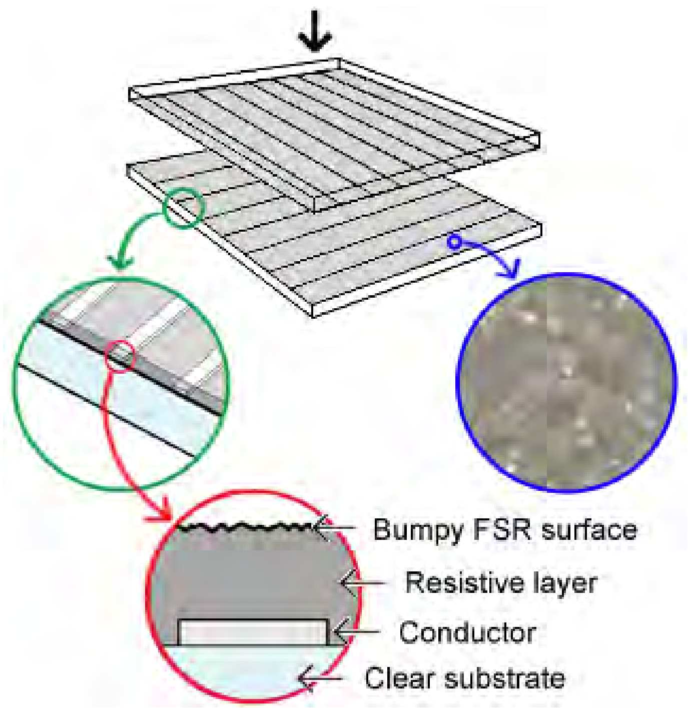

Figure 3.1: Schematic diagram of the Unmousepad [2009]: two arrays of electrodes are sandwiched with a resistive layer that provides interpolation on the analog layer to increase the resolution to >100dpi

By using rows and columns of electrodes instead of continuous layers, several input points can be distinguished (e.g., Joguet and Largillier [2007]). Recently Rosenberg and Perlin [2009] introduced a method for high-resolution multi touch input by introducing a continuous resistive layer between the electrode array which allows an accurate interpolation (see Figure 3.1 and sensing of lightweight objects (5grams). Furthermore, by scanning each column/row intersection consecutively, occlusion effects are eliminated and arbitrary combinations of touches can be detected.

The Unmousepad allows high resolution multi touch input.

Optical sensing relies on camera images.

Optical sensing. All vision based approaches rely on infrared light in the area of 850–950nm that is tracked by a camera equipped with an infrared bandpass filter to block out all other wavelengths. Spatial and temporal resolution of these systems solely depends on the camera used.

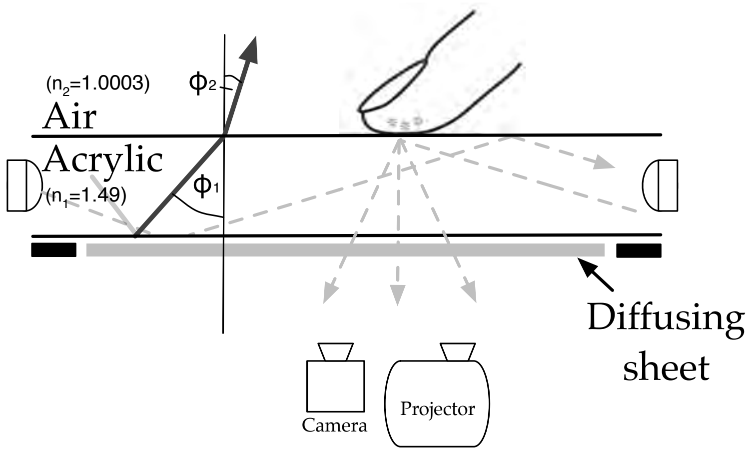

Figure 3.2: Schematic diagram illustrating the total internal reflection of the light rays for a certain angle ϕ. Source: Wagner [2009]

FTIR is cheap and scales well.

Frustrated Total Internal Reflection (FTIR). Multi touch interaction probably had its biggest increase in popularity with a talk of Jefferson Han at TED2006. The approach he presented is cheap and scalable, and can be reproduced by everyone with basic electronic skills. Since then, a very active community arose, developing simple to follow guides on how to build such a system for a small amount of money. It is also a popular technique for researchers who want to track fingers in a very reliable way. The setup guarantees bright blobs and high contrast for everything that is pressed on the surface.

The basic setup as proposed by Han [2005] is shown in Figure 3.2. An optical waveguide (usually a sheet of acrylic) is flooded with infrared light through its edges. The light reflects internally and fills the plate with infrared light. Upon a touch, the internal reflection is frustrated and the light falls out of the acrylic plate to a camera pointed at the sheet from below.

Moeller and Kerne [2010] proposed a design that measures the loss of light at the edge opposite to the infrared LEDs due to touches. The approach allows a thin form factor as no camera from below is necessary. But it is also prone to false negatives due to occlusion as are all systems that compute touch information from a side view.

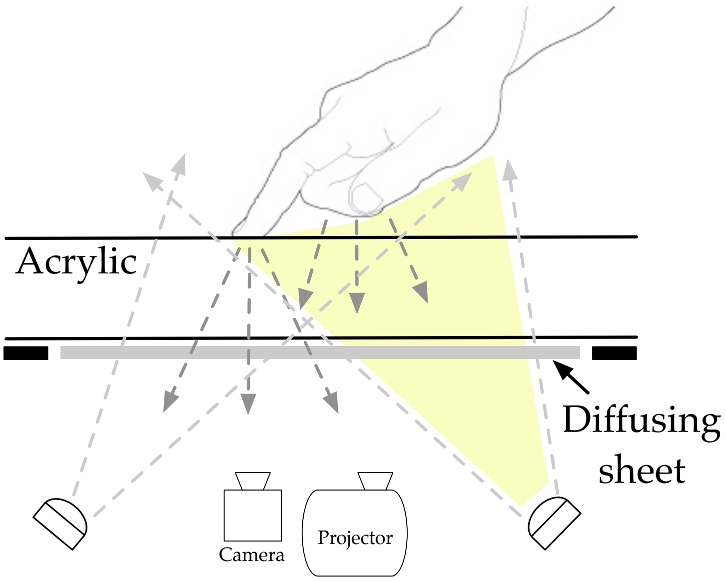

Figure 3.3: Schematic diagram for diffuse illumination setup. Source: Wagner [2009]

Diffuse Illumination (DI). In contrast to the FTIR technology, diffuse illumination is perfect for tracking arbitrary objects placed on the surface. The whole rear of the touch surface is flooded with infrared light and objects can be detected by light reflecting from their bottom. This is the technology used for the *Microsoft Surface*. Due to the flooded light, the contrast is much lower than with FTIR. Hence, it is a bit harder to reliably track fingers. But the main advantage of this technology is to be able to track other objects besides fingers such as the fiducial markers used for the *ReacTable* project (Jordà et al. [2007])

DI setup is very sensitive to ambient light.

Hodges et al. [2007] proposed a more compact design by using a discrete array of light sensors and infrared light emitters mounted behind a LCD. Although the thin form factor is a big advantage, it is also less scalable than a projector/camera based setup. Hofer et al. [2009] built upon this setup by adding a FTIR sheet in front of the LCD. A flat but better scalable design was proposed by Jackson et al. [2009]. They incorporated an array of fibers to channel the light from below the touch surface to the side thereby gaining a thin form factor. This design works with both DI and FTIR setups.

Improvements on basic FTIR and DI setups.

Hilliges et al. [2008] proposed a system depending on a malleable pouch containing inked water. The pouch is flooded with phosporescent light that bounces off the inner surface of the pouch when a touch occures.

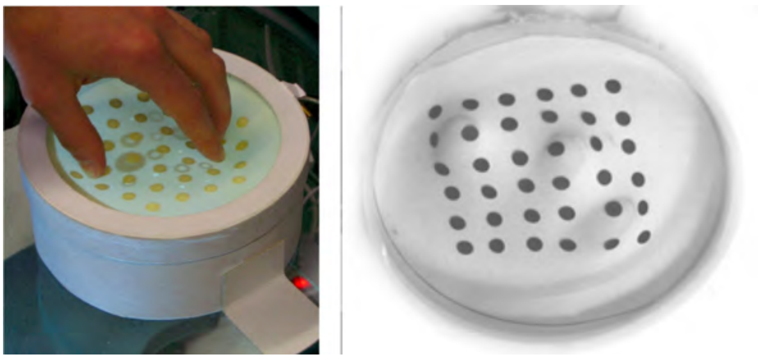

Figure 3.4: Left: Milczynski et al. proposed to place a marked malleable surface over a camera. The *elastable* surface also provides passive tactile feedback. Right: camera image of a very similar system by Vogt et al. [2004]

Other vision based approaches. Both, Vogt et al. [2004] and Milczynski et al. [2006] (Figure 3.4) placed a malleable surface marked with a point grid above a camera to detect finger input by analyzing the distortion of the surface. Although multiple fingers could be detected, the approach is not scalable as the distortion decreases with increasing surface size.

Other techniques. Steurer and Srivastava [2003] proposed a discrete grid of either metal plates or hall sensors to detect objects placed on the surface. The objects were passive, i.e., did not contain any power sources but were equipped with either metal plates or magnets depending on the type of grid used for sensing.

Figure 3.5 shows the market share of the different technologies for 2010. Resistive and capacitive sensing clearly dominate the market.

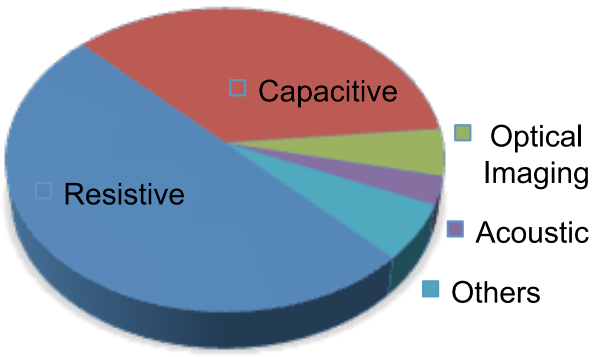

Figure 3.5: World-wide market share for 2010 of current touch input technology. Note that almost all of those resistive touch screens are single-touch panels. Source: DisplaySearch 2010 Touch Panel Market Analysis

Technology	Physical affordances	Interaction restrictions	Reliability	Availability	Sensing of ...	Implementations
Capacitive	transparent array of electrodes, limited scaling	direct finger contact necessary	high	high	only fingers and special pens	Apple iPhone
Resistive	thin layer above display	–	high	high for single, low for multi touch	any object	Unmousepad [2009]
FTIR	camera behind light carrier	sufficient pressure necessary	rather high	high, popular DIY technique	anything that exerts force	Han [2005]
DI	camera and diffuse infrared lighting from behind	–	depends on ambient light	high, popular DIY technology	anything	ReacTable [2007], Microsoft Surface
Fluid pouch	camera behind	top projection	medium	low	anything that exerts enough force	Hilliges et al. [2008]
Metal or hall sensor grid	passive objects	no finger input	high	low	metal or magnets	Steurer and Srivastava [2003]

Table 3.1: Comparison of touch input techniques

3.2 Haptic Feedback

In this section we give a quick survey on previously proposed systems. The area of haptic interaction is very wide and researchers from many different disciplines contribute to it. We focus here on haptic devices that are somewhat related to touch interaction. The structure is similar to the previous section as we will first present proposed systems and then studies that investigate possible advantages of tactile feedback.

Force-feedback vs. tactile feedback.

Haptic interfaces can be classified as either providing tactile or force-feedback devices. Force-feedback is directed to proprioceptive perception, i.e., the devices communicates force information to a user by limiting and resisting her movements. Tactile feedback targets the mechanoreceptors in the outer skin layers through slight pressure or vibration. The most common actuators for tactile feedback are voice-coil motors and piezoelectric actuators. While the possible feedback with voice-coils is rather limited – they have a fixed vibration frequency – piezoelectric devices are more flexible and can be thought of as small speakers utilizing the sound pressure to create vibrations. Although it is possible to combine both in one system, there are currently no available implementations (Roberts and Paneels [2007]).

Commercial systems provide high resolution force-feedback haptic space.

Commercially available products. By now there are a few commercially available high resolution force-feedback devices like the SensAble Phantom Massie and Salisbury [1994] (see Figure 3.6) or the HapticMaster by FCS Control Systems Van der Linde et al. [2002]. These systems can simulate forces and resistance for movements, i.e., they usually present a user with a pen like handle for interaction and when a user moves it or exerts force on it, she can feel forces and resistance. Due to the handle used, these systems only provide a one-point access to a high-resolution haptic space. Unfortunately they provide only force-feedback and do not provide tactile feedback. Such a combination would allow for much richer feedback.

Figure 3.6: The SenseAble Phantom force-feedback device depicted in a 3-D modeling scenario (Massie and Salisbury [1994]).

Haptic Pen.

Actuated workbench.

Indirect devices. Poupyrev et al. [2004] gives indirect tactile feedback through a pen by embedding TouchEngine actuators underneath the screen of a Wacom tablet. Similar to the Phantom, a user also holds a handle/pen whereas here she gets tactile instead of force-feedback. The Haptic Pen for stylus based touch screens Lee et al. [2004] is a similar approach although here the actuator is embedded in the pen instead of the screen. These design suffer from their indirectness. Tactile properties are perceived less precise when explored through a probe instead of directly with a finger Lederman and Klatzky [2009]. With actuated workbench (Pangaro et al. [2002]) and more recently Madgets (Weiss et al. [2010]), the focus is different. They both use an array of electromagnets to move tangibles on a tabletop surface. While the technical setup is similar to ours, these can rather be interpreted as force-feedback devices although the focus of this work lies on arranging tangibles according to a visual overlay, e.g., for remote collaboration purposes.

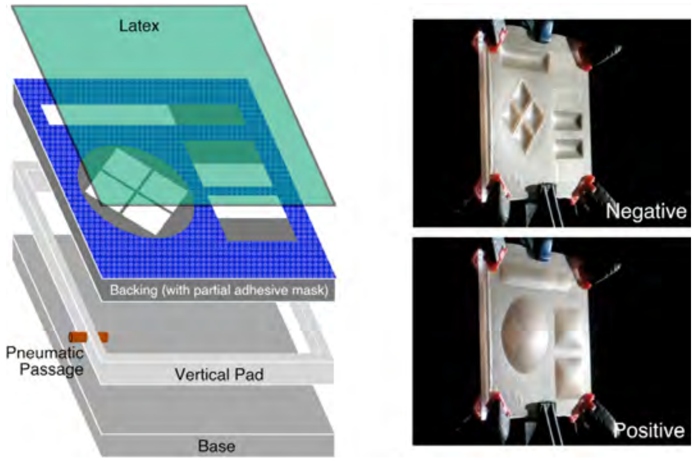

Figure 3.7: Pneumatically created positive and negative shapes to make buttons more tangible (Harrison and Hudson [2009]).

Shape and height displays. Several systems that can be seen as shape or height displays aim to improve button interaction by giving these GUI elements a physical form. Harrison and Hudson [2009] introduced dynamically changeable, i.e., pneumatically inflatable buttons on a touch display (see Figure 3.7). While the level of inflation can be adapted dynamically the overall outline of the button interface is fixed once the system is assembled.

