

What You See Is What You Feel

A catalogue record is available from the Eindhoven University of Technology Library

ISBN: 978-90-386-1672-8

© 2009 Koert van Mensvoort

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronical or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission from the author.

WHAT YOU SEE IS WHAT YOU FEEL

On the simulation of touch in graphical user interfaces

Proefschrift

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven, op gezag van de
Rector Magnificus, prof.dr.ir. C.J. van Duijn, voor een
commissie aangewezen door het College voor
Promoties in het openbaar te verdedigen
op woensdag 29 april 2009 om 16.00

door

Koert Martinus van Mensvoort

geboren te Veldhoven

Dit proefschrift is goedgekeurd door de promotoren:

prof.dr.ir. J.H. Eggen

en

prof.dr.ir. J.J. van Wijk

Contents

PROLOGUE: A SOCIETY OF SIMULATIONS	9
1 INTRODUCTION	17
1.1 In this Thesis	17
1.2 Urge for Physical Interaction	17
Touch in Graphical User Interfaces	19
1.3 Physical Computing	21
Haptic Technology	22
Ubiquitous Computing, Tangible Interaction	23
1.4 The Persistence of existing Technology	24
The Failure of Force Feedback	25
A different approach	26
1.5 Research questions and goals	27
1.6 Terminology	27
1.7 Overview of the chapters	30
2 OPTICALLY SIMULATED HAPTIC FEEDBACK	31
2.1 Renaissance Tricks	32
2.2 Animated Graphical User Interfaces	34
Learning from Disney Animators	35
2.3 Introducing the Active Cursor Technique	38
Two-way Communication via Cursor Position	39
Contextual feedback	41
Algorithm	44
2.4 Discussion	45
3 MEASURING THE ILLUSION	47
3.1 Related Work	48
Application in WIMP interfaces	49
3.2 Hypotheses	51
3.3 Methods	51
Subjects	51
Apparatus	51
Procedure	53
Design	53
3.4 Results	55

	Experiment I	57
	Experiments II and III	58
	Comparison of Experiment I, II, and III	61
	Prior Knowledge	63
	Result Summary	63
3.5	Discussion	64
	Conclusions	65
4	USABILITY OF OPTICALLY SIMULATED HAPTIC FEEDBACK	67
4.1	Related Work	67
	Hypotheses	69
4.2	Experiment	69
	Subjects	69
	Apparatus	70
	Procedure	71
	Design	71
4.3	Results	73
	Movement time	73
	Efficiency: Index of Performance	75
	Effectiveness: Error Rate	76
	Satisfaction: User Preference	77
4.4	Discussion	78
	Drawbacks & Limitations	80
	Conclusions	80
5	DESIGNING INTERFACES YOU CAN TOUCH	83
5.1	Powercursor – Goals & Design Considerations	84
	Target user group: Interaction Designers	84
	Choice of Platform: Adobe Flash	85
5.2	Basic Elements of the PowerCursor Toolkit	88
	Essential elements: Engine and Cursor	88
	Structure: Generic & Mergeable Components	89
	Basic Ready Made Objects	90
	Behaviours	95
	Event Handling	98
	Sound Design Support	100
	Debug Tools	101
5.3	Combining elements: like LEGO	102
	Example 1: Hole-Hill	102
	Example 2: The Pit	103

	Example 3: Dynamic Button	104
5.4	Exploring the Applicability	105
	Assisted Navigation	105
	Mixed initiative and persuasive interfaces	106
	Interpassive Interfaces	107
	Aesthetical Interactions	108
	Outside the desktop	110
5.5	Discussion	111
6	GENERAL CONCLUSIONS & FUTURE DIRECTIONS	115
6.1	Research Results	115
	Drawbacks of optically simulated haptic feedback	117
6.2	Future Directions	118
	Simulate your new computer on your old computer	119
	There is more in the box than there is in the box	119
	New media create new perceptual spaces	120
	Simulations: inferior derivatives or catalysts of change?	120
6.3	In Conclusion	121
	EPILOGUE: FROM NEW MEDIA TO NEXT NATURE	123
	REFERENCES	127
	SUMMARY	139
	SAMENVATTING	143
	ACKNOWLEDGMENTS	147
	PHD COMMITTEE	148
	PROPOSITIONS	149
	STELLINGEN	150
	CURRICULUM VITAE	151

An interviewer once asked Pablo Picasso why he paints such strange pictures instead of painting things the way they are. Picasso asks the man what he means. The man then takes out a photograph from his wallet and says, "This is my wife!" Picasso looks at the photo and then says: "isn't she rather short and flat?"

PROLOGUE¹

A Society of Simulations

Before diving into the specifics of my research, I wish to briefly explore the social-cultural context in which this project was conducted. Following on some personal observations regarding the dominant role of visual representations in our culture, I will argue that we are now living in a society, in which simulations are often more influential, satisfying and meaningful than the things they are presumed to represent. Media technologies play a fundamental role in our cycle of meaning construction. This is not necessarily a bad thing, nor is it entirely new. Yet, it has consequences for our concepts of *virtual* and *real*, which are less complementary, than they are usually understood to be.

In the research presented further on in this thesis, I aim to take advantage of these observations regarding simulations and visual dominance, applying them practically and positively towards a richer and more physical paradigm of graphical computer interfaces.

VISUAL POWER

Before you read on, a personal anecdote from my youth: when I was a child, I thought the people I saw on TV were really living inside the television. I wondered where they went when the TV was turned off and I also remember worrying it would hurt the TV, when I switched it off. Obviously, I am a grown man now and I've long learned that the television is just a technological device, created to project distant images into the living room of the viewers and that those flickering people weren't actually living inside the cathode ray tube.

Now I return to my argument. Over the last century or so, the technological reproduction of images has grown explosively. Each of us is confronted with more images every day than a person living in the Middle Ages would have seen in their whole lifetime. If you open a 100-year-old newspaper you will be amazed by the volume of text and the absence of pictures. How different things are today: the moment you are born, covered in

¹ This essay is a combined and extended version of two earlier published texts: Mensvoort, Koert van (2007) *The Picture Bubble*, in Gerritzen et al. *Style First*, Birkhauser Basel, ISBN-13: 978-3764384388, pp 48-52; and Mensvoort, Koert van & Grievink, Hendrik-Jan (2008) *Fake for Real*, AllMedia / Bis Publishers, ISBN 978-9063691776

womb fluid, not yet dressed or showered, your parents are already there with the digital camera, ready to take your picture. And of course the pictures are instantly uploaded to the family website, where the whole world can watch and compare them with the medical ultrasound photographs already shared before you were born.

Images occupy an increasingly important place in our communication and transmission of information. More and more often, it is an image that is the deciding factor in important questions. Provocative logos, styles and icons are supposed to make us think we are connected to each other, or different from each other. Every schoolchild nowadays has to decide whether he or she is a skater, a jock, a preppie, or whatever. Going to school naked is not an option. But no matter which T-shirt you decide to wear, they are inescapably a social communication medium. Your T-shirt will be read as a statement, which your classmates will use to stereotype you.

I remember the strange feeling of recognition I had when I was in Paris for the first time and saw the Eiffel Tower. There it was, for real! I felt as if I was meeting a long-lost cousin. Of course, you take a snapshot to show you've been there: 'Me and the Eiffel Tower'. Thousands of people take this same picture every year. Every architect dreams of designing such an icon. Today, exceptional architecture often wins prizes before the building is finished; their iconic quality is already recognized on the basis of computer models².

PICTURE THIS!

Does anyone still remember the days when a computer was a complex machine that could only be operated by a highly trained expert using obscure commands? Only when the graphical user interface (GUI) was introduced did computers become everyday appliances; suddenly anyone could use them. Today, all over the world, people from various cultures use the same icons, folders, buttons and trash cans. The GUI's success is owed less to the cute pictures than to the metaphor that makes the machine so accessible: the computer desktop as a version of the familiar, old-fashioned kind. This brings us to an important difference between pictures and pictures – it is indeed awkward that we use the same word for two different things. On the one hand, there are pictures we see with our eyes. On the other, there are mental pictures we have in our heads – pictures as in "I'm trying to picture it."

Increasingly, we are coming to realize that 'thinking' is fundamentally connected to sensory experience. In *Metaphors We Live By*, Lakoff and Johnson (1980) argue that human thought works in a fundamentally metaphorical way. Metaphors allow us to use physical and social experiences to understand countless other subjects. The world we live in has become so complex; we continuously search for mental imagery to help us help us understand things. Thus politicians speak in clear sound bites. Athletic

2 Examples of architectural structures that are already famous and celebrated before being build are the Freedom Tower by Liebeskind/Childs in New York and the CCTV building by Rem Koolhaas in Beijing.

shoe companies do not sell shoes, they sell image. Thoracic surgeons wander around in patients' lungs like rangers walking through the forest, courtesy of head-mounted virtual-reality displays.

You would expect that this surfeit of images would drown us. It is now difficult to deny that a certain visual inflation is present, and yet our unslakeable hunger for more persists. We humans, after all, are extremely visually oriented animals. From cave paintings to computers, the visual image has helped the human race to describe, classify, order, analyze and grow our understanding of the world around us (Bright, 2000). Perhaps the most extraordinary thing about our visual culture (Mirzoeff, 1999) is not the number of pictures being produced but our deeply rooted need to visualize everything that could possibly be significant. Modern life amid visual media compels everyone and everything to strive for visibility (Winkel, 2006). The more visible something is, the more real it is, the more genuine (Oosterling, 2003). Without images, there seems to be no reality.

VIRTUAL FOR REAL

When considering simulations, one almost immediately thinks of videogames. Nowadays, the game industry has grown bigger than the film industry and its visual language has become so accepted that it is almost beyond fictional. Virtual computer worlds are becoming increasingly 'real' and blended with our physical world. In some online roleplaying games, aspiring participants have to write an application letter in order to be accepted to a certain group or tribe. We still have to get used to the fact that you can earn an income with gaming nowadays (Heeks, 2008), but how normal is it anyway, that at the bakery round the corner, you can trade a piece of paper – called money – for a bread?³

Most people would denounce spending too much time in virtual worlds, but which world should be called virtual then? Simply defining the virtual as opposite to physical is perhaps too simple. The word 'virtual' has different meanings that are often entangled and used without further consideration. Sometimes we use the word virtual to mean 'almost real,' while at other times we mean 'imaginary'. This disparity in meaning is almost never justified: fantasy and second rank realities are intertwined. It would be naïve to think simulations are limited to video games, professional industrial or military applications. In a sense, all reality is virtual; it is constructed through our cognition and sensory organs. Reality is not so much 'out there', rather it is what we pragmatically consider to be 'out there'. Our brain is able to subtly construct 'reality' by combining and comparing sensory perceptions with what we expect and already know (Dennett, 1991; Gregory, 1998; Hoffman, 1998; IJsselstein, 2002).

3 We usually do not realize that 'money' is in many respects a virtual phenomenon: a symbolic representation of value constructed to replace the awkward, imprecise trading of physical goods. Indeed, paying \$50 for a pair of sneakers is much easier than trading two chickens or a basket of apples for them. As long as we all believe in it, the monetary system works fine.

Even the ancient Greeks talked about the phenomenon of simulation. In the Allegory of the Cave, Plato describes human beings as being chained in a cave and watching shadows on the wall, without realizing that they are ‘only’ representations of what goes on behind them – outside of the scope of their sensory perception. In Plato’s teaching, an object such as a chair, is just a shadow of the idea Chair. The physically experienced chair we sit on is thus always a copy, a simulation, of the idea Chair and always one step away from reality.

Today, the walls of Plato’s cave are so full of projectors, disco balls, plasma screens and halogen spotlights that we do not even see the shadows on the wall anymore. Fakeness has long been associated with inferiority – fake Rolexes that break in two weeks, plastic Christmas trees, silicone breast implants, imitation caviar –, but as the presence of media production evolves, the fake seems to gain a certain authenticity. Modern thinkers agree that because of the impasto of simulations in our society, we can no longer recognize reality. In *The Society of the Spectacle*, Guy Debord (1967) explains how everything we once experienced directly has been replaced in our contemporary world by representations. Another Frenchman, Jean Baudrillard (1981), argues that we live in a world in which simulations and imitations of reality have become more real than reality itself. He calls this condition ‘hyperreality’: the authentic fake. In summer we ski indoors; in winter we spray snow on the slopes. Plastic surgeons sculpt flesh to match retouched photographs in glossy magazines. People drink sports drinks with non-existent flavors like “wild ice zest berry”. We wage war on video screens. Birds mimic mobile–phone ring tones.⁴ At times, it seems the surrealists were telling the truth after all. And though you certainly cannot believe everything you see, at the same time, images still count as the ultimate evidence. Did we really land on the moon? Are you sure? How did it happen? Or was it perhaps a feat of Hollywood magic? Are we sure there is no Loch Ness Monster? A city girl regularly washes her hair with pine-scented shampoo. Walking in the forest with her father one day, she says, “Daddy, the woods smell of shampoo.” Do we still have genuine experiences at all, or are we living in a society of simulations?

MEDIA SCHEMAS

A hundred years ago, when the Lumière brothers showed their film ‘L’arrivée d’un train’ (1895), people ran out of the cinema when they saw the oncoming train. Well, of course – if you see a train heading towards you, you get out of the way. Today, we have adapted our media schemas. We remain seated, because we know that the medium of cinema can have this effect.

4 The Superb Lyrebird living in Southern Australia sings and mimics all the calls of other birds, as well as other sounds he hears in the forest – even cellphone ring-tones, chainsaws and camera shutters – to attract females (Attenborough, 1998).

Media schemas⁵ are defined as the knowledge we possess about what media are capable of and what we should expect from them in terms of their depictions: representations, translations, distortions, etc (IJsselsteijn, 2002; Mensvoort & Duyvenbode, 2001; Nevejan, 2007). This knowledge enables us to react to media in a controlled way (“Don’t be scared, it’s only a movie.”). A superficial observer might think media schemas are a new thing. This would be incorrect. For centuries, people have been dealing with developments in media. Think of carrying on a telephone conversation, painting with perspective, or composing a letter with the aid of writing technology – yes, even the idea that you can set down the spoken word in handwriting was new once.

Let’s face it. Our brains actually have only limited capabilities for understanding media. When our brain reached its current state of evolutionary development in Africa some 200,000 years ago (Hedges, 2000; Goodman et al., 1990), what looked like a lion, actually was a lion! And if contemplating the nature of reality at that point would have been a priority, one would have made for an easy lion’s snack (IJsselsteijn, 2002). Although we do seem to have gained some media awareness over the years, some part of this original impulse – in spite of all our knowledge – still reacts automatically and unconsciously to phenomena, as we perceive them. When we see the image of an oncoming train, we physically still are inclined to run away, even though cognitively we know it is not necessary.

Our media schemas are thus not innate but culturally determined. Every time technology comes out with something new, we are temporarily flummoxed, but we carry on pretty well. We are used to a world of family photographs, television and telephone calls. Imagine if we were to put someone from the Middle Ages into a contemporary shopping street. He would have a tough job refreshing his media schemas. But to us it is normal, and a lucky thing, too. It would be inconvenient indeed if with every phone call you thought, “How strange – I’m talking to someone who’s actually far away.” We are generally only conscious of our media schemas at the moment when they prove inadequate and we must refresh them, as those people in the 19th century had to do when they saw the Lumière brothers’ filmed train coming at them.

MEDIA SPHERE

I once took part in an experiment in which I was placed in an entirely green room for one hour. In the beginning everything seemed very green, but after some time

5 The term *media schemas* stems from the concept of *schemas*, which in psychology and cognitive sciences is described as a mental structure that represents some aspect of the world (Piaget, 1997). According to schema theory, all human beings possess categorical rules or scripts that they use to interpret the world. New information is processed according to how it fits into these rules. These schemas can be used not only to interpret but also to predict situation occurring in our environment.

the walls became grey. The green was not informative any more and I automatically adjusted. Something similar seems to be going on with our media. Like the fish, who do not know they are wet; we are living in a technologically mediated space. We have adjusted ourselves, for the better because we know we will not be leaving this room any time soon. Today, media production has expanded by such leaps and bounds that images and simulations are often more influential, satisfying and meaningful than the things they simulate. We consume illusions. Images have become part of the cycle in which meanings are determined. They have bearing on our economy, our judgments and our identities. In other words: we are living the simulation.

A disturbing thought, or old news? In contrast to Plato, his pupil Aristotle believed imitation was a natural part of life. Reality reaches us through imitation (Aristotle calls it *mimesis*): this is how we come to know the world. Plants and animals too, use disguises and misleading appearances to improve their chances of survival (think of the walking stick, an insect that looks like a twig). Now then, the girl that says that “the woods smell of shampoo”, should we consider this a shame and claim that this young child has been spoiled by media? Or is this child merely fine-tuning herself with the environment she grows up in? In the past, the woods used to smell of woods. But how interesting was that anyway?

OUR INTERFACED WORLD-VIEW

Four centuries ago, when Galileo Galilei became the first human being in history to aim a telescope at the night sky, a world opened up to him. The moon turned out not to be a smooth, yellowish sphere but covered with craters and mountains. Nor was the sun perfect: it bore dark spots. Venus appeared in phases. Jupiter was accompanied by four moons. Saturn had a ring. And the Milky Way proved to be studded with hundreds of thousands of stars. When Galileo asserted, after a series of observations and calculations, that the sun was the center of our solar system, he had a big problem. No one wanted to look through his telescope to see the inevitable.

While some dogs have such limited intelligence that they chase their own tails or shadows, we humans like to think we are smarter; we are used to living in a world of complex symbolic languages and abstractions. While a dog remains fooled by his own shadow, a human being performs a reality check. We weigh up the phenomena in our environment against our actions to form a picture of what we call *reality*. We do this not only individually, but also socially (Searl, 1995). Admittedly, some realities are still rock solid – simply try and kick a stone to feel what I mean. However, this is not in conflict with the point I am trying to make, which is that the concepts of *reality* and *authority* are much more closely related to one another than most people realize. Like the physical world, which authority is pretty much absolute, media technologies are gradually but certainly attaining a level of authority within in our society that consequently increases their realness.

Today the telescope is a generally accepted means of observing the universe. The earth is no longer flat. We have long left the dark ages of religious dogma and have experienced great scientific breakthroughs, and yet there are still dominant forces shaping our world-view. As we are descending into the depths of our genes, greet webcam-friends across the ocean, send probes to the outskirts of the universe, find our way using car navigation, inspect our house's roof with Google earth and as it is not unusual for healthy, right-minded people to inform themselves about conditions in the world by spending the evening slouched in front of the television, we come to realize that our world-view is fundamentally being shaped through interfaces. Surely, the designers of these interfaces have an important responsibility in this regard. As media technologies evolve and are incorporated within our culture, our experience of reality changes along. This process is so profound – and one could argue, successful – it almost goes without notice, that to a large extent, we are living in a virtual world already.

In the research presented hereafter, I aim to positively take advantage of the fluid border between the virtual and the real, in proposing that it is possible to leverage the reality-constructing abilities of the human mind to simulate touch through purely optical means.

1

Introduction

1.1 IN THIS THESIS

We explore the role of simulations in our society and specifically we investigate the application of simulated touch in visual interfaces. As part of this research, we present optically simulated haptic feedback, an approach to simulate touch percepts in a standard graphical user interface without resorting to special and scarcely available haptic input/output devices. We investigate the perceptual experience of optically simulated haptic feedback, establish the usability benefits of the technique and present a prototyping toolkit that enables designers to seamlessly apply visual force feedback in their interfaces. Our aim is to contribute to a richer and more physical paradigm of graphical user interfaces. Moreover, we aim to increase our awareness and understanding of simulations in general. Our scientific research results are therefore deliberately presented in a socio-cultural context that reflects the dominance of the visual modality in our society and the ever-increasing role of media and simulations in people's everyday lives.

1.2 URGE FOR PHYSICAL INTERACTION

In our physical world, the kinetic behavior of objects is self-explanatory. It informs us about the physical properties of an object. If you open a door you will feel a certain resistance that tells you something about the door, how it is placed and what it is made of. When you lift a box you feel whether the box is full or empty. Everyday expressions such as, 'hands on experience', 'get the feel of it', 'feel the rhythm', 'have a feeling for', 'handy', 'hold on', 'get a grip on it' reflect the closeness of touch in interaction with our immediate environment (Keyson, 1996). Touch can play a powerful role in communication. It can offer an immediacy and intimacy unparalleled by words or images. Touch can be pleasurable or painful – it is one of our most intimate senses. In



Figure 1-1 Children demonstrating the richness and pleasure of physical interaction (image: Barbara Derksen).

the physical world, touch can further serve as a powerful mechanism for reinforcing trust and establishing group bonding (Burgoon et al., 1984; Burgoon, 1991). The firm handshake, an encouraging pat on the back, a comforting hug, all speaks to the profound expressiveness of physical contact.

Although few doubt this intrinsic value of touch perception in everyday life, examples in modern technology where human-machine communication utilizes the tactile and kinesthetic senses as additional channels of information flow are scarce. Our digital age is primarily a visual age; the visual modality is dominating our culture. According to Jean Baudrillard (1988), *“we live in the imaginary world of the screen, of the interface and the reduplication of contiguity and networks. All our machines are screens. We too have become screens, and the interactivity of men has become the interactivity of screens.”* While screens were originally found only in offices, nowadays they have made their way into homes, phones, shops, public squares, railway stations – they are more or less everywhere. We use these flat rectangular objects to inform ourselves about the state of our world. We use screens to check our e-mail, screens to monitor safety on the streets, screens to follow fashion. Scientists use screens to explore the outer limits of the universe and to descend into the structures of our genes. A painful truth: many of us spend more time with computer monitors than with our own friends and families (Massaro, 2007).



Figure 1-2 Physical scrollbar, an installation created by Dutch artist Jan Robert Leegte. According to the artist, most of us consider the scrollbar to be a virtual object – but in its use it triggers reactions such as frustration, which suggests a subconscious acceptance of the inherent “reality” of these objects.

Touch in Graphical User Interfaces

Since the invention of the mouse (English, Engelbart, Berman 1967) and the direct manipulation interface (Shneiderman, 1983), desktop metaphor interfaces based on windows, icons, menus and pointing – so called WIMP interfaces – have become the dominant paradigm in human-computer interaction. They are used while typing a letter, handling a spreadsheet, playing a game, doing 3D modeling, updating your social network, or watching Youtube videos. Whether using a PC, Mac, Linux, Desktop, Laptop or other, millions of people spend a significant part of their lives in front of a WIMP interface.

The conventions of use and interaction with computers have been accepted and adopted to relatively rapidly. All over the world, people from different cultures and social backgrounds have come to work with the same interface elements; windows, buttons, trashcan and folders that emulate, and have steadily replaced, the physical writing desk. The onscreen desktop displays an imaginary reality in which the user seemingly controls the machine. But, everyone who regularly works with a computer knows the other scenario: all of a sudden the computer can halt, display obscure error messages and do all kinds of things you were not planning on. Although this desktop metaphor is just an illusion – a rhetorical facade of otherwise incomprehensible technology – these conventions ease the use of a computer for almost everyone.

The average desktop computer setup consists of a mouse, keyboard, a flat 2D screen and two small speakers. The vast majority of current graphical user interfaces involve



Figure 1-3 Homo desktopus, an optimised human for desktop computing (Image Mensvoort 2002).

manipulation of onscreen artifacts with a mouse-controlled cursor (Myers, 1998). The mouse is the dominant pointing and selecting device and has become one of the most frequently handled devices in many people's daily lives. More frequent than cash, the steering wheel, doorknobs, pens, hammers, or screw drivers (Zhai and MacKenzie, 1998). Its design has not been altered much since its invention by English, Engelbart and Berman in 1967. There have been some improvements in the ergonomics of the mouse device. Many manufacturers place tiny wheels on the front of their mice and trackballs that users can roll to move vertically on-screen through documents and web pages. Some companies place pointing sticks

between the buttons of their mice to allow both vertical and horizontal scrolling. Improvements have been made in its shape and degrees of freedom. Mice have become optical and wireless.




Figure 1-4: First computer mouse, designed by Douglas Engelbart (English, Engelbart, Berman 1967)

From a sensorial point of view, the computers we use are extremely limited machines with hardly any physicality to them. They engage only a fraction of our human sensory bandwidth. If evolution were to naturally select the human race based solely on desktop computer use, people would evolve towards one-eyed blobs with tiny ears, a small mouth, no nose and a large click finger, and with no other sensory organs (Figure 1-3). Obviously this is not a probable future for the human race, but a future physical anthropologist who knew nothing about the human race might, upon digging up a contemporary desktopcomputer, conclude that *homo desktopus* must have been the users of the device (Buxton 1986).



Figure 1-5 Schöne Aussichten (Nice Views), photomontage created by artist Wibke Pausch (2005).

1.3 PHYSICAL COMPUTING

Before the prevalent use of computers, almost all human tasks involved the use of exquisite sensory-motor skills. By and large computer interfaces have not taken advantage of these deep-seated human capabilities. The touch feedback that did exist in older analog technologies through mechanical mechanisms such as knobs, switches and dials have for the most part been replaced by digital electronics and visual displays. The objects on your computer screen are completely lacking in bodily properties. Although this weightlessness of cyberspace has some major advantages, few would dispute the intrinsic value of touch perception in everyday interactions.



Figure 1-6 KlimaKontrolle, A fan in front of a computer screen accelerates until it blows away the entire desktop. The video subverts our preconceptions of the computer screen, and allows human physicality and atmospheric conditions to affect this normally closed digital space (Maurer and Wouters 2002).



Figure 1-7 Examples of force feedback devices: from left to right: a) The IPO force feedback trackball. Not only the user, but also the system can reposition the trackball, through the two added servo motors. (Engel et al., 1994). b) The Logitech wingman force feedback mouse (Rosenberg 1996), c) The Sensable phantom a force feedback enabled 3D pointing device (Massie and Salisbury, 1994).

Haptic Perception

The word haptic is based on the Greek word, “haptesthai,” meaning touch. Gibson (1966) defines haptics as “*The sensibility of the individual to the world adjacent to his body by use of his body*”. The haptic sensory modality consists of various mechanoreceptors (detecting skin deformations), proprioceptors (providing information about joint angle, muscle length, and tension) and thermoreceptors (coding absolute and relative changes in temperature), that work together with the primary sensory cortex (Mather 2006). Contrary to vision and hearing, which are passive (input only) senses that can not act upon the environment, the haptic channel is a bi-directional (input and output) communication channel that can be used to actively explore our environment and inform us about pressure, texture, stretch, motion, vibration, temperature in our surroundings. Gibson (1966) emphasized the close link between haptic perception and body movement: haptic perception is active exploration.

Haptic Technology

It has often been suggested that the use of haptic perception in human-computer interaction could lead to more natural interactions (Baecker et al., 1995; Bevan, 1995). Researchers have addressed this issue with the development and evaluation of several mechanical haptic devices (Akamatsu and Sato, 1994; Akamatsu et al., 1994; Engel et al., 1994; Kerstner et al., 1994; Massie and Salisbury, 1994; Rosenberg, 1996; Ramstein, 1995; Rosenberg, 1996). Haptic technology refers to technologies that communicate with the user via haptic feedback, which is typically evoked by applying forces, vibrations and/or motions to the user using a force-feedback device. These devices enable people to experience a sense of touch while using a hardware device such as a joystick or a mouse to interact with a digital display (Figure 1-7). They are used to simulate a wide range of object dynamics such as mass, stiffness, viscosity, textures, pulses, waveforms, vibrations and simultaneous compound effects, that provide the user with haptic feedback while interacting with a system.



Figure 1-8 An example of digital data made physical: The datafountain translates online currency rates of Yen, Euro and Dollar to waterjets. Through an internet connection the currency rate data displayed in the water jets is refreshed every five seconds (Mensvoort 2003).




Figure 1-9 The marble answering machine, incoming messages are represented by physical marbles that can be manipulated (Crampton Smith, 1995).

Ubiquitous Computing, Tangible Interaction

While several researchers seek to improve visual interfaces by adding haptic technology to the desktop, others pursue a more radical approach. It has often been suggested that the clickable atmosphere of the WIMP interface has seen its best days, thanks to speech technology, gesture recognition, etc. Numerous researchers, such as Mark Weiser (1994) and Don Norman (1999), have been calling for the end of the desktop metaphor. Following Mark Weiser's vision of the invisible computer (1991), the ubiquitous computing research field aims to integrate computation into the environment, rather than having computers as distinct objects. This paradigm is also referred to as Ambient Intelligence (Aarts and Marzano, 2003) or, more recently, Everyware (Greenfield, 2006). Promoters of this idea expect that embedding computation into the environment and everyday objects enables people to interact with information-processing devices more naturally and casually than they currently do, and in whatever place or condition they find themselves.

Mark Weiser envisioned that once computing is everywhere, this could easily lead to a restless environment. He proposed that the invisible computer should be, what he called, calm technology. A classic example of calm technology is Natalie Jeremijenko's Live Wire (Weiser & Brown, 1996), a piece of plastic cord that hangs from a small electric motor mounted on the ceiling and is connected to the area Ethernet network, such that each passing packet of information causes a small twitch of the motor. Bits flowing through the wires of a computer network become tangible through motion, sound, and even touch. Other examples are the Datafountain, an internet enabled waterfountain connected to real time currency rates, (Mensvoort, 2003, Figure 1-8) and the commercially available Ambient Orb, a multicolored lightbulb that changes color according to fluctuations of the stock market (Ambient Devices, 2005).

Building upon the vision of ubiquitous computing, Ishii and Ullmer (1997) introduced a framework for tangible interaction. Tangible interaction tries to bridge the gaps between both cyberspace and the physical environment, by coupling of bits with graspable physical objects. A classic example of a tangible user interface is Durrell Bishop's marble answering machine (Crampton Smith, 1995, Figure 1-9) in which each incoming voice message is represented by a physical marble that pops out of the machine. To listen to a message, you place the marble on the speaker. To delete the message, you recycle the marble into the machine. Since the introduction of tangible user interfaces, numerous studies in the field of tangible computing have been conducted (Harrison et al. 1998; Ljungstrand et al., 2000; Djajadiningrat et al., 2004; Van den Hoven & Eggen 2004).

1.4 THE PERSISTENCE OF EXISTING TECHNOLOGY

As a result of the rapidly increasing role of digital technology in our society, it seems obvious that computing activities will no longer be limited within one device. Computing is everywhere and has become intrinsic to our daily lives. However, the growing use of additional computing devices like smart phones, PDA's, digital camera's, GPS-trackers, RFID-readers, etc. does not necessarily imply the end of the desktop computing model.

Despite the promises of force feedback, tangible interactions and the disappearing computer, millions of people around the globe are still working behind a WIMP-based desktop computer every day. While computer chips have become smaller, cheaper, more powerful and readily available, the interface advances seem to fall behind. Mobile devices are used everywhere, but not for everything. Fingers on tiny keyboards are a major obstacle to mobile productivity. Speech recognition has not improved much in the last decade, due to human inter- and intrapersonal variations in speech and disruptive background noise, not to mention the drawback of others listening to you dictating email. Arguably, most of us will still find that much of our work is best performed in a desktop setup with a large flat screen, ergonomic keyboard and mouse. While some activities are gradually moving away from the desktop

computer environment towards a growing number of niches, others activities are being incorporated within the desktop computing paradigm; consider, for example, that watching video online has been steadily stealing market share from traditional TV viewing (BBC News, 2006). During the thirty years of its existence, the desktop has become much more than the office machine it was in the beginning: it has evolved into a versatile production / communication / entertainment device.

We seem to have more of a ‘both/and’ instead of an ‘either/or’ situation. While other devices have proven to be more suitable for certain specific tasks, the WIMP based desktop computer remains the ‘Swiss army knife’⁶ - the generic all purpose device - of our digital age (Buxton, 2001; Mensvoort, 2001). Even as the next generation of desktops like Microsoft’s Surface (Rowell, 2007), consisting solely of a table size multi-touch screen and thereby eliminating mechanical intermediaries like the mouse and keyboard between you and your computer, are expected to provide a great leap forward in collaborative computing and the exchange of digital data, they will be less practical for conducting a simple individual task like writing a letter.

In spite of its obvious drawbacks, the desktop computing model penetrated deeply into our society and cannot be expected to disappear overnight. In general, once a technology reaches a certain critical mass of social acceptance, competing technologies have to be more than significantly better to take over. A classic example of this mechanism is the QWERTY keyboard. The order of letters on the keyboard was chosen to reduce the probability that different hammers of the mechanical typewriter would get entangled. Over time, mechanical typewriters were replaced by computers. Although various alternative keyboards layouts were developed from a user-centered perspective, enabling more comfortable and faster typing, the QWERTY layout remains the dominant standard up to today.

The Failure of Force Feedback

Researchers tend to be overenthusiastic about their newly developed technologies. Over fifteen years ago, Tognazinie (1992) predicted force feedback would be implemented in the graphical user interface within three years time: “*We will undoubtedly see commercially available force feedback devices added to visual interfaces, letting users directly feel as well as see the object on the display... consider feeling the cell walls as you slip from cell to cell in your 1995 spreadsheet application.*” Although today, many force-feedback devices are commercially available in specialist- and gaming environments and commonly applied to support visually-impaired people, they never became part of the standard desktop computer setup. Some of the force feedback devices that were principally developed to enhance WIMP interfaces, e.g. the Logitech Wingman Force

6 A Swiss army knife is useful as a multi-propose device, but you would not use it daily to butter your bread, if you have the dedicated device – known as the knife – available. When having soup, you will switch to using a spoon. Likewise the load on desktop computing is relieved by a variety of specialized digital devices.

Feedback mouse and iFeel mouse, were even withdrawn from the market. Maybe feeling the cell walls as you slip from cell to cell in your spreadsheet application is not useful after all?