Inflatable buttons.

With BubbleWrap (Figure 3.8), a textile-based electromagnetic haptic display (Bau et al. [2009]), the physical alignment of the buttons is also fixed. This system uses a small electromagnet in each button that repels a small permanent magnet to present a user with force-feedback when she tries to push the button.

BubbleWrap

Relief, a scalable actuated shape display (Leithinger and Ishii [2010]), presents a user with an array of 120 aluminum pins whose height can be manipulated by commercially available electric slide potentiometers that also sense user input (Figure 3.9). The pinarray is covered by a flexible sheet to top-project visual information. Although localized actuation is possible with this design, the possible feedback is very limited due to the motor sliders used.

Relief

Wagner et al. [2004] proposed a small 6x6 pin-array shape display mounted on optical mouse tracker. While this setup

Figure 3.8: BubbleWrap haptic display prototype (Bau et al. [2009]).

Haptic tabletop puck.

can give localized feedback, it depends on indirect interaction through the mouse tracking. A similar approach was taken by Marquardt et al. [2009] with the Haptic Tabletop Puck. Here, the user gets only one interaction point per puck. As it is designed for use on an interactive tabletop, the interaction is more direct than with the previous system but still requires an additional device between the touch screen and the user's finger.

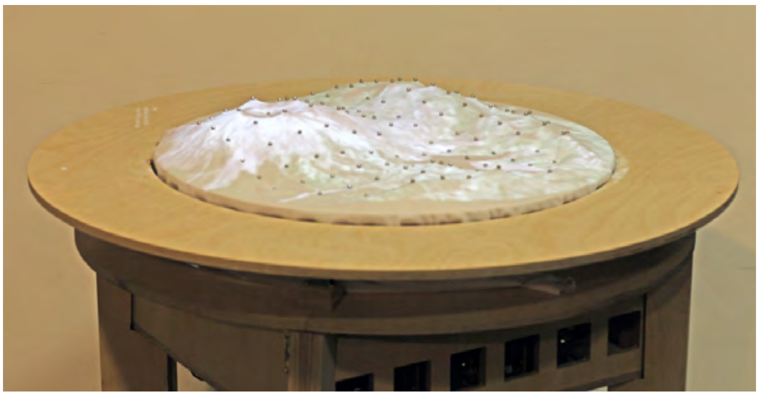

Figure 3.9: Relief system with pins covered with Lycra surface and top projected landscape (Leithinger and Ishii [2010]).

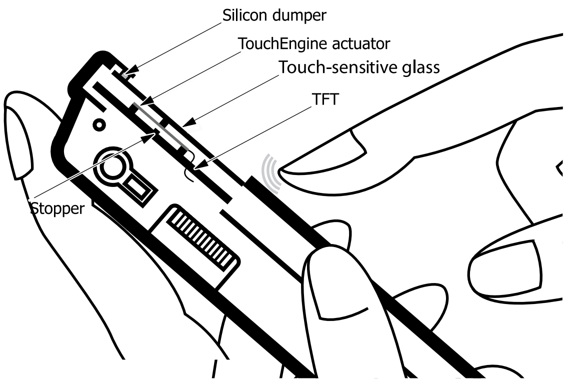

Figure 3.10: Haptic display for small touch screens with a piezo actuator underneath the screen to give feedback by slightly displacing the whole screen (Poupyrev and Maruyama [2003]).

Mobile devices. For mobile devices current research can be divided into two groups: systems that apply the feedback directly to the screen and devices which are equipped with additional actuators at the sides or the back. One of the first systems that suggested a design for vibrotactile feedback for PDAs and touch screens was Active Click (Fukumoto and Sugimura [2001]). It included different implementations for both direct screen application and additional actuators at the back of an PDA. By mounting actuators underneath the screen and hence displacing the whole screen, they were able to produce simple affirmative feedback. A more sophisticated approach (see Figure 3.10) is Ambient Touch and its follow-up TouchEngine (Poupyrev et al. [2002], Poupyrev and Maruyama [2003]). The PreSenseII (Rekimoto and Schwesig [2006]) system is a pressure sensitive input pad with an actuator mounted underneath the touch surface. It aims to model tactile feedback to acknowledge button presses (see Figure 3.11 for details on how tactile feedback was mapped to user actions). The mounting principle is the same as suggested for Active Click, but they used layers of piezoceramic films to create more complex signals and patterns. The additional feedback proved to be an advantage but still, the feedback was globally applied to the whole pad and hence limited in expressiveness. ComTouch, a vibrotactile communica-

Active click.

Ambient touch and TouchEngine.

PreSenseII.

ComTouch.

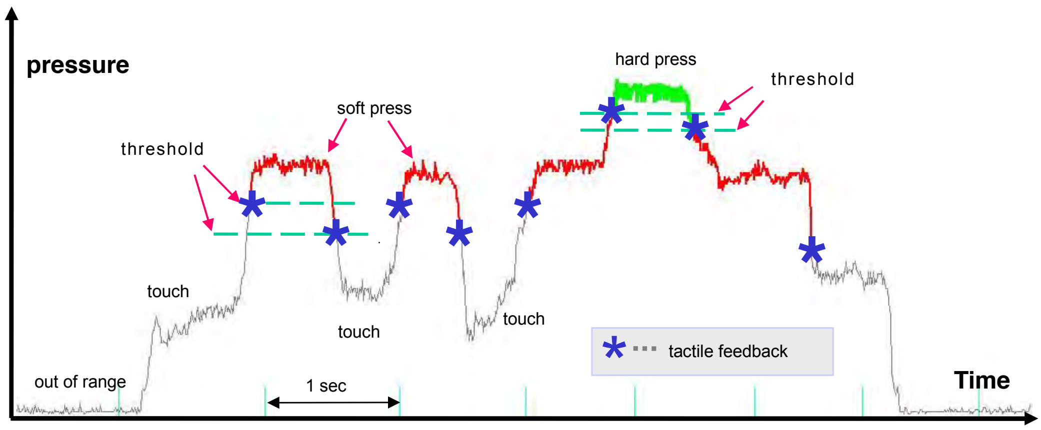

Figure 3.11: Pre-SenseII multi-level button operation and corresponding pressure values (Rekimoto and Schwesig [2006]).

Multi-actuator PDAs.

tion device (Chang et al. [2002]) creates directional feedback through several actuators attached to the device. Hoggan et al. [2007] also suggested a mobile device with multiple high quality piezo actuators to locate feedback at the device and to encode more information in the feedback patterns than just simple affirmations (see Figure 3.12). Yatani

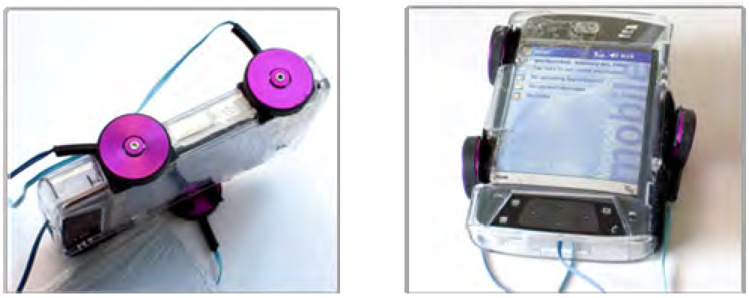

Figure 3.12: Multi-actuator PDA as proposed by Hoggan et al. [2007]

SemFeel.

and Truong [2009] improved on this with SemFeel, a user interface with semantic tactile feedback for mobile touch screens. By attaching five actuators at the back of a PDA, they looked into possible spatial feedback patterns and whether users are able to distinguish them reliably. While these systems already provide promising results the next step is to continue these studies with a device that can provide spatial feedback directly on the screen.

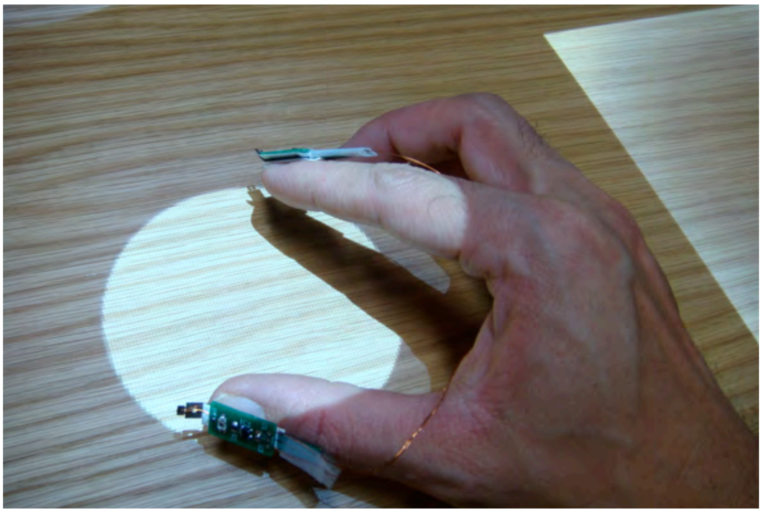

Figure 3.13: Two Senseable Rays actuators attached to forefinger and thumb to feel the size of the projected object (Rekimoto [2009]).

Nail-mounted actuators.

Nail-mounted devices. Instead of arranging actuators in a spatial fashion, they can also be mounted directly on the finger of the user. Then, feedback can be provided according to the sensed position of the finger. SmartFinger by Ando et al. [2002] is a nail-mounted tactile display build around a voice-coil actautor. It depends on a photodetector at fingertip and a fingernail sensor recording pressure changes through color changes underneath the nail to create corresponding bump patterns. Note that it is a self-contained general purpose device that reacts to edges, texture changes and finger pressure data, and is otherwise not bound to a specific application. A similar approach was taken by FingerSight (Galeotti et al. [2008]) although here a camera is used to gain information from the visual environment. Rekimoto [2009] designed with SenseableRays (see Figure 3.13) a similar feedback modality but used time-modulated structured light from a projector to allow each finger to determine its absolute 2-D position within the interaction space. With a second projector the system can even be extended to a 3-D haptic space. While all of these systems enable localized feedback depending on the actual position of interaction, they require users to wear a special hardware devices for each point of tactile feedback and are limited to simple vibrotactile patterns.

Figure 3.14: The liquid haptic output pad by White [1998]. Left: final prototype with liquid filled bladder on top. Right: the underlying array of electromagnets without the bladder.

Haptic displays using magnetorheological fluid

Rheology-based devices. The following systems all depend on either electrorheological or magnetorheological fluid. We will introduce the properties of these fluids in chapter 4.1—"Magnetorheological Fluid". This fluid can be used for both force-feedback as well as tactile feedback and we consider them particularly interesting for use in human computer interaction. The first systems were designed solely as haptic displays (White [1998] used MR fluid while Taylor et al. [1996] used ER fluid) without any input sensing. White's approach (see Figure 3.14) is already very close to the design used for MudPad. Bicchi et al. [2002] and Sgambelluri et al. [2006] took the design a step further and developed haptic displays in which a whole hand could be inserted to explore a 3-D haptic model. But without user tracking these systems were not interactive.

TeslaTouch provides localized feedback on a touch screen.

Electrotactile devices. A different way of creating tactile sensations is by directly applying static current to the skin. The sensation is the same as the static electricity that can be felt with an ungrounded electric device. Still, the sensation is very subtle and it is only perceived upon movement but not for stationary fingers. For TeslaTouch Bau et al. [2010]

(see Figure 3.15) combined a conductive foil with a touch tablet to provide localized feedback to a user's moving finger. Currently, it is only able to present the same signal on the whole screen.

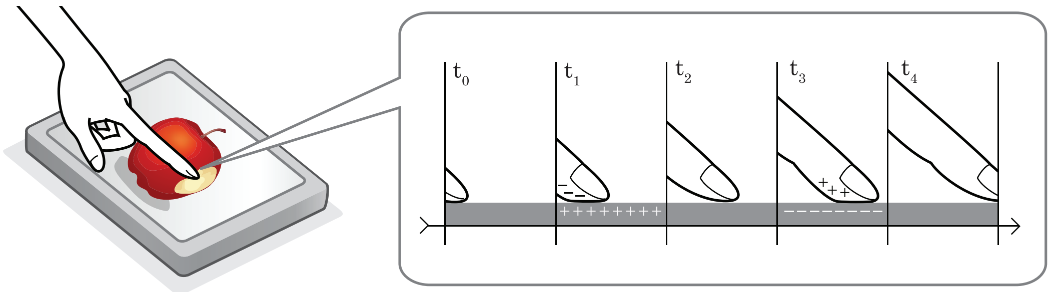

Figure 3.15: TeslaTouch by Bau et al. [2010] uses electrovibration to control electrostatic friction between a touch surface and the user's finger.

Ultrasound devices. The nonlinear phenomenon of ultrasound – acoustic radiation pressure – can be used to produce tactile sensations in thin air (Iwamoto et al. [2008]). Although this also requires 3-D tracking, it is preferable to design a tactile feedback in a system that does not require a user to wear a special device.

Conclusion. This body of related work shows that there is a definite interest in localized tactile feedback for touch interaction. Some of the proposed systems already allow the design of more complex feedback patterns which allow to evaluate semantic qualities of spatially distributed feedback. But all these systems are limited in different ways (cf. Table 3.2). Therefore we are proposing a design that combines multi touch input with localized active haptic feedback simultaneoulsy at arbitrary screen locations. The system allows us to enrich an entire GUI with a haptic layer to allow for more natural interaction by emulating parameters of widgets like buttons that are well-known from everyday life.

	# of actuation points	location of actuation	input(I) and/or output(O)	feedback type	versatility of feedback
mobile systems					
Fukumoto and Sugimura [2001]	1	global	I/O	vibrotactile	limited
Poupyrev et al. [2002] [2003]	1	global	I/O	vibrotactile	piezo
Rekimoto and Schwesig [2006]	1	global	I/O	vibrotactile	piezo
Chang et al. [2002]	4	distributed	I/O	vibrotactile	limited
Hoggan et al. [2007]	3	distributed	I/O	vibrotactile	piezo
Yatani and Truong [2009]	5	back of device	I/O	vibrotactile	piezo
nail-mounted devices					
Ando et al. [2002]	1/nail	localized	O	vibrotactile	limited
Galeotti et al. [2008]	1/nail	localized	I/O	vibrotactile	limited
Rekimoto [2009]	1/nail	localized	I/O	vibrotactile	limited
rheology-based systems					
White [1998]	64	localized	O	viscosity	limited
Taylor et al. [1996]	25	localized	O	viscosity	limited
Bicchi et al. [2002]	whole hand	localized	O	viscosity	immersive
Sgambelluri et al. [2006]	whole hand	localized	O	viscosity	immersive
MudPad	**84**	**localized**	**I/O**	**viscosity**	**high**
electro-tactile device					
Bau et al. [2010]	1	localized	I/O	static electricity	subtle attraction forces

Table 3.2: Comparison of haptic feedback systems

3.3 Performance studies

3.3.1 Evaluation of touch interaction

Touch screens are generally well liked and perceived as intuitive input devices. They take the concept of *direct manipulation* (Shneiderman [1983]) to an even more direct level as they the visual and physical interaction space fall together. In the beginning of touch input research, the performance in terms of speed and error rate was broadly investigated in comparison to mouse and keyboard input. Potter et al. [1988] evaluated different pointing and selecting techniques. The main difference to mouse based input is that for selecting a target it is usually unnecessary to navigate to the object as it can be touched and selected directly when no dragging is involved. Only if the wrong target gets selected, which happens often with small and densely displayed targets, users need a mechanism to correct their choice. Potter et al. therefore suggested different selection techniques such as land-on, first-contact, and take-off. They were able to show that the take-off method lead to significantly less errors than other methods.

Touch screen interaction faster than mouse for large distances but slower for small targets.

Albinsson and Zhai [2003] proposed new methods for high-precision touch input but came to similar conclusions for small targets. Here, the actual performance in terms of speed and error rate is worse than for traditional input devices such as mouse (Sears and Shneiderman [1991]) and keyboard (Sears [1991]). To acquire small targets on a touch screen tactile feedback can be advantageous as it does not suffer from occlusion like visual feedback and especially if it is localized. For bigger targets it is similar or better as long as the touchscreen is sufficiently precise.

3.3.2 Evaluations of tactile touch interaction

As until recently no devices capable of localized tactile feedback were available, there are no studies comparing their effectiveness to regular touch screens. There are a few studies employing work-arounds where several actuators are attached to mobile devices to create spatial patterns (e.g., Leung et al. [2007], Hoggan et al. [2008], Yatani and Truong [2009]). Poupyrev et al. [2002] found that through different vibration patterns information such as a user's scrolling rate and position on the screen can be communicated and that interaction in general was faster with tactile feedback. Leung et al. [2007] evaluated using haptic touch screens under cognitive load and found both performance and subjective benefits. Hoggan et al. [2008] also showed that there is a significant difference between standard touch screens and haptically augmented touch screens in terms of subjective workload. They also found that for their setup there is no significant performance difference between a tactile touch screen and a physical keyboard.

Tactile feedback has advantages in certain situations.

Then there are several studies on advantages of force-feedback – as provided by the SenseAble PHANTOM – especially for navigating menus (Oakley et al. [2001]) or visualizations Pan¨ eels et al. [2009]). Other studies started to look into the semantics that can be encoded by tactile feedback (MacLean and Enriquez [2003] for knobs, Brewster and Brown [2004] on the design of *tactons* (tactile icons), Yatani and Truong [2009] for PDAs with five actuators at the back.

All of these studies point to possible advantages on several layers such as reduced error rates and memory load. Hence, we propose a system that allows more extensive studies in these areas.

Chapter 4

Hardware Design

"The three most dangerous things in the world are a programmer with a soldering iron, a hardware type with a program patch and a user with an idea."

—Rick Cook

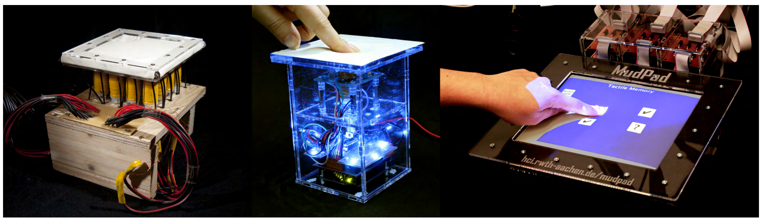

Figure 4.1: The three prototypes.

We built three prototypes.

The hardware development went through several iterations in which prototypes in different sizes for different purposes were built. The first one is a proof-of-concept model with which we investigated the properties of the fluid to evaluate if it is suitable for the intended purpose of providing tactile feedback on a touch screen. The second prototype is a first fully functional one-button travel version to demonstrate the possible feedback characteristics during the ACM Conference on Human Factors in Computing 2010 where we submitted the project for the student research competition (Jansen [2010]). The final prototype is a tablet sized version that was eventually used to evaluate the design.

Before we elaborate on the actual system design, we explain the properties of *magnetorheological fluid* and why we chose this material.

4.1 Magnetorheological Fluid

Smart fluids can be manipulated.

MudPad depends on the unique properties of MR fluid (magnetorheological fluid) – a *Smart Fluid*. The main characteristic of a smart fluid is that one of its properties can be manipulated by applying a magnetic oder electric field. For magnetorheological fluid that property is its viscosity which can be switched from liquid to semisolid by applying a magnetic field. In the semisolid state it then behaves like a *Bingham plastic*[1] meaning that it is viscoplastic but deforms under stress.

Mixture of our MR fluid.

Most industrial magnetorheological fluids are suspensions of coated micron-sized iron particles in a carrier liquid. The carrier fluid is usually based on mineral oil. The coating helps to prevent the settling of the particles. The fluid used for our first prototype was such an industrial fluid supplied by Lord Corporation[2] . While it was sufficient for a proof-of-concept prototype, we looked for chemically more compliant alternatives as we decided for a latex sheet as top cover. Eventually we made our own fluid as none of the commercial available ones met our needs sufficiently. The self-made fluid consists of 80% (by weight) carbonyl iron particles[3] suspended in glycerin which is easily available through pharmacies. We can adjust the viscosity range between *off* and *on state* by varying the mixing ratio of the fluid.

4.1.1 How it works

The main principle is very easy: the iron particles can freely flow in the carrier (see Figure 4.2) and as soon as a magnetic

[1] A well-known Bingham plastic is, e.g., toothpaste.
[2] www.lord.com
[3] The carbonyl iron powder was kindly supplied by BASF.

field is applied the particles align along the magnetic flux lines, form chains, and the fluid stiffens. Figure 4.3 shows how the particles behave when exposed to a homogeneous field – they form lines which account for the increase in viscosity. The base viscosity of the fluid in its *off state* depends on the amount of particles. Up to about 80% (our standard mixing ratio) the viscosity is comparable to olive oil.

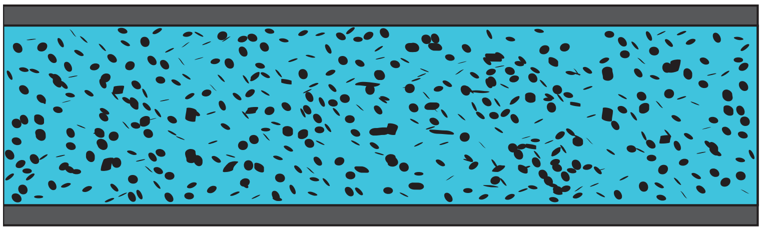

Figure 4.2: MR fluid in its off state. The particles can freely flow within the carrier and the viscosity is accordingly low.

When in *on state*, the viscosity changes to an equivalent of peanut butter. The effect is only local though and reversed as soon as the field is removed. The reaction time, i.e., the time span from on state to off state and vice versa, is in the range of 1–2ms. Accordingly the state can be switched with a rate of up to 800Hz which is well above the temporal acuity threshold for human perception of vibrations (see Table 2.1).

MR fluid shows increased viscosity in magnetic field.

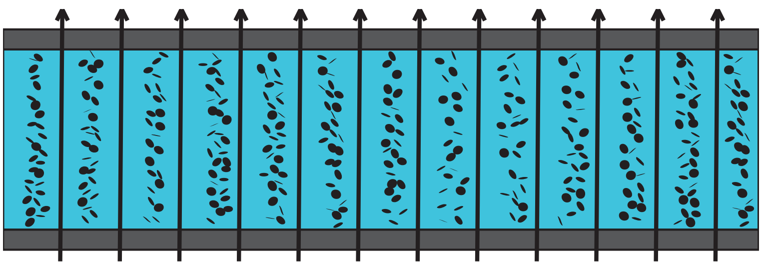

Figure 4.3: MR fluid when exposed to a homogeneous magnetic field. The particles align align along the flux lines and form chains.

As we mentioned earlier the particles used for industrial-strength fluids are coated to prevent settling. We found

Particle settling.

that the settling of particles in our fluid was acceptable as during use no settling occurred. When used daily the settling was so minimal that the particles automatically re-dispersed after short use. Noticeable settling only occurred when the fluid was sitting for at least a few days. But even then the particles easily re-dispersed upon shaking of the fluid container. Hence, for our application the high price for industrial-strength fluid (approx. 800$ per liter) is unjustified.

4.2 First Prototype

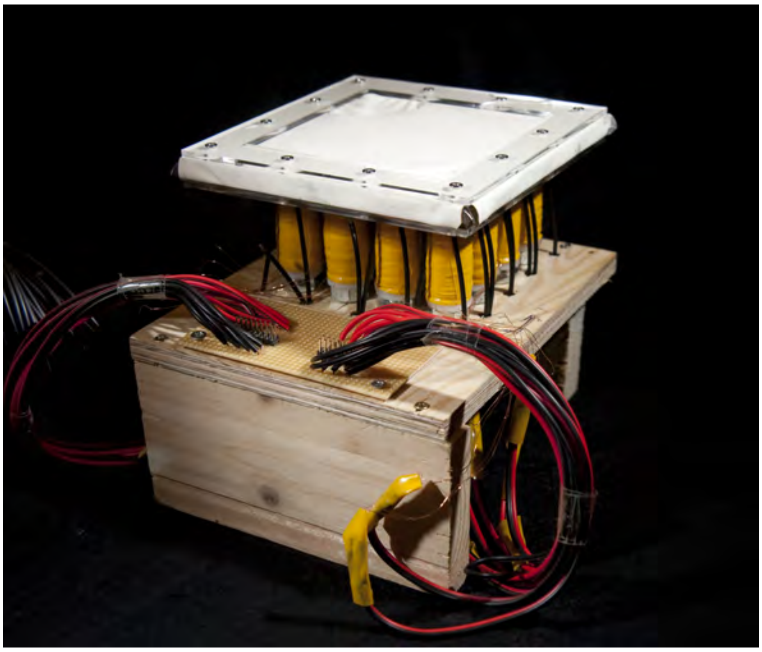

Figure 4.4: The first prototype featuring a 4x4 array of electromagnets. It depends on optical input sensing.

For our first tests, we ordered pre-manufactured magnetorheological fluid from Lord Inc. (MRF-122-2ED). To investigate the behavior of the fluid in a magnetic field, an already existing array of electromagnets from a different research project (Madgets Weiss et al. [2010]) was reused.

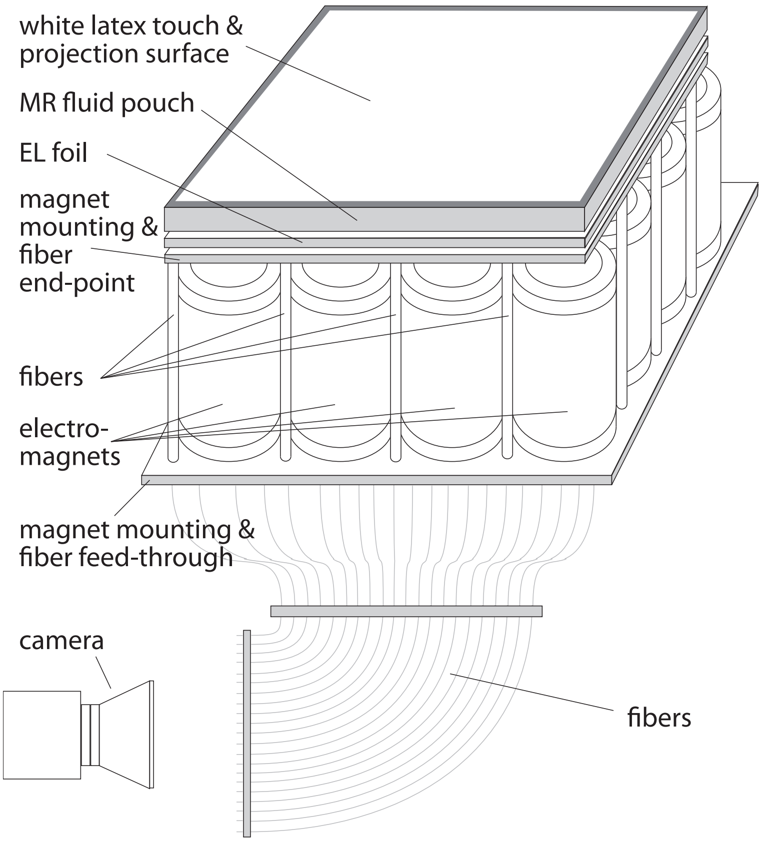

Figure 4.5: Diagram of the first system design.

Figure 4.5 shows a schematic overview of the system design. We used laser-cut acrylic to create a small basin of 7mm depth, filled it with MR fluid and covered it with a sheet of latex. The array of electromagnets is interweaved with an array of optical fibers so that a camera positioned below the magnets can peek through them to detect fingers and markers on the surface (see 3.1—"Diffuse Illumination (DI)"). Due to the clear acrylic bottom and the white latex cover, a white blob is visible on the bottom when a finger presses down on the latex. However, as Figure 4.6 illustrates, optical sensing proved to be unreliable. Mainly a finger's outline can be seen from below and due to the rather low distribution of optical fibers such a setup leads to a very low tracking resolution and many false-negatives.

First prototype as proof-of-concept. Issues with optical tracking.

The electromagnets are switched by two Arduino Mega boards. As one board only has 14 PWM outputs of which

Controlling the electromagnets.

Figure 4.6: A touch from a finger on a MudPad overlay as it is seen from below. The particles directly underneath the finger do not get displaced completely.

two are occupied by the serial USB connection, we used two boards to drive all 16 magnets independently from each other.

4.2.1 Findings

MR fluid suitable for tactile feedback.

This first prototype suggested that the fluid is a suitable material to provide real-time tactile feedback. Its reaction time is well below the human perception threshold so that not only different levels of softness are feasible but also vibrations through the full frequency range of human tactile receptors (see Table 2.2), i.e., frequencies up to 400Hz are no problem and signals can be further varied through manipulation of duty cycles. For exact result extensive psychophysical measurements would be necessary. However, the optical sensing in combination with the small outline of an finger/object placed on the surface proved to be too low in resolution to be useful for touch screen interactions. Hence, we redesigned the system and switched to a different input sensing technique for the second prototype.

Optical sensing not suitable with MR fluid.

4.3 Second Prototype

Figure 4.7: The second prototype is a fully functional travel version and simulates one button to demonstrate the possible feedback modes.

By taking the findings from the first prototype into account, we considered different sensing techniques that might be suitable for input tracking with a MudPad overlay.

4.3.1 Discussion of Available Input Techniques

As we illustrated earlier (3.1—"Touch Technologies") there are many available techniques for multi touch input. Unfortunately, most of them are unsuitable in conjunction with a MudPad overlay.

No optical sensing due to low resolution.

The first prototype showed that a diffuse illumination approach in combination with light reflected by the latex cover is not a reliable sensing technique due to the low number of fibers embedded in the array of electromagnets. Hence, all optical approaches that depend on a line of sight from below the touch area can be dismissed.

No capacitive sensing due to the fluid.

The easiest way to achieve multi touch sensing for a tablet sized area is a capacitive touch screen as used for the Apple iPad. But a MudPad overlay interferes with the capacitive change induced by a finger touching the surface. A quick test with an iPad showed that even for a thin version ($\approx$3mm) placed on the surface, no touches were detected.

Pressure/resistive sensing combines several interesting properties.

We eventually decided on pressure based sensing as it is robust against magnetic fields and, if constructed as proposed by Rosenberg and Perlin [2009], capable of multi touch sensing. Furthermore, due to it being pressure sensitive, it allows to explore the haptic representation of an interface by lightly resting a hand on the surface. Only touches above a certain pressure threshold are interpreted as such. Unfortunately, it is currently almost impossible to acquire such a screen[4].