An explanation for this failure of force feedback to become a standard feature in WIMP interfaces might be the fact that they have never been truly integrated in the interface. Keyson (1996) already observed in 1996: *“The lack of auditory and touch information in human-computer interaction may be largely attributed to the emergence of the graphical user interface as the de facto-standard platform for supporting human-computer communication. Graphical concepts such as windows, the mouse and icons date back to the Xerox Star of the early seventies. Even the universally accepted mouse, which utilizes human motor skills, exploits primarily visual feedback. This is in contrast to the sense of touch feedback in grasping real objects in everyday life. In short, to be successful, new human interface technologies, utilizing more than the visual sense alone, will not only have to demonstrate performance gains but will also have to be integrated with existing graphical user interface styles in a compatible and consistent manner.”* Adding a layer of touch feedback to an existing interface might already be a killer application for visually impaired people, but in order to be accepted by a larger audience a more profound integration is needed. Here a vicious circle becomes apparent: A). Force feedback devices are not part of the standard computing setup, because there are hardly any interaction styles developed that utilize haptic feedback as a primary communication channel. B). There are hardly any interaction styles developed that utilize haptic feedback as a primary communication channel, because force feedback devices are not part of the standard setup.

A different approach

Taking the preceding to consideration we decided to pursue a pragmatic approach towards physical computing. We believe that, until alternative interaction models for WIMP based interfaces have been developed and socially accepted (we speak of models because we expect the successor of the desktop computer will not be one general purpose device, but rather a cocktail of various devices), WIMP will remain the standard for millions. Due to its omnipresent use, every small improvement in WIMP interfaces can effectively be considered a huge improvement in design. Therefore, we decided that the core effort of this thesis should be to improve the physicality of existing WIMP interfaces without resorting to special hardware. In the current study we introduce a novel method of simulating touch within a cursor-controlled graphical user interface. This so called optically simulated haptic feedback is evoked through active cursor displacements (Figure 1-10).⁷ This technique might be applicable in various types of graphical user interfaces. In order to limit the scope of our research, we

7 Our Active Cursor technique was first presented at the International Browserday New York 2001. (Mirapaul, Matthew, Arts Online: Innovative Webmasters Chase Fame at Browserday, New York Times, April 2, 2001).

choose to focus on the WIMP environment. We anticipate that this research will lead to a higher awareness of the potential of touch in digital applications and eventually to the development of novel haptic based interaction styles. These might possibly open a future road for dedicated haptic technologies and devices that enhance the desktop setting altogether and enable a richer paradigm of human computer interaction.

1.5 RESEARCH QUESTIONS AND GOALS

This study started with a personal fascination with the simulated ‘reality’ of the graphical user interface and a desire to enhance the materiality of this virtual environment. While working with professional force feedback devices in a specialist research environment, the thought emerged that the perception of touch was not entirely generated by the mechanical haptic device alone, but that ‘what was seen’ on the screen played a role in the perception as well. Was the perception of force feedback not entirely mechanical, but in fact already partly optically induced? The idea that tactile effects could be evoked through an optical illusion alone was inspired from Renaissance painters, who already centuries ago invented illusionary techniques like perspective and *trompe d’oeil* to increase the presence of their paintings (discussed in detail in chapter 2).

The general questions of this dissertation are:

1. Can a perception of touch be evoked visually? How can such an optical simulation of haptic feedback be implemented in a WIMP-based interface?
2. How does optically simulated haptic feedback perceptually compare to mechanically simulated haptic feedback?
3. Can optically simulated haptic feedback increase the usability of graphical user interfaces?
4. How can optically simulated force feedback be applied in interface design by non-programmer interaction designers?
5. What is the expected applicability of optically simulated haptic feedback?
6. How does this application of a simulated experience relate to other developments taking place in our society? What is the larger role of simulations in our current society?

1.6 TERMINOLOGY


Within haptic- and human computer interaction literature, the terminology has been somewhat fluid over time. Haptics is often used as a catchall term to cover a variety of distinct sub-types, including *proprioceptive* (general sensory information about the bodily position and relative positions of neighboring bodyparts), *vestibular* (the perception of head motion), *kinaesthetic* (the feeling of motion in the body), *cutaneous* (sensory information from the skin), and *tactile* (the sense of pressure experienced through the skin) (Oakley et al., 2001). The term *haptic feedback* can refer to various

types of input on the haptic sensory modality: e.g. pressure, texture, stretch, motion, vibration, temperature or combinations. In the context of computers, haptic feedback can refer to the simple feel of pressing buttons on a keyboard to more sophisticated forms of force feedback by mechanical devices.


When the first computer-controlled mechanical force feedback devices were invented, the term ‘force feedback’ was still reserved for direct haptic feedback, resulting from direct contact between the human body and some object in the physical environment – consider, for instance of the pressure of a steering wheel when operating a car. At that time, device generated haptic feedback was referred to as ‘simulated force feedback’ or ‘virtual force feedback’, emphasizing that the device *simulated* a haptic sensation; not the real thing, but a surrogate. With the acceptance of such mechanical force feedback devices, the adjectives ‘simulated’ and ‘virtual’ were dropped in literature. Nowadays, when discussing ‘force feedback’, one usually means the haptic feedback generated by a computer controlled mechanical device. In literature and popular language, these devices are usually described as ‘force feedback devices’, ‘haptic devices’, or sometimes redundantly as ‘haptic force feedback devices’. Although one might have expected the term ‘mechanical force feedback devices’, this is only rarely used.

Regarding techniques that aim to evoke haptic percepts by optical means, various terms have been suggested: sticky icons (Worden, 1997), simulated force feedback (Mensvoort, 2002), pseudo haptic feedback (Lecuyer, 2001, 2004), force fields (Ahlström, 2006), gravity (Park et al., 2006). Although terms like ‘sticky’, ‘force’ and ‘gravity’ are easy to understand, they overlook that the haptic perception is *simulated*. Terms like ‘simulated’ and ‘pseudo’ are more precise in this regard, but still lack for not specifying the means of the simulation. Lack of precise terminology becomes especially problematic when techniques are compared across modalities. Given that in the current study haptic feedback is simulated both mechanically as well as optically, we need to use a terminology descriptive enough to distinguish between the two techniques. In order to meet this requirement, we speak of ‘mechanically simulated haptic feedback’ and ‘optically simulated haptic feedback’ (Figure 1-10). The adjective ‘simulated’, which we use according to Oxford American dictionary as “*to imitate the appearance or character of*”, is added to emphasize that the applied techniques are reproductions. This emphasizes that, although that both optical and mechanical techniques can be used to simulate haptic feedback, they are only capable of reproducing a portion of the haptic spectrum; they do not have the same sensory richness of unmediated haptics (think of an embrace or a kiss). We chose this terminology because it precisely and transparently describes the technique: haptic feedback is simulated by mechanical/optical means.


Occasionally, in later chapters, we abbreviate these terms to the shorter ‘haptic force feedback’ and ‘visual force feedback’, describing the technology from the user’s perspective: force feedback is *experienced* by the user via the haptic- or visual sensory modality. This terminology complies best with existing jargon used in literature as well as popular language.



Haptic exploration of a bump shaped object, directly with the hand.



Haptic exploration of a bump is simulated via a mechanical device. The quality of the simulation is defined by the expressiveness of the force feedback device.



Haptic exploration of a bump is simulated via active cursor displacements. The quality of the simulation is defined by the strength of the optical illusion, evoked by the interactive animations.

Figure 1-10. The haptic experience of actively exploring a slope (a) can be simulated both mechanically (b) and optically (c). The mechanical technique simulates the bump shape via force feedback, asserted by a mechanical device, which the user straightforwardly senses and perceives via the haptic sensory modality. The optical technique simulates the haptic feedback via an optical illusion, which is evoked by displacing the cursor on the screen *as if* there are forces asserted on the mouse. In Chapter 3 both techniques are experimentally compared.

1.7 OVERVIEW OF THE CHAPTERS

The prologue aims to embed the research in a broader social and cultural context by reflecting upon the role of simulations in our society at large. In Chapter 1 we specifically introduce the subject matter and define the scope of our research. In Chapter 2 we work towards a first design. We introduce optically simulated haptic feedback and describe its basic implementation. In the following two chapters we empirically test the technique developed in Chapter 2 in comparison with haptic feedback as generated by a mechanical force feedback device. In Chapter 3 we compare the perceptual experience and in Chapter 4 we compare the usability in a pointing task of both types of haptic feedback. In Chapter 5 we describe a software prototyping toolkit that enables designers to create novel interaction styles using visual force feedback and describe the possibilities for new interaction styles. In Chapter 6 we draw conclusions and discuss future directions. Finally, the thesis returns to the larger social and cultural context and concludes with a short epilogue containing a number of philosophical reflections – already initiated in the prologue – and a vision towards the future. An overview of the various types of activities in the chapters is depicted in Figure 1-11.

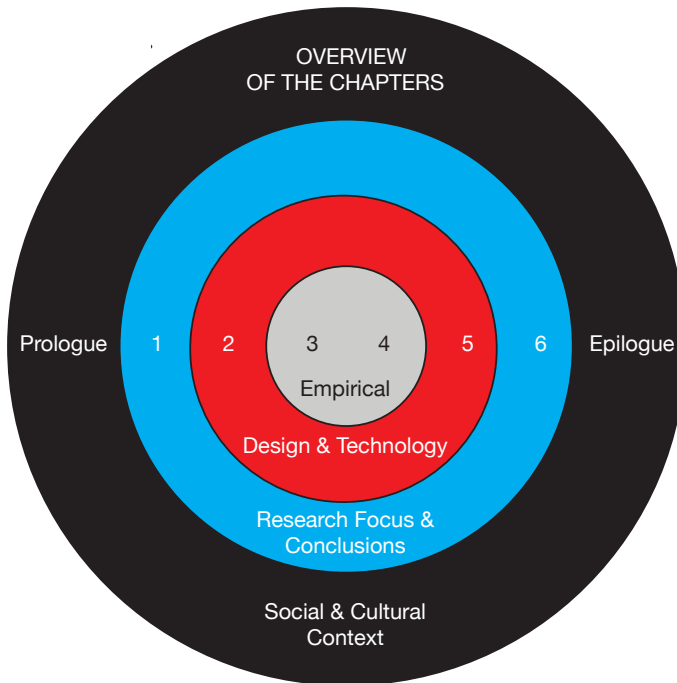


Figure 1-11 Overview of the type of activities conducted throughout the chapters. From left to right the chapter numbers are displayed, showing how the thesis from a reflection on the larger social/culturally context zooms in on concrete design, technological and empirical work, in order to return to the larger context in the final chapters.

2

Optically Simulated Haptic Feedback

In this chapter⁸, we present an approach to design a physically richer user interface without taking resort to special haptic input/output devices. We will show that interactive animations can be used to simulate the functioning of force-feedback devices. Renaissance painters invented various techniques to increase the presence of their paintings. We aim at doing similar work for the contemporary graphical user interface. We discuss the use of interactive animations towards a richer and more physical interface. The role of movement in interactive applications is still underestimated. In the early days of graphical user interfaces, the use of interactive animation was cost inefficient because of the scarce processing power. Nowadays, interactive animations can be implemented without significant performance penalty. Whereas animation of independent objects is properly studied and applied in motion cinema, only a few researches focused on animation in direct interaction with a user. We designed and implemented a series of experimental interaction styles that manipulate the cursor position to communicate with the user. By applying tiny displacements upon the cursor's movement, haptic sensations like slickness, pressure, texture or mass can be simulated. Optically simulated force feedback exploits the domination of the visual over the haptic domain. This perceptual illusion of touch will be experimentally tested in detail in chapter 3.

8 This chapter is an extended and updated version of Mensvoort, K. van (2002) What you see is what you feel: exploiting the dominance of the visual over the haptic domain to simulate force-feedback with cursor displacements, Proceedings of the conference on Designing interactive systems: processes, practices, methods, and techniques, June 25-28, 2002, London, England.



Figure 2-1 From left to right: Egyptian, Greek and early medieval paintings were predominantly symbolic. The figures on the paintings are not drawn to exactly resemble the things they represent, rather they function as language elements in the visual story told by the painting.


2.1 RENAISSANCE TRICKS

If we compare the computer screen with the renaissance canvas, the limitations and possibilities show some remarkable similarities. Both painters and interface designers are constrained to a flat and rectangle canvas. Their goal is to represent or reflect our rich world of experiences and sensations within these limitations. Pre-renaissance paintings were in general symbolic; the objects on the paintings do not *look like* the things they represent (Figure 2-1). In the renaissance presence⁹ gains importance, paintings aim to *reflect* reality on the canvas.


⁹ Presence has been a subject of discussion throughout centuries. Recently, virtual reality researcher Lombard & Ditton (1997) defined presence as ‘the perceptual illusion of non-mediation.’



Figure 2-2 Mathematical Perspective: Niccolò da Tolentino Leads the Florentine Troops. Paolo Uccello 1450s, Tempera on wood, 182 x 320 cm



Material expression (van Eyck)



Sfumato (Da Vinci)



Object shading (Giotto)


Figure 2-3 Examples of painting techniques invented by Renaissance painters, to increase the expressiveness of their paintings. From left to right: Material Expression (Van Eyck), Sfumato (Da Vinci) and Shading (Giotto).

Renaissance painters invented various techniques to capture the three-dimensional world on the flat painting canvas: shading, perspective, sfumato, trompe d'oeil and material expression. At first, these techniques were developed separately. For instance, in the battlefield painting by Uccello the mathematical perspective is applied within the finest detail, but the characters still have a rather simplistic appearance similar to the medieval paintings (Figure 2-2). This style resembles some of the early computer 3d renderings from the eighties; perspective is modeled well, but material expression is lacking. In the same period, Giotto started adding shadings to model his characters within space. The Van Eyck brothers enhanced the presence of the landscape with detailed material expressions. The Mona Lisa was provided with her mysterious look using sfumato – which in Italian means smoky –, a painting technique blends of colours so subtly that there is no perceptible transitions. (Figure 2-3).

At the peak of the renaissance period, Leonardo Da Vinci combines the different painting techniques in his masterpiece ‘The Last Supper’ (Figure 2-4). Perspective, material expression, sfumato and trompe d’oeil were applied to give the visitors of the dining room of Santa Maria delle Grazie in Milan a virtual reality like experience – avant la lettre – of dining together with Jesus and his apostles. Leonardo tried to ‘extend the room’ by means of trompe d’oeil, a technique in which the perspective in the painting is ingeniously devised as an expansion of the perspective of the space in which it is set, to make it look like Jesus and his apostles were sitting at the end of the dining hall (Kobovy 1988).



Tromp d’oeil



Mathematical Perspective

Figure 2-4 Leonardo Da Vinci applied trompe d’oeil, mathematical and atmospheric perspective to enhance the presence of ‘The Last Supper’ in the dining room of Santa Maria delle Grazie in Milan (The last Supper, Leonardo Davinci 1495-1497, tempera on gesso, pitch and mastic, 460 × 880 cm).

Arguably as a result of automated imaging techniques like photography and film that emerged in the last two centuries, painting has developed itself in a different direction. Away from the visual realism, which reached its summit in the Renaissance, towards non-photographable styles like impressionism, cubism, abstractionism, hyperrealism and surrealism. Since then, techniques aiming to enhance the realism of visual representations have mostly been developed further in other media than painting. Especially film has proven to be a highly effective immersive medium; according to the classical anecdote, people ran out of the cinema when the Lumiere brothers projected their film ‘l’arrive du train’ featuring a single shot of a train arriving at the station (Lumiere 1895). In the next section we describe how visual realism has developed itself into today’s computer interfaces.

2.2 ANIMATED GRAPHICAL USER INTERFACES

Early computer interfaces were command-line driven. Users had to learn codes and commands to control the system. With the introduction of Graphical User Interfaces (GUI’s) together with the development of the mouse (English, Engelbart, Berman 1997), the transition from command manipulation to direct manipulation was made.

Direct manipulation permitted novice users to access powerful facilities without the burden of learning to use a complex syntax and lengthy lists of commands. Direct manipulation involves three interrelated techniques (Shneiderman, 1983):

1. Provide a physically direct way of moving a cursor or manipulating the objects of interest.
2. Present a concrete visual representation of the objects of interest and immediately change the view to reflect operations.
3. Avoid using a command language and depend on operations applied to the cognitive model which is shown on the display.

In the first graphical user interfaces the movements of the objects on the screen were abrupt and unrefined. Use of animation was cost inefficient because of the scarce processing power. After the successful application of animated computer visualizations and with the increasing processing power of computers, use of animation techniques as a means of making the interface easier to understand and more pleasant to use came within focus.




Figure 2-5 Sequential Animation Drawings from a Mickey Mouse Anniversary Promo.

Learning from Disney Animators

Many of the principles of traditional animation were developed in the 1930's at the Walt Disney studios. These principles were developed to make animation, especially character animation, more realistic and entertaining. Cartoon Animators use a broad range of effects to enhance the illusion of the animation. Often, animators mimic physical effects, such as inertia and friction, to reinforce the illusion of substance (Laybourne 1979). These basic animation techniques are still applied in today's computer generated animations of Disney and Pixar. They also made their way into applications of computer visualization outside the entertainment realm.



Figure 2-6 Walt Disney with his main character Mickey Mouse. © Disney Corp.

In their paper “Animation: From Cartoons to User Interface”, Chang and Ungar (1993) list three principles from Disney animators Thomas & Johnston (1981) that apply to interface animation: solidity, exaggeration, and reinforcement. They can be characterized as follows:

1. Characters and objects should seem *solid*.
2. *Exaggerating* the behavior of objects makes the user interface more engaging.
3. The interface should *reinforce* the illusion of reality.

Principles of traditional animation were first applied to 3D computer visualization, as suggested by Lasseter (1987). Robertson (1991) showed that animated 3D-visualizations can shift some of the user’s cognitive load to the human perceptual system, where it can be subconsciously processed which effectively reduces the cognitive load of the user. Bederson et al. (1999) examined how animating a viewpoint change in a spatial information system affects a user’s ability to build a mental map of the information in the space. They found that animation improves users ability to reconstruct the information space, with no penalty on task performance time (Benjamin, et al., 1999). It has also been suggested that animation in user interfaces improves decision making. Gonzales (1996) investigated the relative effects of images, transitions and interactivity styles in animated interfaces and found that subjects performed better with animated interfaces based on realistic and smooth rather than abstract and abrupt images. Use of animated icons for 2D graphical user interfaces was suggested by Baecker (1991), but these early animated desktop icons were distracting because they were always running, resulting in a blinking screen of ten or twenty canned animations going on simultaneously on the desktop. Motion is known to be an attention grabbing phenomenon (Lu & Sperling, 1995). Having a screen full of motion only needlessly distracted the user from his task. In the next paragraph we see that connecting animations to direct object manipulation, is more beneficial.




Figure 2-7 When clicked, the cursor is glued to the stickybutton. The user has to pull the mouse to release it.

Animate the manipulated object

If judiciously applied, the techniques of cartoon animation can enhance the illusion of direct manipulation that many human computer interfaces strive to present. In particular, animation can convey physical properties of the objects that a user manipulates; strengthening the sense that real work is being done. Various people experimented with interactive animations as a means of making the GUI more tactile.




Figure 2-8 An interface which behaves based upon the characteristics of fluidity and elasticity (Maurer, 2001).




Figure 2-9 An animated menu (Thomas & Calder 2001).




Figure 2-10 Enhance the illusion of direct manipulation by deforming objects as they are manipulated (Thomas & Calder 2001).

Thomas and Calder (1995) suggested some techniques that application programmers can use to animate direct manipulation interfaces. Their approach is based on suggesting a range of animation effects by distorting the view of the manipulated object.

Our work on interactive animations started with an experiment with a stickybutton (Mensvoort, 1999). A seemingly normal button that, when clicked, turns the cursor into gum. The user has to pull the mouse to release it (Figure 2-7). Similar work was done by Ording who developed prototypes of 3D buttons with a 'rubbery feel' (Poppe, 1999). Maurer (2001) prototyped an experimental interface (actually, more of a thought provoking performance than a functional interface) in which 'the logic of the material overrides the logic of the system'. Based upon principles of liquid material Maurer transforms the familiar desktop into an elastic experience (Figure 2-8).

In later research, Thomas & Calder (2001) extend the visual feedback for direct manipulation interfaces, by smoothing the changes of interface components (Figure 2-9), animating manipulated objects (Figure 2-10) and providing cues that anticipate the result of a manipulation. They also show the effects to be effective and enjoyable for users. Recently Agarawala & Balakrishnan (2006) experimented with virtual desktops that behave in a more physically realistic manner by adding physics simulation and using piling instead of filing as the fundamental organizational structure. In BumpTop (Figure 2-11), an experimental pen-based virtual desktop, objects can be casually



Figure 2-11 Two recent applications of interactive animations in screen based interfaces. On the left: BumpTop (Agarwala and Balakrishnan 2006), on the right: Apple's iPhone (2007).

dragged and tossed around, influenced by physical characteristics such as friction and mass, much like one would manipulate lightweight objects in the real world.

More recently, Apple (2007) in its iPhone and iTouch uses fine tuned interactive animations and well orchestrated transitions to increase the quality of the interaction. The devices are equipped with a multi-touch screen and a tilt sensor which are used to operate the device through physically richer interactions like sliding, pushing and shaking besides the traditional button clicks known from earlier computer interfaces. Figure 2-11 shows an example of how the iPhone provides its users a nearly physical experience of browsing through their virtual record collection.

2.3 INTRODUCING THE ACTIVE CURSOR TECHNIQUE

Within computer interfaces based on windows, icons, menus and pointing (WIMP), the cursor is one of the most important assets. It is the representation of the user within the interface. The point/click task is the primary operation in WIMP interfaces.¹⁰ The cursor channel is used intensely in the interaction with the system. Within the reality of the desktop computing metaphor the cursor represents your body. Marshall McLuhan described people's tendency to extend their identities in animate objects, when interacting with them (McLuhan 1964). For instance, when driving a car the vehicle becomes an extension of our body. It absorbs our sense of identity and if one car hits another, the driver of the vehicle being struck is more likely to say: 'Hey! You hit *me*!', than 'You hit my car' or 'your car hit my car', to be accurate. Likewise, in desktop computing, we narrow down our bodies into the tiny arrow of the computer cursor.

According to Heidegger (1927), tools should be understood as connections or linkages between humans and reality. Heidegger indicated the way in which tools are present to human beings when they are used as "readiness-to-hand". Tools that are

¹⁰ As an informal experiment, we distributed a program that counts the mouse clicks among colleagues, finding that the average user clicked well over a 1000 times in one working day.

used for doing something typically withdraw from peoples attention: the attention of, for example, a person who drives a nail into a wall is not directed at the hammer but at the nail. A persons involvement with reality takes place through the ready-to-hand artefact. Only when it breaks down does it ask attention for itself again. The artefact is than, in Heidegger' words, "present-at-hand" and is not able to facilitate a relationship between a user and his or her world anymore.



Figure 2-12 Changing the cursor icon has proven to be effective and intuitive way to communicate properties of the system (Muller, 1988)

Despite its important role in WIMP interfaces, cursor behavior has not altered much since its invention by Engelbart in the sixties (Engelbart, 1968; English, 1967). An early improvement is the use of a dynamic cursor icon to inform about the status of the system or the effects of the next mouse action (Figure 2-12). Changing the cursor icon to an hourglass, hand or I-beam has proven to be effective and intuitive (Muller, 1988).


Another approach to enhance navigation in GUI's is to focus on the input device. In chapter 1 we already discussed how various force-feedback devices were developed in order to introduce touch feedback into digital interfaces (Brooks et al., 1990; Akamatsu et al., 1994; Engel et al., 1994; Akamatsu et al., 1995). Touch can play a unique role in communication -unparalleled by words, sound or images-, and it has often been suggested that improvements in this domain could lead to richer and more natural computer interfaces (Baecker et al., 1995; Bevan, 1995). Force feedback devices are used to simulate a wide range of material object properties such as elasticity, hardness, stiffness and textures. Although force-feedback devices are commercially available, they have not become part of the standard desktop set-up. Not much software is, therefore, developed that utilizes direct haptic feedback as a primary communication channel; haptic feedback remains an 'add-on' for existing interfaces. Knowing there are potential benefits of force feedback devices, we raised the question whether it would be possible to simulate the functioning of force feedback in a standard WIMP GUI setting. This research question resulted in the development of the active cursor technique, as described in the remainder of this section.

Two-way Communication via Cursor Position


As discussed previously, the cursor behavior plays an important role in the communication between the user and the system in the graphical user interface. In her book, *Computers as Theatre*, Laurel (1991), argues that both the computer and the human are active agents working together to achieve some common goal. It is the goal of the designer to facilitate these two active agents in their effort to collaborate.



One-way communication (standard setup without haptic feedback)



Two-way communication (mechanically simulated haptic feedback)



Two-way communication (optically simulated haptic feedback)

Figure 2-13 Communication dialog for WIMP interface setup, force-feedback setup and active-cursor setup. In the last setup force feedback is simulated visually.

Note that only the active communication channels are depicted. Passive feedback, like the linear placement of the cursor according to the users movement which is also perceived by the user, takes place in all cases and is omitted in the figure. See also Figure 1-10.

In the standard WIMP interface setup, the cursor position is used for input only. We expected that it might be beneficial to use the cursor not only as an input channel for the human, but also as an output channel for the computer. We developed a cursor interface in which the system manipulates the cursor position to give feedback to the user. The user still has main control over the cursor movements, but the system is allowed to apply tiny displacements to the cursor position. This system has a lot in common with existing force-feedback systems, except for the fact that in force-feedback systems the location of the cursor is manipulated as a result of the force sent to the haptic display (force-feedback mouse, trackball, etc), whereas in our system the cursor location is directly manipulated (Figure 2-13; Figure 1-10). Since direct two-way communication through the pointing device has proved successful with haptic devices, it seems reasonable to expect benefits from direct communication through cursor positions.

The active cursor displacements result in interactive animations that induce haptic sensations like stickiness, stiffness, or mass. The cursor is displaced as if there are real forces working on the mouse. The disparity between the visual feedback, i.e. slowing down of cursor on screen, and the user applied motion of the input device to compensate provides an optical illusion of haptic feedback. The domination of the visual over the haptic domain induces the illusion that the input device experiences a force in the direction of these additional cursor displacements. We know that humans tend to integrate multi-modal sensations into single meaningful events in

the external world (Gibson, 1966). The user tends to ‘feel’ what he ‘sees’. Similar to tricks and techniques developed by renaissance painters and cartoon animators with the intent to enhance the expressiveness of a limited medium, like the painting canvas or cinema, our technique can be applied to enhance the expressiveness of WIMP interface. Haptic percepts like slickness, pressure, texture, or mass can be evoked within the constrained mouse operated computer screen. Apparently, haptic feedback can be simulated optically. The optical illusion at work here is described further and empirically investigated in chapter 3.

Contextual feedback

The active cursor displacements can be used to create various (dynamic) slopes as well as textures and material properties, which can provide the user with feedback while navigating through the screen. In this section, we describe the first design ideas related to the active cursor technique. The application domain is developed further in the remainder of this study (chapter 5 and 6) and is extensively discussed in section 5.4.

‘Holes’ and ‘bumps’ were among the first virtual haptic objects we created. When the cursor moves over a hole, it is dragged towards the centre. When moving over a hill, the cursor is dragged away from the centre (Figure 2-14). Due to these cursor displacements a hole becomes an easily accessible part of the screen whereas a hill area is hard to access. Such contextual feedback seems to communicate in an immediate and intuitive way. It is possible to create various slopes as well as dynamic slopes and textures to give contextual feedback to the user (Figure. 2-15).




Figure 2-14 Active cursor simulations of a hole and a hill




Figure 2-15 From left to right: Pocket, Tunnel, hillbuttons (hard to access) holebuttons (easily accessible) and an inaccessible holebutton (occupied by the red dot).

If applied carefully, active cursor displacements might be used to create a dialog between the user and the system, e.g. by guiding the user towards preferred positions or discouraging to go to unadvised locations. While thinking about what we could learn from touch interaction in everyday life, the aspect of navigation emerges as a central area: Knowing where you are, where you came from, where you could go. Figure 2-16 shows a decision graph, which consists of small holes, and gutters that push the cursor in one direction. At every node in the graph the cursor is stopped, because of the hole-shaped force field underneath. The user can choose where to go next with a small mouse movement, after which the system conveys the cursor towards the next node. Dialogs like these could be helpful in guiding a user through a decision dialog, for instance a form or an installation wizard.

Inspired by the material expression in renaissance paintings, we have also experimented with material and texture expression using active cursor displacement. Figure 2-17 shows a few of the textures we have simulated. For instance, the sand structure slows the cursor down like you would expect if you walk on a beach bare feed. The ice texture adds a slippery effect to the cursor behavior. Simple textures can be combined in order to simulate more complex textures and slopes.




Figure 2-16 Cursor is guided through the decision graph. Every black dot represents a decision point.

Figure 2-17 Simulation of textures through cursor displacements.




Figure 2-18 Mapping properties of three dimensional objects




Figure 2-19 Helloworld. Touch & manipulate a 3D globe. (<http://www.koert.com/work/helloworld>)

Another use of the technique could be to convey properties of materials or 3D objects to the user. We have experimented with expressing material properties of three dimensional objects through cursor displacements. In the physical world, a wooden cube would have a certain weight. If one would want to flip it over its side you will have to apply a certain force (Figure 2-18). This property of the material can be simulated by pushing the cursor downwards into the direction of the gravity. Likewise with the ventilator, one would expect to be pushed away if one would touch it. Active cursor displacements are applied to meet this expectation (Figure 2-18). In Figure 2-19, the 3D shape is rotated according to the position of the cursor on the shape. In addition to this, the size and angle of the cursor is manipulated according to the active cursor forces.

Algorithm

The algorithm that we implemented is in many respects analogous to known straightforward ways of implementing force feedback. Instead of expressing force feedback through a specialized device the force vectors are expressed through displacements on the cursor position (see Figure 2-13). These cursor displacements are calculated every display refresh loop based on the force fields defined underneath the cursor position. Just as with real force-feedback devices, it is possible to create any 3D slope as well as dynamic slopes and textures. Also, simple shapes and textures can be combined to compose complex structures.

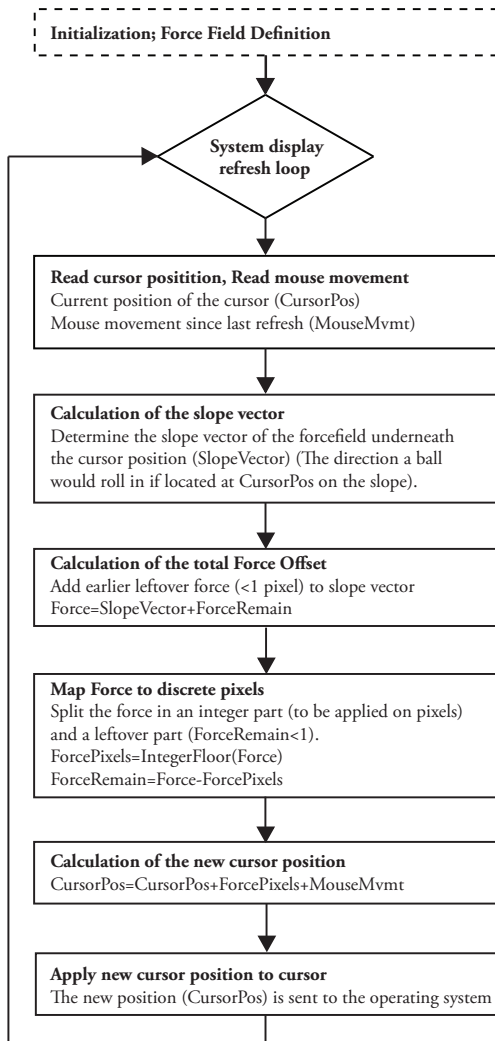


Figure 2-20 The algorithm used to simulate force-feedback.

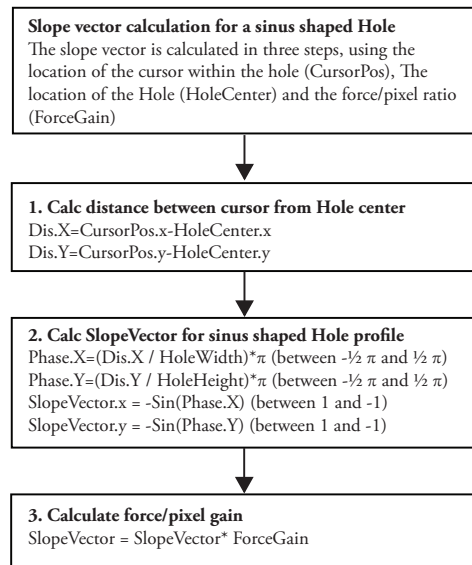


Figure 2-21 algorithm used to calculate the slopevector for a hole.

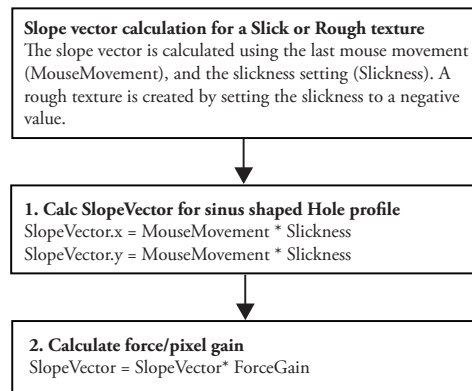


Figure 2-22 algorithm used to calculate a Slick or Rough texture.

Our algorithm is shown in Figure 2-20. A virtual surface is defined and for each system display refresh, the algorithm computes the force that would have been applied on the input device as if it were force-feedback enabled. Depending of the slope of the virtual surface beneath the cursor position this force, the slope vector (Figure 2-21), is calculated. The displacement induced by the simulated force is then added to the cursor movement and rounded to integer pixels. The rest displacement, <1 pixel, is stored for future calculations, which enables the algorithm to translate very small forces into discrete pixel displacements.

The slopevector can be used to calculate slopes like bumps, holes, gutter, slope, etc. It can also take the users mouse movement as an input and actively exaggerate or diminish the mouse movement made by the user resulting in slick or rough areas. Furthermore, the slopevector can also be used to create inaccessible areas. Figure 2-21 shows the calculation of the slope vector for the hole-type structures. Figure 2-22 describes the calculation of a slick or rough texture. Various functions can be combined resulting in complex shapes and behaviors. In chapter 5, we discuss the force fields we have generated in more detail.