4.3.2 System Design.

Force-resistive sensors work best.

The main purpose of this prototype is to demonstrate what kind of tactile feedback is possible. We therefore developed a one-magnet version with a 7x7cm MudPad overlay (see Figure 4.8). The acrylic sheet at the bottom of the fluid pouch which was used for the first prototype is now replaced by a flexible foil such as to allow pressure from above to be transferred to the underlying pressure sensors. Six force sensitive resistors act as pressure sensors and register user input to trigger different feedback patterns. An Arduino board polls these sensors and switches the electromagnet accordingly. The six sensors are mapped to the following six functions:

[4]There is now an open hardware project that provides schematics to build a simplified version. See www.sensibleui.com for details.

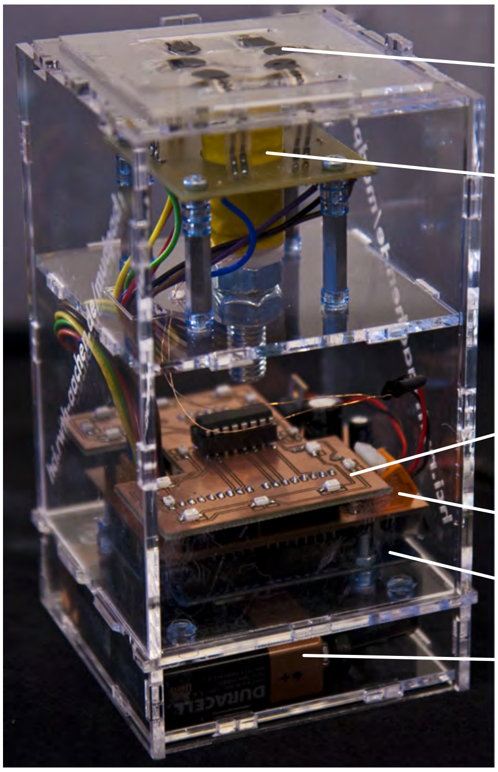

Figure 4.8: Prototype 2 explained without MudPad overlay.

1. magnet off
2. magnet on
3. button function, i.e., switch to off when pressed and turn on again upon release (cf. Table 5.1)
4. slow regular vibration (15Hz)
5. irregular vibration (5-80Hz)
6. fast vibration (200Hz)

4.3.3 Lessons learned

This prototype confirmed that the fluid is indeed suitable to create both force-feedback as well as vibrotactile feedback. The small pouch size also made it easier to evaluate different fluid compositions. We found that fluids with more than 80% iron particles show a noticeable increase in their off state viscosity whereas fluids with less than 60% particles show a not acceptable decrease in their *on state* viscosity.

4.4 Third Prototype

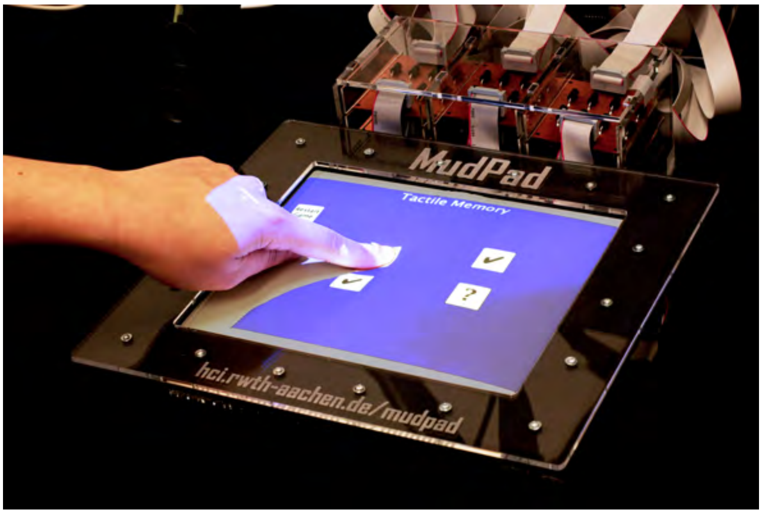

Figure 4.9: The third prototype is a 10″ version incorporating a 12x7 array of electromagnets. It is controlled by an Arduino and uses a projector to top-project the GUI.

Final prototype to evaluate the design.

The third prototype is a fully functional version of the proposed design. It was used for a quantitative study (6.1—"Quantitative Study") and exhibited several times on different venues (6.2—"Exhibitions") to gather qualitative feedback by possible users.

As Figure 4.10 shows, we built the system on an array of 84 electromagnets which can be addressed individually. Hence, we are able to create localized magnetic fields that actuate the fluid only in the direct vicinity of activated magnets. Most feedback patterns cannot be felt by a second finger placed beside the active finger. Only higher frequency signals tend to spread a little so that they can be felt by an adjacent finger. For touch input we eventually decided to employ a 4-wire single touch screen as resistive multi touch

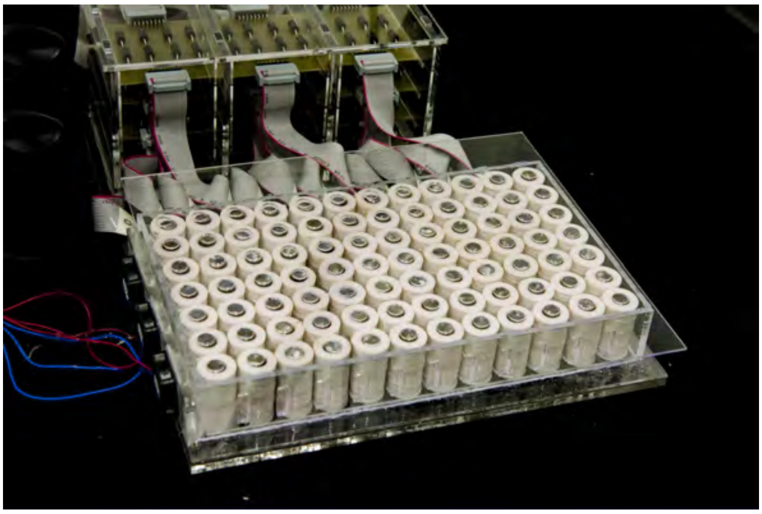

Figure 4.10: The 12x7 array of electromagnets that actuate the fluid.

screens were almost impossible to acquire at the time. As the screen can be easily exchanged and is not part of the proposed design, it sufficed to conduct the planned studies to evaluate the system.

4.4.1 Output Resolution

The magnets are spaced 2cm from each other. Accordingly, we can address 84 different locations on a 2x2cm grid at the same time. Despite the sparsely distributed sources of magnetic fields, it is still possible to create haptic feedback at arbitrary positions by superimposing the fields of several adjacent magnets. Such a group of magnets represents the *actuation domain* of a certain location. For a single finger in such a domain the output resolution is still determined by the input resolution as the feedback given is created according to the input position and changes even when the finger is moved within the actuation domain. However, for several inputs within one actuation domain the individual feedback patterns for these inputs would be mixed.

Upt to 84 concurrent feedback locations.

Chapter 5

Application Scenarios

In this chapter, we describe possible scenarios in which localized tactile feedback can be useful. We start by describing the elementary building blocks out of which we then design more complex feedback patterns. We present a few examples for patterns designed for specific widgets that are presently common on mobile touch screen devices. Finally, we outline some concrete examples for designs of tactile feedback for specific applications.

5.1 Elementary Building Blocks

The elementary blocks listed in Table 5.1 were identified after constructing the second prototype. Informal tests showed that, although the viscosity of the fluid can be manipulated almost in a linear fashion, it is hard for users to distinguish different levels. As one of our objectives is reduced memory load, we decided to exclude this property for our feedback design. Accordingly, the fluid can only be in two different states: fluid vs. stiff.

Possible feedback patterns can be broken down in elementary building blocks.

Interestingly, the transition from stiff to fluid can be perceived distinctly (building block 2). Depending on the fluid composition and the field strength of the magnetic field a force-feedback effect can be produced. By switching the fluid from stiff to fluid as soon as a touch event is recog-

Button feedback.

Magnet Signal	Fluid State	UI Mapping (System View)	UI Mapping (User View)
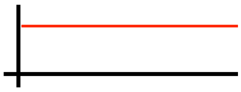	stiff	inactive areas, i.e., no user input possible	prevent interaction
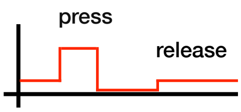	quick on/off transition	active UI elements, e.g., buttons	acknowledge user input
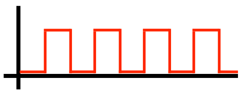	(rapidly) changing, vibrating	active areas, demanding user attention	communicate system processes, e.g., progress bar
	fluid	active areas allow interaction	neutral

Table 5.1: Elementary building block from which feedback patterns can be constructed.

nized, the fluid surrounding the finger gives way and the finger sinks in further. This effect is most suitable to simulate a button press and will be used whenever a button is enriched with tactile feedback. Fast switching of the fluid is interpreted as vibration. The fast reaction time of the fluid allows a fine grained manipulation of the waveforms used to control it.

5.2 Widget Mappings

Possible mappings for standard widgets.

Starting from these elementary building blocks, we designed possible tactile feedback patterns for a few widgets that are part of the standard widget set for Mac OS X or iOS (see Figure 5.1 (Jansen et al. [2010a])). This set could easily be extended to create mappings for all available widgets.

Although a tactile mapping for all widgets within a given set is interesting to investigate possible advantages for standard GUIs, we rather focus on more specific mappings in an application context. Hence, we did not implement these widget mapping designs. Depending on the context different mappings for the same visual representation, i.e., widget, might be sensible and analogously in a certain context the same feedback pattern could be used for different widgets.

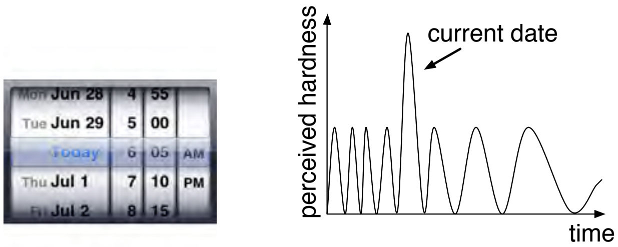

(a) Date picker widget

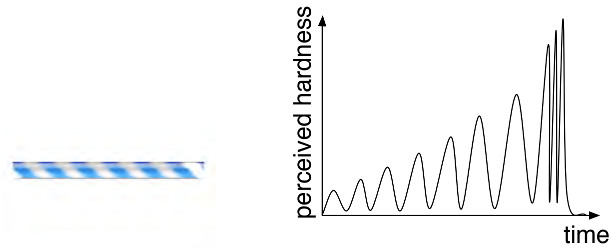

(b) Progress bar widget

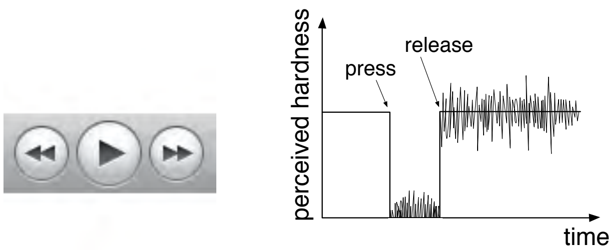

(c) Playback control buttons for a music player application

Figure 5.1: Example widgets and their tactile feedback mappings.

5.3 Possible Applications

Before we present concrete usage scenarios, we want to make a distinction concerning the classification of possible feedback mappings. For MudPad multiple senses are involved during interaction. Accordingly, feedback can be designed in a **multimodal** or **crossmodal** fashion.

Definition: *Crossmodal vs. multimodal feedback*

CROSSMODAL VS. MULTIMODAL FEEDBACK:
Crossmodal feedback presents the same information to different senses, while multimodal feedback splits information to present each part to different modalities. Intersensory integration is usually easier for crossmodal presentation as no conflicts between modalities are possible.

We now introduce a few concrete examples and suggest possible mappings for each one (see Table 5.2 for a summary).

	substitute	**additional**
crossmodal	keyboard	keyboard, graph
multimodal	sequencer, secure keypad	background, graph, error prevention

Table 5.2: Categorization of example applications depending on whether information is substituted or added, and whether feedback is designed cross- or multimodal.

5.3.1 Virtual Keyboard

Virtual keyboard with tactile feedback on the buttons.

Enriching a virtual keyboard with tactile feedback is probably one of the most intuitive applications for localized tactile feedback. Figure 5.2 shows an implementation for the MudPad memory game demo (Jansen et al. [2010c]). Upon finishing the game a player can add her name to the high score by typing on this virtual keyboard. The feedback mapping used here is the force-feedback perceivable when switching the fluid from *on* to *off*. Consequently, the

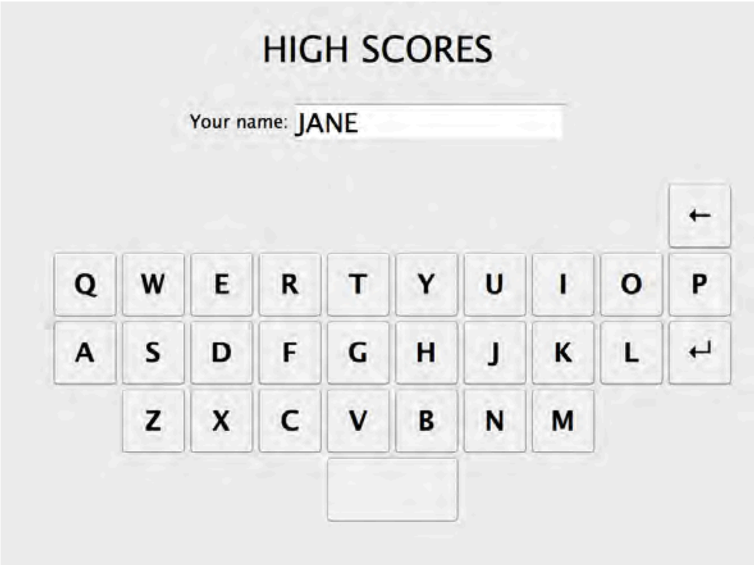

Figure 5.2: Screenshot of a virtual keyboard for MudPad as used within the Memory Game demo to enter high scores (Jansen et al. [2010c]).

keyboard area needs to be *activated* when it gets displayed, i.e., the magnets must be turned on. This kind of feedback is crossmodal as the tactile feedback supports the visually presented information such as a shadowing effect when pressing a key and the appearing letters for each keystroke.

Error prevention by stiffening unlikely buttons.

Error feedback. When the text entry is coupled with a dictionary it is possible to adjust the feedback in a way that it communicates to a user if the current entry is an item in that dictionary (cf. Hoffmann et al. [2009]). The feedback for a letter combination that cannot be completed in a word of the English language could be switched to a buzzing sensation to get a users attention, or the force-feedback could simply be deactivated, i.e., the fluid does not turn to fluid for wrong entries. Note that this kind of feedback does not hinder a user to enter words which are not part of her dictionary but simply draws her attention. However, here the feedback is used in a multimodal way as the tactile information is not (necessarily) presented visually but additional.

5.3.2 Graph Exploration

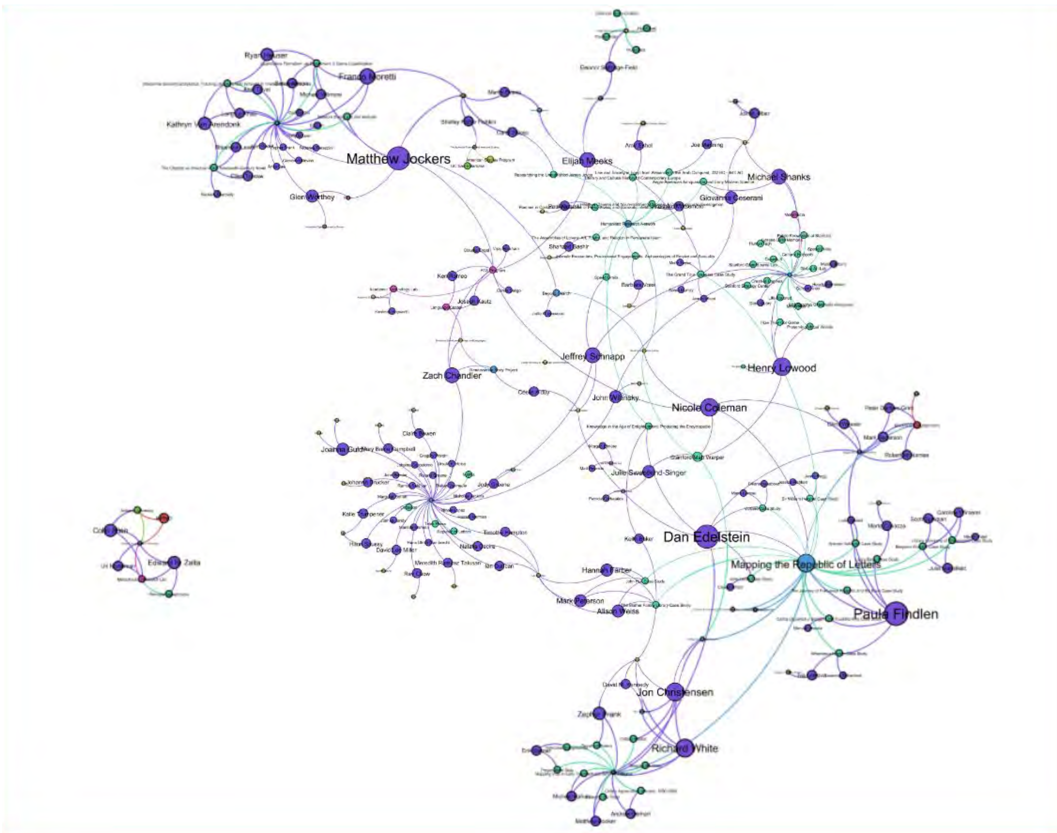

Figure 5.3: Example for a graph visualization that can be navigated with the help of MudPad. Source: Stanford Digital Humanities Network Map[a]

[a] dhs.stanford.edu

Exploring dense graphs by touch.

A graph network as depicted above (Figure 5.3) can be navigated with MudPad both as stand-alone or as additional control. Because of the usual size and density of such visualizations the usage as an additional touch controller would be more sensible (judging by the size of the current prototype). To get an overview over the distribution of nodes one could imagine areas pulsing according the local density of information. Then, when a node gets selected, all its direct neighbor nodes could be actuated to make it easier to locate nodes of interest eyes-free.

5.3.3 Music Sequencer

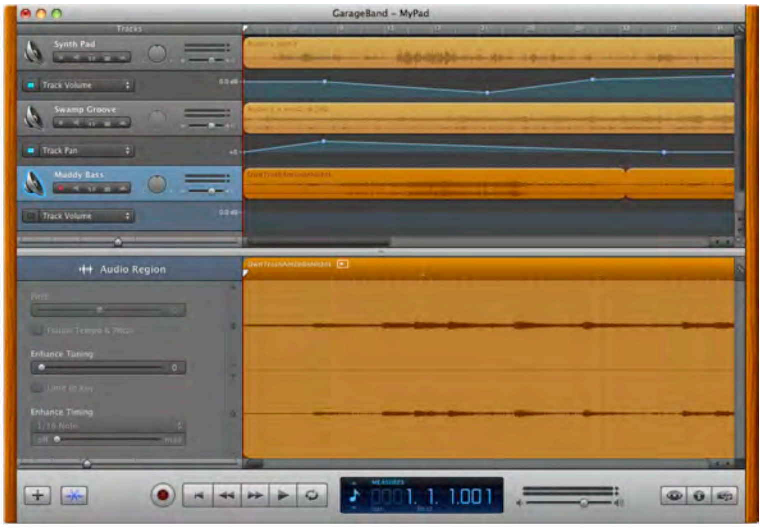

Figure 5.4: Screenshot of the GarageBand sequencer program that comes with each Mac.

Tactile representation of music tracks.

A music composition application such as depicted in Figure 5.4 can be augmented with tactile feedback in way so that each instrument track gives a tactile representation of the audio file it contains. The reaction time of the MR fluid is fast enough to use audio waveforms as base for the feedback pattern. Note though that the tactile perception of a finger tip is much more limited than the ear's. Consequently the audio waveform needs to be remapped to the most receptive range of 0–400Hz. Then, an audio track can be felt when it is touched regardless of it being currently muted. Such a feature can be useful in a beat matching task as is common for DJs. To preserve the metaphor of direct manipulation this would be most useful when MudPad is used as a stand-alone device where interaction happens in-place.

5.3.4 Secure Keypad

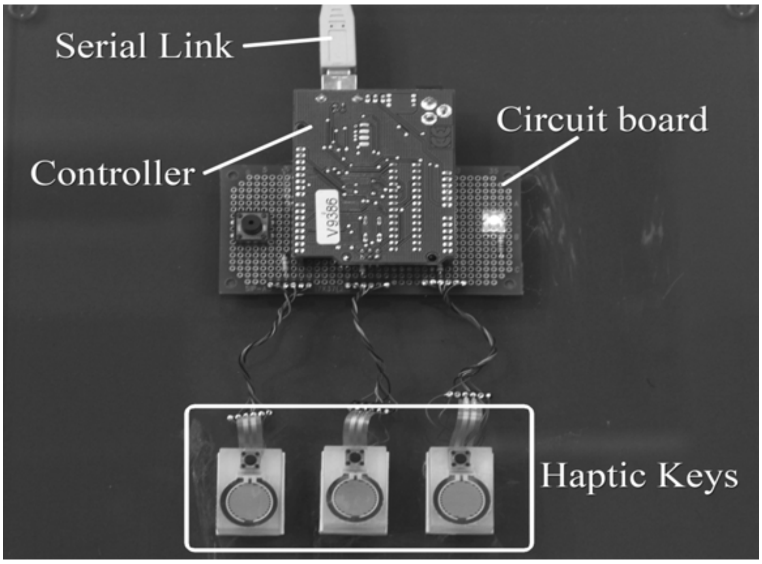

Figure 5.5: The secure keypad as proposed by Bianchi et al. [2010]. It can easily be implemented with MudPad.

Enabling touch screens to enter sensitive information in a secure way.

Touch screen interaction is rather public and prone to shoulder surfing. Therefore users have to be very careful to shield their hand while entering a password, e.g., at an ATM. Due to the private nature of tactile interaction – it is invisible to bystanders – it is most suited to communicate secret information (cf. Bianchi et al. [2010]). One can imagine receiving tactile hints to insert fake letters or numbers when typing a password on a touch screen. That way shoulder surfers cannot spy out passwords or to be more precise, they would need a much larger sample from the same person to extract the real numbers. This can be implemented for both a stand-alone or a separate control device.

5.3.5 Background Processes

Tactile bump instead of system beeps.

Tactile feedback can also be used to communicate information about background processes to users. As long as touch input is recognized by the system–as a measure that a user will actually receive the feedback–information such as incoming mail, finished downloads, or completed computations can be indicated by subtle tactile signals instead of undirected audio or distracting visual feedback. Of course in case of user absence the system needs to switch to a different modality if it is mandatory that the user receives the message.

Chapter 6

Evaluation

"Everything is vague to a degree you do not realize till you have tried to make it precise."

—Bertrand Russell

The interaction with MudPad differs from regular touch screen interaction due to the soft and malleable input surface. The quantitative study we present in this chapter was designed to classify MudPad in regard to its performance as an input device. Accordingly, we compare it to a regular glass surface touch screen and a specialized input pad. After presenting the study and our results in detail, we give an overview over informal user feedback we gained from various demonstrations of our third prototype.

6.1 Quantitative Study

Objectives for this study.

This study was a pilot study to explore the effects of the choice of input device and different feedback combinations on typing speed, error rate and correction attempts. Our intention is not an accurate classification of possible input performance with MudPad. As it is the first system capable of localized active multi-point feedback, we only want to ensure to that there is no significant performance loss with

MudPad to suggest it as a means to investigate the effects of such a feedback device.

6.1.1 Study Design

We designed this study to investigate the following hypotheses:

1. The MudPad overlay has no significant influence on the typing speed.
2. Tactile feedback provided by MudPad leads to fewer errors.
3. Less corrections are necessary when tactile feedback is provided.

Task: enter numbers fast and accurate.

The task. Each user was asked to enter numbers as fast and as accurate as possible. They could correct their entries when they detected an error. There were versions for left- and right handed users with the same distance between the displayed number they had to enter and the virtual keypad to type. The distance was chosen so that the keypad had to be operated in an eyes-free fashion as keypad and displayed number were too wide apart to be kept in visual focus at the same time. All participants were asked to only use their index finger for typing. Each number consisted of six digits and was presented in groups of two (see Figure 6.1). For each condition 30 numbers were presented accumulating to a total of 270 numbers entered by each user. The numbers were randomly distributed between 100000 and 999999 with no consecutive occurrences of the same digits.

The variables. We controlled two independent variables

- IV1: input device used to enter the numbers with three different values:

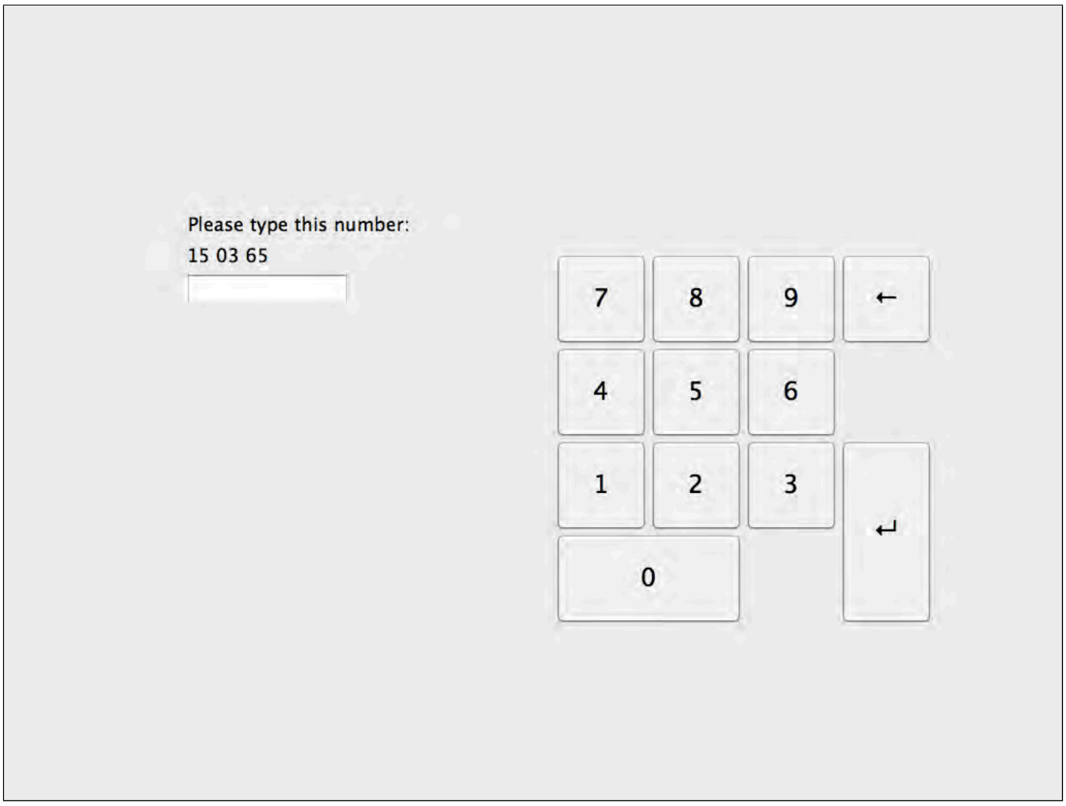

Figure 6.1: Screenshot of the GUI presented to a right-handed user.

  - physical number keypad,
  - MudPad
  - glass-surfaced touch screen

- IV2: feedback type with four different parameter values:
  - none
  - tactile
  - auditive
  - combined (audiotactile)

and measured three dependent variables

- DV1: time necessary to enter a complete number
- DV2: error rate, i.e., amount of mistyped numbers
- DV3: corrections, i.e., how often was backspace used.

All combinations of IV1 and IV2 result in nine conditions.

The conditions. By creating all possible combinations of IV1 and IV2 we get nine different conditions to be tested. The physical keypad provides auditive and tactile feedback by itself, therefore only one condition is created for it. Each user had to perform the task for the following nine different combinations of IV1 and IV2:

1. **Physical** standard number keypad (6.2).
2. MudPad with GUI keypad **without** any kind of additional feedback (baseline condition).
3. MudPad with GUI keypad and **audio** feedback.
4. MudPad with GUI keypad and **tactile** feedback.
5. MudPad with GUI keypad and both, **audio and tactile** feedback.
6. Glass touch screen with GUI keypad **without** any kind of additional feedback (alternate baseline condition).
7. Glass touch screen with GUI keypad and **audio** feedback.
8. Glass touch screen with GUI keypad and **tactile** feedback.
9. Glass touch screen with GUI keypad and both, **audio and tactile** feedback.

With nine conditions, we decided to test 18 users and assigned each one to a group by using a latin square.

Nail-mounted magnet to provide tactile feedback on a standard touch screen.

Feedback design. As the goal of this study is to compare the performance of the system with the MudPad overlay to the use of a standard glass surfaced touch screen, we needed a way to give tactile feedback on a glass touch screen. We decided to use a small permanent magnet attached to the fingernail of the forefinger (cf. Ando et al. [2002], see Figure 6.3) to produce a somewhat comparable sensation when pressing a button on the GUI keypad. The

Figure 6.2: Setup for the physical condition.

Figure 6.3: Nail-mounted permanent magnet for tactile feedback without the MudPad overlay.

array of electromagnets under the touch screen was used to produce a small attractive force upon a button press.

Button feedback is used for MudPad

The tactile feedback with the MudPad overlay makes use of the fluid's ability to change its viscosity. Per default (in the tactile condition), all buttons are stiffened and upon a button down event the fluid gets liquid again thereby producing a small *pop* effect.

The audio feedback uses standard Mac OS X system sounds. Each number button is mapped to the *Tink* sound, the backspace button to *Morse*, and the Enter button to the *Pop* sound.

Participants. The participants for this study were students recruited through a posting in a mailing list. All of them are computer science students aged 21–34. They all participated voluntarily and were not compensated for their time. The study took about 30min to complete.

6.1.2 Results

The data gained from this study was analyzed using a general linear model for repeated measures. To compute average keystroke times only correctly entered numbers were considered, i.e., the dataset was cleaned for this part of the analysis.