2.4 DISCUSSION

Inspired by renaissance painters, who centuries ago already applied various types of optical illusions in order to enhance the expressiveness of their paintings, we invented a technique to optically simulate haptic feedback within a standard graphical user interface. The functioning of a mechanical force feedback device is simulated, with active cursor displacements. These active cursor displacements, influencing the normal cursor movement linked to the users mouse movements, can be applied to generate various slope, texture and force perceptions with the user. This sense of touch is an illusion, which is presumably based upon the domination of the visual over the haptic sensorial modality. This optical illusion of touch is empirically investigated in chapter 3 and 4.

Our method was developed for use with a standard mouse, but should work on any cursor-controlled interface. Optically simulated haptic feedback can be used to display direct contextual feedback to a user. This seems to open up a broad range of interface design possibilities. Contextual feedback through cursor displacements may inspire designers to create a new type of interaction styles.

Before (optically simulated) force feedback can be fully applied in more complex interaction styles, an expressive language of satisfactory and tolerable active cursor behaviors needs to be developed. Interface designers and researchers need to experiment more with the technique in order to explore the affordances and find out what works and what does not. This issue is addressed further in chapter 5, in which a toolkit is presented aimed at designing interfaces with the technique, and in section 5.4 which discusses the expected application domain in further detail.

The active cursor technique is not to be expected to replace mechanical haptic

feedback altogether, since it can be applied only in combination with a visual display and thus will not work for visually impaired people. Rather, we expect that the ability to employ haptic interaction styles in a standard WIMP interface might instigate the acceptance of haptic devices.

3

Measuring the illusion: Perception of optically and mechanically simulated Bumps and Holes

In the previous chapter, we have introduced optically simulated haptic feedback, a technique to evoke a perception of haptic feedback without resorting to special mechanical force feedback devices. The operation of the force feedback device is substituted by tiny displacements on the cursor position relative to the intended force. Apparently, the visual domain dominates the haptic domain in this situation, and this induces the illusion that the input device exerts a force in the direction of these additional cursor displacements. In this way haptic percepts like stickiness, touch, or inertia can be evoked.

In the current chapter¹¹ we investigate the perception of optically simulated haptic feedback. The perception of optically and mechanically simulated bumps and holes was tested experimentally. Results show that people can recognize optically simulated bump and hole structures, and that active cursor displacements influence the haptic perception of bumps and holes. Depending on the simulated strength of the force, optically simulated haptic feedback can take precedence over mechanically simulated haptic feedback and also the other way around. When optically simulated and mechanically simulated haptic feedback counteract each other, however, the weight attributed to each source of haptic information differs from user to user. It is concluded that active cursor displacements can be used to simulate the operation of mechanical force feedback devices.

11 This chapter is based on Mensvoort van, K. Hermes D.J., Vos, P., Liere van, R. 2009. Perception of optically and mechanically simulated bump and holes. Accepted for publication by *Transactions on Applied Perception*.

3.1 RELATED WORK

Numerous studies on human perception indicate that stimuli in one modality can evoke percepts in another (Marks, 1978; Stein, 1993; Welch, 1986). We know that humans exhibit distal attribution, which is the tendency to quickly integrate multi-modal sensations into single meaningful events in the external world. Gibson (1966) describes our senses as active interrelated systems providing information for our perception of the real world. Whereas most classical frameworks of interfaces between perception and action rely on separate coding (Massaro, 1990), plenty of evidence from experimental psychology and psychophysics indicate that perception and action share a common computational code (Prinz, 2005). Building upon work by Sperry (1952), who argued that the perception–action cycle is the fundamental logic of the nervous system, common coding theory claims that perception and action processes are functionally intertwined: perception is a means to action and action is a means to perception (Prinz, 1984).

It is well known that vision can influence haptic perception (Heller et al., 1999; Klatzky et al., 1987; Lederman et al., 1986; Rock and Victor, 1964). A classic and robust example of visual-to-haptic intersensory interaction is the size-weight illusion, documented by Charpentier (1894) and Flourney (1891) over 100 years ago (Murray, 1999). When lifting two objects of different volumes but equal weights, people judge the smaller object to be heavier. In this example, haptic feedback still plays a role, since the volume of the object is not only seen, but also felt by the hand. Runeson and Frykholm (1981) showed that an external observer, watching another person handling a heavy box, is able to infer the weight of the lifted object. They concluded that visual information passed through the optic array and representing the kinetic pattern of the movement can also play a role in extracting higher-order properties within the haptic domain. Possibly, mirror neurons play a role in such intersensory interactions. A mirror neuron is a neuron, which fire both when one acts and when one observes the same action performed by someone else; as though the observer itself were acting. Mirror neuron systems simulate observed actions, and are thought to be important for understanding the actions of other people, and for learning new skills by imitation (Rizzolatti & Craighero, 2004).

Vision is thus assumed to contribute to what is generally taken to be the privileged domain of the haptic sense combining tactile and proprioceptive cues. Carr and Lederman (1995) have demonstrated the dominance of vision over haptics in various experiments. Research by Miner (1996) demonstrates that visual stimuli can influence haptic perception in virtual environments. More recently, Ernst and Banks (2002) showed that humans integrate visual and haptic information in a statistically optimal fashion.

Application in WIMP interfaces

The active cursor technique, introduced in the previous chapter, aims at evoking haptic effects, while using a normal mouse not capable of producing any force feedback except that involved in resistance as the mouse is moving over a surface. Comparable techniques of simulating touch through manipulation of the graphical element that represents the user have been intuitively applied earlier in videogames. For example, in the classic racing game *Outrun* (Suzuki, 1986) the players must, when the road bends, exert force on their input devices to keep the car in the middle of the road. This effect provides the players with the sensation of being “pushed” off the road.

In WIMP based graphical user interfaces, an early application of what in retrospect can be considered as optically simulated haptic feedback is the use of *sticky icons* introduced by Keyson (1997) and Worden (1997). With sticky icons the cursor’s control/display ratio, which determines the mapping between the physical mouse movement and the cursor movement on the screen, is reduced as the cursor enters a target, and then returns to normal after passing the target. Inside the target, equal mouse movements result in smaller cursor movements due to the change in cursor gain. In this way, the cursor speed diminishes, when the user enters a target, though keeping the mouse speed constant. Like with the active cursor technique (Mensvoort, 2002) this reduction effectively results in an enlargement of the motor space underneath the target, while the visual space remains unchanged. The decoupling of motor space and visual space induces the ‘sticky’ feeling. Ahlström (2002) suggested that the cursor gain technique, manipulation of cursor gain, could also be used to simulate more complex slopes like holes and hills and that these could be applied to guide a user in a graphical user interface. Lécuyer et al. (2004) conducted a perceptual experiment confirming that subjects indeed could identify various types of bumps and holes by seeing the variation of the cursor gain.

Various experiments have been conducted to assess the benefits of cursor manipulation techniques in WIMP interface. Keyson (1997) compared mechanical force feedback, consisting of a pulling force towards the centre of a target, with sticky targets, consisting of a reduction of the cursor gain within the target. The results of the experiment showed that target acquisition performance was generally higher in the tactile-feedback condition, followed by cursor-gain feedback, and then normal cursor control. In research by Worden (1997), the sticky icons had no effect on accuracy, but substantially improved the speed of performance over the traditional pointer. Older users especially benefited from the adaptive technique. Given the pervasiveness of pointing in graphical interfaces, every small improvement in the target-acquisition task, represents a substantial improvement in usability. Blanch et al. (2004) formalized the cursor gain technique and showed its performance in a pointing task is given by Fitts’ index of difficulty in motor rather than visual space. Baudisch et al. (2005) showed the benefits of the cursor gain technique in a snapping task. The benefits of the

active cursor technique have been established experimentally in targeting tasks (Park et al., 2006; Mensvoort et al., 2008) and steering tasks (Ahlström, 2005).

While pointing in the physical world is governed by Fitts' law and constrained by physical laws, pointing in the virtual world does not necessarily have to abide by the same constraint (Balakrishnan, 2004). Both the cursor gain and the active cursor technique are aimed at decoupling the visual space and motor space through cursor manipulation. The active cursor technique differs from the cursor gain technique in that the direction of adjustment is not necessarily parallel to the direction of the mouse movement. Although the cursor gain technique (Keyson, 1997; Worden, 1997; Ahlström, 2002; Blanch et al., 2004; Lécuyer et al., 2004; Baudish et al., 2005) is easy to implement, it only works when the mouse is being moved by the user and is limited to the direction of the users' movement. Just as for mechanical force feedback devices, the active cursor technique (Mensvoort, 2002; Ahlström, 2005; Park et al., 2006; Mensvoort et al., 2008) also works when, the user is not moving the mouse. From a mathematical perspective, the manipulation of cursor gain is a more restricted way to simulate force feedback than actively displacing the cursor; the active cursor algorithm can, by constraining the cursor displacements within the vector of the user's mouse movement, be set to generate the same effects as the cursor gain technique but not the other way round. Ahlström (2006) compared the active cursor technique (force fields) with the cursor gain technique (sticky targets) in two realistic pointing situations which involve several closely placed targets and found that the force fields improve pointing performance and that the sticky target technique does not. However, this does not mean that the active cursor technique's more realistic way of simulating force feedback will in all cases lead to advantages for the user; having the cursor move without the user's action may also lead to drawbacks in some cases. Further research, which falls outside the scope of the current study, is required in this regard. This chapter focuses on a comparison between optically and mechanically simulated force fields.

In the current study we assess the perception of bumps and holes, generated with active cursor displacements in comparison with the perception of bumps and holes generated by a mechanical force feedback mouse device. The experimental setup was inspired by the work of Robles-De-La-Torre and Hayward (2001), who assessed the perception of bumps and holes and, by using an ingenious robotic device, were able to uncouple force feedback from geometric information. Through this uncoupling they were able to generate stimuli with the geometric properties of a hole, but with the force feedback properties of a bump, and *visa versa*. Think for instance of a contradictory structure that, when explored geometrically by moving your finger, makes your finger go down just as it would when the finger passes into a hole, but at the same time pushing it away from its centre with force feedback, just as it would when a bump structure would be passed. They found that force cues — not geometric cues — determine perceived shape. Likewise, we in our study uncouple the visual force information, i.e. cursor displacements seen on the screen, from the haptic force information generated by the mechanical haptic device. In a factorial design we compared the perceptual

effect and interaction of both types of simulated haptic feedback, and determined how optical and mechanical haptic feedback independently and jointly contributed to this topographical experience.

3.2 HYPOTHESES

Our first hypothesis is that optically simulated haptic feedback can be applied to create perceivable 'bump' and 'hole' structures and that people are able to judge the heights of the bumps and the depths of the holes to an extent comparable to that obtained with mechanically simulated haptic feedback.

Our second hypothesis is that optically simulated haptic feedback can be used to enhance or decrease the perceived height of 'bump' and 'hole' structures generated with a mechanical force feedback device. In other words, our second hypothesis states that optically simulated haptic feedback can influence the perception of mechanically simulated haptic feedback.

3.3 METHODS

Subjects

Thirty volunteer subjects participated. There were 19 male and 11 female subjects, ranging in age from 20 to 36 years. Of the 30 participants 23 were right-handed and seven were left-handed. All subjects were regular users of mice in their daily work. The subjects were not informed about the goal of the experiments in advance. During the experiment, subjects were presented with various bump/hole structures which they could explore with the mouse in order to determine their heights/depths. The slopes of some bump/hole structures were generated through optically simulated haptic feedback, those of others through mechanically simulated haptic feedback, and some through, matching or conflicting, combinations of both techniques. The subjects were not informed about the different techniques used to generate the haptic structures. We divided the subjects randomly into three groups of ten people; each group being assigned a different combination of ranges of the mechanical and optical nominal force setting.

Apparatus


The experiment was conducted using the Logitech Wingman force feedback mouse, a mouse attached to a mouse pad replacing the mouse mat and with two motors supplying force feedback to the user (Rosenberg, 1997). This mouse was used in all experimental conditions. The host computer was a Pentium III class PC with a screen resolution of 1024x768 pixels on a 17-inch monitor. The default Windows XP cursor was used. The experiment was implemented in C++. The data were collected with

1-pixel and 1-ms resolution and saved in output files for subsequent analysis. The subjects sat in a quiet, isolated room. During the session the experimenter waited at the other side of the room. For the mechanically simulated haptic feedback condition, the motors in the Logitech Wingman force feedback mouse were used to create hole-shaped and bump-shaped force-fields, pushing the cursor towards the centre of the target or away from the centre. In the optically simulated force feedback condition the same force field was simulated with cursor displacements.

For the experiment we needed a formula that could render fluent bumps and holes and without artefacts that could function as unintended cues for the subjects to recognise the shapes. We tried different mathematical means of rendering the bumps and holes: linear, polynomial, Gaussian, and sinusoid. The polynomial shapes were not chosen because they have discontinuous derivatives at their boundaries at the zero plane that could become an unintended cue for the subjects. The linear shape has discontinuities both at the zero plane and at the top. The Gaussian shape is completely continuous, but is zero nowhere. For the current experiment we chose to use a squared cosine shape since this shape results in a bump or hole with a clear but not too abrupt boundary and a smooth top (Figure 3-1).

The circular area where the force field was applied had a diameter of 240 pixels, the same as the diameter of the area occupied by the visually displayed target. The range of the mechanical forces applied to the mouse and the range of the force gains of the optically simulated haptic feedback was set by a committee of four people that were involved in similar projects and had knowledge of the techniques used. They preset the optically and mechanically simulated strength so that they were, in their perception, individually equal. This was done by conducting a series of mini-experiments in which these four people compared different mechanical strengths to different optical strengths up to the point where they believed the most extreme slopes (deep hole and high bump) to be equal across the two conditions. The values resulting from this cross-modal matching are called the nominal values. During the experiment this setting was varied between the three different groups.

$$f(x) = \cos^2(x)$$



$$f(x, y) = \cos^2(\sqrt{x^2 + y^2})$$




Figure 3-1. Sinusoid calculation of the structures in 2D and 3D.




Figure 3-2. Screenshot of the experiment.

Procedure

In the experiment the subjects were presented with a series of ‘bump’ and ‘hole’ shaped force feedback fields centered in an area indicated by a circle 240 pixels in diameter (Figure 3-2). The force fields were generated as a combination of mechanically and optically simulated force feedback in independently varying strengths. Subjects were instructed to move the cursor over the circle on the screen and asked to indicate how deep or high the structure was they perceived within the outlined circle. They were asked to do this at their own pace. They were not informed in advance about the different feedback conditions. After the experiment the subjects were asked if they experienced different ways of representing the bumps and holes and if so, which ones. In addition, they were asked what strategy they used to determine the height or depth of a field and, finally, if they had had previous experiences with force feedback before the experiment.

Design

The experiment was a 7x7 within-subjects design. The two factors, optically simulated haptic feedback (OSHF) and mechanically simulated haptic feedback (MSHF), were varied over seven levels:

OPTICALLY SIMULATED HAPTIC FEEDBACK	-3,-2,-1,0,+1,+2,+3
MECHANICALLY SIMULATED HAPTIC FEEDBACK	-3,-2,-1,0,+1,+2,+3




Figure 3-3 All 49 combinations of optical and mechanical structures presented in the experiment.



Figure 3-4. The Wingman Force Feedback Mouse used.

The mechanically and optically simulated components both have seven different height settings: deep hole (-3), medium hole (-2), shallow hole (-1), flat (0), low bump (1), medium bump (2), high bump (3). Combining these settings into a factorial design, results in 49 combinations of optical and mechanical force field stimuli (Figure 3-3).

In the first half of the experiment, all 49 combinations were presented to the subjects in random order. In the second half the same series of combinations was shown but in reversed order, resulting in a total of 98 trials. For each of these trials the user had to estimate the height of the force feedback texture underneath the disk on an integer scale from -3 to +3, where -3 represents a strong hole and +3 a strong bump. Note that in a part of these combinations the various types of feedback reinforce each other, whereas in other combinations they counteract. For instance, a medium bump in mechanical force feedback (+2) combined with a deep hole (-3) with optically simulated force feedback will result in a contradictory hole/bump situation (+2,-3) in which the subject has to integrate between the different modalities. We introduced a test phase at the beginning of the experiment, to let the users know what kind of heights they would encounter during the experiment. During this phase the nine most extreme values (+3,+3),(+3,0),(+3,-3), (0,+3),(0,0),(0,-3), (-3,+3),(-3,0),(-3,-3) from the main experiment were presented to the subjects in a setting that is identical to that of the main experiment. Likewise these values were displayed in a randomized order. In order to gain insight into a possible turning point between the dominance of the different modalities, we divided the subjects into three groups of ten people; each group conducted the experiment with different ranges of the mechanical and optical nominal force settings.

- Exp. 1: 100% nominal optical strength, 80% nominal mechanical strength
- Exp. 2: 100% nominal optical strength, 100% nominal mechanical strength
- Exp. 3: 80% nominal optical strength, 100% nominal mechanical strength

As mentioned, these settings were determined on the basis of four informed judges who tried to balance the perceived relative strengths of the two feedback conditions in such a way that, on average, they would play an equal role.

Another more fundamental reason to carry out the experiment for three different ranges of the optical and the mechanical strength is to control for the possible strategy participants may adopt to adapt their estimations to these ranges. So, their strategy might be, after the practice trials, to rate the depth of a hole with -3 when both the optical and the mechanical strength are minimum, and to rate the depth of a hole with +3 when both optical and mechanical strength are maximum, more or less independent of the actual settings of these ranges. If the participants would follow this strategy the results for the three experiments will be the same. If not, the estimations by the participants will vary in accordance with the different ranges of the optical and the mechanical strengths of the simulation.

3.4 RESULTS

It appeared that participants easily recognized the condition in which both optical and mechanical feedback were 0; in all cases they indicated the height of the object as 0. This means that the variance for this data point is zero, and so it was excluded from the following analyses of variance. The data were subjected to a 3-way analysis of variance with the estimations (ESTIM) of the height of the virtual object as dependent variable, and the factor EXP, experiment, representing the three nominal settings of the ranges of the two kinds of feedback, the factor OSHF, the strength of the optically simulated haptic feedback, and the factor MSHF, that of the mechanically simulated haptic feedback MSHF. The results are presented in Table I. There are significant effects of MSHF ($F(6, 2736) = 143.152$; $p < 0.001$) and OSHF ($F(6, 2736) = 958.367$; $p < 0.001$), but not of EXP ($F(2, 2736) = 0.272$; $p = 0.762$). On the other hand, there is no significant first-order interaction between MSHF*OSHF ($F(35, 2736) = 0.619$; $p = 0.961$); but there are significant first-order interactions between MSHF*EXP ($F(12, 2736) = 8.622$; $p < 0.001$) and OSHF*EXP ($F(12, 2736) = 5.205$; $p < 0.001$). This shows that the effect of both mechanically and optically simulated haptic feedback depends on the factor EXP, experiment. So, the different ranges used for the three experimental set-ups lead to significantly different estimations of the height of the virtual objects. The data of the three experimental conditions will, therefore, be analysed separately. Finally, there was no significant three-way interaction between MSHF, OSHF, and EXP. We will discuss the detailed results of Experiment I, conducted with group 1, and then only deal with the differences found for the other experiments II and III, which were conducted with group 2 and 3.

Tests of Between-Subjects Effects(b)

Dependent Variable: ESTIM

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	6292,388(a)	143	44,003	47,858	,000
Intercept	28,631	1	28,631	31,139	,000
EXP	,500	2	,250	,272	,762
MSHF	789,719	6	131,620	143,152	,000
OSHF	5286,993	6	881,165	958,367	,000
EXP * MSHF	95,126	12	7,927	8,622	,000
EXP * OSHF	57,433	12	4,786	5,205	,000
MSHF * OSHF	19,934	35	,570	,619	,961
EXP * MSHF * OSHF	42,738	70	,611	,664	,986
Error	2515,600	2736	,919		
Total	8838,000	2880			
Corrected Total	8807,988	2879			

a R Squared = ,714 (Adjusted R Squared = ,699)**b** Weighted Least Squares Regression - Weighted by MSHF \sim 0 | OSHF \sim 0 (FILTER)

Table 1: ANOVA table for the 3-way analysis of variance with the factor EXP for Experiment, the factor MSHF for the strength of the mechanically simulated haptic feedback, and OSHF for the strength of the optically simulated feedback.




Figure 3-5: Average estimations of the height of the virtual bumps and holes for the feedback conditions of Experiment I, in which the mechanically simulated haptic feedback (MSHF) varied over 80% of its maximum range and optically simulated haptic feedback (OSHF) over 100%.

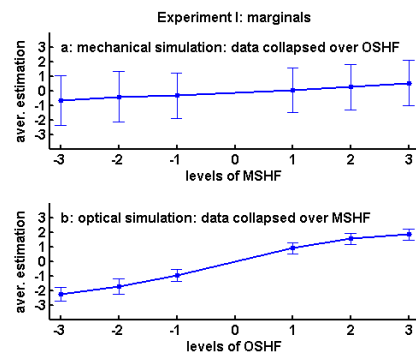


Figure 3-6: Marginal means of the estimations shown in Figure 3-5. In a the results for mechanically simulated haptic feedback are presented; hence, the data are collapsed over optically simulated haptic feedback. In b the results are collapsed over mechanically simulated haptic feedback, showing the results for optically simulated haptic feedback. The vertical lines represent one standard deviation up and one down.

Experiment I

The results of Experiment I, in which the participants received 100% nominal optically simulated haptic feedback and 80% nominal mechanically simulated haptic feedback, are presented in Figure 3-5. The vertical lines indicate for each combination of mechanically- and optically simulated haptic feedback the standard deviations of 20 estimations, 2 estimations for all 10 participants.

The estimations by the participants of Experiment I were subjected to a two-way analysis of variance with MSHF and OSHF as fixed factors. There were significant main effects of both MSHF ($F(6, 912) = 27.672$; $p < 0.001$) and OSHF ($F(6, 912) = 427.526$; $p < 0.001$). There was no significant interaction ($F(35, 912) = 0.503$; $p = 0.993$). Hence, we can collapse the data over MSHF and over OSHF, resulting in two marginal means. The results are shown in Figure 3-5. It can be seen that, for mechanically simulated haptic feedback, with the data collapsed over optically simulated haptic feedback, the range of the estimations is much smaller than for optically simulated haptic feedback, while the standard deviations of the result are much larger. This shows that in this experimental configuration the participants attributed more weight to the values of the optically simulated haptic feedback than to those of the mechanically simulated haptic feedback. This is substantiated by the linear regression analysis on the data which shows that the regression coefficient was 0.739 for OSHF and 0.188 for MSHF; the intercept was 0.101. The correlation coefficient between OSHF and the estimations of the participants was 0.831, whereas it was 0.211 for MSHF, which is highly significant according to the difference test for correlations coefficients based on Fisher's z-transform ($z = 21.37$, $N = 960$; $p < 0.001$).




Figure 3-7: Average estimations for the monosensory simulation conditions, which means that either optically simulated haptic feedback was zero (a) or mechanically simulated haptic feedback was zero (b). Note that these graphs are the same as the vertical cross sections of figure 3-5 along the two horizontal axes.




Figure 3-8: Average estimations of the data shown in Figure 3-5 for which the optically simulated haptic feedback was equal to the mechanically simulated haptic feedback, the coherent condition, compared with the incoherent condition for which the optically simulated haptic feedback was opposite to the mechanically simulated haptic feedback. Note that these represent the diagonals shown in Figure 3-5. In (a) the results for the coherent feedback are presented; in (b) the results for the incoherent feedback.

The finding that in this experimental configuration the participants attributed more weight to the values of the optically simulated haptic feedback than to those of the mechanically simulated haptic feedback can also be seen by comparing the conditions in which there is only mechanically simulated feedback with those in which there is only optically simulated feedback. The average results for these mono-sensory conditions are shown in Figure 3-7a and b.

Another interesting aspect becomes apparent when one looks at the diagonals of Figure 3-5. The diagonal from the front corner, the point $(-3, -3)$, to the back corner left to right, the point $(3, 3)$, shows the estimations by the participants for those feedback conditions in which the numerical category of the mechanically simulated haptic feedback was equal to that of the optically simulated haptic feedback; so, mechanical and optical feedback reinforce each other. This will be called *coherent feedback*. The other diagonal, running from the left to the right corner, on the other hand, shows the conditions in which mechanical and optical simulations oppose each other. This will be called *incoherent feedback*. Notice that the standard deviations for the values on the latter diagonal (Figure 3-8b) are larger than for those on the former (Figure 3-8a). This is substantiated by the amount of explained variance in both conditions as determined by the correlation coefficients between feedback and the estimations by the participants. In the coherent feedback condition it is 70%, whereas in the incoherent condition it is only 4%. This is highly significant ($z = 5.70$, $N = 120$; $p < 0.001$).

It is concluded that in the incoherent condition the variability of the responses is much larger than in the coherent condition. In addition, a three-way ANOVA with MSHF and OSHF as fixed factors, and participant (PP) as random factor resulted in no main effect of PP, but there was a significant interaction between both MSHF and PP ($F(54, 315)=5.285$; $p < 0.001$), and OSHF and PP ($F(54, 315)=9.026$; $p < 0.001$). This shows that different participants reacted differently to the different combinations of feedback. The much smaller amount of explained variance in the incoherent stimulus condition shows that this is to a large extent due to the different responses in the incoherent stimulus conditions, whereas participants are more consistent with each other in the coherent condition.

Experiments II and III

The main difference between the results of the three experiments corresponded with the larger range covered by the strength of the mechanical simulation relative to that of the optical simulation. In Experiment II the nominal range of the mechanical simulation was increased from 80% to 100%, while in Experiment III the range of the optical simulation was reduced from 100% to 80%. This expressed itself in a higher weight the participants attributed to the mechanical component of the stimulus in Experiment II and III. Hence, the results were similar to those of Experiment I except for shifts according to the different balance between the mechanical and optical contribution to the simulations. This will be discussed in detail now.

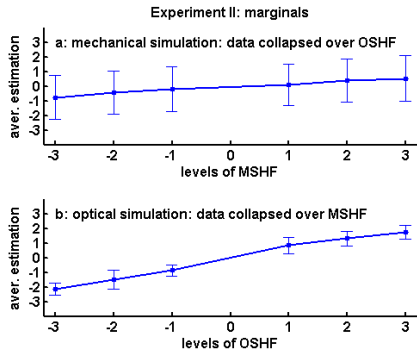


Figure 3-9. Marginal means of the estimations of Experiment II. In (a) the results for mechanically simulated haptic feedback are presented; hence, the data are collapsed over optically simulated haptic feedback. In (b) the results are collapsed over mechanically simulated haptic feedback, showing the results for optically simulated haptic feedback. The vertical lines represent one standard deviation up and one down.

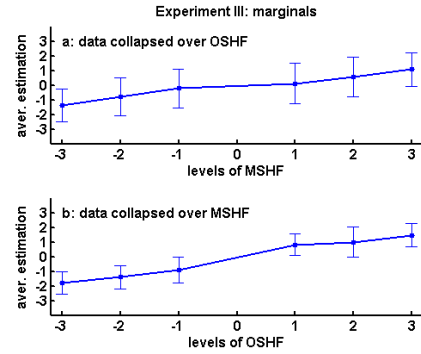


Figure 3-10. Marginal means of the estimations of Experiment III. In (a) the results for mechanically simulated haptic feedback are presented; hence, the data are collapsed over optically simulated haptic feedback. In (b) the results are collapsed over mechanically simulated haptic feedback, showing the results for optically simulated haptic feedback. The vertical lines represent one standard deviation up and one down.

Indeed, in Experiment II there were significant main effects of both MSHF ($F(6, 912) = 32.847$; $p < 0.001$) and OSHF ($F(6, 912) = 346.737$; $p < 0.001$), while there was no significant interaction ($F(35, 912) = 0.672$; $p = 0.928$). In Experiment III, with similar results, the statistics were $F(6, 912) = 91.905$; $p < 0.001$, for MSHF, $F(6, 912) = 217.535$; $p < 0.001$, for OSHF, and $F(35, 912) = 0.749$; $p = 0.855$ for the interaction. Hence, we can collapse the data over MSHF and over OSHF, resulting in two marginal means. The results are shown in Figure 3-9 for Experiment II and in Figure 3-10 for Experiment III. It can be seen that, again, for optically simulated haptic feedback, with the data collapsed over mechanically simulated haptic feedback, the range of the estimations is smaller, and the size of the standard deviations of the estimations is much larger, than for mechanically simulated haptic feedback, with the data collapsed over OSHF. This shows that, although the nominal contribution of the mechanically simulated haptic feedback was larger now, also in this experimental configuration the participants attributed more weight to the values of the optical simulation than to those of the mechanical simulation.

Showing the results for the monosensory simulations gives a similar picture as for experiment I. This is shown in Figure 3-11 and Figure 3-12. Again the variance of the monosensory conditions with purely mechanical simulations, shown in Figure 3-11a and Figure 3-12a, is much less than that of the average conditions shown in Figure 3-9a and Figure 3-10a, which contrasts with the conditions of purely optical simulations. But note that the difference is less for Experiment III (Figure 3-12) than for Experiment II (Figure 3-11), which in its turn is less than for Experiment I (Figure 3-7). This must be due to the range of the mechanical levels now being relatively




Figure 3-11: Average estimations of experiment II for the monosensory stimulation conditions, which means that either optically simulated haptic feedback was zero (a) or mechanically simulated haptic feedback was zero (b). Note that these graphs are the same as the vertical cross sections of figure 3-5 along the two horizontal axes.




Figure 3-12: Average estimations of experiment III for the monosensory stimulation conditions, which means that either optically simulated haptic feedback was zero (a) or mechanically simulated haptic feedback was zero (b). Note that these graphs are the same as the vertical cross sections of Figure 3-5 along the two horizontal axes.

wider in Experiment III than in Experiment II, where it is wider than in Experiment I, this in relation to the range of optical levels. In other aspects the results are similar. Furthermore, the standard deviations in Figure 3-12b are now larger when compared with Figure 3-11b, where it is larger than in Figure 3-7b. This shows that the effect on the participants' estimations of the mechanical simulation is now stronger than in experiment II, where it is stronger than in Experiment I.

This is substantiated by a linear regression analysis on the data which for Experiment II yielded a regression coefficient of 0.678 for OSHF and 0.209 for MSHF; the intercept was 0.086. The correlation coefficient between OSHF and the estimations of the participants was 0.801, whereas it was 0.246 for MSHF, which difference is highly significant according to the difference test for correlations coefficients based on Fisher's z -transform ($z = 18.60$, $N = 960$; $p < 0.001$). Observe that the difference between the weights attributed to optically simulated haptic feedback and to mechanically simulated haptic feedback diminishes in correspondence with the larger range of the mechanically simulated haptic feedback which from Experiment I to II was increased from 80% to 100%. For Experiment III these regression coefficients were 0.581 for OSHF and 0.376 for MSHF, while the intercept was 0.119. The correlation coefficient between OSHF and the estimations of the participants was 0.675, whereas it was 0.437 for MSHF, which difference is again highly significant ($z = 7.69$, $N = 960$; $p < 0.001$). The smaller difference between the weights now corresponds with the smaller range of the optically simulated haptic feedback which from Experiment II to III was decreased from 100% to 80%.

The results for the coherent and the incoherent stimulus conditions are presented in Figure 3-13 for Experiment II and in Figure 3-14 for Experiment III. A comparison




Figure 3-13: Average estimations of experiment II for the monosensory stimulation conditions, which means that either optically simulated haptic feedback was zero (a) or mechanically simulated haptic feedback was zero (b). Note that these graphs are the same as the vertical cross sections of Figure 3-5 along the two horizontal axes.




Figure 3-14 : Average estimations of the data shown in Figure 3-13 for which the optically simulated haptic feedback was equal to the mechanically simulated haptic feedback, the coherent condition, compared with the incoherent condition for which the optical simulation was opposite to the mechanical simulation. In (a) the results for the coherent feedback are presented. In (b) the results for the incoherent feedback are presented.

between the coherent and the incoherent stimulus conditions in Experiment II gives a percentage of explained variance of 64% in the coherent feedback condition and of 6% in the incoherent feedback, which is highly significant ($z = 6.87$, $N = 120$; $p < 0.001$). For Experiment III the results are 46% for the coherent feedback condition and 19% for the incoherent feedback, which is highly significant ($z = 12.30$, $N = 120$; $p < 0.001$). So, it is concluded again that in the incoherent condition the variability of the responses is much larger than in the coherent condition. And as in Experiment I, another three-way ANOVA with MSHF and OSHF as fixed factors and PP as random factor resulted in no main effect of PP, but in significant interactions between MSHF and PP (Experiment II: $F(54, 315) = 6.405$; $p < 0.001$; Experiment III: $F(54, 315) = 4.548$; $p < 0.001$), and OSHF and PP (Experiment II: $F(54, 315) = 10.843$; $p < 0.001$; Experiment III: $F(54, 315) = 4.945$; $p < 0.001$). This shows that participants did not respond in the same way to the various feedback conditions. The much smaller amount of explained variance in the incoherent stimulus condition shows that this is largely due to the different responses in the incoherent stimulus conditions. Observe that the difference in explained variance between the coherent and the incoherent condition is larger, again, than in Experiment II, where it was already larger than that in Experiment I.

Comparison of Experiment I, II, and III

The results show that the three different experimental conditions, in which the range of the mechanically simulated haptic and the optically simulated haptic feedback was varied over 80% mechanical feedback and 100% optical feedback, 100% mechanical

feedback and 100% optical feedback, and 100% mechanical feedback and 80% optical feedback, correspondingly changes the weight the participants attributed to the mechanically and the optically simulated haptic feedback. This follows from the decreasing correlation coefficients between the estimations and the optically simulated haptic feedback condition, and the increasing correlation coefficients between the estimations and the mechanically simulated haptic feedback condition. Although initially the choice of the values for the ranges of the optical and the mechanical feedback was based on intuitive judgments by the committee of four experts that, in the conditions of Experiment III, optical feedback would not dominate, the results show that even then the participants on the average attributed more weight to the optical simulation. We do not have a single explanation for the unexpected result that participants attributed more weight to the optical simulation than the committee of four, who set the standard and were familiar to the different techniques. It could be that the experimenters, who had set the standard, subconsciously overweighed the optical simulation to make it more recognisable by the participants. Or it could be that uninformed participants who are not familiar with the two different techniques attribute more weight to the optical simulation than people who know the used techniques. Despite that we did not find a clear turning point where the optical outweighs the mechanical stimulation, the results clearly show that optically simulated haptic feedback is a powerful tool to complement or replace mechanical feedback, and that the relative role of each form of feedback can be adjusted.