Exclusion of extreme outliers.

Of our 18 participants we had to exclude the data from one user as he/she was an extreme outlier who put up with an extreme number of errors for the sake of a slight increase in speed. While this hardly influenced the means for typing speed, it had a huge influence on the results for corrections and errors as his/her results were three times higher than the average value.

Main effect for input device. To compute the main effect of an input device, all individual measurements made with that device are averaged irrespective the feedback type provided. For the number pad it is not possible to manipulate the different feedback components, i.e., it always provides tactile feedback and the mechanical movement of the keys is audible. Thus, the *combined feedback* condition of MudPad and the touch screen could be seen as the "corresponding" conditions for analysis. Therefore, we decided to evaluate the data gained from the number pad twice – first in contrast to the respective *combined feedback* conditions for the touch screen and MudPad and then again to the combined means of all feedback conditions of MudPad or the touch screen. In the next paragraph, *combined condition* always refers to either condition 5 or 9 defined on page 56 whereas *combined means* refers to averaged values over conditions 2–5 and 6–9 respectively.

No significant main effect between all three input devices.

In direct comparison to the *combined condition* the data does not show a significant effect between the three input devices for either average keystroke time, errors or corrections. However, when the combined means of the other input devices are used for evaluation, the results for errors and corrections are around p=0.5 (see Table 6.1). As

	Numpad vs *combined condition*		Numpad vs average over all conditions	
Measure	F(2,15)	p	F(2,15)	p
avg keystroke	.75	0.489	2.058	.162
errors	1.665	.222	3.811	.046
corrections	1.421	.272	3.352	.063

Table 6.1: F values and significance level for main effect of input device on average keystroke time, errors, and corrections.

Keystrokes with MudPad take 20ms longer than with a number pad.

Figure 6.4 illustrates, the average keystroke times hardly differ. But a pairwise comparison[1] of the input devices shows an increase of about 20ms per keystroke for MudPad when compared to the number pad (p=0.056, see Table A.4). However, there is no significant effect between MudPad and the touch screen for average keystroke times (see next paragraph for a detailed comparison between MudPad and touch screen).

Less errors and more corrections with MudPad.

For the other two dependent variables we found that when compared with the *combined condition* the effect is not significant, whereas it is for *number of errors* when compared to the *combined means* (see Table 6.1). A pairwise comparison (see Table A.5) shows that with MudPad the error rate per 30 entered numbers is in average 0.4 lower than with the touch screen (p=0.012) while one more correction is made (p=0.019, see Table A.6). The different results are illustrated

[1]Complete tables for all pairwise comparisons can be found in appendix A—“Statistical Data”.

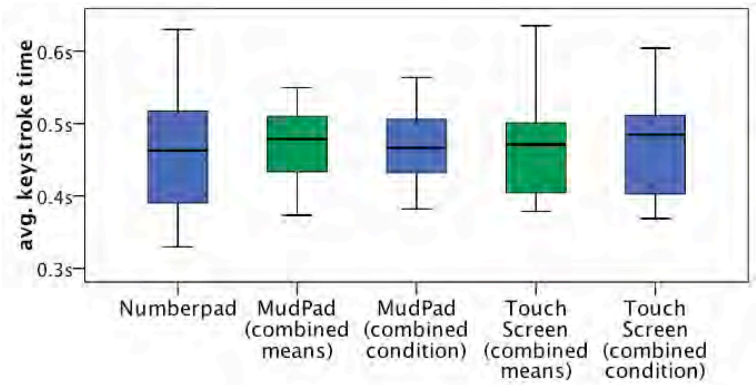

Figure 6.4: Average keystroke times for the three input devices tested. We included two boxes for MudPad and the touch screen to compare the results for the *combined condition* with those averaged over all feedback conditions.

by Figure 6.5 for errors and by Figure 6.6 for corrections.

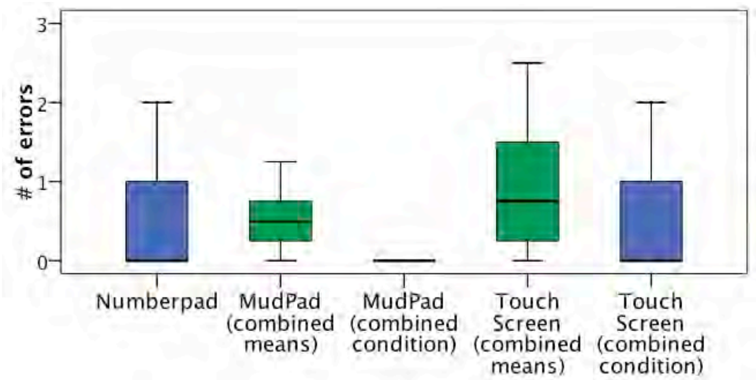

Figure 6.5: Number of errors for all three input devices with separate boxes for MudPad and touch screen for *combined condition* and averaged values over all feedback conditions.

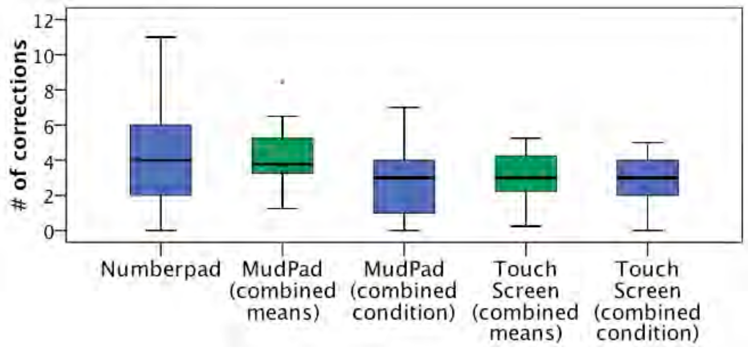

Figure 6.6: Number of corrections made for all three input devices with separate boxes for MudPad and touch screen for *combined condition* and averaged values over all feedback conditions.

Effects for different feedback types. We now present our results for the combined effects of different feedback conditions for MudPad and the touch screen. As the feedback for the number pad cannot be controlled individually, we compare from now on only MudPad with a glass surfaced touch screen.

Comparing various feedback types

Table 6.2 compiles the results of the multivariate tests. The differences for the average keystroke time are highly significant (p=0.007), and again a pairwise comparison reveals interesting details. As Figure 6.7 illustrates, *audio feedback* led to shortest keystroke times while the *tactile feedback* condition was the slowest. A pairwise comparison shows no significant effect between *tactile* and *no feedback* (p=0.149)–an unexpected result–and also no significant difference between *no* and *combined feedback* (p=0.817)–an even more unexpected result. As *audio feedback* accounts for a significant advantage against *no feedback*, we expected to see this advantage also for the *combined* condition.

Tactile condition slowest.

The effect of both input device and feedback type on the number of errors made is almost highly significant (p=0.012 for both) and shows that with MudPad more errors are detected and corrected (see Figures 6.8 and 6.9).

The complete analysis can be found in appendix A—"Statistical Data".

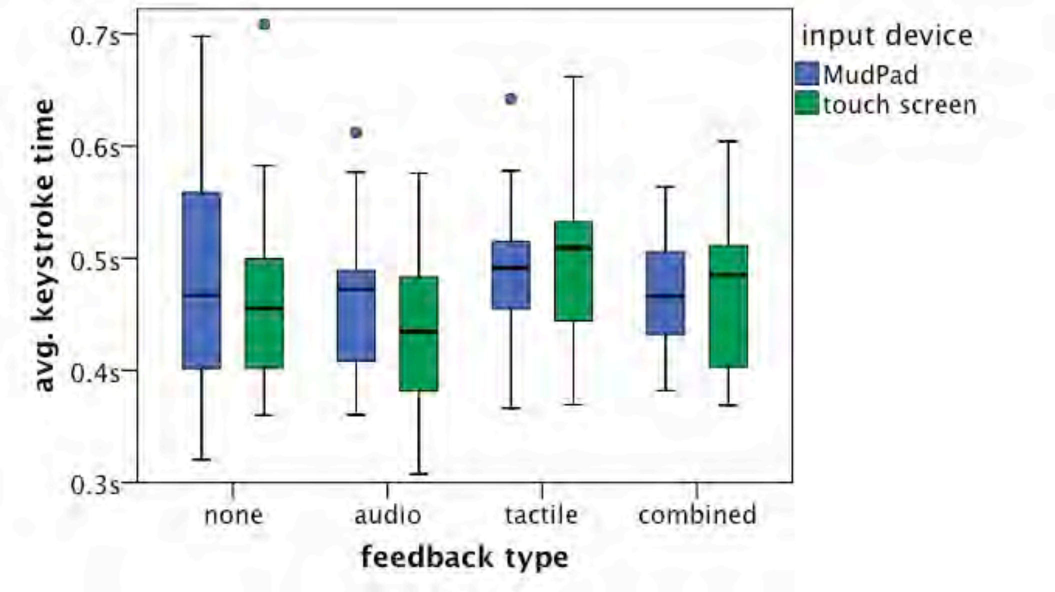

Figure 6.7: Average keystroke times for MudPad and touch screen grouped per feedback condition.

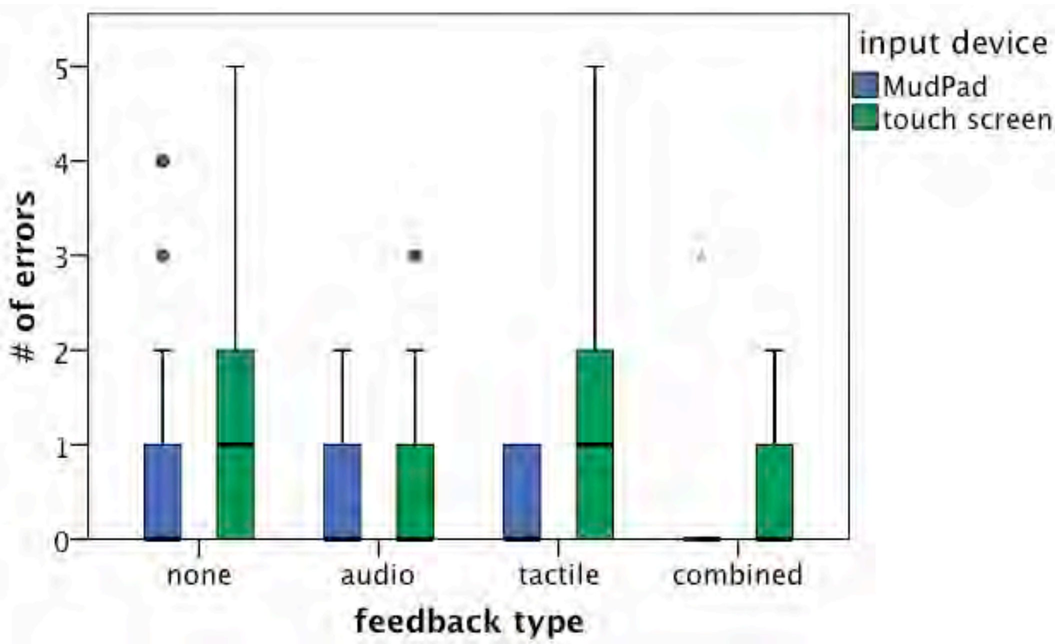

Figure 6.8: Number of errors for MudPad and touch screen grouped per feedback condition.

	keystroke		errors		corrections	
Effect	F	p	F	p	F	p
input device	F(1,16)=1.2	.285	F(1,16)=8.1	.012	F(1,16)=6.8	.019
feedback type	F(3,14)=4.6	.007	F(3,14)=4.6	.007	F(3,14)=9.4	.000
input device ∗ feedback type	F(3,14)=1.4	.255	F(3,14)=1.8	.161	F(3,14)=2.7	.058

Table 6.2: F values and significance level for effects of input device, feedback type, and cross effects on average keystroke time, errors, and corrections.

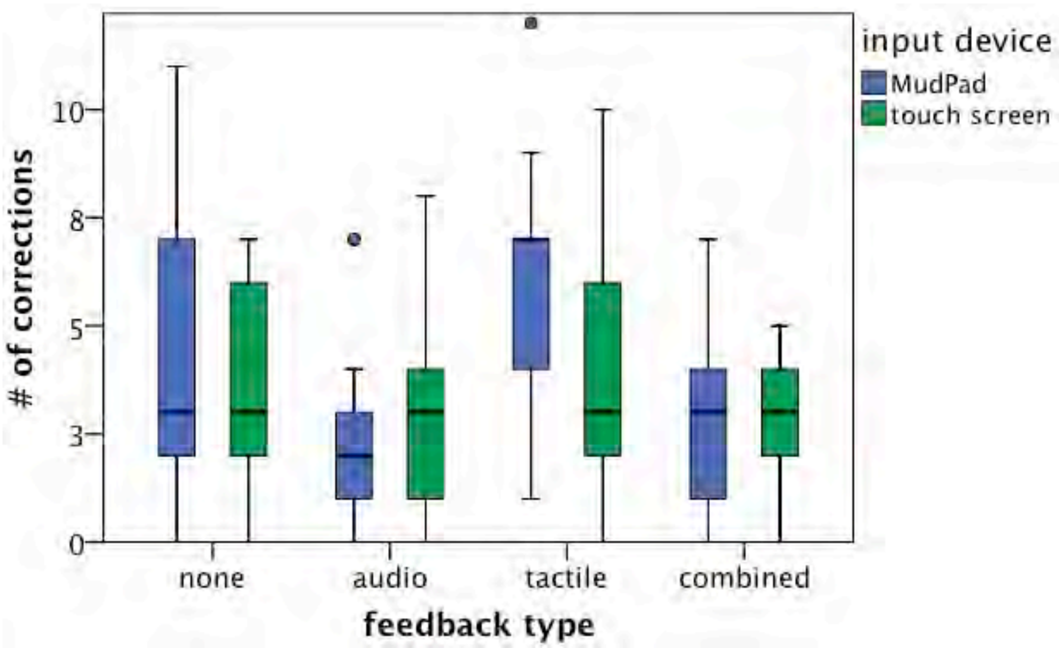

Figure 6.9: Number of corrections for MudPad and touch screen grouped per feedback condition.

6.1.3 Discussion

Our data supports the first hypothesis. We did not find a significant effect of the choice of input device on the average keystroke time. In a direct comparison, typing with MudPad is slightly slower (∼5%) than with a number pad. But as number pads provide numerous haptic clues during use, e.g., keys can easily be distinguished by touch, we expected the number pad to perform better. The main pur-

First hypothesis supported by our data.

pose of this test was to establish if MudPad has a negative influence on user performance in comparison to a standard touch screen. We did not find a significant difference. Our results support the first hypothesis, but as this test was only a small pilot study additional tests would be necessary to further validate this result.

Data suggests that errors are detected more often when using MudPad.

The second hypothesis on less errors with MudPad is also supported while the third one is contradicted. MudPad does not prevent users from making errors–the tactile feedback is only provided *in response to* a user action. But the higher number of corrections suggests that these errors get detected more often when using MudPad. An additional benefit could be gained by implementing an *error prevention* mechanism as proposed by Hoffmann et al. [2009]. As MudPad allows to stiffen the surface, a dictionary based error prevention is possible. Hoffmann et al. proved such a mechanism useful for their haptic keyboard. A future study with MudPad could repeat their study.

In conclusion we can say that MudPad exhibits no improvements for input speed but an increase in accuracy due to easier detection of errors.

6.2 Exhibitions

We had the opportunity to demonstrate our third prototype at several occasions. In direct interaction with visitors coming by to try it, we gained some insights into what most of them liked about it and possible issues to improve. For these presentations we developed a *tactile memory game*. It is based on the common *memory* game where pairs from a set of cards spread out on a table are successively turned to find matching pairs. For our game we presented users with an initial set of four buttons (see Figure 6.10 for a screenshot and Figure 4.9 for a running game, two buttons are already matched and display a tick mark as a visual reminder). Each button is associated with a specific tactile

Tactile memory game demonstrates tactile feedback patterns.

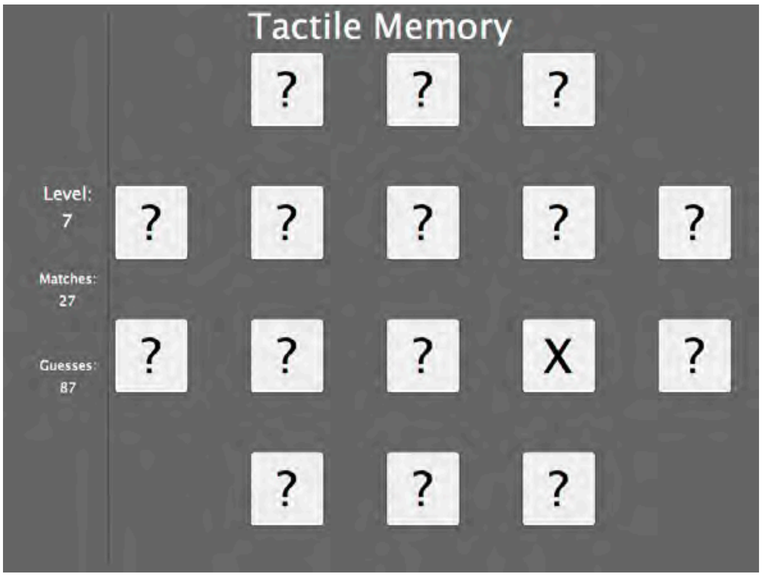

Figure 6.10: Screenshot of the *tactile memory game* with currently 16 buttons displayed. The button marked with an X is currently selected and the game now expects a user to press the button with the matching feedback pattern.

feedback pattern when pressed. Within each set users have to find the pairs of buttons with the same pattern. When all pairs are found the next level is reached and the number of buttons is increased up to 20 buttons.

6.2.1 Ars Electronica Linz September 2010

The *Ars Electronica Festival* is an annual event dedicated to electronic art. With more than 90.000 visitors in total and alone 15.000 during the open house day it is one of the most important festivals for digital art world-wide.

MudPad was on display for three days including the open house day. Especially during the open house day, we attracted a very diverse group of visitors. We had many children and their parents or grand-parents stopping by to play the game. Even though the game is visually not very attracting–apart from the buttons nothing is displayed–is was well liked by all children and some spent up to 20min playing it. The most common remark was that MudPad has a very pleasant touch sensation. The soft surface is unfamiliar for a touch screen but was well accepted.

6.2.2 ACM Symposium on User Interfaces and Software Technology New York October 2010

MudPad was then accepted as a demo at UIST'10 (Jansen et al. [2010a]). It was presented for four hours during the demo reception. We added a high score to our game (see Figure 5.2) to demonstrate a keyboard implementation where the *button feedback* (cf. Table 5.1) was used. As UIST is a conference on user interfaces most visitors at our booth were experts. Again the overall feedback was very positive. Most testers enjoyed touching the surface and explored the different feedback patterns. The keyboard draw special attention, although many testers wished for a stronger feedback. In general we noticed that the tactile sensitivity varies hugely among users. Some attendees reported being unable to detect anything.

6.2.3 ACM Conference on Interactive Tabletops and Surfaces Saarbrücken November 2010

We submitted a tech note on MudPad to the ITS'10 conference (Jansen et al. [2010b]) and also presented it there as an accompanying demo (Jansen et al. [2010c]). The system was again well received and won both the best note and best demo award. Attendees showed great interest in localized tactile feedback for touch screens. They appreciated MudPad as a first step to be able to explore this field. However, to take full advantage of the variable viscosity they asked for a broader range, i.e., the fluid should be more rigid in its *on state*. Further experiments with fluid compositions are necessary to fully explore the possible range.

Chapter 7

Summary and future work

"The beginning of knowledge is the discovery of something we do not understand."

—Frank Herbert

7.1 Summary and contributions

Touch screens are widespread as of today. It can be expected that they become even more ubiquitous not only for mobile devices but also for all kinds of everyday appliances. While they are very popular and well liked, they lack tactile feedback for user interactions. Their screens are smooth glass surfaces and do not exhibit any tactile clues whether user input was recognized.

Touch screens lack tactile feedback.

In this thesis we presented MudPad, a touch interface with localized active haptic feedback. Each screen area of MudPad can be controlled individually and simultaneously to provide a specific tactile feedback pattern. As of today MudPad is the only system capable of localized active haptic feedback on a touch screen.

MudPad provides localized active haptic feedback.

We developed the system design over several iterations and

User study suggests positive effects.

eventually evaluated the system formally. Our user study suggests that the soft latex cover has no significant negative influence on user performance (measured as average keystroke time). By interacting with MudPad users are less likely to leave errors undetected. Furthermore, the soft touch surface was well received and often preferred over a glass surface.

We implemented a *tactile memory game* that makes use of this feature and demonstrates the range of possible feedback characteristics. Additionally we suggested several applications that would benefit from this kind of feedback.

We believe that tactile feedback can improve touch screen interactions on several layers. In a mobile situation, e.g., while walking, it can lower the mental load necessary to operate touch screens. Confirmations about user actions can be provided by tactile feedback thereby allowing eyes-free interaction. For sensitive information such as passwords tactile feedback can be used to implement a secure haptic keypad (Bianchi et al. [2010]).

MudPad is a research prototype allowing to investigate the range of possible applications and the usefulness of localized tactile feedback. Due to its system design it is highly unlikely for it to become ever commercialized. Nonetheless, it is valuable for research purposes as long as there are no systems available to investigate localized active haptic feedback for touch screens.

7.2 Future work

To fully explore the range of applications for which MudPad could be used, the current single touch screen needs to be replaced by a multi touch screen. The availability of such screens is currently extremely low with only very few manufacturers asking very high prices. Recently, a Korean start-up announced kits based on the ideas presented by Rosenberg and Perlin [2009]. The ability to simultaneously present different tactile feedback patterns on the same screen–possibly for different users–can only be employed with a multi touch screen.

Upgrade touch screen to multi touch.

As we already mentioned in the conclusion on our exhibitions, the full range of possible viscosities need to be explored. We decided to prepare our fluid ourselves to achieve a chemically compliant fluid for our latex surface. There are probably better suited fluids but the focus of this thesis was not on specific properties of the fluid. We stopped investigating different compositions once we found a suitable mixture.

Find better suited fluid composition.

Most interestingly, there are many studies that can be conducted with MudPad. First, it would be useful to measure the actual physical forces exhibited by MudPad. Then psychophysical studies on limits and ranges of *perceivable* feedback patterns would be helpful in designing appropriate patterns for use in applications.

Psychophysical experiments to measure range of perceivable feedback.

Then, new development tools are necessary to enable application developers to easily add tactile feedback to their applications. Currently there is no *industry standard*. Each vendor of tactile feedback devices supplies an individual API. There are efforts to establish common open-source APIs but they are currently not equipped to design a haptic representation for an entire GUI. They are rather focused on single actuators or on virtual haptic spaces. As of today we are not aware of a suitable solution for touch screen interfaces.

Tools to help developer add tactile feedback to touch screen applications.

Appendix A

Statistical Data

A.1 Descriptive Statistics

	Mean	Std. Deviation	N
AvgKeystrokeTime. MudPad.audio	461.244763	71.1862142	17
AvgKeystrokeTime. MudPad.combined	481.535692	81.8374423	17
AvgKeystrokeTime. MudPad.none	482.914234	98.5774510	17
AvgKeystrokeTime. MudPad.tactile	488.885796	72.1796192	17
AvgKeystrokeTime. TouchScreen.audio	440.630529	80.6294079	17
AvgKeystrokeTime. TouchScreen.combined	464.776216	69.4270174	17
AvgKeystrokeTime. TouchScreen.none	467.859899	88.9349099	17
AvgKeystrokeTime. TouchScreen.tactile	505.391452	95.7469107	17

Table A.1: Descriptive statistics for average keystroke time.

	Mean	Std. Deviation	N
Typos.MudPad.audio	.59	.712	17
Typos.MudPad. combined	.18	.728	17
Typos.MudPad.none	.94	1.435	17
Typos.MudPad.tactile	.35	.493	17
Typos.TouchScreen. audio	.59	1.004	17
Typos.TouchScreen. combined	.41	.712	17
Typos.TouchScreen. none	1.35	1.455	17
Typos.TouchScreen. tactile	1.41	1.460	17

Table A.2: Descriptive statistics for number of errors.

	Mean	Std. Deviation	N
Corrections.MudPad. none	4.18	3.486	17
Corrections.MudPad. audio	2.53	2.095	17
Corrections.MudPad. tactile	7.12	4.484	17
Corrections.MudPad. combined	2.88	1.867	17
Corrections. TouchScreen.none	3.53	2.427	17
Corrections. TouchScreen.audio	2.53	2.183	17
Corrections. TouchScreen.tactile	3.88	3.100	17
Corrections. TouchScreen.combined	2.71	1.724	17

Table A.3: Descriptive statistics for number of corrections.

A.2 Pairwise Comparisons

A.2.1 Main effect of input device

Measure:keystroke

(I) InputDevice	(J) InputDevice	Mean Difference (I−J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	-20.228	9.799	.056	-41.002	.546
	3	-11.247	10.127	.283	-32.716	10.221
2	1	20.228	9.799	.056	-.546	41.002
	3	8.981	8.125	.285	-8.243	26.204
3	1	11.247	10.127	.283	-10.221	32.716
	2	-8.981	8.125	.285	-26.204	8.243

Table A.4: Pairwise comparison of main effect on average keystroke time. Input device 1 = numerpad, 2 = MudPad, 3 = touch screen.

Measure:typos

(I) InputDevice	(J) InputDevice	Mean Difference (I−J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	.132	.238	.586	-.372	.637
	3	-.294	.261	.276	-.847	.259
2	1	-.132	.238	.586	-.637	.372
	3	-.426*	.150	.012	-.744	-.109
3	1	.294	.261	.276	-.259	.847
	2	.426*	.150	.012	.109	.744

Table A.5: Pairwise comparison of main effect on number of errors. Input device 1 = numerpad, 2 = MudPad, 3 = touch screen.

Measure:corrections

(I) InputDevice	(J) InputDevice	Mean Difference (I−J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	-.176	.633	.784	-1.519	1.166
	3	.838	.628	.200	-.492	2.169
2	1	.176	.633	.784	-1.166	1.519
	3	1.015*	.389	.019	.190	1.840
3	1	-.838	.628	.200	-2.169	.492
	2	-1.015*	.389	.019	-1.840	-.190

Table A.6: Pairwise comparison of main effect on number of corrections. Input device 1 = numerpad, 2 = MudPad, 3 = touch screen.

A.2.2 Effect of feedback type

Measure:keystroke

(I) FeedbackType	(J) FeedbackType	Mean Difference (I−J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	−22.218	12.288	.089	−48.267	3.830
	3	−24.449*	8.635	.012	−42.755	−6.144
	4	−46.201*	17.391	.017	−83.069	−9.333
2	1	22.218	12.288	.089	−3.830	48.267
	3	−2.231	9.477	.817	−22.322	17.860
	4	−23.983*	10.483	.036	−46.207	−1.759
3	1	24.449*	8.635	.012	6.144	42.755
	2	2.231	9.477	.817	−17.860	22.322
	4	−21.752	14.341	.149	−52.153	8.650
4	1	46.201*	17.391	.017	9.333	83.069
	2	23.983*	10.483	.036	1.759	46.207
	3	21.752	14.341	.149	−8.650	52.153

Table A.7: Pairwise comparison of feedback type on average keystroke time. Feedback type 1 = audio, 2 = combined, 3 = none, 4 = tactile.

Measure:errors

(I) FeedbackType	(J) FeedbackType	Mean Difference (I−J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	.294	.192	.145	−.113	.701
	3	−.559	.323	.103	−1.244	.127
	4	−.294	.219	.198	−.758	.170
2	1	−.294	.192	.145	−.701	.113
	3	−.853*	.249	.003	−1.381	−.325
	4	−.588*	.183	.005	−.976	−.200
3	1	.559	.323	.103	−.127	1.244
	2	.853*	.249	.003	.325	1.381
	4	.265	.265	.332	−.296	.826
4	1	.294	.219	.198	−.170	.758
	2	.588*	.183	.005	.200	.976
	3	−.265	.265	.332	−.826	.296

Table A.8: Pairwise comparison of feedback type on number of corrections. Feedback type 1 = audio, 2 = combined, 3 = none, 4 = tactile.

Measure:corrections

(I) FeedbackType	(J) FeedbackType	Mean Difference (I–J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	1.324*	.586	.038	.081	2.566
	3	-1.647*	.732	.039	-3.200	-.095
	4	1.059	.578	.086	-.167	2.285
2	1	-1.324*	.586	.038	-2.566	-.081
	3	-2.971*	.762	.001	-4.585	-1.356
	4	-.265	.391	.508	-1.093	.564
3	1	1.647*	.732	.039	.095	3.200
	2	2.971*	.762	.001	1.356	4.585
	4	2.706*	.606	.000	1.420	3.991
4	1	-1.059	.578	.086	-2.285	.167
	2	.265	.391	.508	-.564	1.093
	3	-2.706*	.606	.000	-3.991	-1.420

Table A.9: Pairwise comparison of feedback type on number of errors. Feedback type 1 = audio, 2 = combined, 3 = none, 4 = tactile.

Bibliography

Pär-Anders Albinsson and Shumin Zhai. High precision touch screen interaction. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '03, pages 105–112, New York, NY, USA, 2003. ACM.

Hideyuki Ando, Takeshi Miki, Masahiko Inami, and Taro Maeda. Smartfinger: Nail-mounted tactile display. In *SIGGRAPH '02: ACM SIGGRAPH 2002 conference abstracts and applications*, pages 78–78, New York, NY, USA, 2002. ACM.

Olivier Bau, Uros Petrevski, and Wendy Mackay. Bubblewrap: A textile-based electromagnetic haptic display. In *CHI EA '09: Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, pages 3607–3612, New York, NY, USA, 2009. ACM.

Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. Teslatouch: electrovibration for touch surfaces. In *UIST '10: Proceedings of the 23nd annual ACM symposium on User interface software and technology*, pages 283–292, New York, NY, USA, 2010. ACM.

Michel Beaudouin-Lafon. Instrumental interaction: an interaction model for designing post-wimp user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '00, pages 446–453, New York, NY, USA, 2000. ACM.

Andrea Bianchi, Ian Oakley, and Dong Soo Kwon. The secure haptic keypad: a tactile password system. In *CHI '10: Proceedings of the 28th international conference on Human factors in computing systems*, pages 1089–1092, New York, NY, USA, 2010. ACM.