Furthermore, the comparisons of the coherent with the incoherent feedback conditions show that the correlation coefficients for the coherent stimulus conditions slightly increase from Experiment I to III, hence with decreasing optically and increasing mechanically simulated haptic feedback, from 0.925, via 0.941 to 0.956. The difference between the first and the last correlation coefficients is statistically significant ($z = 2.10$; $N = 120$; $p < 0.025$). This shows that in the coherent condition, mechanical and optical simulations of haptic feedback enhance each other. On the other hand, the correlation coefficients for the incoherent stimulus conditions decrease highly significantly from Experiment I to III from 0.705, through 0.690, to 0.281. The difference between the last correlation coefficient and the first is statistically significant ($z = 4.50$; $N = 120$; $p < 0.001$), as is the difference between the last and the second ($z = 4.27$; $N = 120$; $p < 0.001$). This shows that in the incoherent condition the 'disturbing' influence of the increasing contribution of the mechanical feedback on the participants' responses increases.

The three experiments were carried out by three different groups. In this between-groups design it would have been possible that each group would adapt the range of estimations to the range of the optical and the mechanical feedback. Since this did not happen we can conclude that the estimations were based on a real percept of the depth of a hole or the height of a bump, and not on a perceptual normalization of the

estimations over the stimulus ranges. Note that in a between-subject design in which the trials of the three experiments would have been randomly mixed, this strategy is not possible, and the differences found would only have been more significant.

Prior Knowledge

After the experiment the participants were asked if they were experienced or had heard of force feedback. Twenty eight participants said that they knew about mechanical force feedback either through computer games or through professional use. Among the subjects there were eight who said that they knew about optically simulated force feedback before the experiment. We compared the answers of these participants with their experimental results. We found no correlation between prior knowledge of the applied techniques and the outcome of the experiment.

Result Summary

The results clearly show that mechanically and optically simulated haptic feedback complement each other in supplying information as to the depths of holes and heights of bumps in a GUI. Noticeably, in the ranges applied in the present study, optically simulated haptic feedback gives a stronger haptic illusion of force feedback than anticipated by the researchers, even to a degree that it can replace mechanical force feedback. In all three experimental conditions the variation of the optical simulation contributes more to the variance of the participants' estimations than does mechanical simulation.

The statistical interaction between mechanically and optically simulated haptic feedback was not significant, indicating that *on average* the haptic effects of mechanical and optical simulations are simply additive. On the other hand, for incoherent stimuli when one simulation method gives the haptic illusion of a hole and the other of a bump, the variance of the estimations is much higher than for coherent stimuli. This is due to the fact that the relative weight each subject attributes to the information coming in through the visual and the haptic modality is different. Most subjects attributed more weight to the visual information, but a minority of subjects, one in Experiment I, two in experiment II and three in Experiment III, predominantly paid attention to the haptic information generated by the mechanical force feedback device. Together with the findings discussed in section 4.3, this increase, though statistically not significant, again corresponds with the relatively increasing contribution of the mechanically simulated feedback and its larger range in Experiment II and III. Hence, the subjects did not adapt their estimations of the heights of the bumps and the depths of the holes to the ranges of the mechanical and the optical strength, but really based their estimations on the percept of depth and height as induced by the optical and the mechanical stimulus components. In the case of conflicting information some participants attributed other weights to the two stimulus components than others.

3.5 DISCUSSION

We compared the perception of virtual bump and hole structures in a graphical user interface, simulated mechanically with a force feedback device and optically by applying cursor displacements as if there was force feedback. The present study points out that, since active cursor displacements can induce strong haptic illusions, they are well suited for simulating bump and hole structures. Hence, the results are in line with our first hypothesis that optically simulated haptic feedback can be applied to create perceivable bump and hole structures and that people are able to judge the relative heights of the bumps and the depths of the holes to an extent comparable to that of mechanically simulated force feedback.

Our second hypothesis, that optically simulated haptic feedback can be used to enhance or decrease the perceived height of bump and hole structures generated with a mechanical force feedback device, is confirmed in the sense that, on average, the haptic illusions induced by mechanical and optical simulations are additive. On the other hand, for incoherent stimuli for which the illusions induced by mechanical and optical simulations oppose each other, a minority of subjects attributes more weight to haptic information than to visual information, as mentioned one for Experiment I, two for Experiment II and three for Experiment III. In situations in which the mechanically simulated haptic feedback will be much stronger the relative weight attributed to the optically simulated haptic feedback will be much less.

These findings do not demonstrate that force feedback can be simulated optically in all situations in which mechanical force feedback has shown to be beneficial. Optically simulated haptic feedback can naturally be applied only in situations in which the users of the GUI can direct their attention to the visual information presented on the screen, even more specifically to the position of the cursor on the screen where the haptic illusions are induced.

Furthermore, it is possible that optical simulations are especially beneficial for small targets, or for weak mechanical forces. In the next chapter we show, using a Fitts' type target-acquisition task, that optically simulated haptic feedback has a significantly higher usability than mechanically simulated haptic feedback, but the benefit of the optical simulation expressed itself mainly for smaller targets, less than 20 pixels in diameter. In the experiments presented here, the holes and bumps were 240 pixels in diameter, which is in the range where this usability benefit of optically simulated force feedback is no longer significant. Hence, it may very well be that it is not so much size of targets but strengths of forces that matters. In other words, the enhancing effects of optically simulated haptic feedback may express themselves especially when mechanical feedback is relatively weak.

Conclusions

We conclude that optically simulated haptic feedback is a good alternative for mechanically simulated force feedback when rendering holes and bumps. The research presented here showed that users are able to interpret the simulated topographical structures correctly. Active cursor displacements, influencing the normal cursor movement linked to the user's mouse movements, can be applied to generate bump and hole structures. Participants in the experiments successfully identified the bump and hole slopes and estimated their sizes in a consistent way, in both the optically simulated haptic feedback condition, and in the condition using a mechanical force feedback device. Optically simulated haptic feedback can further be applied to alter the perception of mechanically simulated haptic structures. In some respects, e.g., for more subtle forces, optically simulated haptic feedback is likely to be even more expressive than mechanical simulations of haptic feedback, at least for the ranges tested in the present study. Optically and mechanically simulated haptic feedback must be applied in a coherent way. If not, different users will react differently and, hence, unpredictably.

Optically simulated haptic feedback can be applied to generate perceivable haptic structures in a standard cursor controlled graphical user interfaces, without resorting to special mechanical input/output devices. This technique of simulating haptic feedback optically opens up an additional communication channel with the user. Optically simulated haptic feedback is not expected, however, to replace mechanically simulated haptic feedback in general. Rather we expect that, since optically simulated haptic feedback can be implemented in a standard desktop set-up without special hardware, it could catalyse the development of novel physical interaction styles and the acceptance of mechanical force feedback devices.

More research into the potential and limitations of optically simulated haptic feedback is still needed. Further perceptual experiments might deal with the recognition of various objects, from simple forms, like ramps, squares, gutters, and triangles, to dynamic complex scenes. In the next chapter we will move on to assess the usability of optically simulated haptic feedback in a simple pointing task. Another important research path, is the design of novel interaction styles based upon optically simulated haptic feedback. The current interfaces are not designed with tactility in mind. Before optically simulated force feedback can be fully applied in more complex interaction styles, designers and researchers need to experiment more with the technique; explore its affordances and find out what works and what does not. This issue will be addressed in chapter 5 where we will present a software toolset that enables designers to add optically simulated haptic feedback to their interfaces without difficult programming.

4

Usability of Optically Simulated Haptic Feedback

The current chapter¹² describes an experiment in which we determined the usability of optically simulated haptic feedback (OSHF) and compared it with the usability of a mouse provided with mechanically simulated haptic feedback (MSHF) and, as a reference, with a mouse operating as usual with neither mechanically nor optically simulated haptic feedback. In order to quantify and compare the usability of the mouse in the three conditions, subjects carried out a Fitts'-type target-acquisition task. This means that the subjects were presented with targets on a screen and instructed to move the mouse cursor to the target and click. The targets were varied in size and in respect of the distance to the current cursor position. We will use Fitts' law (Fitts, 1954; Fitts and Peterson, 1964), which describes the influence of size and distance of the target on the movement time, to measure the efficiency of the target-acquisition task. The error rate will be measured to determine the effectiveness, and a questionnaire will be presented to determine the satisfaction of the users under the three feedback conditions. Following the ISO 9241 (1998) standard on usability, we will establish the usability of the mouse under the three conditions by combining the measurements of efficiency, effectiveness, and user satisfaction.

4.1 RELATED WORK

Given the pervasiveness of pointing in graphical interfaces, every small improvement in the target-acquisition task, stands for a substantial improvement on user productivity.

¹² This chapter is based on Usability of optically simulated haptic feedback, Mensvoort van, K. Hermes D.J., Monfort van, M., 2008, published in International Journal of Human-Computer Studies 66 (2008), pp. 438–451.

Consequently, various researchers focused on improving pointing performance. Bier and Stone (1986) introduced “snap-dragging”. This direct-manipulation technique uses the ruler and compass metaphor to provide constrained cursor positioning, and to help the user place his or her next point with precision. In graphical applications, cursor snap-dragging is traditionally used as a tool for positioning the input-device cursor relatively to the position of existing markers and axis-aligned grids. Kabbash and Buxton (1995) introduced *area cursors* that have a activation area that is larger than the single pixel of standard cursors, and thus make it easier to select, especially smaller, targets. A drawback of the use of area cursors is that it becomes impossible to select one target if several targets are closely grouped together. Others proposed to dynamically change the size of interface widgets near the cursor position to offer the user a larger target area to interact with their presumed focus of attention. This so called fish-eye technique (Furnas, 1986; Bederson, 2000; McGuffin and Balakrishnan, 2002) has been implemented in the Apple Mac OSX dock, where icons in the desktop toolbar expand when the cursor is positioned over them. The cursor can also be warped to the eye-gaze area which encompasses the targets, when using an eye-tracking system (Zhai, 1999). Another promising improvement is the use of *sticky icons* introduced by Keyson (1997) and Worden (1997). With sticky icons the cursor’s control/display ratio, which determines the mapping between the physical mouse movement and the cursor movement on the screen, is automatically reduced as the cursor approached a target, and then returns to normal after passing the target. User movement in entering a target is thus slowed down by the change in cursor gain. In research by Worden (1997), the sticky icons had no effect on accuracy, but substantially improved the speed of performance over the traditional pointer. Especially older users benefited from the sticky icon technique.

Another approach to enhance navigation in GUI’s has focused on the input device. Akamatsu and MacKenzie (1996) analyzed the acquisition of targets with tactile feedback, produced by a solenoid-driven pin that stimulated the user’s index finger, and with force-feedback, produced by an electromagnetic friction on the mouse. They found that the feedback modes reliably reduced overall targeting time and error rates. The combination of tactile and force-feedback produced the lowest mean acquisition times. Keyson (1997) conducted a study involving a trackball with force-feedback, comparing force-feedback, consisting of a pulling force towards the centre of a target, with sticky targets, consisting of a reduction of the cursor gain within the target. Results of the experiment showed that target acquisition performance was generally higher in the tactile-feedback condition, followed by cursor-gain feedback, and then no cursor-gain feedback. Oakley et al. (2001) examined how haptic feedback affected simple user-interface controls, finding that it did not consistently decrease task time, but significantly reduced the error rate. Dennerlein et al. (2000) showed that force feedback can improve performance for steering and a combined steering–targeting task. They conducted an experiment in which participants moved the cursor through a

small tunnel with varying indices of difficulty using a conventional and force feedback mouse. For the force feedback condition, the mouse displayed force that pulled the cursor to the center of the tunnel. Movement times were on average 52% faster during the force-feedback condition when compared with the conventional mouse.

In chapter 2 we presented the active cursor technique, a method aimed at optically simulating haptic feedback in a standard WIMP GUI setting without resorting to special input/output devices. Active cursor displacements are applied to evoke a percept of haptic feedback, while using a normal mouse not enabled with mechanical force feedback. In our realization, interactive animations are used to simulate the haptic operation of mechanical force feedback devices. In chapter 3 we have showed that the disparity between the optical feedback, i.e., slowing down the speed of the cursor, and the increasing reaction force applied to the input device to compensate for this, induces an illusion of haptic feedback. Like force-feedback devices, the active cursor technique can guide the user towards preferred positions or communicate properties of the interface to the user. Due to these cursor displacements a hole becomes an easily accessible part of the screen, whereas a hill area is hard to access. In the current experiment we assess visually simulated haptic feedback in comparison with haptic feedback, generated by a mechanical force feedback mouse. Since direct two-way communication through the pointing device has proved beneficial for haptic devices (Akamatsu et al., 1995; Keyson, 1997; Oakley et al., 2001), it seems reasonable to expect benefits from two-way communication through cursor displacements.

Hypotheses

Our first hypothesis is that OSHF as presented by the active cursor technique will increase the usability of the mouse. This will result in an increase of the efficiency of the device as measured by the index of performance based on Fitts' law, an increase in effectiveness as measured by a decrease in the numbers of errors, and an increase in satisfaction as measured by a questionnaire presented to the participants.

Our second hypothesis is that the increase in usability by applying OSHF will be similar in extent to the increase in usability obtained by MSHF.

4.2 EXPERIMENT

Subjects

Thirty-four volunteer subjects participated. There were 22 male and 12 female subjects, ranging in age from 18 to 34. All subjects were regular users of mice in their daily work. None of the subjects had experience with force feedback mice or active cursor feedback. The subjects were not informed about the goal of the experiments in advance.

Apparatus

The experiment was conducted using the Logitech Wingman force-feedback mouse, a mouse attached to a mouse pad replacing the mouse mat and with two motors supplying force-feedback to the user (Rosenberg, 1997). This mouse was used in all experimental conditions. The host computer was a Pentium II class PC with a screen resolution of 1024x768 pixels on a 17-inch monitor. The data were collected with 1-pixel and 1-ms resolution and saved in output files for subsequent analysis. The subjects sat in a quiet, isolated room while the experimenter sat in an adjacent room.

For the normal feedback condition, the condition with neither mechanically nor optically simulated haptic feedback, target entry was indicated only through the displayed image of the cursor path. For the mechanically simulated haptic feedback (MSHF) condition, the motors in the Logitech Wingman force feedback mouse were used to create a hole-shaped force-field over the target, pushing the mouse towards the centre of the target. In the optically simulated haptic feedback (OSHF) condition the same force field was simulated with cursor displacements. Both the mechanically and optically simulated haptic feedback were calculated using the algorithm described in Figure 2-20, except that for the MSHF the calculated force vector was sent to the motor of the force feedback mouse and for OSHF the force resulted in a cursor displacement. Prior to the experiment the force feedback mouse was investigated for potential anomalies in its force rendering. This was tested by subsequently sending the entire range of software forces to the mouse, while measuring the actual mechanical force employed by the mouse in Newton using a strain gauge. We have conducted this test for both the x and the y axes. The force rendering of the device was ascertained to be linear and no anomalies were found.

To calculate the hole structures we needed a formula that could render fluent holes with a clear but not too abrupt boundary and a smooth bottom. We tried different mathematical means of rendering the bumps and holes: linear, polynomial, Gaussian, and sinusoid. The polynomial shapes were not chosen because they have discontinuous derivatives at their boundaries at the zero plane that could become an unintended cue for the subjects. The linear shape has discontinuities both at the zero plane and at the top. The Gaussian shape is completely continuous, but is zero nowhere. We chose to use a sinusoid shape, and calculated the hole shape according the algorithm described in Figure 2-21. The area where the force field was applied was the same as the area occupied by the visually displayed target (Figure 4-1).

Finally, the mechanically simulated forces had to be calibrated with the optically simulated forces. The balancing between the forces applied to the mouse and the force gains of the optically simulated haptic feedback was set by a committee of four people that were involved in similar projects and had knowledge of the techniques used. They preset the optically and mechanically simulated strengths so that they were, in their perception, individually equal. This was done by conducting a series of




Figure 4-1 The target sizes range from 10 to 80 pixels.

mini-experiments in which these four people compared different mechanical strength adjustments to different optical strengths up to the point where they believed the different holes to be of equal dept. This fixed force calibration setting was used for all target sizes and all subjects throughout the experiment.

Procedure

Subjects performed a simple target-acquisition task. The experimental screen started with a circle in the middle of the display. Once this circle was clicked on, it disappeared and the first target appeared. This target in its turn disappeared when clicked on and the second target appeared, etc., etc. The position of the targets in respect to the preceding targets was varied over eight directions: right above the target, right below it, to its left, its right or diagonally under angles of 45 degrees. The directions were divided into three sets, horizontal, vertical, and diagonal, each of which was applied just as often. Since the direction of a movement is of minor importance with circular targets (Boritz et al., 1991; Jagacinski and Monk, 1985; MacKenzie and Buxton, 1992), the results over these directions were merged. Subjects were instructed to move the cursor to the target and select the target by pressing the left mouse button. They were asked to do this in a normal tempo, as they would do in a normal desktop setting. They were not informed in advance about the different feedback conditions. After the experiment they were asked if they experienced different types of feedback and what ones they most preferred.

Design

The experiment was a 3x4x3 within-subject design. The factors and levels were as follows:

TARGET DISTANCE	72, 144, 288 pixels
TARGET SIZE	10, 20, 40, 80 pixels
HAPTIC FEEDBACK	mechanically simulated, normal, optically simulated

Each subject participated in three groups of three sessions, one session for each feedback condition. The first group of three sessions served to familiarize the subjects with the experiment. Only the data from the second and third group were used for data analysis. All six possible orders of the three feedback conditions were used to counterbalance order and learning effects. Subjects were assigned randomly to two orders, performing one order in the second group of sessions and one in the third group. In each session the subjects underwent, in random order, the three feedback conditions. In one session 72 targets were presented, three distances, four sizes, three directions and two repetitions. The target, the size of which was varied between 10, 20, 40 and 80 pixels, is shown in Figure 4-1. The default Windows 98 cursor was used.

These conditions created a range of task difficulties typical of point-acquisition tasks. The movement time MT , i.e., the time it takes to move the pointer from one target to the next and select (click on) it, can to a good approximation be described by Fitts' law (Card et al., 1978, MacKenzie, 1989; 1992). Fitts' law assumes that, for one specific application, the movement time only depends on the proportion of the distance A between the original and the target circle and the diameter W of the target circle. The latter because smaller targets require a higher precision of the movement than larger targets. In point-acquisition tasks by mice cursors on computer screens, the independent variables A and W are mostly expressed in pixels. Fitts' law in its original formula is: the time MT to move to and select a target of width W at distance A is

$$MT = a + b ID = a + b \log_2(2A/W) \text{ (sec)} \quad (1)$$

where a and b are constants for that application determined through linear regression. The log term is the *index of difficulty* ID and carries the unit bits. In this formulation the ID can become negative as A/W approaches zero. Hence, we will not use Fitts' original formula for the ID , nor Welford's (1960) but, following MacKenzie (1989), we will use the more accurate formulation by Shannon, $ID = \log_2(A/W + 1)$:

$$MT = a + b \log_2(A/W + 1) \text{ (sec)}. \quad (2)$$

Using Equation 2 in our experimental design, the task ID ranges from

$$ID = \log_2(72/80 + 1) = 0.93 \text{ (bits)} \quad (3)$$

for the easiest task with the maximum target size of 80 and the minimum distance of 72 pixels, to

$$ID = \log_2(288/10 + 1) = 4.90 \text{ (bits)} \quad (4)$$

for the hardest task with the minimum target size of 10 pixels at the maximum distance of 288 pixels. The performance of a device is often expressed as the inverse of the

constant b in equation (1) which has bits per second as unit. This quantity is often called the index of performance IP :

$$IP = 1/b \text{ (bits/sec). (5)}$$

This quantity represents the inverse of the time in seconds that the MT increases with every increment in bits of the ID . Hence, if for a certain device the MT increases rapidly with every bit of ID , the IP is low. It is high, otherwise. We will use the IP to measure the performance, and hence the efficiency, of the various feedback conditions quantitatively.

4.3 RESULTS

Movement time

The movement times averaged over the 34 participants are presented in Figure 4-2 for the three feedback conditions, the three distances and the four target sizes. The size of the target and the distance, both in pixels, between the targets are shown below the graphs, as are the feedback conditions, 'm' for mechanically simulated haptic feedback, 'n' for normal feedback, and 'o' for optically simulated haptic feedback. The vertical lines through the top of the bars give the 5% two-sided confidence intervals for the means of the movement times.

When the regression constants a and b in equation 1 were determined for the data points through the average movement times plotted as functions of the ID , $\log_2(A/W + 1)$, a yielded 0.316 s, while b was 0.191 s/bits which corresponds to an average index of performance over the three feedback conditions of 5.25 bits/s. The movement times as calculated according to these regression constants are presented as horizontal lines through the three-column bars in Figure 4-2. With these values a proportion of 0.977 of the variance of these average movement times could be explained. The two large effects corresponding with Fitts' law can clearly be distinguished. First, the movement times decrease as the target size increases from 10, to 20, 30 and 40 pixels. Second, the movement times increase as the distance between the start position and the targets increases from 72, to 144 and 288 pixels. Only for the targets at the smallest two distances, small but significant deviations from Fitts' law can be found. These deviations correspond with the more generally found tendency that the MT is higher than predicted by Fitts' law for small ID s (MacKenzie, 1992; Friedlander et al., 1998).

These findings are substantiated by a 3-way analysis of variance on the movement times with $DISTANCE$ and $SIZE$ as fixed factors and $FEEDBACK$ as random factor. The Anova table is presented in Table 1. Indeed, this analysis showed highly significant effects for $SIZE$ ($F(3)=822.54$, $p<0.0001$) and $DISTANCE$ ($F(2)=1080.33$, $p<0.0001$), and a significant effect for the interaction between $SIZE$ and $DISTANCE$

($F(6)=4.02$, $p=0.0194$). But more interesting is the significant interaction between SIZE and FEEDBACK, $F(6)=2.65$, $p=0.0052$. This interaction can be seen in the relative heights of the three adjacent columns of a bar in Figure 4-2 giving the results for the three feedback conditions. It can be seen that, for the smaller targets of size 10 and 20, the right column in the bar, representing the optically simulated force feedback condition, is lower than the middle column and, with one exception, than the left column. Indeed, the effect of FEEDBACK in itself was not significant in the 3-way analysis of variance ($F(2)=1.67$, $p=0.254$), and neither was the interaction between DISTANCE and FEEDBACK. But, as mentioned, the interaction between SIZE and FEEDBACK was significant, $F(6)=2.65$, $p=0.0052$.

This result will be analysed more globally by calculating and comparing the index of performance separately for the three feedback conditions.

Source	Sum of Sq.	df	Mean Sq.	F	Prob.>F
SIZE	42.5671	3	14.189	822.54	0
DISTANCE	14.4279	2	7.214	1080.33	0
FEEDBACK	0.0699	2	0.035	1.67	0.254
SIZE*DISTANCE	0.0731	6	0.0122	4.02	0.0194
SIZE*FEEDBACK	0.1035	6	0.0173	5.69	0.0052
DISTANCE*FEEDBACK	0.0267	4	0.0067	2.2	0.1301
SIZE*DISTANCE*FEEDBACK	0.0364	12	0.0030	0.2	0.9984
Error	17.7344	1188	0.0149		
Total	75.039	1223			

Table 1: Anova table for the 3-way analysis of variance on the movement times with SIZE and DISTANCE as fixed factors and FEEDBACK as random factor.




Figure 4-2 Average movement times for the four different sizes (SIZE) of the target circle, the three distances (DIST) between origin and target, and the three feedback (FB) conditions, 'm', mechanically simulated haptic feedback, 'n', normal feedback, and 'o', optically simulated haptic feedback. The vertical lines represent the two-sided 95% confidence intervals for the means of the movement times. The horizontal lines through the three columns of each bar represent the movement times as estimated on the basis of Fitts' equation fitted through the data points.

Efficiency: Index of Performance

As mentioned, by fitting a regression line through all 36 data points as a function of ID , the coefficients of Fitts' law were calculated. This gave an index of performance IP averaged over the three conditions of 5.25 bits/s. Next the results were analysed for the three conditions separately. For each participant separately a regression line was fitted through each set of twelve data points corresponding with each condition. A boxplot of the thus obtained participants' IP s is presented for the three conditions in Figure 4-3. The boxes present the median and the quartiles, the whiskers the extreme values, while the plusses are outliers. The average was 5.09 for the MSHF condition, 5.14 for the normal feedback condition, and 5.52 for the OSHF condition. This indicates that, contrary to our expectations, the largest difference is between mechanically simulated haptic feedback and normal feedback on the one hand and optically simulated haptic feedback on the other. Indeed, a repeated measures analysis of variance with Huynh-Feldt's correction for the lack of sphericity showed a significant effect of FEEDBACK on the IP s ($F(1.535) = 5.063$, $p = 0.017$).




Figure 4-3 Boxplot of the Indices of Performance (IPs) of the participants for each of the three conditions. The boxes represent the median and the quartiles, the whiskers the lowest and highest normal values, while the plusses indicate the outliers.




Figure 4-4 Error rate for the three feedback conditions 'f', force-feedback, 'n', normal feedback, and 'v', visual simulated force-feedback. The total number of clicks was 10368. The error bars represent the 95% confidence intervals under the assumption that the number of errors has a binomial distribution.

A post-hoc analysis with Dunn-Šidák's correction for multiple comparisons showed that, surprisingly, there was a significant difference between the mechanically (m) and the optically simulated haptic feedback (o) condition ($p = 0.048$). Hence, contrary to our second hypothesis that the benefits of optically simulated haptic feedback would be about the same as mechanically simulated haptic feedback, we found that optically simulated haptic feedback actually appeared to be more efficient than mechanically simulated haptic feedback. The data discussed above show that this is mainly due to the shorter movement time realised for the smaller targets. Our first hypothesis, that optically simulated haptic feedback would be more efficient than normal feedback, was not directly supported by a significant difference in IPs ($p = 0.093$).

Effectiveness: Error Rate

The error rate is defined as the proportion of mouse clicks that missed the target. Due to a technical problem the error-rate data for 10 subjects were not saved and unavailable for analysis. We measured it for 24 of the 34 subjects. The results are shown in Figure 4-4. The number of errors was lowest for the visually simulated condition, 318 errors in a total of 10368 clicks or 3.07%. With mechanically simulated haptics the subjects made 349 error clicks or 3.37%, whereas the error rate in the normal-feedback condition was highest, 3.94%, or 409 out of 10368 clicks. An analysis of proportions based on log-odd ratios showed that the difference in error rate under the MSHF condition was significantly higher in normal-feedback than in the OSHF condition ($p < 0.001$). It was not significantly different between mechanically and optically simulated haptic feedback ($p = 0.206$), while the error rate was about significantly lower than under the

normal-feedback condition ($p=0.026$). (The p -value of 0.026 is not significant when Bonferroni's or Dunn-Šidák's adjustment is applied.)

This shows that the effectiveness of optically simulated haptic feedback is at least as great as that of mechanically simulated haptic feedback, and that both mechanically and optically simulated haptic feedback enable higher levels of target acquisition than normal feedback, i.e. the normal operating of a mouse without feedback. This corresponds with both our first and our second hypothesis.

Satisfaction: User Preference

After the experiment the subjects filled in a questionnaire in which they were asked if they had experienced a difference between the various parts of the experiment. They were asked to describe the differences in their own words and were asked which part they preferred. We did not ask the subjects whether they adopted different strategies for the different conditions.

From one of the 34 subjects no preference judgment was obtained; this subject reported not to have noticed any differences between the various trials. One other only expressed dislike for mechanical force feedback devices, but was indifferent in respect of the optically simulated haptic feedback or the normal feedback. The remaining 32 subjects were able to distinguish among the three different feedback conditions. We want to accentuate that this is no surprise, since the experiment consisted of three separate blocks of trials. Although the subjects were allowed to describe the conditions in their own words, it was unambiguous what condition they meant when their preference was asked. Of the 32 remaining subjects, 22 expressed their preference for the optically simulated haptic feedback condition. For the rest, five had a preference for the mechanically simulated haptic feedback, while the last five subjects expressed a preference for the no-feedback condition. This difference between optically simulated haptic feedback on the one hand and the other two conditions is highly significant (with 10.67, 10.67, and 10.67 as expected values and 22, 5, 5, as obtained values, $\chi^2 = 18.68$, $v = 2$, $p < 0.001$). Obviously, the difference in preference for the mechanically simulated haptic feedback condition is not significantly different from that for the normal feedback condition.

Thus, optically simulated haptic feedback outperforms both mechanically simulated haptic feedback and normal feedback in respect of satisfaction as measured by asking the subjects for their preferences. This outcome regarding satisfaction corresponds with our first hypothesis, but not with our second hypothesis: the satisfaction score of optically simulated haptic feedback is not similar, but exceeds the satisfaction reported with mechanically simulated haptic feedback. Like with efficiency, the satisfaction of optically simulated haptic feedback is higher than expected.

4.4 DISCUSSION

We have compared the usability of mechanically simulated haptic feedback with optically simulated haptic feedback, i.e., by applying active cursor displacements in the direction of the simulated force. The usability of the mouse-controlled cursor under these feedback conditions was compared in a target acquisition experiment. The usual condition without any additional feedback was included as a reference. In the three feedback conditions, target distance and size were varied systematically. Earlier research (Akamatsu and Sato, 1994; Epps, 1986) pointed out that force feedback can improve the efficiency and effectiveness of a target-acquisition task.

The present study points out that optically simulated haptic feedback can replace mechanically simulated haptic feedback in target acquisition tasks and with respect to efficiency and satisfaction outperforms mechanically simulated haptic feedback. An analysis of the movement times under the three conditions showed that the higher efficiency was mainly realized for the smaller targets, in this case the targets of 10 and 20 pixels wide. The reduced error rates for both mechanically and optically simulated haptic feedback are likely due to the fact that for both these forms of feedback the probability of going beyond the target and clicking past it is prevented by the slower cursor movements when the target is passed. Both factors then can contribute to the higher satisfaction of optically simulated haptic feedback over both normal feedback and mechanically simulated haptic feedback.

The results are all in line with our first hypothesis that, in target-acquisition tasks, the application of optically simulated haptic feedback increases the usability of the mouse controlled cursor. It came as a surprise that, for smaller targets, the efficiency is even higher than that of mechanically simulated haptic feedback. This was not in line with our second hypothesis that the usability of optically simulated haptic feedback would be comparable to mechanically simulated haptic feedback. It may seem remarkable that “seeing” the forces acting on a manipulated object can, for small objects, be more beneficial than actually feeling the forces through a mechanical device. We have to realize here, that while the subjects are moving the target, the optically simulated haptic feedback as provided on the screen is not just a matter of ‘seeing forces’ through visual information only, but is actually an integration between on the one hand the force information as computed by the active cursor technique, and on the other hand the information the users have about their own movements. It is expected that both aspects play a role in the ‘visual force feedback’ as perceived by the subject.

The result that mechanically simulated haptic feedback is less efficient for smaller targets, may be explained from the fact that, in common situations, the forces acting on small objects can be quite weak. The same (hole) shape formula is used with different target sizes. This implies that the absolute force (and visual force) differs between the different sizes of targets; bigger targets are deeper than small ones. The inertia of the haptic device, in combination with the weight of the user’s hand holding it, could have caused the smaller forces to diminish in the mechanically simulated condition.



Figure 4-5 An example application of Active Cursor in a graphical user interface. Once the cursor enters the target it is dragged towards the centre of the target, resulting in easier target acquisition.

Since optically simulated haptic feedback is a software based method, inertia does not play a role. Very weak forces will be summed up and over time result in a cursor displacement. Of course inertia could be simulated in the active cursor algorithm, but this was not done in the current experiment. This might have caused the higher efficiency for smaller targets. A related explanation can be sought in the perceptual domain. Since human perception is adapted to the physical world in which, forces being equal, the change in velocity for smaller, lighter objects is relatively larger than that for bigger, heavier objects. One may, therefore, derive the forces acting on a small object better from seeing its changing velocity on its trajectory than from touch. This may certainly be relevant when one only looks at objects and does not touch them oneself. If the manipulations of an object are observed by a person other than the user, the observer can infer the presence of forces from the movements of the object. For these situations it is beneficial to have the facility to infer the presence of forces acting on an object indirectly from its visually observed movements, and not directly from the haptically felt forces. Runeson and Frykholm (1981) already concluded that visual information available in the kinetic pattern of the movement is also available as higher order properties of the optic array. Vision is thought to play a role in what is generally taken to be the privileged domain of the haptic sense combining tactile and proprioceptive cues.

Wickens and Carswell's (1995) results indicate a potential cost of dealing with information presented in more than one modality, which may explain the benefit of optically simulated haptic feedback over mechanically simulated haptic feedback for smaller targets in our experiment. Wickens, Sandry and Vidulich (1983) developed a *cognitive resource theory*. Their theory states that there can be competition between modalities during tasks, such that the human attention and processing required during both input and output result in better performance if information is distributed across modalities. Wickens, Sandry and Vidulich demonstrated that supporting tasks by different input modalities can be beneficial. Potentially, the benefits of cross-modality information may only be seen in well practiced tasks. The full benefits of optically or mechanically simulated haptic feedback may, therefore, only become apparent after extensive use.

Further research is needed to assess visual force-feedback in more complex GUI's. The active cursor technique can be applied in today's graphical user interfaces. Figure 4-5 shows a simple application of the technique on the standard window buttons in a graphical user interface. This can be tested online at www.koert.com/work/windowpoll/ (Mensvoort, 2003).