A. Bicchi, E.P. Scilingo, N. Sgambelluri, and D. De Rossi. Haptic interfaces based on magnetorheological fluids. In *Proceedings 2th International Conference Eurohaptics 2002*, pages 6–11. 2002.

Stephen Brewster and Lorna M. Brown. Tactons: Structured tactile messages for non-visual information display. In *AUIC '04: Proceedings of the fifth conference on Australasian user interface*, pages 15–23, Darlinghurst, Australia, Australia, 2004. Australian Computer Society, Inc.

Stephen Brewster, Faraz Chohan, and Lorna Brown. Tactile feedback for mobile interactions. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 159–162, New York, NY, USA, 2007. ACM.

C.M. Butter, H.A. Buchtel, and R. Santucci. Spatial attentional shifts: Further evidence for the role of polysensory mechanisms using visual and tactile stimuli. *Neuropsychologia*, 27(10):1231–1240, 1989.

Angela Chang, Sile O'Modhrain, Rob Jacob, Eric Gunther, and Hiroshi Ishii. Comtouch: design of a vibrotactile communication device. In *DIS '02: Proceedings of the 4th conference on Designing interactive systems*, pages 312–320, New York, NY, USA, 2002. ACM.

V.G. Chouvardas, A.N. Miliou, and M.K. Hatalis. Tactile displays: Overview and recent advances. *Displays*, 29(3): 185 – 194, 2008.

Paul Dietz and Darren Leigh. Diamondtouch: a multi-user touch technology. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, UIST '01, pages 219–226, New York, NY, USA, 2001. ACM.

M.O. Ernst and M.S. Banks. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870):429–433, 2002.

Franklin N. Eventoff. Electronic pressure sensitive force transducer. U.S. Patent No. 4489302, 1984.

Masaaki Fukumoto and Toshiaki Sugimura. Active click: tactile feedback for touch panels. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 121–122, New York, NY, USA, 2001. ACM.

John Galeotti, Samantha Horvath, Roberta Klatzky, Brock Nichol, Mel Siegel, and George Stetten. Fingersight^{TM}: Fingertip control and haptic sensing of the visual environment. In *SIGGRAPH '08: ACM SIGGRAPH 2008 new tech demos*, pages 1–1, New York, NY, USA, 2008. ACM.

E.P. Gardner and J.H. Martin. Coding of sensory information. *Principles of neural science*, 4:411–429, 2000.

E.P. Gardner, J.H. Martin, and T.M. Jessell. The bodily senses. *Principles of Neural Science*, pages 430–450, 2000.

Jefferson Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *UIST '05: Proceedings of the 18th annual ACM Symposium on User Interface Software and Technology*, pages 115–118, New York, NY, USA, 2005. ACM.

Chris Harrison and Scott E. Hudson. Providing dynamically changeable physical buttons on a visual display. In *CHI '09: Proceedings of the 27th international conference on Human factors in computing systems*, pages 299–308, New York, NY, USA, 2009. ACM.

O. Hilliges, D. Kim, and S. Izadi. Creating malleable interactive surfaces using liquid displacement sensing. In *Horizontal Interactive Human Computer Systems, 2008. TABLETOP 2008. 3rd IEEE International Workshop on*, pages 157–160, Oct. 2008.

Steve Hodges, Shahram Izadi, Alex Butler, Alban Rrustemi, and Bill Buxton. Thinsight: versatile multi-touch sensing for thin form-factor displays. In *Proceedings of the 20th annual ACM symposium on User interface software and technology*, UIST '07, pages 259–268, New York, NY, USA, 2007. ACM.

Ramon Hofer, Daniel Naeff, and Andreas Kunz. Flatir: Ftir multi-touch detection on a discrete distributed sensor array. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction*, TEI '09, pages 317–322, New York, NY, USA, 2009. ACM.

Alexander Hoffmann, Daniel Spelmezan, and Jan Borchers. Typeright: a keyboard with tactile error prevention. In *CHI '09: Proceedings of the 27th international conference* *on Human factors in computing systems*, pages 2265–2268, New York, NY, USA, 2009. ACM.

Eve Hoggan, Sohail Anwar, and Stephen A. Brewster. Mobile multi-actuator tactile displays. In *HAID'07: Proceedings of the 2nd international conference on Haptic and audio interaction design*, pages 22–33, Berlin, Heidelberg, 2007. Springer-Verlag.

Eve Hoggan, Stephen A. Brewster, and Jody Johnston. Investigating the effectiveness of tactile feedback for mobile touchscreens. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1573–1582, New York, NY, USA, 2008. ACM.

E.L. Hutchins, J.D. Hollan, and D.A. Norman. Direct manipulation interfaces. *Human–Computer Interaction*, 1(4): 311–338, 1985.

Takayuki Iwamoto, Mari Tatezono, and Hiroyuki Shinoda. Non-contact method for producing tactile sensation using airborne ultrasound. In *EuroHaptics '08: Proceedings of the 6th international conference on Haptics*, pages 504–513, Berlin, Heidelberg, 2008. Springer-Verlag.

Daniel Jackson, Tom Bartindale, and Patrick Olivier. Fiberboard - compact multi-touch display using channeled light. In *ITS '09: Proceedings of ACM International Conference on Interactive Tabletops and Surfaces*, pages 25 – 28, 2009.

Yvonne Jansen. Mudpad: fluid haptics for multitouch surfaces. In *CHI EA '10: Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems*, pages 4351–4356, New York, NY, USA, 2010. ACM.

Yvonne Jansen, Thorsten Karrer, and Jan Borchers. Mudpad: localized tactile feedback on touch surfaces. In *Adjunct proceedings of the 23nd annual ACM symposium on User interface software and technology*, UIST '10, pages 385–386, New York, NY, USA, October 2010a. ACM.

Yvonne Jansen, Thorsten Karrer, and Jan Borchers. Mudpad: tactile feedback and haptic texture overlay for touch surfaces. In *ACM International Conference on Interactive*

Tabletops and Surfaces, ITS '10, pages 11–14, New York, NY, USA, November 2010b. ACM.

Yvonne Jansen, Thorsten Karrer, and Jan Borchers. Mudpad: a tactile memory game. In *ACM International Conference on Interactive Tabletops and Surfaces*, ITS '10, pages 306–306, New York, NY, USA, November 2010c. ACM.

Pascal Joguet and Guillaume Largillier. Devices and methods of controlling manipulation of virtual objects on a multi-contact tactile screen. U.S. Patent Application No. 2007/0198926, 2007.

Sergi Jordà, Günter Geiger, Marcos Alonso, and Martin Kaltenbrunner. The reactable: Exploring the synergy between live music performance and tabletop tangible interfaces. In *TEI '07: Proceedings of the 1st international conference on Tangible and embedded interaction*, pages 139–146, New York, NY, USA, 2007. ACM.

S J Lederman and R L Klatzky. Haptic perception: A tutorial. *Atten Percept Psychophys*, 71(7):1439–59, Oct 2009.

Johnny C. Lee, Paul H. Dietz, Darren Leigh, William S. Yerazunis, and Scott E. Hudson. Haptic pen: A tactile feedback stylus for touch screens. In *UIST '04: Proceedings of the 17th annual ACM symposium on User interface software and technology*, pages 291–294, New York, NY, USA, 2004. ACM.

SK Lee, W. Buxton, and KC Smith. A multi-touch three dimensional touch-sensitive tablet. *ACM SIGCHI Bulletin*, 16(4):21–25, 1985.

Daniel Leithinger and Hiroshi Ishii. Relief: A scalable actuated shape display. In *TEI '10: Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction*, pages 221–222, New York, NY, USA, 2010. ACM.

Rock Leung, Karon MacLean, Martin Bue Bertelsen, and Mayukh Saubhasik. Evaluation of haptically augmented touchscreen gui elements under cognitive load. In *ICMI '07: Proceedings of the 9th international conference on Multimodal interfaces*, pages 374–381, New York, NY, USA, 2007. ACM.

K. MacLean and M. Enriquez. Perceptual design of haptic icons. In *Proceedings of Eurohaptics*, pages 351–363. 2003.

Nicolai Marquardt, Miguel A. Nacenta, James E. Young, Sheelagh Carpendaled, Saul Greenberg, and Ehud Sharlin. The haptic tabletop puck: Tactile feedback for interactive tabletops. In *ITS '09: Proceedings of ACM International Conference on Interactive Tabletops and Surfaces*, pages 85 – 92, 2009.

T.H. Massie and J.K. Salisbury. The phantom haptic interface: A device for probing virtual objects. In *Proceedings of the ASME Winter Annual Meeting, Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, volume 55, pages 295–300, 1994.

M. Milczynski, T. Hermann, T. Bovermann, and H. Ritter. A malleable device with applications to sonification-based data exploration. In *Proceedings of the 12th Meeting of the International Conference on Auditory Display*, 2006.

Jon Moeller and Andruid Kerne. Scanning ftir: unobtrusive optoelectronic multi-touch sensing through waveguide transmissivity imaging. In *Proceedings of the fourth international conference on Tangible, embedded, and embodied interaction*, TEI '10, pages 73–76, New York, NY, USA, 2010. ACM.

M. Nielsen, M. St
"orring, T.B. Moeslund, and E. Granum. A procedure for developing intuitive and ergonomic gesture interfaces for HCI. *Gesture-Based Communication in Human-Computer Interaction*, pages 105–106, 2004.

Ian Oakley, Stephen Brewster, and Philip Gray. Solving multi-target haptic problems in menu interaction. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 357–358, New York, NY, USA, 2001. ACM.

Sabrina Pan¨ eels, Jonathan Roberts, and Peter Rodgers. Haptic interaction techniques for exploring chart data. In M. Altinsoy, Ute Jekosch, and Stephen Brewster, editors, *Haptic and Audio Interaction Design*, volume 5763 of *Lecture Notes in Computer Science*, pages 31–40. Springer Berlin / Heidelberg, 2009.

Gian Pangaro, Dan Maynes-Aminzade, and Hiroshi Ishii. The actuated workbench: Computer-controlled actuation in tabletop tangible interfaces. In *UIST '02: Proceedings of the 15th annual ACM symposium on User interface software and technology*, pages 181–190, New York, NY, USA, 2002. ACM.

R. L. Potter, L. J. Weldon, and B. Shneiderman. Improving the accuracy of touch screens: an experimental evaluation of three strategies. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '88, pages 27–32, New York, NY, USA, 1988. ACM.

Ivan Poupyrev and Shigeaki Maruyama. Tactile interfaces for small touch screens. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 217–220, New York, NY, USA, 2003. ACM.

Ivan Poupyrev, Shigeaki Maruyama, and Jun Rekimoto. Ambient touch: designing tactile interfaces for handheld devices. In *UIST '02: Proceedings of the 15th annual ACM symposium on User interface software and technology*, pages 51–60, New York, NY, USA, 2002. ACM.

Ivan Poupyrev, Tatsushi Nashida, Shigeaki Maruyama, Jun Rekimoto, and Yasufumi Yamaji. Lumen: Interactive visual and shape display for calm computing. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Emerging technologies*, page 17, New York, NY, USA, 2004. ACM.

Jun Rekimoto. Smartskin: An infrastructure for freehand manipulation on interactive surfaces. In *CHI '02: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 113–120, New York, NY, USA, 2002. ACM.

Jun Rekimoto. Senseablerays: Opto-haptic substitution for touch-enhanced interactive spaces. In *CHI EA '09: Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*, pages 2519–2528, New York, NY, USA, 2009. ACM.

Jun Rekimoto and Carsten Schwesig. Presenseii: bidirectional touch and pressure sensing interactions with tactile feedback. In *CHI '06: CHI '06 extended abstracts on Human factors in computing systems*, pages 1253–1258, New York, NY, USA, 2006. ACM.

J.C. Roberts and S. Paneels. Where are we with haptic visualization? In *EuroHaptics Conference, 2007 and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems. World Haptics 2007. Second Joint*, pages 316–323. IEEE, 2007.

Ilya Rosenberg and Ken Perlin. The UnMousePad: an interpolating multi-touch force-sensing input pad. In *SIGGRAPH '09: Proceedings of the 36th annual conference on Computer graphics and interactive techniques*, pages 1–9, New Orleans, Louisiana, 2009. ACM.

Andrew Sears. Improving touchscreen keyboards: design issues and a comparison with other devices. *Interacting with Computers*, 3(3):253 – 269, 1991.

Andrew Sears and Ben Shneiderman. High precision touchscreens: design strategies and comparisons with a mouse. *Int. J. Man-Mach. Stud.*, 34:593–613, April 1991.

N. Sgambelluri, E.P. Scilingo, A. Bicchi, R. Rizzo, and M. Raugi. Advanced modelling and preliminary psychophysical experiments for a free-hand haptic device. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1558–1563, 2006.

B. Shneiderman. The future of interactive systems and the emergence of direct manipulation. *Behaviour & Information Technology*, 1(3):237–256, 1982.

B. Shneiderman. Direct manipulation: A step beyond programming languages. *Computer*, 16:57–69, August 1983.

P. Steurer and M.B. Srivastava. System design of smart table. In *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications*, page 473. IEEE Computer Society Washington, DC, USA, 2003.

PM Taylor, A. Hosseini-Sianaki, and CJ Varley. An electrorheological fluid-based tactile array for virtual environments. In *1996 IEEE International Conference on Robotics and Automation, 1996. Proceedings.*, volume 1, 1996.

RQ Van der Linde, P. Lammertse, E. Frederiksen, and B. Ruiter. The hapticmaster, a new high-performance haptic interface. In *Proc. Eurohaptics*, pages 1–5. 2002.

Florian Vogt, Timothy Chen, Reynald Hoskinson, and Sidney Fels. A malleable surface touch interface. In *SIGGRAPH '04: ACM SIGGRAPH 2004 Sketches*, page 36, New York, NY, USA, 2004. ACM.

CR Wagner, SJ Lederman, and RD Howe. Design and performance of a tactile shape display using rc servomotors (short paper). *Haptics-e: The Electronic Journal for Haptics Research*, 3(4), 2004.

Julie Wagner. Slap, silicone illuminated active peripherals. Master's thesis, RWTH Aachen University, 2009.

Malte Weiss, Florian Schwarz, Simon Jakubowski, and Jan Borchers. Madgets: actuating widgets on interactive tabletops. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, UIST '10, pages 293–302, New York, NY, USA, 2010. ACM.

T. White. Introducing liquid haptics in high bandwidth human computer interfaces. Master's thesis, MIT, May 1998.

Jacob O. Wobbrock, Meredith Ringel Morris, and Andrew D. Wilson. User-defined gestures for surface computing. In *Proceedings of the 27th international conference on Human factors in computing systems*, CHI '09, pages 1083–1092, New York, NY, USA, 2009. ACM.

Koji Yatani and Khai Nhut Truong. Semfeel: a user interface with semantic tactile feedback for mobile touch-screen devices. In *UIST '09: Proceedings of the 22nd annual ACM symposium on User interface software and technology*, pages 111–120, New York, NY, USA, 2009. ACM.

Thomas G. Zimmerman, Joshua R. Smith, Joseph A. Paradiso, David Allport, and Neil Gershenfeld. Applying electric field sensing to human-computer interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '95, pages 280–287, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.

Index

Typeset March 14, 2011