Drawbacks & Limitations

We want to emphasize that the laboratory setting in which we compared the various modalities is much simpler than a real graphical user interface, where a cluttering of intervening targets that attract the cursor could cause a drawback. Since both optically simulated haptic feedback and mechanically simulated haptic feedback will suffer from this drawback, and our current research deals with a *comparison* between the two techniques, we decided not to implement any of the proposed solutions to the problem. Instead, we assured that intervening targets do not occur in our setup. The disadvantage of having intervening force targets that can push the user in the wrong direction, has been discussed quite extensively, and various techniques have been proposed to counter this problem (Worden, 1997; Blanch, 2004; Balakrishnan, 2004). A straightforward approach is to make the stickiness or force fields underneath the targets dependent on cursor speed, assuming that when the cursor is moving at high speed the user is not likely to be anywhere near the desired target. Although this has been relatively successful, it does not entirely solve the problem and an analysis indicates that this may be a fundamentally intractable problem unless cursor trajectories can be accurately predicted (McGuffin, 2002; Balakrishnan, 2004). Analyses of cursor trajectories shows that they tend to be relatively straight and thus predictable (Oirschot & Houtsma, 2000) and a number of algorithms have been proposed to predict the target of the cursor's movement based on the cursor's trajectory (Murata, 1995, 1998; Mensvoort & Oirschot, 2001). The problem has, however, not been entirely solved. Especially smaller targets in close proximity will give problems. Until further solutions on trajectory prediction are established, application of both optically and mechanically simulated haptic feedback underneath small and closely connected targets will be limited. For instance, a scrollbar will be hard to use if the cursor is continually warped toward the window border.

Furthermore, in the current experiment we have compared usability only in the form of efficiency, effectiveness, and satisfaction of the mouse as a pointing device under the three feedback conditions. Factors like involvement or presence are omitted from the current experiment. The communicative role of haptic feedback in our target-acquisition task is limited.

Conclusions

We conclude that optically simulated haptic feedback is a good alternative for mechanically simulated haptic feedback in target-acquisition tasks. In some respects, e.g. for smaller targets, it works even better than mechanically simulated haptics. Optically simulated haptic feedback results in lower error rates, more satisfaction, and a higher index of performance, which can be attributed to the shorter movement times realized for the smaller targets. For larger targets, optically simulated haptic feedback resulted in comparable movement times as mechanically simulated haptic feedback.

Although the majority of the subjects preferred optically simulated haptics over mechanically simulated haptics and the condition without feedback, optically simulated haptic feedback is not expected to be capable of replacing mechanical force feedback devices in general. Rather we expect that, since optically simulated haptic feedback can be implemented in a standard desktop set-up without special hardware, it could catalyse the development of novel physical interaction styles and the acceptance of force feedback devices. The current WIMP interfaces are not designed with tactility in mind. Therefore, an important research path is the design of novel interaction styles based upon visual force-feedback. In chapter 5 we present a software toolkit that enables designers to create new interaction styles using visually simulating force feedback, without difficult programming.

5

Designing Interfaces You Can Touch

In the previous chapters we have measured the perception of optically simulated haptic feedback in comparison with mechanically simulated haptic feedback – showing that people are well able to perceive optically simulated holes and hills – and investigated the usability of the active cursor technique in a simple pointing task – showing that its application can be beneficial in graphical user interfaces. Furthermore we have discussed how visual force feedback transforms the cursor channel from an input only, to an input/output channel and we have hypothesized this would bring opportunities for the design of richer interaction styles.

The current chapter focusses on the design of interfaces that apply visual force feedback. Before visual force feedback can be fully brought into play in more complex graphical user interfaces, interaction designers and researchers need to further experiment with the technique. Since the current WIMP interfaces are not designed with tactility in mind, the development of novel interaction styles based upon visual force feedback should be an important research path. Unfortunately, visual force feedback is relatively difficult to program. Interaction designers with creative ideas on applying visual force feedback in graphical user interfaces might not be able to prototype their ideas due to technical barriers. It is for this reason, that we have developed a software toolkit, called PowerCursor¹³, which aims to provide interaction designers with a means to easily design and experiment with the active cursor technique. The toolkit enables interaction designers to apply visual force feedback without difficult programming. Through the distribution of the toolkit, we aim to provide interaction designers with

13 The Powercursor toolkit can be downloaded at www.powercursor.com. The software was programmed by Koert van Mensvoort, Koen Hendrix and Pascal de Man. Development of the software toolkit was supported by TU/e, Fonds BKVB and Digitale Pioniers.

a means to experiment with visual force feedback, allowing them to prototype their ideas for tactile interfaces and find out what works and what does not. Furthermore, in section 5.5 we will describe some of the interactive experiments created with the beta-version of the PowerCursor toolkit and speculate on the expected application domain of visual force feedback.

5.1 POWERCURSOR – GOALS & DESIGN CONSIDERATIONS

The goal of the PowerCursor software toolkit is to bring the active cursor technique within reach of a larger group of designers, enabling them to explore the possibilities for applications of visual force feedback. Since the spectrum of suitable applications of visual force feedback is yet to be fully identified, the toolkit should be a flexible environment that leaves room for creativity. We aim to make the toolkit itself as usable as possible for the target audience, and several core aspects of the toolkit are geared towards this purpose.

Target user group: Interaction Designers

The definition of the target user group came naturally. After having presented the active cursor technique at various international conferences¹⁴ and on the web, we were frequently contacted by interaction designers who informed us they would want to work with visual force feedback, but were unable to do the programming themselves. It quickly became clear there was a group of people eager to experiment with the active cursor technique, but for whom the technical complexity was a cut-off point. This made us realize we had to create a software toolkit, in order to remove the technical barrier for these people that kept them from designing interfaces with visual force feedback.

We define the target users of the toolkit as ‘interaction designers who are able to conceptualize and design creative applications of visual force feedback, but who are unable to program the active cursor displacements that generate visual force feedback effects from scratch. We wanted to provide this group with a software toolkit that enables them to realize their ideas, without being forced to do elaborate and complicated programming.

Through conversations with the interaction designers who had indicated they were interested to work with visual force feedback, we were able to set up a list of criteria for the toolkit. The first and most important criterion for the users, already discussed, is that they would have to do (1) ‘no elaborate programming’ to generate the visual force feedback effects. Although some interaction designers are very good programmers, the majority prefers to constrain themselves to drag and drop behaviors and simple scripting.

14 Among the events where the active cursor technique was presented were: International Browserday New York (Mirapaul, 2001), Designing Interactive systems, London (Mensvoort, 2002) and ISOC awards: best of on internet & arts, Amsterdam 2001.

This criterion is already incorporated in the starting idea (creating a software toolkit for non-programmers). A second important criterion is (2) ‘graphical flexibility’. If we want to enable designers to design more tactile interaction styles and think beyond the existing GUI language of buttons, sliders and scrollbars, it is crucial they are able to create their own visual look and graphical elements. A third criterion is what we called (3) ‘designer friendliness’, meaning that the toolkit should be a flexible environment for designers, allowing them to quickly sketch, prototype and implement their own creative ideas; the toolkit should not impose a certain fixed visual or interactive style upon its users, that would limit their room for experimentation and creative expression. A fourth criterion mentioned was (4) ‘portability’. Designers want to be able to easily distribute their newly made interfaces to others; preferably platform independent and over the web. Furthermore the (5) ‘learning curve’ of designing interfaces with visual force feedback should not be too steep. It would be preferred if the users could use the toolkit as a plug-in in their existing interface development environment. Finally, since some designers are programmers, and require control on a programming level, it would be preferred if they can have (6) ‘programmers control’ on the toolkit, meaning that they have access to lower levels of control and the internal workings of the software.

Choice of Platform: Adobe Flash

Various programming environments were considered as a potential platform to develop the toolkit; C++, Flash, Visual Basic, Java, Director, Javascript combined with DHTML. We have looked at the various platforms and derived the most suitable platform, from the users criteria discussed in the previous section. Table 2 shows the benefits of different development platforms on the different criteria.

	C++	Flash	Java	Visual Basic	Director	Javascript DHTML
1. No programming	-	+++	+	+++	++	+
2. Graphical Flexibility	+	+++	+	+	++	+
3. Designer Friendly	-	+++	+	+	++	+
4. Portability	+	+++	+++	+	+	+++
5. Learning curve	-	+++	+	+++	+++	+
6. Programmers Control	+++	+	++	++	++	+
Total	+++++	+++++++ +++++++	+++++++	+++++++ +	+++++++ ++++	+++++++

Table 5-1 Criteria in choosing a platform for the Powercursor toolkit.

The first criterion, (1) ‘no programming’ can in theory be met in all programming environments by creating an extra software layer that conceals the program code from the user. Nevertheless, clearly some environments are more opportune in this regard than others. C++ for instance is a low level hardware oriented environment and

building a software layer to conceal all the technicalities would be a huge investment. Development environments like Flash, Visual Basic and Director are more high level, presume less programming knowledge of the users and are more suitable as design environments for our target user group.

Regarding the second criterion (2) 'graphical flexibility', we learned that, unfortunately most programming languages only offer support for existing standard widgets. If one wants to create custom graphics on platforms like Visual Basic or C++ this is possible, but it has to be programmed from scratch. Since a minimum programming effort is an important requirement, we can not expect our users to take on the burden of having to program their own graphics from scratch. This means, that if we would choose a development platform with limited graphical support, we would have to implement an entire graphic design suite within the toolkit. Not practical from a development perspective and also not practical for the users who would have to learn to use this graphical suite in order to build interfaces with visual force feedback. Only two potential development platform scores high on graphical flexibility: Adobe Flash and Director. These environments specialize in the easy creation of rich interactive graphical content without difficult programming. Flash furthermore provides a rich object-oriented graphical design environment in conjunction with a programming environment, allowing users to easily create their own graphics on screen and add behaviors to them as well. Although our first active cursor demo's were programmed in Macromedia Director, it does not offer an object oriented handling of graphical elements like Flash. Director was very popular among interaction designers ten years ago, but has become less used over the last few years and seems to have been replaced by Adobe Flash. Flash is currently widely used in the (interface) design community in the development of rich interactive content and as the de facto video platform on the web.

The fact that so many interaction designers have chosen Flash as their development environment of choice is one but not the only signifier for 'designer friendliness' (3). Flash offers an very visually based development environment that can be used both at a programmers- and a designers level, meaning that you can create animations and interactive behaviors by means of drag and dropping of behaviors, but also by more low level programming.

Another benefit of Flash (and also Java and Javascript) is their extreme 'portability' (4). Whereas C++ and Visual Basic applications have to be compiled on the platform on which they are running, Flash projects are designed to run platform independently and over the web. It seems also attractive to develop the toolkit in an environment that its users are already familiar with. Many of the designers in our target user group already work with Flash, wich means they don't have to climb a steep (5) 'learning curve' before being able to create interfaces with visual force feedback.

A drawback of Flash is the relatively limited (6) 'programming control'. At first we were unsure if it was even possible to implement visually simulated haptic feedback in the Flash environment. Especially in earlier versions programming control is far from

perfect, but since Actionscript 2.0 was introduced in Flash MX it is possible to develop object-oriented projects in Flash. Nevertheless programming control in Flash is limited. It is less powerful and convenient than object-oriented programming languages such as C++ or Java. Flash has evolved into an, at times, awkward combination of programming language and graphical editor, and anyone who programs Flash's symbols in an object-oriented manner is bound to encounter some difficulties. Compilation and runtime execution are relatively slow, variables that refer to symbols have to be declared in both the code and the authoring environment, code execution is tied to Flash's symbols and time lines, and its syntax is very tolerant which makes syntax checking and debugging facilities rather limited.

Another limit in programmers control is due to Flash's security restrictions. Whereas in C++ it is possible to directly address system variables and resources, Flash uses an extensive sandbox security system to limit the transfer of information that might pose a risk to security or privacy. A Flash movie always executes inside a sandbox, and access to information outside the sandbox is severely limited. A Flash movie cannot access the user's operating system, hard disk, screen resolution, or webcam (unless, in some cases, explicit permission is granted by the local user). This poses limitations to our toolkit, particularly the fact that Flash cannot change the system cursor's position. This restriction also holds for Javascript/DHTML and online Java and is directly related to the portability and platform independent nature of these programming languages. Fortunately, Flash does offer easy ways to hide the real mouse cursor and replace it with a custom one. In principle, it is also possible to work around Flash's sandbox system by setting up a socket connection with a dedicated C++ application that can take care off the restricted operations unavailable from Flash's sandbox system – Of course, this decreases the portability, since people would have to install the C++ client. Taking all of the preceding in consideration, we concluded Flash to be the best development environment for the PowerCursor toolkit.

5.2 BASIC ELEMENTS OF THE POWERCURSOR TOOLKIT

The PowerCursor toolkit is set up to seamlessly integrate within the Flash design environment and consists of a set of library elements that can be included in any Flash project, next to the movieclips created by the user. By adding the PowerCursor movieclips on the stage¹⁵, optically simulated haptic feedback can be added to any Flash project. Each element has a set of parameters (like for instance force strength, color, etc.) that can be set to control its functioning. The powercursor engine and cursor object are required for the toolkit to function, all other objects are optional.

Essential elements: Engine and Cursor

PowerCursor Engine

As its name implies, the engine is the heart of the PowerCursor toolkit. Without it, none of the other parts will operate properly. One of the few requirements for using the toolkit is that the Engine

should always be placed in the highest level stage – called the ‘root’– of the Flash document. This way the engine can always be located by all objects and cursors¹⁶. The engine is the only entity that has a list of all the cursors, and is therefore always the link between objects and cursors. It acts as a mediator between objects and cursors, both providing the cursors’ information to the objects and propagating the object’s forces to the cursors. However, it does not simply pass on these forces but compiles them into one single displacement per frame. Objects such as hills and gutters, described further on, request every cursor’s position continuously (e.g., at every frame) through the engine and test whether any cursors are touching the object’s area. If so, the object will calculate a (2-dimensional) force for each of those cursors, and sends these forces to the engine. Every frame, the engine may thus receive a multitude of forces from different objects and targeted at different cursors. It is the engine’s task to combine these forces with each other (which can be done in several ways) and with other data. The Engine object has various parameters that determine how the forces are calculated.




Figure 5-1 Parameters for the powercursor engine.

15 The terms ‘movieclips’, ‘library’ and ‘stage’ are Flash specific terms. Flash uses the theater metaphor of the ‘stage’, on which graphical symbols – called movieclips – enter to become visible for the user. Movieclip symbols are reusable pieces of flash animation that reside in the library and can be animated on stage. The ‘timeline’ determines when a certain movieclip is on stage. This system works recursively, meaning that a movieclip itself contains a stage with the ability to animate other graphical symbols.

16 Internally, the engine is set up to be able to handle multiple cursors at a time. Although in most cases only one cursor will be on stage, this feature was added for design flexibility.

Like for instance, whether forces should be scaled along with the visibility of objects, whether forces should be ignored if the cursor reaches a certain speed, whether forces of overlapping objects are added, or only the strongest or topmost object should be able to exert a force to the cursor. These parameters can be set in the component panel of the PowerCursor Engine object (Figure 5-1). Taking into account the forces as well as the cursor's current speed, grip, inertia, and mouse movements, the engine then calculates a displacement vector for every cursor and sends it to the cursors involved. The forces are calculated according to the algorithm described in section 2.3.

ArrowCursor, CursorMaker

In order for the software to function a cursor has to be present on the stage. As described earlier, the Flash sandbox system does not allow displacing the system cursor. To bypass this limitation, the system cursor is usually made invisible and a mockup cursor that can be displaced is created on stage. The toolkit offers two types of cursors: ArrowCursor and CursorMaker. The simplest is ArrowCursor which mimics the operating system cursor arrow. Like a regular system cursor, it can change its appearance according to systems state (finger icon while hovering over buttons, hand icon during dragging, etc). Alternatively, the CursorMaker could be used. This object does not have a graphic defined, but by placing it within an existing movieclip a cursor behavior can be added to that clip. With this element it is possible to create a custom graphic that behaves like a cursor, including all forces exercised by the engine.



Figure 5-2 The toolkit allows you to choose between using the readymade ArrowCursor object (left) or to create your own custom cursor by adding the CursorMaker behavior to a movieclip (right).

Structure: Generic & Mergeable Components

In the design of the various objects in the PowerCursor toolkit, it would have been relatively easy to create objects for specific purposes, such as a square hole the size of a desktop icon, or an object that exactly fits Windows' default minimize-maximize-close buttons. However, in order to comply with the 'graphical flexibility' and 'design friendliness' criteria, we chose to make PowerCursor objects very basic, generic and combinable. The collection of basic components was chosen to facilitate the expression of countless more complex and richer cursor behaviors when combined. Their shapes are generic (round or rectangular), rotatable and scalable, their graphics can be modified independent of their functionality, and the forces they apply to cursors are configurable. Moreover, the 'behaviour' objects are designed to expand existing symbols of any shape. This functionality has been added to allow for the adding of PowerCursor functionality to any shape, logo, or other graphic in Flash. Such graphics can be made into a hole, or a wall, or a slick area by adding a 'behaviour' object to that graphic.

It is also possible to combine PowerCursor objects with other Flash symbols, including other PowerCursor objects, into new compound symbols. And by adding a special Dimmer object to other objects, specific sections of objects can be deactivated. This makes it possible to produce new shapes such as gutters cut in half or quarter-hills. Through this mechanism, one can also combine parts of objects with each other, which could for example result in a circular area which is half hill and half hole. In the following paragraphs we describe the basic ready made objects and the behaviors. Since it is the main goal of this section to describe the PowerCursor functionality and provide the reader with an overview of what can be made with the toolkit, we deliberately chose to describe these elements from a users perspective, rather than from a programmers perspective.

Basic Ready Made Objects

The PowerCursor toolkit includes a wide variety of force field objects, divided into ready-made objects and behaviours. As their name implies, ready-made objects can be added to a Flash document simply by dragging them from the library onto the stage. Ready-made objects have some default graphics and are ideal for quick prototyping or a first exploration of the toolkit. Some objects actively assert a force upon the mouse cursor (Hill, Gutter, Ramp), some only alter the mouse movement made by the user (Sand, Slick), others manipulate the functioning of other objects (Dimmer) or provide functionality to script interactive behaviors (Button). The size, rotation and visualization of the objects, as well as parameters that determine the sound and cursor icon when entering and leaving an object can be set by the designer. Below we introduce the objects available in the toolkit.

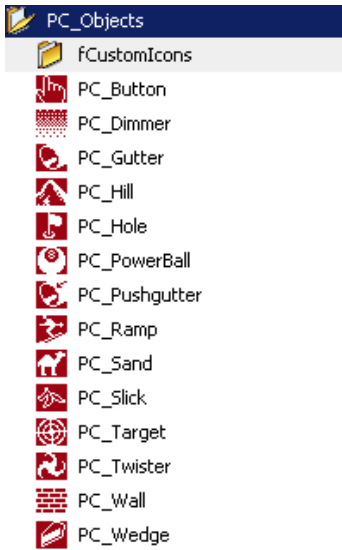


Figure 5-3 Objects can be dragged directly onto the stage from the library.

Ramp

The simplest object of all is the ramp. It simply pushes cursors in a certain direction with constant force. Such an effect could also be described as 'wind', or 'gravity' (when going down). We have chosen to consistently describe the force fields in terms of physical slopes and thus named this effect a 'ramp'. The direction of the force is specified by an angle parameter (in degrees), which is relative to the object, so when the object is rotated, the force direction rotates along.




Figure 5-4 Visualization of the Ramp object.

Hill & Hole

The Hill is a circular object that pushes cursors away from its center. A round, force-free ‘hilltop’ can be created in its center using the ‘plateau’ parameter. Depending on its Strength, the Hill can push harder or weaker, and the shape parameter determines how its force changes with distance from the center. The different slope-shapes are illustrated in Figure 5-6.




Figure 5-5 The visualizations of the Hill and Hole as they are dragged on the stage.




Figure 5-6 The available slope-shapes for the Hill, Hole, Twister and Gutter are V-shaped (constant slope) U-shaped (linear slope) Sine-shaped (sine slope).

The Hole works exactly like the Hill, except it pulls cursors towards its center instead of pushing them away. When the Hill or Hole is assigned a negative strength parameter, it will act like the other type (a Hill will become a Hole and vice-versa). Although from a programmers point of view, it might seem not logical to create two separate objects, but from a user point it is friendlier. The Hole also looks the same as the Hill, with the exception of its graphic shading.



Gutter, Pushgutter

The Gutter is a rectangular area that draws cursors towards its bottom, which is its center line. In this respect, it works like a hole, but only in one dimension instead of two. The Gutter is useful for guiding cursors along a line. If the Gutter's Strength parameter were to be set to a negative value, the Gutter turns into a dam. The different shapes available for Gutters are the same as those for hills and holes (see in Figure 5-6). The Pushgutter is a combination of the Gutter and the Slope, creating a rectangular area that draws the cursor towards its center line, while pushing it through.



Figure 5-7 The visualization of the Gutter (upper) and the Pushgutter (lower).



Twister

The twister acts much like a Hill or Hole, but instead of affecting cursors straight to or from its center, it spins them around. This can give a 'drainy' effect, where cursors twirl down the vortex like water down a drain. This can be especially useful in combination with the 'pc_plateau' event (discussed in the event section further on), which is activated whenever a cursor hits the twister's center – in other words, the bottom of the drain.



Figure 5-8 Visualization of the Twister.



Sand

The Sand is an area that is hard for cursors to get through. It can be used to simulate a sandy or sticky texture. It has both a Roughness and a Stickyness parameter: the first controls the semi-random forces this object applies to the cursor, the second governs the additional speed decrease this object imposes. The forces this object applies should hinder the cursor when it moves, but be absent when the cursor is still. To achieve this, a friction force calculated is in proportion to the cursor's current speed and in the opposite direction. Because a random factor is added as well, the exact force applied is different per frame, which suggests a rough texture.




Figure 5-9 Visualization of the Sand object.

Slick

The Slick object does not actually exert forces on cursors. Instead, it is an area with a low surface grip, which causes cursors to lose speed more slowly. Cursors will maintain higher speeds on a Slick object, giving the impression of a slippery or 'icy' surface. Setting the Strength parameter of this object to a negative value effectively turns it into a sticky surface, where cursors lose speed more quickly than normal.



Figure 5-10 Visualization of the Slick object.

Wall

While other objects act on any cursor that touches it, the Wall is an object that cannot be touched at all. When the Engine is about to displace a cursor, it first checks whether this displacement would position it on a wall. If this is the case, the displacement is canceled and the Engine will try to find an alternative movement that will not put the cursor on any wall (see Figure 7). Although this makes it impossible to move a cursor onto a wall area, it is possible to 'jump' over a wall by moving the mouse briskly. This can cause the cursor to move from one side of the wall object to the other without touching it.



Figure 5-12 Visualization of the Wall object.

Wedge

The wedge object acts like a piece of pie taken out of a hill or hole (depending on force settings). The cursor will be guided towards or away from, or towards the tip of the wedge object.



Figure 5-11 Visualization of the Wedge.

Target

The target is designed specifically to ease target acquisition of icons, buttons and other clickable interface elements. When the force enters the target it is pushed towards the center. Once the cursor reaches the inner (dish) area the force field is dimmed. The result is an object that is easy to enter and leave. Target object is a dynamic version of the hole object.



Figure 5-13 Visualization of the Target.



Dimmer

The Dimmer is a special object, since it does not apply forces itself. Instead, it can create areas that weaken or disable forces created by other objects. Dimmers add their dimming effect to a cursor just like other objects add forces to them. When forces are combined in the Engine, it inspects whether any of those forces should be weakened by a dimmer, and scales the forces appropriately. Dimmers have either a global or a local scope. Global dimmers weaken every force on the cursor they affect, while local dimmers only weaken forces from objects with the same parent symbol as the dimmer itself. Dimmers can also be used to weaken or disable the effect of the user's mouse-movements upon the cursor.



Figure 5-14 Visualization of the Dimmer object.



Button

The Button object does not apply any forces. Instead, it creates a simple button that can be used to create interactive behaviors. The button visualization can be set by the user.




Figure 5-15 Visualization of the Button object.



PowerBall

The Powerball doesn't apply any forces itself. It is a passive free floating object which movements are determined by the forces of underlying objects. It is a variant on the cursor object, except that it is not connected to the mouse movements of the user.

Behaviours

Behaviours do not have standard graphics attached to them, but can be linked to existing movieclips. Where ready-made objects compare the cursors' coordinates against their own standard graphics, behaviors compare the cursors' positions to the parent object they are part of. Thus, to use a behaviour, it should be included in an existing Flash symbol. The icon of the behaviour is made invisible during runtime. In this way, behaviours can add PowerCursor functionality to any symbol or shape in Flash. For example, it is possible to draw any custom shape in Flash using the Brush tool and subsequently make that surface sticky or rough or impassable by adding the appropriate behaviour to it (Figure 5-17).

It is also possible to add multiple behaviors to one symbol. Some behaviors, like the HillMaker, RoughMaker, WallMaker, apply tactile properties to objects, whereas others like for instance the ButtonMaker, Synchronizer and EventDispatcher add functional properties to existing movieclips in order to determine their interactive behaviour. The different behaviours are discussed below.

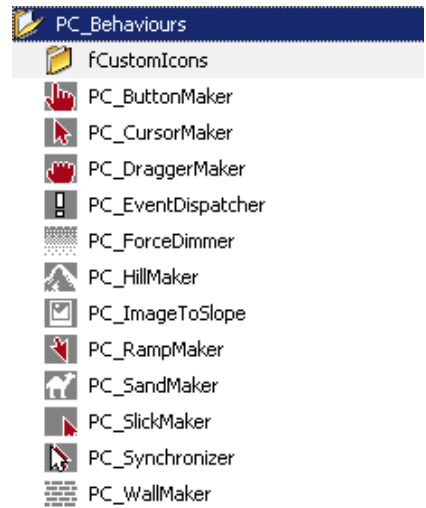


Figure 5-16 Behaviors can be dragged from the library into existing movieclips.



Figure 5-17 An example of a shape that's turned into a wall by adding a wall behavior.

HillMaker

The PC_HillMaker is a behavior with an effect similar to the PC_Hill object. Because the PC_Hill object is of a fixed size, it can displace the cursor according to various interesting slope-shapes, but the HillMaker does not have that advantage. Therefore, the PC_HillMaker simply detects the nearest edge and pushes the cursor in that direction. Because the HillMaker has no way of knowing the shape of its parent object, it has to find the nearest edge through trial-and-error. This is rather inefficient, and because it happens every frame for every cursor touching it, the HillMaker can be quite a calculation-intensive component. Having numerous HillMaker objects with several cursors touching them may lead to performance issues.

RampMaker

The PC_RampMaker works exactly like the PC_Ramp object, except that the RampMaker has no surface itself but applies the roughness to another symbol. It can be used to create a slippery or icy surface in a custom shape.

SandMaker

The SandMaker works like the Sand object, except that the SandMaker has no surface itself but applies the roughness to another symbol. It can be used to create a sandy or sticky surface in a custom shape.

SlickMaker

The SlickMaker works exactly like the Slick object, except that the SlickMaker has no surface itself but applies the roughness to another symbol. It can be used to create a slippery or icy surface in a custom shape.

WallMaker

The WallMaker is the behaviour version of the Wall object. It works exactly like that object, except that the WallMaker has no shape itself and is meant to create a wall in a custom shape. The WallMaker applies its wall-making ability to its parent object. In the example given in Figure 5-18 this object is a maze structure of which the drawn parts become inaccessible after the WallMaker behaviour is added.



Figure 5-18 An example of a simple maze created using a WallMaker that makes the red wall parts inaccessible for the cursor.

ButtonMaker

The PC_ButtonMaker does not displace cursors. Instead, it uses PowerCursor's events and an event listener to create button-like behaviour. When the ButtonMaker is touched by a cursor, it can jump to a different frame of its parent MovieClip. On a mouse click, the ButtonMaker can move to yet another frame. These target frames can be set using parameters. The ButtonMaker was included in the toolkit because the button-like use of event dispatchers it implements, will probably be a frequently used application of PowerCursor's events. This behaviour could also be created by coding an appropriate event listener in an EventDispatcher, but to accommodate for non-programmers this modifier was included separately. An example of using a buttonmaker is discussed further on in section 5.4.

DraggerMaker

With the DraggerMaker behaviour, a symbol can be made draggable, meaning that once the left mouse button is pressed while the cursor is inside the object, it will be connected to the cursor.



ForceDimmer

The ForceDimmer is a special modifier that does not use forces. Instead, it can create areas that weaken or disable forces. Dimmers add their dimming effect to cursors just like other objects add forces to them. When forces are combined in the Engine, it inspects whether any of those forces should be weakened by a dimmer, and scales the forces appropriately.

Dimmers have either a global or a local scope. Global dimmers weaken every force on the cursor they affect, while local dimmers only weaken forces from objects with the same parent symbol as the dimmer itself. Figure 5-19 shows an example of a ForceDimmer being used to disable the forces underneath the right part of a hill object.



Figure 5-19 Due to the drawn ForceDimmer area overlapping the right side of the hill, only the left part of the hill will exercise force,



ImageToSlope

The PC_ImageToSlope is a very powerful object. It can be used to transform bitmap images to force objects. This allows users to draw their own height maps, where the brightest regions of the image are interpreted as the highest, while the darkest regions are the lowest. The result is a force that pushes the cursor towards the darker regions of the bitmap.



Figure 5-20 The ImageToSlope object used to create a hill type forcefield.



EventDispatcher

The PC_EventDispatcher is an empty subclass of PC_Modifier. It does nothing more than dispatch the standard PowerCursor events that every object dispatches when its parent its parent MovieClip is touched by cursors. The Events dispatched by this object are detailed in the Event handling section furtheron. The behavior was included to accommodate programmers in tracking and handling PowerCursor events at an actionscript level.



Synchronizer

The PC_Synchronizer is a special behaviour that synchronizes PowerCursor's mock cursors with the system cursor. As discussed in section 5.1, Flash sandbox security prohibits manipulation of the system cursor position. Therefore, Powercursor hides the system cursor and shows a mock cursor of which the position can be manipulated. Over time the dislocation of the mock cursor and the hidden system cursor will increase, which can, for instance when the system cursor reaches the screen border, lead to problems. In some rare cases, for instance when working with third party

components that are programmed to respond on the system cursor, it can be necessary to synchronize the location of the mock cursor with the real system cursor. The synchronizer object aims to smoothly synchronize the two cursor locations. This is accomplished by applying forces aimed at the system cursor. This will push the cursor towards the system cursor until their disparity is negligible. The subtlety parameter can make the synchronization process more unobtrusive, but a low strength parameter also goes a long way in achieving this.

Event Handling

Many programming languages feature events, generated when a certain incident occurs, such as a key being pressed or a mouse being moved. Every time an event is triggered, code that is specifically connected to that event (a so-called event handler) is executed. Flash has some standard built-in events, such as `keyPress`, `mouseDown`, and `rollover`.

Flash furthermore features an interface panel with which numerous standard actions can be tied to such events. This so-called ‘behaviours panel’ allows Flash users to use the power of events without the need to manually code event handlers. For example, through a few mouse clicks in the behaviours panel, a designer could link the action “Go to frame 7” to the event “Mouse roll over”. This would cause the selected `MovieClip` to jump to frame #7 on its timeline whenever the mouse rolls over it.

The built-in mouse events that Flash uses in the Behaviours panel, such as `mouseDown` and `rollover`, are related to the system cursor. As the `PowerCursor` system uses mock cursors instead of the system cursor, the built-in events do not function properly on `PowerCursor` cursors and objects. We do still want to offer Flash designers the convenience of using events and behaviours, especially without coding. Therefore, several events were recreated in `PowerCursor` versions; specifically `pc_rollin`, `pc_rollover`, `pc_rollout`, `pc_mousedown`, `pc_mouseup`, and `pc_plateau`.

All `PowerCursor` events are generated consistently across the toolkit; every ready-made object and every modifier in the `PowerCursor` toolkit works with them. Because the dispatching of events and all related administration is implemented in the `PC_Object` parent class, all ready-mades and modifiers also send out the same events with the same data at the same incidents. The only difference is that the ready-made objects create the events themselves, while the modifiers dispatch the events from the symbols they modify. So whether a hill comes directly out of the `PowerCursor` toolkit or was hand-drawn and modified into a hill, they both dispatch the same events. The specifications of these events are listed below.

`PowerCursor` events are dispatched in a universal way. Some `PowerCursor` events are dispatched from cursors, others from objects and behaviours. Because the events are implemented in the `PC_Object` and `PC_Cursor` superclasses, all individual subclasses handle events in the same way.

Cursor Events

Cursors dispatch three different events, related to the forces applied on the cursor. These events can be send out once per frame per cursor at maximum.

Event Name	Description
PC_ForceStart	Triggered when objects try to apply forces to this cursor, and this did not happen in the previous frame.
PC_Force	Triggered when objects try to apply forces to this cursor
PC_ForceEnd	Triggered when no object tries to apply forces to this cursor, although this did happen in the previous frame.

Table 5-2 Cursor events available in the Powercursor toolkit.

Object & Behaviour Events

Six distinct events are dispatched by objects when a cursor is entering, within, leaving, or when the mousebutton is pressed or released above the object and lastly, when a cursor is rolling over a force-free plateau that some objects have in the centre. These events can be dispatched once per frame per object per cursor at maximum.

Event Name	Description
PC_Rollin	Triggered when the number of cursors touching this object increases.
PC_RollOver	Triggered whenever any cursor touches this objects.
PC_RollOut	Triggered when the number of cursors touching this object decreases.
PC_MouseDown	Triggered when, at the moment the mouse button is pressed, any cursors touch this object.
PC_MouseUp	Triggered when, at the moment the mouse button is released, any cursors touch this object.
PC_Plateau	Triggered whenever any cursor hits the force-free plateau in the centre of a certain object.

Table 5-3. Object & Behaviour events available in the Powercursor toolkit.

Events Object

All object events carry an event object with them, which contains information on the cursor that triggered the event. (The cursor events have no event object.) The event object contains information about the event type, the related object, the related cursor and its position and last displacement.

Property	Type	Description
Type	String	The event type, such as PC_Rollin or PC_MouseDown.
target	String	The path to the object that dispatched this event.
X	Number	The x-coordinate of the cursor that caused this event.
Y	Number	The x-coordinate of the cursor that caused this event.
xLastDisp	Number	The last horizontal displacement of the cursor that triggered this event (a positive numer denotes a rightward movement).
yLastDisp	Number	The last vertical displacement of the cursor that triggered this event (a positive numer denotes a downward movement).
Path	String	The path to the cursor symbol that triggered this event.

Table 5-4. Elements of the event object.

To access the event object, an event listener needs to be coded in Flash. The code required for such an event listener may differ for distinct Flash versions but instructions can be found in the accompanying Flash documentation.

Sound Design Support

In everyday haptic communications, sound is usually directly connected to the interaction between the user and the touched object – think for instance of the sound caused by scratching a rough surface and how that enhances, or at least emphasizes, the haptic perception. Like in the physical world, sound can be a crucial component in graphical user interfaces, thus it is no surprise Flash offers extensive support for interactive sound design. However this support is limited to dynamically starting and stopping sounds from the library in response to interactive events like mouse clicks and rollovers, and not fine-tuned to haptic feedback. In order to simulate the typical haptic sound experience, we have added a special sound module to the powercursor engine that allows for the creation of sound effects that dynamically change along with the forces exerted on the cursor

Every object can be linked with a sound effect that is enabled once the cursor is over the specific object and changes its volume along with the forces caused by the object. This makes it possible to create realistic rolling and scratching sounds that are synchronized with the forces applied by a specific surface, like for instance, a rolling sound that becomes louder when the cursor *rolls* into a hole object with an increasing speed. A library of various standard sounds is distributed along with the powercursor engine; rolling, bouncing, etc. In addition to using the library sounds, it is possible to create and link your own sound effects to engine objects.

Debug Tools

Besides objects and behaviours that assert forces on the cursor or facilitate interactivity, the toolkit also includes some objects for debugging purposes.



Watcher

The Watcher object retrieves information about a single cursor from the engine and displays it on screen. It shows the cursor's position and speed, which objects try to affect the cursor, and the current forces acting on it.



PC_ForceSensor

The ForceSensor is a simple object used to monitor a force at a fixed position. The arrow shaped ForceSensor object rotates and resizes itself according to the exiting forces underneath. It can be placed at a fixed location on the stage to monitor the forces at this location.



PowerGrid

The PowerGrid is a much more powerful object than the ForceSensor, since it can monitor forces not only at a specific location, but also at a selected area of the stage. It visualizes all the forces underneath the grid area by distorting the grid according to the forces active underneath. Users can set the size and resolution of the grid and drag it during runtime. Especially when working with complex composed force objects the PowerGrid is a handy tool to visualize the active forces.




Figure 5-21 The Watcher object, showing the cursor position and speed as well as forces applied by the object underneath.




Figure 5-22 A powergrid object showing the forces of a whirl, hole and hill positioned underneath.

5.3 COMBINING ELEMENTS: LIKE LEGO

The PowerCursor objects and behaviours have been designed to provide the basic visual force feedback functionality. The modularity and flexibility of the objects and the Flash development environment makes it possible to compose more complex object from the objects and behaviours in the toolkit.

Besides the ability to scale and rotate objects to the desired size and position, forces over a desired part an object can be dimmed using a custom drawn forcedimmer allowing new objects shapes. Furthermore forces exercised by objects and behaviours can be added or – depending on the engine setting – overruled using only the strongest force or the force of the topmost object. In combination with Flash's functionality of timelines an almost endless spectrum of dynamic force objects can be created. In this paragraph we describe some examples of force objects that can be created using basic powercursor elements.

Example 1: Hole-Hill

A simple example of a combined object is the hole-hill. Suppose you would want to create an object of which the border functions as a hole – meaning that it attracts the cursor towards its center – while the center of the object should have the properties of a hill – pushing the cursor away from the hilltop. Such an object is not included in the PowerCursor toolkit's ready made objects, but can be created by combining a hole and a scaled hill object. Figure 5-23 shows how the force fields of the hole and hill objects can be combined to create the desired hole-hill behavior.




Figure 5-23 a) The Hole-Hill forcefield can be created by combining a hill and a hill forcefields (dotted lines). b) Shows how the hole and the hill ready made objects are combined. c) Shows a shaded rendering of the resulting force field.

Example 2: The Pit

Suppose you would want to make an object with the opposite behaviour of the wall – a readymade object in the powercursor library that allows for the creation of regions that cannot be entered by the cursor. Inverting the metaphor of the wall, would result into a cliff, or a pit; an object you can enter but cannot leave. The toolkit doesn't have the pit among its readymade objects or behaviours. Still it is quite easy to create the desired pit object using the buttonmaker and wallmaker behaviours. Instead of applying a wallmaker behaviour to the shape of the object, we will have to create a wall object, which forms a ring around the pit. Figure 5-24 shows how the wall area around the pit is disabled by default (left picture) and enabled once the cursor enters the pit (right picture). The buttonmaker object is used to move the playhead to the timeline position that contains the wall to achieve this. In order to avoid being trapped in the pitt forever, the user can click while inside the pitt, which makes the buttonmaker transference the playhead to the 'down' label, where the impenetratable area around the pit is removed again.

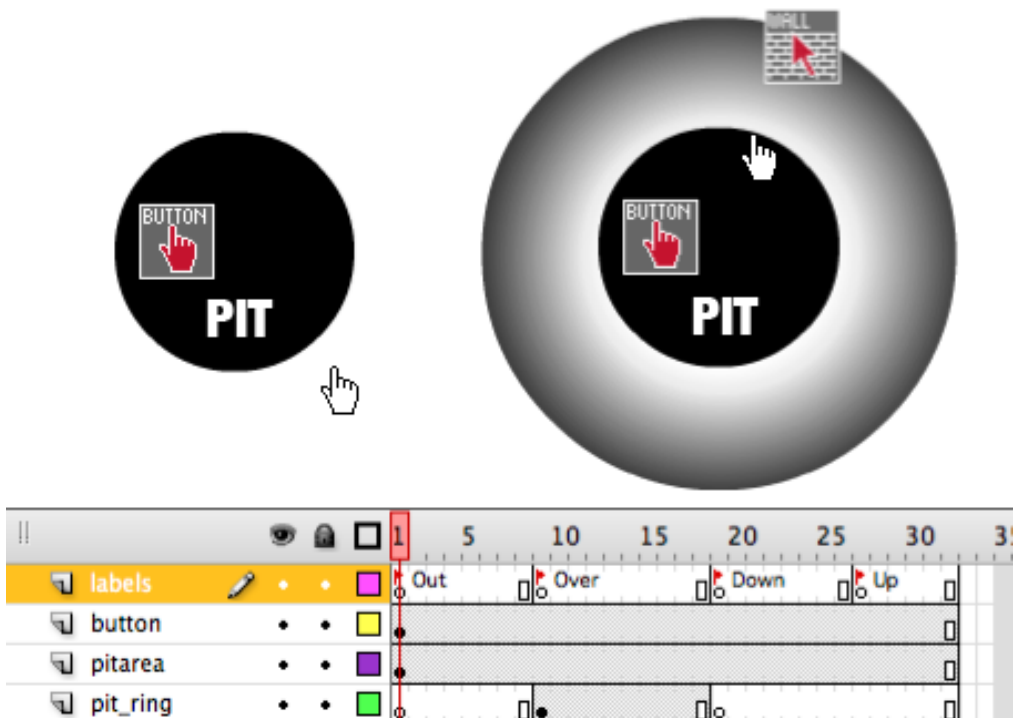


Figure 5-24 Creation of a Pit object using a Wallmaker and a Buttonmaker: Once the cursor enters the pit, the buttonmaker moves the playhead to the 'over' frame, where the pit_ring keeps the cursor within the pit area. After pressing the mouse the playhead is transferred to the 'down' part of the timeline and the cursor can move out again.

Example 3: Dynamic Button

As a last example we show how to create a dynamic button that changes graphics and forces as the cursor moves over, and the button is pressed. A buttonmaker behaviour is used to make the playhead jump at the 'out', 'over', 'down' and 'up' labels at the various stages of interacting with the object. As the mousebutton is pushed down while the cursor is over the button area, the shape of the button is changed and a hole force field is activated underneath the button, which pulls the cursor towards the center of the button. (Figure 5-25, Frame 15). As the mousebutton is released, the button graphics are changed to its original shape and a hill shaped force field, which pushes the cursor outside the button area, is activated (Figure 5-25, Frame 23).




Figure 5-25 How to create a button that attracts the cursor when the mouse rolls over, but pushes the cursor away after the button is pushed. From left to right the 'out', 'down' and 'up' stages of the button are displayed.

5.4 EXPLORING THE APPLICABILITY

In the current section we discuss some of the experimental interaction styles that have been prototyped with PowerCursor so far and explore the expected application domain of visual force feedback. Although the 1.0 version of the PowerCursor toolkit is not yet released – as we write – various interaction designers have already been experimenting with a beta-version¹⁷ of the toolkit. The preliminary design experiments that were conducted during the development stages of the toolkit, not only helped us to iteratively develop the features of the toolkit, but also generated ideas for novel interaction styles that use visual force feedback. Since the final version of the toolkit is yet to be released, the overview of the applicability presented here, is likely to be incomplete. However, it already provides a preliminary framework of the various uses of visual force feedback and might bring new ideas for future applications.

Assisted Navigation

While thinking about what we could learn from touch interaction in everyday life, the aspect of navigation emerges as a central area. Knowing where you are, where you came, where you could go. In chapter 2 we already discussed that with the active cursor technique, the cursor channel is no longer an input channel only, but is transformed into an input/output channel. Visual force feedback can be used to communicate textures and slopes to the user, which may represent material properties or states of the system, that provide the user with contextual feedback while navigating the interface. Among the first virtual haptic objects we created are ‘holes’ and ‘hills’ (Figure 2-14). If the cursor rolls over a hole, it is dragged towards the centre. When rolling over a hill, the cursor is dragged away from the centre. Due to these cursor displacements a hole becomes an easily accessible part of the screen whereas a hill area is hard to access. In chapter 4 we showed that placing hole shaped force fields underneath targets, enhances the usability of a pointing task in a graphical user interface (Figure 4-5). In this case, the user’s navigation in the graphical interface is assisted on a fairly rudimentary level; there is no intention from the side of the system, it just aims to help the user reach its goal.

Assisted navigation can also be applied on a semantic level, whereby the system not just supports the user to simply reach its goal, but also provides contextual feedback to communicate properties of the system. As an example we describe a drop down menu. In current WIMP interfaces the unavailability of an item in a drop down menu is usually communicated by rendering the button gray. Since this button is disabled and cannot be clicked, it seems to make sense to also make them less reachable than the enabled buttons in the menu. In chapter 2 we already described that visual force

17 Since the launch of the first beta in April 2007, over 360.000 people (counting: 1-2009) have visited the website. Although by far the largest majority of the visitors only looks at the online demo’s, about 2% of the visitors, i.e. 7.200 people, have downloaded the toolkit.

feedback can be applied to communicate textural properties. The disabled menu item can be made less reachable by simply placing a slick texture underneath the disabled button, which causes the mouse to slide over with more speed while navigating the menu. Such a slick area effectively decreases the motor space underneath the disabled buttons, while keeping the visual space unchanged. Figure 5-26 shows a drop down menu in which two options are made both grey and slick.

Mixed initiative and persuasive interfaces

In chapter 2 we discussed how, according to Laurel (1991), both the computer and the human can be seen as active agents working together to achieve some common goal. When assisting the user in a pointing task, as in the previous section, the computer merely facilitates the user's behavior. However, visual force feedback can also be put to work to give the computer a more active and intelligent role. Once both the user and the system can take the initiative to undertake actions, Horvitz (1999) proposes to speak of *mixed initiative* interfaces.

Visual force feedback increases the possibilities for implementing mixed initiative interfaces. Hendrix (2006) proposes visual force feedback could be applied in so-called 'wizards', which are often used to guide users through a complex process such as configuring a network connection or installing an application. In such wizards, the user is sequentially presented with a number of choices or input fields. If there is only one choice or input field on screen, the system could take the initiative to guide the cursor towards it, and in the case of more input options the cursor could be guided to a suggested or most likely option. Computer-initiated mouse movement might be especially helpful to assist users when they try to do things that require some other action elsewhere first; for instance, trying to log in without entering a name or password, or trying to continue an installation before entering a necessary serial number. In such cases, the computer could censor access to the inactive dialog box and point the cursor to the required input field, signaling the user that it requires some action there before continuing.

A subset of the mixed initiative dialog is the so-called *persuasive computing interface* (Fogg, 1998) in which the system aims to influence the user's behavior. With visual force feedback this can be realized literally by autonomously pushing the cursor into a recommended direction that the user can passively confirm or actively overrule. In this case the system is actively trying to impose decisions upon the user. By pushing the cursor, the system can lead the user to the next point where it wants his or her input, or suggest particular options by moving the cursor over them. Particularly in situations where the user has to make a choice (such as Yes/No/Cancel confirm dialogs), cursor displacements could steer the cursor toward the recommended option. Figure 5-27 shows an example of a dialog where – after the user has pressed a button that would launch an ethically irresponsible nuclear attack – the system tries to persuade the




Figure 5-26 In dropdown menus, disabled items are often made grey to communicate they can not be selected. Using visual force feedback, they can also be made slick, which makes the cursor slide over them quickly, towards selectable items. Test online at: www.powercursor.com/examples/?example=menu.swf




Figure 5-27 After clicking the red button – which initiates a highly unrecommended action – the system asks for a confirmation, whereby the cursor is pushed towards the ‘No’ button which cancels the action. An example of a persuasive dialog.

user to cancel his command, by pushing the cursor towards the ‘no’ button that will abort that action when clicked. If sure of the intended action the user can still move the cursor towards ‘yes’ and continue with the action, but this will take a larger effort than simply clicking on the ‘no’ button already underneath the cursor. Obviously the example of a ‘nuke’ button is fairly extreme, however advisory confirmation might also be useful when initiating possibly dangerous operations like deleting a program or formatting a hard drive. Through the use of visual force feedback the system can in a subtle way communicate to the user that an unrecommended action is being conducted. If consistently and mindfully added, various low level persuasive interface elements can possibly be combined to higher level persuasive dialogs, which coach the user towards a certain behavior.

Of course one can easily imagine these kind of ‘mixed-initiative’ or ‘persuasive’ interaction styles could be bothersome as well; think for instance of a web banner that catches the cursor and forces the user to click on an advertisement. It is clear that the question ‘who is in control’, is an important issue here. From force feedback literature [referenties!], we know that while people may appreciate advisory haptic feedback, they find it important to have the feeling of being in control and are empowered to overrule the forces as they wish. One would expect the same with visual force feedback.

Interpassive Interfaces

Taking the principles of mixed initiative to an extreme, one could use visual force feedback to create an interface that does not require any action from the user at all and allow the user to toggle between a ‘lean forward’ controlling attitude and a more passive ‘lean back’ mode in which the system leads. As an example, one could think

of an *interpassive* video application that combines interactivity with classic linearity, resulting in a hybrid between TV and PC. Potentially such an interpassive approach could lead to a calmer and more fluent interaction; the user is still in control, but isn't forced to be active and make decisions all the time.

It could furthermore be applied to have video material interactively respond with visual force feedback to the users actions. Figure 5-28 shows an example in which an actor in a video responds once the cursor hovers over his face and literally pushes it away.



Figure 5-28 The actor in the video responds when the mouse (like a fly) touches its face and waves the cursor away. © Olivier Otten, www.selfcontrolfreak.com.

Aesthetical Interactions

When aiming to improve the quality of graphical user interfaces, one by default thinks of techniques that make them more efficient and effective. However there are more elements that can also increase the quality of an interface. Computers are traditionally perceived as office machines, but over the last few decades, they have expanded to become general purpose communication and entertainment machines and have penetrated domains, like for instance the home, where efficiency and effectiveness aren't the sole criteria of quality. In these emerging computing domains, other virtues like aesthetics and style are also highly appreciated and this change of context has reflected on the design of both the graphical user interface and the computer hardware, which in recent years have become much more aesthetically refined¹⁸. This increasing weight of aesthetics also influences the interaction. According to Media theorist Lev Manovich, interaction increasingly becomes an *aesthetic event*. As an example, he describes how he switches his LG Chocolate phone on and off much more than is 'functionally' necessary – being so mesmerized by the simple act of switching the device on (Manovich, 2006; 2007). Phones like the LG Chocolate are not marketed on their technological features or their functional usability, but on their high quality styling and interaction. In an experience economy (Pine & Gilmore, 1999), aesthetics can have a decisive marketing

18 An illustrious example of the computer as an aesthetical object was the Apple iMac – introduced in 1998 – of which Apple famously declared that “the back of our computer looks better than the front of anyone else’s”.

value. According to Djajadiningrat et al. (2004) an important factor for aesthetics of interaction is the involvement of other sensory modalities next to the visual, i.e., the haptic and auditory quality of materials and controls. As examples they give the push and feel of a button or the sound of a compact disk opening.




Figure 5-29 Interaction as an aesthetic event: The mundane task of setting a switch is aesthetically enhanced by a subtle force under the switch that drags the cursor along as the button is being switched.




Figure 5-30 Digital version of the classical 'ball in pocket' game, simulating the physics of balls and holes.

The simulated tangibility introduced with visual force feedback can be applied to create richer interactions that are aesthetically pleasing for the user. A simple yet elegant example is the tactile switch depicted in Figure 5-29, which slightly drags the cursor along when pressed, providing the user with a subtle form of haptic feedback, which has no functional goal besides being aesthetically pleasing. Figure 5-30 shows a digital version of the classical 'ball in the pocket' game, which uses the ability to express material properties for entertainment purposes.



Figure 5-31: Nintendo Wii Remote controller and an onscreen keyboard selection tool. Visual force feedback could be applied to create richer interactions or improve the rather imprecise pointing in this setup.

Outside the desktop

In concordance with the research scope, defined in chapter 1, this study predominantly focuses on improving the interaction of WIMP based interfaces. However, this does not mean visual force feedback cannot be applied in contexts, beyond the desktop model of Windows, Icons, Menus and Pointing (WIMP). The general idea of optically simulated haptic feedback is to utilize the decoupling of the input and output space to evoke a haptic illusion. In the WIMP-setup the input space is defined through the mouse and the output through the visual display. In digital environments, where the input and output spaces are directly connected – think for instance of a touch screen – the opportunities of applying visual force feedback will probably be limited. However, the principle of the technique can be applied in any graphical user interface setup where the input and output space are decoupled. One could for instance think of specific VR contexts like the Nintendo Wii system, which uses physical gestures of a remote controller, rather than a mouse, to control the interface (Figure 5-31). The active cursor technique could be used to create richer interactions or improve pointing in this interactive setup, which tends to be rather imprecise. In general, screen based interaction styles that detect and respond to the users movements could be enriched with visual force feedback.

Another application environment beyond the desktop computing model, might be mobile devices with embedded tilt sensors, where visual force feedback could be applied to evoke tactile effects in response to the movements the user is making with the device. Figure 5-33 shows an example of material expression on a cellphone with a simulated fluid that moves along with the phone's movements. Figure 5-32 shows an application of the 'ball in the pocket' game on tilt sensor enabled cellphone.



Figure 5-32: Photomontage of a digital version of the classical 'ball in the pocket' game, already described in Figure 5-30, on a tilt sensor enabled phone.



Figure 5-33: Water in your phone would normally be rather devastating, but with the NEC N702iS, it is turned into a joy. Thanks to the embedded tilt sensor the fluid moves on the screen as the phone moves. The water level drops as your battery level goes down. Tilt sensor enabled devices, like the iPhone & iTouch, are becoming increasingly available on the market.

5.5 DISCUSSION

In the current chapter we have presented a Flash software toolkit that allows interaction designers to apply optically simulated haptic feedback without difficult programming. The toolkit contains different readymade objects, such as hills, holes, gutters, rough and slick surfaces, as well as behaviours, which can be added to connect haptic effects to existing movieclips. The toolkit has been built to integrate into the Flash design environment, which is commonly used by interaction designers. It allows non-programmers to add visual force feedback to their own interface designs. We have created this toolkit to give interaction designers a means to develop a language of suitable visual force feedback interfaces, without the limit of having to do any elaborate programming. The toolkit consists of a collection of readymade objects and behaviours that represent force objects, interactive elements and debug objects. These objects and behaviours from the PowerCursor library can be combined and merged in order to create more complex visual force feedback objects and behaviours. This setup allows the users of the toolkit to create their own force objects and apply them in their own interfaces, or even redistribute them to other users of the toolkit.

Although the toolkit is yet to be released among a larger community of interaction designers, the beta version already helped in generating ideas and demonstrators for possible applications of visual force feedback.

Among the ideas currently developed, assisted navigation is presumably the most important application. In chapter 4, we have already experimentally shown that the technique can be applied to assist the user in a standard pointing task. It is expected that assisted navigation is also applicable in other interface elements – in paragraph 5.4 we have described the example of a dropdown menu, where disabled buttons are made slick and less accessible.

Once the system takes a more active role in its feedback, we speak of ‘mixed initiative’. Visual force feedback could help implement principles of mixed initiative, for instance when the system recommends a certain action by pushing the cursor towards the button that initiates the recommend action – we have prototyped an example of a dialog where the system aims to persuade the user to cancel an unrecommended action, by moving the cursor towards the cancel button. Mixed interaction could also be useful in installation wizards, where the user is guided through a series of actions in which the system advises the user the default option.

Besides functional improvements in interface dialogs, visual force feedback can also be applied to create richer interactions. While computers have moved from the office into other domains like the home and entertainment domain, other issues of interaction like aesthetics, expressiveness, and playfulness have become increasingly important. Visual force feedback is expected to have a range of applications that bring richness into the interaction between the system and the user. Besides its ability to express material properties like texture, elasticity, stiffness and mass, visual force feedback can be used to convey properties of 3-dimensional objects. Among the examples we have prototyped are a flipable wooden cube and a globe that can be turned around (see chapter 2).

Finally we have shortly explored some of the possible applications of visual force feedback outside the desktop computing model. Although our research deliberately focused on the WIMP interface domain, our technique to optically simulated haptic perception might also have applications outside the desktop. One could for instance think of systems like the Nintendo Wii – where input and output are also decoupled – or a mobile phone with an embedded tilt sensor, for which we have proposed a simple gaming concept in paragraph 5.4.

From the successful prototypes created so far, we can conclude that the PowerCursor toolkit is a usable tool to prototype visual force feedback enabled interfaces. The explorations of the application domain furthermore learned that visual force feedback is applicable in WIMP interfaces and possibly also beyond the desktop computing model. However we want to emphasize that only a few applications have been prototyped until now, so we will make no claims on the completeness of the range of applications as presented in this chapter. The toolkit is published under a creative commons license and can be downloaded for free on the website www.powercursor.com. With the release of the toolkit we enable a larger community of interaction designers to experiment with visual force feedback without having to do elaborate programming. We hope that, on the long term, this contributes to the development of a richer and more physical paradigm of graphical user interfaces.



Figure 5-34 Screenshot of the website www.powercursor.com.

6

General Conclusions & Future Directions

The current chapter aims to bring together the research conducted in earlier chapters towards a general conclusion and a vision towards the future. This study started with a personal fascination with the simulated ‘reality’ of the graphical user interface and a desire to enhance the materiality of this virtual environment. The idea for our research project emerged when, while working with mechanical force feedback devices, the thought occurred that ‘what was seen on the screen’ played a role in the haptic feedback simulated by these devices. This understanding led to the question whether it might be possible to generate a haptic experience, with solely visual means. And if so, whether this could bring opportunities to increase the quality of existing graphical user interfaces, without resorting to special hardware.

6.1 RESEARCH RESULTS

It is concluded haptic feedback can be optically simulated. Users are able to recognize haptic structures simulated by applying active cursor displacements upon the users mouse movements. The active cursor algorithm, described in chapter 2, can be used to optically simulate various slopes as well as dynamic slopes, textures and properties of 3D objects. This technique of optically simulating haptic feedback opens up an additional communication channel with the user in a standard graphical interface, to give the user additional feedback while navigating through the screen.

In chapter 3 we empirically researched the perception of optically simulated bumps and holes in comparison with the bumps and holes generated with a mechanical force feedback mouse. Participants in the experiment successfully identified the bump and hole slopes, in both the optically simulated haptic feedback condition, as well as in the condition using a mechanical force feedback device. Optically simulated haptic

feedback can further be applied to alter the perception of mechanically simulated haptic structures. In some respects, e.g., for more subtle forces, optically simulated haptic feedback is likely to be even more expressive than mechanical simulation of force feedback, at least for the ranges tested in our study. Furthermore, we have learned from the experiment that optically and mechanically simulated haptic feedback must be applied in a coherent way. If not, different users will react differently and hence, unpredictably.

Regarding our technique of simulating haptics visually, we have to realize, that while the subjects are moving the target, the optically simulated haptic feedback as provided on the screen is not just a matter of ‘seeing forces’ through visual information only, but is actually an integration between on the one hand the force information as computed by the active cursor technique, and on the other hand the information the users have about their own movements. It is expected that both aspects play a role in the ‘visual force feedback’ as perceived by the subject.

In chapter 4 we tested and established the usability benefits of the active cursor technique in a simple Fitt’s law style pointing task. It is concluded that putting hole-shaped force fields underneath targets increases the usability of pointing in a graphical interface. Optically simulating the force fields is a good alternative for mechanical simulating them with a dedicated force feedback device. In some respects, e.g., for smaller targets, the optical simulation is even more effective than the mechanical force feedback. Optically simulated haptic feedback results in lower error rates, more satisfaction, and a higher index of performance, which can be attributed to the shorter movement times realized for the smaller targets. For larger targets, optically simulated haptic feedback resulted in comparable movement times as mechanically simulated haptic feedback. Furthermore the majority of the subjects preferred optically simulated haptics over mechanically simulated haptics and the condition without feedback.

Our technique of simulating haptic feedback optically opens up an additional communication channel with the user that promises opportunities for novel interaction styles in a standard graphical user interface setting. In chapter 5 we have presented a software toolkit, called PowerCursor, which allows interaction designers to add visual force feedback to their interfaces, without having to do elaborate programming. The toolkit consists of a collection of ready-made objects and behaviours that represent force objects, interactive elements and debug objects, which can be combined in order to creatively design novel visual force feedback behaviours. The toolkit can be downloaded for free at www.powercursor.com.

Using the beta-version of the toolkit, we have furthermore explored the applicability of optically simulated haptic feedback. Besides the results from chapter 4, which provided evidence of the technique being applicable in assisted navigation, demonstrators were created that suggest visual force feedback might be applied to create so-called

mixed initiative interfaces. Furthermore we have experimented with using visual force feedback to express material properties, to create aesthetically pleasing interactions – which with the migration of computers into other domains than the office environment are becoming more relevant. Unlike techniques of simulating haptic effects based on manipulation of the cursor gain (Keyson, 1997; Worden, 1997; Ahlström, 2002; Lécuyer et al., 2004), our active cursor technique also works when the user is not moving the mouse; just as for actual force feedback devices that can generate force when the user is only passively holding them. This creates a more realistic simulation of device generated force feedback, that can possibly give more freedom in interface design. However, before optically simulated haptic feedback can be fully applied in more complex interface settings, the application domain still needs to be explored and evaluated further. The creation and distribution of the PowerCursor toolkit is geared towards this purpose.

Drawbacks of optically simulated haptic feedback

When considering the drawbacks of optically simulated haptic feedback, one constraint is immediately obvious. Optically simulated haptic feedback only makes sense when the user can look at the cursor. The effect of optically simulated touch can be applied only in combination with a visual display on which the user must fix his or her attention. One of the major benefits of force feedback devices is their ability to relieve the overloaded visual perception channel. In today's visual culture (Mirzoeff, 2003; Gerritzen et al., 2004) this is not only an interest of people with visual disabilities. It is evident that optically simulated haptic feedback does not share this advantage.

Another expected drawback of optically simulated haptic feedback is its incapacity to simulate a static force feedback situation, in which the user overrules the movement of the force feedback device by exerting a static force in the opposite direction, without moving. With optically simulated haptic feedback this is impossible; the user will always have to make an active movement to compensate for the force applied to the cursor. This might result in less control and satisfaction.

Notwithstanding the outcome of the current study, we do not expect optically simulated haptic feedback will outperform mechanically simulated haptic feedback in all cases, let alone be able to replace haptic feedback. Touch can play a powerful role in communication. It can offer an immediacy and intimacy unparalleled by words, sound or images. The firm handshake, an encouraging pat on the back, a comforting hug, all speaks to the profound expressiveness of physical contact. In the real world, touch can further serve as a powerful mechanism for reinforcing trust and establishing group bonding (Burgoon et al., 1984; Burgoon, 1991). Various researchers have explored the opportunities to enhance communication between people in a computer mediated environment through mechanical force feedback devices (Brave and Dahley, 1997; Rovers and Van Essen, 2004). The role optically simulated haptic feedback may play in more complex haptic interactions remains to be investigated.

Although both optical and mechanical techniques are replicating haptics up to a certain level, they are only capable to reproduce a portion of the wide spectrum of haptic perceptions people can experience. The limited expressiveness of technologically mediated haptics is often omitted in virtual reality research, we feel it is important to emphasize that both mechanically as well visually based haptic simulation techniques do not have the same sensorial richness of unmediated haptics – think of an embrace or a kiss from a lover, or the warmth of the sunlight on your skin. Haptic technologies, up til now, are not really sophisticated in comparison with our wide range of human sensorial abilities.

6.2 FUTURE DIRECTIONS

Considering the phenomenon of optically simulated haptic feedback, further research into the possibilities and limitations is needed. Our research shows that people can recognize optically simulated bump and hole structures generated through active cursor displacements and that, depending on the simulated strength of the force, optically simulated haptic feedback can take precedence over mechanically simulated haptic feedback and also the other way around. We need to learn more on the relation between mechanically and optically simulated haptics. Possibly, optically simulated haptic feedback can be applied in combination with mechanical haptic devices, whereby the optical simulation would be used to enhance, or nuance its mechanical counterpart. This idea is especially promising since we have found that optically simulated haptic feedback is more effective in expressing the more subtle, weaker forces than mechanically simulated haptic feedback (see chapter 3).

In general we need to learn more on the haptic expressiveness of optically simulated haptic feedback. Further perceptual experiments which might deal with the recognition of various objects, from simple forms, like ramps, squares, gutters, and triangles, to dynamic complex scenes. Furthermore it would be helpful to learn more about the psychophysical processes causing the haptic percepts. Although prior research literature is available to explain the effect (see chapter 3), it would be good to obtain further and more specific knowledge on the human perceptual processes responsible for the experienced haptic percepts.

Regarding its applicability, we have scientifically shown that optically simulated haptic feedback can improve the usability of a pointing task. However, unlike most real mouse operated graphical interfaces, our experiment, presented in chapter 4, was conducted in a very simple environment. If we are to employ optically simulated haptic feedback in more complex interaction styles, the applicability needs to be researched further. With the active cursor technique, the cursor channel is no longer an input channel only, but is transformed into an input/output channel. The current interfaces are not designed with tactility in mind. The most important research path, therefore, is the design of novel interaction styles based upon optically simulated haptic feedback. In chapter 5 we have made a modest start with these explorations. Interface designers

and researchers need to experiment more with the technique in order to explore the affordances of the created objects and find out what works and what does not. The PowerCursor toolkit, which enables designers to add optically simulated haptic feedback to their interfaces without difficult programming, is geared towards this purpose. Future work may consist of the design and evaluation of applications that use optically simulated haptic feedback, in particular in more complex graphical interfaces, both within and beyond the WIMP paradigm. It is evident that the principles of optically simulated haptic feedback could also be explored in other computing settings than the WIMP interface we have focused on in this study. In chapter 5 we have argued the technique should certainly work in any interface setting where the input and output are decoupled. Furthermore, we expect interactive animations can also be employed to reinforce the illusion of substance in other computing environments, like for instance tilt sensor enabled mobile phones or touch screens (see also section 2.2).

Simulate your new computer on your old computer

We think manufacturers of haptic devices can benefit from optically simulated haptic feedback. Although the advantages are clear, force-feedback devices have not made it to the average desktop. It might be because of the lack of software applications for these devices. And software is scarce because people do not have force feedback devices at their homes. The optically simulated haptic feedback technique could break this vicious cycle. Software developers and manufacturers of mechanical haptic devices could create software drivers, that are flexible to visually or mechanically simulate haptic feedback – depending on the capabilities of the users input device the haptic feedback can be evoked via the mechanical device, with cursor displacements, or both. If interaction designers can assume the availability of a haptic device (simulated by cursor displacements or not), the use of haptic information can grow to become a serious factor in human/computer interaction. Once interface designers can count on its presence, haptic feedback can grow to become a default communication channel in human/computer interaction.

There is more in the box than there is in the box

Inspired by renaissance painters, who did similar work by inventing illusionary techniques like mathematical perspective, *sfumato* and *tromp d'oeil* to enrich the expressiveness of their painting canvas (Figure 2-3) (Kubovy, 1988), we have taken advantage of the imaginative abilities of the human mind to optically evoke an haptic experience in a standard graphical user interface. In general, we want to call for more human computer interaction research and applications that exploit the elasticity of the human mind (Aldersey-Williams et al., 2008). Interaction designers could still become more aware that they are not entirely bound by the physical spaces of a pre-digital environment, but rather by imaginative spaces. Although knowledge and awareness of the human sensorial constitution lies on the basis of any successful interaction

design, we could use more artful interaction design, that takes advantage of the reality constructing abilities of the human mind, rather than interaction design that simply aims to recreate perceptual situations and limitations from the physical world.

New media create new perceptual spaces

Although the psychology of perception and action has a vast literature, the advent of computers has created new challenges and opportunities. With computer interfaces, humans are no longer exposed to the physical world governed by the laws of physics, but to a synthetic world whose laws can be programmed at will. For example, overlapping windows with scrollable contents do not correspond to anything in the physical world, yet most users understand them easily. Pointing with a mouse is also unnatural because the non-linear control-display ratio, i.e., the mapping between mouse movement and cursor motion – which we manipulated, using the active cursor technique, to evoke an optically simulated haptic experience that could not have existed outside the realm of digital interfaces.

Computer based interactions open a wide space of human experiences not covered by traditional psychology. This brings opportunities for both psychological researchers – that can employ digital technologies to learn more about the human perceptual system – and interaction designers – that can deliberately create new perceptual experiences, that did not exist in the pre-digital world, but nonetheless make sense from a human perspective. In order to successfully employ media technologies that tweak the human perceptual system, both groups need each others knowledge and experience.

Simulations: inferior derivatives or catalysts of change?

Simulations are commonly perceived as inferior copies of some original. However this judgment is perhaps too simple and negative. In the prologue of this study we observed that simulations can at times be more influential, satisfying and meaningful than the things they are presumed to represent. Arguably, simulations can also play a transformational role towards change. In *Language of New Media*, Lev Manovich (2001) describes how new media initially often mimic some older medium, yet as time passes the older medium is transformed. For example, flipping through my record collection on my iPhone (Figure 2-11), provides me with the familiar feeling of having a record collection and being able to browse it. This makes the idea of storing music digitally, more acceptable from a users perspective. However, at the same time, the model of storing music digitally completely revolutionizes the record industry. In the long run digital music may even cause the disappearance of physical records altogether! Likewise the first cars were designed as horseless carriages. Likewise you still click on an 'envelope' icon to open your email application, which caused an enormous decrease in the use of actual envelopes. These examples illustrate how simulations can play a catalyzing role in the transition from old to newer media.

When considered in terms of media schemas – defined in the prologue as the knowledge we possess about what media are capable of and what we should expect from them – the notion of a new medium mimicking some older medium can be seen as a strategy to softening the changes in our media schemas. Rather than the shock people experienced when, according to the anecdote, the Lumière brothers (1895) showed their film of a train arriving at the station and people ran out of the theatre, the new medium politely introduces itself as a simulation of some older, familiar medium, in order to smoothen its acceptance. No surprise the desktop metaphor, with its folders, buttons and trashcan made the computer accessible to millions¹⁹.

Similarly, our active cursor technique emerged from an attempt to simulate the functioning of force-feedback devices. Although we have argued that our technique – since it introduces the possibility of haptic communication in the familiar desktop setup – could catalyze the introduction of mechanical haptic devices, this does not necessarily mean that the optical simulation will eventually be replaced by mechanical haptics. Even though the technique emerged in an attempt to replace a mechanical haptic device, it might also bring opportunities and possibilities that would not have been possible with a mechanical haptic device.

6.3 IN CONCLUSION

Overall, we conclude that a perception of touch can be simulated optically. In particular, within a standard graphical interface, applying active cursor displacements upon the user's mouse movements can evoke such optically simulated haptic percepts.

Optically simulated haptic feedback is in many regards a good alternative for mechanically simulated haptic feedback. People are well able to recognize optically simulated bump and hole structures and, depending on the simulated strength of the force, optically simulated haptic feedback can take precedence over mechanically simulated haptic feedback. Furthermore it can be used to enhance the usability of pointing and assist the user while navigating the interface, express textures or material properties of (three-dimensional) objects, create mixed initiative interfaces – whereby both the systems and the user can initiate actions – and possibly allows for richer, more aesthetical interactions that increase the quality of the users' experience.

We do not expect optically simulated haptic feedback to be able to replace mechanical force feedback devices altogether – especially since it will always require a visual display. However, we do think optically simulated haptic feedback could catalyze the acceptance of force feedback devices. Possibly the technique can also be applied in combination with mechanical force feedback devices. The PowerCursor toolkit should make our technique of optically simulated haptic feedback available for a larger group

19 At the beginning of the digital era, several metaphors from the physical world were transferred to the digital environment in order to make, otherwise incomprehensible, technology understandable. Meanwhile the digital environment has become so accepted that concepts from the digital realm may soon be transferred to our physical environment.

of interaction designers, allowing them to further explore the applicability of the technique within more complex settings.

With this study we hope to have contributed to a richer and more physical paradigm of graphical user interfaces and to a better understanding of simulations and media technologies in general. We have tried to use simulations positively, and have argued, not merely represent, but at times also constitute our reality. Some further philosophical observations on this notion can be found in the epilogue of this study.

From New Media to Next Nature

Nearing the end of this thesis, I will briefly return to the larger social-cultural context in which this project was conducted. In the prologue I discussed the ability of the human mind to pragmatically construct ‘reality’ by combining and weighing sensory perceptions with what we expect and already know. Within my research, I have taken advantage of the human tendency to integrate the various senses and to simulate haptic percepts with merely optical means.

Continuing on the thoughts set out in the prologue, I want to share some of my philosophical observations about the functioning of media at large. But before laying out my argument and vision for the future, I think it is fair and relevant to tell you something about myself, a personal confession: I am an utterly clumsy car parker. All around me, I see cars slide into parking spaces almost effortlessly. It seems so simple, even illiterates and uneducated people can do it! I myself, however, can barely manage to legally – that is, without shocking, bumping and bumper bashing – leave my car in the parking lot. So now you know, I admit it: a terrible parker wrote this thesis. Any further reading is entirely at your own risk.

Now, returning to our reflections regarding media. If we ask people what they think of as *media*, the current top three are: Internet, telephone and television. However, for someone who – like myself – is a deplorable car parker, things are a bit more complicated. I prefer the broader definition: Media is anything that functions as an extension of the human body. Our senses and organs are being remolded by the media we use. Television as an extension of your eyes, a warm jacket as an extension of your skin, the car as an extension of your legs (McLuhan, 1964).

I am a clumsy car parker, because I am too conscious of the car as a medium. I maneuver my car with the mindset of a pedestrian, which doesn’t work. Other people do it better, fortunately, although it still sometimes goes wrong. However, when two cars bump into each other, chances are the driver of the car being hit will shout: “Hey! You’ve hit me!”, instead of “You’ve hit my car” or “Your car has hit my car!”, which would be the more accurate accusation. Hence, the car absorbs our sense of identity; we extend our senses into the vehicle (McCloud, 1994).

20 A Dutch version of this essay has been published earlier under the title *Als de producten tot je spreken*, Mensvoort, Koert van (2007) in *Nieuwe Media Cultuur in Nederland*, Lovink et al. (editors) (2007) uitgave van het Instituut voor Netwerkcultuur, Lectoraat van het Instituut voor Interactieve Media, Hogeschool van Amsterdam.

When people speak of ‘media’ they usually only speak of the media that didn’t exist yet in the lives of their parents or grandparents; primarily, we think and speak of the media which arising we are aware of²¹. These ‘new media’ are still in the front of our mind, we are not yet used to their existence and we are dealing with them consciously. Luckily, we deal quite differently with the older media! Imagine how clumsy it would be if, every time I write a message list, tied my shoes, or pressed a light switch, I would have to be aware of my media use. It is only once we start reflecting on ‘media’ as such, that we realize they are of all times.

For centuries, we have been living in sort of a ‘virtual reality’ that we sometimes also denote with the word ‘culture’. We try to control and improve our environment by way of the roofs above our heads, electric light, dishwasher, central heating, etcetera. The problem we have nowadays is not that we are living in an era of media, that is nothing new. The problem we have is that we are living in an era with so many poorly designed media. On a daily basis, millions of people are moving a mouse cursor on a computer screen. We point at icons and buttons. We click on them. We type letters on a so-called ‘keyboard’ – of which the signs are ordered as such that the chance that the hammers of the typewriter will hit each other is minimal. Future historians will probably sweep together the information age and the industrial age on one pile and define it as an era in which humankind was stooped by its own technology. Arguably, the ‘modern’ knowledge worker is really no more than the conveyor belt operator of our time. I can only hope that in a hundred years or so we can tastefully laugh about this utterly absurd situation.

Contrary to most people, who deem the amount of media to be increasing, I assert that the media are actually disappearing. That is to say, the media are slowly disappearing from the front of our minds where we are consciously aware of them. Truly sophisticated technology makes itself invisible (Weisser, 1991). You don’t recognize it as technology as such; rather, you just use it. Successful media applications are a lot like natural phenomena: self explanatory, ecologically and evolutionary adjusted to their environment. Now that I think about it; my vacuum-cleaning robot does have a lot of similarities with a animal pet. At times, the device behaves somewhat silly, but the charm with which it swifts through our home makes up for that. It knows all the stairways and the cat is its friend. Did you know the hyper allergic cat – it doesn’t make you sneeze – is already on the market? Now we only have to wait for the company that integrates the vacuum cleaner and the cat into one single product. Strange idea? Don’t be too sure. The extent to which media technologies are intervening in the constructive, material, aesthetic and social practice of our everyday lives, can hardly be underestimated. Fifteen percent of Japanese ten- to fourteen year olds – as a result of an

21 See also the prologue of this study, which discusses the concept of ‘media schemas’ as mental structures that represent what media are capable of and what we should expect from them.

education with videogames and Tamagochi – already don't know anymore that 'death' is irrevocable (Kyodo News, 2005). Admittedly Japan is the most technological society in the world, however they certainly are not the only people living in a technological culture (Schwarz & Jansma, 1989).

Ambient intelligence, nanotechnology, biotechnology, augmented reality and tissue engineering are just a few of the new research fields with which we are speeding into our future. All of these young research fields radically interfere with our sense of what is 'natural'. It is for that reason I have proposed to connect them under the label *next nature* (Mensvoort et al., 2005). While our natural environment is being replaced by a world of media and design, at the same time, our technological world is so intricate and complex that it is becoming a nature of its own. In summary: The media are disappearing, nature takes over. It is important to realize that I don't mean the sweet, beautiful nature, we all know from Disney films. On the contrary, this is real nature.

In my view, this changing concept of nature is among the most important themes of our time. Wild systems, genetic surprises, calm technology, autonomous machines and beautiful black flowers. We seem to enter a magical media garden that may surprise and astonish us, that may also knock us down or be kindly disposed to us. As far as my future research and design efforts are concerned, I am eager to revolve them further around this somewhat daunting, but also fascinating development. These are exhilarating times we are living in, unsure what the future may bring. However, I am sure of one thing: we will get the next nature we deserve. Good luck parking!

References

- Aarts, E., Marzano, S., (editors) (2003). *The New Everyday: Visions of Ambient Intelligence*, 010 Publishers.
- Agarawala, A. & Balakrishnan, R. (2006). Keepin' it Real: Pushing the Desktop Metaphor with Physics, Piles and the Pen. *Proceedings of CHI 2006 - the ACM Conference on Human Factors in Computing Systems*. p. 1283-1292. (10 pages)
- Ahlström, D., Hitz, M. and Leitner, G. (2002). Improving Mouse Navigation - A Walk Through the Hilly Screen Landscape. In *Proc. DSV-IS 2002*, Springer-Verlag, 185-195.
- Ahlström, David (2005). Modeling and improving selection in cascading pull-down menus using Fitts' law, the steering law and force fields, *Proceedings of the SIGCHI conference on Human factors in computing systems*, April 02-07, 2005, Portland, Oregon, USA
- Ahlström, D., Hitz, M., and Leitner, G. (2006). An evaluation of sticky and force enhanced targets in multi target situations. In *Proceedings of the 4th Nordic Conference on Human-Computer interaction: Changing Roles* (Oslo, Norway, October 14 - 18, 2006). A. Mørch, K. Morgan, T. Bratteteig, G. Ghosh, and D. Svanaes, Eds. NordiCHI '06, vol. 189. ACM, New York, NY, 58-67. DOI=<http://doi.acm.org/10.1145/1182475.1182482>
- Akamatsu, M. & Sato, S. (1994). A multi-modal mouse with tactile and force feedback, *International Journal of Human-Computer Studies*, 40, 443-453.
- Akamatsu, M., Sato, S. & Mackenzie, I. S. (1994). Multimodal mouse: A mouse-type device with tactile and force display, *Presence*, 3, 73-90.
- Akamatsu, M. and MacKenzie, I. S. (1996). "Movement Characteristics Using a Mouse with Tactile and Force Feedback." *International Journal of Human-Computer Studies* 45: 483-493.
- Akamatsu, M., MacKenzie, I.S., & Hasbroucq, T., (1995). A comparison of tactile, auditory, and visual feedback in a pointing task using a mouse-type device. *Ergonomics* 38, 816-827.
- Aldersey-Williams, Hugh., Hall, Peter., Sargent, Ted., Antonelli, Paola., (2008). *Design and the elastic mind*. The Museum of Modern Art, New York. ISBN-13: 978-0870707322
- Ambient Devices (2004). Ambient Orb <http://www.ambientdevices.com/cat/orb/> (retrieved august 2004)
- Attenborough, D. (1998). *Life of Birds*. DVD from the BBC Series, 2 Entertain Video, ASIN: B00004CXKJ

- Aristotle, (350 BC). *Poetics*, Translated by Stephen Halliwell. Chapel Hill: University of North Carolina Press, 1987. Pg 34.
- Baecker, R. M., Grudin, J., Buxton, W. A. S. & Greenberg, S. (1995). *Readings in human-computer interaction: Toward the year 2000* [Chap. 7: Touch, Gesture, and Marking]. San Francisco: Kaufmann.
- Balakrishnan, R., (2004). "Beating" Fitts' law: virtual enhancements for pointing facilitation, *International Journal of Human-Computer Studies*, v.61 n.6, p.857-874, December 2004
- Baudisch, P., Cutrell, E., Hinckley, K., and Eversole, A. (2005). Snap-and-go: helping users align objects without the modality of traditional snapping. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Portland, Oregon, USA, April 02 - 07, 2005). CHI '05. ACM, New York, NY, 301-310. DOI= <http://doi.acm.org/10.1145/1054972.1055014>
- Baudrillard, J. (1981). *Simulacra and Simulation*. Translated by Sheila Faria Glaser, University of Michigan Press (1995), ISBN: 0472065211
- Baudrillard, J. (1988). *Xerox and infinity* (London: Agitac, Trans.). Paris: Touchepas.
- BBC News (2006). Online video 'eroding TV viewing', 27 November 2006.
- Bederson, B., (2000). Fisheye menus. *Proceedings of the ACM UIST Symposium on User Interface Software and Technology*, pp. 217–225.
- Bevan, M., (1995). Force-feedback technology. *VR NEWS – Virtual Reality Worldwide* 4, 23-29.
- Bevan, M. (1995). Force-Feedback technology. *VR NEWS – Virtual Reality Worldwide*, 4(6):23-29, July 1995
- Bier, Eric A., Stone, Maureen C., (1986). Snap-dragging, *ACM SIGGRAPH Computer Graphics*, v.20 n.4, p.233-240, Aug. 1986
- Blanch, R., Guiard, Y., Beaudouin-Lafon, M. (2004). Semantic pointing: improving target acquisition with control-display ratio adaptation, *Proceedings of the SIGCHI conference on Human factors in computing systems*, p.519-526, April 24-29, 2004, Vienna, Austria.
- Boritz, J., Booth, K.S., & Cowan, W.B., (1991). Fitts' law studies of directional mouse movement. *Proceedings of Graphical Interface '91*, 216-223. Toronto: CIPS.
- Brave S, Dahley D. (1997). inTouch: a medium for haptic interpersonal communication. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'97)*, Atlanta, GA, 22–27 March 1997.
- Bright, R. (2000). 'Uncertain Entanglements' in Sian Ede (red.), *Strange and Charmed Science and the contemporary Visual Arts*, Calouste Gulbenkian Foundation, London 2000, p.120-143
- Brooks, Jr., F.P., Ouh-Young, M., Batter, J.J., & Kilpatrick, P.J., (1990). Project

- GROPE - Haptic Displays for Scientific Visualization. Proceedings SIGGRAPH '90, August, 177-185, Dallas, TX, Computer Graphics 24, 177-185.
- Burgoon, J.K. (1991). Relational message interpretations of touch, conversational distance, and posture. *Journal of Nonverbal Behavior*, 15(4), 233-259.
- Burgoon, J.K., Buller, D.B., Hale, J.L., & deTurk, M.A. (1984). Relational messages associated with nonverbal behaviors. *Human communication Research*, 10(30), 351-378.
- Buxton, W. (1986). There's More to Interaction than Meets the Eye: Some Issues in Manual Input. In Norman, D. A. and Draper, S. W. (Eds.), (1986), *User Centered System Design: New Perspectives on Human-Computer Interaction*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 319-337.
- Buxton, W. (2001). Less is More (More or Less), in P. Denning (Ed.), *The Invisible Future: The seamless integration of technology in everyday life*. New York: McGraw Hill, 145 – 179.
- Card, S.K., English, W.K., & Burr, B.J., (1978). Evaluation of a mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT. *Ergonomics* 21, 601-613.
- Carr, K.T. & Lederman, S.J., (1995). *Simulated and Virtual Realities: Elements of Perception*. Taylor and Francis Ltd, London, UK.
- Chang, B., & Ungar, D., (1993). "Animation: From Cartoons to User Interface", *UIST1993*, pp. 45-55. (November 1993)
- Charpentier, A., (1891). Analyse expérimentale de quelques éléments de la sensation de poids [Experimental study of some aspects of weight perception]. *Arch. Physiol. Normales Pathologiques* 3, 122-135.
- Crampton Smith, G. (1995). *The Hand That Rocks the Cradle*. I.D., May/June 1995, pp. 60-65.
- Debord, Guy (1967). *Society of the Spectacle*, Zone Books, 1995, ISBN 0942299795
- Dennerlein, J.T., Martin, D.B., Hasser, C., (2000). Force-feedback improves performance for steering and combined steering-targeting tasks. In: *Proceedings of the Conference on Human Factors in Computing Systems*, April 1–6, 2000, Hague, The Netherlands. Association for Computing Machinery, New York.
- Dennett, D. (1991). *Consciousness Explained*. Boston: Little, Brown, and Co.
- Djajadiningrat, T., Wensveen, S., Frens, J., Overbeeke, K. (2004). Tangible products: redressing the balance between appearance and action, *Personal and Ubiquitous Computing*, v.8 n.5, p.294-309, September 2004.
- Engel, F. L., Goossens, P. & Haakma, R. (1994). Improved efficiency through I- and E-feedback: A trackball with contextual force feedback. *International Journal of*

Human-Computer Studies, 41, 949-974.

- Engelbart, D. & English, W. 1968. Research Center for Augmenting Human Intellect. Reprinted in ACM SIGGRAPH Video Review 106 (1994), video made in 1968. English, W.K., Engelbart, D.C., & Berman, M.L. 1967. Display selection techniques for text manipulation. *IEEE Transactions on Human Factors in Electronics HFE-8* 1 (1967), pp. 5-15.
- English, W. K., Engelbart, D. C. & Berman, M. L (1967). Display selection techniques for text manipulation. *IEEE Transactions on Human Factors in Electronics*, HFE-8, 21-31.
- Epps, B.W., 1986. Comparison of six cursor control devices based on Fitts' law models. Proceedings of the 30th Annual Meeting of the Human Factors Society, September 29 - October 3, 1986, Dayton, Ohio, pp. 327-331.
- Ernst, M. O. and Banks, M. S. 2002. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature (London)* 415, 429--433.
- Fitts, P.M., 1954. The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology* 47, 381-391.
- Fitts, P.M., & Peterson, J.R., 1964. Information capacity of discrete motor responses. *Journal of Experimental Psychology* 67, 103-112.
- Flourney, T., 1894. De l'influence de la perception visuelle des corps sur leur poids [The influence of visual perception on the apparent weight of objects]. *L'Année Psychologique* 1, 1989-208.
- Fogg, B.J. 1998. Persuasive computers: perspectives and research directions, Proceedings of the SIGCHI conference on Human factors in computing systems, p.225-232, April 18-23, 1998, Los Angeles, California, United States.
- Friedlander, N., Schlueter, K., & Mantei, M. (1998). Bullseye! When Fitts' Law doesn't Fit. Proceedings of the CHI '98 Conference on Human Factors in Computing Systems, ACM Press, New York, NY, pp. 257-264.
- Furnas, G., 1986. Generalized fisheye views. Proceedings of the ACM CHI Conference on Human Factors in Computing Systems, pp. 16-23.
- George G. Robertson , Jock D. Mackinlay , Stuart K. Card, Cone Trees: animated 3D visualizations of hierarchical information, Proceedings of the SIGCHI conference on Human factors in computing systems: Reaching through technology, p.189-194, April 27-May 02, 1991, New Orleans, Louisiana, United States.
- Goodman M., Tagle D.A., Fitch D.H., Bailey W., Czelusniak J., Koop B.F., Benson P., Slightom J.L. (1990) Primate evolution at the DNA level and a classification of hominoids. *Journal of molecular evolution* 1990;30(3):260-6.
- Greenfield, Adam (2006). Everywhere: the dawning age of ubiquitous computing. New Riders, p.11-12. ISBN 0321384016.

- Gregory, R L. (1998) Eye and brain: The psychology of seeing (5th Edition). Oxford University Press.
- Gibson, J.J., 1966. The senses considered as perceptual systems. Boston, MA: Houghton Mifflin.
- Harrison, B.L., Fishkin, K.P., Gujar, A., Mochon, C., Want, R., (1998) Squeeze me, hold me, tilt me! An exploration of manipulative user interfaces, Proceedings of the SIGCHI conference on Human factors in computing systems, p.17-24, April 18-23, 1998, Los Angeles, California, United States
- Hedges, S. B. (2000) Human evolution: A start for population genomics, *Nature* 408, 652-653 (7 December 2000)
- Heeks, R. (2008). Current Analysis and Future Research Agenda on “Gold Farming”: Real-World Production in Developing Countries for the Virtual Economies of Online Games. Working Paper Series, Development Informatics Group, University of Manchester.
- Heidegger, M. 1927. Sein und Zeit. Tübingen, Germany: Max Niemeyer Verslag.
- Heller, M. A., Calcaterra, J. A., Green, S. L., and Brown, L. 1999. Intersensory conflict between vision and touch: The response modality dominates when precise, attention-riveting judgments are required. *Perception Psychophys.* 61, 1384--1398.
- Hoffman, D.D. (1998). Visual Intelligence: How We Create What We See. New York: W.W. Norton and Co.
- Horvitz, Eric, 1999. Principles of mixed-initiative user interfaces, Proceedings of the SIGCHI conference on Human factors in computing systems: the CHI is the limit, May 15-20, 1999, Pittsburgh, PN, pp.159-166.
- Ishii, H. Ullmer, B. (1997) Tangible bits: towards seamless interfaces between people, bits and atoms, Proceedings of the SIGCHI conference on Human factors in computing systems, p.234-241, March 22-27, 1997, Atlanta, Georgia, United States.
- ISO 9241, 1998. Ergonomic requirements for office work with visual display terminals (VDTs) - Part 11: Guidance on usability.
- Jagacinski, R.J., & Monk, D.L., 1985. Fitts' law in two dimensions with hand and head movement. *Journal of Motor Behavior* 17, 77-95.
- Kabbash, P. & Buxton, W., 1995. The “Prince” technique: Fitts' law and selection using area cursors. *Proceedings of the CHI '95 Conference on Human Factors in Computing Systems* (pp. 273-279) New York: ACM
- Kerstner, W. Pigel, G. and Tscheligi, M. (1994). The FeelMouse: Making Computer Screens Feelable. In W. Zagler et al., editors, *Computers for handicapped Persons*. Proc. Of the ICCHP'94. Springer-Verlag, 1994

- Keyson, D.V. (1996). Touch in User Interface Navigation, thesis Technische Universiteit Eindhoven, 1996 ISBN 90-368-0288-X
- Keyson, D.V. (1991). Dynamic control gain and tactile feedback in the capture of cursor movements, *Ergonomics*, (12), Taylor & Francis Ltd. (1991), 1287--1298.
- Klatzky, R. L., Lederman, S., and Reed, C. (1987). There's more to touch than meets the eye: The salience of object attributes for haptics with and without vision. *J. Exp. Psychol. General*, 116, 356--369.
- Kyodo News. (2005). Informal survey on awareness of the concept of 'death' among Japanese teenagers.
- Kubovy M. (1988). *Psychology of perspective and Renaissance art*. Cambridge University Press, Cambridge.
- Lakoff, G. & Johnson, M. (1980). *Metaphors we live by*, The University of Chicago Press, Chicago and London, 1980.
- Laurel, B. (1993) *Computers as Theatre*, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, 1993
- Lécuyer, A., Coquillart, S., Kheddar, A., Richard, P., and Coiffet, P. (2000). Pseudo-Haptic Feedback: Can Isometric Input Devices Simulate Force Feedback?. In *Proceedings of the IEEE Virtual Reality 2000 Conference* (March 18 - 22, 2000). VR. IEEE Computer Society, Washington, DC, 83.
- Lécuyer, A., Burkhardt, J.-M., & Etienne, L., (2004). Feeling bumps and holes without a haptic interface: The perception of pseudo-haptic textures. *Proceedings of the CHI 2004 Conference*, Vienna, Austria, April 24-29 2004, pp. 239-246.
- Lederman, S. J., Thorne, G., and Jones, B. (1986). Perception of texture by vision and touch: Multidimensionality and intersensory integration. *J. Exp. Psychol. Hum. Perception Performance* 12, 169--180.
- Ljungstrand, P. Redström J., Holmquist, L.E. (2000). WebStickers: using physical tokens to access, manage and share bookmarks to the Web, *Proceedings of DARE 2000 on Designing augmented reality environments*, p.23-31, April 2000, Elsinore, Denmark
- Lombard, M., Ditton, T., (1997). At the heart of it all: the concept of presence. *Journal of Computer-Mediated Communication* 3 (2), online Journal.
- Lu ZL, Sperling G (1995) The functional architecture of human visual motion perception. *Vision Res* 35:2697--2722.[
- Lumière, Auguste & Lumière, Louis. (1895). *L'Arrivée d'un train à la Ciotat*
- MacKenzie, I.S., (1989). A note on the information theoretic basis for Fitts' law. *Journal of Motor Behavior* 21, 323-330.
- MacKenzie, I.S., (1992). Fitts' law as a research and design tool for human-computer

- interaction. *Human-Computer Interaction* 7, 91-139.
- MacKenzie I.S., & Buxton, W., (1992). Extending Fitts' law to two-dimensional tasks. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, May 3-7, 1992, Monterey, California, pp. 219-226.
- Manovich, L. (2001). *The language of new media*, The MIT Press, ISBN-13: 978-0262632553
- Manovich, L. (2006). Interaction as aesthetic event, *Receiver magazine*, Vodafone Group. www.receiver.vodafone.com
- Manovich, L. (2007). Mobile phone as gesamtkunstwerk, in Gerritzen et Al. *Style First*, Birkhauser Basel, ISBN-13: 978-3764384388, pp 45-47.
- Marks, L.E., (1978). *The unity of the senses*. New York, NY: Academic Press.
- Massaro, D. W. (1990). An information-processing analysis of perception and action. In O. Neumann & W. Prinz (Eds.), *Relationships between perception and action: Current approaches* (pp. 133-166). Berlin: Springer.
- Massaro, J. (2007). 65 Percent of Americans Spend More Time with Their Computer than Their Spouse, SupportSoft. http://supportsoft.mediaroom.com/index.php?s=press_releases&item=408 (retrieved July 2008)
- Massie, T. and Salisbury, K. (1994). The PHANTOM Haptic Interface: A Device for Probing Virtual Objects. In *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environments and teleoperator systems* (Chicago, IL), 1994.
- Mather, G. (2006). *Foundations of Perception*. Psychology Press, ISBN 978-0863778346.
- Maurer, L. Wouters, R. (2002). Klima Kontrolle, Video presented at International Browserday, Paradiso, 2002. <http://www.xelor.nl/xelor/media/html/klima.html> (retrieved July 2008)
- Maurer, L. Elastic Reality (2001). in Garcia, D. *The end of the browser*. *Interactions* 8, 1 (Jan. 2001), 42-52. p48
- McCloud, S. (1994). *Understanding Comics: The Invisible Art*, Harper Paperbacks; Reprint edition, 1994, ISBN 006097625X.
- McGuffin, M., (2002). Fitts' law and expanding targets: an experimental study and applications to user interface design. M.Sc. Thesis, Department of Computer Science, University of Toronto.
- McGuffin, M. & Balakrishnan, R., (2002). Acquisition of expanding targets, *Proceedings of the SIGCHI conference on Human factors in computing systems: Changing our world, changing ourselves*, April 20-25, 2002, Minneapolis, Minnesota, USA
- McLuhan, Marshall, (1964). *Understanding Media: The Extensions of Man*, Gingko

Press, ISBN 1-58423-073-8.

Mensvoort van, K. & Duyvenbode van, M. (2001). Het Bos ruikt naar Shampoo, Documentary. VPRO Television, April 2001.

Mensvoort van, K. (2001). Blow your Browser. in Garcia, D (editor). The end of the browser. Interactions 8, 1 (Jan. 2001), 42-52. p49

Mensvoort van, K., (2002). What you see is what you feel. Exploiting the dominance of the visual over the haptic domain to simulate force-feedback with cursor displacement. Proceedings of the Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques, June 25-28, 2002, London, England. New York: ACM Press, pp. 345-348

Mensvoort van, K. (2003). Datafountain, online currency rates translated to water. <http://www.koert.com/work/datafountain/> (retrieved jan. 2008).

Mensvoort van, K., (2003). Windowpoll. Design for cursor attracting buttons. <http://www.koert.com/work/windowpoll/>.

Mensvoort van, K. (2007). The Picture Bubble, in Gerritzen et Al. Style First, Birkhauser Basel, ISBN-13: 978-3764384388, pp 48-52.

Mensvoort van, K. (2007). 'Als de producten tot je spreken' in Nieuwe Media Cultuur in Nederland, Lovink et Al. (editors) (2007) uitgave van het Instituut voor Netwerkcultuur, Lectoraat van het Instituut voor Interactieve Media, Hogeschool van Amsterdam.

Mensvoort van, K., Hermes, D.J., Montfort, M. (2008). Usability of optically simulated haptic feedback. International Journal of Human-Computer Studies, Volume 66, Issue 6, June 2008, Pages 438-451.

Mensvoort van, K. & Grievink, H.J. (2008) Fake for Real: A Memory game on reality, AllMedia / Bis Publishers, ISBN 978-9063691776

Mensvoort van, K. M., & Oirschot, H. K., (2001). Cursor trajectory prediction using a genetic algorithm. <http://www.koert.com/work/cursorpredictor/>

Mensvoort van, K., Gerritzen, M., Schwarz, M., (Eds.) (2005), Next Nature, BIS Publishers, ISBN 90-636-9093-2, pp. 4-43.

Mensvoort van, K., Hermes, D.J., Vos, P., Liere van, R. (2009). Perception of optically simulated haptic feedback. Transactions on Applied Perception (accepted).

Miner, N., Gillespie, B., & Caudell, T., (1996). Examining the influence of audio and visual stimuli on a haptic display. Proceedings of the 1996 IMAGE Conference, Phoenix, AZ, June 23-25, 1996.

Mirapaul, M., (2001). Arts Online: Innovative Webmasters Chase Fame at Browserday, New York Times, April 2, 2001.

Mirzoeff, N., (1999). An introduction to visual culture, Routledge, London, ISBN: 0415158761.

- Muller, M. J., (1988). Multifunctional cursor for direct manipulation user interfaces, Proceedings of the SIGCHI conference on Human factors in Computing Systems, May 15-19, 1988, Washington, D.C., pp. 89-94.
- Murata, A. (1995). Discussion on method for predicting targets in pointing by mouse. Paper presented at the Proceedings of the Sixth International Conference on Human-Computer Interaction, Tokyo, Japan.
- Murata, A. (1998). Improvement of pointing time by predicting targets in pointing with a PC mouse. *International Journal of Human-Computer Interaction*, 10, 23-32.
- Murray, D.J., Ellis, R.R., Bandomir, C.A. & Ross, H.E., (1999). Charpentier (1891) on the size-weight illusion. *Percept. Psychophys*, 61, 1681-1685.
- Myers, B (1998). A Brief History of Human Computer interaction Technology. *ACM Interactions*, 5, 2, 1998, pp. 44-54
- Nevejan, C. (2007) Presence and the Design of Trust. PhD. Dissertation, University of Amsterdam, 2007.
- Norman, D. (1999) *The Invisible Computer: Why Good Products Can Fail, the Personal Computer Is So Complex, and Information Appliances Are the Solution*. MIT Press. ISBN-13: 978-0262640411
- Oakley, I., McGee, M.R., Brewster, S. & Gray, P. (2000) Putting the Feel in 'Look and Feel', In *Proceedings of ACM CHI 2000*, The Hague, ACM Press, Addison-Wesley, pp.415-422.
- Oirschot, H. K.-v., & Houtsma, A. J. M. (2000). Cursor trajectory analysis. *Lecture Notes in Computer Science*, 2058, 127-134.
- Oosterling, H. (2000). *Radicale Middelmattigheid ('Radical Mediocrity')*, uitgeverij Boom, Amsterdam. P23. ISBN:9053526218
- Oosterling, H. (2003). Act your actuality, in Gerritzen, M. et al. 2004. *Visual Power: News*. Gingko Press, Corte Madera, CA, 2004, ISBN: 9063690568.
- Piaget, J. (1997). *Jean Piaget: Selected Works*. Routledge, 1997, ISBN 9780415168892
- Park, J., Han, S.H., Yang, H., (2006). Evaluation of cursor capturing functions in a target positioning task, *International Journal of Industrial Ergonomics* 36 (2006) pp721-730
- Pine, J. and Gilmore, J. (1999). *The Experience Economy*, Harvard Business School Press, Boston, 1999.
- Poppe, Inc. (1999). Finger on the Button: Bas Ording's interface design. In: *If/Then*, Vol. 1, no. 1, ISBN: 9072007522 pp. 258-259.
- Prinz, W. (1984). Modes of linkage between perception and action. In W. Prinz & A.-F. Sanders (Eds.), *Cognition and motor processes* (pp. 185-193). Berlin:

Springer.

- Prinz, W. (2005). Experimental approaches to action. In J. Roessler and N. Eilan (Eds.), *Agency and self-awareness* (pp. 165-187). New-York, Oxford University press.
- Ramstein, C. (1995). A Multimodal User Interface System with Force Feedback and Physical Models. *IFIP Interact'95*, (Lillehammer, Norway) Chapman & Hall, 157-162.
- Rizzolatti, G., & Craighero, L. (2004). The mirror-neuron system. *Annual Review in Neuroscience*, 27, 169-92.
- Robles-De-La-Torre, G., Hayward, V. (2001). Force can overcome object geometry in the perception of shape through active touch. *Nature*, Volume 412, Issue 6845, pp. 445-448.
- Rock, I. and Victor, J. (1964). Vision and touch: An experimentally created conflict between the senses. *Science* 143, 594--596.
- Rosenberg, L., & Brave, S. (1996). Using Force feedback to enhance human performance in graphical user interfaces. *Companion Proceedings of CHI'96: Conference on Human Factors in Computing Systems*, 291-292.
- Rosenberg, L.B. (1997). FEELit mouse: Adding a realistic sense of FEEL to the computing experience. <http://www.force-feedback.com/feelit> (retrieved July 2008).
- Rovers, A.F. & Essen van, H.A. (2004). HIM: a framework for haptic instant messaging, CHI '04 extended abstracts on Human factors in computing systems, April 24-29, 2004, Vienna, Austria.
- Rowell, L. (2007). *Scratching the Surface*. NetWorker, ACM Press, Volume 10 Issue 3.
- Runeson S, & Frykholm G. (1981). Visual perception of lifted weight. *Journal of Experimental Psychology: Human Perception and Performance*, 7, 733-740.
- Searl, John (1995). *The Construction of Social Reality*, Free Press (1997), ISBN: 0684831791.
- Shneiderman, B. (1981). Direct manipulation: A step beyond programming languages. In *Proceedings of the Joint Conference on Easier and More Productive Use of Computer Systems. (Part - II): Human interface and the User interface - Volume 1981* (Ann Arbor, MI, May 20 - 22, 1981). L. Borman, Ed., New York, NY, ACM Press, p. 143.
- Shneiderman, B. (1983). Direct manipulation: a step beyond programming languages, *IEEE Computer* 16, 8 (August 1983), 57-69.
- Stein, E., & Meredith, A., (1993). *The merging of the senses*, London: MIT Press.
- Suzuki, Yu, (1986). *Outrun*, classic arcade racing game, Sega 1986.

- Schwarz, M., Jansma, R., eds. (1989). *The Technological Culture*. Amsterdam: De Balie, 1989.
- Thomas, B.H, Calder, P. (1995). Animating direct manipulation interfaces, *Proceedings of the 8th annual ACM symposium on User interface and software technology*, p.3-12, November 15-17, 1995, Pittsburgh, Pennsylvania, United States
- Thomas, B.H., Calder, P. (2001). Applying cartoon animation techniques to graphical user interfaces. *TOCHI* 8(3):198–222, September 2001.
- Tognazzini, B. (1992). *TOG on Interface*. Addison-Wesley Professional. ISBN-13: 978-0201608427
- Van den Hoven, E., Eggen, B. (2004). *Tangible Computing in Everyday Life: Extending Current Frameworks for Tangible User Interfaces with Personal Objects*. *EUSAI 2004*: 230-242
- Weiser, M. (1991). *The computer of the 21st century*. *Scientific American*, pages 94–100, September 1991.
- Weiser, M. (1994). The world is not a desktop, *interactions*, v.1 n.1, p.7-8, Jan. 1994 [doi>10.1145/174800.174801]
- Weiser, M., Brown, J. S. (1996). *Designing calm technology*. *PowerGrid Journal*, Vol. 1, No. 1, 1996.
- Welch, R. B., & Warren, D. H., (1986). Intersensory Interactions. In *Handbook of Perception and Human Performance*, edited by K. R. Boff, L. Kaufman, and J. P. Thomas, chap. 25. New York: Wiley.
- Welford, A.T., (1960). The measurement of sensory-motor performance: Survey and reappraisal of twelve years progress. *Ergonomics* 3, 189-230.
- Wickens, C., Sandry, D., Vidulich, M., (1983). Compatibility and resource competition between modalities of input, central processing, and output. *Human Factors* 25(2), 227-248.
- Wickens, C.D., & Carswell, C.M., (1995). The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors* 37, 473-494.
- Winkel van, C. (2006) *Het primaat van de zichtbaarheid*, NAI Uitgevers, 2006, ISBN 90-5662-424-5
- Worden, A., Walker, N., Bharat, K., & Hudson, S., (1997). Making computers easier for older adults to use: Area cursors and sticky icons, *Proceedings of the SIGCHI conference on Human factors in computing systems*, March 22-27, 1997, Atlanta, Georgia, United States, p.266-271.
- Ijsselstein, W. A. (2002). Elements of a multi-level theory of presence: Phenomenology, mental processing and neural correlates. In *Proceedings of PRESENCE 2002* (pp. 245-259). Porto, Portugal.

- Zhai, S., Morimoto, C., & Ihde, S., (1999). Manual and gaze input cascaded (MAGIC) pointing. In Proceedings of the SIGCHI Conference on Human factors in Computing Systems: the CHI is the limit, May 15-20, 1999, Pittsburgh, PN, pp. 246-253.
- Zhai, S., & MacKenzie, I.S., (1998). Teaching old mice new tricks: Innovations in computer mouse design. In Proceedings of Ergon-Axia '98 - the First World Congress on Ergonomics for Global Quality and Productivity, pp. 80-83.

Summary

This study introduces a novel method of simulating touch with merely visual means. Interactive animations are used to create an optical illusion that evokes haptic percepts like stickiness, stiffness and mass, within a standard graphical user interface. The technique, called *optically simulated haptic feedback*, exploits the domination of the visual over the haptic modality and the general human tendency to integrate between the various senses.

The study began with an aspiration to increase the sensorial qualities of the graphical user interface. With the introduction of the graphical user interface – and in particular the desktop metaphor – computers have become accessible for almost anyone; all over the world, people from various cultures use the same icons, folders, buttons and trashcans. However, from a sensorial point of view this computing paradigm is still extremely limited.

Touch can play a powerful role in communication. It can offer an immediacy and intimacy unparalleled by words or images. Although few doubt this intrinsic value of touch perception in everyday life, examples in modern technology where human-machine communication utilizes the tactile and kinesthetic senses as additional channels of information flow are scarce. Hence, it has often been suggested that improvements in the sensorial qualities of computers could lead to more natural interfaces.

Various researchers have been creating scenarios and technologies that should enrich the sensorial qualities of our digital environment. Some have developed mechanical force feedback devices that enable people to experience haptics while interacting with a digital display. Others have suggested that the computer should ‘disappear’ into the environment and proposed tangible objects as a means to connect between the digital and the physical environment.

While the scenarios of force feedback, tangible interactions and the disappearing computer are maturing, millions of people are still working with a desktop computer interface every day. In spite of its obvious drawbacks, the desktop computing model penetrated deeply into our society and cannot be expected to disappear overnight. Radically different computing paradigms will require the development of radically different hardware. This takes time and it is yet unsure when, if so, other computing paradigms will replace the current desktop computing setup.

It is for that reason, that we pursued another approach towards physical computing. Inspired by renaissance painters, who already centuries ago invented illusionary techniques like perspective and *trompe d’oeil* to increase the presence of their paintings, we aim to improve the physicality of the graphical user interface, without resorting to special hardware.

Optically simulated haptic feedback, described in this thesis, has a lot in common with mechanical force-feedback systems, except for the fact that in mechanical force-feedback systems the location of the cursor is manipulated as a result of the force sent to the haptic device (force-feedback mouse, trackball, etc), whereas in our system the cursor location is directly manipulated, resulting in an purely visual force feedback. By applying tiny displacements upon the cursor's movement, tactile sensations like stickiness, touch, or mass can be simulated. In chapter 2 we suggest that the active cursor technique can be applied to create richer interactions without the need for special hardware. The cursor channel is transformed from an input only to an input/output channel. The active cursor displacements can be used to create various (dynamic) slopes as well as textures and material properties, which can provide the user with feedback while navigating the on-screen environment.

In chapter 3 the perceptual illusion of touch, resulting from the domination of the visual over the haptic modality, is described in a larger context of prior research and experimentally tested. Using both the active cursor technique and a mechanical force feedback device, we generated bumps and hole structures. In a controlled experiment the perception of the slopes was measured, comparing between the optical and the mechanical simulation. Results show that people can recognize optically simulated bump and hole structures, and that active cursor displacements influence the haptic perception of bumps and holes. Depending on the simulated strength of the force, optically simulated haptic feedback can take precedence over mechanically simulated haptic feedback, but also the other way around. When optically simulated and mechanically simulated haptic feedback counteract each other, however, the weight attributed to each source of haptic information differs between users. It is concluded that active cursor displacements can be used to optically simulate the operation of mechanical force feedback devices.

An obvious application of optically simulated haptic feedback in graphical user interfaces, is to assist the user in pointing at icons and objects on the screen. Given the pervasiveness of pointing in graphical interfaces, every small improvement in a target-acquisition task, represents a substantial improvement in usability. Can active cursor displacements be applied to help the user reach its goal? In chapter 4 we test the usability of optically simulated haptic feedback in a pointing task, again in comparison with the force feedback generated by a mechanical device. In a controlled Fitts'-law type experiment, subjects were asked to point and click at targets of different sizes and distances. Results learn that rendering hole type structures underneath the targets improves the effectiveness, efficiency and satisfaction of the target acquisition task. Optically simulated haptic feedback results in lower error rates, more satisfaction, and a higher index of performance, which can be attributed to the shorter movement times realized for the smaller targets. For larger targets, optically simulated haptic feedback resulted in comparable movement times as mechanically simulated haptic feedback.

Since the current graphical interfaces are not designed with tactility in mind, the development of novel interaction styles should also be an important research path. Before optically simulated haptic feedback can be fully brought into play in more complex interaction styles, designers and researchers need to further experiment with the technique. In chapter 5 we describe a software prototyping toolkit, called PowerCursor, which enables designers to create interaction styles using optically simulated haptic feedback, without having to do elaborate programming. The software engine consists of a set of ready force field objects – holes, hills, ramps, rough and slick objects, walls, whirls, and more – that can be added to any Flash project, as well as force behaviours that can be added to custom made shapes and objects. These basic building blocks can be combined to create more complex and dynamic force objects. This setup should allow the users of the toolkit to creatively design their own interaction styles with optically simulated haptic feedback. The toolkit is implemented in Adobe Flash and can be downloaded at www.powercursor.com.

Furthermore, in chapter 5 we present a preliminary framework of the expected applicability of optically simulated haptic feedback. Illustrated with examples that have been created with the beta-version of the PowerCursor toolkit so far, we discuss some of the ideas for novel interaction styles. Besides being useful in assisting the user while navigating, optically simulated haptic feedback might be applied to create so-called mixed initiative interfaces – one can for instance think of an installation wizard, which guides the cursor towards the recommended next step. Furthermore since optically simulated haptic feedback can be used to communicate material properties of textures or 3D objects, it can be applied to create aesthetically pleasing interactions – which with the migration of computers into other domains than the office environment are becoming more relevant. Finally we discuss the opportunities for applications outside the desktop computer model. We discuss how, in principle, optically simulated haptic feedback can play a role in any graphical interface where the input and output channels are decoupled.

In chapter 6 we draw conclusions and discuss future directions. We conclude that optically simulated haptic feedback can increase the physicality and quality of our current graphical user interfaces, without resorting to specialistic hardware. Users are able to recognize haptic structures simulated by applying active cursor displacements upon the users mouse movements. Our technique of simulating haptic feedback optically opens up an additional communication channel with the user that can enhance the usability of the graphical interface.

However, the active cursor technique is not to be expected to replace mechanical haptic feedback altogether, since it can be applied only in combination with a visual display and thus will not work for visually impaired people. Rather, we expect the ability to employ tactile interaction styles in a standard graphical user interface, could catalyze the development of novel physical interaction styles and on the long term might instigate the acceptance of haptic devices.

With this research we hope to have contributed to a more sensorial and richer graphical user interface. Moreover we have aimed to increase our awareness and understanding of media technology and simulations in general. Therefore, our scientific research results are deliberately presented within a social-cultural context that reflects upon the dominance of the visual modality in our society and the ever-increasing role of media and simulations in people's everyday lives.

Samenvatting

Dit proefschrift introduceert een nieuwe techniek om tast (haptiek) te simuleren met enkel visuele middelen. Via interactieve animaties wordt een optische illusie gecreëerd welke haptische percepties – zoals plakkerigheid, stijfheid en massa – oproept in een standaard grafische gebruikersomgeving. De techniek, genaamd *optisch gesimuleerde haptische feedback*, maakt gebruik van de dominantie van de visuele over de haptische modaliteit en de menselijke neiging om percepties in verschillende zintuiglijke modaliteiten te integreren tot één consistente ervaring.

Het onderzoek begon met een aspiratie de zintuiglijke kwaliteiten van grafische interfaces te verbeteren. Sinds de introductie van de grafische gebruikersomgeving – in het bijzonder de bureaublad metafoor – zijn computers voor vrijwel iedereen toegankelijk geworden; overal ter wereld gebruiken mensen uit allerlei culturen dezelfde iconen, mappen, knoppen en prullenbak. Zintuiglijk is dit computer paradigma echter nog extreem beperkt: de objecten op het bureaublad zijn plat en gewichtloos.

Tast kan een belangrijke rol spelen in communicatie. Tast kan een directheid en intimiteit overbrengen welke onvergelijkbaar is met beeld of geluid. In onze alledaagse fysieke ervaringswereld is het kinetische gedrag van objecten vanzelfsprekend. Het geeft informatie over de lichaamseigenschappen van een object. Als je een deur opent, zul je een zekere weerstand voelen die je iets over de deur zegt, hoe deze geplaatst is en waar hij van gemaakt is. Als je een doos optilt, voel je of de doos vol of leeg is. Hoewel de intrinsieke waarde van haptische perceptie in het alledaagse leven onbetwist is, zijn de toepassingen waarbij tast wordt ingezet als communicatiemiddel in de interactie tussen mensen en computers relatief schaars. Er is dan ook al vaak gesuggereerd dat verbeteringen op het haptische vlak kunnen leiden tot een meer natuurlijke omgang met digitale technologie.

Diverse onderzoekers hebben inmiddels scenario's en technologieën gecreëerd die de zintuiglijkheid van onze digitale omgeving zouden moeten verbeteren. Sommige onderzoekers hebben mechanische force feedback apparaten ontwikkeld welke mensen in staat stellen haptische terugkoppeling te ervaren in de interactie met de digitale omgeving. Anderen hebben gesuggereerd dat de computer volledig moet 'verdwijnen' in de omgeving en dat de interactie tussen de digitale en de fysieke omgeving moet verlopen via fysieke, tastbare (tangible) objecten.

Terwijl deze scenario's van force feedback, tastbare objecten en de verdwijnende computer langzaam aan volwassen worden, werken iedere dag nog altijd miljoenen mensen met de desktop computer interface. Ondanks de overduidelijke minpunten is het desktop computer paradigma diep in onze maatschappij gepenetreerd; het ligt niet in de verwachting dat het op stel en sprong zal verdwijnen. Radicaal andere interface

paradigma's vereisen de ontwikkeling van radicaal andere hardware. Dit kost tijd en het is nog onzeker of en wanneer alternatieve interactie paradigma's het huidige desktop model zullen vervangen.

Het is om die reden dat wij een andere benadering met betrekking tot de behoefte aan zintuiglijk rijkere, meer tastbare interacties hebben gevolgd. Geïnspireerd door Renaissance schilders, welke eeuwen geleden optische illusie-technieken – zoals mathematisch perspectief, sfumato en tromp d'oeil – ontwikkelden om de 'presence' van hun schilderijen te verbeteren, hebben we getracht de haptische kwaliteit van de grafische gebruikersomgeving te verhogen, zonder gebruik te maken van speciale en slechts mondjesmaat beschikbare hardware.

Optisch gesimuleerde haptische feedback, zoals omschreven in dit proefschrift, heeft veel overeenkomsten met mechanische force feedback apparaten. Het belangrijke verschil is dat, terwijl bij mechanische force feedback systemen, de positie van de cursor wordt gemanipuleerd als gevolg van de kracht die naar het haptisch apparaat (force-feedback muis, trackball, etc.) gezonden wordt, terwijl in ons systeem de cursor direct gemanipuleerd wordt; een puur visuele vorm van feedback. Tactiele ervaringen, zoals plakkerigheid, slipperigheid of massa kunnen worden gesimuleerd door minieme manipulaties van de door de gebruiker gemaakte cursor bewegingen.

In hoofdstuk 2 suggereren we dat onze *active cursor* techniek kan worden toegepast om rijkere, meer tactiele interactiestijlen te creëren zonder dat daarvoor speciale hardware nodig is. Het cursor kanaal wordt getransformeerd van een enkelvoudig input tot een input/output kanaal. De actieve cursor verplaatsingen kunnen worden toegepast om verschillende (dynamische) hellingen alsmede textuur en materiaaleigenschappen te genereren, welke de gebruiker contextuele feedback geven, terwijl deze door de schermomgeving navigeert.

In hoofdstuk 3 wordt de optische illusie van tast, welke een gevolg is van de dominantie van de visuele over de haptische modaliteit, beschreven in de context van eerder onderzoek en verder experimenteel onderzocht. Gebruikmakend van zowel de active cursor techniek als van een mechanisch force feedback apparaat, werden heuvels en kuilen van diverse hoogte en diepte gegenereerd. De perceptie van de glooiingen werd gemeten, waarbij de optische en de mechanische simulatie vergeleken werden. Uit de resultaten blijkt dat de proefpersonen zowel de optisch alsmede de mechanisch gesimuleerde heuvels en kuilen kunnen herkennen. Afhankelijk van de gesimuleerde kracht kan de optisch gesimuleerde haptische feedback de mechanisch gesimuleerde haptische feedback domineren, maar ook andersom. Wanneer optisch en mechanisch gesimuleerde haptische feedback met elkaar in conflict zijn, verschilt het gewicht dat aan de verschillende bronnen van haptische informatie wordt toegekend van gebruiker tot gebruiker. Uit het experiment concluderen we dat de werking van mechanische force feedback apparaten via actieve cursor verplaatsingen kan worden gesimuleerd.

In hoofdstuk 4 onderzoeken we de bruikbaarheid van optisch gesimuleerde haptische feedback in een aanwijstaak, wederom in vergelijking met mechanisch gesimuleerde haptische feedback, gegenereerd met behulp van een mechanische force feedback muis. Een voor de hand liggende toepassing van optisch gesimuleerde haptische feedback in grafische gebruikersomgevingen is het assisteren van de gebruiker bij het aanwijzen en selecteren van iconen en objecten op het scherm. Gezien de frequentie waarmee de aanwijstaak in de grafische gebruikersomgeving door gebruikers wordt uitgevoerd, zal iedere kleine verbetering van de aanwijstaak een substantiële verbetering in bruikbaarheid teweegbrengen.

In een gecontroleerd op Fitts'-law gebaseerd experiment, werd aan de proefpersonen gevraagd om verschillende doelen, met wisselende groottes en op wisselende afstanden, aan te klikken. De resultaten leren ons dat het aanbrengen van een kuil structuur onder de doelen de effectiviteit, efficiency en bevrediging van de aanklik taak in grafische gebruikers interfaces verbetert. Optisch gesimuleerde haptische terugkoppeling resulteert in minder fouten, meer bevrediging en een hogere prestatie index, vooral voor de kleinere doelen. Naarmate de doelen groter worden is de bewegingstijd van mechanische en optisch gesimuleerde haptische feedback vergelijkbaar.

In hoofdstuk 5 beschrijven we een software toolkit, PowerCursor genaamd, welke ontwerpers in staat stelt interactie stijlen met optisch gesimuleerde haptische feedback te ontwikkelen, zonder zelf uitgebreid te hoeven programmeren. Omdat de huidige grafische gebruikersomgevingen niet zijn ontworpen vanuit de mogelijkheid met haptische feedback te werken, is de ontwikkeling van nieuwe interactie stijlen een belangrijke onderzoekslijn. Voordat optisch gesimuleerde haptische feedback volledig kan worden uitgebuit in complexere interactieve dialogen, moeten ontwerpers en onderzoekers verder experimenteren met de techniek. De PowerCursor software bestaat uit een verzameling van krachtveldobjecten – kuilen, heuvels, hellingen, ruwe en gladde objecten, muren, draaikolken, etc – welke aan een Flash project kunnen worden toegevoegd, alsmede krachtgedragingen die aan zelf gemaakte vormen en objecten kunnen worden gekoppeld. Deze primaire bouwstenen kunnen worden gecombineerd tot complexe en dynamische tactiele objecten. Deze opzet stelt de gebruikers van de toolkit in staat creatieve zelf bedachte interactiestijlen te ontwikkelen met optisch gesimuleerde haptische feedback. De toolkit is geïmplementeerd in Adobe Flash en kan worden gedownload via www.powercursor.com.

Tevens verkennen we in hoofdstuk 5 de toepasbaarheid van optisch gesimuleerde haptische feedback. Ideeën voor nieuwe interactie stijlen worden bediscussieerd en geïllustreerd met voorbeelden, welke in een beta-versie van de PowerCursor toolkit zijn ontworpen. Naast de eerder genoemde toepassing in het ondersteunen van de bewegingen van de gebruiker, kan optisch gesimuleerde haptische feedback gebruikt worden om zogenaamde 'mixed-initiative interfaces' te realiseren – denk bijvoorbeeld aan een installatie wizzard, waarbij het systeem de cursor steeds beweegt in de richting van de geadviseerde volgende stap. Daarnaast kan optisch gesimuleerde

haptische feedback gebruikt worden om materiaaleigenschappen van texturen of driedimensionale objecten te communiceren. Ook kan de techniek gebruikt worden om esthetische meer plezierige interacties te realiseren – welke met de verplaatsing van de computer naar andere domeinen dan het traditionele office domein, steeds relevanter worden. Tot slot bediscussiëren we de mogelijkheden voor toepassing van de techniek buiten het desktop computer paradigma.

In hoofdstuk 6 verzamelen we onze bevindingen en speculeren we over de toekomst. We concluderen dat optisch gesimuleerde haptische feedback de tastbaarheid en kwaliteit van grafische gebruikersomgevingen kan verhogen. Gebruikers zijn goed in staat om, met de *active cursor* techniek gesimuleerde, haptische structuren in te schatten. Onze techniek om haptische feedback optisch te simuleren, opent een extra communicatiekanaal met de gebruiker dat de bruikbaarheid van de grafische gebruikersomgeving kan verhogen.

Omdat optisch gesimuleerde haptische feedback alleen toepasbaar is in combinatie met een visueel display, ligt het niet in de lijn der verwachting dat de actieve cursor techniek mechanisch gesimuleerde haptische feedback volledig zal gaan vervangen – de techniek is bijvoorbeeld geheel niet toepasbaar voor mensen met visuele beperkingen. Wel verwachten we dat de mogelijkheid om tactiele interactie stijlen te implementeren in een standaard grafische gebruikersomgeving, een katalyserende functie zou kunnen hebben voor de ontwikkeling en acceptatie van haptische apparaten.

Meer algemeen betogen we dat onderzoekers en ontwerpers van interactieve media beter gebruik zouden kunnen maken van de realiteit scheppende vermogens van de menselijke hersenen, in plaats van zich te beperken tot het nabouwen van situaties uit de pre-digitale omgeving: Nieuwe media creëren nieuwe perceptuele ruimtes.

Met dit onderzoek hebben we bijgedragen aan een meer zintuiglijk en rijker paradigma voor de grafische gebruikersomgeving. Daarnaast hebben we getracht ons algehele bewustzijn en begrip met betrekking tot media technologie en simulaties te vergroten. Het is om die reden dat ons wetenschappelijk onderzoek is gepresenteerd in een sociaal-culturele context die reflecteert op de dominantie van beelden en de toenemende rol van media en simulaties in ons dagelijkse leven.

Acknowledgements

Following a swimming degree, typing diploma, high school degree, a drivers licence and a masters degree, obtaining a PhD is perhaps the highest degree-wise goal you can set for yourself. Why would you? Undertaking a PhD research is in many regards a very personal endeavor. Luckily, it is not something you do alone. I want to thank:

Berry Eggen, my first promoter, who provided me with the opportunity to conduct a part-time PhD research at the TU/e and who has been the ideal promoter throughout the process – leaving me with a lot of freedom while staying involved at the same time. Thank you Berry, I am eagerly looking forward to our future collaborations.

Jack van Wijk, my second promoter, who became involved only in the later stages of the project, but nonetheless strongly influenced and helped in defining its shape. Thank you Jack, for your sharp feedback, guidance and support.

Dik Hermes, who provided me with my first scientific position at the Center for User System Interaction and who coached me in the early phases of my research activities. Thank you for the good collaboration and for stimulating my interest in real science.

Furthermore I want to express my gratitude to Gerrit van der Veer and Pieter-Jan Stappers for their inspiring comments on the draft version of this document. My big brother Martijn van Mensvoort, who has been my most profound critic throughout – I advise you to move into science yourself. My interns Koen Hendrix, Maurice van Montfort, Peter Vos and Sander Coenen, who have done such valuable work in the context of this project. Pascal de Man and Niels Schrader for their contribution to the PowerCursor engine. Digitale Pioniers and Fonds BKVB for financially supporting PowerCursor.com. Martin Boschman for technical assistance in the experimental setup. Carolyn F. Strauss who cleaned up some of my non-native English. Jim Juola who provided helpful comments on the empirical research chapters. Yvonne de Kort and Don Bouwhuis who helped with the statistics.

Staff and students of the TU/e department of Industrial Design, for enabling a vibrant and stimulating climate at the university that makes me want to kiss the future.

My dear parents Ad & Wil van Mensvoort, for exemplifying how to be good people.

My wife, Mieke Gerritzen, for her continuous advise, support, inspiration and making it look so very easy to do a PhD research.

Finally, the Lavazza coffee company, for providing me with such stimulating drugs throughout the process. Couldn't have done it without you.

PhD Comittee

Promotors

Prof. Dr. ir. Berry Eggen,
Prof. Dr. ir. Jack van Wijk

Core commitee

Prof. Dr. Gerrit van der Veer
Prof. Dr. Pieter-Jan Stappers
Dr. Dik Hermes

Additional members

Prof. Dr. Bas Haring
Prof. Dr. Jean-Bernard Martens
Prof. Dr. ir. Jeu Schouten (chair)

Propositions

1. Interactive animations can evoke an illusion of touch (Chapter 2).
2. Optically simulated haptic feedback can influence the perception of bump and hole structures, generated with a mechanical force feedback device (Chapter 3).
3. Optically simulated haptic feedback can increase the usability of a graphical user interface (Chapter 4).
4. Media technologies can evoke experiences that did not occur in the pre-mediated environment but nonetheless feel intuitive.
5. Today the walls of Plato's cave are so full of beamers, disco balls, plasma screens and halogen spotlights, that we don't even recognise the shadows on the wall anymore (Prologue).
6. The fact that the average western person worries more about financial meltdowns and mortgage deductions than about hurricanes or floods, supports for the argument that we have long traded basic elements of our reality for something more virtual.
7. Most children know more logos and brands than bird or tree species.
8. While our natural environment is being replaced by a world of design, at the same time our technological world becomes so intricate and complex that it becomes a nature of its own (Epilogue).
9. 'Attention' is the scarcest resource of the information age.
10. A new medium often mimics a familiar older medium. Yet, over time the older medium is superseded and transformed into a cultural relic. Examples of this principle are electronic mail, the horseless carriage – nowadays better known as the car – and the record collection on your mp3-player.
11. Keeping a strict border between a professional and a private life is unnatural and, in general, the diminishing of this border is a positive development.
12. The things we design often end up designing us.
13. Since our democratic system, by definition, is based on *representation*, the increasing influence of *images* in the political process is not necessarily a negative development.
14. Context is content.

Stellingen

1. Interactieve animaties kunnen een illusie van tast oproepen (Hoofdstuk 2).
2. Optisch gesimuleerde haptische feedback kan de perceptie van met een mechanisch force feedback apparaat gegenereerde heuvels en kuilen beïnvloeden (Hoofdstuk 3).
3. Optisch gesimuleerde haptische feedback kan de bruikbaarheid van de grafische gebruikersomgeving verhogen (Hoofdstuk 4).
4. Mediatechnologieën kunnen ervaringen oproepen, die in de pre-mediale omgeving niet voorkwamen, maar niettemin intuïtief aanvoelen.
5. Inmiddels is de grot van Plato zo volgehangen met projectoren, discoballen, plasmaschermen en halogeenspotjes, dat we zelfs de schaduwen op de muur niet meer herkennen (Proloog).
6. Het gegeven dat de gemiddelde westerse mens zich meer zorgen maakt over de financiële crisis en de hypotheekrenteaftrek, dan over orkanen of overstromingen, ondersteunt het argument dat we al lang geleden basale elementen van onze realiteit hebben ingeruild voor iets meer virtueels.
7. Tegenwoordig, kennen de meeste kinderen meer logo's en merken, dan vogel- of boomsoorten.
8. Terwijl onze natuurlijke omgeving gestaag wordt vervangen door een ontworpen omgeving, wordt tegelijkertijd onze technologische omgeving zo complex en veelomvattend dat we deze als een natuur op zichzelf moeten gaan beschouwen (Epiloog).
9. 'Aandacht' is de meest schaarse grondstof van het informatietijdperk.
10. Nieuwe media imiteren vaak een ouder vertrouwd medium, dat na verloop van tijd echter wordt overstemd en eindigt als cultureel relikwie. Voorbeelden van dit principe zijn: de elektronische post, de paardloze wagen – inmiddels beter bekend als de auto – en de platencollectie op je mp3-speler.
11. Het aanbrengen van een strikte grens tussen werk en privé is onnatuurlijk en het vervagen van deze grens is over het algemeen een positieve ontwikkeling.
12. De dingen die we ontwerpen, ontwerpen uiteindelijk vaak ook ons.
13. Aangezien ons democratische systeem per definitie is gebaseerd op 'representatie', is de toenemende invloed van 'beeldvorming' in het politieke proces niet noodzakelijk een negatieve ontwikkeling.
14. Context is content.

Curriculum Vitae

Koert van Mensvoort (1974) started his career in the late eighties with the creation of videogames – belonging to the first generation of whizkids who are now no longer kids. In the nineties he moved on and studied computer science, philosophy and art. He received a M.Sc in computer sciences from Eindhoven University of Technology (1997) and a MFA from the Sandberg Institute, Masters of Rietveld Academy, Amsterdam (2000).

Currently Van Mensvoort is co-director of the All Media Foundation (2005-), an Amsterdam based non profit organization that conceives, researches and visualizes current cultural issues. Furthermore he is a part-time assistant Professor at the Eindhoven University of Technology (Industrial Design Department) (2003-). Earlier Van Mensvoort worked as an associate researcher at the Center for User-System Interaction (1998-2003), as a teacher at the Sandberg Institute (2002-2006) and as a Visionary in Residence at Art Center College of Design in Pasadena (2008).

Much of his work revolves around the relation between people and media. Among his works are the Datafountain (an internet enabled water fountain connected to money currency rates), the online interactive dancefilm 'Drift' (featuring a dancer without a body), the TV documentary 'Daddy! The Woods smell of Shampoo', the Fake for Real memory game (on the tensed relation between reality and simulation) and the 'Biggest Visual Power Show', an intellectual spectacle blending between a scientific conference and a pop concert, held in Amsterdam (NL), Zeche Zollverein (DE) and Los Angeles (USA).

Van Mensvoort is (co)author of numerous books and publications; among them Next Nature, Visual Power, Natuur 2.0, Masters of Rietveld, Entry Paradise – New Worlds of Design, Artvertising, States of Nature, Style First and Nieuwe Media Cultuur in Nederland.

Van Mensvoort does not work in one specific media or style, but rather uses all media to visualize his ideas. His most profound experience in life, so far, has been the discovery of next nature. Which revolves around the idea that our technological world is so complex, that it has become a nature of its own. Many of his current activities relate to the exploration of this nature caused by human culture.

Van Mensvoort is married and lives in Amsterdam.

Websites: www.koert.com, www.nextnature.net, www.all-media.eu, www.powercursor.com, www.visual-power.com, www.fakeforreal.com, www.naonsdemens.nl

What You See Is What You Feel

On the simulation of touch in graphical user interfaces

With the introduction of the desktop metaphor, computers have become accessible for almost anyone; all over the world, people from various cultures use the same icons, folders, buttons and trashcans. From a sensorial point of view, however, this computing paradigm is still extremely limited.

This study began with an aspiration to increase the sensorial qualities of the graphical user interface. Inspired by renaissance painters – who centuries ago already applied various types of optical illusions in order to enhance the expressiveness of their paintings – we introduce a method of simulating touch with merely visual means. Interactive animations are used to create an optical illusion that evokes haptic percepts like stickiness, stiffness and mass, within a standard graphical user interface.

Our technique, called *optically simulated haptic feedback*, exploits the domination of the visual over the haptic modality and the general human tendency to integrate between the various senses. Experiments have shown that this method of simulating haptic feedback visually, can be a good alternative for mechanical force feedback devices.

Besides confirming that people are well able to recognize optically simulated bump and hole structures and having proven the usability benefits of the technique in a pointing task, we present a toolkit that allows interaction designers to apply the technique without elaborate programming.

With this study we hope to contribute to a more sensorial and richer graphical user interface. Moreover we have aimed to increase our awareness and understanding of media technology and simulations in general. Therefore, our scientific research is purposefully presented within a social-cultural context that reflects upon the dominance of the visual modality in our society and the ever-increasing role of media and simulations in people's everyday lives.

